

An introduction to the categorical p -adic Langlands program

Matthew Emerton, Toby Gee, and Eugen Hellmann

ABSTRACT. We give an introduction to the “categorical” approach to the p -adic Langlands program, in both the “Banach” and “analytic” settings.

CONTENTS

1. Introduction	2
1.1. A rapid overview of the Langlands program	2
1.2. The p -adic perspective	3
1.3. Proofs of reciprocity	4
1.4. Unramified geometric Langlands and the Fargues–Scholze conjecture	6
1.5. Some differences between the ℓ -adic and p -adic cases	8
1.6. Motivation from Taylor–Wiles patching	9
1.7. The “Galois to automorphic” direction	11
1.8. Topics we omit	11
1.9. A brief guide to the notes	11
1.10. Acknowledgements	12
2. Notation	12
2.1. Fields	12
2.2. p -adic Hodge theory	13
2.3. Hodge–Tate weights vs. highest weights	13
3. Taylor–Wiles patching as a motivation for categorical p -adic local Langlands	14
3.1. A very brief introduction to patching	14
3.2. p -adic Langlands: an overview and history	18
4. Moduli stacks of (φ, Γ) -modules: the Banach case	31
4.1. An overview of [EG22]	32
4.2. Fixed determinant variant	34
5. Moduli stacks of (φ, Γ) -modules: the analytic case	34
5.1. (φ, Γ) -modules over the Robba ring	34
5.2. Stacks of de Rham objects	47

M.E. was supported in part by the NSF grants DMS-1902307, DMS-1952705, and DMS-2201242. T.G. was supported in part by an ERC Advanced grant. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 884596). E.H. was supported by Germany’s Excellence Strategy EXC 2044-390685587 “Mathematics Münster: Dynamics–Geometry–Structure” and by the CRC 1442 Geometry: Deformations and Rigidity of the DFG.

5.3. Drinfeld-style compactifications	52
6. Categorical p -adic local Langlands conjectures	74
6.1. Expectations: The Banach Case	74
6.2. Expectations: the analytic case	83
7. Some known cases of categorical p -adic Langlands	97
7.1. GL_1	97
7.2. The Banach case for $GL_2(\mathbf{Q}_p)$ — I. The structure of \mathcal{X}	102
7.3. The Banach case for $GL_2(\mathbf{Q}_p)$ — II. Proof sketches	111
7.4. The Banach case for $GL_2(\mathbf{Q}_p)$ — III. Examples	121
7.5. Semiorthogonal decompositions	127
7.6. Preliminaries in the Banach case for $GL_2(\mathbf{Q}_{p^f})$	133
7.7. Partial results and examples in the Banach case for $GL_2(\mathbf{Q}_{p^2})$	139
7.8. The adjoint functor	158
7.9. Global functions and Hecke algebras	163
8. Categorical local Langlands in the case $\ell \neq p$	164
8.1. C -groups	164
8.2. Langlands parameters in the local case — $\ell \neq p$	165
8.3. Categorical local Langlands	166
9. Moduli stacks of global Langlands parameters and the cohomology of Shimura varieties	168
9.1. Global moduli stacks	168
9.2. The global-to-local map	169
9.3. Computing the cohomology of Shimura varieties via categorical Langlands	170
9.4. An example in the global case: $(PGL_2)_{/\mathbf{Q}}$	172
9.5. The structure of $\mathcal{X}_{GL_2, \mathbf{Q}, S}^{\text{odd}}$	173
9.6. Eigenvarieties and sheaves of overconvergent p -adic automorphic forms	174
Appendix A. Background on category theory	186
Appendix B. Coherent, Ind-Coherent, and Pro-Coherent sheaves on Ind-algebraic stacks	199
Appendix C. Localization and Morita theory	207
Appendix D. Topological abelian groups, rings and modules	217
Appendix E. Representations of p -adic analytic groups	221
Appendix F. Derived moduli stacks of group representations	228
References	230

1. Introduction

The goal of these notes are to discuss some of the p -adic aspects of the Langlands program, and especially the emerging “categorical” perspective. We begin with an overview of some history and our goals in these notes; we refer the reader to Section 2 for any unfamiliar notation.

1.1. A rapid overview of the Langlands program. The Langlands program began with Langlands’s celebrated letter to Weil (reproduced e.g. in [DS15, §6]), the contents of which were elaborated on in several subsequent writings of Langlands, including his article “Problems in the theory of automorphic forms” [Lan70]

and his Yale lectures on Euler products [Lan71]. The article focuses on the *functoriality* conjecture for automorphic representations and its consequences, while the Yale lectures explain his construction of automorphic L -functions, and discuss the problem of proving their expected analytic properties (analytic continuation and functional equation).¹

In his Yale lectures Langlands already touches, if only tangentially, on the idea that the automorphic L -functions may include all motivic L -functions. Subsequently, this idea developed into the *reciprocity* conjecture. (See [Lan79] for one articulation of this conjecture.) Roughly speaking, the reciprocity conjecture articulates a correspondence between certain automorphic representations (those which are *algebraic*) of $G(\mathbf{A}_F)$ (for a connected reductive algebraic group G over a number field F), and motives (with coefficients in $\overline{\mathbf{Q}}$, say) over F whose motivic Galois group is closely related to the C -group² cG of G . Conjectures on Galois representations (especially the Fontaine–Mazur conjecture [FM95]) suggest that such motives in turn may be identified with certain compatible systems of ℓ -adic Galois representations $\mathrm{Gal}_F \rightarrow {}^cG(\overline{\mathbf{Q}}_\ell)$.

There are many subtleties involved in trying to formulate a precise reciprocity conjecture. For example, in the case that G is not some GL_d , one has to worry about L -packets; and one should restrict to automorphic representations which are not *anomalous* in the sense of [Lan79] (or else replace Gal_F by $\mathrm{Gal}_F \times \mathrm{SL}_2$ and work in the framework of Arthur’s conjectures regarding non-tempered endoscopy [Art89]). We refer to [BG14] for a more precise discussion of the conjecture.

We note that there is a relationship between reciprocity and functoriality: namely, since functoriality is tautologically true for (compatible systems of) Galois representations, cases of reciprocity can be used to deduce cases of functoriality. Also, since L -groups and C -groups involve Galois groups in their definition, representations of Galois groups into algebraic groups can sometimes be related to L -homomorphisms of L -groups or C -homomorphisms of C -groups, and thus some cases of reciprocity can be subsumed into functoriality. (This is Langlands’s original perspective on the Artin conjecture, as explained in [Lan70].) If one replaces the Galois group Gal_F by the hypothetical Langlands group L_F , and compatible systems of Gal_F -representations by representations of L_F , then one can also extend the reciprocity conjecture to non-algebraic automorphic representations, or (using the “ L_F -form” of the L -group or C -group) entirely subsume reciprocity into functoriality. (A version of this last-mentioned perspective was already adopted at times by Langlands, by using the “Weil group form” of the L -group.)

1.2. The p -adic perspective. Our intention is not so much to focus on reciprocity in the manner described above (the relationship between automorphic representations and motives, or compatible systems of Galois representations), but rather to fix a prime p and consider the relationship between automorphic representations and p -adic Galois representations (for this fixed choice of p). At first, this

¹Since the standard L -functions associated to automorphic representations on GL_d coincide with the L -functions constructed by Godement and Tamagawa, which are known to admit analytic continuations and functional equations, the expected properties would follow from functoriality. Langlands’s Yale lectures discuss another approach, via constant terms of Eisenstein series.

²The C -group, introduced in [BG14], is a refinement of the L -group introduced by Langlands, which is better adapted to the problem of relating automorphic forms and Galois representations; we also recall the definition below.

doesn't much change the problem, since (at least for representations valued in GL_d) a compatible system of semisimple Galois representations is determined by any one of its members, and the Fontaine–Mazur conjecture gives a purely Galois-theoretic condition for a p -adic Galois representation to be motivic.

But focusing on a particular prime p brings to the fore certain aspects of the theory of automorphic forms and Galois representations which are absent in the more motivic perspective on reciprocity. For example, since at least the work of Ramanujan, it has been known that automorphic forms can satisfy interesting congruences modulo powers of p . Since the work of Swinnerton–Dyer [SD73] and Serre [Ser73], it has been understood that these congruences are related to (or manifestations of, if one prefers) analogous congruences between p -adic Galois representations. Extending the notion of congruences between automorphic forms, one is led to take various p -adic completions of spaces of automorphic forms to obtain the notion of p -adic automorphic forms, with associated Galois representations that need not be motivic. Related to this, one has notions of mod p automorphic forms, to which one might associate mod p Galois representations; and reciprocity conjectures have been formulated in this context, famously by Serre [Ser87] in the context of classical modular forms, and more recently in greater generality by others (e.g. [BDJ10]).

We do not intend at all to survey these developments; rather, our aim is simply to indicate that the p -adic perspective emerged naturally over a long period of time, and has led to a natural collection of problems and concerns: e.g. how to arrive at *some* notion of automorphic form which admits p -adic integral or mod p coefficients, and which allows for genuinely p -adic objects? And how to phrase reciprocity in a manner which allows for p -adic integral or mod p Galois representations, which will be associated to the p -adic integral or mod p objects that one introduces on the automorphic side?

We make one last remark on the p -adic theory for now: once one allows oneself to p -adically interpolate automorphic forms, one sees that automorphic forms naturally lie in families (e.g. Hida families, Coleman families, . . . ; see [Eme11b] for a survey). And Galois representations also lie in families (e.g. via Mazur's deformation theory [Maz89], which he was at least partly motivated to develop in response to Hida's theory of ordinary families of p -adic modular forms). These phenomena of continuous families of objects — on both the automorphic and Galois side — are not a feature of reciprocity in its more motivic formulation. (Cuspidal automorphic representations are rigid objects, and so are motives, if one has fixed a particular number field as the field of definition.) Before passing to our next topic, we note that families of Galois representations *are* a feature of the *geometric* Langlands program. We will return to this point below.

1.3. Proofs of reciprocity. Some of the very first results on reciprocity were proved in the case of representations with solvable image, by combining class field theory (which implies reciprocity in the abelian context) with techniques more classically automorphic in nature (see in particular the work of Langlands and Tunnell, [Lan80, Tun81]). In 1995, a breakthrough in our understanding of reciprocity was achieved with the proof by Wiles [Wil95] and Taylor–Wiles [TW95] of the modularity of semistable elliptic curves over \mathbf{Q} , a result which was soon improved

to handle all elliptic curves over \mathbf{Q} [BCDT01]. Building on these methods, reciprocity has since been proved in many further interesting contexts; see Calegari’s recent survey [Cal21] for some of the highlights of the last 25 years.

The crux of the method of [Wil95, TW95] and the many subsequent results that build on their ideas is to prove the modularity — or, more generally, automorphy — of a p -adic Galois representation. More precisely, the Taylor–Wiles method gives a way to deduce from the modularity of a single p -adic Galois representation ρ_1 the modularity of another p -adic Galois representation ρ_2 which is congruent to ρ_1 modulo p . It turns out that there are many such congruences between Galois representations, and modularity can be propagated from a single representation to many others in this way (either using a fixed prime p , or using several different primes p). Of course, it is necessary to have an initial supply of representations whose modularity is already known, such as CM forms, or 2-dimensional Artin representations with solvable image (the latter being used in Wiles’ proof of Fermat’s Last Theorem [Wil95]).

Modularity lifting theorems are also known as “ $R = \mathbf{T}$ ” theorems, because they are proved by identifying a \mathbf{Z}_p -algebra R (a “Galois deformation ring”) parameterizing Galois representations congruent to a fixed representation ρ_1 with another \mathbf{Z}_p -algebra \mathbf{T} , a “Hecke algebra”, the endomorphism algebra of an appropriate space of (p -adic) modular forms. Over the last 25 years it has become apparent that such theorems should hold in great generality, although finding appropriate definitions of R and \mathbf{T} so that we literally have $R = \mathbf{T}$ is still something of an art. In particular, it has become clear (following in particular Skinner–Wiles [SW99], Wake–Wang-Erickson [WWE17], Newton–Thorne [NT19]) that in general R should be taken to be a so-called pseudodeformation ring, i.e. a deformation ring for pseudorepresentations, rather than a deformation ring for literal Galois representations.

Roughly speaking, pseudorepresentations capture the information given by characteristic polynomials of representations; for this reason they are also known as pseudocharacters. They were originally introduced for GL_2 by Wiles [Wil88], and were considerably developed for GL_d by many authors (we highlight in particular Chenevier’s notion of a determinant [Che14], which showed how to define them for arbitrary primes p). A theory valid for general reductive groups was introduced by Vincent Lafforgue [Laf18].

From Lafforgue’s point of view, as recently made precise (in the setting of local Galois representations) by Fargues–Scholze [FS21, §VIII], a pseudodeformation ring R is the ring of global functions on a moduli stack of L -parameters, so that the moduli space of pseudorepresentations is a coarse moduli space for the moduli stack of L -parameters; and for general reductive groups, the Hecke algebra \mathbf{T} is replaced with a larger algebra, the algebra of “excursion operators” (which agrees with the usual Hecke algebra for GL_d).

From this perspective, it is natural to wonder whether it is possible to upgrade the putative equality $R = \mathbf{T}$ to a statement on the actual moduli stack of L -parameters, rather than on its coarse moduli space; for example, one could ask whether spaces of p -adic modular forms admit an interpretation in terms of sheaves on the stack of L -parameters. The possibility of such statements, and their local analogues, is the main subject of these notes; but before describing them further, we

turn to the geometric Langlands program, which was also a substantial motivation for Vincent Lafforgue’s work.

1.4. Unramified geometric Langlands and the Fargues–Scholze conjecture. Very roughly, the unramified geometric Langlands correspondence for a connected reductive group G and a curve X takes the form of an equivalence of categories

$$(1.4.1) \quad D(\text{const. Bun}_G(X)) \cong \text{QCoh}(\text{LocSys}_{\widehat{G}}(X)),$$

where the left hand side is a derived category of “constructible sheaves” on the stack Bun_G of G -bundles on X , and the right hand side is a derived category of “quasi-coherent sheaves” on a stack of “ \widehat{G} -local systems” on X ; see for example [Gai17].

The recent work of Fargues–Scholze [FS21] geometrizes the classical local Langlands correspondence by “taking X to be the Fargues–Fontaine curve”. (One motivation for doing this is that the fundamental group of the Fargues–Fontaine curve is Gal_F .) Slightly more precisely (but still only approximately), if F/\mathbf{Q}_p is a finite extension, and X is the Fargues–Fontaine curve for F , then for each prime $\ell \neq p$, one expects an equivalence of the approximate form (1.4.1), where the left hand side is replaced by an appropriate derived category of (solid) ℓ -adic sheaves on Bun_G , and the right hand side is a derived category of (Ind-)coherent sheaves on a moduli stack of \widehat{G} -valued ℓ -adic representations of Gal_F .

The stack Bun_G of G -bundles on the Fargues–Fontaine curve admits a stratification into locally closed substacks, with one of the open strata being a quotient stack $[\cdot/G(F)]$. The ℓ -adic sheaves on this quotient stack correspond to smooth ℓ -adic representations of $G(F)$, so by considering an appropriate pushforward of sheaves from this quotient stack to Bun_G , an equivalence of the form (1.4.1) implies in particular that there is a fully faithful functor

$$(1.4.2) \quad D(\text{sm. } G(F)) \rightarrow \text{QCoh}(\text{LocSys}_{\widehat{G}})$$

where the left hand side is a derived category of ℓ -adic representations of $G(F)$, and the right hand side is a derived category of quasicohherent sheaves on the stack of \widehat{G} -valued ℓ -adic representations of Gal_F .

Specialising to the case $G = \text{GL}_d$, we have the following rough conjecture, where we now allow the case $\ell = p$.

CONJECTURE 1.4.3. *If F is a finite extension of \mathbf{Q}_p , if \mathcal{O} is the ring of integers in a finite extension of \mathbf{Q}_ℓ (p and ℓ each denoting some fixed prime), and if d is a positive integer, then the category of smooth $\text{GL}_d(F)$ -representations on locally torsion \mathcal{O} -modules admits a fully faithful embedding into the category of quasicohherent sheaves on \mathcal{X} , an appropriate moduli stack parameterizing d -dimensional ℓ -adic representations of the absolute Galois group Gal_F .*

We also aim for a version of the conjecture where smooth p -power torsion representations of $\text{GL}_d(F)$ are replaced by locally analytic representations (i.e. the kind of p -adic representation theory relevant for the theory of overconvergent p -adic automorphic forms, Coleman families, ...), and the stack \mathcal{X} is replaced by a stack of equivariant vector bundles on the Fargues–Fontaine curve. As there is not even an informal way to incorporate this case into the above rough conjecture (that simultaneously treats the cases $\ell \neq p$ and $\ell = p$) we leave the formulation of the conjecture in the p -adic locally analytic case to the body of the notes. We only

remark that the conjecture in this case should be seen as an overconvergent version of the p -adic limit of Conjecture 1.4.3 (given by passing from torsion \mathcal{O} -modules to complete torsion-free \mathcal{O} -modules) in the $\ell = p$ case.

We have described this conjecture as “rough” for several reasons: one should be precise about the stack \mathcal{X} appearing in the statement of the conjecture — when $\ell \neq p$, it should be understood to be the stack of Weil–Deligne representations introduced in [Zhu20] (see also [DHKM20] and [FS21] for alternative constructions of the underlying classical stacks), while in the case $\ell = p$ it should be one of the moduli stacks constructed in [EG22] (and one should *a priori* consider these stacks as derived stacks — although in the $\ell \neq p$ context, it is shown in [Zhu20] that the resulting stacks are in fact classical, and we anticipate that the same will be true in the $\ell = p$ context); one should also be precise about what sort of categories are being considered — the envisaged embedding will not be compatible with the natural abelian category structure on source and target, so one should work with appropriate derived (or, probably better, stable infinity) categories (where “appropriate” is also doing some work — e.g. one should work with the triangulated/stable ∞ -category of Ind-coherent complexes on \mathcal{X} , and then make an analogous modification to the triangulated category of representations as well, as in [Zhu20, §4.1]).

REMARK 1.4.4. As far as we know, the possibility of a conjecture along the lines of Conjecture 1.4.3 with $\ell = p$ was first raised by Michael Harris (see [Har16, Question 4.7]). We first learned that a p -adic version of the conjecture of Fargues–Scholze in the form (1.4.1) should hold from Laurent Fargues in the summer of 2016. Peter Scholze explained some related conjectures to T.G. in 2018, and a rough formulation of a version of Conjecture 6.1.14 was discussed by M.E., T.G., E.H. and Scholze at the Hausdorff school in Bonn in 2019.

REMARK 1.4.5. We motivated Conjecture 1.4.3 by the work of Fargues–Scholze. In fact, in the case $\ell \neq p$ various conjectures along the lines of Conjecture 1.4.3 were independently proposed by Ben-Zvi–Chen–Helm–Nadler, Fargues–Scholze, E.H., and Zhu. The history of such conjectures is discussed in the introductions to the papers [FS21, Zhu20, Hel20, BZCHN20]. It is natural to guess that a similar conjecture should hold in the case $\ell = p$, but it is less clear what the precise formulation should be. We propose two precise conjectures in Section 6, one in the “Banach” case (focusing on representations on torsion \mathcal{O} -modules and in the limit on lattices in Banach space representations), and one in the “analytic” setting (focusing on locally analytic representations), and explain some relationships between them. However, it does not seem that either conjecture implies the other, and we do not know whether they admit a common refinement.

The two cases correspond to two classical notions of p -adic automorphic forms, or p -adic modular forms. Serre defined p -adic modular forms as limits of p -adic modular forms with torsion coefficients (or more precisely: as the p -adic limits of the reductions modulo powers of p of the q -expansions of modular forms). These spaces of p -adic modular forms certainly belong to the “Banach case”, whereas the subspace of overconvergent modular forms, as introduced by Katz, belong to the “analytic case”. In the discussion in these notes the space of p -adic modular forms is replaced by the completed cohomology of modular curves (or more general towers of Shimura varieties), as in Sections 3.2.11 and 9.3.1, and the space of overconvergent modular forms is then replaced by the space of locally analytic vectors underlying the completed cohomology, which is used to construct eigenvarieties, see Section 9.6.

REMARK 1.4.6. The reader might ask why we are only looking for fully faithful functors, rather than equivalences of categories as in the work of Fargues–Scholze described above. It does seem reasonable to hope that there are such equivalences of categories (even in the global setting, as speculated in [FS22, §7]). One difficulty is to even define categories of p -adic sheaves on Bun_G , but this difficulty is likely to be resolved by the recent work of Lucas Mann [Man22]. This work should then play an important role in formulating a precise conjecture.

Nevertheless, at the time of writing there is no such precise conjecture in the literature (let alone a theorem — and we note that as far as we are aware, it is not expected that the Fargues–Scholze construction of the spectral action will go over to the p -adic setting in any simple manner). In contrast, we can make precise conjectures about the fully faithful embeddings of categories of representations of $\mathrm{GL}_n(F)$, and even prove theorems in the case of $\mathrm{GL}_2(\mathbf{Q}_p)$. Furthermore, at least for $\mathrm{GL}_2(\mathbf{Q}_p)$ and $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ (and their inner forms), we produce semiorthogonal decompositions of the categories of (Ind-)coherent sheaves, which we expect to correspond to a semiorthogonal decomposition of the appropriate category of sheaves on Bun_G induced by the closure relations on Bun_G itself. (See Section 7.5 for some results in this direction.)

REMARK 1.4.7. There is no particular reason to restrict to the group GL_d in the statement of Conjecture 1.4.3, rather than considering dual groups or L -groups or C -groups of more general reductive groups over F (as in [BZCHN20, Hel20, Zhu20]). However, since we are interested in the $\ell = p$ case, we will for the most part focus on the case of GL_d , since the relevant “stacks of parameters” have been constructed in this case. (Since the literature on Taylor–Wiles patching, which sits in the background as an important source of intuition and motivation, also focuses on this case, this provides another reason for us to focus on it in this paper as well.)

1.5. Some differences between the ℓ -adic and p -adic cases. While it is possible to formulate ℓ -adic and p -adic versions of Conjecture 1.4.3 in a somewhat uniform fashion, there are significant differences between the two cases.

That there should be differences is not surprising from the point of view of representation theory. Indeed, already for $\mathrm{GL}_2(F)$, the classification of irreducible smooth mod ℓ representations is well-understood for $\ell \neq p$, and can be formulated uniformly for any F (these results are due to Vignéras, see [Vig89]). In contrast, for $\ell = p$, there is only a classification when $F = \mathbf{Q}_p$ (due to Barthel–Livné and Breuil [BL94, Bre03a]), and already for F a non-trivial unramified extension of \mathbf{Q}_p there are many complications, for example infinite families of admissible irreducible representations (as constructed by Breuil–Paškūnas [BP12]), and (absolutely) irreducible representations which are not admissible (constructed by Le [Le19]). Furthermore, supersingular irreducible admissible representations are never finitely presented (see Schraen’s paper [Sch15c] and Wu’s [Wu21b]).

These differences and difficulties are mirrored on the Galois side of the correspondence. For $\ell \neq p$ the moduli stacks of Galois representations are 0-dimensional ind-algebraic stacks of a rather mild type: the ind-structure arises from there being infinitely many connected components (controlled by the conductors of the Galois representations), and each connected component is a quotient of an affine scheme by a reductive group. Accordingly these stacks admit many global functions, and

have coarse moduli spaces which can be described in terms of pseudorepresentations (at least away from a small number of bad primes, this has been proved by Fargues–Scholze [FS21, Thm. VIII.3.6]).

In the Banach case for $\ell = p$, the moduli stacks we consider are those defined in [EG22]. These are also ind-algebraic stacks, but their dimension grows with $[F : \mathbf{Q}_p]$, and the ind-structure is much more complicated: indeed, they are naturally formal algebraic stacks with only finitely many connected components (just coming from a condition on central characters). Furthermore, we do not expect them to admit any interesting global functions, and beyond the case of $\mathrm{GL}_2(\mathbf{Q}_p)$ (and the case of tori) we do not expect them to have non-trivial associated moduli spaces.

In the analytic case the moduli stacks are defined in these notes (a more detailed discussion is going to appear in the forthcoming paper [HHS]) as the moduli stacks of equivariant vector bundles on the Fargues–Fontaine curve (or equivalently as moduli stacks of (φ, Γ) -modules over the Robba ring). As far as the dimension is concerned the analytic case behaves similar to the Banach case: the (expected) dimension grows with $[F : \mathbf{Q}_p]$ and should coincide with the dimension of the generic fiber of the formal stacks in the Banach case (and in fact conjecturally this generic fiber should be closely related to an open substack of the analytic moduli stack). However, these stacks now should not have an ind-structure but rather should be rigid analytic Artin stacks. Though there are a few more interesting global functions on these stacks we expect that they all arise from the theory of Hodge–Tate weights and hence the algebra of global functions still loses a lot of information about the moduli stack.

REMARK 1.5.1. Another difference between the cases $\ell = p$ and $\ell \neq p$ is the dependence on the field F . In the case $\ell \neq p$, the (in general conjectural) functors admit a qualitatively uniform description over all F , and it does not seem to be any easier to establish instances of the conjecture for $F = \mathbf{Q}_p$ or for a general F . In contrast, if $\ell = p$ then we only know how to construct a functor for $\mathrm{GL}_2(\mathbf{Q}_p)$ in the Banach case, in which case it shares some qualitative similarities with the case $\ell \neq p$ (for example, it is “not very derived”, with the only derived phenomena relating to finite-dimensional representations, and supersingular irreducible admissible representations roughly correspond to skyscraper sheaves supported on irreducible Galois representations).

However already for $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ we expect it is essential to consider the functor on the derived level, and supersingular irreducible admissible representations no longer correspond to skyscraper sheaves, but instead should correspond to complexes with non-coherent H^{-1} . (See Section 7.7.)

REMARK 1.5.2. The differences between the ℓ -adic and p -adic settings are reminiscent of the differences between the Betti and de Rham versions of geometric Langlands (see for example [BZN18] for the former and [Gai15] for the latter), although for reasons of ignorance we are unsure of how seriously one should take this analogy.

1.6. Motivation from Taylor–Wiles patching. Fargues has emphasised that the conjectures described in Section 1.4 did not arise from attempting to “copy and paste” from the geometric Langlands program to the setting of the Fargues–Fontaine curve, but rather from a study of p -adic period morphisms, the

p -adic geometry of Shimura varieties and Rapoport–Zink spaces, and the classical local Langlands correspondence (in particular the phenomena of L -packets and endoscopy). Similarly, while it may be possible to arrive at the conjectures presented here by an appropriate procedure of “setting $\ell = p$ in the Fargues–Scholze conjecture”, we instead came to them via a roundabout process over many years, largely motivated by considerations coming from p -adic automorphic forms and the Taylor–Wiles patching method.

As well as its importance in shaping our ideas, the Taylor–Wiles patching method has suggested several important properties of the conjectural functors as in Conjecture 1.4.3, even in the case $\ell \neq p$. For example, the patching method suggests that the kernel of such a functor should be a genuine sheaf (rather than a complex thereof), a property that was not at all apparent in the first versions of the constructions of Ben-Zvi–Chen–Helm–Nadler and Zhu. (There is also a strong connection between the Taylor–Wiles method and the classical p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$, for which see [CEG⁺18].)

In brief, the connection between Conjecture 1.4.3 and the Taylor–Wiles method is as follows. A typical example of a smooth representation of $\mathrm{GL}_d(F)$ is a compact induction $c\text{-Ind}_{\mathrm{GL}_d(\mathcal{O}_F)}^{\mathrm{GL}_d(F)} V$, where V is a continuous (i.e. smooth) representation of $\mathrm{GL}_d(\mathcal{O}_F)$ on a literally finite \mathcal{O} -module V . One envisages that the embedding of Conjecture 1.4.3 will map such representations to genuine coherent sheaves (as opposed to complexes of such), say $\mathfrak{A}(V)$, and that these coherent sheaves will be “globalizations” over \mathcal{X} of the patched modules $M_\infty(V)$ that arise in the Taylor–Wiles–Kisin patching method, in the sense that if R_∞ is one of the rings usually denoted this way in the theory of patching, and if $\mathrm{Spf} R_\infty \rightarrow \mathcal{X}$ is the natural versal morphism, then the pullback of $\mathfrak{A}(V)$ along this map coincides with $M_\infty(V)$. In particular, the conjecture entails that patched modules should be “purely local” (where now “local” is understood to pertain to the local Galois-theoretic situation, in opposition to the global automorphic/Galois-theoretic choice that was made in order to perform patching).

Ascertaining the truth or otherwise of this “pure locality” in general is well-known to be a major open problem in the arithmetic theory of automorphic forms. It is worth noticing that, building upon [CEG⁺16], the patching strategy has been extended to the “analytic case” in [BHS17b]. Using roughly the same formulation as above one can associate patching modules $M_\infty^{\mathrm{rig}}(V)$ to locally analytic principal series representations $V = \mathrm{Ind}_B^G(\delta)^{\mathrm{an}}$. The corresponding coherent sheaf should then be the pullback of the evaluation of the conjectural “analytic” functor on V (and is related to the fiber over δ of the sheaf of overconvergent p -adic automorphic forms of finite slope on eigenvarieties in a similar way as the classical patching modules are related to modules of classical automorphic forms).

In the opposite direction, since the patched modules $M_\infty(V)$ are constructed from the cohomology of Shimura varieties (with coefficients in a local system corresponding to V), it is natural to ask whether the cohomology of Shimura varieties can in some way be computed in terms of the (conjectural) functors of Conjecture 1.4.3. This should indeed be the case, as we explain in Section 9.2 (and as will be more thoroughly explained in the paper in preparation of M.E. and Zhu [EZ]). In parallel to this, Section 9.6 gives a precise computation of the sheaves of overconvergent automorphic forms on eigenvarieties (that can be defined in terms of

completed cohomology of Shimura varieties) in terms of the conjectural functor in the analytic case.

1.7. The “Galois to automorphic” direction. A variant of the expectation that patched modules are “purely local” is the hope that it should be possible to canonically associate an admissible (smooth) representation of $\mathrm{GL}_n(F)$ to each n -dimensional p -adic representation of Gal_F ; and that this association should be realised in the cohomology of Shimura varieties. This expectation is borne out by a good deal of evidence in the case of Shimura curves, the most recent example being the spectacular results on GK-dimensions due to Breuil, Herzig, Hu, Morra, and Schraen [BHH⁺20, BHH⁺21].

While our focus is mostly on functors from smooth representations to sheaves on stacks of Galois representations, our conjectures also entail the existence of adjoint functors from sheaves to representations of p -adic groups. In particular, applying these to “skyscraper” sheaves (that is, to sheaves of finite length, supported at particular Galois representations), we are able to show that our conjectures are consistent with the above-mentioned hope. (See e.g. Remark 7.8.5.)

1.8. Topics we omit. For reasons of time and ignorance, we do not discuss many recent (and even not so recent) developments in the p -adic Langlands program, which we nonetheless expect have important connections to the ideas discussed here. In particular we say nothing about the cohomology of the Drinfeld upper half plane and related p -adic locally symmetric spaces, as studied by Colmez, Dospinescu, Le Bras, Niziol and Pan [CDN20, CDN22, CDN21, DLB17, Pan17] (the interested reader could see <https://mathoverflow.net/questions/432932/how-does-the-cohomology-of-the-lubin-tate-drinfeld-tower-fit-into-categorical-p> for an explanation of the expected connections); about Scholze’s p -adic Jacquet–Langlands functor [Sch15b]; or about the coherent cohomology of Shimura varieties, and in particular about any relationship to Boxer and Pilloni’s higher Hida and Coleman theories [Pil20, BP20, BP21] or to Diamond and Sasaki’s geometric Serre weight conjectures [DS22].

1.9. A brief guide to the notes. These notes cover many topics, and we hope that it is possible to read many of the sections more or less independently from each other, referring back for definitions where necessary. In Section 3 we give an idiosyncratic introduction to the Taylor–Wiles method, guided by its connections to the classical p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$, and its role in motivating our conjectures. In Section 4 we very briefly recall the definitions and basic properties of the stacks of (φ, Γ) -modules from [EG22]. The rather longer Section 5 explores analogues of these constructions in the “analytic” context of (φ, Γ) -modules over Robba rings.

Section 6 states our main (local) p -adic Langlands conjectures, in both the Banach and analytic settings, and explains their relationships to other topics in the literature, e.g. the Breuil–Mézard conjecture. We then explain what is known about these conjectures in Section 7; as well as the relatively straightforward case of GL_1 , we explain forthcoming work of Andrea Dotto with M.E. and T.G. which proves the Banach conjecture for $\mathrm{GL}_2(\mathbf{Q}_p)$, and we explain some evidence for the Banach conjecture for $\mathrm{GL}_2(F)$ for F/\mathbf{Q}_p unramified. Section 8 briefly recalls some of the results and expectations in the case $\ell \neq p$. This material is used in Section 9,

which explains conjectures on the (completed) cohomology of Shimura varieties and eigenvarieties.

The appendices recall various results on (infinity-) category theory and the representation theory of p -adic analytic groups, and establish others for which we do not know of a reference.

REMARK 1.9.1. The categorical p -adic Langlands correspondence is still in its infancy, and (as these notes make plain) it is unclear to us what the ultimate form of many of the conjectures we make should be. There is also a considerable amount of technical machinery needed to get off the ground. While we have endeavoured to make precise and correct statements, our aim has been where possible to explain and illustrate phenomena rather than give a detailed development of all of this machinery. In particular, we for the most part treat infinity categories and derived algebraic geometry as a black box.

It seems extremely likely that the future development of the theory will rely on condensed mathematics, and in particular the theory of solid rings (and modules, sheaves. . .). It is also likely that some of the definitions and constructions would be streamlined by phrasing them in these terms. However, due to our own limitations, we have for the most part not attempted to do this.

1.10. Acknowledgements. We would like to thank the organisers (Pierre-Henri Chaudouard, Wee Teck Gan, Tasho Kaletha, and Yiannis Sakellaridis) and participants of the IHES Summer School on the Langlands Program for the opportunity and encouragement to write these notes.

We would like to thank the many colleagues, collaborators and friends with whom we have discussed the ideas presented in these notes. In particular we would like to thank Johannes Anschütz, David Ben-Zvi, Roman Bezrukavnikov, George Boxer, Christophe Breuil, Frank Calegari, Ana Caraiani, Pierre Colmez, Gabriel Dospinescu, Andrea Dotto, Laurent Fargues, Tony Feng, Dennis Gaitsgory, David Geraghty, Ian Grojnowski, David Hansen, Michael Harris, David Helm, Christian Johansson, Arthur-César Le Bras, Bao Le Hung, Brandon Levin, Jacob Lurie, Lucas Mann, David Nadler, James Newton, Vytautas Paškūnas, Vincent Pilloni, Alice Pozzi, Joaquin Rodrigues Jacinto, David Savitt, Peter Scholze, Benjamin Schraen, Jack Sempliner, Sug Woo Shin, James Timmins, Pol van Hoften, Akshay Venkatesh, Carl Wang-Erickson, and Xinwen Zhu.

We would also like to thank Pierre Colmez, Fred Diamond, Andrea Dotto, David Hansen, Claudius Heyer, Vytautas Paškūnas, and David Savitt for their comments on earlier versions of these notes.

2. Notation

2.1. Fields. We fix an algebraic closure $\overline{\mathbf{Q}}$ of \mathbf{Q} , and algebraic closures $\overline{\mathbf{Q}}_p$ of \mathbf{Q}_p for each p . If M is a number field, we write $\mathrm{Gal}_M := \mathrm{Gal}(\overline{\mathbf{Q}}/M)$ for the absolute Galois group of M . If S is a set of places of M then we write $M(S)/M$ for the maximal extension unramified outside of S , and write $\mathrm{Gal}_{M,S} := \mathrm{Gal}(M(S)/M)$. We fix embeddings $\overline{\mathbf{Q}} \hookrightarrow \overline{\mathbf{Q}}_p$ for each p and $\overline{\mathbf{Q}} \hookrightarrow \mathbf{C}$, so that if v is a place of M , there is a homomorphism $\mathrm{Gal}_{M_v} \rightarrow \mathrm{Gal}_M$.

If F/\mathbf{Q}_p is a finite extension, write Gal_F for the absolute Galois group $\mathrm{Gal}(\overline{\mathbf{Q}}_p/F)$. Write I_F for the inertia subgroup of Gal_F , W_F for the Weil group, and k_F for the

residue field of the ring of integers \mathcal{O}_F of F . We normalise local class field theory so that a uniformizer corresponds to a geometric Frobenius element.

Let \mathcal{O} denote the ring of integers in a fixed finite extension L of \mathbf{Q}_p , let k be the residue field of \mathcal{O} , and let ϖ denote a uniformizer of \mathcal{O} .

2.2. p -adic Hodge theory. If ρ is a de Rham representation of Gal_F on an L -vector space W , then we will write $\mathrm{WD}(\rho)$ for the corresponding Weil–Deligne representation of W_F , and if $\sigma : F \hookrightarrow L$ is a continuous embedding of fields then we will write $\mathrm{HT}_\sigma(\rho)$ for the multiset of Hodge–Tate numbers of ρ with respect to σ , which by definition contains i with multiplicity $\dim_L(W \otimes_{\sigma, F} \widehat{F}(i))^{\mathrm{Gal}_F}$; for example, if ε denotes the p -adic cyclotomic character, then $\mathrm{HT}_\sigma(\varepsilon) = \{-1\}$.

Suppose that L contains the images of all continuous embeddings $F \hookrightarrow \overline{L}$. By a d -tuple of labeled Hodge–Tate weights $\underline{\lambda}$, we mean a tuple of integers $\{\lambda_{\sigma, i}\}_{\sigma: F \hookrightarrow L, 1 \leq i \leq d}$ with $\lambda_{\sigma, i} \geq \lambda_{\sigma, i+1}$ for all σ and all $1 \leq i \leq d-1$. We will also refer to $\underline{\lambda}$ as a *Hodge type*. We say that $\underline{\lambda}$ is *regular* if $\lambda_{\sigma, i} > \lambda_{\sigma, i+1}$ for all σ and all $1 \leq i \leq d-1$. We say that a de Rham representation ρ has Hodge type $\underline{\lambda}$ (or labeled Hodge–Tate weights $\underline{\lambda}$) if for each $\sigma : F \hookrightarrow L$ we have $\mathrm{HT}_\sigma(\rho) = \{\lambda_{\sigma, i}\}_{1 \leq i \leq d}$.

By an *inertial type* τ we mean a representation $\tau : I_F \rightarrow \mathrm{GL}_d(L)$ which extends to a representation of W_F with open kernel (so in particular, τ has finite image). We say that a de Rham representation ρ has inertial type τ if $\mathrm{WD}(\rho)|_{I_F} \cong \tau$.

2.3. Hodge–Tate weights vs. highest weights. Under the p -adic local Langlands correspondence, Hodge–Tate weights on the Galois-theoretic side of the correspondence will be related to highest weights on the representation theoretic side, and it is useful to have notation adapted to either side of the correspondence.

DEFINITION 2.3.1. If $\underline{\lambda}$ is a regular Hodge type, then we write $\xi_{\sigma, i} = i - 1 - \lambda_{\sigma, d+1-i}$, so that $\xi_{\sigma, 1} \geq \dots \geq \xi_{\sigma, d}$.

DEFINITION 2.3.2. If $\underline{\lambda}$ is a regular Hodge type, then we view each $\xi_\sigma := (\xi_{\sigma, 1}, \dots, \xi_{\sigma, d})$ as a dominant weight of the algebraic group GL_d (with respect to the upper triangular Borel subgroup). We let $W_{\underline{\lambda}}$ be the corresponding \mathcal{O} -representation of $\mathrm{GL}_d(\mathcal{O}_F)$, defined as follows: for each $\sigma : F \hookrightarrow L$, we write W_{ξ_σ} for the algebraic \mathcal{O}_F -representation of K (more precisely, the dual Weyl module) of highest weight ξ_σ . Then we define

$$W_{\underline{\lambda}} := \otimes_{\sigma} W_{\xi_\sigma} \otimes_{\mathcal{O}_F, \sigma} \mathcal{O},$$

and write $L(\underline{\lambda}) = W_{\underline{\lambda}}[1/p]$, a representation of $\mathrm{GL}_d(F)$ (or of the algebraic group $\mathrm{Res}_F^{\mathbf{Q}_p} \mathrm{GL}_d$).

REMARK 2.3.3. The preceding notation may not be optimal, but serves our purposes. The “ W ” stands for (dual) *Weyl*, and emphasizing the dependence on $\underline{\lambda}$ (rather than on ξ) facilitates the comparison with other considerations of p -adic Hodge theory. The notation $L(\underline{\lambda})$ is chosen because it is traditional in the theory of category \mathcal{O} . Indeed, for any (not necessarily dominant, not necessarily integral) weight ξ (of the Lie algebra of $\mathrm{Res}_F^{\mathbf{Q}_p} \mathrm{GL}_d$), we use $M(\xi)$ to denote the Verma module of highest weight ξ , and $L(\xi)$ to denote the unique simple quotient of $M(\xi)$. When ξ is dominant integral, one finds that $L(\xi)$ is the usual highest weight representation $W_{\underline{\lambda}}[1/p]$, explaining our choice of notation in this case.

3. Taylor–Wiles patching as a motivation for categorical p -adic local Langlands

In this section we review some of the history of the Taylor–Wiles patching method and of the p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$, emphasising where possible the connections between them. We do not pretend to give a thorough overview, and we view everything through the lens of the categorical Langlands program. The reader may wish to consult Calegari’s survey [Cal21] for a more mainstream account of developments in the Taylor–Wiles method since its inception.

3.1. A very brief introduction to patching. Patching was first introduced by Taylor and Wiles as a technique for proving modularity lifting theorems [TW95]. It was then further developed by Kisin, who showed that it provides a mechanism for relating local and global aspects of the theory of p -adic Galois representations [Kis09c, Kis07a, Kis09a, Kis10b, GK14].

We place ourselves in the following simple but illustrative context. Fix a continuous, absolutely irreducible, modular representation $\bar{r} : \mathrm{Gal}_{\mathbf{Q},\{p,\infty\}} \rightarrow \mathrm{GL}_2(k)$ with determinant $\bar{\varepsilon}^{-1}$, and suppose that $\bar{\rho} := \bar{r}|_{\mathrm{Gal}_{\mathbf{Q}_p}}$ and $\bar{r}|_{\mathrm{Gal}_{\mathbf{Q}(\zeta_p)}}$ are also absolutely irreducible.

For any set S of primes containing p and ∞ , we may consider the global deformation space $\mathcal{X}_S(\bar{r})$ of \bar{r} over \mathcal{O} , which parameterizes deformations of \bar{r} to representations of $\mathrm{Gal}_{\mathbf{Q},S}$ over complete Noetherian local \mathcal{O} -algebras, with fixed determinant ε^{-1} . We may also consider the local deformation space $\mathcal{X}(\bar{\rho})$, which parameterizes deformations of $\bar{\rho}$ to representations of $\mathrm{Gal}_{\mathbf{Q}_p}$ over complete Noetherian local \mathcal{O} -algebras. These are each the Spf of a complete Noetherian local \mathcal{O} -algebra, denoted $R_S(\bar{r})$ and $R(\bar{\rho})$ respectively. Restriction of Galois representations induces a morphism

$$(3.1.1) \quad \mathcal{X}_S(\bar{r}) \rightarrow \mathcal{X}(\bar{\rho}),$$

which is a finite morphism (see [EP20] and [AC14]).

The reduced tangent space to the closed point of $\mathcal{X}_S(\bar{r})$ (resp. $\mathcal{X}(\bar{\rho})$) is equal to the global Galois cohomology group $H^1(G_{\mathbf{Q},S}, \mathrm{Ad}^0(\bar{r}))$ (resp. to the local Galois cohomology group $H^1(G_{\mathbf{Q}_p}, \mathrm{Ad}^0(\bar{\rho}))$), and the restriction map between deformation spaces induces the restriction map

$$(3.1.2) \quad H^1(G_{\mathbf{Q},S}, \mathrm{Ad}^0(\bar{r})) \rightarrow H^1(G_{\mathbf{Q}_p}, \mathrm{Ad}^0(\bar{\rho}))$$

on reduced tangent spaces, which is given by the usual pullback of cohomology classes.

In order to simplify the situation, we will assume that (3.1.1) is in fact a closed immersion when $S = \{p, \infty\}$; equivalently, we assume that the Selmer group³ $\mathrm{Sel}(G_{\mathbf{Q},\{p,\infty\}}, \mathrm{Ad}^0(\bar{r}))$ vanishes. Under this assumption (and the running assumption that $\bar{r}|_{\mathrm{Gal}_{\mathbf{Q}(\zeta_p)}}$ is absolutely irreducible), the formula of Greenberg–Wiles ([DDT97, Thm. 2.18]) shows that

$$\dim_k \mathrm{Sel}(G_{\mathbf{Q},\{p,\infty\}}, \mathrm{Ad}^0(\bar{r})(1)) = \dim_k \mathrm{Sel}(G_{\mathbf{Q},\{p,\infty\}}, \mathrm{Ad}^0(\bar{r})) + 1 = 1.$$

³This is a “strict” Selmer group, i.e. defined by the condition of a cohomology class vanishing locally at p .

The Taylor–Wiles method then requires us to consider primes $q \equiv 1 \pmod p$ such that the restriction morphism $H^1(G_{\mathbf{Q}, \{p, \infty\}}, \text{Ad}^0(\bar{\rho})(1)) \rightarrow H^1(\text{Gal}_{\mathbf{Q}_q}, \text{Ad}^0(\bar{\rho})(1))$ is non-zero (equivalently, injective), and such that $\bar{\rho}(\text{Frob}_q)$ has distinct eigenvalues; these are the so-called *Taylor–Wiles primes*. Another application of the Greenberg–Wiles formula then shows that

$$\text{Sel}(\text{Gal}_{\mathbf{Q}, \{p, q, \infty\}}, \text{Ad}^0(\bar{\rho})) = 0$$

for such a Taylor–Wiles prime q , and thus that the restriction morphism of deformation spaces

$$(3.1.3) \quad \mathcal{X}_{\{p, q, \infty\}}(\bar{\rho}) \rightarrow \mathcal{X}(\bar{\rho})$$

is again a closed immersion.

Write $\bar{\rho}_q := \bar{\rho}|_{\text{Gal}_{\mathbf{Q}_q}}$. We now consider the restriction morphism of deformation spaces

$$\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho}) \rightarrow \mathcal{X}(\bar{\rho}_q).$$

Under our hypotheses that $q \equiv 1 \pmod p$ and that $\bar{\rho}(\text{Frob}_q)$ has distinct eigenvalues, it can be shown that $\mathcal{X}(\bar{\rho}_q) = \text{Spf } R(\bar{\rho}_q)$ takes a very simple form: namely, $R(\bar{\rho}_q)$ is of the form $\mathcal{O}[[x, y]]/((1+x)^{p^n} - 1)$, where p^n is the largest power of p dividing $(q-1)$. (More precisely, one checks that all deformations of $\bar{\rho}_q$ are reducible, and the variables x, y then correspond respectively to a generator of the tame inertia group and to a Frobenius element.) In particular, we obtain a map

$$(3.1.4) \quad \mathcal{X}_{\{p, q, \infty\}}(\bar{\rho}) \rightarrow \text{Spec } \mathcal{O}[[x]]/((1+x)^{p^n} - 1),$$

whose fibre over the locus $x = 0$ is equal to $\mathcal{X}_{\{p, \infty\}}(\bar{\rho})$; in this way we may regard $\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho})$ as a kind of thickening⁴ of $\mathcal{X}_{\{p, \infty\}}(\bar{\rho})$. (This interpretation of the situation would be optimal if the morphism (3.1.4) were flat. This is not always the case, but what comes out of the Taylor–Wiles method is that $\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho})$ supports a faithful module which is flat over $\mathcal{O}[[x]]/((1+x)^{p^n} - 1)$, namely the completed homology group $\tilde{H}_1(q)$ introduced below, so the situation is always fairly close to the optimal one.)

Now it turns out that $\mathcal{X}_{\{p, \infty\}}(\bar{\rho})$ has relative dimension 2 over \mathcal{O} , while $\mathcal{X}(\bar{\rho})$ has relative dimension 3 over \mathcal{O} . Thus the closed immersion (3.1.1) has codimension 1, and the closed immersion (3.1.3) then realizes $\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho})$ as a thickening (in the “ x direction”, if we refer to (3.1.4)) of $\mathcal{X}_{\{p, \infty\}}(\bar{\rho})$ along the transverse direction to $\mathcal{X}_{\{p, \infty\}}(\bar{\rho})$ in $\mathcal{X}(\bar{\rho})$.

A key point is that for any given power p^n of p , we can find plenty of Taylor–Wiles primes q (ultimately via a Čebotarev argument), such that $q \equiv 1 \pmod{p^n}$. Thus we can thicken up $\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho})$ as much as we like inside $\mathcal{X}(\bar{\rho})$, and the rather surprising idea at the heart of the patching method is to try to realize $\mathcal{X}(\bar{\rho})$ as some kind of limit of all these thickenings $\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho})$.

To make this precise, we first regard each deformation ring $R(\bar{\rho}_q)$ as an $S_\infty := \mathcal{O}[[x]]$ -algebra (this gives a uniform meaning to the “ x direction” referred to above as q varies), and then try to glue together the various $\text{Spec } S_\infty$ -schemes $\mathcal{X}_{\{p, q, \infty\}}(\bar{\rho})$ as q varies in a manner compatible with their embeddings into $\mathcal{X}(\bar{\rho})$; equivalently, we try to glue together the various $(S_\infty, R(\bar{\rho}))$ -bialgebras $R(\bar{\rho}_q)$. As part of the

⁴Since if we look modulo the uniformizer of \mathcal{O} , then we find that $\text{Spec } k[[x]]/(1+x)^{p^n} - 1 = \text{Spec } k[x]/x^{p^n}$ is a nilpotent thickening of $\text{Spec } k$.

patching process, we simultaneously glue to together various (completed) homology groups.

Recall that we define completed (first) homology of the modular curve as

$$(3.1.5) \quad \tilde{H}_1 := \varprojlim_r H_1(Y(p^r), \mathcal{O}) \cong \varprojlim_r H_c^1(Y(p^r), \mathcal{O}).$$

Here $Y(p^r)$ is the modular curve of level p^r (or more precisely a congruence quotient of PGL_2 -symmetric spaces, which is a union of connected components of the usual modular curves for GL_2).

This is acted on by a Hecke algebra \mathbf{T} generated by the prime-to- p Hecke operators, as well as by $\mathrm{PGL}_2(\mathbf{Q}_p)$ and $\mathrm{Gal}_{\mathbf{Q}, \{p, \infty\}}$. All of these actions commute with one another. There is a maximal ideal \mathfrak{m} of \mathbf{T} corresponding to the modular Galois representation \bar{r} . We can localize \tilde{H}_1 at \mathfrak{m} to get a $\mathbf{T}_{\mathfrak{m}}$ -module $\tilde{H}_{1, \mathfrak{m}}$. Similarly, for each q as above we have the completed homology group

$$\tilde{H}_1(q) := \varprojlim_r H_1(Y(p^r q), \mathcal{O}) \cong \varprojlim_r H_c^1(Y(p^r q), \mathcal{O}),$$

which has actions of $\mathrm{PGL}_2(\mathbf{Q}_p)$, of $\mathrm{Gal}_{\mathbf{Q}, \{p, q, \infty\}}$, of prime-to- pq Hecke operators, and of certain Hecke operators at q . We can and do localize at \mathfrak{m} to get S_{∞} - and $R(\bar{\rho})$ -modules $\tilde{H}_{1, \mathfrak{m}}(q)$. In fact, since the action of S_{∞} on $\tilde{H}_{1, \mathfrak{m}}(q)$ is by definition via the composite of our chosen morphism $S_{\infty} \rightarrow R(\bar{r}_q)$ and the natural morphism $R(\bar{r}_q) \rightarrow R_{\{p, q, \infty\}}(\bar{r})$, we can use the formal smoothness of S_{∞} and the closed immersion (3.1.3) to produce a morphism $S_{\infty} \rightarrow R(\bar{\rho})$, compatible with the actions of both rings on $\tilde{H}_{1, \mathfrak{m}}(q)$.

Following [Sch15b, §9], we make an ultraproduct construction to produce an $R(\bar{\rho})$ -module M_{∞} with an action of $\mathrm{PGL}_2(\mathbf{Q}_p)$, with the property that

$$M_{\infty} \otimes_{R(\bar{\rho})} R_{\{p, \infty\}}(\bar{r}) = \tilde{H}_{1, \mathfrak{m}}.$$

By construction (and local-global compatibility at the primes q), M_{∞} is a finite projective $S_{\infty}[[K]]$ -module, where $K = \mathrm{PGL}_2(\mathbf{Z}_p)$.

In general one cannot assume that the Selmer group $\mathrm{Sel}(G_{\mathbf{Q}, \{p, \infty\}}, \mathrm{Ad}^0(\bar{r}))$ vanishes (equivalently, we cannot assume that (3.1.1) is a closed immersion) and the Taylor–Wiles method goes as follows. Let $g := \dim_k \mathrm{Sel}(G_{\mathbf{Q}, \{p, \infty\}}, \mathrm{Ad}^0(\bar{r})) + 1$. The Čebotarev argument alluded to above shows that for each n there is a set Q_n of Taylor–Wiles primes with the properties that

$$(3.1.6) \quad \dim_k \mathrm{Sel}(G_{\mathbf{Q}, \{p, \infty\} \cup Q_n}, \mathrm{Ad}^0(\bar{r})) = \dim_k \mathrm{Sel}(G_{\mathbf{Q}, \{p, \infty\}}, \mathrm{Ad}^0(\bar{r})),$$

and for each $q \in Q_n$, we have $q \equiv 1 \pmod{p^n}$. We write $S_{\infty} := \mathcal{O}[[x_1, \dots, x_g]]$ and $R_{\infty} := R(\bar{\rho})[[y_1, \dots, y_{g-1}]]$ for formal variables x_i, y_i . By (3.1.6), we may lift (3.1.1) and (3.1.3) to (non-canonical) closed immersions $\mathcal{X}_{\{p, \infty\}}(\bar{r}) \rightarrow \mathrm{Spf} R_{\infty}$ and $\mathcal{X}_{\{p, \infty\} \cup Q_n}(\bar{r}) \rightarrow \mathrm{Spf} R_{\infty}$, and straightforward generalizations of the constructions explained above in the case $g = 1$ yield a finite projective $S_{\infty}[[K]]$ -module M_{∞} with an action of R_{∞} , and we can again produce a morphism $S_{\infty} \rightarrow R_{\infty}$ compatible with the actions of each ring on M_{∞} . By construction, we again have

$$(3.1.7) \quad M_{\infty} \otimes_{R_{\infty}} R_{\{p, \infty\}}(\bar{r}) = \tilde{H}_{1, \mathfrak{m}}.$$

The patched module M_{∞} can be thought of as being a way to extract the local at p information from completed cohomology. Its construction depends on many choices, but the perspective first taken in [CEG⁺16] is that it should be purely local in an obvious sense, and that the same construction for more general Shimura

varieties should give a candidate for a p -adic local Langlands correspondence beyond the case of $\mathrm{GL}_2(\mathbf{Q}_p)$. This expectation has yet to be verified, but it is one that has been important for us in finding the conjectures presented in these notes, and the modules M_∞ appear in the statements of our conjectures (see in particular Remark 6.1.30).

3.1.8. *Patching at finite level.* If V is any finitely generated \mathcal{O} -module with a continuous action of K , then we set

$$(3.1.9) \quad M_\infty(V) := M_\infty \otimes_{\mathcal{O}[[K]]} V.$$

Then $V \mapsto M_\infty(V)$ is an exact functor, and (3.1.7) implies that

$$M_\infty(V) \otimes_{R_\infty} R_{\{p, \infty\}}(\bar{r}) = H_c^1(V)_\mathfrak{m},$$

where $H_c^1(V)_\mathfrak{m}$ denotes the (localized at \mathfrak{m}) compactly supported cohomology group with coefficients in the local system determined by V . (Since \mathfrak{m} is non-Eisenstein, the natural map $H_c^1(V)_\mathfrak{m} \rightarrow H^1(V)_\mathfrak{m}$ is an isomorphism, but we have used the H_c^1 notation to emphasize that we are working with *homology*.)

REMARK 3.1.10. As we indicate below, the definition just given of $M_\infty(V)$ in terms of M_∞ is historically backwards. Indeed the original Taylor–Wiles method fixed a V and constructed an associated patched ring — denoted R_∞ in the original literature, but which we might denote $R_\infty(V)$ — which in retrospect can be interpreted as that quotient of the ring R_∞ introduced above which acts faithfully on $M_\infty(V)$. In the contexts originally considered by Taylor and Wiles, the patched module $M_\infty(V)$ is cyclic over R_∞ , and hence free over $R_\infty(V)$, and so only appears implicitly (in the use of mod p multiplicity one results that show that certain Hecke modules arising from the cohomology of modular curves are free over the Hecke algebras that act on them). The patched modules $M_\infty(V)$ themselves were introduced (independently) by Diamond [Dia97] and Fujiwara [Fuj06], and it was the work of Kisin [Kis09a] (as interpreted by [EGS15]) that then recast patching in terms of the exact functor $V \mapsto M_\infty(V)$. That this functor is determined by a single “large” patched module M_∞ was observed in [CEG⁺16]. (Note though that the formula (3.1.9) does not literally appear in [CEG⁺16], where instead one finds essentially equivalent (but visually more involved) formulas involving Homs and duals.)

The reader may wish to compare the formula (3.1.9) for the patching functor to Conjecture 6.1.14 (2), which suggests the mnemonic “ $M_\infty = M_\infty(\mathcal{O}[[K]])$ ”. The jump from considering the functor $V \mapsto M_\infty(V)$ to the conjectures of these notes is roughly to move from considering representations of $\mathrm{PGL}_2(\mathbf{Z}_p)$ to representations of $\mathrm{PGL}_2(\mathbf{Q}_p)$, by replacing V by the compact induction $c\text{-Ind}_{\mathrm{PGL}_2(\mathbf{Z}_p)}^{\mathrm{PGL}_2(\mathbf{Q}_p)} V$, and then to move from modules over Galois deformation rings to sheaves on stacks of (φ, Γ) -modules.

Actually, in the context of patching, the passage from $\mathrm{PGL}_2(\mathbf{Z}_p)$ to $\mathrm{PGL}_2(\mathbf{Q}_p)$ was one of the main contributions of [CEG⁺16], and so there is natural progression of ideas from [EGS15] (patching as a functor on $\mathrm{PGL}_2(\mathbf{Z}_p)$ -representations) to [CEG⁺16] (patching as a functor on $\mathrm{PGL}_2(\mathbf{Q}_p)$ -representations) to Expected Theorem 7.2.1 below (patching is promoted to a functor from $\mathrm{PGL}_2(\mathbf{Q}_p)$ -representations to (complexes of) sheaves on a stack of (φ, Γ) -modules).

As already indicated above, the patching method can be extended from the case of modular curves to more general Shimura varieties and related congruence

quotients. However, the case of $\mathrm{PGL}_2(\mathbf{Q}_p)$ considered here is particularly nice, because a p -adic local Langlands correspondence has been constructed in this case (in fact for $\mathrm{GL}_2(\mathbf{Q}_p)$). The idea of a p -adic local Langlands correspondence was motivated by Taylor–Wiles patching, but the correspondence itself was constructed by quite different techniques. We review some of the history in the rest of this section.

3.2. p -adic Langlands: an overview and history. Following preliminary and for the most part unpublished investigations by many researchers (including in approximate historical order Fontaine–Langlands, Serre, Mazur, Harris, Vignéras, and Schneider–Teitelbaum), the investigation of a possible p -adic Langlands correspondence, relating p -adic representations of the group $\mathrm{GL}_2(\mathbf{Q}_p)$ to 2-dimensional p -adic representations of the local Galois group $\mathrm{Gal}_{\mathbf{Q}_p}$, was instigated by Christophe Breuil [Bre03a, Bre03b]. He was motivated to a considerable extent by various considerations that arose out of the proof of modularity lifting theorems [Wil95, TW95, CDT99, BCDT01]; in particular, he was directly motivated by considerations arising from the Breuil–Mézard conjecture (see [Bre08, §1.4] for an account of these motivations). We begin our overview by discussing some ideas and methods related to these theorems, before turning to a discussion of the correspondence itself, its application to the study of the Breuil–Mézard conjecture, and its relationship to completed (co)homology via local-global compatibility.

3.2.1. *Taylor–Wiles–Kisin patching.* Taylor–Wiles patching, as introduced in [TW95] and briefly reviewed above, is a process that constructs modules over local deformation rings out of the (co)homology of Shimura varieties (thought of as Hecke modules).

The patching method was originally developed to prove modularity lifting theorems, showing that certain p -adic Galois representations (in the original applications, those associated to elliptic curves) correspond to modular forms. Originally the patching arguments required as inputs a number of deep results about congruences between modular forms, including mod p multiplicity one theorems, level raising theorems, cases of the weight part of Serre’s conjecture, and Serre’s modularity conjecture itself. Subsequently, the situation has reversed, and such results can now be deduced from patching arguments.

While patching arguments are still used to prove that Galois representations are automorphic, they are now also used to study the internal properties of Galois representations and of congruences between automorphic representations, and it is this application which interests us. We will not attempt to give a full account of these shifts, and will instead discuss the contemporary perspective on the patching construction, but in brief we note that the dependence on mod p multiplicity one theorems was removed independently by Diamond [Dia97] and Fujiwara [Fuj06], level raising was addressed by Taylor’s “Ihara avoidance” argument [Tay08], the weight part of Serre’s conjecture by T.G. [Gee11], and Serre’s conjecture was proved by Khare–Wintenberger [KW09] and Kisin [Kis09b]. Many of these developments, and much of what we explain below, rely crucially on Kisin’s improvement of the Taylor–Wiles method [Kis09c].

We maintain the notation above, but to avoid clutter we write R_p for $R(\bar{\rho})$ from now on. We also go back to considering GL_2 rather than PGL_2 , so that in particular the exact functor $V \mapsto M_\infty(V)$ is a functor from the category of continuous $\mathrm{GL}_2(\mathbf{Z}_p)$ -representations on finitely generated \mathcal{O} -modules to the category

of finitely generated modules over R_∞ , where R_∞ is a power series ring over R_p . There are certain choices of $\mathrm{GL}_2(\mathbf{Z}_p)$ -representation that are of particular interest, namely those that are lattices in a *locally algebraic type*, and those that are irreducible representations of $\mathrm{GL}_2(\mathbf{F}_p)$ over k . Evaluating the patching functor M_∞ on the first kind of representation relates patching to the study of modularity lifting; and this provided the original motivation for the patching construction. Evaluating the patching functor on the second kind of representation establishes a relationship between patching and the weight part of Serre’s conjecture.

To be more precise, recall that by a locally algebraic type σ , we mean an irreducible $\overline{\mathbf{Q}}_p$ -representation of $\mathrm{GL}_2(\mathbf{Z}_p)$ obtained by tensoring an algebraic representation with a finite-dimensional representation corresponding to an inertial type τ via Henniart’s inertial local Langlands correspondence [BM02, App. A]; we say that a lift $\overline{\rho}$ is of type σ if it has Hodge–Tate weights corresponding to the highest weight of the algebraic representation, and inertial type corresponding to τ . Suppose that σ is a locally algebraic type, and let σ° be a $\mathrm{GL}_2(\mathbf{Z}_p)$ -invariant lattice in σ . If $R_p(\sigma)$ denotes the quotient of R_p parameterizing lifts of type σ , then (as a consequence of local-global compatibility) the R_p -action on $M_\infty(\sigma^\circ)$ factors through $R_p(\sigma)$, and a modularity lifting theorem for lifts of type σ (i.e. a theorem proving that lifts of $\overline{\rho}$ of type σ are necessarily modular) can be interpreted in terms of patching as saying that the action of the corresponding quotient $R_\infty(\sigma)$ of R_∞ on $M_\infty(\sigma)$ should be faithful.

As for the weight part of Serre’s conjecture, this can be rephrased as the statement that $M_\infty(\overline{\sigma})$ is non-zero for precisely those Serre weights $\overline{\sigma}$ of $\mathrm{GL}_2(\mathbf{F}_p)$ that lie in a specified set $W(\overline{\rho})$ (where by definition a Serre weight is an irreducible k -representation of $\mathrm{GL}_2(\mathbf{F}_p)$).

3.2.2. An observation. There is an important observation we can make that motivated much of the further development of the subject. If V is a $\mathrm{GL}_2(\mathbf{Z}_p)$ -representation on a finite rank free \mathcal{O} -module (e.g. a lattice in a locally algebraic type), and $\overline{V} := V/\varpi V$ is the associated residual representation, then $M_\infty(V)$ is \mathcal{O} -torsion free (by exactness), and so $M_\infty(V) \neq 0$ iff $M_\infty(\overline{V}) \neq 0$ iff $M_\infty(\overline{\sigma}) \neq 0$ for some Jordan–Hölder factor $\overline{\sigma}$ of \overline{V} . Thus, if the weight part of Serre’s conjecture is known, then the question of whether or not $M_\infty(V)$ is non-zero can be answered via an examination of the set of Jordan–Hölder factors of \overline{V} .

3.2.3. The Breuil–Mézard conjecture. As already noted, in order to prove a modularity lifting theorem for Galois representations of locally algebraic type σ , one must show that $R_\infty(\sigma)$ acts faithfully on $M_\infty(\sigma^\circ)$, for some (equivalently, any) invariant lattice σ° contained in σ . It follows from the patching construction that $M_\infty(\sigma^\circ)$ is maximal Cohen–Macaulay over $R_\infty(\sigma)$; so if $R_p(\sigma)$ is a domain, then $R_\infty(\sigma)$ (which is then also a domain) necessarily acts faithfully on $M_\infty(\sigma^\circ)$. Perhaps the simplest way to establish that $R_p(\sigma)$ is a domain is to show that it is formally smooth, and the arguments of [CDT99] and [BCDT01] were devoted to establishing such a smoothness result (for appropriately chosen σ). In analyzing the conditions introduced in these papers related to ascertaining such results, Breuil and Mézard were led to formulate their eponymous conjecture [BM02]. This conjecture describes the Hilbert–Samuel multiplicity of $R_p(\sigma)/\varpi$ in terms of the Serre weights of $\overline{\rho}$ that appear as constituents of σ°/ϖ .

A key breakthrough in understanding the Breuil–Mézard conjecture was made by Kisin (see e.g. [Kis10b]), who strengthened the observation of (3.2.2) to note

that it follows from the weight part of Serre’s conjecture, the exactness of the patching functor, and the maximal Cohen–Macaulay nature of patched modules, that the formula of the Breuil–Mézard conjecture precisely describes the Hilbert–Samuel multiplicity of R/ϖ , where R denotes that quotient of $R_\infty(\sigma)$ that acts faithfully on $M_\infty(\sigma^\circ)$. Thus the Breuil–Mézard conjecture, when combined with the weight part of Serre’s conjecture, is seen to be essentially equivalent to the faithfulness of the action of $R_p(\sigma)$ on $M_\infty(\sigma^\circ)$, i.e. to a modularity lifting theorem. (See [EG14, §5] for a more precise formulation of this equivalence in a more general setting.)

By combining this observation with an argument using the p -adic Langlands correspondence, Kisin was in fact able to prove the Breuil–Mézard conjecture (and so also to deduce corresponding modularity lifting theorems) [Kis09a]. We recall more of the details of this argument below. For now, we merely remark that, even without knowing whether or not $R_p(\sigma)$ acts faithfully on $M_\infty(\sigma^\circ)$, there is an obvious inequality of Hilbert–Samuel multiplicities $e_{R_p(\sigma)/\varpi} \geq e_{R/\varpi}$, and so Kisin’s observation always yields a weak form of the Breuil–Mézard conjecture, in which equality is replaced by an inequality in one direction.

3.2.4. A locally analytic variant. Above we discussed an approach to studying Serre weights via patching functors. There is a related story in the analytic setup which is discussed in [Bre15] and [BHS19]. The question of which Serre weights are associated to a residual Galois representation $\bar{\rho}$ can be understood as the question of which irreducible $\mathrm{GL}_2(\mathbf{Z}_p)$ -representations (more generally $K = \mathrm{GL}_n(\mathcal{O}_F)$ -representations) embed into the (conjectural) mod p representation $\Pi(\bar{\rho})$. (Note that here we are adopting the more traditional perspective on p -adic and mod p local Langlands; the connection between this perspective and the categorical conjecture emphasised in these notes is discussed in Remark 7.8.5). In some analogy with this question, Breuil [Bre15, Conj. 6.1] has formulated a conjecture about the irreducible constituents of locally analytic principal series representation that embed into the p -adic local Langlands correspondence $\Pi(\rho)$ of a p -adic Galois representation ρ . While the question about Serre weights is motivated by the global question about the weights of modular forms whose associated Galois representation has fixed reduction modulo p , Breuil’s socle conjecture is motivated by the global question about (the weights of) non-classical overconvergent p -adic modular forms with prescribed associated Galois representation.

In [BHS19] many cases of Breuil’s socle conjecture are proven using a similar strategy (patching functors and a multiplicity formula in the spirit of the Breuil–Mézard conjecture) as discussed above: Instead of studying $V \mapsto M_\infty(V)$ as a functor on certain locally algebraic K -representations, one studies a related functor on certain locally analytic representations (or the BGG category \mathcal{O} and a functor from that category to locally analytic representations). In this case the functor produces coherent sheaves on the rigid analytic generic fiber of $\mathrm{Spf}R_\infty$ that are supported on closed subvarieties defined in terms of trianguline Galois representations (see also Section 3.2.6 below for a discussion of trianguline representations). For these trianguline deformation spaces there is a locally analytic analogue of the Breuil–Mézard conjecture formulated in [BHS19, Conj. 4.3.4]. Roughly this conjecture predicts that the multiplicities of irreducible constituents in locally analytic representations that are parabolically induced from a character δ match the multiplicities of certain cycles in trianguline deformation spaces whose parameters

correspond to δ . In fact this conjecture can be proven in many cases [BHS19, Theorem 4.3.8] by relating it to multiplicity formulas in geometric representation theory.

3.2.5. *The p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$.* The p -adic local Langlands correspondence, associating admissible continuous unitary p -adic Banach space representations $\Pi(\rho)$ of $\mathrm{GL}_2(\mathbf{Q}_p)$ to two dimensional continuous representations $\rho : \mathrm{Gal}_{\mathbf{Q}_p} \rightarrow \mathrm{GL}_2(L)$, was first proposed by Breuil for those representations ρ that are potentially semistable with distinct Hodge–Tate weights [Bre03a, Bre03b].

More precisely, Breuil proposed in this context that $\Pi(\rho)$ should be constructed as a certain completion of the locally algebraic representation $\Pi(\rho)_{\mathrm{alg}}$ obtained by tensoring together the smooth representation of $\mathrm{GL}_2(\mathbf{Q}_p)$ attached to $\mathrm{WD}(\rho)$ (the Weil–Deligne representation underlying the Dieudonné module of the potentially semistable representation ρ) with a finite dimensional algebraic representation of $\mathrm{GL}_2(\mathbf{Q}_p)$ encoding the Hodge–Tate weights of ρ . (The representation $\Pi(\rho)_{\mathrm{alg}}$ contains a locally algebraic type σ , and Breuil’s proposal was motivated by the Breuil–Mézard conjecture.) Furthermore, in the case when ρ is crystabelline or semistabelline (i.e. becomes crystalline, or semistable, upon restriction to a finite abelian extension of \mathbf{Q}_p), he made his proposal completely explicit: namely, in the irreducible crystabelline case, he proposed that $\Pi(\rho)$ should be the universal unitary completion of $\Pi(\rho)_{\mathrm{alg}}$, while in the irreducible genuinely semistabelline (i.e. non-crystabelline) case, he proposed that $\Pi(\rho)$ should be a certain explicit completion of $\Pi(\rho)_{\mathrm{alg}}$ depending on the \mathcal{L} -invariant of ρ .

The motivation for proposing the universal unitary completion as the candidate for $\Pi(\rho)$ in the irreducible crystabelline case is that in this case there is no “extra” p -adic Hodge theoretic data carried by ρ besides that already carried by $\Pi(\rho)_{\mathrm{alg}}$ (that is, there is no \mathcal{L} -invariant or the like in this case; more precisely, there is up to isomorphism a unique weakly admissible filtration with fixed Hodge–Tate weights on $D_{\mathrm{cris}}(\rho)$ such that the corresponding Galois representation is irreducible). The representation $\Pi(\rho)_{\mathrm{alg}}$ determines (and is determined by) $D_{\mathrm{cris}}(\rho)$ and the Hodge–Tate weights, and the universal unitary completion is the only evident choice of completion available that doesn’t depend on any additional choices.

The major initial difficulty in investigating Breuil’s conjectured correspondence was in establishing any non-trivial properties of his proposed completions; it was not even evident that these completions were non-zero. One exception to this situation was in the case when ρ was the restriction to $\mathrm{Gal}_{\mathbf{Q}_p}$ of a global two-dimensional representation $r : \mathrm{Gal}_{\mathbf{Q}} \rightarrow \mathrm{GL}_2(L)$ arising from a Hecke eigenform. In this case the theory of completed cohomology, together with local-global compatibility for the classical Langlands correspondence, gives rise to a non-zero unitary completion of $\Pi(\rho)_{\mathrm{alg}}$. In the semistabelline case, Breuil was furthermore able to show, using methods arising from the theory of p -adic L -functions (and we remind the reader that the notion of \mathcal{L} -invariant had its origin in that same theory) that this completion factored through his proposed completion of $\Pi(\rho)_{\mathrm{alg}}$ [Bre10b].

Inspired by the approach of Kato, Perrin-Riou and others to the theory of p -adic L -functions via Galois cohomology, explicit reciprocity laws, and p -adic Hodge theory, and his own reinterpretation and further developments of these ideas in terms of (φ, Γ) -modules, Colmez was then able to make a major breakthrough, and establish, in the irreducible and genuinely semistabelline case, that Breuil’s

proposed completion $\Pi(\rho)$ is non-zero, topologically irreducible, and admissible. (See [Col10a] for the culmination of Colmez’s investigations in this direction.) Soon thereafter, using similar techniques, Berger and Breuil were able to prove the analogous result in the crystabelline case [BB10]. A key consequence of their result is that, for crystabelline ρ , the universal unitary completion $\Pi(\rho)$ of $\Pi(\rho)_{\text{alg}}$ is the unique non-zero unitary completion of $\Pi(\rho)_{\text{alg}}$.

Since the theory of (φ, Γ) -modules applies to arbitrary representations ρ , not just potentially semistable ones, Colmez’s results led him to propose that the p -adic local Langlands correspondence should exist for arbitrary continuous representations $\rho : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(L)$. (In the case when ρ is not potentially semistable with distinct Hodge–Tate weights, one no longer has a representation $\Pi(\rho)_{\text{alg}}$ available from which to even conjecturally construct $\Pi(\rho)$ as some kind of completion.) Reverse-engineering his (φ, Γ) -module approach to studying Breuil’s completions, Colmez discovered the first of his two functors, the functor $\Pi \mapsto V(\Pi)$ mapping topologically finite length admissible unitary continuous Banach $\text{GL}_2(\mathbf{Q}_p)$ -representations (satisfying a finite length condition on their reductions mod ϖ , which was shown by Paškūnas to be superfluous [Paš13]) to continuous representations of $\text{Gal}_{\mathbf{Q}_p}$ [Col10c, §IV].

In the light of his investigations of the Breuil–Mézard conjecture, Kisin then made the crucial suggestion that Colmez’s proposed general p -adic local Langlands correspondence should be regarded as taking place over local Galois deformation spaces, and that deformation-theoretic tools could play a role in its construction. Following this suggestion, Colmez was able to extend his (φ, Γ) -module methods to construct $\Pi(\rho)$ for general ρ [Col10c]. The functors are related via a natural isomorphism $V(\Pi(\rho)) \cong \rho$. Furthermore, he was able to prove that when ρ is potentially semi-stable with distinct Hodge–Tate weights, the representation $\Pi(\rho)$ is indeed a completion of $\Pi(\rho)_{\text{alg}}$, as Breuil had originally proposed. In the cases when ρ is crystabelline or semistabelline, this followed from the results already mentioned due to him and Berger–Breuil (indeed, the general construction of $\Pi(\rho)$ is motivated by the constructions in those cases). However, in general (i.e. in cases that are not crystabelline, or semistabelline, or, more generally, trianguline), the action of the entirety of $\text{GL}_2(\mathbf{Q}_p)$ on $\Pi(\rho)$ is constructed in a very indirect way (it is only the action of the upper triangular Borel subgroup of $\text{GL}_2(\mathbf{Q}_p)$ that can be seen directly in terms of (φ, Γ) -modules) by deformation-theoretic methods (or, one could say, by rigid analytic interpolation), and so the nature of the locally algebraic vectors in $\Pi(\rho)$ is not obvious. Colmez’s proof that they contain (and in fact coincide with) $\Pi(\rho)_{\text{alg}}$ ultimately rests on a comparison with the global context of completed cohomology, mediated via the work of one of us (M.E.) on local-global compatibility.

The theory of the p -adic local Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$ was completed by Paškūnas, who showed that the functor $\Pi \mapsto V(\Pi)$ induces a bijection between isomorphism classes of supersingular topologically irreducible admissible unitary continuous Banach space representations of $\text{GL}_2(\mathbf{Q}_p)$, and irreducible continuous two-dimensional representations of $\text{Gal}_{\mathbf{Q}_p}$ [Paš13]. In fact, more was established: essentially, V induces an equivalence between the category of topologically finite length admissible Banach space representations of $\text{GL}_2(\mathbf{Q}_p)$, and the category whose objects are deformations of two-dimensional continuous representations of $\text{Gal}_{\mathbf{Q}_p}$. Furthermore, this equivalence can be described Morita-theoretically: there

are certain injective objects in the category of $\mathrm{GL}_2(\mathbf{Q}_p)$ -representations, whose endomorphism rings may be identified with Galois deformation rings, that mediate the equivalence. Since these injective objects (or, equivalently, their topological duals, which are then projective objects in an appropriate category) play an important role in the work of Andrea Dotto, M.E. and T.G. which we explain in Sections 7.2–7.4, we will discuss them further below.

3.2.6. *The locally analytic p -adic Langlands correspondence.* As we will discuss further below, the Banach space representations occurring in the p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$ are ultimately constructed by passage to the limit of constructions modulo p^n (and then by inverting p), and in particular they naturally fit into the “Banach” context in these notes.

On the other hand, an important tool for studying Banach space representations such as $\Pi(\rho)$ is to study their space of analytic vectors $\Pi(\rho)^{\mathrm{an}}$ (in the sense of [ST03, §7]), which are locally analytic representations of $\mathrm{GL}_2(\mathbf{Q}_p)$. This goes in the opposite direction to the consideration of mod p^n representations, since $\Pi(\rho)^{\mathrm{an}}$ forgets about the Banach space (and hence p -adic integral) structure on $\Pi(\rho)$. On the other hand, Colmez’s construction shows that the locally analytic representation $\Pi(\rho)^{\mathrm{an}}$ can be constructed directly from the rigid analytic (φ, Γ) -module $D_{\mathrm{rig}}(\rho)$ over the Robba ring \mathcal{R} associated to ρ , so that it still makes sense to study the p -adic local Langlands correspondence in this context. Motivated by this observation Colmez [Col16] extended the p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$ to the setup of (possibly non-étale) rank two (φ, Γ) -modules over the Robba ring and locally analytic representations of $\mathrm{GL}_2(\mathbf{Q}_p)$. This naturally fits into the “analytic” context in these notes.

We point out that usually p -adic Hodge-theoretic properties of the representation ρ are not so easy to read off from the associated (φ, Γ) -module $D(\rho)$, but it is much easier to read them off from the rigid analytic variant $D_{\mathrm{rig}}(\rho)$. Hence Hodge-theoretic information, e.g. the Hodge filtration on $D_{\mathrm{dR}}(\rho)$ (assuming that ρ is de Rham) should be connected to properties of the locally analytic representation $\Pi(\rho)^{\mathrm{an}}$ (and it should not be easy to understand this directly in terms of the Banach space representation $\Pi(\rho)$). We now elaborate on this point.

One of the original motivations in the p -adic Langlands program was to improve the understanding of local-global compatibility in the cohomology of modular curves (or more general Shimura varieties). To each system of Hecke eigenvalues that occurs in the cohomology $H^1(Y(N), \mathbf{Q}_p)$ of the modular curve $Y(N)$ there corresponds a continuous 2-dimensional representation ρ of $\mathrm{Gal}_{\mathbf{Q}}$ on a \mathbf{Q}_p vector space. The restriction $\rho|_{\mathrm{Gal}_{\mathbf{Q}_\ell}}$ to decomposition groups at primes $\ell \neq p$ is completely determined (at least up to Frobenius semisimplicity, which is expected to be automatic) by the eigenvalues of the Hecke operators at ℓ . However the action of the Hecke operators at p only determines the Weil–Deligne representation associated to the de Rham representation $\rho|_{\mathrm{Gal}_{\mathbf{Q}_p}}$. The missing piece to determine $\rho|_{\mathrm{Gal}_{\mathbf{Q}_p}}$ is the Hodge filtration which is mysteriously determined by the Hecke action away from p but not captured by the Hecke action at p . The idea is that p -adic local global compatibility involving completed (co-)homology and the p -adic local Langlands correspondence will make it possible to read off the Hodge filtration from the $\mathrm{GL}_2(\mathbf{Q}_p)$ -representation in the corresponding system of Hecke eigenvalues on completed (co-)homology $\tilde{H}^1(Y(N), \mathbf{Q}_p)$ without using the action of Hecke

operators away from p . In fact, the Hodge filtration should be determined by the locally analytic vectors in $\widetilde{H}^1(Y(N), \mathbf{Q}_p)$. A generalization of this philosophy, which goes back to Breuil's initial ideas about the p -adic local Langlands correspondence, is visible for example in Breuil's socle conjecture (see Section 9.6.29), where the relative position of the Hodge filtration with respect to Frobenius stable flags is determined by the socle of a certain locally analytic representation. We will discuss the p -adic local global compatibility involving completed (co-)homology of modular curves in more detail in Section 3.2.11 below.

Passing from a (classical) (φ, Γ) -module D to its variant D_{rig} over the Robba ring \mathcal{R} has the advantage that the (φ, Γ) -module can become reducible, even though D was not reducible. For example this happens if D comes from a crystalline or semi-stable (or more generally crystabelline or semistabelline) representation. The (φ, Γ) -modules that are completely reducible are called trianguline and were studied extensively by Colmez [Col08]. In terms of the p -adic local Langlands correspondence a representation ρ is trianguline (i.e. $D_{\text{rig}}(\rho)$ is trianguline) if and only if the locally analytic vectors $\Pi(\rho)^{\text{an}}$ contain (irreducible constituents of) locally analytic principal series representations. In this case usually $\Pi(\rho)^{\text{an}}$ is an extension of two parabolically induced representations, and the extension class of these representations is determined by the extension class of (φ, Γ) -modules of rank 1 defined by $D_{\text{rig}}(D)$ (and vice versa). This implies in particular that in the case of crystalline or semi-stable representations ρ the Hodge filtration on the Weil–Deligne representation $\text{WD}(\rho)$ is determined by the locally analytic vectors $\Pi(\rho)^{\text{an}}$, as it should be according to the above discussion.

3.2.7. *The mod p local Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$.* In our discussion so far, we have set the p -adic local Langlands correspondence in the context of p -adic representations of $\text{GL}_2(\mathbf{Q}_p)$. However, from the beginning of the theory, mod p considerations have played a fundamental role. In his paper [Bre03a] (and building on earlier work of Barthel and Livné [BL94]), Breuil classified the absolutely irreducible admissible smooth representations of $\text{GL}_2(\mathbf{Q}_p)$ over a finite field k , and showed *grosso modo* that they are of the form $(c\text{-Ind}_{\text{GL}_2(\mathbf{Z}_p)_Z}^{\text{GL}_2(\mathbf{Q}_p)} \bar{\sigma}) \otimes_{\mathcal{H}(\bar{\sigma}), x} k$, where Z denotes the centre of $\text{GL}_2(\mathbf{Q}_p)$, $\mathcal{H}(\bar{\sigma}) = \text{End}((c\text{-Ind}_{\text{GL}_2(\mathbf{Z}_p)_Z}^{\text{GL}_2(\mathbf{Q}_p)} \bar{\sigma}))$ is the relevant Hecke algebra, which admits an explicit isomorphism $\mathcal{H}(\bar{\sigma}) \cong k[S^{\pm 1}, T]$, and where $x : \mathcal{H}(\bar{\sigma}) \rightarrow k$ is some given character of $\mathcal{H}(\bar{\sigma})$. The key contribution of [Bre03a] is to show that such representations are irreducible in the so-called *supersingular* case, i.e. the case when $x(T) = 0$ (and it is this statement that is special to the context of $\text{GL}_2(\mathbf{Q}_p)$).

Building on this classification, Breuil was able to define a *mod p semisimple local Langlands correspondence*, by attaching to each semisimple two-dimensional representation $\bar{\rho} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(k)$ a corresponding semisimple admissible smooth representation $\kappa(\bar{\rho})$ of $\text{GL}_2(\mathbf{Q}_p)$ over k . Breuil further conjectured that this correspondence was compatible with the conjectured p -adic local Langlands correspondence, in the sense that if $\rho : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(L)$ is continuous, then one should have $\overline{\Pi(\rho)}^{\text{ss}} \cong \kappa(\bar{\rho}^{\text{ss}})$. (Here, on either side of the conjectured isomorphism, the notation $\overline{(-)}^{\text{ss}}$ indicates that we reduce an invariant lattice modulo the uniformizer ϖ , and then semisimplify. Note also that on either side, the reductions are only uniquely determined after semi-simplifying, and so it is natural to pass to semi-simplifications when making such a comparison between the p -adic and mod p settings.)

Subsequently, Berger [Ber10] was able to obtain a description of Breuil’s mod p correspondence in terms of (φ, Γ) -modules, and using this, was able to verify Breuil’s conjectured compatibility in the context of the paper [BB10]. The same (φ, Γ) -module arguments apply to show this compatibility naturally extends to the general p -adic local Langlands correspondence of [Col10c].

Colmez was subsequently able to define a fully fledged *mod p local Langlands correspondence*, i.e. to define the correspondence in such a fashion that, in the case when $\bar{\rho}$ is reducible but indecomposable, it takes into account the corresponding extension class of characters. More precisely, for each representation $\bar{\rho} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(k)$, one may define an associated representation $\kappa(\bar{\rho})$ of $\text{GL}_2(\mathbf{Q}_p)$, with the property that $V(\kappa(\bar{\rho})) = \bar{\rho}$. If $\bar{\rho}$ is *not* a twist of an extension of the mod p cyclotomic character by the trivial character, then $\kappa(\bar{\rho})$ can in fact be characterized by this last property, together with the condition that it admits no finite-dimensional $\text{GL}_2(\mathbf{Q}_p)$ -invariant subobject or quotient. (See e.g. [Eme11a, Thm. 3.3.2].)

In the exceptional case, when $\bar{\rho}$ is a twist of an extension of the mod p cyclotomic character by the trivial character, then the “correct” definition of $\kappa(\bar{\rho})$ (by which we mean the one that is compatible with correspondence that we explain in Section 7.2) is *not* compatible with semi-simplification. In fact, even in this case, it is possible to alter the definition of $\kappa(\bar{\rho})$ so that it is so compatible, and Colmez did indeed adopt this alternative definition; but this alternative choice of $\kappa(\bar{\rho})$ fails to satisfy compatibility with deformations, in the form discussed in more detail below. (See [CEG⁺18, Rem. 6.22] for a more precise description of the difference between Colmez’s definition and ours.)

3.2.8. *p -adic local Langlands over deformation space and proofs of the Breuil–Mézard conjecture.* As already remarked, the construction of the p -adic local Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$ is intimately related to describing the variation of the correspondence over deformation spaces of Galois representations, and we now turn to discussing the various forms this description has taken in the literature. The first key input to all of these descriptions is that Colmez’s functor V is well-defined on the mod p and p -adic integral level, so that it can be utilized to study the p -adic local Langlands correspondence in a deformation-theoretic context.

In [Eme11a], a deformation-theoretic description of the correspondence is given for deformations of a representation $\bar{\rho} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(k)$ that is *not* a twist of an extension of cyclotomic character by the trivial character; in what follows we assume for simplicity that $\bar{\rho}$ also admits only trivial endomorphisms, so that we can and do take R_p to be the universal deformation \mathcal{O} -algebra of $\bar{\rho}$ with some fixed determinant (rather than having to work with framed deformations). For simplicity of exposition, we furthermore assume that this fixed determinant is ε^{-1} ; correspondingly, we will assume that all of the representations of GL_2 that we consider have trivial central character, and accordingly we will regard them as representations of PGL_2 .

In this case, the arguments of Colmez [Col10c] and Kisin [Kis10a] can be encapsulated as follows: there exists a representation π^{univ} of $\text{PGL}_2(\mathbf{Q}_p)$ on an \mathfrak{m}_{R_p} -adically complete and orthonormalizable R_p -module, deforming the smooth representation $\kappa(\bar{\rho})$, with the property that $V(\pi^{\text{univ}}) \cong \rho^{\text{univ}}$, the universal deformation of $\bar{\rho}$. The $\text{PGL}_2(\mathbf{Q}_p)$ -representation attached by the p -adic local Langlands

correspondence to any particular deformation of $\bar{\rho}$ is then obtained by specializing π^{univ} over this deformation.

In [Paš13], a deformation-theoretic description of the correspondence is given that is more flexible and more conceptual than the formulation just recalled. In order to make the description as simple as possible, we assume that $\bar{\rho}$ is *not* a twist of an extension of the trivial character by the cyclotomic character. (Note that this is different to the exceptional case considered above; the forbidden extension has the characters in the opposite order.) We then let \tilde{P} denote a projective envelope of the Pontryagin dual $\kappa(\bar{\rho})^\vee$ in a suitable category of representations of $\text{PGL}_2(\mathbf{Q}_p)$ on compact \mathcal{O} -modules. One of the main results of [Paš13] is that there is a canonical isomorphism

$$(3.2.9) \quad R_p \cong \text{End}(\tilde{P}),$$

so that \tilde{P} may be naturally regarded as “living over” $\text{Spf } R_p$, and that the consequent R_p -structure on $V(\tilde{P}^d)$ realizes $V(\tilde{P}^d)$ as the universal deformation r^{univ} of $\bar{\rho}$. If $\rho : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(L)$ admits a lattice deforming $\bar{\rho}$, corresponding to a morphism $x : R_p \rightarrow \mathcal{O}$, then the unitary Banach space representation $\Pi(\rho)$ is determined via the isomorphism

$$\Pi(\rho) \cong L \otimes_{\mathcal{O}} (\tilde{P} \otimes_{R_p, x} \mathcal{O})^d,$$

where the superscript d denotes Schikhof duality on pseudocompact \mathcal{O} -torsion free \mathcal{O} -modules, i.e. $\text{Hom}_{\mathcal{O}}^{\text{cont}}(-, \mathcal{O})$.

If $\bar{\rho}$ is not a twist of an extension of the cyclotomic character by the trivial character (so that we are in neither of the exceptional cases considered above), then it is furthermore proved in [Paš13] that \tilde{P} is topologically flat over R_p (in the sense that the functor $-\hat{\otimes}_{R_p} \tilde{P}$ on pseudo-compact R_p -modules is exact). In this case, the representation π^{univ} associated to $\bar{\rho}$ as above is related to \tilde{P} via the isomorphism (of $R_p[\text{PGL}_2(\mathbf{Q}_p)]$ -modules)

$$\pi^{\text{univ}} \cong \text{Hom}_{R_p}^{\text{cont}}(\tilde{P}, R_p).$$

In [Kis09a], Kisin gives quite a different description of the p -adic local Langlands correspondence, not over all of $\text{Spf } R_p$, but over the locus $\text{Spf } R_p(\sigma)$ parameterizing lifts of a fixed locally algebraic type σ . For deformations ρ of $\bar{\rho}$ that are parameterized by $\text{Spf } R_p(\sigma)$, Colmez’s proof of Breuil’s conjecture describing $\Pi(\rho)$ as a completion of the locally algebraic representation $\Pi(\rho)_{\text{alg}}$ shows that the restriction to $\text{Spf } R_p(\sigma)[1/p]$ of the p -adic local Langlands correspondence, which we denote by $\pi(\sigma)$, may be described as a certain completion of $c\text{-Ind}_{\text{PGL}_2(\mathbf{Z}_p)}^{\text{PGL}_2(\mathbf{Q}_p)} \sigma$. This description of $\pi(\sigma)$ allows Kisin to deduce the Breuil–Mézard conjecture for $\bar{\rho}$, as we now recall.

A choice of Jordan–Hölder filtration in the reduction $\bar{\sigma}$ of some chosen $\text{PGL}_2(\mathbf{Z}_p)$ -invariant lattice σ° in σ determines a corresponding filtration on $\overline{\pi(\sigma)}$ (the reduction mod ϖ of some invariant lattice in $\pi(\sigma)$) by completions of representations of the form $c\text{-Ind}_{\text{PGL}_2(\mathbf{Z}_p)}^{\text{PGL}_2(\mathbf{Q}_p)} \bar{\sigma}_i$, where $\bar{\sigma}_i$ runs through the Jordan–Hölder factors of $\bar{\sigma}$. For each $\bar{\sigma}_i$, we let σ_i be an algebraic L -representation of $\text{PGL}_2(\mathbf{Q}_p)$ with $\sigma_i^\circ \otimes_{\mathcal{O}} k \cong \bar{\sigma}_i$, and set $R_p(\bar{\sigma}_i) = R_p(\sigma_i) \otimes_{\mathcal{O}} k$ (which is easily checked to be independent of the choice of σ_i). Now the explicit classification of the irreducible representations of $\text{PGL}_2(\mathbf{Q}_p)$, and the fact that Colmez’s functor V is compatible with reduction,

shows that the only $\bar{\sigma}_i$ that can contribute a non-zero completion are those that appear as Serre weights of $\bar{\rho}$, and that if a given Serre weight $\bar{\sigma}_i$ contributes a completion, this completion must be the \mathfrak{m}_x -adic completion of $c\text{-Ind}_{\text{PGL}_2(\mathbf{Z}_p)}^{\text{PGL}_2(\mathbf{Q}_p)} \bar{\sigma}_i$, where \mathfrak{m}_x is the kernel of the character $x : \mathcal{H}(\bar{\sigma}_i) \rightarrow k$ for which $(c\text{-Ind}_{\text{PGL}_2(\mathbf{Z}_p)}^{\text{PGL}_2(\mathbf{Q}_p)} \bar{\sigma}_i) \otimes_{\mathcal{H}(\bar{\sigma}_i), x} k$ is a constituent of $\kappa(\bar{\rho})$. Furthermore, this completion, when thought of as an $R_p(\sigma)$ -module, in fact is supported on $R_p(\bar{\sigma}_i)$, the \mathfrak{m}_x -adic completion of $R_p(\bar{\sigma}_i)$ being explicitly identified with the \mathfrak{m}_x -adic completion of $\mathcal{H}(\bar{\sigma}_i)$. In this way, one obtains an inequality of Breuil–Mézard type, but in the opposite direction to the one given by the argument described in (3.2.3) above. Combining these two inequalities gives an equality, which is the Breuil–Mézard conjecture. As noted in (3.2.3), this conjecture then implies that each patched module $M_\infty(\sigma^\circ)$ is faithful over the ring $R_p(\sigma)$, giving rise to a modularity lifting theorem for global deformations of type σ . This is the approach to the Fontaine–Mazur conjecture taken in [Kis09a].

In the argument that is sketched in the preceding paragraph, the reason that only an inequality (rather than the precise equality of the Breuil–Mézard conjecture) is obtained is more-or-less because completions (being a form of projective limit) are not *a priori* right exact, so that it is not clear *a priori* that all the possible completions of the representations $c\text{-Ind}_{\text{PGL}_2(\mathbf{Z}_p)}^{\text{PGL}_2(\mathbf{Q}_p)} \bar{\sigma}_i$ that might occur in $\bar{\pi}(\bar{\sigma})$ actually do so occur.

Issues of exactness of this kind are ultimately of a homological nature, and so it is perhaps not surprising that having stronger homological information available allows the argument to be strengthened. This was achieved by Paškūnas in the paper [Paš15]. Using the construction of \tilde{P} as described above, together with Colmez’s result on locally algebraic vectors, one finds that $R_p(\sigma)$ is precisely the support of the R_p -module $\text{Hom}_{\text{PGL}_2(\mathbf{Z}_p)}(\sigma^\circ, \tilde{P}^\vee)$ (for any choice of invariant lattice σ° in σ). In fact, it is more convenient to work with the R_p -module

$$(3.2.10) \quad M(\sigma^\circ) := \tilde{P} \otimes_{\mathcal{O}[[\text{PGL}_2(\mathbf{Z}_p)]]} \sigma^\circ,$$

since this turns out to be a coherent R_p -module. Since \tilde{P} is projective as a $\text{PGL}_2(\mathbf{Q}_p)$ -representation, it is in particular projective as a $\text{PGL}_2(\mathbf{Z}_p)$ -representation, and so the functor $M(-)$ implicitly defined by the preceding formula is *exact*. Arguing with this functor one can essentially combine the two prongs of Kisin’s argument into a single argument that yields the Breuil–Mézard equality directly.

3.2.11. *Local-global compatibility.* In [Eme11a] a local-global compatibility relating the p -adic local Langlands correspondence to completed cohomology is proved, giving a p -adic analogue of the classical local-global compatibility [Lan73, Del73, Car83] relating cohomology at infinite level and the classical local Langlands correspondence. In order to state the result, we fix as above a modular irreducible residual representation $\bar{r} : \text{Gal}_{\mathbf{Q}} \rightarrow \text{GL}_2(k)$, setting $\bar{\rho} := \bar{r}|_{\text{Gal}_{\mathbf{Q}_p}}$, and work at a tame level N divisible by the prime-to- p Artin conductor of \bar{r} .

Similar to the completed homology (3.1.5) we let $\tilde{H}^1 := \tilde{H}_{\text{ét}}^1(Y_1(N), \mathcal{O})$ denote the p -adically completed étale cohomology of the “modular curves” (again: more precisely, congruence quotients of PGL_2 -symmetric spaces, which are unions of the connected components of the usual modular curves for GL_2) of tame level N ; by definition we have

$$\tilde{H}^1 := \varprojlim_s \varinjlim_n H_{\text{ét}}^1(Y_1(Np^n), \mathcal{O}/\varpi^s),$$

and \tilde{H}^1 is naturally equipped with commuting actions of the groups $\mathrm{Gal}_{\mathbf{Q}}$ and $\mathrm{PGL}_2(\mathbf{Q}_p)$, and of a Hecke algebra \mathbf{T} (generated by the usual Hecke operators S_l, T_l at primes $l \nmid Np$). The representation \bar{r} gives rise to a maximal ideal $\mathfrak{m}(\bar{r})$ of \mathbf{T} , and we may form the corresponding localisation $\tilde{H}_{\mathfrak{m}(\bar{r})}^1$ of \tilde{H}^1 . There is also a universal pro-modular deformation ρ^{univ} of \bar{r} over $\mathbf{T}_{\mathfrak{m}(\bar{r})}$, and if we set $\bar{\rho} := \bar{r}|_{\mathrm{Gal}_{\mathbf{Q}_p}}$, then restriction from $\mathrm{Gal}_{\mathbf{Q}}$ to $\mathrm{Gal}_{\mathbf{Q}_p}$ induces a morphism $R_p \rightarrow \mathbf{T}_{\mathfrak{m}(\bar{r})}$.

Assume now that $\bar{\rho}$ is not a twist of an extension of the cyclotomic character by the trivial character; the main result of [Eme11a] then states that there is an isomorphism (of $\mathbf{T}_{\mathfrak{m}(\bar{r})}[\mathrm{Gal}_{\mathbf{Q}} \times \mathrm{PGL}_2(\mathbf{Q}_p)]$ -modules)

$$\tilde{H}_{\mathfrak{m}(\bar{r})}^1 \cong (\rho^{\mathrm{univ}} \otimes_{R_p} \pi^{\mathrm{univ}}) \hat{\otimes}_{\mathbf{T}_{\mathfrak{m}(\bar{r})}} O,$$

where π^{univ} is the orthonormalizable $\mathrm{PGL}_2(\mathbf{Q}_p)$ -representation over R_p discussed in Section 3.2.8 above, and where O is a p -torsion free cofinitely generated $\mathbf{T}_{\mathfrak{m}(\bar{r})}$ -module⁵ of “oldforms” related to the chosen tame level.⁶ If $\bar{\rho}$ is furthermore not a twist of an extension of the trivial character by the cyclotomic character, then passing to continuous R_p -duals, and taking into account the relationship between π^{univ} and \tilde{P} in this case, we may rephrase this as an isomorphism

$$(3.2.12) \quad (\tilde{H}_1)_{\mathfrak{m}(\bar{r})} \cong \tilde{P} \hat{\otimes}_{R_p} (\rho^{\mathrm{univ}})^* \otimes_{\mathbf{T}_{\mathfrak{m}(\bar{r})}} O^*,$$

where $\tilde{H}_1 = \mathrm{Hom}_{\mathcal{O}}(\tilde{H}^1, \mathcal{O})$ denotes the completed homology of the modular curves as in (3.1.5) at tame level N , and O^* (the \mathcal{O} -dual of O) is now a finitely generated $\mathbf{T}_{\mathfrak{m}(\bar{r})}$ -module. The proof of the isomorphism is a combination of a Morita-theoretic argument, similar to the more sophisticated such arguments that are developed in [Paš13], and an interpolation argument, related to the “capture” arguments that appear in [CDP14]. The key input is a combination of classical local-global compatibility and the result of Berger–Breuil (that $\Pi(\rho)$ is the unique unitary completion of $\Pi(\rho)_{\mathrm{alg}}$ when ρ is crystabelline): classical local-global compatibility gives rise to copies of $\Pi(\rho)_{\mathrm{alg}}$ inside \tilde{H}^1 , and then the result of Berger–Breuil shows that the closures of these in \tilde{H}^1 are necessarily isomorphic to $\Pi(\rho)$.

Local-global compatibility gives another approach to the Fontaine–Mazur conjecture. Namely, if $x : \mathbf{T}_{\mathfrak{m}(\bar{r})} \rightarrow L$ is a homomorphism corresponding to a pro-modular deformation r of \bar{r} , with kernel \mathfrak{p}_x , and if $\rho := r|_{\mathrm{Gal}_{\mathbf{Q}_p}}$, then passing to \mathfrak{p}_x -torsion in the local-global compatibility isomorphism (3.2.12) yields an isomorphism

$$(3.2.13) \quad \tilde{H}^1[\mathfrak{p}_x] \otimes_{\mathcal{O}} E \cong r \otimes \Pi(\rho).$$

If ρ is potentially semi-stable with distinct Hodge–Tate weights, so that by Colmez’s results $\Pi(\rho)_{\mathrm{alg}} \neq 0$, then one concludes that the system of Hecke eigenvalues given by x appears in $\tilde{H}_{\mathrm{alg}}^1$. The comparison between locally algebraic vectors in completed cohomology and the classical cohomology of local systems on modular curves then shows that x , and thus r , arises from a classical modular form.

⁵I.e. the \mathcal{O} -dual of a finitely generated $\mathbf{T}_{\mathfrak{m}(\bar{r})}$ -module.

⁶If we work at full level N , and then pass to a colimit over all levels N , the resulting limiting space of oldforms acquires an action of $\mathrm{GL}_2(\mathbf{A}_f^p)$, and can be described in terms of the local Langlands correspondence in families of [EH14] and [HM18]. This is related to the $\ell \neq p$ case of categorical local Langlands that we discuss briefly in Section 8, and which plays a key role in the global conjecture (Conjecture 9.3.2 below). For simplicity of exposition, though, we don’t say anything more about it here.

The Fontaine–Mazur conjecture for deformations of $\bar{\rho}$ therefore reduces to the problem of showing that all such deformations are pro-modular. Under standard hypotheses on $\bar{\rho}$ needed for the Taylor–Wiles method, this can be accomplished by proving a “big $R = \mathbf{T}$ ” theorem, identifying $\mathbf{T}_{\mathfrak{m}(\bar{\rho})}$ with the universal deformation \mathcal{O} -algebra for deformations of $\bar{\rho}$ which are minimally ramified away from p . Such theorems were first proved in [Böc01], using the infinite fern of Gouvêa and Mazur together with modularity lifting theorems in small weight (where the deformation rings $R_p(\sigma)$ can be described explicitly via Fontaine–Laffaille theory). Moreover, we note that (3.2.13) and the construction of the p -adic Langlands correspondence imply that $\tilde{H}^1[\mathfrak{p}_x] \otimes_{\mathcal{O}} L$ is non-zero for every $x : \mathbf{T}_{\mathfrak{m}(\bar{\rho})} \rightarrow L$. This is in fact a stronger statement than the *pro-modularity result* of the “big $R = \mathbf{T}$ theorem”.

The results of [Eme11a] therefore give an alternative proof of the results of [Kis09a] (with slightly different hypotheses); while both proofs rely on Taylor–Wiles–Kisin patching and the p -adic Langlands correspondence, they differ in that the approach of [Kis09a] is to use the p -adic Langlands correspondence to prove enough about the ring $R_p(\sigma)$ to deduce that the patched $R_{\infty}(\sigma)$ -module $M_{\infty}(\sigma^{\circ})$ is faithful, while the approach of [Eme11a] can be viewed as computing the functor M_{∞} in terms of the p -adic Langlands correspondence.

The rephrasing (3.2.12) of the main result of [Eme11a] in terms of completed homology is highly suggestive of two things. Firstly, from the optic of these notes, where we replace Galois deformation rings with moduli stacks of L -parameters, it suggests that we can compute the completed cohomology of modular curves by pulling back some canonical sheaf of $\mathrm{PGL}_2(\mathbf{Q}_p)$ -representations from a stack of local parameters to a stack of global parameters; we pursue this suggestion in Section 9. Secondly, it suggests the possibility of using the Taylor–Wiles patching method to extract \tilde{P} from the completed homology \tilde{H}_1 ; this possibility was realized in [CEG⁺18]. In particular, one finds that the patching functor M_{∞} “is” the functor M implicitly described by (3.2.10). (Note, though, that it seems hopeless to construct the sheaf L_{∞} of Section 7.2 in this way, or from any other global construction.)

3.2.14. *Comparison to the proof of classical local Langlands.* Our construction of the p -adic local Langlands correspondence via global means may be compared with the construction of the classical local Langlands correspondence, which also proceeds by realising a local representation globally [HT01, Hen00, Sch13]. However, as discussed in [CEG⁺16, §1.3], there are some significant differences between our construction and the classical one. Classically, one first uses the classification of Weil–Deligne representations, and the parallel classification of irreducible smooth representations of $\mathrm{GL}_n(\mathbf{Q}_p)$ [BZ77, Zel80] to reduce to constructing a bijection between irreducible n -dimensional Weil–Deligne representations and cuspidal representations of $\mathrm{GL}_n(\mathbf{Q}_p)$. A trace formula argument allows one to realize any cuspidal representation (up to twist) as the local component of a cuspidal automorphic representation, and the associated Weil–Deligne representation is then constructed in the cohomology of a Shimura variety (in particular, this construction *presumes* that local-global compatibility will hold, just as that of [CEG⁺18] does). In order to verify that the correspondence so constructed is in fact purely local, one ultimately relies on the prescription for the local Langlands correspondence in terms of L - and ε -factors.

In the p -adic context, most of these steps aren't available (at least at present). For example, the only known classification of admissible unitary Banach representations of $\mathrm{GL}_2(\mathbf{Q}_p)$ is that of [Paš13], which is *in terms of* the p -adic local Langlands correspondence, rather than being independent of it. Furthermore, we can't expect to realize a local p -adic Galois representation globally, even if we require it to be potentially semistable of some prescribed weights and type: the family of such representations is one-dimensional (thus uncountable), while there are only countably many modular representations of $\mathrm{Gal}_{\mathbf{Q}}$. This is why Taylor–Wiles patching is so crucial to the approach of [CEG⁺16, CEG⁺18]: it provides a mechanism for relating the local and the global contexts in the p -adic setting, in the absence of the Bernstein–Zelevinsky classification/trace formula argument that relates the local and global contexts in the classical setting.

Perhaps the biggest difference between the p -adic and classical contexts is that, in the p -adic context, we don't have any *a priori* prescription for the correspondence in terms analogous to the classical prescription in terms of L - and ε -factors; the closest analogue is the requirement that there should be a compatibility between the p -adic and classical correspondence mediated by the passage to locally algebraic vectors (which should parallel the passage from a potentially semistable Dieudonné module to the underlying Weil–Deligne representation on the Galois side). However, it is only in the principal series case, for the group $\mathrm{GL}_2(\mathbf{Q}_p)$, that this requirement has the possibility of providing enough information to determine the p -adic correspondence. This is why the consideration of the (unramified) principal series case is so crucial in all known proofs of a p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$.

3.2.15. *Patching and the analytic setup.* In a series of papers [BHS17b], [BHS17a] and [BHS19] it was shown that the patching construction can be used not only to prove results about automorphic forms and Galois representations (for example, to prove the existence of an automorphic form whose associated Galois representations has prescribed local properties), but also to prove results about overconvergent p -adic automorphic forms and their associated Galois representations.

A key observation in the theory of overconvergent p -adic automorphic forms is that the Hecke eigenvalues of overconvergent p -adic automorphic eigenforms nicely vary in families and form a so-called eigenvariety Y . The p -adic automorphic forms (of which the eigenvariety parameterizes the Hecke eigenvalues) themselves form a coherent sheaf \mathcal{M} on this eigenvariety. In particular we note that in the theory of eigenvarieties, spaces of (overconvergent finite slope) p -adic automorphic forms can be described as the (dual of) global sections of a coherent sheaf on a rigid analytic space, that is closely related to the generic fiber of a deformation space of Galois representations.

Characterizing the support Y of this coherent sheaf \mathcal{M} purely in terms of Galois representations can be regarded as an *overconvergent version* of the Fontaine–Mazur conjecture: it amounts to giving a necessary and sufficient condition on a Galois representation ρ that ensures that ρ is associated to a overconvergent p -adic automorphic form of finite slope. Kisin [Kis03] constructed a rigid analytic subvariety (the *finite slope space* of loc. cit.) of (the product of \mathbf{G}_m with) a deformation space of 2-dimensional $\mathrm{Gal}_{\mathbf{Q}}$ -representations that is defined in purely

Galois-theoretic terms (asking that the Galois representations are trianguline⁷ at p and unramified almost everywhere) and that conjecturally equals the eigencurve of Coleman–Mazur. This construction in fact inspired the definition of the trianguline variety in [BHS17b] and the following overconvergent version of the Fontaine–Mazur conjecture (that we state slightly informally, and which was first proposed by Kisin in the case $d = 2$; see [Kis03, 11.7(2), 11.8]): a p -adic d -dimensional odd global Galois representation is associated to an overconvergent p -adic automorphic form of finite slope on GL_d if and only if it is unramified almost everywhere and trianguline at places dividing p . (In the case $d = 2$, this conjecture was proved under auxiliary technical hypotheses in [Eme11a, Thm. 1.2.4]; a slightly weaker version, under slightly different hypotheses, is established in [Kis09a, Cor., p. 642].)

However, these constructions and conjectures only give a characterization of the support Y of the coherent sheaf \mathcal{M} . We refer to Section 9.6 for a precise conjecture characterizing the coherent sheaf \mathcal{M} purely in terms of local Galois representations and a global-to-local map of stacks of Galois representations. The idea that a characterization like this should exist originates in the insight that the patching construction can also be carried out for eigenvarieties.

More precisely, in a similar way as M_∞ over R_∞ (and $M_\infty(\sigma)$ over $R_\infty(\sigma)$) is related (via pullback along a local-global map) to a space of automorphic forms over a global Galois deformation ring, there is (in the setup of a definite unitary group G over a CM field F satisfying some assumptions necessary for the patching construction) a *patched* version [BHS17b] of the eigenvariety $X_p(\bar{\rho})$ and a coherent sheaf \mathcal{M}_∞ on $X_p(\bar{\rho})$ such that \mathcal{M} and Y may be reconstructed from \mathcal{M}_∞ and $X_p(\bar{\rho})$ via pullback along a local-global map. The main idea of the construction is to mimic the construction of eigenvarieties of M.E. [Eme06b] (using a locally analytic Jacquet functor), but using a big patched module M_∞ (respectively its dual Π_∞ , which is a big Banach space representation of $G(F_p)$), instead of the space of p -adic automorphic forms.

The key idea of [BHS17a] and [BHS19] is to use Galois-theoretic tools (i.e. the geometry of the space $\mathfrak{X}_{G,\mathrm{tri}}$ from Section 5.3) to analyze the local geometry of the space $X_p(\bar{\rho})$, which in turn gives information about the behavior of the sheaf \mathcal{M}_∞ .

4. Moduli stacks of (φ, Γ) -modules: the Banach case

We now review the (sometimes conjectural) constructions of stacks of (φ, Γ) -modules that we will work with. We work in two different settings: in this section we cover the “Banach” case treated in [EG22], which (conjecturally) relates to smooth p -adic representations of $G(\mathbf{Q}_p)$. In Section 5 we discuss the “analytic” case, relating to locally analytic representations of $G(\mathbf{Q}_p)$.

In this section we very briefly recall some of the main results of [EG22]. The reader may find it useful to refer to [EG20] for a more extended introduction to and overview of [EG22], and to Sections 7.2 and 7.6 below for a more explicit description of the stacks in some simple cases.

The stacks that we consider below are formal algebraic stacks in the sense of [Eme]; colloquially, they are “the formal analogue of Artin stacks”, in the sense that they admit smooth covers by formal schemes (or more precisely by formal algebraic spaces). We note that they are in general not p -adically formal; equivalently,

⁷The name *trianguline* was only invented later by Colmez using a different but equivalent definition as the one in Kisin’s paper.

their special (i.e. mod p) fibres are genuinely formal, rather than algebraic. We refer the reader to [EG22, App. A] for a basic overview of formal algebraic stacks and their properties.

4.1. An overview of [EG22]. We have our fixed finite extension F/\mathbf{Q}_p , and we work with coefficient rings which are \mathcal{O} -algebras, where $\mathcal{O} = \mathcal{O}_L$ is sufficiently large. We let \mathcal{X}_d denote the moduli stack of projective étale (φ, Γ) -modules over F of rank d , in the sense of [EG22, Defn. 3.2.1].

REMARK 4.1.1. There are several different possible descriptions of the objects that are parameterized by \mathcal{X}_d . For example, one can consider them as (φ, Γ) -modules with respect to either the full cyclotomic extension of Gal_F , or its \mathbf{Z}_p subextension; but it is also sometimes helpful to think of them as étale $(\mathrm{Gal}_F, \varphi)$ -modules over $W(\mathbf{C}^b)$, in the sense of [EG22, Defn. 2.7.3]. More generally it would probably be helpful to give a description in terms of prismatic F -crystals following [BS21, Wu21c, DL21b].

By its very definition (and by Fontaine’s theory of (φ, Γ) -modules [Fon90]), the groupoid of $\overline{\mathbf{F}}_p$ -points of \mathcal{X}_d , which coincides with the groupoid of $\overline{\mathbf{F}}_p$ -points of the underlying reduced substack $\mathcal{X}_{d, \mathrm{red}}$, is naturally equivalent to the groupoid of continuous representations $\overline{\rho} : \mathrm{Gal}_F \rightarrow \mathrm{GL}_d(\overline{\mathbf{F}}_p)$. More generally, if A is any finite \mathcal{O} -algebra, then the groupoid $\mathcal{X}_d(A)$ is canonically equivalent to the groupoid of continuous representations $\mathrm{Gal}_F \rightarrow \mathrm{GL}_d(A)$. It follows easily that universal Galois lifting rings are versal rings to \mathcal{X}_d at finite type points; so we can think of \mathcal{X}_d as an algebraization of Mazur’s Galois deformation rings.

REMARK 4.1.2. If A is not a finite \mathcal{O} -algebra (e.g. if $A = k[X]$), then there is typically no Galois representation corresponding to an object of $\mathcal{X}_d(A)$. There is a maximal substack $\mathcal{X}_d^{\mathrm{Gal}}$ of \mathcal{X}_d over which the universal (φ, Γ) -modules comes from a Gal_F -representation; this stack was originally constructed and studied by Wang-Erickson [WE18].

The inclusion $\mathcal{X}_d^{\mathrm{Gal}} \subset \mathcal{X}_d$ induces a bijection on $\overline{\mathbf{F}}_p$ -points (and is versal at such points), and one can roughly imagine that $\mathcal{X}_d^{\mathrm{Gal}}$ is the union of the formal completions of \mathcal{X}_d at its closed $\overline{\mathbf{F}}_p$ -points (which correspond to semisimple Gal_F -representations). Thus although the stack $\mathcal{X}_d^{\mathrm{Gal}}$ is the more obvious definition of a moduli stack of “Langlands parameters”, the stack \mathcal{X}_d has richer geometric structure, and although for example the restriction maps from the stacks of global Galois representations to the stacks \mathcal{X}_d factor through the stacks $\mathcal{X}_d^{\mathrm{Gal}}$, we anticipate that it is the stack \mathcal{X}_d (or rather, its derived categories of quasicoherent sheaves) that is the natural setting for the p -adic local Langlands correspondence.

In the p -adic Langlands program it is often important to impose conditions on the Gal_F -representations under consideration, such as demanding that they be crystalline of fixed Hodge–Tate weights, or more generally of being potentially crystalline or semistable of a fixed inertial type. For any⁸ Hodge type λ and inertial type τ , there is a corresponding substack $\mathcal{X}_d^{\mathrm{crys}, \lambda, \tau}$ (and similarly $\mathcal{X}_d^{\mathrm{ss}, \lambda, \tau}$), which are characterized as being the unique closed substack of $(\mathcal{X}_d)_{\mathcal{O}}$ which is flat over \mathcal{O} and whose groupoid of A -valued points, for any finite flat \mathcal{O} -algebra A , is equivalent to

⁸For the following construction, we need to assume that L is “large enough”, e.g. so that τ is defined over L . We largely suppress this point in what follows.

the groupoid of A -valued Gal_F -representations satisfying the corresponding p -adic Hodge theoretic property (of being potentially crystalline or semistable of Hodge type $\underline{\lambda}$ and inertial type τ). In the below if τ is trivial then we drop it from the notation.

These substacks are constructed in [EG22] using (among other techniques) Kisin's results on Breuil–Kisin modules and crystalline representations [Kis06]; again, they are probably more naturally defined in terms of (log-) prismatic F -crystals, following [BS21].

The following theorem summarises the main results of [EG22].

THEOREM 4.1.3.

- (1) \mathcal{X}_d is a Noetherian formal algebraic stack.
- (2) Each substack $\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$ and $\mathcal{X}_d^{\mathrm{ss}, \underline{\lambda}, \tau}$ is a p -adic formal algebraic stack.
- (3) The underlying reduced substack $\mathcal{X}_{d, \mathrm{red}}$ of \mathcal{X}_d (which is an algebraic stack) is of finite type over \mathbf{F}_p , and is equidimensional of dimension $[F : \mathbf{Q}_p]d(d-1)/2$.
- (4) The special fibre of each of $\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$ and $\mathcal{X}_d^{\mathrm{ss}, \underline{\lambda}, \tau}$ (which is an algebraic stack) is equidimensional of dimension equal to that of the flag variety determined by $\underline{\lambda}$. In particular, if $\underline{\lambda}$ is regular, then each of these stacks is equidimensional of dimension $[F : \mathbf{Q}_p]d(d-1)/2$.
- (5) The irreducible components of $\mathcal{X}_{d, \mathrm{red}}$ admit a natural labelling by Serre weights.

REMARK 4.1.4. The notion of a Serre weight is recalled in Section 6.1.37. The precise meaning of the labelling of the irreducible components of $\mathcal{X}_{d, \mathrm{red}}$ by Serre weights is somewhat involved; see [EG22, §5.5]. Roughly, the description is as follows. Each irreducible component of $\mathcal{X}_{d, \mathrm{red}}$ admits a dense open substack, whose $\overline{\mathbf{F}}_p$ -points correspond to Gal_F -representations which are maximally nonsplit of niveau one; that is, they admit a unique filtration by characters. The ordered tuple of inertial weights of these characters determines the irreducible component, and these ordered tuples correspond to Serre weights, because Serre weights are by definition (isomorphism classes of) irreducible $\overline{\mathbf{F}}_p$ -representations of $\mathrm{GL}_d(k)$, and as such are indexed by their highest weights. (There is one subtlety in making this identification, which is that it is necessary to distinguish Serre weights whose highest weights are congruent modulo $p-1$; this manifests itself on the Galois side as a distinction between peu and très ramifiée representations.)

REMARK 4.1.5. The two appearances of $[F : \mathbf{Q}_p]d(d-1)/2$ in Theorem 4.1.3 — as the dimension of the underlying reduced substack $\mathcal{X}_{d, \mathrm{red}}$, and of the dimension of the special fibres of potentially crystalline/semistable substacks of regular Hodge type — is not a coincidence. As well as being an important ingredient in the proof of Theorem 4.1.3, it underlies the Breuil–Mézard conjecture (see Section 6.1.37).

REMARK 4.1.6. As far as we are aware, there is no general dimension theory for formal algebraic stacks in the literature. It is, however, reasonable to think that such a theory exists, and that this dimension can be computed from the dimensions of the versal rings. This being the case, the results of [BIP21] (which in particular compute the dimensions of unrestricted local Galois lifting rings) imply that \mathcal{X}_d will be equidimensional of dimension $1 + [F : \mathbf{Q}_p]d^2$. Since its underlying reduced substack $\mathcal{X}_{d, \mathrm{red}}$ is equidimensional of dimension $[F : \mathbf{Q}_p]d(d-1)/2$, we see that \mathcal{X}_d

should have $1 + [F : \mathbf{Q}_p]d(d+1)/2$ directions which are purely formal (including the p -adic direction). In particular, it is (provably) never a p -adic formal algebraic stack.

REMARK 4.1.7. Again, as far as we are aware, there is no notion in the literature of what it means for a formal algebraic stack to be lci. We anticipate that reasonable definitions exist, and that \mathcal{X}_d is lci; again, this should follow from the results of [BIP21] (which show that the unrestricted local Galois lifting rings are lci).

4.2. Fixed determinant variant. It is often convenient to consider a variant of these stacks where we fix the determinant. If $\chi : \mathrm{Gal}_F \rightarrow \mathcal{O}^\times$ is a character, then we let $\mathcal{X}_d^\chi(A)$ be the groupoid of pairs (D, θ) where D is a rank d projective étale (φ, Γ) -module with A -coefficients, and θ is an identification of $\wedge^d D$ with χ . These stacks (and their potentially crystalline and semistable variants) are studied in an appendix to [DEGa], where the analogue of Theorem 4.1.3 is established (the only difference being that in the case of “Steinberg” Serre weights, there can be multiple corresponding irreducible components, indexed by the d th roots of unity in $\overline{\mathbf{F}}_p$). In particular, the underlying reduced substack of \mathcal{X}_d^χ is again equidimensional of dimension $[F : \mathbf{Q}_p]d(d-1)/2$.

5. Moduli stacks of (φ, Γ) -modules: the analytic case

Similarly to Section 4, we want to define and study stacks of (φ, Γ) -modules on the category of rigid analytic spaces. It is sometimes more convenient to think of these objects in terms of equivariant vector bundles on the Fargues–Fontaine curve. Most of the results in this section will appear in the forthcoming paper [HHS]. We start by recalling the setup and the comparison with (φ, Γ) -modules.

5.1. (φ, Γ) -modules over the Robba ring. Associated with the field F (or rather with the cyclotomic extension F_∞ and the algebraic closure \overline{F}) we can define two versions of the Fargues–Fontaine curve. Let $\varpi^b \in \mathcal{O}_{\hat{F}_\infty^b}$ denote the choice of a pseudo-uniformizer. We set

$$\begin{aligned} Y_{F_\infty} &= \mathrm{Spa}(W(\mathcal{O}_{\hat{F}_\infty^b}), W(\mathcal{O}_{\hat{F}_\infty^b})) \setminus V(p[\varpi^b]), \\ Y_{\overline{F}} &= \mathrm{Spa}(W(\mathcal{O}_{\hat{F}^b}), W(\mathcal{O}_{\hat{F}^b})) \setminus V(p[\varpi^b]). \end{aligned}$$

These adic spaces over \mathbf{Q}_p come equipped with an automorphism φ and a continuous action of Γ , respectively Gal_F , that commutes with φ . The Fargues–Fontaine curves X_{F_∞} and $X_{\overline{F}}$ are the quotients

$$\begin{aligned} X_{F_\infty} &= Y_{F_\infty} / \varphi^{\mathbf{Z}}, \\ X_{\overline{F}} &= Y_{\overline{F}} / \varphi^{\mathbf{Z}}, \end{aligned}$$

which inherit a continuous action of Γ , respectively Gal_F .

Note that the curves X_{F_∞} and $X_{\overline{F}}$ have a distinguished point ∞ defined by the kernel of the usual θ maps of Fontaine

$$\begin{aligned} \theta_{F_\infty} : W(\mathcal{O}_{\hat{F}_\infty^b}) &\longrightarrow \mathcal{O}_{\hat{F}_\infty}, \\ \theta_{\overline{F}} : W(\mathcal{O}_{\hat{F}^b}) &\longrightarrow \mathcal{O}_{\hat{F}}. \end{aligned}$$

In fact ∞ is the unique fixed point for the Γ , respectively Gal_F -action.

We define the following categories of (semi-)linear algebra objects:

- the category $\mathrm{VB}_{X_{F_\infty}}^\Gamma$ of Γ -equivariant vector bundles on X_{F_∞} ,
- the category $\mathrm{VB}_{X_{\bar{F}}}^{\mathrm{Gal}_F}$ of Gal_F -equivariant vector bundles on $X_{\bar{F}}$,
- the category $\mathrm{VB}_{Y_{F_\infty}}^{\varphi, \Gamma}$ of (φ, Γ) -equivariant vector bundles on Y_{F_∞} ,
- the category $\mathrm{VB}_{Y_{\bar{F}}}^{\varphi, \mathrm{Gal}_F}$ of $(\varphi, \mathrm{Gal}_F)$ -equivariant vector bundles on $Y_{\bar{F}}$.

In fact in all the above categories we ask that the action of Γ resp. Gal_F on the vector bundles is continuous in the evident sense (that we will not spell out explicitly). By the work of Berger, Fargues–Fontaine and Kedlaya, all these various categories are known to be canonically equivalent.

THEOREM 5.1.1.

(i) *The projection $X_{\bar{F}} \rightarrow X_{F_\infty}$ is equivariant with respect to the canonical projection $\mathrm{Gal}_F \rightarrow \Gamma$ and induces via pullback an equivalence of categories*

$$\mathrm{VB}_{X_{F_\infty}}^\Gamma \rightarrow \mathrm{VB}_{X_{\bar{F}}}^{\mathrm{Gal}_F}.$$

The corresponding statement for $Y_{\bar{F}} \rightarrow Y_{F_\infty}$ also holds true.

(ii) *The projection $Y_{F_\infty} \rightarrow X_{F_\infty}$ is Γ -equivariant and induces, via pullback, an equivalence of categories*

$$\mathrm{VB}_{X_{F_\infty}}^\Gamma \rightarrow \mathrm{VB}_{Y_{F_\infty}}^{\varphi, \Gamma}.$$

The corresponding statement for $Y_{\bar{F}} \rightarrow X_{\bar{F}}$ also holds true.

REMARK 5.1.2. In Scholze’s world of diamonds one has $X_{F_\infty}/\Gamma = X_{\bar{F}}/\mathrm{Gal}_F$, giving a more geometric interpretation of the first equivalence.

Furthermore there are obvious compatibilities between these equivalences. There is a further description of these categories in terms of the Beauville–Laszlo gluing lemma [BL95]. This gives rise to Berger’s category of B -pairs [Ber08a]. We set

$$B_e = \Gamma(X_{\bar{F}} \setminus \{\infty\}, \mathcal{O}_{X_{\bar{F}}}), \quad B_{\mathrm{dR}}^+ = \hat{\mathcal{O}}_{X_{\bar{F}}, \infty}, \quad B_{\mathrm{dR}} = B_{\mathrm{dR}}^+[1/t].$$

Each of these rings is equipped with a continuous Gal_F -action. Using the Beauville–Laszlo gluing lemma, we can identify $\mathrm{VB}_{X_{\bar{F}}}^{\mathrm{Gal}_F}$ with the category of triples (M, Λ, ξ) consisting of a finite projective B_e -module M and a finite projective B_{dR}^+ -module Λ equipped with continuous semi-linear Gal_F -actions and a Gal_F -equivariant isomorphism

$$\xi : M \otimes_{B_e} B_{\mathrm{dR}} \xrightarrow{\sim} \Lambda \otimes_{B_{\mathrm{dR}}^+} B_{\mathrm{dR}}.$$

There is yet another description of $\mathrm{VB}_{X_{\bar{F}}}^{\mathrm{Gal}_F} \cong \mathrm{VB}_{X_{F_\infty}}^\Gamma$ using (φ, Γ) -modules over imperfect period rings. We write \mathcal{R}_F for the Robba ring

$$\mathcal{R}_F = \varprojlim_{r \rightarrow 1} \varinjlim_{s \rightarrow 1} B_F^{[r, s]},$$

where $B_F^{[r, s]}$ is the ring of rigid analytic functions on the closed annulus $\mathbf{B}_{F_{\infty, 0}}^{[r, s]}$ of inner radius r and outer radius s in the closed unit disc over $F_{\infty, 0}$, the maximal unramified subextension of F_∞ . For $s > r \gg 0$ there is a canonical way to define a continuous action of Γ on $B_F^{[r, s]}$ and to define a ring homomorphism

$$\varphi : B_F^{[r, s]} \rightarrow B_F^{[r^{1/p}, s^{1/p}]}$$

commuting with the Γ -action. These data induce a ring homomorphism $\varphi : \mathcal{R}_F \rightarrow \mathcal{R}_F$ and a continuous Γ -action on \mathcal{R}_F commuting with φ (we refer e.g. to [KPX14, 2.2] for the details of the definition of the φ - and Γ -actions, and for the variants of the ring \mathcal{R}_F with coefficients in an affinoid algebra below).

DEFINITION 5.1.3. The category $\text{Mod}_{\mathcal{R}_F}^{\varphi, \Gamma}$ is the category of finite projective \mathcal{R}_F -modules together with continuous Γ -action and an isomorphism $\varphi^*D \rightarrow D$ commuting with the Γ -action.

REMARK 5.1.4. We recall the relation of (φ, Γ) -modules over \mathcal{R}_F with the (φ, Γ) -modules used in Section 4. The (φ, Γ) -modules of Section 4 are defined over a p -adically complete ring \mathbf{A}_F . This ring contains a subring $\mathbf{A}_F^\dagger \subset \mathbf{A}_F$ of *overconvergent elements* that is stable under the action of φ and Γ . Moreover \mathbf{A}_F^\dagger embeds (equivariantly for φ and Γ) into the \mathbf{Q}_p -algebra \mathcal{R}_F . By a theorem of Cherbonnier–Colmez [CC98] every étale (φ, Γ) -module over $\mathbf{B}_F = \mathbf{A}_F[1/p]$ is overconvergent, where by definition the étale (φ, Γ) -modules over \mathbf{B}_F are those that come by base change from a (φ, Γ) -module over \mathbf{A}_F . More precisely, the scalar extension from $\mathbf{B}_F^\dagger = \mathbf{A}_F^\dagger[1/p]$ to \mathbf{B}_F induces an equivalence between the corresponding categories of étale (φ, Γ) -modules. Moreover, the extension of scalars from \mathbf{B}_F^\dagger to \mathcal{R}_F induces a fully faithful embedding of the category of étale (φ, Γ) -modules over \mathbf{B}_F^\dagger to the category of (φ, Γ) -modules over \mathcal{R}_F . By definition the (φ, Γ) -modules over \mathcal{R}_F in the essential image of this functor are called étale.

By Kedlaya’s slope filtration theorem [Ked08] the étale (φ, Γ) -modules are precisely the objects that are semi-stable of slope zero, in the following sense: If D is a (φ, Γ) -module D of rank r , then by the classification of the (φ, Γ) -modules of rank 1 (see Section 7.1 below), we may write $\bigwedge^r D = \mathcal{R}_F(\delta)$, where δ is a continuous character of F^\times . We then define the *slope* of D to be

$$\text{slope}(D) = \frac{\text{val}_p(\delta(\varpi_F))}{r},$$

where ϖ_F is a choice of a uniformizer in F . We furthermore say that D is *semi-stable* if $\text{slope}(D') \geq \text{slope}(D)$ for all (φ, Γ) -stable subobjects $D' \subseteq D$.

THEOREM 5.1.5. *There is an equivalence of categories*

$$\text{Mod}_{\mathcal{R}_F}^{\varphi, \Gamma} \longrightarrow \text{VB}_{Y_{F_\infty}}^{\varphi, \Gamma}.$$

PROOF. (Sketch.) The rough idea behind this theorem is the following: any (φ, Γ) -module over \mathcal{R}_F admits a model over some closed annulus $\mathbf{B}_{F_{\infty,0}}^{[r,s]}$. Using the choice of the coordinate function $[\varpi^b]$ we can map a corresponding closed annulus in Y_{F_∞} to $\mathbf{B}_{F_{\infty,0}}^{[r,s]}$, equivariantly for the action of Γ (and compatibly with φ in an evident way). The pullbacks of the φ -bundles with Γ -action on the various $\mathbf{B}_{F_{\infty,0}}^{[r,s]}$ to the closed annuli in Y_{F_∞} spread out to a Γ -equivariant φ -bundle on Y_{F_∞} (as φ is an automorphism of Y_{F_∞} that shifts the radii of closed annuli). The hard part of the theorem is the essential surjectivity, sometimes referred to as *deperfection*. It is proved using the additional structure on the cyclotomic tower given by Tate’s normalized traces. A proof of the theorem can be given by combining [FF18, 10.1.2] and [Ber08a, Theorem 2.2.7]. We refer to [BC08, 4.2] for a variant of the descent argument using Tate’s normalized traces. \square

We will be interested in families of these (equivalent) objects with coefficients in rigid analytic spaces. More precisely we want to define a stack \mathfrak{X}_d whose $\text{Sp}(\mathbf{Q}_p)$ -valued points are given by the groupoid of rank d objects in $\text{VB}_{X_F}^{\text{Gal}_F}$ or any of the equivalent categories described above. We briefly describe the variants of the above categories over an affinoid rigid space $\text{Sp}(A)$, which suffices to define our stacks.

Let A be an affinoid algebra topologically of finite type over \mathbf{Q}_p . By a result of Kedlaya [Ked16] the Fargues–Fontaine curves $X_{F_\infty}, X_{\bar{F}}$ are strongly Noetherian, and the same applies (locally) to their coverings $Y_{F_\infty}, Y_{\bar{F}}$. Hence the fiber products

$$X_{F_\infty, A} := X_{F_\infty} \times_{\mathrm{Sp}(\mathbf{Q}_p)} \mathrm{Sp}(A)$$

etc. are well defined in the category of adic spaces, and come equipped with a canonical continuous action of Γ , respectively Gal_F .

Similarly to the absolute case considered above we can define

- the category $\mathrm{VB}_{X_{F_\infty, A}}^\Gamma$ of Γ -equivariant vector bundles on $X_{F_\infty, A}$,
- the category $\mathrm{VB}_{X_{\bar{F}, A}}^{\mathrm{Gal}_F}$ of Gal_F -equivariant vector bundles on $X_{\bar{F}, A}$,
- the category $\mathrm{VB}_{Y_{F_\infty, A}}^{\varphi, \Gamma}$ of (φ, Γ) -equivariant vector bundles on $Y_{F_\infty, A}$,
- the category $\mathrm{VB}_{Y_{\bar{F}, A}}^{\varphi, \mathrm{Gal}_F}$ of $(\varphi, \mathrm{Gal}_F)$ -equivariant vector bundles on $Y_{\bar{F}, A}$.

Again we note that the Γ - respectively Gal_F -action on the vector bundles is supposed to be continuous in the evident way. Moreover, there is a variant with coefficients of the category of (φ, Γ) -modules over the imperfect Robba ring \mathcal{R}_F . We define

$$\mathcal{R}_{F, A} = \varinjlim_{r \rightarrow 1} \varprojlim_{s \rightarrow 1} B_F^{[r, s]} \widehat{\otimes}_{\mathbf{Q}_p} A$$

which is again equipped with a continuous Γ -action and an endomorphism φ commuting with the Γ -action. Then the category $\mathrm{Mod}_{\mathcal{R}_{F, A}}^{\varphi, \Gamma}$ is defined to be the category of finite projective $\mathcal{R}_{F, A}$ -modules D together with an isomorphism $\varphi_D : \varphi^* D \rightarrow D$ and a continuous Γ -action commuting with φ_D .

The conclusions of Theorems 5.1.1 and 5.1.5 still hold true for the categories with coefficients in A , i.e. there are equivalences of categories

$$\begin{array}{ccccc} \mathrm{VB}_{X_{F_\infty, A}}^\Gamma & \xrightarrow{\sim} & \mathrm{VB}_{Y_{F_\infty, A}}^{\varphi, \Gamma} & \xleftarrow{\sim} & \mathrm{Mod}_{\mathcal{R}_{F, A}}^{\varphi, \Gamma} \\ \downarrow \sim & & \downarrow \sim & & \\ \mathrm{VB}_{X_{\bar{F}, A}}^{\mathrm{Gal}_F} & \xrightarrow{\sim} & \mathrm{VB}_{Y_{\bar{F}, A}}^{\varphi, \mathrm{Gal}_F} & & \end{array}$$

induced by similar functors as in the case without coefficients. Moreover, given a morphism $A \rightarrow B$ of affinoid algebras, there are obvious base change functors

$$(-) \widehat{\otimes}_A B : \mathrm{VB}_{X_{F_\infty, A}}^\Gamma \longrightarrow \mathrm{VB}_{X_{F_\infty, B}}^\Gamma$$

etc. and the above equivalences of categories are obviously compatible with the base change from A to B .

REMARK 5.1.6. It is also worth noticing that the description of these equivalent categories using the Beauville–Laszlo gluing lemma generalizes to families. This allows us to modify a given family of equivariant vector bundles over $X_{\bar{F}, A}$ by a given family of invariant lattices over the completion of $X_{\bar{F}, A}$ along $\{\infty\} \times \mathrm{Sp} A$.

5.1.7. Rigid analytic Artin stacks. We introduce basic notions in the framework of Artin stacks on a category of rigid analytic spaces to set the ground for the definition of the stacks of rigid analytic (φ, Γ) -modules. Denote by Rig_L the category of rigid analytic spaces over a fixed base field L that is a finite extension of \mathbf{Q}_p . We will equip Rig_L with the Tate-fpqc topology [CT09, 2.1]. The coverings in this topology are generated by the usual (admissible) Tate coverings and the

morphisms $\mathrm{Sp}(A) \rightarrow \mathrm{Sp}(B)$ of rigid spaces for faithfully flat maps $B \rightarrow A$ of affinoid algebras. With respect to this topology all representable functors are sheaves, and coherent sheaves satisfy descent [Con06, Theorem 4.2.8]. In the following a stack on Rig_L will be a category fibered in groupoids over Rig_L that satisfies descent for the Tate-fpqc topology. As in the case of Artin stacks on schemes, we have to start with the definition of an analogue of algebraic spaces.

DEFINITION 5.1.8. A quasi-analytic space is a sheaf \mathcal{F} on Rig_L such that the diagonal $\mathcal{F} \rightarrow \mathcal{F} \times \mathcal{F}$ is representable and such that there exists an étale surjection $U \rightarrow \mathcal{F}$ from a representable sheaf U onto \mathcal{F} .

REMARK 5.1.9.

(i) In fact below we do not need this level of generality: by a result of Conrad and Temkin [CT09, Theorem 1.2.2] every separated quasi-analytic space is representable by a rigid analytic space. Below we will only meet separated quasi-analytic spaces.

(ii) The PhD thesis of Evan Warner [War17] also develops a theory of Artin stacks on adic spaces. The main difference to our set up is that Warner works with all strongly Noetherian adic spaces (and uses the étale topology), whereas we allow only rigid analytic spaces (and use the Tate-fpqc topology); moreover Warner uses a different terminology. We mainly restrict ourselves to rigid analytic spaces as, at least for the time being, the theory of (φ, Γ) -modules with coefficients [KPX14] and in particular the finiteness results about their cohomology, is limited to coefficients in affinoid algebras of classical rigid analytic geometry.

DEFINITION 5.1.10. A rigid analytic Artin stack is a stack \mathfrak{X} on Rig_L such that the diagonal $\Delta_{\mathfrak{X}} : \mathfrak{X} \rightarrow \mathfrak{X} \times_L \mathfrak{X}$ is representable by quasi-analytic spaces and such that there exists a rigid analytic space U and a smooth surjection $U \rightarrow \mathfrak{X}$.

Given a rigid analytic Artin stack \mathfrak{X} we can define as usual its category of coherent sheaves as

$$\mathrm{Coh}(\mathfrak{X}) = \varprojlim_{\mathrm{Sp}(A) \rightarrow \mathfrak{X}} \mathrm{Coh}(\mathrm{Sp}(A)),$$

where the limit is taken over all maps of affinoid rigid spaces $\mathrm{Sp}(A)$ to \mathfrak{X} . As usual, given a smooth surjection $U \rightarrow \mathfrak{X}$ from an (affinoid) rigid analytic space U to \mathfrak{X} this can be computed as the limit of the diagram

$$\mathrm{Coh}(U) \rightrightarrows \mathrm{Coh}(U \times_{\mathfrak{X}} U) \rightrightarrows \dots$$

Similarly the stable ∞ -category of coherent sheaves on \mathfrak{X} can be defined as the homotopy limit of ∞ -categories

$$\mathbf{D}_{\mathrm{coh}}(\mathfrak{X}) = \varprojlim_{\mathrm{Sp}(A) \rightarrow \mathfrak{X}} \mathbf{D}_{\mathrm{coh}}(\mathrm{Sp}(A)),$$

and for a given covering it can be computed as the homotopy limit of the diagram of ∞ -categories

$$\mathbf{D}_{\mathrm{coh}}(U) \rightrightarrows \mathbf{D}_{\mathrm{coh}}(U \times_{\mathfrak{X}} U) \rightrightarrows \dots$$

For a morphism of rigid analytic spaces $f : \mathfrak{X} \rightarrow \mathfrak{Y}$, or more generally for a morphism between rigid analytic Artin stacks, the definition implies that we have a canonical derived pullback $Lf^* : \mathbf{D}_{\mathrm{coh}}(\mathfrak{Y}) \rightarrow \mathbf{D}_{\mathrm{coh}}(\mathfrak{X})$. If the map f is proper we also have (derived) pushforward $Rf_* : \mathbf{D}_{\mathrm{coh}}(\mathfrak{X}) \rightarrow \mathbf{D}_{\mathrm{coh}}(\mathfrak{Y})$ which is right adjoint to Lf^* . More precisely, for general f the pushforward Rf_* can be defined on the

subcategory of $\mathbf{D}_{\text{coh}}(\mathfrak{X})$ consisting of sheaves with proper support over \mathfrak{Y} . For the rigid analytic Artin stacks that we consider there always exists a dualizing complex $\omega_{\mathfrak{X}}$, and the internal duality

$$\mathbf{D}_{\mathfrak{X}}(-) = R\mathcal{H}om_{\mathcal{O}_{\mathfrak{X}}}(-, \omega_{\mathfrak{X}})$$

allows us to use

$$f^! = \mathbf{D}_{\mathfrak{X}} \circ Lf^* \circ \mathbf{D}_{\mathfrak{Y}} : \mathbf{D}_{\text{coh}}(\mathfrak{Y}) \rightarrow \mathbf{D}_{\text{coh}}(\mathfrak{X})$$

as an ad hoc definition of the upper shriek pullback $f^!$.

REMARK 5.1.11. The definitions above are a bit restrictive, as we only defined categories of coherent and not of quasicoherent sheaves. This is mainly due to the fact that quasicoherent sheaves on rigid analytic spaces are (for topological reasons) a bit delicate. In the world of condensed rings and modules of Clausen–Scholze it is possible to develop a theory of a (derived) category of quasicoherent sheaves together with a full six-functor formalism that gives a more uniform approach than the above ad hoc definition of $f^!$.

We mention two sources of rigid analytic Artin stacks: Given an Artin stack \mathbf{X} on the category Sch_L of L -schemes such that \mathbf{X} is (locally) of finite type over L (i.e. \mathbf{X} can be covered by a scheme \mathbf{U} (locally) of finite type over L), its analytification can be defined as the sheafification of

$$\mathbf{X}^{\text{an}}(\text{Sp}(A)) = \mathbf{X}(\text{Spec}(A)).$$

In particular: for a presentation $\mathbf{V} = \mathbf{U} \times_{\mathbf{X}} \mathbf{U} \rightrightarrows \mathbf{U} \rightarrow \mathbf{X}$ the analytification of \mathbf{X} can be defined as the stack quotient of $\mathbf{V}^{\text{an}} \rightrightarrows \mathbf{U}^{\text{an}}$. Given \mathbf{X} we have an obvious analytification functor on the (derived) category of coherent sheaves which commutes with Lf^* , $f^!$, Rf_* whenever they are defined.

On the other hand, if \mathcal{X} is a formal algebraic stack over $\text{Spf}(\mathcal{O}_L)$, then we can define its generic fiber as

$$(5.1.11) \quad \mathcal{X}_{\eta}^{\text{rig}}(\text{Sp}(A)) = \varinjlim_{\overline{Y}} \mathcal{X}(Y),$$

where the limit is taken over all formal models (in the sense of Raynaud) Y of $\text{Sp}(A)$.

PROPOSITION 5.1.12. *Let \mathcal{X} be a formal algebraic stack. Then $\mathcal{X}_{\eta}^{\text{rig}}$ is a rigid analytic Artin stack.*

5.1.13. *Rigid analytic stacks of (φ, Γ) -modules.* We now return to the main objective of this subsection and define stacks of (φ, Γ) -modules.

DEFINITION 5.1.14. For $d \geq 1$ the category fibered in groupoids $\mathfrak{X}_{\text{GL}_d} = \mathfrak{X}_d$ over Rig_L is the groupoid that assigns to an affinoid space $\text{Sp } A$ the groupoid of (φ, Γ) -modules of rank d over $\mathcal{R}_{F,A}$.

Using the fact that a (φ, Γ) -module over $\mathcal{R}_{F,A}$ admits a model over an annulus over $\text{Sp } A$, descent for vector bundles on rigid analytic spaces implies:

PROPOSITION 5.1.15. *The category fibered in groupoids \mathfrak{X}_d is a stack.*

REMARK 5.1.16. For the expected relationship between \mathfrak{X}_d and the stack \mathcal{X}_d of Section 4, see (5.1.27) and the surrounding discussion.

REMARK 5.1.17. By the above discussion the stack \mathfrak{X}_d can be described equivalently as the stack of Gal_F -equivariant vector bundles on the Fargues–Fontaine curve and in the following we will freely choose whichever description fits better with our purposes. Hence there are some obvious similarities with the classical theory of (stacks of) vector bundles on an algebraic curve. We list the most important similarities and differences:

- (i) There is a cohomology theory

$$R\Gamma_{\mathrm{Gal}_F}(X, -) := R\Gamma(\mathrm{Gal}_F, R\Gamma(\mathrm{Sp} A \times X_{\bar{F}}, -))$$

for Gal_F -equivariant vector bundles on $\mathrm{Sp} A \times X_{\bar{F}}$. In terms of (φ, Γ) -modules this cohomology theory coincides with the usual (φ, Γ) cohomology that is computed by the Herr complex [KPX14, Def. 2.2.3]. More precisely, given a (φ, Γ) -module D over $\mathcal{R}_{F,A}$, the corresponding Herr complex is

$$R\Gamma_{\varphi, \Gamma}(D) = [D^\Delta \xrightarrow{(\varphi^{-1}, \gamma^{-1})} D^\Delta \oplus D^\Delta \xrightarrow{(\gamma^{-1}) \oplus (1-\varphi)} D^\Delta],$$

where $\Delta \subset \Gamma$ is the torsion subgroup and where $\gamma \in \Gamma/\Delta$ is a topological generator. This complex is (quasi-isomorphic to) a perfect complex of A -modules concentrated in degrees 0, 1, 2, see [KPX14, Theorem 4.4.5].

- (ii) As a consequence of (i) an analogue of Grothendieck’s cohomology and base change theorem for coherent sheaves holds true in the context of equivariant vector bundles on the Fargues–Fontaine curve, respectively in the context of (φ, Γ) -modules. More precisely, for a given (φ, Γ) -module D over $\mathcal{R}_{F,A}$ and a point $x \in \mathrm{Sp} A$ the base change map

$$\mathrm{bc}_x^i : H_{\varphi, \Gamma}^i(D) \otimes_A k(x) \longrightarrow H_{\varphi, \Gamma}^i(D \otimes_{\mathcal{R}_{F,A}} \mathcal{R}_{F, k(x)})$$

is surjective if and only if it is an isomorphism; and if bc_x^i is surjective, then $H_{\varphi, \Gamma}^i(D)$ is a vector bundle in a neighborhood of x if and only if bc_x^{i-1} is surjective.

- (iii) There is no analogue of an ample line bundle: for a given equivariant vector bundle \mathcal{V} the Euler characteristic formula (see op. cit.)

$$\sum_{i \geq 0} (-1)^i \dim H_{\mathrm{Gal}_F}^i(X_{\bar{F}}, \mathcal{V}) = -\mathrm{rk} \mathcal{V}[F : \mathbf{Q}_p]$$

implies that there is always a non-vanishing H^1 (unless of course $\mathcal{V} = 0$). In particular we cannot twist away higher cohomology. (Without taking the Gal_F -equivariance into account, there is an analogue of an ample line bundle on $X_{\bar{F}}$. This implies that we can at least twist the vector bundle \mathcal{V} such that its H^2 vanishes.)

- (iv) Local Tate duality [KPX14, Theorem 4.4.5 (3)] asserts that given a (φ, Γ) -module over $\mathcal{R}_{F,A}$ there is a canonical isomorphism

$$R\Gamma_{\varphi, \Gamma}(D) \xrightarrow{\sim} R\mathrm{Hom}_A(R\Gamma_{\varphi, \Gamma}(D^\vee(\varepsilon)), A)[-2].$$

Here $D^\vee = \mathrm{Hom}_{\mathcal{R}_{F,A}}(D, \mathcal{R}_{F,A})$ and ε is the cyclotomic character.

- (v) There is a theory of Hodge–Tate–Sen weights, which we explain in Section 5.1.30 below. It introduces substantial differences in the discussion of P -structures, for a parabolic subgroup $P \subseteq \mathrm{GL}_d$, in Section 5.3 below.

CONJECTURE 5.1.18. *The stack \mathfrak{X}_d is a rigid analytic Artin stack of dimension $d^2[F : \mathbf{Q}_p] = \dim \operatorname{Res}_{F/\mathbf{Q}_p} \operatorname{GL}_d$, which is furthermore normal and a local complete intersection.*

We now describe some partial results in the direction of this conjecture.

THEOREM 5.1.19.

- (i) *The diagonal $\Delta_{\mathfrak{X}_d} : \mathfrak{X}_d \rightarrow \mathfrak{X}_d \times \mathfrak{X}_d$ of \mathfrak{X}_d is representable by rigid spaces.*
- (ii) *Let L' be a finite extension of L and let $x \in \mathfrak{X}_d(L')$ be an L' -valued point. Then there exists an open neighborhood \mathfrak{U}_x of x in \mathfrak{X}_d which is a rigid analytic Artin stack. Moreover, if $d = 2$, then \mathfrak{U}_x has the expected dimension and is normal and a local complete intersection.*

SKETCH OF PROOF: (i) This basically follows from the fact that $R\Gamma_{\varphi,\Gamma}(D)$ is universally computed by a perfect complex. More precisely, given two families D_1, D_2 of (φ, Γ) -modules over an affinoid rigid analytic space $\operatorname{Sp} A$, we have to show that the functor

$$\underline{\operatorname{Isom}}(D_1, D_2) : \operatorname{Sp} B \longmapsto \operatorname{Isom}_{\varphi,\Gamma}(D_1 \widehat{\otimes}_A B, D_2 \widehat{\otimes}_A B)$$

on the category of affinoid spaces over $\operatorname{Sp} A$ is representable by a rigid analytic space.

By looking at the induced map on the top exterior powers it follows that $\underline{\operatorname{Isom}}(D_1, D_2)$ is an open subfunctor of

$$\begin{aligned} \underline{\operatorname{Hom}}(D_1, D_2) &= \underline{H}_{\varphi,\Gamma}^0(D_1^\vee \otimes D_2) : \operatorname{Sp} B \longmapsto \operatorname{Hom}_{\varphi,\Gamma}(D_1 \widehat{\otimes}_A B, D_2 \widehat{\otimes}_A B) \\ &= H_{\varphi,\Gamma}^0((D_1^\vee \otimes D_2) \widehat{\otimes}_A B). \end{aligned}$$

Let us write $D = D_1^\vee \otimes D_2$. Then there is a complex of vector bundles \mathcal{E}^\bullet concentrated in degrees $[0, 2]$ together with an isomorphism $\underline{H}_{\varphi,\Gamma}^0(D) \cong \underline{\ker}(\mathcal{E}^0 \rightarrow \mathcal{E}^1)$, where the functor on the right hand side maps $\operatorname{Sp} B$ to $\ker(\mathcal{E}^0 \otimes_A B \rightarrow \mathcal{E}^1 \otimes_A B)$. As \mathcal{E}^0 and \mathcal{E}^1 are vector bundles, this functor is representable by the preimage of the zero section under the induced map on the corresponding geometric vector bundles over $\operatorname{Sp} A$.

(ii) We sketch the proof in the case $d = 2$. Obviously we can reduce to the case $L' = L$. Let D be a (φ, Γ) -module of rank 2 over $\mathcal{R}_{F,L}$. In order to construct a smooth surjection from a rigid analytic space to a neighborhood of D , we distinguish two possible cases:

- (a) There exists a character $\delta : F^\times \rightarrow L^\times$ such that $D' = D \otimes \mathcal{R}_{F,L}(\delta)$ is étale.
- (b) No such character exists.

Case (a): As D' is étale there exists a Galois representation $\rho' : \operatorname{Gal}_F \rightarrow \operatorname{GL}_2(L)$ with $D_{\operatorname{rig}}(\rho') = D'$. After conjugating ρ' by some element of $\operatorname{GL}_2(L)$ if necessary, we may assume that ρ' takes values in $\operatorname{GL}_2(\mathcal{O})$ (i.e. we choose a stable lattice in the Galois representation). Let $\bar{\rho}' : \operatorname{Gal}_F \rightarrow \operatorname{GL}_2(k)$ denote the reduction of ρ' and let $R_{\bar{\rho}'}$ denote the universal framed deformation ring of $\bar{\rho}'$. Then ρ' defines a point of the rigid analytic generic fiber $(\operatorname{Spf} R_{\bar{\rho}'})^{\operatorname{rig}}$ and we are left to show that the morphism

$$\begin{aligned} (\operatorname{Spf} R_{\bar{\rho}'})^{\operatorname{rig}} &\longrightarrow \mathfrak{X}_2 \\ \rho &\longmapsto D_{\operatorname{rig}}(\rho) \otimes \mathcal{R}(\delta^{-1}) \end{aligned}$$

that maps ρ' to D is smooth in a neighborhood of ρ' . This follows from the infinitesimal lifting criterion together with the fact that the complete local ring of $(\mathrm{Spf} R_{\bar{\rho}'})^{\mathrm{rig}}$ at ρ' pro-represents the universal framed deformation functor of $\rho' : \mathrm{Gal}_F \rightarrow \mathrm{GL}_2(L)$, see [Kis09c, Lemma 2.3.3, Prop. 2.3.5].

Case (b): After again twisting with some rank 1 object we may assume that all slopes of subobjects of D are non-negative. By Lemma 5.1.21 below there exists a character $\delta : F^\times \rightarrow L^\times$ and a non-split extension

$$0 \longrightarrow D \longrightarrow D' \longrightarrow \mathcal{R}_{F,L}(\delta) \longrightarrow 0$$

such that D' is étale and such that $\mathrm{Ext}_{\varphi,\Gamma}^0(\mathcal{R}_{F,L}(\delta), D) = \mathrm{Ext}_{\varphi,\Gamma}^2(\mathcal{R}_{F,L}(\delta), D) = 0$, as well as $\dim \mathrm{Hom}_{\varphi,\Gamma}(D', \mathcal{R}_{F,L}(\delta)) = 1$. Again we choose an \mathcal{O} -lattice ρ' in the Galois representation associated to D' and consider ρ' as a point of $(\mathrm{Spf} R_{\bar{\rho}'})^{\mathrm{rig}}$ the rigid analytic generic fiber of the universal framed deformation ring of the reduction $\bar{\rho}'$ of ρ' .

We construct a locally closed subset $Y \subseteq (\mathrm{Spf} R_{\bar{\rho}'})^{\mathrm{rig}}$ containing ρ' and a smooth morphism $Y \rightarrow \mathfrak{X}_2$ as follows: Let \tilde{D} denote the universal (φ, Γ) -module over $(\mathrm{Spf} R_{\bar{\rho}'})^{\mathrm{rig}}$ and let Y_1 denote the scheme-theoretic support of $H_{\varphi,\Gamma}^2(\tilde{D}(\delta^{-1}\varepsilon))$. As

$$H_{\varphi,\Gamma}^2(\tilde{D}(\delta^{-1}\varepsilon)) \otimes k(\rho') = H_{\varphi,\Gamma}^2(D'(\delta^{-1}\varepsilon)) = (\mathrm{Hom}_{\varphi,\Gamma}(D', \mathcal{R}_{F,L}(\delta)))^\vee$$

is of dimension 1, we can find a neighborhood $Y_2 \subseteq Y_1$ of ρ' and a lift $\tilde{f} \in \Gamma(Y_2, H_{\varphi,\Gamma}^2(\tilde{D}(\delta^{-1}\omega)))$ of the dual basis to a chosen basis vector

$$f^\vee \in \mathrm{Hom}_{\varphi,\Gamma}(D', \mathcal{R}_{F,L}(\delta)).$$

After replacing Y_2 by a smaller open neighborhood of ρ' if necessary, we may assume that $H_{\varphi,\Gamma}^2(\tilde{D}(\delta^{-1}\varepsilon))|_{Y_2}$ is free of rank 1 with basis \tilde{f} . Localizing further we find that there is a neighborhood $Y \subseteq Y_2$ of ρ' together with a surjection

$$\tilde{f}^\vee : \tilde{D}|_Y \longrightarrow \mathcal{R}_{F,Y}(\delta)$$

that specializes to the chosen surjection $f^\vee : D' \rightarrow \mathcal{R}_{F,L}(\delta)$ at ρ' and we moreover may assume that at all points $y \in Y$ we have

$$(5.1.20) \quad \mathrm{Ext}^0(\mathcal{R}_{F,k(y)}(\delta), \tilde{D} \otimes k(y)) = \mathrm{Ext}^2(\mathcal{R}_{F,k(y)}(\delta), \tilde{D} \otimes k(y)) = 0.$$

We claim that the map $Y \rightarrow \mathfrak{X}_2$ defined by $\tilde{D}|_Y \mapsto \ker \tilde{f}^\vee$ is smooth. Again it is enough to check this on the complete local ring of a point $y \in Y$. The complete local ring at y pro-represents the functor of deformations of $\tilde{D}_y = \tilde{D} \otimes k(y)$ together with a deformation of the map $f_y^\vee : \tilde{D}_y \rightarrow \mathcal{R}_{F,k(y)}(\delta)$. By our choices of Y this functor may be identified with the functor of deformations of $D_y = \ker f_y^\vee$ together with a deformation of the extension

$$0 \longrightarrow D_y \longrightarrow \tilde{D}_y \longrightarrow \mathcal{R}_{F,L}(\delta) \longrightarrow 0.$$

By the vanishing of the Ext-groups (5.1.20) this functor is formally smooth over the deformation functor of D_y .

We are left to prove the statement about the dimension and the claims on the local structure. This is Proposition 5.1.25 below. \square

LEMMA 5.1.21. *Let D be a (φ, Γ) -module of rank 2 over $\mathcal{R}_{F,L}$ and assume that all slopes of subobjects of D are non-negative and that $D \not\cong \mathcal{R}_{F,L}(\delta)^{\oplus 2}$ for any character $\delta : F^\times \rightarrow L^\times$. Let $\eta : F^\times \rightarrow L^\times$ be a continuous character such that*

- for all embeddings $\tau : F \hookrightarrow L$ one has $(\text{wt}_\tau(\eta) + \mathbf{Z}) \cap \text{wt}_\tau(D) = \emptyset$, where we refer to Section 5.1.30 below for the definition of the weights.
- the slope of η is $\text{slope}(\eta) = -\text{slope}(D)$.

Then there exists a non-split extension

$$(5.1.22) \quad 0 \longrightarrow D \longrightarrow D' \longrightarrow \mathcal{R}_{F,L}(\eta) \longrightarrow 0$$

such that D' is étale and moreover

$$\begin{aligned} \text{Ext}_{\varphi,\Gamma}^0(\mathcal{R}_{F,L}(\eta), D) &= \text{Ext}_{\varphi,\Gamma}^2(\mathcal{R}_{F,L}(\eta), D) = 0, \\ \dim_L \text{Hom}_{\varphi,\Gamma}(D', \mathcal{R}_{F,L}(\eta)) &= 1. \end{aligned}$$

REMARK 5.1.23. Given D as in the lemma, a character η satisfying the assumptions in the lemma obviously always exists.

PROOF. As $D(\eta^{-1})$ and $D^\vee(\eta\omega)$ have non-integral Hodge–Tate–Sen weights, the vanishing of the Ext-groups is a direct consequence of the identifications

$$\begin{aligned} \text{Ext}_{\varphi,\Gamma}^0(\mathcal{R}_{F,L}(\eta), D) &= H_{\varphi,\Gamma}^0(D(\eta^{-1})) \\ \text{Ext}_{\varphi,\Gamma}^2(\mathcal{R}_{F,L}(\eta), D) &= H_{\varphi,\Gamma}^0(D^\vee(\eta\omega)^\vee). \end{aligned}$$

Choose any non-split extension (5.1.22). The claim on the Hodge–Tate weights also implies the claim on the dimension of $\text{Hom}_{\varphi,\Gamma}(D', \mathcal{R}_{F,L}(\eta))$, so we only need to show that D' can be chosen to be étale. By the condition on the slope of η we automatically have $\text{slope}(D') = 0$ and by Kedlaya’s theorem we are left to show that for all (saturated) subobjects $D'' \subseteq D$ we have $\text{slope}(D'') \geq 0$. By assumption this is true if $D'' \subseteq D$. It therefore suffices to show that we can choose D' such that all saturated (φ, Γ) -stable subobjects of D' are contained in D . Assume to the contrary that there is some D'' such that the map $D'' \rightarrow \mathcal{R}_{F,L}(\eta)$ is non-zero. We distinguish the cases $\text{rk } D'' = 1$ and $\text{rk } D'' = 2$.

If $\text{rk } D'' = 1$, then $D'' \rightarrow \mathcal{R}_{F,L}(\eta)$ is an isomorphism after inverting t , which implies that the extension (5.1.22) is split after inverting t . However, the assumption $(\text{wt}_\tau(\eta) + \mathbf{Z}) \cap \text{wt}_\tau(D) = \emptyset$ implies that the canonical map

$$\text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}, D) \longrightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}[1/t], D[1/t])$$

is an isomorphism, which is a contradiction, as we chose (5.1.22) to be non-split.

If $\text{rk } D'' = 2$, then $E = D'' \cap D \subseteq D$ is a (φ, Γ) -submodule of rank 1. By assumption, D has only finitely many saturated subobjects of rank 1. Each such subobject E gives rise to a short exact sequence

$$(5.1.24) \quad 0 \longrightarrow \mathcal{R}_{F,L}(\delta_1) \longrightarrow D \longrightarrow \mathcal{R}_{F,L}(\delta_2) \longrightarrow 0$$

and our assumptions on the Hodge–Tate–Sen weights implies the vanishing of certain Ext^0 and Ext^2 groups and hence a short exact sequence

$$\begin{aligned} 0 \rightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}(\eta), \mathcal{R}_{F,L}(\delta_1)) &\rightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}(\eta), D) \rightarrow \\ &\rightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}(\eta), \mathcal{R}_{F,L}(\delta_2)) \rightarrow 0. \end{aligned}$$

Moreover, as in the rank 1 case, the condition $(\text{wt}_\tau(\eta) + \mathbf{Z}) \cap \text{wt}_\tau(D) = \emptyset$ implies that inverting t induces an isomorphism between this exact sequence and the exact sequence

$$\begin{aligned} 0 \rightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}(\eta)[1/t], \mathcal{R}_{F,L}(\delta_1)[1/t]) &\rightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}(\eta)[1/t], D[1/t]) \rightarrow \\ &\rightarrow \text{Ext}_{\varphi,\Gamma}^1(\mathcal{R}_{F,L}(\eta)[1/t], \mathcal{R}_{F,L}(\delta_2)[1/t]) \rightarrow 0. \end{aligned}$$

As D has only finitely many subobjects, we can choose the extension

$$D' \in \text{Ext}_{\varphi, \Gamma}^1(\mathcal{R}_{F,L}(\eta), D)$$

such that for all possibilities to write D as an extension (5.1.24), the extension D' is not in the image of $\text{Ext}_{\varphi, \Gamma}^1(\mathcal{R}_{F,L}(\eta), \mathcal{R}_{F,L}(\delta_1))$.

In this case, if $D'' \subseteq D'$ is a saturated subobject with $D'' \cap D = E = \mathcal{R}_{F,L}(\delta_1)$ and such that $D'' \rightarrow \mathcal{R}_{F,L}(\eta)$ is non-zero, then

$$\mathcal{R}_{F,L}(\eta)[1/t] \cong (D''/E)[1/t] \subseteq (D'/E)[1/t]$$

implies that the image $(D'/E)[1/t]$ of D' in $\text{Ext}_{\varphi, \Gamma}^1(\mathcal{R}_{F,L}(\eta)[1/t], \mathcal{R}_{F,L}(\delta_2)[1/t])$ vanishes. This contradicts our choice of D' above. \square

The existence of local charts implies the existence of versal rings at rigid analytic points of \mathfrak{X}_d .

PROPOSITION 5.1.25. *Let $d = 2$ and $x \in \mathfrak{X}_d$ and let R be a versal ring at x . Then R is a local complete intersection of dimension $d^2[F : \mathbf{Q}_p]$ and is normal.*

SKETCH OF PROOF. As any two versal rings are smoothly equivalent, we may use any choice of a versal ring to prove the assertion. We use the notation from the proof of Theorem 5.1.19. If x is as in case (a) of that proof, then R can be chosen to be isomorphic to the complete local ring of $(\text{Spf } R_{\rho'})^{\text{rig}}$ at the point ρ' . The result then follows from [BIP21].

In case (b), where x is not the twist of an étale (φ, Γ) -module, standard deformation theory arguments give us a presentation of a versal ring R such that $\text{Spf } R \rightarrow \mathfrak{X}_d$ is formally smooth of dimension m and such that every irreducible component of R has dimension at least $m + [F : \mathbf{Q}_p]d^2$, and in order to show that R is lci, we need to show that for all irreducible components equality holds. By the existence of local charts (i.e. Theorem 5.1.19 (2)), we may view R as the complete local ring of a point $y \in \text{Sp } A$, where $\text{Sp } A \rightarrow \mathfrak{X}_d$ is smooth. We claim that each Z be an irreducible component of $\text{Sp } A$ that contains y is such that the (φ, Γ) -module on Z is generically on Z the twist of an étale (φ, Γ) -module. Then the argument in case (a) implies that Z has the expected dimension $m + [F : \mathbf{Q}_p]d^2$. It is enough to show that there are no irreducible components Z of $\text{Sp}(A)$ containing y on which the (φ, Γ) -module is (generically) reducible (because otherwise it is generically irreducible and thus étale up to twist).

To see this, given an extension (5.1.24) we consider the groupoid of deformations \mathfrak{D} of this short exact sequence and denote by S a versal ring to \mathfrak{D} such that $\text{Spf } S \rightarrow \mathfrak{D}$ is formally smooth of dimension m . Then it is enough to show that all irreducible components of S are of dimension strictly less than $m + [F : \mathbf{Q}_p]d^2$. In the case $d = 2$ this can easily be checked by a direct computation (compare also the discussion of the stack of B -bundles in 5.3 below).

For normality we need to check in addition (if $F = \mathbf{Q}_p$) that for a generic choice of a reducible (φ, Γ) -module

$$0 \rightarrow \mathcal{R}_{F,L}(\delta_1) \rightarrow D \rightarrow \mathcal{R}_{F,L}(\delta_2) \rightarrow 0$$

the stack \mathfrak{X}_2 is smooth at D . Again this follows from a tangent space computation (resp. a computation of Ext-groups) together with the fact that generically on the stack of B -bundles, the ratio δ_1/δ_2 is very regular in the sense that

$$\text{Ext}_{\varphi, \Gamma}^0(\mathcal{R}_{F,L}(\delta_2), \mathcal{R}_{F,L}(\delta_1)) = \text{Ext}_{\varphi, \Gamma}^2(\mathcal{R}_{F,L}(\delta_2), \mathcal{R}_{F,L}(\delta_1)) = 0. \quad \square$$

REMARK 5.1.26. Basically the same strategy (using not just the Borel B but all parabolic subgroups) should settle the claim on the dimension and the local structure for general d . Only the computation of the dimension of spaces/stacks of extensions becomes more involved and is not finished yet. We note that similar computations of dimensions of Ext-groups also show up in the proof of the main theorem of [BIP21].

In connection with the integral theory in Section 4, there is a morphism

$$(5.1.27) \quad \pi_d : \mathcal{X}_{d,\eta}^{\text{rig}} \longrightarrow \mathfrak{X}_d$$

from the rigid analytic generic fiber of \mathcal{X}_d to the stack \mathfrak{X}_d that is basically given by forgetting the étale lattice. This morphism should be formally étale, but not representable by rigid analytic spaces (nor by quasi-analytic spaces), compare the $d = 1$ case in Section 7.1. It is not clear from the definitions that such a map exists, but the existence is a consequence of the following conjecture, which recently has been proved by Gal Porat [Por22]:

CONJECTURE 5.1.28. *Let A be a p -adically complete \mathbf{Z}_p -algebra topologically of finite type over \mathbf{Z}_p . Then every étale (φ, Γ) -module over $\mathbf{A}_{F,A}$ is overconvergent.*

It is worth noticing that the conjecture follows from work of Berger–Colmez [BC08] if the étale (φ, Γ) -module arises from a family of Galois representations over A . The work of Porat [Por22] also treats the case that the étale (φ, Γ) -module is not attached to a Galois representation, which can happen if A is not finite over \mathbf{Q}_p .

REMARK 5.1.29. We point out that the proof of Theorem 5.1.19 only gives charts locally around rigid analytic points. As a set theoretic cover of a rigid analytic space by admissible open subsets is not necessarily an admissible cover, this is not enough to give a full proof of Conjecture 5.1.18. Extending the argument in the above proof to all points of the corresponding adic space (not just those points x such that $L' = k(x)$ is finite over L) would prove Conjecture 5.1.18. The given argument would directly generalize if one had an analogue of Kedlaya’s slope filtration theorem for (φ, Γ) -modules over $\mathcal{R}_{F,L'}$, where L' is allowed to be the (completed) residue field of any point in an adic space of finite type over \mathbf{Q}_p and if moreover one could prove that the map π_d (that exists due to [Por22]) is formally étale. While étaleness is not expected to be difficult to prove, we have no idea whether the generalization of Kedlaya’s theorem to general coefficient fields is true or not.

5.1.30. *Hodge–Tate–Sen weights and global sections.* The theory of Hodge–Tate–Sen weights for representations of Gal_F on finite dimensional \mathbf{Q}_p -vector spaces generalizes to families of (φ, Γ) -modules, respectively to families of Gal_F -equivariant vector bundles on $X_{\widehat{F}}$. Let A denote an affinoid algebra and let D_A be a (φ, Γ) -module D_A over $\mathcal{R}_{F,A}$. We write $\mathcal{E}_A = \mathcal{E}_{F_\infty}(D_A)$ for the corresponding Γ -equivariant vector bundle on $X_{F_\infty,A}$. We then consider the special fiber

$$\widehat{D}_{\text{Sen}}(D_A) = \mathcal{E}_A|_{\text{Sp } A \times \{\infty\}} = D_A \widehat{\otimes}_{\mathcal{R}_{F,A}} (\widehat{F}_\infty \widehat{\otimes}_{\mathbf{Q}_p} A)$$

where the tensor product on the right hand side is induced by a certain (explicit) map $\mathcal{R}_F \rightarrow \widehat{F}_\infty$. The map $\mathcal{R}_F \rightarrow \widehat{F}_\infty$ factors through $F_\infty \subset \widehat{F}_\infty$ and the module $\widehat{D}_{\text{Sen}}(D_A)$ has a decompleted version

$$D_{\text{Sen}}(D_A) = D_A \otimes_{\mathcal{R}_{F,A}} (A \otimes_{\mathbf{Q}_p} F_\infty)$$

which is a finite projective $A \otimes_{\mathbf{Q}_p} F_\infty$ -module that carries a continuous action of the 1-dimensional p -adic Lie group Γ . On the decompleted version $D_{\text{Sen}}(D_A)$ the Γ -action is locally analytic and hence one can derive this action, defining a morphism

$$\text{Lie } \Gamma \longrightarrow \text{Lie}(\text{GL}(D_{\text{Sen}}(D_A))) = \text{End}(D_{\text{Sen}}(D_A)).$$

The cyclotomic character induces a trivialization $\mathbf{Z}_p \cong \text{Lie } \Gamma$ and hence the image of $1 \in \mathbf{Z}_p$ defines an endomorphism

$$\Theta : D_{\text{Sen}}(D_A) \longrightarrow D_{\text{Sen}}(D_A).$$

This endomorphism commutes with the Γ -action, so the coefficients of its characteristic polynomial take values in $F \otimes_{\mathbf{Q}_p} A$ (though we caution the reader that this does not mean that the pair (D_A, Θ) descends to an object over $F \otimes_{\mathbf{Q}_p} A$). If A is an L -algebra and $|\text{Hom}(F, L)| = [F : \mathbf{Q}_p]$, then $F \otimes_{\mathbf{Q}_p} A \cong \prod_{\tau: F \rightarrow L} A_\tau$, where $A_\tau = A$. For each embedding τ we obtain an endomorphism Θ_τ of the finite projective A -module

$$D_{\text{Sen}, \tau}(D_A) = D_{\text{Sen}}(D_A) \otimes_{A \otimes_{\mathbf{Q}_p} F} A_\tau$$

with characteristic polynomial χ_{Θ_τ} with coefficients in A . The zeros of this polynomial (i.e. the eigenvalues of Θ_τ) are called the Hodge–Tate–Sen weights of D labeled by τ . If $A = L$ is a field containing all the zeros of χ_{Θ_τ} , then we write $\text{wt}_\tau(D) \subseteq L$ for the set of all these zeros.

We write $\mathfrak{gl}_n // \text{GL}_n$ for the GIT quotient of the Lie algebra \mathfrak{gl}_n by the adjoint action of GL_n . Mapping a (φ, Γ) -module D_A over $\mathcal{R}_{F,A}$ to the coefficients of the characteristic polynomial χ_Θ of Θ defines a map

$$(5.1.31) \quad \omega_d : \mathfrak{X}_d \longrightarrow (\text{Res}_{F/\mathbf{Q}_p}(\mathfrak{gl}_n // \text{GL}_n))_L = \text{WT}_{d,L}.$$

If $|\text{Hom}(F, L)| = [F : \mathbf{Q}_p]$, then of course $\text{WT}_{d,L} = \prod_{\tau: F \rightarrow L} \text{WT}_{\tau,d,L}$ and each $\text{WT}_{\tau,d,L} = (\mathfrak{gl}_n // \text{GL}_n)_L$ is an affine space.

Taking the top exterior power of an equivariant vector bundle defines a map

$$\det : \mathfrak{X}_d \longrightarrow \mathfrak{X}_1 \longrightarrow \mathcal{T},$$

where \mathcal{T} is the space of continuous characters of F^\times (see Section 7.1 below for a description of \mathfrak{X}_1 and for the definition of the last map). Similarly, the top exterior power defines a map $\text{WT}_{d,L} \rightarrow \text{WT}_{1,L}$ and ω_1 induces a map $\mathcal{T} \rightarrow \text{WT}_{1,L}$. We expect that these morphisms are the source of all global functions on the stack \mathfrak{X}_d .

CONJECTURE 5.1.32. *The maps ω_d and \det induce an isomorphism*

$$\Gamma(\mathfrak{X}_d, \mathcal{O}_{\mathfrak{X}_d}) \cong \Gamma(\text{WT}_{d,L}, \mathcal{O}_{\text{WT}_{d,L}}) \widehat{\otimes}_{\Gamma(\text{WT}_{1,L}, \mathcal{O}_{\text{WT}_{1,L}})} \Gamma(\mathcal{T}, \mathcal{O}_{\mathcal{T}}).$$

COROLLARY 5.1.33. *Assume Conjecture 5.1.32. Then the connected components of \mathfrak{X}_d are in bijection (via the determinant) with the connected components of \mathfrak{X}_1 and hence with the connected components of the space*

$$\mathcal{W} = (\text{Spf } \mathcal{O}[[\mathcal{O}_F^\times]])^{\text{rig}}$$

of continuous characters of \mathcal{O}_F^\times .

Note that the connected components of \mathcal{W} are in turn in bijection with the characters $\mu(F) \rightarrow \mathcal{O}^\times$. We point out that that a similar result for the components of universal deformation rings was proven by Böckle–Iyengar–Paškūnas [BIP21].

5.2. Stacks of de Rham objects. The stack of (φ, Γ) -modules \mathfrak{X}_d has closed substacks defined in terms of p -adic Hodge theory, similarly to the closed substacks $\mathcal{X}_d^{\text{ss}, \lambda, \tau} \subset \mathcal{X}_d$. In the case of (φ, Γ) -modules over the Robba ring these closed substacks can be studied in terms of (filtered) Weil–Deligne representations. In the conjectural relation of locally analytic representations with sheaves on the stack \mathfrak{X}_d these closed substacks of de Rham objects (of given weight and inertial type) will play a role in the comparison with the smooth categorical Langlands conjectures. We will now define these “de Rham loci”.

Let F' be a finite Galois extension of F with Galois group $\text{Gal}(F'/F)$ and write F'_0 respectively F_0 for the maximal unramified subextension of F' respectively F . Moreover, let σ denote the Frobenius lift on F'_0 . We consider the stack $\text{Mod}_{d, \varphi, N, F'/F}$ on the category of L -schemes that maps an L -algebra A to the groupoid of finite projective $A \otimes_{\mathbf{Q}_p} F'_0$ -modules D of rank d together with

- an $\text{id} \otimes \sigma$ -linear automorphism $\varphi_D : D \xrightarrow{\sim} D$.
- an $A \otimes_{\mathbf{Q}_p} F'_0$ -linear endomorphism $N : D \rightarrow D$ satisfying

$$N \circ \varphi_D = p\varphi_D \circ \sigma^* N.$$

- an action of $\text{Gal}(F'/F)$ on D , commuting with φ_D and N , such that

$$g((a \otimes f) \cdot d) = (a \otimes g(f)) \cdot g(d)$$

for all $a \in A$, $f \in F'_0$, $d \in D$ and $g \in \text{Gal}(F'/F)$.

We further recall the stack of d -dimensional Weil–Deligne representations $\text{WD}_{d, F}$ of F which is an Artin stack on the category of L -schemes, see [DHKM20], [Zhu20, §3]. Recall that for a given L -scheme S the groupoid $\text{WD}_{d, F}(S)$ parameterizes vector bundles \mathcal{E} on S together with a smooth action ρ of the Weil group W_F and an endomorphism $N : \mathcal{E} \rightarrow \mathcal{E}$ satisfying $Ng = q^{\|g\|} gN$. Given F' we write $\text{WD}_{d, F'/F} \subset \text{WD}_{d, F}$ for the open and closed substack of $\text{WD}_{d, F}$ consisting of those Weil–Deligne representations (\mathcal{E}, ρ, N) such that the restriction of ρ to the inertia subgroup $I_{F'}$ of F' is trivial.

We note that Fontaine’s construction of the Weil–Deligne representation associated to a $(\varphi, N, \text{Gal}(F'/F))$ -module [Fon94] (see also [BS07, Prop. 4.1]) implies that these two stacks become (non-canonically) isomorphic over a large enough field L :

LEMMA 5.2.1. *Assume that L is large enough such that $[F'_0 : \mathbf{Q}_p] = |\text{Hom}(F'_0, L)|$ and fix an embedding $F'_0 \hookrightarrow L$. Then there is an isomorphism of stacks*

$$\text{WD}_d : \text{Mod}_{d, \varphi, N, F'/F} \xrightarrow{\sim} \text{WD}_{d, F'/F}.$$

We will spell out this isomorphism in detail in Section 5.2.7 below, and also give explicit descriptions of the stacks involved (at least for $F' = F$).

In order to link a $(\varphi, N, \text{Gal}(F'/F))$ -module $D \in \text{Mod}_{d, \varphi, N, F'/F}(A)$ to equivariant vector bundles, we need to introduce a filtration on D . Let us write $G = \text{Res}_{F/\mathbf{Q}_p} \text{GL}_d$ and $G' = \text{Res}_{F'/\mathbf{Q}_p} \text{GL}_d$ for the moment. Moreover, we write $G'_0 = \text{Res}_{F'_0/\mathbf{Q}_p} \text{GL}_d$.

Let $T \subseteq G$ denote the Weil restriction of the diagonal torus of GL_d and let $B \subseteq G$ denote the Weil restriction of the upper triangular matrices. For a choice of a dominant cocharacter $\underline{\lambda} \in X_*(T_{\bar{L}})_+$ we write $L_{\underline{\lambda}}$ for the reflex field of $\underline{\lambda}$ and $P_{\underline{\lambda}} \subseteq G_{L_{\underline{\lambda}}}$ for the parabolic subgroup (containing $B_{L_{\underline{\lambda}}}$) defined by $\underline{\lambda}$. The flag variety

G_{L_λ}/P_λ then parameterizes filtrations on $L_\lambda^d \otimes_{\mathbf{Q}_p} F$ of type λ . More precisely, using the standard identification

$$X_*(T)_+ = \prod_{\tau: F \hookrightarrow \bar{L}} \mathbf{Z}_+^d$$

we write

$$\lambda = (\lambda_{\tau,1} \geq \lambda_{\tau,2} \geq \cdots \geq \lambda_{\tau,d})_{\tau: F \hookrightarrow \bar{L}}.$$

Then an \bar{L} -valued point of G_{L_λ}/P_λ is given by a product of descending, exhaustive and separated \mathbf{Z} -filtrations on

$$\bar{L}^d \otimes_{\mathbf{Q}_p} F = \prod_{\tau: F \hookrightarrow \bar{L}} \bar{L}_\tau^d$$

(with $\bar{L}_\tau = \bar{L}$) such that $(\lambda_{\tau,1} \geq \lambda_{\tau,2} \geq \cdots \geq \lambda_{\tau,d})$ are precisely the weights (counted with multiplicity) of the filtration on \bar{L}_τ^d .

Using the canonical map $G \rightarrow G'$ we can also view λ as a cocharacter of a torus in G' (more precisely again of the Weil restriction of the diagonal torus in GL_d) and we obtain a similar parabolic subgroup P'_λ and a flag variety $G'_{L_\lambda}/P'_\lambda$.

Note that there is a forgetful map

$$\mathrm{WD}_{d,F} \longrightarrow \mathrm{BGL}_d = */\mathrm{GL}_d$$

to the stack of rank d vector bundles. And similarly there is a canonical map

$$\mathrm{Mod}_{d,\varphi,N,F'/F} \longrightarrow */G'_0$$

forgetting the action of φ, N and $\mathrm{Gal}(F'/F)$ on a projective $A \otimes_{\mathbf{Q}_p} F'_0$ -module of rank d .

The scheme G/P_λ is equipped with a (diagonal) left translation action of GL_d and we can define the fiber product

$$\mathrm{Fil}_\lambda \mathrm{WD}_{d,F} = \mathrm{WD}_{d,F} \times_{*/\mathrm{GL}_d} \mathrm{GL}_d \backslash G_{L_\lambda}/P_\lambda$$

that is called the stack of filtered Weil–Deligne representations (of Hodge–Tate weight λ). Its S -valued points parameterize Weil–Deligne representations on vector bundles \mathcal{E} over S together with a filtration on $F \otimes_{\mathbf{Q}_p} \mathcal{E}$ of type λ .

Similarly, the scheme G'/P'_λ is equipped with a (diagonal) left translation action of G'_0 and we can consider the closed substack

$$\mathrm{Fil}_\lambda \mathrm{Mod}_{d,\varphi,N,F'/F} \subset \mathrm{Mod}_{d,\varphi,N,F'/F} \times_{*/G'_0} G'_0 \backslash G'_{L_\lambda}/P'_\lambda$$

parameterizing (for a given object $D \in \mathrm{Mod}_{d,\varphi,N,F'/F}$) those filtrations (of type λ) on $D \otimes_{F'_0} F$ that are stable under the action of $\mathrm{Gal}(F'/F)$.

REMARK 5.2.2. We point out that $\mathrm{Gal}(F'/F)$ -stable filtrations can only exist if the cocharacter λ of (our maximal torus in) G' comes from (our maximal torus in) G via the map $G \rightarrow G'$.

The following lemma then generalizes Lemma 5.2.1.

LEMMA 5.2.3. *Assume that L is large enough such that $[F' : \mathbf{Q}_p] = |\mathrm{Hom}(F', L)|$ and fix an embedding $F' \hookrightarrow L$. Then there is a canonical isomorphism*

$$\mathrm{Fil}_\lambda \mathrm{Mod}_{d,\varphi,N,F'/F} \xrightarrow{\sim} \mathrm{Fil}_\lambda \mathrm{WD}_{d,F'/F}$$

lying above the isomorphism WD_d of Lemma 5.2.1.

The stack $\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,F'/F}$ plays an important role in the characterization of de Rham (φ, Γ) -modules and via the above isomorphism we can hence link filtered Weil–Deligne representations to (φ, Γ) -modules, at least if the base field is large enough. More precisely we have the following theorem:

THEOREM 5.2.4. *For fixed $\underline{\lambda} \in X_*(T_{\bar{L}})_+$ there is a canonical closed embedding*

$$(\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,F'/F})^{\mathrm{an}} \longrightarrow \mathfrak{X}_{d,L_{\underline{\lambda}}}.$$

Moreover, the image of this morphism coincides with the closed substack of de Rham (φ, Γ) -modules of Hodge–Tate–Sen weight $\underline{\lambda}$ that become semi-stable after base change to F' .

The construction of a family of (φ, Γ) -modules of rank d over the rigid analytic stack $(\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,F'/F})^{\mathrm{an}}$ is the generalization of the work of Berger [Ber08b] to families of the objects involved.

To make sense of the “moreover” part of Theorem 5.2.4 we need to define the notion of a family of (potentially) semi-stable or of de Rham (φ, Γ) -modules over $\mathcal{R}_{F,A}$. Given a (φ, Γ) -module D_A over $\mathcal{R}_{F,A}$ we can consider the base change

$$D_{\mathrm{dif}}(D_A) = D_A \otimes_{\mathcal{R}_{F,A}} F_{\infty}((t))$$

which is endowed with a derivation

$$\partial : D_{\mathrm{dif}}(D_A) \longrightarrow D_{\mathrm{dif}}(D_A)$$

above $t \frac{d}{dt}$ that is defined by the derivative of the Γ -action. Then D_A is defined to be de Rham if $D_{\mathrm{dif}}(D_A)^{\partial=0}$ is a projective $A \otimes_{\mathbf{Q}_p} F_{\infty}$ -module of rank d and if the canonical map

$$D_{\mathrm{dif}}(D_A)^{\partial=0} \otimes_{F_{\infty}} F_{\infty}((t)) \longrightarrow D_{\mathrm{dif}}(D_A)$$

is an isomorphism. The proof of the “moreover” part of Theorem 5.2.4 then follows closely the proof of [BC08, Thms. 5.3.1, 5.3.2]

REMARK 5.2.5. Recall that the (φ, Γ) -modules in the analytic set up are not required to be étale. Hence the stacks of de Rham (φ, Γ) -modules can be described by filtered (φ, N) -modules (with descent data) as done in Theorem 5.2.4 without putting an additional “weak admissibility” condition on the filtration. It is however possible to pass to an open substack of $(\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,F'/F})^{\mathrm{an}}$ where the (φ, Γ) -modules associated to the filtered $(\varphi, N, \mathrm{Gal}_F)$ -modules are étale, see [HH20] and [Hel13] for that construction in the case of a trivial Gal_F -action. This étale locus then agrees with the (union over appropriate τ of the) image of the generic fibers of the stacks $\mathcal{X}_d^{\mathrm{ss}, \underline{\lambda}, \tau}$ under the map (5.1.27).

REMARK 5.2.6. We point out that it is important to fix the cocharacter $\underline{\lambda}$ in the theorem. If $\underline{\lambda}$ is allowed to vary arbitrarily, then the union of the corresponding substacks $(\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,F'/F})^{\mathrm{an}}$ is not closed in \mathfrak{X}_d but expected to be Zariski-dense. This should be seen as an analogue of the density of crystalline representations in local deformation rings, see [Che13, Nak14, BIP22].

5.2.7. The role of the embedding $F'_0 \hookrightarrow L$. In Lemma 5.2.1 above, we have identified the stacks $\mathrm{Mod}_{d,\varphi,N,F'/F}$ and $\mathrm{WD}_{d,F'/F}$ over a field L containing F'_0 using the choice of an embedding $\sigma_0 : F'_0 \hookrightarrow L$. We elaborate a bit further about the choice of this embedding. For simplicity (and in order not to overload the notation) we limit ourselves to the case $F' = F$. We refer to [BG19, §2.6] and [HH20, §3] for the explicit description of the stacks that we are using in the following.

Assume that L contains $F'_0 = F_0$ (the equality holding because we have assumed that $F' = F$); then for an L -algebra A we have a decomposition

$$A \otimes_{\mathbf{Q}_p} F_0 = A \otimes_L (L \otimes_{\mathbf{Q}_p} F_0) = A \otimes_L \prod_{\varsigma: F_0 \hookrightarrow L} L_\varsigma = \prod_{\varsigma: F_0 \hookrightarrow L} A_\varsigma,$$

where $L_\varsigma = L$ and $A_\varsigma = A$ for all ς . The action of $\text{Gal}(F_0/\mathbf{Q}_p)$ of course just permutes the factors in the product. For the remainder of this section we fix an embedding ς_0 and rewrite $A \otimes_{\mathbf{Q}_p} F_0 = \prod_{i=1}^f A_i$ with $A_i = A_{\varsigma_0^i} = A$. We deduce that we have isomorphisms

$$\begin{aligned} (\text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d)_L &= \prod_{i=1}^f \text{GL}_d \\ (\text{Res}_{F_0/\mathbf{Q}_p} \text{Mat}_{d \times d})_L &= \prod_{i=1}^f \text{Mat}_{d \times d}. \end{aligned}$$

Note that $\text{WD}_{d,F/F} = X_d/\text{GL}_d$, where

$$X_d = \{(\varphi, N) \in \text{GL}_d \times \text{Mat}_{d \times d} \mid N\varphi = q\varphi N\}/\text{GL}_d,$$

and where GL_d acts on GL_d and $\text{Mat}_{d \times d}$ by conjugation. On the other hand $\text{Mod}_{d,\varphi,N,F/F} = Y_d/\text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d$, where

$$Y_d = (\{(\varphi, N) \in \text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d \times \text{Res}_{F_0/\mathbf{Q}_p} \text{Mat}_{d \times d} \mid N\varphi = p\varphi\sigma^* N\})/\text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d,$$

and where $\sigma^*(N_1, \dots, N_f) = (N_f, N_1, \dots, N_{f-1})$ and $\text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d$ acts by conjugation on $\text{Res}_{F_0/\mathbf{Q}_p} \text{Mat}_{d \times d}$ and acts on $\text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d$ by σ -conjugation, i.e. $g \cdot (\varphi, N) = (g^{-1}\varphi\sigma(g), g^{-1}Ng)$. In other words

$$\begin{aligned} &(g_1, \dots, g_f) \cdot ((\varphi_1, \dots, \varphi_f), (N_1, \dots, N_f)) \\ &= ((g_1^{-1}\varphi_1 g_1, g_2^{-1}\varphi_2 g_1, \dots, g_f^{-1}\varphi_f g_{f-1}), (g_1^{-1}N_1 g_1, \dots, g_f^{-1}N_f g_f)). \end{aligned}$$

Then there is a canonical morphism $X_d \rightarrow Y_d$ given by

$$(\varphi, N) \mapsto ((\varphi, 1, \dots, 1), (N_1, p^{-1}N_1, \dots, p^{-(f-1)}N_1))$$

which is equivariant for the map of group schemes $\text{GL}_d \rightarrow \text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d$ given by

$$g \mapsto (g, g, \dots, g)$$

and which induces the isomorphism of stacks

$$f_{\varsigma_0} : \text{WD}_{d,F/F} = X_d/\text{GL}_d \xrightarrow{\sim} \text{Mod}_{d,\varphi,N,F/F} = Y_d/\text{Res}_{F_0/\mathbf{Q}_p} \text{GL}_d$$

in Lemma 5.2.1. The inverse of this isomorphism is induced by the map $Y_d \rightarrow X_d$ given by

$$((\varphi_1, \dots, \varphi_f), (N_1, \dots, N_f)) \mapsto (\varphi_f \cdots \varphi_2 \varphi_1, N_1).$$

These formulas show that the isomorphism f_{ς_0} depends on the choice of ς_0 : for a different choice ς'_0 the induced automorphism $f_{\varsigma'_0} \circ f_{\varsigma_0}^{-1}$ is not the identity. Moreover, we easily see that the isomorphism f_{ς_0} does not descend to a subfield $L' \subset L$ not containing F_0 : if $\tau : L \rightarrow L$ is a field automorphism of L , then it commutes with f_{ς_0} if and only if its restriction to F_0 is the identity.

On the other hand we note that the map

$$\text{WD}_{d,F/F} \rightarrow \text{GL}_d//\text{GL}_d \cong \mathbf{A}^{d-1} \times \mathbf{G}_m$$

mapping (φ, N) to the coefficients of its characteristic polynomial induces an isomorphism on global sections, see e.g. [Hel20, Remark 3.4]. We deduce that f_{σ_0} induces an isomorphism on the (invariant) global sections

$$(5.2.7) \quad \Gamma(\mathrm{Mod}_{d,\varphi,N,F/F}, \mathcal{O}_{\mathrm{Mod}_{d,\varphi,N,F/F}}) \longrightarrow \Gamma(\mathrm{WD}_{d,F/F}, \mathcal{O}_{\mathrm{WD}_{d,F/F}})$$

that is independent of the choice of ς_0 and that, moreover, this isomorphism descends to any subfield $L' \subseteq L$. We now make this descent explicit in the case $L = F_0$ and $L' = \mathbf{Q}_p$. Then the action σ_{Y_d} of the generator $\sigma \in \mathrm{Gal}(F_0/\mathbf{Q}_p)$ on the φ -part in Y_d is given by

$$(\varphi_1, \dots, \varphi_f) \longmapsto (\sigma^* \varphi_f, \sigma^* \varphi_1, \dots, \sigma^* \varphi_{f-1}),$$

whereas the action σ_{X_d} of σ on the φ -part in X_d is clearly given by $\varphi \mapsto \sigma^* \varphi$. We deduce that (spelling out only the φ -part of the formulas)

$$(f_{\varsigma_0} \circ \sigma_{Y_d})(\varphi_1, \dots, \varphi_f) = (\sigma^* \varphi_f)^{-1}(\sigma_{X_d} \circ f_{\varsigma_0}((\varphi_1, \dots, \varphi_f)))(\sigma^* \varphi_f)$$

and hence f_{ς_0} does not descend to \mathbf{Q}_p . On the other hand the characteristic polynomials of the φ -parts of $f_{\varsigma_0} \circ \sigma_{Y_d}$ and $\sigma_{X_d} \circ f_{\varsigma_0}$ agree, and hence the isomorphism between the global sections (5.2.7) descends to \mathbf{Q}_p .

In fact this observation generalizes to $\mathrm{WD}_{d,F'/F}$ as follows:

PROPOSITION 5.2.8. *Let L be any finite extension of \mathbf{Q}_p and let F'/F be a finite extension. For a field extension M of L containing F'_0 and any choice of an embedding $\varsigma_0 : F'_0 \hookrightarrow M$ the isomorphism*

$$\Gamma((\mathrm{Mod}_{d,\varphi,N,F'/F})_M, \mathcal{O}_{(\mathrm{Mod}_{d,\varphi,N,F'/F})_M}) \longrightarrow \Gamma((\mathrm{WD}_{d,F'/F})_M, \mathcal{O}_{(\mathrm{WD}_{d,F'/F})_M})$$

induced by f_{ς_0} is independent of the choice of ς_0 and descends to a canonical isomorphism

$$\Gamma(\mathrm{Mod}_{d,\varphi,N,F'/F}, \mathcal{O}_{\mathrm{Mod}_{d,\varphi,N,F'/F}}) \longrightarrow \Gamma(\mathrm{WD}_{d,F'/F}, \mathcal{O}_{\mathrm{WD}_{d,F'/F}}).$$

5.2.9. Invariant functions and the Bernstein center. Recall that the local Langlands correspondence for $\mathrm{GL}_d(F)$ defines a bijection $\tau \mapsto \Omega_\tau$ between the d -dimensional inertial types τ (defined over L) of I_F and the Bernstein blocks Ω (defined over L) of the category $\mathrm{Rep}^{\mathrm{sm}} \mathrm{GL}_d(F)$ (we view our representations as representations on \bar{L} -vector spaces for an algebraic closure \bar{L} of L). More precisely, if (r, N) is the Weil–Deligne representation associated to an irreducible smooth representation π of $\mathrm{GL}_d(F)$, then $r|_{I_F} \cong \tau$ if and only if π lies in the Bernstein block Ω_τ .

Moreover, note that over an algebraic closure \bar{L} of L there is a bijection $\tau \mapsto \mathrm{WD}_{d,F,\tau}$ between the d -dimensional inertial types τ (defined over \bar{L}) and the connected components of the stack of Weil–Deligne representations, such that a Weil–Deligne representation (r, N) lies in $\mathrm{WD}_{d,F,\tau}$ if and only if $r|_{I_F} \cong \tau$, see [BZCHN20, §5.2]. As we restrict ourselves to inertial types that become trivial on $I_{F'}$ we may choose L to be large enough such that all inertial types τ over \bar{L} (with $\tau|_{I_{F'}}$ being trivial) are defined over L and such that all the geometric connected components $\mathrm{WD}_{d,F,\tau}$ are defined over L , and (by slight abuse of notation) we will still write $\mathrm{WD}_{d,F,\tau}$ for this stack on L -schemes.

The following proposition (which is a consequence of [BZCHN20, Theorem 5.13], see [HM18] for an integral version) relates the Bernstein center $\mathcal{Z}_{\Omega_\tau}$ of Ω_τ (that we view as a block of the category of smooth representations on L -vector spaces, so in particular $\mathcal{Z}_{\Omega_\tau}$ is an L -algebra) to the invariant functions on $\mathrm{WD}_{d,F,\tau}$.

PROPOSITION 5.2.10. *Let τ be an inertial type defined over L . Then there is a natural isomorphism*

$$\mathcal{Z}_{\Omega_\tau} \xrightarrow{\sim} \Gamma(\mathrm{WD}_{d,F,\tau}, \mathcal{O}_{\mathrm{WD}_{d,F,\tau}}).$$

REMARK 5.2.11. Note that there are two possibilities to normalize the isomorphism of Proposition 5.2.10, corresponding to two possible normalizations of the local Langlands correspondence. In the usual (unitary) normalization the central character of a smooth representation π matches the determinant of the corresponding Weil–Deligne representation (ρ, N) . Sometimes it is more convenient to choose a normalization such that (ρ, N) corresponds to $\pi \otimes |\det|^{-\frac{d-1}{2}}$. This normalization is better suitable for local-global compatibility, and has the advantage that it does not depend on a choice of $q^{1/2}$ (and hence has better properties when considered over non-algebraically closed fields, compare [BS07]). In these notes we will from now on always use the *non-unitary* normalization.

For a fixed inertial type τ we can consider the rigid analytic generic fiber $(\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}}$ of the formal stack $\mathcal{X}_d^{\mathrm{ss},\Delta,\tau}$ defined in Section 4. By the very construction of the stacks $\mathcal{X}_d^{\mathrm{ss},\Delta,\tau}$ this generic fiber admits a map

$$(5.2.12) \quad \pi_d^{\mathrm{ss},\Delta,\tau} : (\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}} \longrightarrow (\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,\tau})^{\mathrm{an}} \subset \mathfrak{X}_d,$$

that is a generalization of the period map [PR09, (5.37)] of Pappas–Rapoport in the case $\tau = 1$ and $\underline{\lambda}$ minuscule. Moreover, it fits into a commutative diagram

$$\begin{array}{ccc} (\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}} & \longrightarrow & \mathcal{X}_{d,\eta}^{\mathrm{rig}} \\ \pi_d^{\mathrm{ss},\Delta,\tau} \downarrow & & \downarrow \pi_d \\ (\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,\tau})^{\mathrm{an}} & \longrightarrow & \mathfrak{X}_d \end{array}$$

with the map π_d from (5.1.27). In particular the isomorphism

$$(5.2.13) \quad \mathcal{Z}_{\Omega_\tau} \xrightarrow{\sim} \Gamma(\mathrm{Mod}_{d,\varphi,N,\tau}, \mathcal{O}_{\mathrm{Mod}_{d,\varphi,N,\tau}})$$

obtained by combining Propositions 5.2.8 and 5.2.10, induces a map

$$(5.2.14) \quad \mathcal{Z}_{\Omega_\tau} \longrightarrow \Gamma((\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}}, \mathcal{O}_{(\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}}})$$

that is defined over any chosen base field L over which τ is defined (not necessarily containing the field F'_0 over which we can identify the stack of $(\varphi, N, \mathrm{Gal}(F'/F))$ -modules with a substack of the stack of Weil–Deligne representations).

REMARK 5.2.15. After composing with the map to the functions on the generic fiber of a potentially semi-stable deformation ring this map coincides with the morphism of [CEG⁺16, Theorem 4.1, Prop. 4.2]. In fact the result of [CEG⁺16] is stronger: it asserts that the map takes values in the ring of bounded functions

$$\Gamma(\mathcal{X}_d^{\mathrm{ss},\Delta,\tau}, \mathcal{O}_{\mathcal{X}_d^{\mathrm{ss},\Delta,\tau}})[\frac{1}{p}]) \subset \Gamma((\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}}, \mathcal{O}_{(\mathcal{X}_d^{\mathrm{ss},\Delta,\tau})_\eta^{\mathrm{rig}}}).$$

5.3. Drinfeld-style compactifications. We discuss variants of the stack \mathfrak{X}_d parameterizing flags on a family of equivariant vector bundles of rank d . These stacks are naturally associated with a choice of a parabolic subgroup $P \subseteq G = \mathrm{GL}_d$. Given a parabolic subgroup P with Levi quotient $M \cong \mathrm{GL}_{d_1} \times \cdots \times \mathrm{GL}_{d_r}$, we define

\mathfrak{X}_P to be the stack of Gal_F -equivariant P -bundles on the Fargues–Fontaine curve $X_{\bar{F}}$, which comes with two projections

$$\begin{array}{ccc} & \mathfrak{X}_P & \\ \beta_P \swarrow & & \searrow \alpha_P \\ \mathfrak{X}_d & & \mathfrak{X}_M \end{array}$$

induced by the maps $P \hookrightarrow \mathrm{GL}_d$ and $P \twoheadrightarrow M$, where $\mathfrak{X}_M \cong \mathfrak{X}_{d_1} \times \cdots \times \mathfrak{X}_{d_r}$ denotes the stack of Gal_F -equivariant M -bundles on $X_{\bar{F}}$. It will be convenient to write $\mathfrak{X}_G = \mathfrak{X}_d$ in the following.

In the special case of $P = B$ the (φ, Γ) -modules in the image of the morphism $\mathfrak{X}_B \rightarrow \mathfrak{X}_d$ are, by definition, precisely the (φ, Γ) -modules that are trianguline (in the sense of Colmez, see [Col08, Col14] for the definition and discussion of two dimensional trianguline representations and their role in the p -adic Langlands correspondence). Beside the fact that trianguline (φ, Γ) -modules are very natural objects to consider (as a triangulation is nothing but a B -structure), their definition was inspired by the fact that the Galois representations associated to overconvergent p -adic modular forms of finite slope (i.e. those p -adic modular forms whose eigensystems are parameterized by the Coleman–Mazur eigencurve) are trianguline at p , as discovered by Kisin [Kis03] (before the name “trianguline” was invented and using an equivalent description). This was later extended to more general eigenvarieties (in particular to eigenvarieties for unitary groups, see e.g. [BC09]). We will return to this point in Section 9.6 when we discuss the relation of eigenvarieties and p -adic automorphic forms with sheaves on \mathfrak{X}_d defined using \mathfrak{X}_B and its compactification.

REMARK 5.3.1. There are of course variants, where GL_d is replaced by an arbitrary reductive group, but for simplicity of the exposition we limit ourselves to the case of GL_d here. We define the stacks in a way such that it is clear what the definitions will look like in general.

Following Drinfeld, Braverman and Gaitsgory [BG02] define two types of so-called Drinfeld compactifications of the stack of P -bundles on an algebraic curve, and we will define their analogues $\tilde{\mathfrak{X}}_P$ and $\hat{\mathfrak{X}}_P$ in the context of Gal_F -equivariant vector bundles, respectively in the context of (φ, Γ) -modules, below. Both of these compactifications have their advantages and their disadvantages: the compactification $\tilde{\mathfrak{X}}_P$ has the advantage that it is functorial with respect to inclusions $P' \subseteq P$ of parabolic subgroups, but has the disadvantage that the map $\alpha_P : \mathfrak{X}_P \rightarrow \mathfrak{X}_M$ does not extend to $\tilde{\mathfrak{X}}_P$. Conversely the compactification $\hat{\mathfrak{X}}_P$ has the advantage that α_P extends to $\tilde{\alpha}_P : \tilde{\mathfrak{X}}_P \rightarrow \mathfrak{X}_M$, but $P \mapsto \tilde{\mathfrak{X}}_P$ is not functorial with respect to inclusions $P' \subseteq P$ of parabolic subgroups.

Before giving the formal definitions let briefly explain the motivation in the case $P = B$. A B -structure on a (φ, Γ) -module D_A is a complete (φ, Γ) -stable flag

$$0 = \mathrm{Fil}_0 \subseteq \mathrm{Fil}_1 \subseteq \cdots \subseteq \mathrm{Fil}_d = D_A$$

such that $\mathrm{Fil}_i/\mathrm{Fil}_{i-1}$ is a projective $\mathcal{R}_{F,A}$ -module of rank 1. Using Plücker coordinates the datum of such a filtration is the same as the datum of the (φ, Γ) -stable rank 1 subobjects

$$\mathcal{L}_i = \bigwedge^i \mathrm{Fil}_i \subseteq \bigwedge^i D_A$$

such that the quotient $(\bigwedge^i D_A)/\mathcal{L}_i$ is a projective $\mathcal{R}_{F,A}$ -module. The lines \mathcal{L}_i satisfy some relations (called Plücker relations) that we do not spell out explicitly, but we refer to the more abstract description below that allows a more uniform description of these relations (but has to pay the price that we consider the evaluation of D_A on all algebraic representations of GL_d instead of just the exterior powers $\bigwedge^i D_A$, compare also Remark 5.3.5 below).

The following example illustrates that there are families D_A over rigid spaces $\mathrm{Sp} A$ that admit a B -structure on a Zariski-open and dense subspace, but this B -structure does not extend to all of $\mathrm{Sp} A$. Hence we aim to somehow compactify the situation. The idea is to still parameterize lines $\mathcal{L}_i \subseteq \bigwedge^i D_A$ that satisfy the Plücker relations, but to drop the condition that $(\bigwedge^i D_A)/\mathcal{L}_i$ is projective over $\mathcal{R}_{F,A}$. This obviously fixes the issue in the following example.

EXAMPLE 5.3.2. Let $F = \mathbf{Q}_p$ and consider the following family of filtered φ -modules over the projective line \mathbf{P}^1 . Let

$$M = \mathcal{O}_{\mathbf{P}^1} e_1 \oplus \mathcal{O}_{\mathbf{P}^1} e_2 \text{ and } \varphi_M = \mathrm{diag}(\varphi_1, \varphi_2)$$

for fixed $\varphi_1 \neq \varphi_2 \in \mathbf{Q}_p^\times$. Moreover we consider the \mathbf{Z} -filtration Fil^\bullet on M given by

$$\mathrm{Fil}^i M = \begin{cases} M & i \leq -1 \\ \mathcal{L} & i = 0 \\ 0 & i \geq 1 \end{cases}$$

where $\mathcal{L} \subset M$ is the universal line over \mathbf{P}^1 . We consider the submodule $M' = \mathcal{O}_{\mathbf{P}^1} e_1 \subset M$. Then it is not possible to define a filtration $\mathrm{Fil}^i M'$ on M' (in the sense that $\mathrm{Fil}^i M' \subset M'$ is locally on \mathbf{P}^1 a direct summand) such that for all $x \in \mathbf{P}^1$ the subobject $\mathrm{Fil}^i M' \otimes k(x) \subset M' \otimes k(x)$ is the intersection of $\mathrm{Fil}^i M \otimes k(x)$ with $M' \otimes k(x)$. If we insist that $\mathrm{Fil}^i M' \subset M'$ is (locally) a direct summand, then we can only arrange this on $\mathbf{A}^1 \subset \mathbf{P}^1$ by setting

$$\mathrm{Fil}^i M' = \begin{cases} M' & i \leq -1 \\ 0 & i \geq 0 \end{cases}.$$

However at $x = \infty$ the specialization $\mathrm{Fil}^i M' \otimes k(x)$ is a proper subset of the intersection $\mathrm{Fil}^i M \otimes k(x) \cap M' \otimes k(x)$.

Using the construction of (φ, Γ) -modules associated to filtered φ -modules (see Theorem 5.2.4 and the discussion below it), this translates to the following situation. Let

$$\mathcal{R}_{F, \mathbf{P}^1}^{\oplus 2} \subset D \subset t^{-1} \mathcal{R}_{F, \mathbf{P}^1}^{\oplus 2}$$

be the (φ, Γ) -module over \mathbf{P}^1 associated to M and let \mathcal{L} be the (φ, Γ) -stable line defined by the sub-filtered φ -module $M' \subset M$. Then $\mathcal{L} \subset D$ is a B -structure on $D|_{\mathbf{A}^1}$, but after specializing to $\infty \in \mathbf{P}^1$ the cokernel of this inclusion has t -torsion.

In order to define these compactifications in general we introduce the following notation, compare [BG02]. We fix a Borel subgroup $B \subseteq G = \mathrm{GL}_d$ and write $T = B/U$, where U is the unipotent radical of B . As usual we write $\langle \cdot, \cdot \rangle$ for the canonical pairing between $X_*(T)$ and $X^*(T)$. Given a parabolic subgroup $B \subseteq P \subseteq G$ with unipotent radical U_P and Levi quotient $M = P/U_P$, we write $X^*(T)_M^+ \supset X^*(T)_G^+$ for the semi-groups of dominant weights for M respectively for G . The set of simple coroots $\{\alpha_i \in X_*(T), i \in \mathcal{I}\}$ of G is indexed by the set

$\mathcal{I} = \mathcal{I}_G = \{1, \dots, d-1\}$ of vertices of the Dynkin diagram of G . Given M as above, the Dynkin diagram for M gives rise to a subset $\mathcal{I}_M \subseteq \mathcal{I}$ and we set

$$X^*(T)_{G,P} = \{\check{\lambda} \in X^*(T) \mid \langle \alpha_i, \check{\lambda} \rangle = 0 \text{ for } i \in \mathcal{I}_M\}.$$

Moreover, we write $S_M = M/[M, M]$ for the maximal torus quotient of M .

Given a (family of) Gal_F -equivariant vector bundle(s) \mathcal{E} on $X_{\bar{F}}$ (or a family of (φ, Γ) -modules D over \mathcal{R}_F) and an algebraic representation V of G we write $V(\mathcal{E})$ (respectively $V(D)$) for the Gal_F -equivariant $\text{GL}(V)$ -bundle (respectively a (φ, Γ) -module with $\text{GL}(V)$ -structure) given by the push-forward of \mathcal{E} along $G \rightarrow \text{GL}(V)$. We view $V(\mathcal{E})$ again as a vector bundle instead of a $\text{GL}(V)$ -bundle, and similarly we view $V(D)$ again as a (φ, Γ) -module. For a character $\check{\lambda} \in X^*(T)_G^+$ we write $V^{\check{\lambda}}$ for the irreducible G -representation of highest weight $\check{\lambda}$ and $L^{\check{\lambda}}$ for the canonical representation $\check{\lambda}: T \rightarrow \mathbf{G}_m$ of T .

The following definition of $\bar{\mathfrak{X}}_P^{\text{naive}}$ as well as the definition of $\tilde{\mathfrak{X}}_P^{\text{naive}}$ below are the analogues of the corresponding compactifications of Bun_P defined in [BG02].

DEFINITION 5.3.3. The stack $\bar{\mathfrak{X}}_P^{\text{naive}}$ is the category fibered in groupoids over Rig_L parameterizing triples

$$\bar{\mathfrak{X}}_P^{\text{naive}}(\text{Sp } A) = \left\{ \begin{array}{l} \mathcal{E}_G \in \mathfrak{X}_G(\text{Sp } A), \mathcal{E}_{S_M} \in \mathfrak{X}_{S_M}(\text{Sp } A) \\ (\kappa_{\check{\lambda}} : L^{\check{\lambda}}(\mathcal{E}_{S_M}) \rightarrow V^{\check{\lambda}}(\mathcal{E}_G))_{\check{\lambda} \in X^*(T)_G^+ \cap X^*(T)_{G,P}} \end{array} \right\}$$

such that the restriction of $\kappa_{\check{\lambda}}$ to fibers over $\text{Sp } A$ is injective, the cokernels of the various $\kappa_{\check{\lambda}}$ are flat over $\text{Sp } A$ and such that the following *Plücker relations* are satisfied:

- for $\check{\lambda} = 0$ the map κ_0 is the identity.
- for $\check{\lambda}, \check{\mu} \in X^*(T)_G^+ \cap X^*(T)_{G,P}$ the composition

$$\begin{aligned} L^{\check{\lambda}}(\mathcal{E}_{S_M}) \otimes L^{\check{\mu}}(\mathcal{E}_{S_M}) &= L^{\check{\lambda}+\check{\mu}}(\mathcal{E}_{S_M}) \xrightarrow{\kappa_{\check{\lambda}+\check{\mu}}} V^{\check{\lambda}+\check{\mu}}(\mathcal{E}_G) \\ &\rightarrow (V^{\check{\lambda}} \otimes V^{\check{\mu}})(\mathcal{E}_G) = V^{\check{\lambda}}(\mathcal{E}_G) \otimes V^{\check{\mu}}(\mathcal{E}_G) \end{aligned}$$

coincides with the map $\kappa_{\check{\lambda}} \otimes \kappa_{\check{\mu}}$.

The stack \mathfrak{X}_P of equivariant P -bundles on $X_{\bar{F}}$ can be identified with the open substack of $\bar{\mathfrak{X}}_P^{\text{naive}}$, where the cokernels of the maps $\kappa_{\check{\lambda}}$ are flat over $\text{Sp } A \times X_{\bar{F}}$ (not just over $\text{Sp } A$) for all $\check{\lambda}$. It is obvious from the definition that the maps $\bar{\mathfrak{X}}_P^{\text{naive}} \rightarrow \bar{\mathfrak{X}}_G$ and $\bar{\mathfrak{X}}_P^{\text{naive}} \rightarrow \bar{\mathfrak{X}}_M \rightarrow \bar{\mathfrak{X}}_{S_M}$ extend to morphisms

$$\begin{aligned} \bar{\beta}_P : \bar{\mathfrak{X}}_P &\longrightarrow \bar{\mathfrak{X}}_G \\ \bar{\alpha}_P : \bar{\mathfrak{X}}_P &\longrightarrow \bar{\mathfrak{X}}_{S_M}. \end{aligned}$$

Moreover $P \mapsto \bar{\mathfrak{X}}_P^{\text{naive}}$ is functorial with respect to inclusions $P' \subseteq P$ of parabolic subgroups (compatible with the projection maps $\bar{\alpha}$ and $\bar{\beta}$).

DEFINITION 5.3.4. The stack $\tilde{\mathfrak{X}}_P^{\text{naive}}$ is the category fibered in groupoids over Rig_L parameterizing triples

$$\tilde{\mathfrak{X}}_P^{\text{naive}}(\text{Sp } A) = \left\{ \begin{array}{l} \mathcal{E}_G \in \mathfrak{X}_G(\text{Sp } A), \mathcal{E}_M \in \mathfrak{X}_M(\text{Sp } A) \\ (\kappa_V : V^{U_P}(\mathcal{E}_M) \rightarrow V(\mathcal{E}_G))_V \end{array} \right\}$$

where the maps κ_V are indexed by the finite dimensional G -representations, such that the restrictions of κ_V to the fibers over $\text{Sp } A$ are injective, the cokernels of the

various κ_V are flat over $\mathrm{Sp} A$, and such that the following version of the Plücker relations are satisfied:

- if V is the trivial representation, then the map κ_V is the identity.
- for any G -representations V_1 and V_2 the following diagram commutes:

$$\begin{array}{ccc} (V_1^{U_P} \otimes V_2^{U_P})(\mathcal{E}_M) & \xrightarrow{\kappa_{V_1 \otimes V_2}} & (V_1 \otimes V_2)(\mathcal{E}_G) \\ \downarrow & & \downarrow = \\ (V_1 \otimes V_2)^{U_P}(\mathcal{E}_M) & \xrightarrow{\kappa_{V_1 \otimes V_2}} & (V_1 \otimes V_2)(\mathcal{E}_G) \end{array}$$

- for a morphism of G -representations $V_1 \rightarrow V_2$ the following diagram commutes:

$$\begin{array}{ccc} V_1^{U_P}(\mathcal{E}_M) & \xrightarrow{\kappa_{V_1}} & V_1(\mathcal{E}_G) \\ \downarrow & & \downarrow \\ V_2^{U_P}(\mathcal{E}_M) & \xrightarrow{\kappa_{V_2}} & V_2(\mathcal{E}_G). \end{array}$$

Again \mathfrak{X}_P can be identified with the open substack of $\tilde{\mathfrak{X}}_P^{\mathrm{naive}}$ where the cokernels of the maps κ_V are flat (over $\mathrm{Sp} A \times X_{\bar{F}}$, not just A -flat). Again, it is obvious from the definition that the maps $\mathfrak{X}_P \rightarrow \mathfrak{X}_G$ and $\mathfrak{X}_P \rightarrow \mathfrak{X}_M$ extend to morphisms

$$\begin{aligned} \tilde{\beta}_P : \tilde{\mathfrak{X}}_P^{\mathrm{naive}} &\longrightarrow \mathfrak{X}_G \\ \tilde{\alpha}_P : \tilde{\mathfrak{X}}_P^{\mathrm{naive}} &\longrightarrow \mathfrak{X}_M. \end{aligned}$$

Moreover, there is a forgetful map $\tilde{\mathfrak{X}}_P^{\mathrm{naive}} \rightarrow \bar{\mathfrak{X}}_P^{\mathrm{naive}}$, inducing the identity on \mathfrak{X}_P , that is given by only remembering the maps κ_V for $V = V^\lambda$ with $\lambda \in X^*(T)_G^+ \cap X^*(T)_{G,P}$. In this case the M -representation $(V^\lambda)^{U_P}$ factors through S_M .

REMARK 5.3.5. The stacks $\bar{\mathfrak{X}}_P^{\mathrm{naive}}$ and $\tilde{\mathfrak{X}}_P^{\mathrm{naive}}$ coincide in the case $P = B$, but are distinct otherwise. In our setting with $G = \mathrm{GL}_d$ and in the case $P = B$ it is enough to remember the maps κ_λ for the fundamental weights of GL_d . That is, an $\mathrm{Sp} A$ -valued point of $\bar{\mathfrak{X}}_B$ can be described by tuples $(\mathcal{E}, \mathcal{L}_1, \dots, \mathcal{L}_d, \nu_1, \dots, \nu_d)$, where \mathcal{E} is an equivariant vector bundle of rank d , the \mathcal{L}_i are equivariant line bundles and the ν_i are morphisms

$$\nu_i : \mathcal{L}_i \longrightarrow \bigwedge^i \mathcal{E}$$

that are fiberwise (over $\mathrm{Sp} A$) injective with A -flat cokernel and such that the ν_i satisfy some Plücker relations (that we do not spell out explicitly).

It turns out the the equivariance condition on vector bundles adds a new phenomenon which makes the above definitions too naive (hence the name): the stack \mathfrak{X}_P is not dense in $\bar{\mathfrak{X}}_P^{\mathrm{naive}}$ respectively in $\tilde{\mathfrak{X}}_P$. Roughly the reason for this behavior is the existence of Hodge–Tate–Sen weights. The following example makes this phenomenon more explicit:

EXAMPLE 5.3.6. Let $d = 2$ and $D \in \mathfrak{X}_2(L)$. Then a B -structure on D is given by a subobject $\mathcal{L} = \mathcal{R}_{F,L}(\delta_1) \subset D$ such that the quotient D/\mathcal{L} is a (φ, Γ) -module of rank 1, say $\mathcal{R}_{F,L}(\delta_2)$. It follows that $\mathrm{wt}_\tau(D) = \{\mathrm{wt}_\tau(\delta_1), \mathrm{wt}_\tau(\delta_2)\}$. As the labeled Hodge–Tate weights of both D and the subobject \mathcal{L} vary continuously

on $\overline{\mathfrak{X}}_B^{\text{naive}}$, it follows that on the closure of \mathfrak{X}_B the weight of the subobject \mathcal{L} must be one of the weights of D . On the other hand, by definition of $\overline{\mathfrak{X}}_B$ all the objects $(D, t^m \mathcal{L})$ for $m \geq 0$ lie in \mathfrak{X}_B .

In order to define good compactifications of \mathfrak{X}_P we would like to take its closure inside the naive compactifications in a reasonable way. In the case $P = B$ this is the content of the following theorem that will be proven in [HHS].

THEOREM 5.3.7.

- (i) The stacks \mathfrak{X}_B and $\overline{\mathfrak{X}}_B^{\text{naive}} = \widetilde{\mathfrak{X}}_B^{\text{naive}}$ are rigid analytic Artin stacks.
- (ii) There is a well defined “scheme-theoretic” image $\overline{\mathfrak{X}}_B$ of \mathfrak{X}_B in $\overline{\mathfrak{X}}_B^{\text{naive}}$.

SKETCH OF PROOF: We sketch the proof of the fact that \mathfrak{X}_B is an Artin stack in the 2-dimensional case. The proof can be divided into two steps: in a first step we construct the restriction $\mathfrak{X}_B^{\text{wreg}}$ of \mathfrak{X}_B to a *weakly regular* subset $\mathfrak{X}_T^{\text{wreg}} \subset \mathfrak{X}_T$. This weakly regular subset ensures that all Ext^2 -terms vanish which makes it possible to construct explicit charts. In a second step we use the fact that we can twist away the second cohomology $H_{\varphi, \Gamma}^2$ and use the Beauville–Laszlo gluing lemma to reduce the general case to the first step.

Step 1: Let us write $\mathfrak{X}_T^{\text{wreg}} \subset \mathfrak{X}_T$ for the open subset of characters (δ_1, δ_2) such that $H_{\varphi, \Gamma}^2(\mathcal{R}(\delta_1/\delta_2)) = 0$. If the base field L contains all Galois conjugates of F , then $\mathfrak{X}_T^{\text{wreg}}$ may be described as the set of all (δ_1, δ_2) such that

$$\delta_1/\delta_2 \notin \{\varepsilon z^\lambda \mid \lambda \in \prod_{\tau: F \hookrightarrow L} \mathbf{Z}_+\}.$$

Let $U = \text{Sp}(A)$ be an affinoid open subset of $\mathfrak{X}_T^{\text{wreg}}$. We will write (δ_1, δ_2) for the restriction of the universal characters to U , and use the same notation for the pullback of δ_i along maps from rigid spaces to U . As $H_{\varphi, \Gamma}^2(\mathcal{R}_{A, F}(\delta_1/\delta_2))$ vanishes pointwise on U , there is a quasi-isomorphism

$$[\mathcal{E}^0 \xrightarrow{d} \mathcal{E}^1] \xrightarrow{f^\bullet} (C_A^\bullet, \partial^\bullet)$$

from a perfect complex concentrated in degrees $[0, 1]$ to the Herr-complex C_A^\bullet from Remark 5.1.17 (i) computing $H_{\varphi, \Gamma}^\bullet(\mathcal{R}_{A, F}(\delta_1/\delta_2))$. Moreover, this quasi-isomorphism stays a quasi-isomorphism after any pullback. The map $f^1 : \mathcal{E}^1 \rightarrow \ker \partial^1 \subset C_A^1$ gives rise to a universal extension

$$0 \longrightarrow \mathcal{R}_{Y, F}(\delta_1) \longrightarrow D_Y \longrightarrow \mathcal{R}_{Y, F}(\delta_2) \longrightarrow 0$$

over the geometric vector bundle $Y \rightarrow U$ associated to \mathcal{E}^1 . This extension induces a morphism

$$g : Y \longrightarrow \mathfrak{X}_B|_U.$$

We claim that g is a smooth surjection. Indeed, g is surjective by construction, and relatively representable as the diagonal of \mathfrak{X}_B is representable. It is left to show that g satisfies the infinitesimal lifting criterion. Let $B \twoheadrightarrow \bar{B}$ be a surjection of affinoid A -algebras and consider a commutative diagram

$$\begin{array}{ccc} \text{Sp}(\bar{B}) & \xrightarrow{\bar{h}} & Y \\ \downarrow & & \downarrow g \\ \text{Sp}(B) & \xrightarrow{h} & \mathfrak{X}_B. \end{array}$$

We need to find a lift $\mathrm{Sp}(B) \rightarrow Y$ making the diagram commute. The morphism \bar{h} defines an element $s_{\bar{B}} \in \Gamma(\mathrm{Sp}(\bar{B}), \mathcal{E}_{\bar{B}}^1)$ (where we write $\mathcal{E}_{\bar{B}}^1$ for the pullback of \mathcal{E}^1 to $\mathrm{Sp}(\bar{B})$, and similarly for the other objects involved) and its image $f^1(s_{\bar{B}})$ defines an extension

$$0 \longrightarrow \mathcal{R}_{\bar{B},F}(\delta_1) \longrightarrow D_{\bar{B}} \longrightarrow \mathcal{R}_{\bar{B},F}(\delta_2) \longrightarrow 0.$$

The morphism h gives a lift of the extension $D_{\bar{B}}$ to an extension D_B of $\mathcal{R}_{B,F}(\delta_2)$ by $\mathcal{R}_{B,F}(\delta_1)$ and (after fixing appropriate bases) we may view D_B as a cocycle $c_B^1 \in C_B^1$ lifting $f^1(s_{\bar{B}}) \in C_{\bar{B}}^1$. To find a lift $\mathrm{Sp}(B) \rightarrow Y$ we need to find an element $s_B \in \Gamma(\mathrm{Sp}(B), \mathcal{E}_B^1)$ such that the extension defined by $f^1(s_B)$ is isomorphic to the extension D_B via an isomorphism lifting the canonical identification of $D_{\bar{B}}$ with the extension defined by $f^1(s_{\bar{B}})$. In other words we need to find s_B and $c_B^0 \in C_B^0$ reducing to 0 in $C_{\bar{B}}^0$ such that $c_B^1 - f^1(s_B) = \partial^0(c_B^0)$.

As $\mathcal{E}_{\bar{B}}^1 \rightarrow H^1(\mathcal{E}_{\bar{B}}^\bullet) = H_{\varphi,\Gamma}^1(\mathcal{R}_{\bar{B},F}(\delta_1/\delta_2))$ is surjective, there exists an element $s'_B \in \mathcal{E}_{\bar{B}}^1$ lifting s_B that represents the cohomology class $[c_B^1]$ defined by the extension D_B . Hence there exists some element $c_B^{\prime 0} \in C_B^0$ such that $c_B^1 - f^1(s'_B) = \partial^0(c_B^{\prime 0})$. We need to modify s'_B and $c_B^{\prime 0}$, as a priori $c_B^{\prime 0}$ does not necessarily reduce to 0 over \bar{B} . However, it follows from the construction that its reduction maps to zero under ∂^0 , i.e.

$$c_B^{\prime 0} \in \ker(\partial^0 : C_B^0 \rightarrow C_{\bar{B}}^1) = \ker(\mathcal{E}_{\bar{B}}^0 \rightarrow \mathcal{E}_{\bar{B}}^1).$$

In particular we may regard $c_B^{\prime 0}$ as an element e_B^0 of $\mathcal{E}_{\bar{B}}^0$. Let $e_B^0 \in \mathcal{E}_{\bar{B}}^0$ be a lift of $c_B^{\prime 0}$ and set

$$\begin{aligned} s_B &= s'_B + d(e_B^0), \\ c_B^0 &= c_B^{\prime 0} - f^0(e_B^0). \end{aligned}$$

Then $c_B^0 = c_B^{\prime 0} - f^0(e_B^0) = 0$ and $c_B^1 - f^1(s_B) = \partial^0(c_B^0) - f^1(d(e_B^0)) = \partial^0(c_B^0)$ as desired.

Step 2: Let $U \subset \mathfrak{X}_T$ be an open neighborhood of a point (δ_1, δ_2) that is not weakly regular and that admits a smooth surjection from an affinoid space $\mathrm{Sp}(A)$. Then there exists some $N \gg 0$ such that $H_{\varphi,\Gamma}^2(\mathcal{R}_{A,F}(z^N \delta_1/\delta_2)) = 0$, compare Remark 5.1.17 (iii). We consider the neighborhood $U' = \{(\eta_1, z^N \eta_2) \mid (\eta_1, \eta_2) \in U\}$ of $(\delta'_1, \delta'_2) = (\delta_1, z^N \delta_2) \in \mathfrak{X}_T^{\mathrm{wreg}}$ and the morphism $a : U \rightarrow U'$ given by $(\eta_1, \eta_2) \mapsto (\eta_1, z^N \eta_2)$. This morphism fits in a diagram

$$\begin{array}{ccc} & & V' \\ & & \downarrow \\ \mathfrak{X}_B|_U & \xrightarrow{b} & \mathfrak{X}_B|_{U'} \\ \downarrow & & \downarrow \\ U & \xrightarrow{a} & U', \end{array}$$

where $V' \rightarrow \mathfrak{X}_B|_{U'}$ is a smooth surjection from a rigid analytic space (that exists by Step 1) and where b is the morphism defined by the pullback

$$(5.3.8) \quad \begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{R}(\eta_1) & \longrightarrow & D & \longrightarrow & \mathcal{R}(\eta_2) \longrightarrow 0 \\ & & \parallel & & \uparrow \iota & & \uparrow \\ 0 & \longrightarrow & \mathcal{R}(\eta_1) & \longrightarrow & D' & \longrightarrow & z^N \mathcal{R}(\eta_2) \longrightarrow 0. \end{array}$$

We define $V \rightarrow \mathfrak{X}_B|_U$ as the projection to the first factor of the fiber product

$$V = V' \times_{\mathfrak{X}_B|_U} \mathfrak{X}_B|_U \longrightarrow \mathfrak{X}_B|_U.$$

As $V' \rightarrow \mathfrak{X}_B|_U$ is a smooth surjection we are left to show that b is relatively representable. This is the content of Lemma 5.3.9 below. This finishes the proof that \mathfrak{X}_B is an Artin stack.

We refer to Remark 5.3.10 below for some comments on how to adapt the strategy to prove the corresponding claims for $\overline{\mathfrak{X}}_B$ and $\overline{\mathfrak{X}}_B^{\text{naive}}$. \square

LEMMA 5.3.9. *The morphism $b : \mathfrak{X}_B \rightarrow \mathfrak{X}_B$ defined by mapping an extension D to the pullback D' defined by (5.3.8) is relatively representable.*

SKETCH OF PROOF: Let $V' = \text{Sp}(A) \rightarrow \mathfrak{X}_B$ be a morphism defining the extension D' . For any D as in (5.3.8) the morphism ι induces an isomorphism $D'[1/t] \xrightarrow{\sim} D[1/t]$. Let us write \hat{D} for the scalar extension $\hat{D} = D \otimes_{\mathcal{R}_F} F_\infty[[t]]$ and similarly \hat{D}' . By the Beauville–Laszlo gluing lemma, the (φ, Γ) -module D may be reconstructed from D' and the datum of the Γ -stable $(A \otimes_{\mathbf{Q}_p} F_\infty)[[t]]$ -lattice in the free $(A \otimes_{\mathbf{Q}_p} F_\infty)((t))$ -module $\hat{D}'[1/t]$. We may assume that D' is free and choose a basis of D' such that the first basis vector spans the subobject $\mathcal{R}_{A,F}(\eta_1)$. Using this basis we may view \hat{D} as a Γ -stable A -valued point of the affine Grassmannian

$$\text{Grass}_\infty = L(\text{Res}_{F_\infty/\mathbf{Q}_p} \text{GL}_2)/L^+(\text{Res}_{F_\infty/\mathbf{Q}_p} \text{GL}_2).$$

The lattice \hat{D}' is then identified with the base point $e_0 \in \text{Grass}_\infty$. Moreover the fact that D is an extension as in (5.3.8) translates to the fact that the lattice \hat{D} lies in the closed subspace

$$t^\nu L^{<0} U_\infty t^{-\nu} e_0 \subset \text{Grass}_\infty,$$

where $\nu = (1, -N)$ and where $U_\infty \subset \text{Res}_{F_\infty/\mathbf{Q}_p} \text{GL}_2$ is the Weil restriction of the unipotent upper triangular matrices.

Conversely, the Γ -stable lattices in $t^\nu L^{<0} U_\infty t^{-\nu} e_0$ precisely describe the pullback of $V' \rightarrow \mathfrak{X}_B$ along b .

There are two technical difficulties in this proof:

- (a) As F_∞ is not finite over \mathbf{Q}_p , the Weil restriction $\text{Res}_{F_\infty/\mathbf{Q}_p} \text{GL}_2$ is not representable by a rigid analytic space.
- (b) The subspace $t^\nu L^{<0} U_\infty t^{-\nu} e_0$ is an ind-scheme (or ind-rigid analytic space) rather than a scheme (or a rigid analytic space).

By working a bit more carefully, we can overcome (a) by descending to some finite subextension $F_n = F(\mu_{p^n}) \subset F_\infty$. To overcome (b), we need to show that the Γ -invariant lattices in fact lie inside some closed rigid analytic subspace of $t^\nu L^{<0} U_\infty t^{-\nu} e_0$. In the case $F = \mathbf{Q}_p$ the main ingredients of this argument are as follows (for general F a variant of this argument works as well): on L -valued points either \hat{D} is uniquely determined by \hat{D}' and the diagram (5.3.8), or D' is de Rham up to twist. In the latter case both \hat{D}' and \hat{D} lie in certain affine Schubert cells with respect to a third base point \hat{D}'' in the affine Grassmannian (in the de Rham case the lattice \hat{D}'' can be chosen such that it defines a (φ, Γ) -module with Hodge–Tate weights $(0, 0)$ and the Schubert cells containing \hat{D}' respectively \hat{D} depend on their respective Hodge–Tate weights). \square

REMARK 5.3.10.

- (i) In order to prove that $\overline{\mathfrak{X}}_B^{\text{naive}}$ and $\overline{\mathfrak{X}}_B$ are rigid analytic Artin stacks we follow the same idea as in the above sketch. In order to construct charts for $\overline{\mathfrak{X}}_B$ we have to consider the closure of $t^\nu L^{<0} U_\infty t^{-\nu} e_0$ in the affine Grassmannian. To construct charts for $\overline{\mathfrak{X}}_B^{\text{naive}}$ we have to consider the union of certain $t^\nu L^{<0} U_\infty t^{-\nu} e_0$.
- (ii) The construction of smooth surjections from a rigid analytic space to \mathfrak{X}_B (respectively $\overline{\mathfrak{X}}_B^{\text{naive}}$, $\overline{\mathfrak{X}}_B$) crucially uses the fact that \mathfrak{X}_T is an Artin stack. For a more general parabolic subgroup P we can prove the same result for \mathfrak{X}_P and its compactifications $\overline{\mathfrak{X}}_P$ and $\tilde{\mathfrak{X}}_P$, if we assume that \mathfrak{X}_M is an Artin stack, i.e. if we can prove Conjecture 5.1.18 for M .
- (iii) The maps $\overline{\alpha}_P, \overline{\beta}_P, \tilde{\alpha}_P, \tilde{\beta}_P$ of course restrict to $\overline{\mathfrak{X}}_P$, respectively $\tilde{\mathfrak{X}}_P$ and will be denoted by the same letters in the following.

One important difference with the classical case of vector bundles on algebraic curves is that the map

$$\overline{\beta}_B : \overline{\mathfrak{X}}_B \longrightarrow \mathfrak{X}_G$$

is neither proper nor surjective (and a similar remark applies to more general parabolic subgroups as well). The reason that the map is not surjective is simply that not every Galois representation (or more generally not every (φ, Γ) -module) is trianguline. The easiest way to see this is to compare the dimensions of the stacks \mathfrak{X}_d and \mathfrak{X}_B (or of some related rigid analytic spaces of Galois representations). The failure of properness is explained in Remark 5.3.14 below. We state the following results/expectations that should clarify differences and similarities with the algebraic situation.

THEOREM 5.3.11. *The morphisms*

$$(5.3.12) \quad \begin{aligned} \overline{\mathfrak{X}}_P &\longrightarrow \mathfrak{X}_G \times \mathfrak{X}_{S_M} \text{ and} \\ \tilde{\mathfrak{X}}_P &\longrightarrow \mathfrak{X}_G \times \mathfrak{X}_M \end{aligned}$$

are representable by rigid analytic spaces and proper.

This result is proven using results on triangulations in families by Liu [Liu15] and Kedlaya–Pottharst–Xiao [KPX14]. The following conjecture relates the stack $\overline{\mathfrak{X}}_B$ (or rather its image in $\mathfrak{X}_d \times \mathfrak{X}_T$) to the *trianguline variety* and hence to eigenvarieties and p -adic automorphic forms. Conjectures 9.6.8 and 9.6.27 will make the relation to spaces of p -adic automorphic forms more precise.

CONJECTURE 5.3.13.

- (i) *The scheme-theoretic images of the morphisms (5.3.12) are equidimensional of dimension $\dim \text{Res}_{F/\mathbb{Q}_p} P$.*
- (ii) *Let $P = B$ and let $R_{\bar{\rho}}$ denote the framed deformation ring of a fixed residual representation $\bar{\rho} : \text{Gal}_F \rightarrow \text{GL}_d(k)$. Then the pullback of the scheme theoretic image of (5.3.12) along the smooth map*

$$(\text{Spf } R_{\bar{\rho}})^{\text{rig}} \times \mathcal{T}^d \longrightarrow \mathfrak{X}_G \times \mathfrak{X}_T$$

coincides with the trianguline variety of [BHS17b].

Here scheme-theoretic image should of course be understood as an analogue of the usual notion of scheme-theoretic images in the context of rigid analytic Artin stacks.

REMARK 5.3.14. In fact one expects that the image of $\overline{\mathfrak{X}}_B$ in \mathfrak{X}_G is Zariski-dense: it should be seen as an analogue of the infinite fern of Gouvea–Mazur [GM98] (and its generalization by Chenevier [Che13] and Nakamura [Nak14]). The infinite fern argument implies directly that the image $\overline{\beta}_B(\overline{\mathfrak{X}}_B)$ is dense in the union of all irreducible components of \mathfrak{X}_G that contain a crystalline point in their interior. We expect that all components of \mathfrak{X}_G contain a crystalline point in their interior (which in fact would be a consequence of the conjectural description of the connected components of \mathfrak{X}_G as a consequence of Conjecture 5.1.32). In light of the relation with eigenvarieties the failure of properness can also be seen as consequence of the fact that the construction of eigenvarieties involves Fredholm hypersurfaces of infinite degree over the weight space.

REMARK 5.3.15. From a purely deformation theoretic viewpoint the expected dimension of the stack \mathfrak{X}_P of equivariant P -bundles is $\dim \operatorname{Res}_{F/\mathbf{Q}_p} P$. In fact we expect that \mathfrak{X}_P is equidimensional of this dimension and is a local complete intersection. The p -adic situation seems to be rather different to the situation with the spaces of L -parameters [Hel20, Remark 2.2], [Zhu20, Remark 2.3.8] where the stacks of L -parameters for parabolic subgroups can fail to be equidimensional (whereas the stacks of L -parameters for reductive groups are always equidimensional and complete intersections). In the situation with stacks of L -parameters one is hence forced to work with derived stacks due to this phenomenon. In the p -adic case we expect that the stacks \mathfrak{X}_P (and similarly their compactifications $\overline{\mathfrak{X}}_P$ and $\widetilde{\mathfrak{X}}_P$) do not have any additional non-trivial derived structure. In fact for fixed d one can show that \mathfrak{X}_B is equidimensional if $[F : \mathbf{Q}_p]$ is large (compared with d). For the time being we do not have an argument for all F , but it seems reasonable to expect equidimensionality independent of the degree $[F : \mathbf{Q}_p]$.

Motivated by Conjecture 5.3.13 we write $\mathfrak{X}_{G,\text{tri}} \subset \mathfrak{X}_G \times \mathfrak{X}_T$ for the scheme theoretic image of $\overline{\mathfrak{X}}_B$ under the proper map $\overline{\beta}_B$. The following theorem gives a partial description of the preimages of the morphism $\mathfrak{X}_{G,\text{tri}} \rightarrow \mathfrak{X}_G$ that will turn out to be important in the construction of companion points in eigenvarieties (i.e. in the construction of overconvergent p -adic automorphic forms of finite slope with prescribed associated Galois representation) in Section 9.6. Before describing the preimages we need some preparation.

Let $D \in \mathfrak{X}_G(L)$ be a crystalline (φ, Γ) -module and assume that the eigenvalues $\varphi_1, \dots, \varphi_d$ of the Frobenius on the associated Weil–Deligne representation $\operatorname{WD}(D)$ lie in L and satisfy $\varphi_i/\varphi_j \neq 1, q$ for all $i \neq j$. Moreover, we write $\underline{\lambda} = (\lambda_{\tau,1} \geq \dots \geq \lambda_{\tau,d})_{\tau:F \hookrightarrow L}$ for the Hodge–Tate weights of D and assume that $\underline{\lambda}$ is regular, i.e. we assume $\lambda_{\tau,i} \neq \lambda_{\tau,j}$ for all τ and $i \neq j$. We fix an ordering of the Frobenius eigenvalues $\varphi_1, \dots, \varphi_d$ or equivalently a complete Frobenius stable flag \mathcal{F}_\bullet on $\operatorname{WD}(D)$. For each embedding $\tau : F \hookrightarrow L$ we define the Weyl group element $w_{\mathcal{F},\tau} \in \mathcal{S}_n$ as the relative position of the τ -part of the Hodge filtration

$$\operatorname{Fil}_\tau^\bullet \subseteq \operatorname{WD}(D) = \operatorname{WD}(D)_\tau \subseteq \prod_{\tau':F \hookrightarrow L} \operatorname{WD}(D)_{\tau'} = \operatorname{WD}(D) \otimes_{\mathbf{Q}_p} L$$

with respect to \mathcal{F}_\bullet (note that by the regularity assumption on $\underline{\lambda}$ the filtration $\operatorname{Fil}_\tau^\bullet$ is indeed a full flag).

We associate to the flag \mathcal{F}_\bullet (equivalently, to the ordering of the Frobenius eigenvalues $\varphi_1, \dots, \varphi_d$) the unramified character $\delta_{\mathcal{F}_\bullet} : T = (F^\times)^n \rightarrow L^\times$ given by

$$(5.3.16) \quad (x_1, \dots, x_n) \mapsto \prod_{i=1}^d \text{unr}_{\varphi_i}(x_i).$$

Moreover, associated to a tuple $(\mu_\tau)_\tau \in \mathbf{Z}^{\text{Hom}(F, L)}$ we have the character

$$(5.3.17) \quad z^\mu : x \mapsto \prod_{\tau} \tau(x)^{\mu_\tau}.$$

More generally for $\underline{\mu} = (\mu_1, \dots, \mu_d)$ with $\mu_i \in \mathbf{Z}^{\text{Hom}(F, L)}$ we write

$$z^\mu : (x_1, \dots, x_d) \mapsto \prod_{i=1}^d z^{\mu_i}(x_i) = \prod_{i=1}^d \prod_{\tau} \tau(x_i)^{\mu_{i,\tau}}.$$

The following characterization of the points in $\mathfrak{X}_{G, \text{tri}}$ above D is [BHS19, Theorem 4.2.3].

THEOREM 5.3.18. *Let $D \in \mathfrak{X}_G(L)$ be a crystalline (φ, Γ) -module of regular Hodge–Tate weight $\underline{\lambda}$ and such that the eigenvalues $\varphi_1, \dots, \varphi_d$ of the Frobenius on $\text{WD}(D)$ lie in L and satisfy $\varphi_i/\varphi_j \neq 1, q$ for $i \neq j$. Then*

$$\begin{aligned} & \{\delta \in \mathfrak{X}_T \mid (D, \delta) \in \mathfrak{X}_{G, \text{tri}}\} \\ &= \bigcup_{\mathcal{F}_\bullet} \{\delta_{\mathcal{F}_\bullet} \cdot z^{w\lambda} \mid w = (w_\tau)_\tau \in W \text{ such that } w_{\mathcal{F}, \tau} \preceq w_\tau w_0 \text{ for all } \tau\}, \end{aligned}$$

where the union runs over all Frobenius stable flags \mathcal{F} of $\text{WD}(D)$ and where $w_0 \in \mathcal{S}_d$ is the longest element.

REMARK 5.3.19.

(i) Following Hansen [Han17, Conjecture 6.2.3], there is a general conjectural description of the fibers of $\mathfrak{X}_{G, \text{tri}} \rightarrow \mathfrak{X}_G$ without assuming that the (φ, Γ) -module D is crystalline (and without the regularity assumptions in the theorem).

(ii) In the case of crystabelline (φ, Γ) -modules this description is more or less the same as in the crystalline case. In the semi-stable case (and more generally in the semistabelline case) we only take those flags \mathcal{F}_\bullet that are stable under the monodromy on the Weil–Deligne representation.

(iii) It is reasonable to ask for a generalization from semistabelline to all de Rham (φ, Γ) -modules D . In this case the image of $\overline{\mathfrak{X}}_B$ in $\mathfrak{X}_G \times \mathfrak{X}_T$ should be replaced by the image of $\overline{\mathfrak{X}}_P$ in $\mathfrak{X}_G \times \mathfrak{X}_{S_M}$, where P is the parabolic subgroup containing B with Levi $M = \prod_{i=1}^r \text{GL}_{d_i}$ such that the Weil–Deligne representation $\text{WD}(D)$ associated to D is a direct sum of indecomposable Weil–Deligne representations D_1, \dots, D_r of respective dimension d_i (compare [BD21]).

The following expectation is motivated by the relation of the trianguline variety with eigenvarieties and with the patched version of the eigenvariety [BHS17b] (and by Conjecture 6.2.4(2) respectively Conjecture 9.6.16).

CONJECTURE 5.3.20. *Let us write $\pi_B : \overline{\mathfrak{X}}_B \rightarrow \mathfrak{X}_G \times \mathfrak{X}_T$ for the canonical (proper) projection.*

(i) *There are isomorphisms*

$$R\pi_{B,*} \mathcal{O}_{\overline{\mathfrak{X}}_B} \cong \mathcal{O}_{\mathfrak{X}_{G, \text{tri}}}, \quad R\pi_{B,*} \omega_{\overline{\mathfrak{X}}_B} \cong \omega_{\mathfrak{X}_{G, \text{tri}}}.$$

(ii) The direct image $R\pi_{B,*}(\mathcal{O}_{\overline{\mathfrak{X}}_B}([F : \mathbf{Q}_p]\rho'))$ is concentrated in degree 0 and is a Cohen–Macaulay module (i.e. its dual $R\pi_{B,*}(\omega_{\overline{\mathfrak{X}}_B}(-[F : \mathbf{Q}_p]\rho'))$ is also concentrated in degree 0).

Here ρ' is the algebraic character $(0, -1, \dots, -d + 1)$ of $T = \mathbf{G}_m^d$ that we view as a line bundle on $*/\mathbf{G}_m^d$ and pull it back to $\overline{\mathfrak{X}}_B$ along the canonical map

$$\overline{\mathfrak{X}}_B \rightarrow \mathfrak{X}_T = \mathfrak{X}_1^d = (\mathcal{T}/\mathbf{G}_m)^d \rightarrow */\mathbf{G}_m^d,$$

see Section 7.1 for the identification of \mathfrak{X}_1 .

Moreover, it makes sense to ask whether the trianguline variety $\mathfrak{X}_{G,\text{tri}}$ is normal and Cohen–Macaulay (note that Cohen–Macaulayness would be a consequence of (i) of the conjecture, if $\overline{\mathfrak{X}}_B$ is Cohen–Macaulay). We point out that the sheaf in (ii) of the conjecture should be regarded as a variant of the coherent Springer sheaf from [BZCHN20] (respectively [Zhu20, 4.4] or [Hel20, 2.3]) and hence (ii) should be seen as a variant of [BZCHN20, Conjecture 4.15] (respectively [Zhu20, Conjecture 4.4.2] or [Hel20, Conjecture 2.17]).

In the case $F = \mathbf{Q}_p$ and GL_2 we can prove the above conjecture and compute the singularities as well as the fiber dimensions of the Cohen–Macaulay modules in (ii) (these fiber dimensions have some interpretation in terms of the conjectural correspondence of locally analytic representations with coherent sheaves on \mathfrak{X}_d , Conjecture 6.2.4).

THEOREM 5.3.21. *Let $F = \mathbf{Q}_p$ and $d = 2$, i.e. $G = \text{GL}_2$.*

- (i) *The stack $\overline{\mathfrak{X}}_B$ is smooth.*
- (ii) *The stack $\mathfrak{X}_{G,\text{tri}}$ is a rigid analytic Artin stack⁹ and normal and Cohen–Macaulay.*
- (iii) *The morphism $\pi_B : \overline{\mathfrak{X}}_B \rightarrow \mathfrak{X}_{G,\text{tri}}$ is birational with geometrically connected fibers and the canonical map*

$$\mathcal{O}_{\mathfrak{X}_{G,\text{tri}}} \rightarrow R\pi_{B,*}\mathcal{O}_{\overline{\mathfrak{X}}_B}$$

is an isomorphism. In particular we also have an isomorphism

$$R\pi_{B,*}\omega_{\overline{\mathfrak{X}}_B} \cong \omega_{\mathfrak{X}_{G,\text{tri}}}.$$

(iv) *Let $\mathcal{F} = R\pi_{B,*}(\mathcal{O}_{\overline{\mathfrak{X}}_B}(\rho'))$ and $\mathcal{G} = R\pi_{B,*}(\omega_{\overline{\mathfrak{X}}_B}(-\rho'))$. Then \mathcal{F} and \mathcal{G} are concentrated in degree zero and in particular are Cohen–Macaulay modules.*

(v) *For $x = (D, \delta_1, \delta_2) \in \mathfrak{X}_{G,\text{tri}}$ the fibers $\mathcal{F} \otimes k(x)$, $\mathcal{G} \otimes k(x)$ and $\omega_{\mathfrak{X}_{G,\text{tri}}} \otimes k(x)$ are one-dimensional unless $D = \mathcal{R}(\delta) \oplus t^{-n}\mathcal{R}(\delta)$ with $n \geq 0$ and $\delta_1 = \delta$. In this case*

$$\begin{aligned} \dim \mathcal{F} \otimes k(x) &= 2 \\ \dim \mathcal{G} \otimes k(x) &= \begin{cases} 2 & n = 0 \\ 1 & n \geq 1 \end{cases} \\ \dim \omega_{\mathfrak{X}_{G,\text{tri}}} \otimes k(x) &= \begin{cases} 3 & n = 0 \\ 2 & n \geq 1. \end{cases} \end{aligned}$$

⁹As we do not know whether $\mathfrak{X}_G = \mathfrak{X}_2$ is a rigid analytic Artin stack this should be interpreted as: there is a rigid analytic Artin stack $\mathfrak{X}_{G,\text{tri}}$ that embeds as a closed substack into $\mathfrak{X}_G \times \mathfrak{X}_T$. The map $\overline{\mathfrak{X}}_B \rightarrow \mathfrak{X}_G \times \mathfrak{X}_T$ factors through $\mathfrak{X}_{G,\text{tri}}$ and is scheme-theoretically dense.

5.3.22. *Local models.* Assume for the remainder of this section that L contains all Galois conjugates of F (inside \bar{F}). We discuss the relation of stacks of (φ, Γ) -modules over \mathcal{R}_F with stacks of (semi-linear) Γ -representations on $F_\infty[[t]]$. We don't know whether it can be expected that the latter stacks are Artin stacks (it seems that they can be rather ill-behaved), but their completions at a fixed Hodge–Tate weight turn out to be well-behaved. At certain “nice” points the morphism from a stack of (φ, Γ) -modules to these stacks is formally smooth and hence we obtain rather explicit local models for stacks like $\bar{\mathfrak{X}}_B$, generalizing the local models for the trianguline variety of [BHS19].

We define the groupoid $\mathfrak{X}_d^{\text{dR},+}$ over Rig_L as follows. For an affinoid algebra A we set

$$\mathfrak{X}_d^{\text{dR},+}(\text{Sp } A) = \left\{ \begin{array}{l} \text{continuous, semi-linear } \Gamma\text{-representations} \\ \text{on finite projective } A \widehat{\otimes} F_\infty[[t]]\text{-modules of rank } d \end{array} \right\},$$

where (by slight abuse of notation) we write $A \widehat{\otimes} F_\infty[[t]] = (A \otimes_{\mathbf{Q}_p} F_\infty)[[t]]$.

REMARK 5.3.23.

- (i) In a similar way we can define a groupoid $\mathfrak{X}_d^{\text{dR}}$ by mapping $\text{Sp } A$ to the groupoid of finite projective $A \widehat{\otimes} F_\infty((t))$ -modules of rank d with a semi-linear Γ -action that (locally on $\text{Sp } A$) admit an $A \widehat{\otimes} F_\infty[[t]]$ -lattice.
- (ii) The motivation to consider these stacks (and the explanation for their name) is Fontaine's theorem [Fon04] that the category of semi-linear Γ -representations on finite free $F_\infty[[t]]$ -modules (respectively on finite dimensional $F_\infty((t))$ -vector spaces) is equivalent to the category of semi-linear Gal_F -representations on finite free B_{dR}^+ -modules (respectively on finite dimensional B_{dR} -vector spaces). We expect that a similar equivalence holds true in families over rigid analytic spaces, but, as we do not have a proof, we restrict our attention to the semi-linear Γ -representations, as they seem to be easier to handle.

EXPECTATION 5.3.24.

- (i) The groupoids $\mathfrak{X}_d^{\text{dR},+}$ and $\mathfrak{X}_d^{\text{dR}}$ are stacks.
- (ii) The diagonal of $\mathfrak{X}_d^{\text{dR},+}$ is representable.
- (iii) The morphism $\mathfrak{X}_d^{\text{dR},+} \rightarrow \mathfrak{X}_d^{\text{dR}}$ is representable.

Though we do not know whether one should expect that $\mathfrak{X}_d^{\text{dR},+}$ and $\mathfrak{X}_d^{\text{dR}}$ are Artin stacks, both stacks admit versal rings at rigid analytic points, and more generally their completions along a fixed Hodge–Tate–Sen weight are formal rigid analytic Artin stacks (a notion that we will not formally introduce here). In order to make sense of this statement, we note that, similarly to the definition of the weight map ω_d in (5.1.31), we have a weight map

$$\omega_d^{\text{dR}} : \mathfrak{X}_d^{\text{dR},+} \longrightarrow \text{WT}_{d,F} \cong \prod_{F \rightarrow L} \mathbf{A}^d / \mathcal{S}_d$$

that maps, on A -valued points, an $A \widehat{\otimes} F_\infty[[t]]$ -module D_A to the characteristic polynomial of the derivative of the Γ -action on D_A/tD_A (recall that we have fixed L large enough such that $[F : \mathbf{Q}_p] = |\text{Hom}(F, L)|$).

REMARK 5.3.25. The Hodge–Tate–Sen weights of an object $D_A \in \mathfrak{X}_d^{\text{dR}}(\text{Sp } A)$ are only well defined mod \mathbf{Z} . More precisely, under the identification of $A \otimes_{\mathbf{Q}_p} F =$

$\prod_{F \hookrightarrow L} A$ we can define the weight of D_A as an element of

$$\prod_{F \hookrightarrow L} ((A/\mathbf{Z})^d / \mathcal{S}_d).$$

Fontaine's theory of almost de Rham representations gives a description of the stacks $\mathfrak{X}_d^{\text{dR}}$ and $\mathfrak{X}_d^{\text{dR},+}$ after completion at a fixed integral Hodge–Tate weight, and similar results hold true for the completion at any fixed Hodge–Tate–Sen weight. We briefly recall Fontaine's classification [Fon04] of B_{dR} -representations (or equivalently $F_\infty((t))$ -representations). Let us write $\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F}$ for the category of continuous semi-linear Gal_F -representations on finite dimensional B_{dR} vector spaces, and similarly we define $\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F}$, $\text{Mod}_{F_\infty((t))}^\Gamma$ and $\text{Mod}_{F_\infty[[t]]}^\Gamma$ in the obvious way. For a Hodge–Tate–Sen weight $a \in F$ we denote by $\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F, a}$ the full subcategory of all objects whose Hodge–Tate–Sen weights are congruent to a modulo \mathbf{Z} . Then

$$\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F} = \bigoplus_{a \in F/\mathbf{Z}} \text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F, a}$$

is a block decomposition of $\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F}$, i.e. the higher Ext-groups between objects in different summands vanish (this is a consequence of [Fon04, Thm. 3.19]).

The block $\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F, 0}$ is called the block of almost de Rham representations (or *presque de Rham* representation in French). It has the following description in terms of linear algebra objects (see [BHS19, Prop 3.1.1]):

THEOREM 5.3.26 (Fontaine). *There is an equivalence of categories*

$$D_{\text{pdR}} : \text{Mod}_{B_{\text{dR}}^+}^{\text{Gal}_F, 0} \longrightarrow \text{FilRep}(\mathbf{G}_a)$$

from the category of almost de Rham representations over B_{dR}^+ to the category of finite dimensional F -vector spaces equipped with an algebraic representation of the additive group \mathbf{G}_a and with a separated and exhaustive \mathbf{Z} -filtration Fil^\bullet stable under the \mathbf{G}_a -action. Under this equivalence, passing from $\text{Mod}_{B_{\text{dR}}^+}^{\text{Gal}_F, 0}$ to $\text{Mod}_{B_{\text{dR}}}^{\text{Gal}_F, 0}$ corresponds to forgetting the filtration.

We recall that (over fields of characteristic 0) a \mathbf{G}_a -representation is nothing but the action of a nilpotent operator. Moreover, an almost de Rham representation V over B_{dR}^+ (or over B_{dR}) is de Rham if and only if the nilpotent operator $N : D_{\text{pdR}}(V) \rightarrow D_{\text{pdR}}(V)$ (defined by the \mathbf{G}_a -action) is zero.

There is a definition of the functor D_{pdR} in the spirit of period rings [BHS19, 3.1] that we do not recall here. On the equivalent category $\text{Mod}_{F_\infty[[t]]}^\Gamma \cong \text{Mod}_{B_{\text{dR}}^+}^{\text{Gal}_F}$ one can characterize almost de Rham representations and the functor D_{pdR} as follows: Given an $F_\infty[[t]]$ -module Λ with continuous Γ -action, we can consider (similarly to the definition of Hodge–Tate–Sen weights) the derivation of the Γ -action at 1 which gives a derivation

$$\partial_\Lambda : \Lambda \longrightarrow \Lambda$$

above the derivation $t \frac{d}{dt}$ on $F_\infty[[t]]$. We can then define

$$D_{\text{pdR}, \infty}(\Lambda) = \bigcup_{N \geq 1} \ker(\partial_\Lambda^N : \Lambda[1/t] \rightarrow \Lambda[1/t])$$

as the sub- F_∞ vector space of $\Lambda[1/t]$ on which ∂_Λ is nilpotent. In particular ∂_Λ induces an F_∞ -linear, nilpotent endomorphism N_∞ of $D_{\text{pdR},\infty}(\Lambda)$. The representation Λ is almost de Rham if and only if the canonical map

$$D_{\text{pdR},\infty}(\Lambda) \otimes_{F_\infty} F_\infty((t)) \longrightarrow \Lambda[1/t]$$

is an isomorphism. In that case we write $\Lambda_0 = D_{\text{pdR},\infty}(\Lambda) \otimes_{F_\infty} F_\infty[[t]]$ and refer to Λ_0 as the standard lattice. Then the lattice $\Lambda \subset \Lambda_0[1/t]$ defines a separated and exhaustive \mathbf{Z} -filtration $\text{Fil}_\infty^\bullet$ on $D_{\text{pdR},\infty}(\Lambda) = \Lambda_0/t\Lambda_0$ by sub- F_∞ vector spaces that are stable under the action of N_∞ , as Λ is stable under ∂_Λ . Again in a similar fashion to the definition of Hodge–Tate–Sen weights, the triple $(D_{\text{pdR},\infty}(\Lambda), N_\infty, \text{Fil}_\infty^\bullet)$ is equipped with a Γ -action which allows us to descend¹⁰ $(D_{\text{pdR},\infty}(\Lambda), N_\infty, \text{Fil}_\infty^\bullet)$ from F_∞ to the desired triple $(D_{\text{pdR}}(\Lambda), N, \text{Fil}^\bullet)$ over F .

The equivalence in Fontaine’s theorem lies at the heart of the following description of the completions of the stacks of B_{dR}^+ -representations. Note that for an L -algebra A we have

$$F_\infty[[t]] \widehat{\otimes}_{\mathbf{Q}_p} A = \prod_{\tau: F \hookrightarrow L} F_\infty[[t]] \widehat{\otimes}_F A,$$

and hence we can apply Fontaine’s description for each embedding τ and for each labeled Hodge–Tate–Sen weight separately. Let us write $G = \text{Res}_{F/\mathbf{Q}_p} \text{GL}_d$ and $\mathfrak{g} = \text{Lie } G$ for the remainder of this subsection. Moreover, we fix an integral labeled Hodge–Tate weight

$$\underline{\lambda} \in \prod_{\tau: F \hookrightarrow L} (\mathbf{Z}^d) / \mathcal{S}_d \subseteq \text{WT}_{d,F}(L)$$

We denote by $[\underline{\lambda}]$ its class modulo \mathbf{Z} (which of course only remembers that the weight was integral but does not distinguish between various $\underline{\lambda}$).

We describe the formal completions $(\mathfrak{X}_d^{\text{dR},+})_{[\underline{\lambda}]}$ and $(\mathfrak{X}_d^{\text{dR}})_{[\underline{\lambda}]}$ of the groupoids $\mathfrak{X}_d^{\text{dR},+}$ and $\mathfrak{X}_d^{\text{dR}}$ along the fixed weight $\underline{\lambda}$ (respectively along the locus where the Hodge–Tate–Sen weight is congruent to $\underline{\lambda}$ modulo \mathbf{Z}). This result is very similar to [BHS19, Cor. 3.1.6, Cor. 3.1.9]. Let

$$\widetilde{\mathfrak{g}}_{P_{\underline{\lambda}}} = \{(A, \mathcal{F}^\bullet) \in \mathfrak{g} \times G/P_{\underline{\lambda}} \mid A\mathcal{F}^\bullet \subset \mathcal{F}^\bullet\} \subseteq \mathfrak{g} \times G/P_{\underline{\lambda}}$$

be the variant of Grothendieck’s simultaneous resolution of singularities that parameterizes pairs (A, \mathcal{F}^\bullet) consisting of an endomorphism A and an A -stable flag \mathcal{F}^\bullet of type $P_{\underline{\lambda}}$ (compare Section 5.2 for the notation). Recall that $\widetilde{\mathfrak{g}}_{P_{\underline{\lambda}}}$ is a subvariety of $\prod_\tau \mathfrak{gl}_d \times \text{GL}_d/P_{\lambda_\tau}$ and again we think of the filtration on $G/P_{\underline{\lambda}}$ as an $[F : \mathbf{Q}_p]$ -tuple of \mathbf{Z} -filtrations that jump in degrees $\underline{\lambda}$.

PROPOSITION 5.3.27. *The functor $D_{\text{pdR},\infty}$ induces isomorphisms*

$$\begin{aligned} (\mathfrak{X}_d^{\text{dR}})_{[\underline{\lambda}]} &\xrightarrow{\sim} \widehat{\mathfrak{g}}_0/G \\ (\mathfrak{X}_d^{\text{dR},+})_{\underline{\lambda}} &\xrightarrow{\sim} (\widetilde{\mathfrak{g}}_{P_{\underline{\lambda}}})_0/G, \end{aligned}$$

where on the right hand side we complete at the closed subspace of all nilpotent endomorphisms (i.e. at the nilpotent cone, respectively its partial Springer resolution).

¹⁰We caution the reader that this descent is not completely formal and requires some work. Recall that we pointed out in the discussion of Hodge–Tate weights that $(D_{\text{Sen}}(D_A), \Theta)$ does not descend to an object over $F \otimes_{\mathbf{Q}_p} A$. However, this can be done after completion at a fixed integral Hodge–Tate weight. The issue here is a similar one.

REMARK 5.3.28.

(a) The proposition implies that the completion of $\mathfrak{X}_d^{\mathrm{dR},+}$ at an integral Hodge–Tate weight is a formal rigid analytic Artin stack (with an appropriate definition of such an object).

(b) In fact, applying Fontaine’s theory for arbitrary (i.e. not necessarily integral) weights $\underline{\lambda}$ one can obtain a similar description of the completion of $\mathfrak{X}_d^{\mathrm{dR},+}$ respectively $\mathfrak{X}_d^{\mathrm{dR}}$ at $\underline{\lambda}$.

In the neighborhood of certain “nice” points, $\mathfrak{X}_d^{\mathrm{dR},+}$ can be thought of as a local model of the stack \mathfrak{X}_d . A sample result for this is the following theorem, compare [HMS18, Theorem 3.1.1]. We will discuss the related situation for B -structures (resp. P -structures) and their Drinfeld compactifications below.

THEOREM 5.3.29. *Let $x \in \mathfrak{X}_d$ be a crystalline point with regular Hodge–Tate weight and Frobenius eigenvalues $\varphi_1, \dots, \varphi_d$ that satisfy $\varphi_i/\varphi_j \notin \{1, q\}$. Then the morphism $\mathfrak{X}_d \rightarrow \mathfrak{X}_d^{\mathrm{dR},+}$ is formally smooth at x .*

Given a parabolic subgroup $P \subseteq \mathrm{GL}_d$ we can consider P -structures (i.e. partial flags) on objects in $\mathfrak{X}_d^{\mathrm{dR}}$ and $\mathfrak{X}_d^{\mathrm{dR},+}$, which allows us to define stacks $\mathfrak{X}_P^{\mathrm{dR}}$ and $\mathfrak{X}_P^{\mathrm{dR},+}$. Parallel to the above discussion of Drinfeld compactifications, we would like to define “compactifications” $\overline{\mathfrak{X}}_P^{\mathrm{dR},+}$ and $\widetilde{\mathfrak{X}}_P^{\mathrm{dR},+}$ of the stack $\mathfrak{X}_P^{\mathrm{dR},+}$ that fit into the diagram

$$\begin{array}{ccc}
 & \widetilde{\mathfrak{X}}_P^{\mathrm{dR},+} & \\
 & \downarrow & \searrow \\
 & \overline{\mathfrak{X}}_P^{\mathrm{dR},+} & \mathfrak{X}_M^{\mathrm{dR},+} \\
 \swarrow & & \downarrow \\
 \mathfrak{X}_d^{\mathrm{dR},+} & & \mathfrak{X}_{S_M}^{\mathrm{dR},+}
 \end{array}$$

We will define such compactifications, and prove finiteness statements about them after formal completion at an integral weight $\underline{\lambda}$.

As above we fix an integral labeled Hodge–Tate weight $\underline{\lambda} = ((\lambda_{\tau,1}, \dots, \lambda_{\tau,d})_{\tau:F \hookrightarrow L})$. We restrict ourselves to the easiest case in which $\underline{\lambda}$ is regular. As $\underline{\lambda}$ is only well defined up to permutation, we choose its dominant representative, i.e. we assume $\lambda_{\tau,1} > \lambda_{\tau,2} > \dots > \lambda_{\tau,d}$ for all τ . In particular the parabolic subgroup $P_{\underline{\lambda}}$ equals the Weil restriction of the Borel subgroup $(\mathrm{Res}_{F/\mathbf{Q}_p} B)_L$. We write $\widetilde{\mathfrak{g}} = \widetilde{\mathfrak{g}}_{P_{\underline{\lambda}}}$ for simplicity.

We now give a description similar to Proposition 5.3.27 of the completion of $\mathfrak{X}_P^{\mathrm{dR},+}$ and use this description to define “compactifications” of the formal completion. Consider the space

$$(5.3.30) \quad X_{B,P} = \widetilde{\mathfrak{g}} \times_{\mathfrak{g}} \widetilde{\mathfrak{g}}_P = \bigcup_{w \in W_P} V_w \subseteq \mathfrak{g} \times G/B \times G/P,$$

where $W = \prod_{\tau} \mathcal{S}_d$ is the Weil group of $\mathrm{Res}_{F/\mathbf{Q}_p} \mathrm{GL}_d$ and $W_P \subseteq W$ is the set of minimal length representatives of elements in W/W_M ($W_M \subseteq W$ being the Weyl group of the Levi $\mathrm{Res}_{F/\mathbf{Q}_p} M$ of $\mathrm{Res}_{F/\mathbf{Q}_p} P$). For $w = (w_{\tau})_{\tau} \in W_P$, the space V_w is defined to be the preimage of the Bruhat cell $G(1, w) \subseteq G/\mathrm{Res}_{F/\mathbf{Q}_p} B \times$

$G/\mathrm{Res}_{F/\mathbf{Q}_p}P$ (recall that we write $G = \mathrm{Res}_{F/\mathbf{Q}_p}\mathrm{GL}_d$). Over V_w , the intersection of the universal flag on $G/\mathrm{Res}_{F/\mathbf{Q}_p}B$ and over $G/\mathrm{Res}_{F/\mathbf{Q}_p}P$ defines a well defined flag in $\mathrm{Res}_{F/\mathbf{Q}_p}M/\mathrm{Res}_{F/\mathbf{Q}_p}B_M$. Given w , we write $w\lambda$ for the permutation of λ by w . Moreover, we can view $w\lambda$ as a weight of M . In particular for each $w \in W$, we obtain a diagram

$$(5.3.31) \quad \begin{array}{ccc} & V_w/G & \\ \swarrow & & \searrow \\ \tilde{\mathfrak{g}}/G & & \tilde{\mathfrak{m}}/\mathrm{Res}_{F/\mathbf{Q}_p}M, \end{array}$$

where by abuse of notation we write $\tilde{\mathfrak{m}}$ for the Lie algebra of $\mathrm{Res}_{F/\mathbf{Q}_p}M$.

PROPOSITION 5.3.32. *The functor D_{pdR} induces an isomorphism of the diagram*

$$\begin{array}{ccc} & (\mathfrak{X}_P^{\mathrm{dR},+})_{ww_0\lambda}^\wedge & \\ \swarrow & & \searrow \\ (\mathfrak{X}_d^{\mathrm{dR},+})_{\lambda}^\wedge & & (\mathfrak{X}_M^{\mathrm{dR},+})_{ww_0\lambda}^\wedge \end{array}$$

with the completion of (5.3.31) at the closed subspace of nilpotent endomorphisms. Here $(\mathfrak{X}_P^{\mathrm{dR},+})_{ww_0\lambda}^\wedge$ is the completion of $\mathfrak{X}_P^{\mathrm{dR},+}$ along the preimage of

$$(\mathfrak{X}_M^{\mathrm{dR},+})_{ww_0\lambda}^\wedge \subseteq (\mathfrak{X}_M^{\mathrm{dR},+})$$

under the canonical map $\mathfrak{X}_P^{\mathrm{dR},+} \rightarrow \mathfrak{X}_M^{\mathrm{dR},+}$.

In the diagram (5.3.31) there are two ways to complete V_w/G into a stack proper over $\tilde{\mathfrak{g}}/G$, and again only one of these compactifications will admit a morphism to $\tilde{\mathfrak{m}}/M$. The easiest way is to consider the Zariski-closure \overline{X}_w of V_w inside the space $X_{B,P} = \tilde{\mathfrak{g}} \times_{\mathfrak{g}} \tilde{\mathfrak{g}}_P$ defined in (5.3.30). In this case the intersection of the two flags over $G/\mathrm{Res}_{F/\mathbf{Q}_p}B$ and $G/\mathrm{Res}_{F/\mathbf{Q}_p}P$ does not induce a well defined family of filtrations on the graded pieces (with respect to the P -filtration), but only on their top exterior powers. We obtain a diagram

$$\begin{array}{ccc} & \overline{X}_w/G & \\ \swarrow & & \searrow \\ \tilde{\mathfrak{g}}/G & & \mathrm{Lie} S_M/S_M. \end{array}$$

Note that the formation of \overline{X}_w is functorial with respect to inclusions $P' \subseteq P$ as can be deduced from the identification

$$X_{B,P}/G = q^{-1}(\mathfrak{p})/\mathrm{Res}_{F/\mathbf{Q}_p}P,$$

where $q : \tilde{\mathfrak{g}} \rightarrow \mathfrak{g}$ is the canonical projection and $\mathfrak{p} \subseteq \mathfrak{g}$ is the Lie algebra of $\mathrm{Res}_{F/\mathbf{Q}_p}P \subseteq G$. The stack \overline{X}_w/G is canonically isomorphic to $\overline{X}'_w/\mathrm{Res}_{F/\mathbf{Q}_p}P$, where $\overline{X}'_w \subseteq q^{-1}(\mathfrak{p})$ is the closure of the locally closed subset $V'_w \subseteq q^{-1}(\mathfrak{p})$ such that $V_w/G = V'_w/\mathrm{Res}_{F/\mathbf{Q}_p}P$.

Note that for fixed w the intersection of the universal flag on $G/\mathrm{Res}_{F/\mathbf{Q}_p}B$ with the partial standard flag (of type P) gives rise to a well defined flag in the graded

pieces of the standard flag. This induces a map $V'_w \rightarrow \tilde{\mathfrak{m}}$ and we can define \tilde{X}_w to be the closure of V'_w inside $q^{-1}(\mathfrak{p}) \times \tilde{\mathfrak{m}}$. Then we obtain a diagram

$$\begin{array}{ccc} & \tilde{X}_w/\mathrm{Res}_{F/\mathbf{Q}_p}P & \\ & \swarrow \quad \searrow & \\ \tilde{\mathfrak{g}}/G & & \tilde{\mathfrak{m}}/\mathrm{Res}_{F/\mathbf{Q}_p}M, \end{array}$$

as well as a canonical projection

$$\pi_P : \tilde{X}_w/\mathrm{Res}_{F/\mathbf{Q}_p}P \longrightarrow \overline{X}_w/G$$

compatible with the projections to $\tilde{\mathfrak{g}}$.

PROPOSITION 5.3.33.

- (i) If $P = B$, then the scheme $\overline{X}_w = \tilde{X}_w$ is normal and Cohen–Macaulay.
- (ii) For a general parabolic subgroup P the scheme \overline{X}_w is unbranched.

These results about the local geometry of \overline{X}_w are proven in [BHS19, Prop. 2.3.3, Thm. 2.3.6] respectively in [Wu21a, Thm. 2.12].

EXPECTATION 5.3.34.

- (i) The schemes \overline{X}_w and \tilde{X}_w are normal and Cohen–Macaulay.
- (ii) The morphism

$$\pi_P : \tilde{X}_w/\mathrm{Res}_{F/\mathbf{Q}_p}P \longrightarrow \overline{X}_w/G$$

satisfies $R\pi_{P,*}\mathcal{O} = \mathcal{O}$ and $R\pi_{P,*}\omega = \omega$.

We can use these objects to define compactifications of the completion $\mathfrak{X}_P^{\mathrm{dR},+}$ by setting

$$(\overline{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda} = (\overline{X}_w/G)_0 \quad \text{and} \quad (\tilde{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda} = (\tilde{X}_w/P)_0.$$

REMARK 5.3.35. Of course, as the notation suggests, we would like to be able to define stacks $\overline{\mathfrak{X}}_P^{\mathrm{dR},+}$ and $\tilde{\mathfrak{X}}_P^{\mathrm{dR},+}$ before completion.

By construction these compactifications fit into a diagram

$$\begin{array}{ccc} & (\tilde{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda} & \\ & \downarrow \quad \searrow & \\ & (\overline{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda} & (\mathfrak{X}_M^{\mathrm{dR},+})_{w\lambda} \\ & \swarrow \quad \searrow & \downarrow \\ (\mathfrak{X}_d^{\mathrm{dR},+})_{\lambda} & & (\mathfrak{X}_{S_M}^{\mathrm{dR},+})_{\det(w\lambda)}, \end{array}$$

where we write $\det(w\lambda)$ for the weight of S_M that is given by the top exterior power on each GL_m -block of the Levi M .

Finally the following result reconnects these stacks, and their explicit local description (respectively: their definition), to the Drinfeld compactifications of \mathfrak{X}_P . Part (ii) of the theorem is a generalization of the local model for the trianguline variety in [BHS19].

THEOREM 5.3.36.

(i) *The canonical map*

$$(\mathfrak{X}_P)_{w\lambda} \longrightarrow (\mathfrak{X}_P^{\mathrm{dR},+})_{w\lambda}$$

induced by $\mathfrak{X}_P \rightarrow \mathfrak{X}_P^{\mathrm{dR},+}$ on completions extends to morphisms

$$\begin{aligned} (\overline{\mathfrak{X}}_P)_{w\lambda} &\longrightarrow (\overline{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda} \\ (\tilde{\mathfrak{X}}_P)_{w\lambda} &\longrightarrow (\tilde{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda}. \end{aligned}$$

(ii) *Let $\tilde{x} \in \tilde{\mathfrak{X}}_P$ and write \bar{x} for its image in $\overline{\mathfrak{X}}_P$. Assume that the image $x_M = (D_1, \dots, D_r)$ of \tilde{x} in \mathfrak{X}_M is regular in the sense that*

- *the map $\mathfrak{X}_M \rightarrow \mathfrak{X}_M^{\mathrm{dR},+}$ is formally smooth at x_M .*
- *for $i \neq j$ the Ext-groups $\mathrm{Ext}^0(D_i, D_j) = 0 = \mathrm{Ext}^2(D_i, D_j) = 0$ vanish.*

Then the morphism

$$(\overline{\mathfrak{X}}_P)_{w\lambda} \longrightarrow (\overline{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda}$$

is formally smooth at \bar{x} .

CONJECTURE 5.3.37. *In the situation of (ii) of the theorem the morphism*

$$(\tilde{\mathfrak{X}}_P)_{w\lambda} \longrightarrow (\tilde{\mathfrak{X}}_P^{\mathrm{dR},+})_{w\lambda}$$

is formally smooth at \tilde{x} .

Of course in combination with the results on the local structure of the varieties \overline{X}_w this gives an explicit description of the local geometry of the Drinfeld compactifications.

In the case $P = B$ the following theorem makes the situation even a bit more precise: Let $\varphi_1, \dots, \varphi_d \in L^\times$ such that $\varphi_i/\varphi_j \neq 1, q$ for $i \neq j$. We write δ_{sm} for the unramified character $(t_1, \dots, t_d) \mapsto \mathrm{unr}_{\varphi_1}(t_1) \cdots \mathrm{unr}_{\varphi_d}(t_d)$. For $w \in W$ let us write

$$\overline{\mathfrak{X}}_{B, z^{w\lambda}\delta_{\mathrm{sm}}} \subseteq \overline{\mathfrak{X}}_B$$

for the preimage of $z^{w\lambda}\delta_{\mathrm{sm}} \in \mathfrak{X}_T$ and let

$$\widehat{\overline{\mathfrak{X}}}_{B, z^{w\lambda}\delta_{\mathrm{sm}}} \subseteq \overline{\mathfrak{X}}_B$$

denote the formal completion of $\overline{\mathfrak{X}}_B$ along this closed substack. We write

$$\mathfrak{X}_{d, (\lambda, \delta_{\mathrm{sm}})\text{-tri}} = \bigcup_{w \in W} \overline{\beta}_B(\overline{\mathfrak{X}}_{B, z^{w\lambda}\delta_{\mathrm{sm}}}) \subseteq \mathfrak{X}_d$$

for the union of the images of $\overline{\mathfrak{X}}_{B, z^{w\lambda}\delta_{\mathrm{sm}}}$ in \mathfrak{X}_d . Note that the assumptions on δ_{sm} imply that the restriction of $\overline{\beta}_B$ to (the formal completion of) $\overline{\mathfrak{X}}_{B, z^{w\lambda}\delta_{\mathrm{sm}}}$ is a closed embedding.

By definition this is the closed substack of \mathfrak{X}_d consisting of all (φ, Γ) -modules D such that D admits a filtration Fil_\bullet with graded pieces

$$\mathcal{R}_F(\mathrm{unr}_{\varphi_i} z^{(\lambda_{w\tau(i)})\tau})$$

(see 5.3.17 for the notation) for some $w \in W$. Over the closed substack $\mathfrak{X}_{d, (\lambda, \delta_{\mathrm{sm}})\text{-tri}}$ this filtration does not glue to a filtration of the universal (φ, Γ) -module \tilde{D} (as the Weyl group element w might vary). However, this turns out to be true after

inverting t , and we write $\mathcal{M}_\bullet \subseteq \tilde{D}[1/t]$ for this filtration. With this notation we can consider the formal substack

$$\widehat{\mathfrak{X}}_{d,(\lambda,\delta_{\text{sm}})-\text{tri}} \subset \widehat{\mathfrak{X}}_d$$

parameterizing deformations \tilde{D}' of D to infinitesimal neighborhoods of $\widehat{\mathfrak{X}}_{d,(\lambda,\delta_{\text{sm}})-\text{tri}}$ such that \mathcal{M}_\bullet lifts to a (by our choice of δ_{sm} : necessarily unique) filtration $\mathcal{M}'_\bullet \subset \tilde{D}'[1/t]$ stable under φ and Γ . The following theorem is basically proven in [BHS19].

THEOREM 5.3.38.

(i) *The morphism*

$$\widehat{\mathfrak{X}}_{B,z^{w\lambda}\delta_{\text{sm}}} \longrightarrow (\widehat{\mathfrak{X}}_B^{\text{dR},+})_{w\lambda} \cong (\overline{X}_{ww_0}/G)_0$$

is formally smooth with relative dualizing sheaf $\mathcal{L}_{2\rho}^{\otimes[F:\mathbf{Q}_p]}$.

(ii) *There is a formally smooth morphism*

$$f_{\lambda,\delta_{\text{sm}}} : \widehat{\mathfrak{X}}_{d,(\lambda,\delta_{\text{sm}})-\text{tri}} \longrightarrow (\tilde{\mathfrak{g}} \times_{\mathfrak{g}} \tilde{\mathfrak{g}})_0$$

with relative dualizing sheaf $\mathcal{L}_{2\rho}^{\otimes[F:\mathbf{Q}_p]}$, whose restriction to $\overline{\beta}_B(\widehat{\mathfrak{X}}_{B,z^{w\lambda}\delta_{\text{sm}}})$ is given by the morphism in (i).

We define the sheaf $\mathcal{L}_{2\rho}$ that occurs in the statement of the theorem: The stack $\widehat{\mathfrak{X}}_{B,z^{w\lambda}\delta_{\text{sm}}}$ has a canonical map to $*/\mathbf{G}_m^d$ (as it sits inside $\widehat{\mathfrak{X}}_B$ which maps to $\widehat{\mathfrak{X}}_T$ and then further projects to $*/\mathbf{G}_m^d$). Hence it makes sense to pull back the line bundle that is associated to the character given by the sum of the positive roots 2ρ . This defines the line bundle in (i). To make sense of this line bundle in (ii) we note that we can write $\mathcal{L}_{2\rho}$ as a pullback of a G -equivariant line bundle on \overline{X}_{ww_0} that is in fact defined on all of $\tilde{\mathfrak{g}} \times_{\mathfrak{g}} \tilde{\mathfrak{g}} \subset \mathfrak{g} \times G/B \times G/B$, as it is given by the pullback of a line bundle (the line bundle corresponding to 2ρ) on one of the G/B -factors.

EXAMPLE 5.3.39. Let us spell out the above result in a bit more detail in the case of $\text{GL}_2(\mathbf{Q}_p)$.

(i) Let $\delta_1 = \text{unr}_{\varphi_1} z^{\lambda_1}$ and $\delta_2 = \text{unr}_{\varphi_2} z^{\lambda_2}$ be two locally algebraic characters with unramified smooth parts, which we assume satisfy $\varphi_1/\varphi_2 \notin \{1, p^{\pm 1}\}$ and $\lambda_1 \neq \lambda_2$. We consider a (φ, Γ) -module that is an extension

$$(5.3.39) \quad 0 \rightarrow \mathcal{R}_L(\delta_1) \rightarrow D \rightarrow \mathcal{R}_L(\delta_2) \rightarrow 0.$$

There are basically two cases to consider: either D is crystalline, or the extension is non-split and $\lambda_1 < \lambda_2$. In the crystalline case we also need to distinguish between $\lambda_1 < \lambda_2$ and $\lambda_1 > \lambda_2$.

- (a) The extension is non-split and $\lambda_1 < \lambda_2$. In this case D is Hodge–Tate but not de Rham and (5.3.39) is the only way to write D as an extension of two rank one (φ, Γ) -modules. In particular (δ_1, δ_2) is the only tuple of characters such that $(D, \delta_1, \delta_2) \in \widehat{\mathfrak{X}}_B$. But if we set $\delta'_1 = \text{unr}_{\varphi_1} z^{\lambda_2}$ and $\delta'_2 = \text{unr}_{\varphi_2} z^{\lambda_1}$, then $(D, \delta'_1, \delta'_2) \in \overline{\mathfrak{X}}_B$. The points

$$x_1 = (D, \delta_1, \delta_2) \text{ and } x_s = (D, \delta_1, \delta_2)$$

are the only points in $\overline{\mathfrak{X}}_B$ above D . Consequently, locally at D the space $\widehat{\mathfrak{X}}_{d,(\lambda,\delta_{\text{sm}})-\text{tri}}$ has two components: one component is the image of a neighborhood (inside $\overline{\mathfrak{X}}_B$) of x_1 , and the other is the image of a neighborhood (inside $\overline{\mathfrak{X}}_B$) of x_s under the projection $\overline{\mathfrak{X}}_B \rightarrow \widehat{\mathfrak{X}}_d$.

(b) The extension D is crystalline. In this case there are again two cases to distinguish:

(b1) We have $\lambda_1 > \lambda_2$. In this case the extension is crystalline no matter whether it is split or not. Moreover, the space $\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}$ is smooth at D .

If the extension D is non-split then there are exactly two points in $\overline{\mathfrak{X}}_B$ above D , and both actually belong to \mathfrak{X}_B . The second point does not arise as in (a) above, but from the fact that D can also be written as an extension

$$0 \rightarrow \mathcal{R}_L(\eta_1) \rightarrow D \rightarrow \mathcal{R}_L(\eta_2) \rightarrow 0,$$

with $\eta_1 = \text{unr}_{\varphi_2} z^{\lambda_1}$ and $\eta_2 = \text{unr}_{\varphi_1} z^{\lambda_2}$. That is: we have changed the *refinement*, i.e. the order of the φ_1 and φ_2 (whereas in (a) we have changed the ordering of the weights λ_1, λ_2), i.e. in Theorem 5.3.38 these points would be addressed by a different choice of the smooth character δ_{sm} .

If the extension is split, then there are three points in $\overline{\mathfrak{X}}_B$ above D . Only one of them corresponds to the refinement given by the ordering φ_1, φ_2 . The other two points correspond to the different ordering and are hence covered by case (b2) below.

(b2) We have $\lambda_1 < \lambda_2$. Then there are three points in $\overline{\mathfrak{X}}_B$, two of which correspond to the ordering (φ_1, φ_2) , i.e. belong to $\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}$. Their description is exactly as in (a). The only difference with (a) is that there is a third point in $\overline{\mathfrak{X}}_B$ mapping to $D \in \mathfrak{X}_2$ corresponding to the refinement given by the ordering (φ_2, φ_1) . Again this point is of course not covered by Theorem 5.3.38.

(ii) Let us describe the local models in the case $d = 2$. The local model $\tilde{\mathfrak{g}} \times_{\mathfrak{g}} \tilde{\mathfrak{g}}$ of Theorem 5.3.38 has two irreducible components both of which are non-singular. To see this, and in order to give explicit equations for the local models, we note that $\tilde{\mathfrak{g}} \times_{\mathfrak{g}} \tilde{\mathfrak{g}}$ is smoothly equivalent to $q^{-1}(\mathfrak{b})$, where $q : \tilde{\mathfrak{g}} \rightarrow \mathfrak{g}$ is the canonical projection. More precisely

$$(\tilde{\mathfrak{g}} \times_{\mathfrak{g}} \tilde{\mathfrak{g}})/G = (q^{-1}(\mathfrak{b}))/B$$

as stacks. Obviously

$$X_1 = \mathfrak{b} \times \{[1 : 0]\} \subset q^{-1}(\mathfrak{b}) \subset \mathfrak{b} \times \mathbf{P}^1$$

is one of the irreducible components. The other component X_s is the closure of

$$V_s = \left\{ \left(\begin{pmatrix} a & b \\ 0 & c \end{pmatrix}, [x : 1] \right) \middle| (a - c)x + b = 0 \right\}$$

where s denotes the non-trivial element of the Weyl group \mathcal{S}_2 of GL_2 . Note that V_s is smooth, open in X_s and disjoint from X_1 . The intersection $X_1 \cap X_s$ is contained in the preimage of the origin $0 = [1 : 0] \in \mathbf{P}^1$ and in the standard affine neighborhood of this point $q^{-1}(\mathfrak{b})$ is described by

$$\left\{ \left(\begin{pmatrix} a & b \\ 0 & c \end{pmatrix}, [1 : y] \right) \middle| y((a - c) + by) = 0 \right\}$$

whereas X_s is given by the equations

$$\left\{ \left(\begin{pmatrix} a & b \\ 0 & c \end{pmatrix}, [1 : y] \right) \middle| (a - c) + by = 0 \right\}.$$

Let us write

$$Z = X_1 \cap X_s = \left\{ \left(\begin{pmatrix} a & b \\ 0 & a \end{pmatrix}, [1 : 0] \right) \right\}$$

for the intersection of the two components. Under the formally smooth morphisms in Theorem 5.3.38, the points of type (a) and (b2) in the first part of the example are mapped to Z , whereas the points of type (b1) are mapped to

$$\left\{ \left(\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}, [x : 1] \right) \right\} \subset V_s$$

More precisely the points of type (a) are mapped to

$$\begin{pmatrix} a & b \\ 0 & a \end{pmatrix}$$

with $b \neq 0$, whereas the points of type (b2) are mapped to the diagonal matrix $\text{diag}(a, a)$.

(iii) The image $\mathfrak{X}_{d,\text{tri}}$ of $\overline{\mathfrak{X}}_B$ in $\mathfrak{X}_d \times \mathfrak{X}_T$ should be seen as a local Galois-theoretic counterpart of an eigenvariety. We refer to [BHS17b] for a mathematically meaningful formulation of this slogan. In the case of the eigencurve we briefly indicate how the points discussed in (i) arise. Let f be a classical modular eigenform of weight $k \geq 2$ level $\Gamma_1(N)$ for some N not divisible by p . Then the Galois representation ρ attached to f is crystalline at p . At level $\Gamma_1(N) \cap \Gamma_0(p)$ there are two *stabilizations* f_α and f_β of f which are eigenforms for the same Hecke eigenvalues away from p but with different U_p -eigenvalues, say α and β . These two different U_p -eigenvalues correspond to the two possible choices of a refinement above. There are at least two branches of the eigencurve that meet at ρ inside the generic fiber of a deformation space of Galois representations: the branch containing the classical modular form f_α and the branch containing the classical form f_β .

In the case that ρ is non-split at p these are the only branches containing the Galois representation ρ . This corresponds to the non-split extensions in case (b1) above (note that both f_α and f_β are of dominant algebraic weight). But if ρ is split crystalline at p there is third branch of the eigencurve containing a non-classical p -adic modular form f' (of weight $1 - k$). This is precisely the case (b2) discussed above. Let us mention explicitly that the two branches of $\overline{\mathfrak{X}}_{d,(\Delta, \delta_{\text{sm}})}\text{-tri}$ in Theorem 5.3.38 meeting at a given (φ, Γ) -module D as in (b2) are not the branches containing the classical forms f_α and f_β , but one of these branches should be thought of as containing a classical modular form, whereas the other branch contains the non-classical form f' (the points on that branch being of non-classical weight). Finally we note that the points in (a) correspond to non-classical p -adic modular forms (that may or may not have classical weight).

REMARK 5.3.40. In [BD21] Breuil and Ding define a variant of the trianguline variety for paraboline (φ, Γ) -modules. In the stacky context discussed above their construction can be described as follows (the notations differ from the ones of Breuil–Ding): Instead of \mathfrak{X}_M we consider its closed substack of \mathfrak{X}'_M of objects that are de Rham of prescribed weight and inertial type, up to a twist with an arbitrary rank 1 object. Then we can define the preimage \mathfrak{X}'_P of \mathfrak{X}'_M in \mathfrak{X}_P and its closure $\overline{\mathfrak{X}'_P}$ in the compactification $\overline{\mathfrak{X}}_P$. Then, in our language, the analogue of the paraboline variety of Breuil–Ding is the image of $\overline{\mathfrak{X}'_P}$ inside $\mathfrak{X}_d \times \mathfrak{X}_{S_M}$. Breuil and Ding also

construct local models for their spaces. These local models are given by the Zariski-closure (inside the space $X_{B,P}$ of (5.3.30)) of the closed subspace in V_w of elements whose image in \mathfrak{m} is a central element of the Lie algebra.

6. Categorical p -adic local Langlands conjectures

In this section we state our main conjectures, and explain some motivation for them (in particular, we explain some motivations coming from the Taylor–Wiles method).

6.1. Expectations: The Banach Case. We now make precise our expectations in the Banach case. In order to do so, we introduce some notation. Fix as usual a finite extension F/\mathbf{Q}_p and a coefficient ring \mathcal{O} . Set $G = \mathrm{GL}_d(F)$, $K = \mathrm{GL}_d(\mathcal{O}_F)$, $Z = Z(G)$.

Our goal, as we’ve already explained, is to make a conjecture relating (the derived category of) smooth representations to coherent sheaves, or, more generally, Ind-coherent complexes, on the formal algebraic stack \mathcal{X}_d . In order to make precise statements, we require some preliminary discussion of the relevant categories of G -representations.

6.1.1. *Preliminaries on categories of representations.* We begin by recalling some material from Appendix E. We let $\mathrm{sm.} G$ denote the abelian category of smooth representations of G on \mathcal{O} -modules, where we define an \mathcal{O} -linear G -representation to be smooth if any element of the representation is fixed by an open subgroup of G , and annihilated by some power of p . (See Definition E.1.2.) We let $D(\mathrm{sm.} G)$ denote the derived stable ∞ -category of $\mathrm{sm.} G$.

Any smooth G -representation admits a canonical structure of $\mathcal{O}[[G]]$ -module, where $\mathcal{O}[[G]]$ is the ring defined in Definition E.1.1, and it’s frequently useful to study smooth G -representations from the perspective of their $\mathcal{O}[[G]]$ -module structures. To this end, Proposition E.2.2 yields a fully faithful t -exact functor

$$(6.1.2) \quad D(\mathrm{sm.} G) \hookrightarrow D(\mathcal{O}[[G]])$$

(the target being the stable ∞ -category of complexes of $\mathcal{O}[[G]]$ -modules) with essential image equal to $D_{\mathrm{sm}}(\mathcal{O}[[G]])$ (the full subcategory of objects with smooth cohomologies).

Since there are many $\mathcal{O}[[G]]$ -modules that are not smooth (e.g. $\mathcal{O}[[G]]$ itself), the functor (6.1.2) is far from being an equivalence. On the other hand, via a consideration of pro-objects, we can construct a functor which is close to being left-inverse to (6.1.2). Indeed, if we let $D_{\mathrm{c.g.}}^-(\mathcal{O}[[G]])$ denote the full subcategory of $D(\mathcal{O}[[G]])$ consisting of complexes whose cohomologies are countably generated and vanish in sufficiently high degree, then (E.4.4) provides a functor

$$(6.1.3) \quad D_{\mathrm{c.g.}}^-(\mathcal{O}[[G]]) \rightarrow \mathrm{Pro} D^b(\mathrm{sm.} G).$$

(For example, $\mathcal{O}[[G]]$ itself maps to the formal pro object “ $\varprojlim_{H,n} c\text{-Ind}_H^G \mathcal{O}/p^n$ ” of $\mathrm{Pro} \mathrm{sm.} G$; here H runs over all compact open subgroups of G and n over all positive integers, and \mathcal{O}/p^n is given the trivial H -action.) Furthermore, the composite

$$D_{\mathrm{c.g.}}^b(\mathrm{sm.} G) \xrightarrow{(E.4.5)} D_{\mathrm{c.g.sm}}^b(\mathcal{O}[[G]]) \hookrightarrow D_{\mathrm{c.g.}}^-(\mathcal{O}[[G]]) \xrightarrow{(6.1.3)} \mathrm{Pro} D^b(\mathrm{sm.} G)$$

is simply the canonical fully faithful functor $D_{\mathrm{c.g.}}^b(\mathrm{sm.} G) \hookrightarrow \mathrm{Pro} D^b(\mathrm{sm.} G)$. (This provides the sense in which we regard (6.1.3) as close to being left-inverse to (6.1.2).)

In order to properly state our categorical Langlands conjecture, we make a preliminary conjecture regarding the abelian category $\mathrm{sm}. G$.

CONJECTURE 6.1.4. *The abelian category $\mathrm{sm}. G$ of smooth \mathcal{O} -linear G -representations is locally coherent.*

REMARK 6.1.5. Recall that this means that $\mathrm{sm}. G$ is generated by its compact objects, and that these compact objects form an abelian subcategory of $\mathrm{sm}. G$. We have seen in Proposition E.3.5 that the compact objects in $\mathrm{sm}. G$ are precisely those representations that are of finite presentation, in the sense of Definition E.3.3, and that $\mathrm{sm}. G$ is generated by these objects. Thus the content of Conjecture 6.1.4 is that the full subcategory of $\mathrm{sm}. G$ consisting of representations of finite presentation is closed under the formation of kernels in $\mathrm{sm}. G$. (It is obviously closed under the formation of cokernels.)

REMARK 6.1.6. Since the smooth representations of finite presentation are precisely those that are finitely presented as $\mathcal{O}[[G]]$ -modules (Lemma E.3.4), Conjecture 6.1.4 certainly holds if $\mathcal{O}[[G]] - \mathrm{Mod}$ is locally coherent, i.e. if $\mathcal{O}[[G]]$ is a coherent ring, i.e. if the category of finitely presented $\mathcal{O}[[G]]$ -modules is stable under the formation of kernels.

In the case $d = 1$ the ring $\mathcal{O}[[G]]$ is even Noetherian, while if $d = 2$ the ring $\mathcal{O}[[\mathrm{SL}_2(F)]]$ is coherent by a result of Shotton [Sho20, Cor. 4.4], and we expect that a variation of his argument (considering the amalgamated product of two copies of KZ over IZ) will show that $\mathcal{O}[[\mathrm{GL}_2(F)]]$ is itself coherent. In either case it thus follows that the category of finitely presented smooth $\mathcal{O}[[G]]$ -modules is abelian (at least after making an additional local Z -finiteness hypothesis if $d = 2$, to deal with the $\mathrm{SL}_2 / \mathrm{GL}_2$ distinction, which could be removed if $\mathcal{O}[[\mathrm{GL}_2(F)]]$ is indeed coherent).

No other cases of Conjecture 6.1.4 are known. It is known that if $d > 2$, then $\mathcal{O}[[G]]$ is not coherent, by a result of Timmins [Tim21]. Our main reason for believing Conjecture 6.1.4 is that we believe that something like Conjecture 6.1.14 must hold, and it is difficult to formulate such a conjecture consistently unless Conjecture 6.1.4 holds.

Suppose now that Conjecture 6.1.4 holds, and let $D_{\mathrm{f.p.}}^b(\mathrm{sm}. G)$ denote the full subcategory of $D(\mathrm{sm}. G)$ consisting of (cohomologically) bounded complexes whose cohomologies are finitely presented $\mathcal{O}[[G]]$ -modules. As noted in the discussion following Proposition E.3.5, $D_{\mathrm{f.p.}}^b(\mathrm{sm}. G)$ is then precisely the full subcategory of $D(\mathrm{sm}. G)$ consisting of coherent objects.

We have a duality theory for $D_{\mathrm{f.p.}}^b(\mathrm{sm}. G)$ defined as follows. We identify $\mathcal{O}[[G]]^{\mathrm{op}}$ with $\mathcal{O}[[G]]$ via $g \mapsto g^{-1}$. In this way, we may canonically convert right $\mathcal{O}[[G]]$ -modules into left G -modules. Then we define an antiequivalence

$$\mathbf{D} : D_{\mathrm{f.p.}}^b(\mathrm{sm}. G) \rightarrow D_{\mathrm{f.p.}}^b(\mathrm{sm}. G)$$

via

$$\mathbf{D}(-) := \mathrm{RHom}_{\mathcal{O}[[G]]}(-, \mathcal{O}[[G]])[d^2[F : \mathbf{Q}_p] + 1].$$

We anticipate that this is of amplitude $[0, d]$. It is equipped with a natural isomorphism $\mathbf{D} \circ \mathbf{D} \xrightarrow{\sim} \mathrm{Id}$, and has the property that

$$(6.1.7) \quad \mathbf{D}(c\text{-Ind}_H^G V) = c\text{-Ind}_H^G V^\vee$$

for any finite length smooth representation V of a compact open subgroup H of G ; here $V^\vee = \mathrm{Hom}_{\mathcal{O}}(V, E/\mathcal{O})$ denotes the Pontryagin dual of V , equipped with its contragredient H -action.

6.1.8. *Serre duality on \mathcal{X}_d .* On the Galois side, we have the derived categories of (Ind-)coherent sheaves $D_{\mathrm{coh}}(\mathcal{X}_d)$ and $\mathrm{IndCoh}(\mathcal{X}_d)$ (as defined in Appendix B.2). There is an involution ι of \mathcal{X}_d given in moduli-theoretic terms by $\rho \mapsto \rho^\vee$. We let $\mathcal{D}_{\mathcal{X}_d}$ denote the antiequivalence of $D_{\mathrm{coh}}(\mathcal{X}_d)$ given by composing Grothendieck–Serre duality with ι^* (i.e. $\iota^* \circ R\mathcal{H}om_{\mathcal{O}_{\mathcal{X}_d}}(-, \omega_{\mathcal{X}_d})$ where $\omega_{\mathcal{X}_d}$ is the dualizing sheaf of \mathcal{X}_d , thought of as pro-coherent sheaf on \mathcal{X} (as discussed in Example B.3.2) and placed in degree 0; recall that we anticipate that \mathcal{X}_d is lci).

6.1.9. *Hodge types, inertial types, components of the Bernstein centre, and Hecke algebras.* By results of Schneider–Zink [SZ99] (the “inertial local Langlands correspondence”), we may associate to any inertial type $\tau : I_F \rightarrow \mathrm{GL}_d(L)$ a finite-dimensional smooth irreducible K -representation $\sigma^{\mathrm{crys}}(\tau)$ as in [EG22, Thm. 8.2.1]. We let $\sigma^{\mathrm{crys}, \circ}(\tau)$ denote an arbitrary choice of K -stable \mathcal{O} -lattice in $\sigma^{\mathrm{crys}}(\tau)$. As is explained in [CEG⁺16], the type τ determines a Bernstein component Ω_τ , and there is a canonical isomorphism

$$(6.1.10) \quad \mathcal{Z}_{\Omega_\tau} \xrightarrow{\sim} \mathrm{End}_G(c\text{-Ind}_K^G \sigma^{\mathrm{crys}}(\tau)).$$

If $\underline{\lambda}$ is a regular Hodge type, then we have the algebraic representation $W_{\underline{\lambda}}$ of K (see Definition 2.3.2). Since $W_{\underline{\lambda}}[1/p]$ is naturally a G -representation, “push-pull” for compact induction shows that

$$\begin{aligned} (c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}}(\tau)^\circ)[1/p] &= c\text{-Ind}_K^G W_{\underline{\lambda}}[1/p] \otimes_E \sigma^{\mathrm{crys}}(\tau) \\ &= W_{\underline{\lambda}}[1/p] \otimes_E c\text{-Ind}_K^G \sigma^{\mathrm{crys}}(\tau). \end{aligned}$$

As $W_{\underline{\lambda}}[1/p]$ is furthermore an irreducible algebraic representation of G , while the compact induction is a smooth G -representation, we find that

$$(6.1.11) \quad \begin{aligned} \mathrm{End}_G\left((c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}}(\tau)^\circ)[1/p]\right) &= \mathrm{End}_G(W_{\underline{\lambda}}[1/p] \otimes_E c\text{-Ind}_K^G \sigma^{\mathrm{crys}}(\tau)) \\ &= \mathrm{End}_G(c\text{-Ind}_K^G \sigma^{\mathrm{crys}}(\tau)) \stackrel{(6.1.10)}{=} \mathcal{Z}_{\Omega_\tau}. \end{aligned}$$

As recalled in Section 5.2.9, the Bernstein centre $\mathcal{Z}_{\Omega_\tau}$ can also be identified with the global functions on a stack of Weil–Deligne representations, and as in (5.2.14), there is a morphism

$$(6.1.12) \quad \mathcal{Z}_{\Omega_\tau} \longrightarrow \Gamma((\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau})_\eta^{\mathrm{rig}}, \mathcal{O}_{(\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau})_\eta^{\mathrm{rig}}}).$$

6.1.13. *Statement of the conjecture.* We are now ready to state our categorical Langlands conjecture.

CONJECTURE 6.1.14. *Assume Conjecture 6.1.4. Then there exists an \mathcal{O} -linear exact fully faithful functor $\mathfrak{A} : D_{\mathrm{f.p.}}^b(\mathrm{sm.} G) \rightarrow D_{\mathrm{coh}}^b(\mathcal{X}_d)$ satisfying the following properties:*

- (1) *The functor \mathfrak{A} has bounded cohomological amplitude.*
- (2) *(L_∞ is a sheaf.) $L_\infty := \mathfrak{A}(\mathcal{O}[[G]])$ is a pro-coherent sheaf, concentrated in degree 0. Furthermore it is flat over $\mathcal{O}[[K]]$.*
- (3) *(Compatibility with duality.) There is a natural equivalence*

$$(6.1.15) \quad \mathfrak{A} \circ \mathbf{D} \xrightarrow{\sim} (\mathbf{D}_{\mathcal{X}_d} \circ \mathfrak{A})[[F : \mathbf{Q}_p]d(d+1)/2 + 1]$$

- of contravariant functors from $D_{\text{f.p.}}^b(\text{sm. } G)$ to $D_{\text{coh}}(\mathcal{X}_d)$.
- (4) (Locally algebraic vectors.) For any regular Hodge type λ and inertial type τ , the scheme-theoretic support of $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys}, \circ}(\tau)\right)^\wedge\right)$ is equal to $\mathcal{X}_d^{\text{crys}, \lambda, \tau}$. (The $\widehat{}$ here denotes p -adic completion.) Furthermore, the action of the Bernstein centre $\mathcal{Z}_{\Omega_\tau}$ on $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys}, \circ}(\tau)\right)^\wedge\right)[1/p]$ induced by the identification (6.1.11) and the functoriality of \mathfrak{A} coincides with the action given by (6.1.12).

REMARK 6.1.16. The notation \mathfrak{A} for the functor of Conjecture 6.1.14 is drawn from the paper [Zhu20], where the same notation is used in the $\ell \neq p$ case.

REMARK 6.1.17. Assuming that it exists, the functor \mathfrak{A} has natural Ind and Pro extensions whose considerations are important (and indeed necessary to make sense of the conjecture as stated). We begin with the Ind-extension of \mathfrak{A} ; this is a cocontinuous functor

$$\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G) \rightarrow \text{Ind Coh}(\mathcal{X}),$$

which we again denote by \mathfrak{A} .

The category $\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$ is the analog, in our context, of the renormalized categories of smooth representations considered in [Zhu20, §4.1]. It admits a canonical functor

$$(6.1.18) \quad \text{Ind } D_{\text{f.p.}}^b(\text{sm. } G) \rightarrow D(\text{sm. } G),$$

which becomes an equivalence if we restrict to bounded below complexes on each side. Indeed, assuming Conjecture 6.1.4, the t -structure on $D(\text{sm. } G)$ is *coherent* in the sense of Definition A.6.9, and $D_{\text{f.p.}}^b(\text{sm. } G)$ is then precisely the subcategory of coherent objects in $D(\text{sm. } G)$, so that the functor is particular instance of (A.6.10). In particular, the heart of $\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$ coincides with the heart of $D(\text{sm. } G)$, i.e. with the abelian category $\text{sm. } G$. Thus, extending \mathfrak{A} to all of $\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$ allows us in particular to evaluate \mathfrak{A} on all smooth G -representations, not just finitely presented ones.

As an aside, we note that in general the t -structure on $D(\text{sm. } G)$ need not be regular (in the sense of Definition A.6.11), and so (6.1.18) need not be an equivalence in general. Indeed, we expect that the t -structure on $D(\text{sm. } G)$ is regular, so that (6.1.18) is an equivalence, precisely when I_1 (the pro- p -Iwahori subgroup of K) is torsion-free.

Since $D^b(\text{sm. } G)$ is a full subcategory of $\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$, we may restrict (the Ind-extension of) \mathfrak{A} to $D^b(\text{sm. } G)$ so as to obtain a functor $D^b(\text{sm. } G) \rightarrow \text{Ind Coh}(\mathcal{X})$. (In fact, taking into account Remark 6.1.27 below, we see that this functor should take values in $D_{\text{coh}}^b(\mathcal{X})$.) We may Pro-extend this functor to obtain a functor

$$\text{Pro } D^b(\text{sm. } G) \rightarrow \text{Pro Ind Coh}(\mathcal{X}).$$

Composing this with the functor (6.1.3), we find that it makes sense to evaluate \mathfrak{A} on objects of $D_{\text{c.g.}}^-(\mathcal{O}[[G]])$. The discussion of Section 6.1.1 shows that the restriction of this functor to $D_{\text{f.p.}}^b(\text{sm. } G) \xrightarrow{\sim} D_{\text{f.p.sm}}^b(\mathcal{O}[[G]])$ (here the target has the obvious meaning, namely it is the full subcategory of $D_{\text{sm}}^b(\mathcal{O}[[G]])$ consisting of complexes whose cohomologies are finitely presented; the indicated equivalence

follows directly from Proposition E.2.2 by restricting the equivalence stated there to the full subcategories of its source and target consisting of complexes whose cohomologies are finitely generated) coincides with the original functor \mathfrak{A} .

REMARK 6.1.19. To make sense of Conjecture 6.1.14 (2), note that, as discussed in Remark 6.1.17, we may evaluate \mathfrak{A} on $\mathcal{O}[[G]]$ (thought of as a module over itself). *A priori* this gives rise to an object of $\text{Pro Ind}(\mathcal{X}_d)$, but part (2) asserts that in fact we obtain an object of $\text{Pro Coh}(\mathcal{X}_d)$.

REMARK 6.1.20. As was already mentioned in Section 1.4, it seems likely that Conjecture 6.1.14 can be viewed as a consequence of a stronger conjecture, to the effect that there should be an equivalence of (derived) categories as in the conjecture of Fargues–Scholze [FS21, Conj. I.10.2], but with p -adic rather than ℓ -adic coefficients. However it does not seem to be clear at the time of writing precisely which category of sheaves on Bun_G should be considered in order to obtain such an equivalence.

REMARK 6.1.21. The full faithfulness of \mathfrak{A} implies in particular that there should be identifications of E_1 -rings between p -adic derived Hecke algebras and certain endomorphism rings of coherent sheaves on \mathcal{X}_d (e.g. p -adic analogues of Feng’s spectral Hecke algebra [Fen19]); see [Zhu20, Conj. 4.3.1] for an analogous conjecture in the case $\ell \neq p$.

REMARK 6.1.22. We could equally well formulate Conjecture 6.1.14 (4) for potentially semistable representations, but for simplicity of exposition we have restricted to the potentially crystalline case.

REMARK 6.1.23. The functor of Conjecture 6.1.14 is necessarily of the form considered in Appendix C.2. That is, it is of the form

$$\pi \mapsto L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi,$$

where, as in Conjecture 6.1.14 (2), $L_\infty := \mathfrak{A}(\mathcal{O}[[G]])$ is a pro-coherent sheaf on the moduli stack \mathcal{X}_d , and is in addition a right $\mathcal{O}[[G]]$ -module (because the endomorphism algebra of $\mathcal{O}[[G]]$ is $\mathcal{O}[[G]]$). (In the case $d = 1$, we see below that L_∞ is the structure sheaf on \mathcal{X}_1 , by (7.1.2), and in the case of $\text{GL}_2(\mathbf{Q}_p)$ we give a construction of L_∞ , but in general we do not have a conjectural candidate for L_∞ .)

Indeed, assuming Conjecture 6.1.4, we may resolve any finitely presented smooth $\mathcal{O}[[G]]$ -module π by a (possibly infinite) complex of direct sums of copies of $\mathcal{O}[[G]]$

$$\cdots \rightarrow \mathcal{O}[[G]]^{\oplus m_n} \cdots \rightarrow \mathcal{O}[[G]]^{\oplus m_0} \rightarrow \pi \rightarrow 0,$$

and we find that $\mathfrak{A}(\pi)$ is computed by

$$\cdots \rightarrow \mathfrak{A}(\mathcal{O}[[G]]^{\oplus m_n}) \cdots \rightarrow \mathfrak{A}(\mathcal{O}[[G]]^{\oplus m_0})$$

and thus by

$$L_\infty \otimes_{\mathcal{O}[[G]]} (\cdots \rightarrow \mathcal{O}[[G]]^{\oplus m_n} \cdots \rightarrow \mathcal{O}[[G]]^{\oplus m_0}) = L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi.$$

(Assumption (1), i.e. the boundedness of the amplitude of \mathfrak{A} , justifies the manipulations with these possibly unbounded below complexes.) Since $\text{Ind } D_{f,p}^b(\text{sm. } G)$ is generated under shifts and colimits by the representations of finite presentation, we have

$$(6.1.24) \quad \mathfrak{A}(\pi) = L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi$$

in general.

REMARK 6.1.25. We note that the only reason for making an *a priori* assumption of bounded amplitude for \mathfrak{A} (i.e. the only reason for imposing condition (1) in the conjecture) is to justify the manipulations that lead to formula (6.1.24), and we should emphasize that the existence of such a formula for \mathfrak{A} is basic to our own point of view on Conjecture 6.1.14. If we grant the validity of this formula, then the boundedness of the amplitude of \mathfrak{A} should follow, as we explain in Remark 6.1.27 below.

REMARK 6.1.26. If V is a smooth $\mathcal{O}[[K]]$ -module, of finite type as an \mathcal{O} -module (so that it is p -adically complete), then we may regard the p -adically completed compactly induced representation $(c\text{-Ind}_K^G V)^\wedge$ as an object of $\text{Pro sm. } G$, and hence form $\mathfrak{A}((c\text{-Ind}_K^G V)^\wedge)$. (And this gives the sense in which we form

$$\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys}, \circ}(\tau)\right)^\wedge\right)$$

in property (4) of Conjecture 6.1.14.)

It then follows from Conjecture 6.1.14 (2) that $\mathfrak{A}((c\text{-Ind}_K^G V)^\wedge)$ is concentrated in degree 0. Indeed, by Remark 6.1.23, we have

$$\mathfrak{A}(c\text{-Ind}_K^G V) = L_\infty \widehat{\otimes}_{\mathcal{O}[[G]]}^L (c\text{-Ind}_K^G V)^\wedge = L_\infty \otimes_{\mathcal{O}[[K]]}^L V$$

(here $\widehat{\otimes}$ indicates the p -adically completed tensor product) and the claim follows from the hypothesized flatness of L_∞ over $\mathcal{O}[[K]]$.

In particular, the sheaves $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys}, \circ}(\tau)\right)^\wedge\right)$ considered in property (4) are concentrated in degree 0.

REMARK 6.1.27. We can be much more precise about the bounded amplitude conjectured in statement (1) of the conjecture. Namely, we anticipate that the functor \mathfrak{A} has amplitude $[1 - d, 0]$. Indeed, in the case $d = 2$ and $p > 2$ this follows from Remark 6.1.26 and the resolution explained in Remark 7.3.11 (at least if we ignore the action of the centre, which we do not expect to contribute to the amplitude), and we anticipate that analogous resolutions exist for all d .

REMARK 6.1.28. Since (by construction) the compact objects in $\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$ are the objects of $D_{\text{f.p.}}^b(\text{sm. } G)$, and the compact objects in $\text{Ind Coh}(\mathcal{X}_d)$ are given by $D_{\text{coh}}^b(\mathcal{X}_d)$, it follows from Lemma A.3.2 (3) that if π is an object of $\text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$, then π actually belongs to $D_{\text{f.p.}}^b(\text{sm. } G)$ if and only if $\mathfrak{A}(\pi)$ is an object of $D_{\text{coh}}^b(\mathcal{X}_d)$.

In particular, if π is concentrated in degree zero, then we see that π is of finite presentation if and only if $\mathfrak{A}(\pi)$ is a bounded complex with coherent cohomology sheaves. (In fact, we expect that $\mathfrak{A}(\pi)$ is automatically bounded, by Remark 6.1.27.)

If π is concentrated in degree zero, and is finitely generated but not of finite presentation, then $\mathfrak{A}(\pi)$ is not coherent. However, we claim that $H^0(\mathfrak{A}(\pi))$ is coherent. Indeed, since π is finitely generated, it admits a surjection $c\text{-Ind}_K^G U \rightarrow \pi$ for some finite length K -representation U . Since \mathfrak{A} is right t -exact (see Remark 6.1.27), we have a surjection $H^0(\mathfrak{A}(c\text{-Ind}_K^G U)) \rightarrow H^0(\mathfrak{A}(\pi))$; and $\mathfrak{A}(c\text{-Ind}_K^G U)$ is coherent (because $c\text{-Ind}_K^G U$ is of finite presentation).

REMARK 6.1.29. We have tried to formulate Conjecture 6.1.14 with a minimal set of useful properties; we explain some motivation for these properties in the following remarks. It is unclear to what extent these properties uniquely determine \mathfrak{A} , but in practice they seem to seriously constrain it.

REMARK 6.1.30. We can think of L_∞ as a “universal”, purely local version of completed cohomology. More precisely, under the assumption that $p \nmid 2d$, for any point $x \in \mathcal{X}_d(\overline{\mathbf{F}}_p)$ the paper [CEG⁺16] constructs a versal morphism $f : \mathrm{Spf} R_\infty \rightarrow \mathcal{X}_d$ at x , together with an R_∞ -module M_∞ with a commuting action of G . Here R_∞ is a power series ring over a universal lifting ring for the Gal_F -representation corresponding to x , with the power series variables being the “patching variables” occurring in the Taylor–Wiles–Kisin method, and M_∞ is obtained by applying the patching method to the completed cohomology of certain unitary groups. (See Section 3 for a very similar construction for modular curves.)

The construction of R_∞ and M_∞ is global and depends on many choices, but we expect that $M_\infty = f^*L_\infty$ (see also Expected Theorem 9.4.2). This would confirm the conjecture in [CEG⁺16, §6] that M_∞ should be purely local; indeed the hypothetical L_∞ in loc. cit. is just the pullback of our L_∞ to a local deformation ring. It also explains some of our expectations for L_∞ ; for example, M_∞ is flat over $\mathcal{O}[[K]]$ by construction, whence our expectation that L_∞ is flat over $\mathcal{O}[[K]]$.

REMARK 6.1.31. Property (2) of Conjecture 6.1.14 is an analogue of the expectation [BZCHN20, Conj. 4.15] that the coherent Springer sheaf is concentrated in degree 0. In our case it is motivated by the expectation of Remark 6.1.30 (which implies that the pullback of L_∞ to any versal ring is concentrated in degree 0).

REMARK 6.1.32. By property (4), for any regular Hodge type $\underline{\lambda}$ and inertial type τ , $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau)\right)^\wedge\right)$ is supported on $\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$. We have the closed immersion $i : \mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau} \hookrightarrow \mathcal{X}_d$, which we expect to be pure of codimension $[F : \mathbf{Q}_p]d(d+1)/2$ (we say “expect” because we are not aware of a dimension theory for formal algebraic stacks in the literature; but note that by [Kis10b, BIP21] we know the corresponding claim on versal rings, i.e. on Galois deformation rings). Granting this expectation, we have $i^!\omega_{\mathcal{X}_d} = \omega_{\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}}[-[F : \mathbf{Q}_p]d(d+1)/2]$, where we abusively write $\omega_{\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}}$ for the dualizing complex shifted by the dimension of $\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$ (so that if $\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$ is Cohen–Macaulay then $\omega_{\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}}$ is a sheaf in degree zero). Thus for any $\mathcal{F} \in D_{\mathrm{Coh}}(\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau})$ we have

$$i_* R\mathrm{Hom}_{\mathcal{O}_{\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}}}(\mathcal{F}, \omega_{\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}}) = R\mathrm{Hom}_{\mathcal{O}_{\mathcal{X}_d}}(i_* \mathcal{F}, \omega_{\mathcal{X}_d})[[F : \mathbf{Q}_p]d(d+1)/2].$$

Write

$$D_{\mathcal{O}_{\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}}} : D_{\mathrm{Coh}}(\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}) \rightarrow D_{\mathrm{Coh}}(\mathcal{X}_d^{\mathrm{crys}, -\underline{\lambda}, \tau^\vee})$$

for the composite of Grothendieck–Serre duality and i^* , where

$$i : \mathcal{X}_d^{\mathrm{crys}, -\underline{\lambda}, \tau^\vee} \xrightarrow{\sim} \mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$$

is induced by the involution ι of \mathcal{X}_d .

Take $\mathcal{F} = \mathfrak{A}\left(\left(c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau)\right)^\wedge\right)$. Note that it follows from (6.1.7) that

$$D\left(\left(c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau)\right)^\wedge\right) = \left(\left(c\text{-Ind}_K^G W_{-\underline{\lambda}}^* \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau^\vee)\right)^\wedge\right)[1],$$

where we abusively write $W_{-\underline{\lambda}}^*$ for the (not dual!) Weyl module, and choose $\sigma^{\mathrm{crys}, \circ}(\tau^\vee)$ to be a lattice in $\sigma^{\mathrm{crys}}(\tau^\vee)$ that is dual to our chosen lattice $\sigma^{\mathrm{crys}, \circ}(\tau)$.

Bearing in mind (6.1.15), we obtain

$$(6.1.33) \quad \mathfrak{A}\left(\left(c\text{-Ind}_K^G W_{-\lambda}^* \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau^\vee)\right)^\wedge\right) \\ \xrightarrow{\sim} \mathbf{D}_{\mathcal{X}_d^{\text{crys},\Delta,\tau}} \left(\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)\right).$$

REMARK 6.1.34. Property (3) of Conjecture 6.1.14 can be thought of explaining the fact that the patched modules (constructed by Taylor–Wiles–Kisin patching at finite level) are maximal Cohen–Macaulay over their supports. In particular we expect that we can strengthen property (4) to say that for any regular Hodge type λ and inertial type τ , $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)$ is concentrated in degree 0, is maximal Cohen–Macaulay over $\mathcal{X}_d^{\text{crys},\Delta,\tau}$, and $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)_\eta^{\text{rig}}$ is locally free of rank one over $(\mathcal{X}_d^{\text{crys},\Delta,\tau})_\eta^{\text{rig}}$. We now sketch an explanation for this.

Since we are assuming that $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)$ is concentrated in degree 0 and is supported on all of $\mathcal{X}_d^{\text{crys},\Delta,\tau}$, and similarly for $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_{-\lambda}^* \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau^\vee)\right)^\wedge\right)$,¹¹ it follows from (6.1.33) that $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)$ is maximal Cohen–Macaulay over $\mathcal{X}_d^{\text{crys},\Delta,\tau}$ (see [Sta, Tag 0B5A]).

It follows that $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)_\eta^{\text{rig}}$ is maximal Cohen–Macaulay over $(\mathcal{X}_d^{\text{crys},\Delta,\tau})_\eta^{\text{rig}}$. Since $(\mathcal{X}_d^{\text{crys},\Delta,\tau})_\eta^{\text{rig}}$ is smooth, this implies that $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau)\right)^\wedge\right)_\eta^{\text{rig}}$ is locally free, and we anticipate that the full faithfulness of \mathfrak{A} , together with [CEG⁺16, Cor. 3.12], implies that it is furthermore locally free of rank one.

The Cohen–Macaulay property of patched modules is usually established by completely different means, using the action of the ring S_∞ of diamond operators; see [Dia97].

REMARK 6.1.35. As already noted when defining our modified version $\mathbf{D}_{\mathcal{X}_d}$ of Grothendieck–Serre duality on \mathcal{X}_d , we have the “main involution” ι on \mathcal{X}_d defined as $\rho \mapsto \rho^\vee$. This induces an auto-involution i^* on $D_{\text{coh}}^b(\mathcal{X}_d)$. We also have the usual main involution $g \mapsto (g^{-1})^t$ of $\text{GL}_d(F)$, and precomposing representations of $\text{GL}_d(F)$ with this involution induces an auto-involution i' of $D_{\text{f.p.}}^b(\text{sm. } G)$. We expect that \mathfrak{A} intertwines the involutions i' and i .

For example, in the statement of Property (3) of Conjecture 6.1.14, rather than replacing Grothendieck–Serre duality \mathbf{D} on \mathcal{X}_d by its twist $\mathbf{D}_{\mathcal{X}_d} := i^* \mathbf{D}$, we could replace the duality \mathbf{D} on $D_{\text{f.p.}}^b(\text{sm. } G)$ by its twist $\mathbf{D}' := i' \circ \mathbf{D}$. Then Property (3) could be rephrased as saying that \mathfrak{A} interchanges \mathbf{D}' and \mathbf{D} (up to a shift). Now (6.1.7) implies (ignoring possible shifts) that $\mathbf{D}'(c\text{-Ind}_K^G V) = c\text{-Ind}_K^G V$ (respectively that $\mathbf{D}'((c\text{-Ind}_K^G V^\circ)^\wedge) = (c\text{-Ind}_K^G V^\circ)^\wedge$) if V is an irreducible mod p representation of K (respectively if V° is a K -invariant lattice in an irreducible locally algebraic representation V of K). Thus we expect that the corresponding

¹¹While we have not made an explicit conjecture involving $W_{-\lambda}^*$, this statement follows from the version for $W_{-\lambda}$, since these representations are isomorphic after inverting p .

coherent sheaves $\mathfrak{A}(\mathrm{c}\text{-Ind}_K^G V)$ (respectively $\mathfrak{A}((\mathrm{c}\text{-Ind}_K^G V^\circ)^\wedge)$) are Grothendieck–Serre self-dual (up to a shift).

REMARK 6.1.36. We also expect a compatibility with parabolic induction, just as in the $\ell \neq p$ and analytic cases (see Section 6.2.13 for the latter).

6.1.37. *The geometric Breuil–Mézard conjecture.* The geometric Breuil–Mézard conjecture for the stack \mathcal{X}_d was formulated in [EG22]; it was motivated by (and indeed is equivalent to) the earlier geometric Breuil–Mézard conjectures for Galois deformation rings of [BM14, EG14], and the reader who is unfamiliar with the Breuil–Mézard conjecture is referred to these papers (and the references therein) for the original motivations for the conjecture, which come from automorphy lifting theorems. In this section we explain how Conjecture 6.1.14 gives a further refinement of the Breuil–Mézard conjecture (lifting the statement from an equality of cycles to an equality in a Grothendieck group of coherent sheaves).

By a *Serre weight* \underline{k} we mean a tuple of integers $\{k_{\bar{\sigma}, i}\}_{\bar{\sigma}: k \rightarrow \overline{\mathbb{F}}_p, 1 \leq i \leq d}$ with the properties that

- $p - 1 \geq k_{\bar{\sigma}, i} - k_{\bar{\sigma}, i+1} \geq 0$ for each $1 \leq i \leq d - 1$, and
- $p - 1 \geq k_{\bar{\sigma}, d} \geq 0$, and not every $k_{\bar{\sigma}, d}$ is equal to $p - 1$.

The set of Serre weights is in bijection with the set of irreducible $\overline{\mathbb{F}}_p$ -representations of $\mathrm{GL}_d(k_F)$, via passage to highest weight vectors (see for example the appendix to [Her09]); for each Serre weight \underline{k} , we write $F_{\underline{k}}$ for the corresponding irreducible k -representation of $\mathrm{GL}_d(k_F)$.

We write $\bar{\sigma}^{\mathrm{crys}}(\lambda, \tau)$ for the semisimplification of the k -representation of $\mathrm{GL}_d(k_F)$ given by $W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau) \otimes_{\mathcal{O}} k$. There are unique integers $n_{\underline{k}}^{\mathrm{crys}}(\lambda, \tau)$ such that

$$\bar{\sigma}^{\mathrm{crys}}(\lambda, \tau) \cong \bigoplus_{\underline{k}} F_{\underline{k}}^{\oplus n_{\underline{k}}^{\mathrm{crys}}(\lambda, \tau)}.$$

Write $Z_{\mathrm{crys}, \underline{\lambda}, \tau}$ for the cycle (i.e. an element of the free abelian group $\mathbf{Z}[\mathcal{X}_{d, \mathrm{red}}]$ on the irreducible components of $\mathcal{X}_{d, \mathrm{red}}$) corresponding to the special fibre of $\mathcal{X}_d^{\mathrm{crys}, \underline{\lambda}, \tau}$. The following is [EG22, Conj. 8.2.2].

CONJECTURE 6.1.38. *There are cycles $Z_{\underline{k}}$ with the property that for each regular Hodge type $\underline{\lambda}$ and each inertial type τ , we have $Z_{\mathrm{crys}, \underline{\lambda}, \tau} = \sum_{\underline{k}} n_{\underline{k}}^{\mathrm{crys}}(\lambda, \tau) \cdot Z_{\underline{k}}$.*

Assume Conjecture 6.1.14. For each regular Hodge type $\underline{\lambda}$ and each inertial type τ we make a choice of lattice $\sigma^{\mathrm{crys}, \circ}(\tau)$ ¹² and set

$$\mathcal{F}_{\underline{\lambda}, \tau} = \mathfrak{A}\left(\left(\mathrm{c}\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau)\right)^\wedge\right),$$

which by Remark 6.1.26 is a coherent sheaf (concentrated in degree zero) on \mathcal{X}_d , and for each Serre weight \underline{k} we set

$$\mathcal{F}_{\underline{k}} := \mathfrak{A}(\mathrm{c}\text{-Ind}_K^G F_{\underline{k}}).$$

An easy induction using [EG14, Lem. 4.1.1] shows that $\mathcal{F}_{\underline{k}}$ is a coherent sheaf (again, concentrated in degree zero) on the special fibre of \mathcal{X}_d .

Since \mathfrak{A} is a functor, it follows from the definitions that we have an equality in $K_0(\mathrm{Coh}(\mathcal{X}_d))$

$$(6.1.39) \quad [\mathcal{F}_{\underline{\lambda}, \tau} \otimes_{\mathcal{O}} k] = \sum_{\underline{k}} n_{\underline{k}}^{\mathrm{crys}}(\lambda, \tau) \cdot [\mathcal{F}_{\underline{k}}].$$

¹²Since we are passing to underlying cycles, the precise choice of lattice is irrelevant.

As explained in Remark 6.1.34, $\mathcal{F}_{\lambda,\tau}$ is maximally Cohen–Macaulay over its support $\mathcal{X}_d^{\lambda,\tau}$, so that the support of $\mathcal{F}_{\lambda,\tau} \otimes_{\mathcal{O}} k$ is equal to $Z_{\text{crys},\lambda,\tau}$. Again, an easy induction using [EG14, Lem. 4.1.1] shows that the support of each $\mathcal{F}_{\underline{k}}$ is a cycle $Z_{\underline{k}}$ in $\mathbf{Z}[\mathcal{X}_{d,\text{red}}]$, so that Conjecture 6.1.38 is an immediate consequence of (6.1.39).

6.2. Expectations: the analytic case. Let $G = \text{GL}_d(F)$ considered as a \mathbf{Q}_p -analytic group. Moreover, we assume in this section that the fixed finite extension L of \mathbf{Q}_p contains the normal closure of F . We want to formulate a conjecture parallel to Conjecture 6.1.14 involving locally analytic representations and sheaves on \mathfrak{X}_d . In order to do so we first have to discuss derived categories in the context of locally analytic representations.

6.2.1. *Derived categories of locally analytic representations.* We write $\mathcal{D}(G)$ for the distribution algebra of G with coefficients in L . This distribution algebra is defined as the dual space of the space of locally analytic functions $G \rightarrow L$, see [ST02] for the precise definition (in our case the field of coefficients will always be understood and we suppress it from the notation).

By a theorem of Schneider–Teitelbaum [ST02, Corollary 3.3] the category of locally analytic representations of G on L -vector spaces of compact type with continuous G -morphisms is anti-equivalent to the category of separately continuous $\mathcal{D}(G)$ -modules on nuclear Fréchet spaces with continuous $\mathcal{D}(G)$ -module homomorphisms. This anti-equivalence is given by mapping an L -vector space V of compact type to its strong dual $V' = V'_b$. Under this anti-equivalence of categories the admissible locally analytic G -representations (by definition) correspond to coadmissible $\mathcal{D}(G)$ -modules, i.e. to those $\mathcal{D}(G)$ -modules M that are coadmissible over the distribution algebra $\mathcal{D}(K)$ of the maximal compact subgroup $K = \text{GL}_n(\mathcal{O}_F)$ of G . More precisely this means that M is an inverse limit of finite type modules over certain Banach completions $\mathcal{D}_\tau(K)$ of $\mathcal{D}(K)$, compare [ST03, §3, 4]. In the following we will write $\text{Rep}_L^{\text{an}} G$ for the category of locally analytic G -representations on L -vector spaces of compact type.

We note moreover that the G -action on a locally analytic representation π automatically extends to a separately continuous action of the distribution algebra $\mathcal{D}(G)$, i.e. the representation π can be regarded as a $\mathcal{D}(G)$ -module [ST02, Prop. 3.2]. Similarly to the category $D(\text{sm. } G)$ we would like to consider the derived category $D(\text{an. } G)$ of complexes of locally analytic representations. The definition of this category (and of derived functors from this category to other derived categories) meets some difficulties, as we are working with topological objects in a homological algebra context. These problems are best dealt with by using condensed structures and solid modules introduced by Clausen–Scholze. A treatment of locally analytic representations in the framework of solid modules was developed by Rodrigues Jacinto and Rodríguez Camargo in [JC21]. We will not go too much into detail concerning these definitions and rather use $D(\text{an. } G)$ as an informal object that will contain all the locally analytic representations that we are interested in. Still we try to give some details as to how a formal definition of $D(\text{an. } G)$ might look.

In terms of [JC21]¹³ the most natural candidate for $D(\text{an. } G)$ is to construct $D(\text{an. } G)$ as a full subcategory of the derived category $D(L_{\blacksquare}[G])$ of solid $L_{\blacksquare}[G]$ -modules. If C is an object in $D(L_{\blacksquare}[G])$, then C is called *derived locally analytic* if the natural map $C^{\text{Rla}} \rightarrow C$ from the derived locally analytic vectors in C to C is a quasi-isomorphism (compare [JC21, Definition 4.40]). In this case (by loc. cit.) C is a (complex of) $\mathcal{D}(G)$ -module(s) extending the $L[G]$ -module structure in a unique way. We define the candidate $D(\text{an. } G)$ for our derived category as the full subcategory of $D(L_{\blacksquare}[G])$ consisting of derived locally analytic objects. At least conjecturally this category contains all locally analytic G -representations [JC21, Proposition 4.43] as well as bounded complexes of those.

In particular we can regard the subcategory $D_{\text{f.p.}}^b(\text{Rep}^{\text{sm}} G)$ of bounded complexes of smooth G -representations on L vector spaces with finitely presented cohomology as a subcategory of $D(\text{an. } G)$. We note that this is not a full subcategory, as the computation of derived homomorphisms involves the computation of Lie-algebra cohomology, compare [JC21, Theorem 5.5].

REMARK 6.2.2. We would like to have a full subcategory of $D(\text{an. } G)$ whose objects satisfy a finiteness condition that will ensure that the hypothetical functor $\mathfrak{A}_G^{\text{rig}}$ in Conjecture 6.2.4 below will produce (bounded complexes of) coherent sheaves. As our use of $D(\text{an. } G)$ is in any case rather informal we will write $D_{\text{f.p.}}^b(\text{an. } G)$ for this full subcategory. In analogy with the Banach case one might speculate that this category is the full subcategory of bounded complexes whose cohomology satisfies some finiteness condition over the distribution algebra $\mathcal{D}(G)$. A precise (and still conjectural) description would involve talking in depth about the condensed set-up used in [JC21]. As in these notes the category $D(\text{an. } G)$ is not properly defined we use $D_{\text{f.p.}}^b(\text{an. } G)$ as a black box for a derived category of locally analytic representations satisfying appropriate finiteness conditions. In any case $D_{\text{f.p.}}^b(\text{an. } G)$ should contain the derived category $D_{\text{f.p.}}^b(\text{Rep}^{\text{sm}} G)$.

6.2.3. *Formulation of the conjecture.* Recall that the center $Z(\mathfrak{g}_L)$ of the enveloping algebra $U(\mathfrak{g}_L)$ of the Lie algebra \mathfrak{g}_L of the \mathbf{Q}_p -analytic group G embeds into $\mathcal{D}(G)$, [ST02, Prop.3.7]. This center can be identified with the global sections $\Gamma(\text{WT}_{d,L}, \mathcal{O}_{\text{WT}_{d,L}})$ of the space of Hodge–Tate–Sen weights of \mathfrak{X}_d using the Harish–Chandra isomorphism.¹⁴

Further recall [BZCHN20, Theorem 5.13], [Zhu20, Conjecture 4.5.1], [Hel20, Conjecture 3.6] (see also Section 8.3) that for smooth representations we have a functor

$$\mathfrak{A}_G^{\text{sm}} : D_{\text{f.p.}}^b(\text{Rep}^{\text{sm}} G) \longrightarrow D_{\text{coh}}^b(\text{WD}_{d,F}).$$

(For a general reductive group this functor is of course conjectural but for $G = \text{GL}_d(F)$ the functor can be constructed from the case of the Iwahori–Hecke algebra using type theory, see [BZCHN20].)

¹³We caution the reader that in [JC21] the group G usually is a uniform pro- p -group. As our reductive groups G are locally uniform pro- p groups and as the condition of being locally analytic only depends on small neighborhoods of 1, we can easily adapt the notions and results from op. cit.

¹⁴Recall that in this section the group G is GL_d and hence we can easily compare the center of the enveloping algebra with the space of functions on the space of Hodge–Tate–Sen weights. The case of an arbitrary reductive group is harder and would rely on similar constructions as in [DPS20, 4.6]. But anyway, we did not define spaces of (φ, Γ) -modules with G -structures for arbitrary reductive groups.

We formulate a conjecture about locally analytic representations and coherent sheaves on \mathfrak{X}_d that roughly parallels Conjecture 6.1.14 in the case of smooth mod p representations and sheaves on \mathcal{X}_d .

CONJECTURE 6.2.4. *There exists an exact $Z(\mathfrak{g}_L)$ -linear functor*

$$\mathfrak{A}_G^{\text{rig}} : D_{\text{f.p.}}^b(\text{an. } G) \rightarrow D_{\text{coh}}^b(\mathfrak{X}_d)$$

satisfying the following conditions:

- (1) *(Compatibility with the smooth case.) Let $\underline{\xi}$ be a dominant algebraic character associated to a regular Hodge–Tate weight $\underline{\lambda}$, and write $W_{\underline{\lambda}}[1/p]$ for the corresponding irreducible algebraic representation of G of highest weight $\underline{\xi}$ as in Definition 2.3.2. Let $\text{Rep}_{\Omega_{\tau}}^{\text{sm}} G$ be the Bernstein block associated to an inertial type τ such that $\tau|_{I_{F'}}$ is trivial, for some finite extension F' of F . After base change to $L' \supset F'$ there is a natural isomorphism*

$$(6.2.5) \quad \mathfrak{A}_G^{\text{rig}}(W_{\underline{\lambda}}[1/p] \otimes -) \cong (\text{pr}_{\text{WD}}^*(\mathfrak{A}_G^{\text{sm}}(-))((- \rho')_{\sigma:F \hookrightarrow L}),$$

where

$$\text{pr}_{\text{WD}} : \text{Fil}_{\underline{\lambda}} \text{Mod}_{d,\varphi,N,\tau} \cong \text{Fil}_{\underline{\lambda}} \text{WD}_{d,F,\tau} \longrightarrow \text{WD}_{d,F,\tau}$$

is the canonical projections.

- (2) *(Compatibility with parabolic induction.) Let $P \subseteq G$ be a parabolic subgroup with Levi quotient M and write \hat{P} respectively \hat{M} for the corresponding dual groups. Then locally analytic parabolic induction*

$$\text{Ind}_{\bar{P}}^G(-)^{\text{an}} : D(\text{an. } M) \rightarrow D(\text{an. } G),$$

where $\bar{P} \subset G$ is the opposite parabolic, satisfies the compatibility

$$(6.2.6) \quad \mathfrak{A}_G^{\text{rig}}(\text{Ind}_{\bar{P}}^G(-)^{\text{an}}) \cong \tilde{\beta}_{\hat{P},*} \circ \tilde{\alpha}_{\hat{P}}^*(\mathfrak{A}_M^{\text{rig}}(-)([F : \mathbf{Q}_p]\rho'_M))$$

whenever this formula makes sense, see Section 6.2.13 for a more precise formulation.

REMARK 6.2.7. Let us comment on the twists by $-\rho'$ respectively by ρ'_M that occur in (6.2.5) respectively (6.2.6).

(i) Let us, as usual, write ρ for the half sum of the positive roots in the dual torus \hat{T} . To fix notations let us fix \hat{T} to be the diagonal torus in GL_d and we canonically identify $\hat{T} = \mathbf{G}_m^d$. Then

$$\rho = \left(\frac{d-1}{2}, \frac{d-3}{2}, \dots, \frac{-d+3}{2}, \frac{-d+1}{2} \right)$$

in general is not a character of \hat{T} , but only lies in $X^*(\hat{T}) \otimes \mathbf{Q}$. Instead we consider the character $\rho' = (0, -1, \dots, -(d-1)) \in X^*(\hat{T})$. Note that ρ' is a shift of ρ by an element of $X^*(\hat{T}) \otimes \mathbf{Q}$ that factors through the determinant and agrees with ρ when evaluated on the coroots of \hat{T} .

If M is a Levi subgroup of G , then we write ρ_M for the half sum of the roots that occur in the \hat{T} representation on $\text{Lie } \hat{U}$, where $\hat{U} \subset \hat{P}$ is the unipotent radical. Then again we can apply a shift to ρ_M to obtain a character ρ'_M of $M = \prod_{i=1}^r \text{GL}_{r_i}$ that factors through the maximal torus quotient $S_{\hat{M}} = \prod_{i=1}^r \mathbf{G}_m$ such that ρ'_M is trivial on the first \mathbf{G}_m factor of $S_{\hat{M}}$.

(ii) Let us now describe the twist in (6.2.5). The space $\mathrm{Fil}_{\underline{\lambda}}\mathrm{WD}_{d,F,\tau}$ has a canonical projection to

$$\mathrm{GL}_d \backslash (\mathrm{Res}_{F/\mathbf{Q}_p} \mathrm{GL}_d)_L / P_{\underline{\lambda}} \cong \mathrm{GL}_d \backslash \left(\prod_{\sigma: F \hookrightarrow L} \mathrm{GL}_d / \hat{B} \right).$$

Here $P_{\underline{\lambda}}$ denotes the parabolic defined in Section 5.2 that agrees with the Weil restriction of the Borel, as $\underline{\lambda}$ is regular. Then we can twist a sheaf on $\mathrm{Fil}_{\underline{\lambda}}\mathrm{WD}_{d,F,\tau}$ by the pullback of a line bundle on the product of flag varieties that we choose to be the line bundle on GL_d / \hat{B} associated with the character $-\rho'$ of \hat{T} in each factor. We point out that twisting with $(-\rho')_{\sigma}$ is compatible with duality (in the spirit of the compatibility in Remark 6.1.32). Again we have an auto-duality $\mathcal{D}_{\mathfrak{X}_d}$ on $D_{\mathrm{coh}}^b(\mathfrak{X}_d)$ induced by composing Grothendieck–Serre duality with the pullback of a map $\mathfrak{X}_d \rightarrow \mathfrak{X}_d$ induced by the main involution on GL_d . If \mathcal{F} denotes a complex of sheaves on $\mathrm{WD}_{d,F,\tau}$, then

$$\mathcal{D}_{\mathfrak{X}_d}(\mathrm{pr}_{\mathrm{WD}}^*(\mathcal{F})(-\rho')_{\sigma}) = \mathrm{pr}_{\mathrm{WD}}^*(\mathcal{D}_{\mathrm{WD}_{d,F,\tau}}(\mathcal{F}))(-\rho')_{\sigma},$$

where $\mathcal{D}_{\mathrm{WD}_{d,F,\tau}}$ is defined the same way as $\mathcal{D}_{\mathfrak{X}_d}$.

(iii) Finally let us describe the twist in (6.2.6). The determinant induces a canonical morphism $\mathfrak{X}_d \rightarrow \mathfrak{X}_1 \rightarrow */\mathbf{G}_m$ (see (7.1.2) for a description of \mathfrak{X}_1 and the canonical projection to $*/\mathbf{G}_m$). In particular we obtain a projection

$$\mathfrak{X}_{\hat{M}} \rightarrow */S_{\hat{M}}$$

and hence it makes sense to twist sheaves on $\mathfrak{X}_{\hat{M}}$ with a character of $S_{\hat{M}}$ (that we view as a line bundle on $*/S_{\hat{M}}$ and pull back to $\mathfrak{X}_{\hat{M}}$).

There should be further compatibilities, namely a compatibility with the Banach case discussed in Section 6.2.10 below and a local-global compatibility that we will talk about in Section 9.6. We now comment a bit on the compatibilities in the conjecture.

REMARK 6.2.8.

(a) Property (1) prescribes the functor $\mathfrak{A}_G^{\mathrm{rig}}$ on locally algebraic representations. This means that in some sense we expect that the extension of the functor $\mathfrak{A}_G^{\mathrm{sm}}$ in smooth representations to locally algebraic representations (which takes values in coherent sheaves supported on de Rham loci) can be interpolated to a functor that takes values in coherent sheaves on \mathfrak{X}_d (without restrictions on the support). In some sense this parallels the point of view that (Hecke eigenvalues on) p -adic automorphic forms interpolate (Hecke eigenvalues on) classical automorphic forms.
 (b) The natural isomorphism in Property (1) depends on the choice of the isomorphism

$$\mathrm{Fil}_{\underline{\lambda}}\mathrm{Mod}_{d,\varphi,N,\tau} \cong \mathrm{Fil}_{\underline{\lambda}}\mathrm{WD}_{d,F,\tau}$$

in Lemma 5.2.3, i.e. on the choice of an embedding $F' \hookrightarrow L'$. In fact we expect that there is a variant of the functor $\mathfrak{A}_G^{\mathrm{sm}}$ that is a functor

$$D_{\mathrm{f.p.}}^b(\mathrm{Rep}^{\mathrm{sm}}G) \longrightarrow D_{\mathrm{coh}}^b\left(\bigcup_{F'/F} \mathrm{Mod}_{d,\varphi,N,F'/F}\right)$$

which agrees with $\mathfrak{A}_G^{\mathrm{sm}}$ after appropriate scalar extensions, see Remark 8.3.3. This version of $\mathfrak{A}_G^{\mathrm{sm}}$ would allow for a smoother formulation of Property (1) (i.e. a version without an auxiliary scalar extension to a field containing F').

(c) We cannot hope that the functor $\mathfrak{A}_G^{\text{rig}}$ is fully faithful without putting any additional conditions on the source $D_{\text{f.p.}}^b(\text{an. } G)$. In the GL_1 -case this is discussed in Remark 7.1.8 below. But the phenomenon is also visible by looking at locally algebraic representations. The functor $\mathfrak{A}_G^{\text{sm}}$ is (conjecturally) fully faithful. Given two smooth representations π_1 and π_2 (concentrated in degree zero) the module of homomorphisms $\text{Hom}_G(\pi_1, \pi_2)$ is the same whether we compute it in the category of smooth representations or in the category of locally analytic representations (of course this is not true for higher Ext-groups). By full faithfulness this module coincides with $\text{Hom}(\mathfrak{A}_G^{\text{sm}}(\pi_1), \mathfrak{A}_G^{\text{sm}}(\pi_2))$, but in general the space of homomorphisms between the analytifications of these sheaves is strictly larger. It seems that this issue will not occur if we either add a *finite slope* condition, or ask that the representations are admissible; compare the discussion in the GL_1 -case. Unfortunately, simultaneously imposing the assumptions of being finitely presented over $\mathcal{D}(G)$ and being admissible is rather restrictive, and would rule out a lot of interesting representations.

(d) We expect that the functor $\mathfrak{A}_G^{\text{rig}}$ can be defined on a larger category like $D(\text{an. } G)$ instead of $D_{\text{f.p.}}^b(\text{an. } G)$ if we allow more general (complexes of) sheaves than just coherent sheaves. Similarly to the case of p -power torsion representations we expect that in this case admissibility of a representation π will not ensure that $\mathfrak{A}_G^{\text{rig}}(\pi)$ is a coherent sheaf without assuming some additional finiteness assumption on the locally analytic representation regarded as a module over $\mathcal{D}(G)$. In the p -torsion case this is discussed in Section 7.7.59, and one knows [Paš10, Thm. 1.1] that, when $G = \text{GL}_2(F)$, there are unitary Banach completions of locally algebraic G -representations whose reduction contains any given irreducible admissible mod p G -representation, so that we expect the mod p phenomena of that discussion to have analogues in characteristic zero.

REMARK 6.2.9. We expect that, similarly to Remark 6.1.23, the functor $\mathfrak{A}_G^{\text{rig}}$ of Conjecture 6.2.4 should be given by

$$\pi \longmapsto \mathcal{L}_\infty \widehat{\otimes}_{\mathcal{D}(G)}^L \pi$$

for a certain derived completed tensor product $-\widehat{\otimes}^L -$ (most likely: defined in some category of solid modules) and a family \mathcal{L}_∞ of $\mathcal{D}(G)$ -modules over the stack \mathfrak{X}_d such that, using the map π_d from (5.1.27), its pullback to the generic fiber of \mathfrak{X}_d should be given by

$$\pi_d^* \mathcal{L}_\infty \cong L_\infty \otimes_{\mathcal{O}_{\mathfrak{X}_d}[[G]]} \mathcal{D}(G, \mathcal{O}_{\mathfrak{X}_d})$$

for a relative version $\mathcal{D}(G, \mathcal{O}_{\mathfrak{X}_d})$ of the distribution algebra defined in similar terms as in [BHS17b, 3.1] or [GRJ18, 1.3].

In the case of $\text{GL}_2(\mathbf{Q}_p)$ we expect that for a point $x \in \mathfrak{X}_2(L)$, given by a (φ, Γ) -module D , the $\mathcal{D}(G)$ -module $\mathcal{L}_\infty \otimes k(x)$ is the dual of the locally analytic representation associated to D by Colmez's extension of the p -adic Langlands correspondence to the case of (not necessarily étale) (φ, Γ) -modules over \mathcal{R}_F , [Col16]. More precisely: the family \mathcal{L}_∞ should interpolate Colmez's p -adic Langlands correspondence for (φ, Γ) -modules over \mathcal{R}_L . Using the object L_∞ from Remark 6.1.23 we can construct \mathcal{L}_∞ (in the case of $\text{GL}_2(\mathbf{Q}_p)$) on the open substack of \mathfrak{X}_2 , where the (φ, Γ) -modules are étale up to twist. The complement of this open subset consists only of trianguline (φ, Γ) -modules. On the other hand Gaisin and Rodrigues Jacquinto [GRJ18] construct an analogue of \mathcal{L}_∞ over (the regular part of) the stack

\mathfrak{X}_B parameterizing trianguline (φ, Γ) -modules (together with a triangulation). At least for the time being it is not clear how to glue \mathcal{L}_∞ from these two cases.

6.2.10. *Compatibility with the Banach case.* We expect a strong link between the conjecture in the Banach case (Conjecture 6.1.14) and the conjecture in the analytic case (Conjecture 6.2.4). On the representation theoretic side, this compatibility should involve passage from a lattice Λ in a Banach space representation $V = \Lambda[1/p]$ (that we view as a pro-object in $D_{\text{f.p.}}^b(\text{sm. } G)$) to the locally analytic vectors (cf. [ST03, §7]) in V . On the side of (coherent) sheaves on stacks of (φ, Γ) -modules, we consider the generic fiber \mathcal{F}_η of a pro-object in $D_{\text{coh}}^b(\mathcal{X}_d)$. This (at least heuristically, but see the discussion following Lemma 7.1.7 for a discussion in the case of GL_1) gives rise to a complex of sheaves on $\mathcal{X}_{d,\eta}^{\text{rig}}$, and we compare this complex of sheaves to the pullback along (5.1.27) of a complex of sheaves on \mathfrak{X}_d . More precisely, for a pro-object $\Lambda \in D_{\text{f.p.}}^b(\text{sm. } G)$ with $V = \Lambda[1/p]$, the comparison should look like

$$(6.2.11) \quad \mathfrak{A}(\Lambda)_\eta^{\text{rig}} \cong \pi_d^* \mathfrak{A}_G^{\text{rig}}(V^{\text{an}}).$$

This formula meets with some difficulties due to the finiteness conditions that we have imposed. We expect that these difficulties can be dealt with by using solid modules on both sides (i.e. solid locally analytic representations and solid quasi-coherent sheaves) but we will not pursue this direction here. Instead we discuss some of the difficulties that one has to deal with.

- We have restricted ourselves to coherent sheaves on rigid analytic spaces (and rigid analytic Artin stacks), though there should be a more flexible ambient category of sheaves on these spaces (defined in terms of solid modules), compare Remark 5.1.11. Though it is possible to put conditions on pro-coherent sheaves \mathcal{F} on \mathcal{X}_d that will assure that \mathcal{F} has a well-defined generic fiber $\mathcal{F}_\eta^{\text{rig}}$ that is a coherent sheaf in $\mathcal{X}_{d,\eta}^{\text{rig}}$, it is less clear which finiteness conditions on the pro-object Λ of $D(\text{sm. } G)$ would assure that $\mathfrak{A}(\Lambda)_\eta^{\text{rig}}$ is an object of $D_{\text{coh}}(\mathcal{X}_{d,\eta}^{\text{rig}})$. Given a single representation Λ (say in cohomological degree zero) the finiteness assumption should at least involve finiteness over $\mathcal{O}[[G]]$.
- In our context of derived categories the passage to the locally analytic vectors $V \mapsto V^{\text{an}}$ should be derived. It turns out that this is not necessary on the category of admissible Banach space representations. More precisely, if V is an admissible Banach space representation (i.e. its dual is finitely generated over $L[[K]]$), then V^{an} is an admissible locally analytic representation and its dual may be described as follows: By [ST03, §4] there is a canonical flat map $\mathcal{O}[[K]] \rightarrow \mathcal{D}(K)$ from the completed group ring $\mathcal{O}[[K]]$ to the distribution algebra $\mathcal{D}(K)$. Given a finitely generated $\mathcal{O}[[K]]$ -module π we can consider the extension of scalars $\pi \otimes_{\mathcal{O}[[K]]} \mathcal{D}(K)$. If π is equipped with a G -action (compatible with the $\mathcal{O}[[K]]$ -module structure), then $\pi \otimes_{\mathcal{O}[[K]]} \mathcal{D}(K)$ naturally becomes a coadmissible $\mathcal{D}(G)$ -module. As $\mathcal{O}[[K]] \rightarrow \mathcal{D}(K)$ is flat, the functor $(-)^{\text{an}}$ is exact on the category of admissible Banach space representations, and one can in fact deduce that V^{an} is even derived locally analytic, see [JC21, Prop. 4.48], [Pan22b, Theorem 2.2.3]. In general, i.e. without invoking admissibility,

we have to consider a derived variant of $V \mapsto V^{\text{an}}$ (that is the derived locally analytic vectors of [JC21]) that can only be defined in the framework of solid modules.

- As discussed in Remark 6.2.8, the admissibility condition is not a good finiteness condition when dealing with the functors \mathfrak{A} and $\mathfrak{A}_G^{\text{rig}}$, i.e. we should not expect that admissible representations are mapped to coherent sheaves in general. On the other hand imposing both finiteness conditions, admissibility and finiteness over $\mathcal{O}[[G]]$, respectively $\mathcal{D}(G)$, is rather restrictive and excludes a lot of interesting representations. Hence it seems that a general formulation of the compatibility (6.2.11) can only be formulated if we enlarge our ambient categories relying on the theory of solid modules.

One instance of the compatibility (6.2.11) of the Banach and the analytic case should be the following. Given a regular Hodge–Tate weight $\underline{\lambda}$ and an inertial type τ , Conjecture 6.1.14 (4) predicts that

$$\mathcal{F} = \mathfrak{A}\left(\widehat{(c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\text{crys}, \circ}(\tau))}\right)$$

is a pro-coherent sheaf concentrated in degree 0 and supported on the p -adic stack $\mathcal{X}_d^{\text{crys}, \underline{\lambda}, \tau}$. We expect that its pullback along any morphism $\text{Spf } A \rightarrow \mathcal{X}_d^{\text{crys}, \underline{\lambda}, \tau}$ for a p -adically complete \mathcal{O} -algebra topologically of finite type over \mathcal{O} is given by a finitely generated A -module. In this case there is a well defined generic fiber \mathcal{F}_η of \mathcal{F} which is a coherent sheaf on $(\mathcal{X}_d^{\text{crys}, \underline{\lambda}, \tau})_\eta^{\text{rig}}$. We then expect that there is an isomorphism

$$(6.2.12) \quad \mathcal{F}_\eta^{\text{rig}} \cong \pi_d^*\left(\mathfrak{A}_G^{\text{rig}}(W_{\underline{\lambda}}[1/p] \otimes c\text{-Ind}_K^G \sigma^{\text{crys}}(\tau))\right).$$

We note that (6.2.12) is not quite of the same form as (6.2.11). However, we note that since $W_{\underline{\lambda}}[1/p]$ is a G -representation, there is an isomorphism

$$c\text{-Ind}_K^G W_{\underline{\lambda}}[1/p] \otimes_{\mathcal{O}} \sigma^{\text{crys}}(\tau) \cong W_{\underline{\lambda}}[1/p] \otimes c\text{-Ind}_K^G \sigma^{\text{crys}}(\tau),$$

and it seems to be reasonable to expect that the locally analytic vectors in

$$\widehat{(c\text{-Ind}_K^G W_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\text{crys}, \circ}(\tau))} \otimes_{\mathcal{O}} L$$

are given by $c\text{-Ind}_K^G W_{\underline{\lambda}}[1/p] \otimes_{\mathcal{O}} \sigma^{\text{crys}}(\tau)$.

Finally let us point out that Colmez [Col19, Rem. 0.2] has described the Jordan–Hölder factors of the locally analytic vectors $\Pi(D)^{\text{an}}$ in the Banach space representation Π associated to an étale (φ, Γ) -module D of rank 2 under the p -adic local Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$. It would be interesting to predict where the representations in [Col19, Corollaire 0.4] are mapped to under the functor $\mathfrak{A}_G^{\text{rig}}$. In the trianguline case this is basically covered by the expected compatibilities with the smooth case and with parabolic induction. But for the remaining representations we do not know how to (conjecturally) describe the image.

6.2.13. Compatibility with parabolic induction. We fix a Borel subgroup $B \subseteq G$ with split maximal torus T . Further let $P \supseteq B$ be a parabolic subgroup with Levi M . Let (π, V) be a locally analytic representation of M on an L -vector space V of compact type. The locally analytic parabolic induction of π is defined by

$$(\text{Ind}_P^G \pi)^{\text{an}} = \{f : G \rightarrow V \text{ locally analytic} \mid f(gp) = \pi(p)^{-1}f(g) \text{ for all } p \in P, g \in G\}.$$

In particular we obtain functors

$$\begin{aligned} \mathrm{Ind}_P^G(-)^{\mathrm{an}} &: \mathrm{Rep}_L^{\mathrm{an}} M \longrightarrow \mathrm{Rep}_L^{\mathrm{an}} G, \\ \mathrm{Ind}_P^G(-)^{\mathrm{an}} &: D(\mathrm{an}. M) \longrightarrow D(\mathrm{an}. G), \end{aligned}$$

where by abuse of notation we use the same notation for the functor on the (informally defined) derived categories. In order to make the discussion less involved we focus on the case $P = B$ in the following.

EXPECTATION 6.2.14. *There are isomorphisms*

$$\mathfrak{A}_G^{\mathrm{rig}}(\mathrm{Ind}_B^G(\pi)^{\mathrm{an}}) \cong \tilde{\beta}_{\hat{B},*} \tilde{\alpha}_{\hat{B}}^*(\mathfrak{A}_T^{\mathrm{rig}}(\iota(\pi)([F : \mathbf{Q}_p]\rho')))$$

functorial in π , whenever π is a representation such that $\mathrm{Ind}_B^G(\pi)^{\mathrm{an}}$ lies in $D_{\mathrm{f.p.}}^b(\mathrm{an}. G)$ and such that the support of $\tilde{\alpha}_{\hat{B}}^(\mathfrak{A}_T^{\mathrm{rig}}(\pi))$ is proper over \mathfrak{X}_d .*

Here $\iota : D(\mathrm{an}. T) \rightarrow D(\mathrm{an}. T)$ twists a T -representation with the character

$$(6.2.15) \quad \delta_B \cdot (1, \dots, (\varepsilon \circ \mathrm{rec})^{i-1}, \dots, (\varepsilon \circ \mathrm{rec})^{d-1}).$$

and ε is the cyclotomic character and rec is the isomorphism of local class field theory.

REMARK 6.2.16. The appearance of the twist with the character (6.2.15) is due to the chosen normalizations. In fact the normalizations we choose depend on the normalization of the functor $\mathfrak{A}_G^{\mathrm{sm}}$, or (which is more or less the same), the normalization of the isomorphism in Proposition 5.2.10, compare Remark 5.2.11. In fact here we use a different normalization than in [Hel20], where compatibility with parabolic induction is stated without involving a twist (but using normalized parabolic induction). The normalization in [Hel20] corresponds to the unitary normalization of the local Langlands correspondence, rather than the non-unitary normalization that we use here. The price we have to pay is the appearance of the twist in Expectation 6.2.14.

Let us stress that locally analytic parabolic induction should not be expected to preserve the finiteness conditions, i.e for $\pi \in D_{\mathrm{f.p.}}^b(\mathrm{an}. T)$ we do not expect that $\mathrm{Ind}_B^G(\pi)^{\mathrm{an}}$ lies in $D_{\mathrm{f.p.}}^b(\mathrm{an}. G)$. However this should be true if π is of bounded slope. We expect that the phenomenon is directly related to the question whether the support of $\alpha_{\hat{B}}^*(\mathfrak{A}_T^{\mathrm{rig}}(\pi))$ is proper over \mathfrak{X}_d or not. As already remarked above we expect that there is a variant of $\mathfrak{A}_G^{\mathrm{rig}}$ defined on $D(\mathrm{an}. G)$ that takes values in a larger category of sheaves on \mathfrak{X}_d (most probably: solid quasi-coherent sheaves). In this context one should have a full 6-functor formalism for these sheaves at hand and the compatibility with parabolic induction 6.2.14 should rather be formulated with $\tilde{\beta}_{\hat{B},!}$ instead of $\tilde{\beta}_{\hat{B},*}$ (of course we don't see the difference if we assume that the support of the corresponding sheaf is proper).

We mention that Orlik and Strauch [OS15] have described the (topologically) irreducible subquotients of some locally analytic principal series representations. More precisely, they describe the irreducible subquotients of $(\mathrm{Ind}_B^G \delta)^{\mathrm{an}}$ for locally algebraic characters $\delta : T \rightarrow L^\times$. The main tool of their study is the construction of certain bi-functors $\mathcal{F}_P^G(-, -)$ that map a pair (M, π) consisting of a Lie-algebra representation M in the BGG category \mathcal{O} (such that the action of $\mathrm{Lie} P$ on M can be integrated to a P -action) and a smooth representation π of the Levi of P to a

locally analytic representation of G . We recall part of their construction. In fact it makes sense to speculate that the construction of Orlik–Strauch extends to functors

$$\begin{aligned}\mathcal{F}_P^G &: \text{Rep}_L^{\text{an}} M \longrightarrow \text{Rep}_L^{\text{an}} G, \\ \mathcal{F}_P^G &: D(\text{an. } M) \longrightarrow D(\text{an. } G),\end{aligned}$$

that behave similarly to parabolic induction, see Expectation 6.2.21 below.

Again we focus on the case $P = B$. In the following we write $U(\mathfrak{g}_L)$ for the universal enveloping algebra of $\mathfrak{g}_L = (\text{Lie } G) \otimes_{\mathbf{Q}_p} L$ (note that here we consider G as a \mathbf{Q}_p -analytic group and hence \mathfrak{g}_L is isomorphic to the product $\prod_{T:F \rightarrow L} \mathfrak{g}_L^T$). Similarly, we write \mathfrak{t}_L and \mathfrak{b}_L for the Lie algebras over L of the \mathbf{Q}_p -analytic groups T respectively B . Let us write $\mathcal{O}(U(\mathfrak{g}_L))$ for the BGG category \mathcal{O} of finitely generated $U(\mathfrak{g}_L)$ -modules on which the Cartan \mathfrak{t}_L acts semi-simply and which are locally finite as $U(\mathfrak{b}_L)$ -modules. We write $\mathcal{O}(U(\mathfrak{g}_L))^{\text{alg}}$ for the subcategory of objects M on which the action of \mathfrak{t}_L is algebraic. On this subcategory one can integrate the \mathfrak{b}_L -action to obtain a B -action on M . Then every object M of $\mathcal{O}(U(\mathfrak{g}_L))^{\text{alg}}$ can be regarded as a module over the sub- L -algebra $\mathcal{D}(\mathfrak{g}_L, B) \subseteq \mathcal{D}(G)$ generated by $U(\mathfrak{g}_L)$ and $\mathcal{D}(B)$. We obtain an exact *contravariant* functor

$$F_B^G(-) : \mathcal{O}(U(\mathfrak{g}_L))^{\text{alg}} \longrightarrow \text{Rep}^{\text{an}} G$$

given by

$$V \longmapsto (\mathcal{D}(G) \otimes_{\mathcal{D}(\mathfrak{g}, B)} V)';$$

compare [OS15, 3.4].

REMARK 6.2.17. On the dual side, in terms of $\mathcal{D}(G)$ -modules, the locally analytic parabolic induction is given by

$$((\text{Ind}_B^G \delta)^{\text{an}})'_b = \mathcal{D}(G) \otimes_{\mathcal{D}(B)} (L_\delta)'_b,$$

where L_δ denotes the 1-dimensional locally analytic representation of T on L via the character δ , see [OS15, Lemma 2.3]. Now let ξ be an algebraic character of T . The Verma module $M(\xi) = U(\mathfrak{g}_L) \otimes_{U(\mathfrak{b}_L)} L_\xi$ is an object of $\mathcal{O}(U(\mathfrak{g}_L))^{\text{alg}}$ and we find that

$$\mathcal{D}(G) \otimes_{\mathcal{D}(\mathfrak{g}_L, B)} M(\xi) = \mathcal{D}(G) \otimes_{\mathcal{D}(B)} L_\xi.$$

In particular, for an algebraic character ξ we find

$$F_B^G(M(-\xi)) = \text{Ind}_B^G(z^\xi)^{\text{an}}.$$

In fact Orlik and Strauch [OS15, 4.4] show that the functor $F_B^G(-)$ can in some sense be twisted by a smooth character of T : for any choice δ_{sm} of a smooth character of T there is an exact contravariant functor

$$F_B^G(-, \delta_{\text{sm}}) : \mathcal{O}(U(\mathfrak{g}_L))^{\text{alg}} \longrightarrow \text{Rep}^{\text{an}} G$$

such that for an algebraic character ξ the representation $F_B^G(M(-\xi), \delta_{\text{sm}})$ is the locally analytic induced representation $(\text{Ind}_B^G z^\xi \delta_{\text{sm}})^{\text{la}}$.

We point out the relation with locally algebraic representations whose smooth part is parabolically induced. Fix an integral, dominant and regular Hodge–Tate weight $\underline{\lambda}$ and write ξ for the corresponding character of \mathfrak{t}_L (via Definition 2.3.1), and $\mathcal{O}(U(\mathfrak{g}_L))_{(\xi)} \subseteq \mathcal{O}(U(\mathfrak{g}_L))$ for the block of the category \mathcal{O} containing the Verma module $M(\xi)$. Then $\mathcal{O}(U(\mathfrak{g}))_{(\xi)} \subseteq \mathcal{O}(U(\mathfrak{g}))^{\text{alg}}$. As $\underline{\lambda}$ is regular the character ξ is automatically dot-regular (i.e. the Weyl group orbit $W \cdot \xi$ of ξ under the dot-action $w \cdot \xi = w(\xi + \rho) - \rho$ consists of $|W|$ elements). The simple objects of

$\mathcal{O}(U(\mathfrak{g}_L))_{(\xi)}$ are the simple quotients $L(w \cdot \xi)$ of the Verma modules $M(w \cdot \xi)$. Moreover, $L(\xi) = W_\lambda[1/p]$ is the irreducible algebraic representation of G defined in Definition 2.3.2.

LEMMA 6.2.18. *There is a canonical isomorphism*

$$F_B^G(L(-\xi), \delta_{\text{sm}}) \cong L(\xi) \otimes_L (\text{Ind}_B^G \delta_{\text{sm}})^{\text{sm}}$$

with the locally algebraic representation given by the tensor product of the irreducible algebraic representation $L(\xi)$ and the smooth parabolic induction of the smooth character δ_{sm} .

Recall that the BGG category has an internal notion of duality $M \mapsto M^\vee$ (usually referred to as BGG duality) that is roughly given by mapping M to the direct sum of the duals of its (finite dimensional) \mathfrak{t} -eigenspaces and using the main involution on \mathfrak{g} (in order to pass from left to right modules and to obtain a $U(\mathfrak{g}_L)$ -module that is locally $U(\mathfrak{b}_L)$ -finite instead of $U(\mathfrak{b}_L)$ -finite, where \mathfrak{b} denotes the Lie-algebra of the opposite Borel \bar{B}). For a simple module M in $\mathcal{O}(U(\mathfrak{g})_L)$ we have $M^\vee \cong M$, but as $(-)^\vee$ is contravariant the functor exchanges subobjects and quotients. Let $\delta = z^\xi \delta_{\text{sm}}$ be a locally algebraic character with smooth part δ_{sm} and algebraic part ξ . We then define the representation

$$(6.2.19) \quad \mathcal{F}_B^G(\delta) = F_B^G(M(-\xi)^\vee, \delta_{\text{sm}}).$$

REMARK 6.2.20. It follows from [OS15] that $\text{Ind}_B^G(\delta)^{\text{an}}$ and $\mathcal{F}_B^G(\delta)$ have the same irreducible subquotients (even with the same multiplicities). Moreover, note that Lemma 6.2.18 together with the fact that there is a canonical surjection $M(-\xi)^\vee \rightarrow L(-\xi)$ implies that for dominant ξ the representation $\mathcal{F}_B^G(z^\xi \delta_{\text{sm}})$ has the locally algebraic representation $L(\xi) \otimes_L (\text{Ind}_B^G \delta_{\text{sm}})^{\text{sm}}$ as a quotient, whereas this locally algebraic representation appears as a subrepresentation in $\text{Ind}_B^G(z^\xi \delta_{\text{sm}})^{\text{an}}$.

We expect that the Orlik–Strauch functors can be used to define a variant of parabolic induction (whose derived versions of course will depend on the definition of the categories $D(\text{an. } T)$ and $D(\text{an. } G)$).

EXPECTATION 6.2.21. *There exist natural functors*

$$\begin{aligned} \mathcal{F}_B^G &: \text{Rep}_L^{\text{an}} T \longrightarrow \text{Rep}_L^{\text{an}} G, \\ \mathcal{F}_B^G &: D(\text{an. } T) \longrightarrow D(\text{an. } G), \end{aligned}$$

such that for a locally algebraic character δ the representation $\mathcal{F}_B^G(\delta)$ is given by (6.2.19).

REMARK 6.2.22. It seems reasonable to speculate that there are internal dualities \mathbf{D}_G resp. \mathbf{D}_T of $D(\text{an. } G)$ resp. $D(\text{an. } T)$ (similar to the dualities discussed in Section 6.1) such that (up to shift)

$$\mathcal{F}_B^G(-) = \mathbf{D}_G \circ \text{Ind}_B^G(-)^{\text{an}} \circ \mathbf{D}_T.$$

The restriction of \mathbf{D}_G to smooth representations should be given by the Zelevinsky involution, and the above expectation suggests that \mathbf{D}_G should be related to Serre-duality under the hypothetical functor $\mathfrak{A}_G^{\text{rig}}$.

Using the conjectural functor \mathcal{F}_B^G we can formulate a variant of Expectation 6.2.14. This expectation is motivated by the properties of Bezrukavnikov’s functor in Theorem 6.2.26 below.

EXPECTATION 6.2.23. *There are isomorphisms*

$$\mathfrak{A}_G^{\text{rig}}(\mathcal{F}_B^G(\pi)) \cong \tilde{\beta}_{\tilde{B},*} \tilde{\alpha}_{\tilde{B}}^! (\mathfrak{A}_T^{\text{rig}}(\iota(\pi)(-[F : \mathbf{Q}_p]\rho')))$$

functorial in π , whenever π is a representation such that $\mathcal{F}_B^G(\pi)$ lies in $D_{\text{f.p.}}^b(\text{an. } G)$ and such that the support of $\tilde{\alpha}_{\tilde{B}}^!(\mathfrak{A}_T^{\text{rig}}(\pi))$ is proper over \mathfrak{X}_d .

REMARK 6.2.24. Let us consider the case $F = \mathbf{Q}_p$ and $d = 2$ and recall from Remark 6.2.9 that we expect $\mathfrak{A}_G^{\text{rig}}$ to be of the form

$$\pi \mapsto \pi \widehat{\otimes}_{\mathcal{D}(G)}^L \mathcal{L}_\infty,$$

where \mathcal{L}_∞ is a family of $\mathcal{D}(G)$ -modules over \mathfrak{X}_2 . For a point $x \in \mathfrak{X}_2(L)$ represented by a (φ, Γ) -module D we expect that

$$\mathcal{L}_\infty \otimes k(x) = (\Pi(D)^{\text{an}})'$$

Here if D is étale, $\Pi(D)^{\text{an}}$ denotes the locally analytic vectors in the representation corresponding to D via the usual p -adic local Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$, and if D is not étale we use Colmez's extension [Col16] of the correspondence. Given a character $\delta : T \rightarrow L$, we are lead to expect that

$$\begin{aligned} H^0(\mathfrak{A}_G^{\text{rig}}(\text{Ind}_B^G(\delta))) \otimes k(x) &\cong (\text{Hom}_G(\text{Ind}_B^G(\delta), \Pi(D)^{\text{an}}))', \\ H^0(\mathfrak{A}_G^{\text{rig}}(\mathcal{F}_B^G(\delta))) \otimes k(x) &\cong (\text{Hom}_G(\mathcal{F}_B^G(\delta), \Pi(D)^{\text{an}}))'. \end{aligned}$$

The dimensions of the vector spaces on the right hand side can be determined explicitly, as Colmez has given a precise description of the locally analytic representations $\Pi(D)^{\text{an}}$ (see [Col14, Col16]), and in particular they are zero unless D is trianguline. Using Expectations 6.2.14 and 6.2.23 the computation of these fiber dimensions precisely matches the computation of fiber dimensions in Theorem 5.3.21 (v).

6.2.25. *Bezrukavnikov's functor on the BGG category.* We elaborate a bit further on the Orlik–Strauch functor and compatibility with parabolic induction. As on the Galois side parabolic induction is defined in terms of the compactification $\tilde{\mathfrak{X}}_B$ of \mathfrak{X}_B and as we can explicitly control the geometry of $\tilde{\mathfrak{X}}_B$ at certain regular points (Theorem 5.3.38), we can conjecturally link the evaluation of the functor $\mathfrak{A}_G^{\text{rig}}$ on certain Orlik–Strauch representations to a functor on the BGG category constructed by Bezrukavnikov [Bez16].

Fix λ and ξ as above. As every object in $\mathcal{O}(U(\mathfrak{g}_L))_{(\xi)}$ is an extension of objects that are (sub-)quotients of Verma modules $M(w \cdot \xi)$ for $w \in W$, the expectations on the functor $\mathfrak{A}_G^{\text{rig}}$ imply that given δ_{sm} , there should be a functor

$$\mathfrak{A}_G^{\text{rig}} \circ F_B^G((-)^\vee, \delta_{\text{sm}}) : \mathcal{O}(U(\mathfrak{g}_L))_{(\xi)} \longrightarrow D_{\text{coh}}^b(\mathfrak{X}_d)$$

such that the sheaves in the image of this functor are supported on the closed substack $\mathfrak{X}_{d,(\lambda, \delta_{\text{sm}})-\text{tri}}$ defined just before Theorem 5.3.38. We use an additional duality $(-)^\vee$ in this functor to make the functor covariant. Using the local description of the stack $\tilde{\mathfrak{X}}_{d,(\lambda, \delta_{\text{sm}})-\text{tri}}$ in Theorem 5.3.38 we can construct this functor, relying on work of Bezrukavnikov.

Let us write

$$H = (\text{Res}_{F/\mathbf{Q}_p} \text{GL}_d)_L = \prod_{\sigma: F \hookrightarrow L} \text{GL}_d$$

for the remainder of this section. We note that $\mathfrak{g}_L \cong \mathrm{Lie}H = \mathfrak{h}$ as Lie algebras over L . In particular $\mathcal{O}(U(\mathfrak{g}_L))$ and $\mathcal{O}(U(\mathfrak{h}))$ are isomorphic. Recall from Section 5.3.22 (where we write G instead of H) that we write

$$\overline{X}_w \subseteq \tilde{\mathfrak{h}} \times_{\mathfrak{h}} \tilde{\mathfrak{h}}$$

for the irreducible component that dominates the closure of the Schubert cell $H(1, w) \subseteq H/B_H \times H/B_H$, where $B_H \subseteq H$ is the Weil restriction of the upper triangular matrices in GL_d .

Moreover, we note that there is a canonical projection

$$\omega : \tilde{\mathfrak{h}} \times_{\mathfrak{h}} \tilde{\mathfrak{h}} \xrightarrow{\mathrm{pr}_1} \tilde{\mathfrak{h}} \longrightarrow \widehat{\mathfrak{t}},$$

(where $\widehat{\mathfrak{t}}$ is the dual Lie algebra of \mathfrak{t}_L) and we write $\overline{X}_{w,0}$ for the intersection of \overline{X}_w with the preimage of 0. Finally note that the zero section gives rise to a copy

$$\begin{aligned} (\mathrm{Res}_{F/\mathbf{Q}_p} \mathrm{GL}_d/P_{\underline{\lambda}})_L \times (\mathrm{Res}_{F/\mathbf{Q}_p} \mathrm{GL}_d/P_{\underline{\lambda}})_L &= H/B_H \times H/B_H \\ &\subseteq \overline{X}_{w_0,0} \subseteq \overline{X}_{w_0} \subseteq \tilde{\mathfrak{h}} \times_{\mathfrak{h}} \tilde{\mathfrak{h}}. \end{aligned}$$

Here $P_{\underline{\lambda}}$ denotes the parabolic subgroup associated to $\underline{\lambda}$ as in Section 5.2 which agrees with the Borel B as we assume $\underline{\lambda}$ to be regular. The following functor has been constructed by Bezrukavnikov in [Bez16, 11.4], see also [BK, Theorem 13.0.1].

THEOREM 6.2.26 (Bezrukavnikov). *There is an exact functor*

$$F_{\xi} : \mathcal{O}(U(\mathfrak{g}_L))_{(\xi)} \longrightarrow \mathrm{Coh}^H((\tilde{\mathfrak{h}} \times_{\mathfrak{h}} \tilde{\mathfrak{h}})_{\widehat{0}})$$

such that

(i) *For $w \in W$ the image of the of the Verma module $M(w w_0 \cdot \xi)$ is the dualizing sheaf of \overline{X}_w :*

$$F_{\xi}(M(w w_0 \cdot \xi)) = \omega_{\overline{X}_w}.$$

(ii) *For $w \in W$ the image of the of the dual Verma module $M(w w_0 \cdot \xi)^{\vee}$ is the structure sheaf of \overline{X}_w :*

$$F_{\xi}(M(w w_0 \cdot \xi)^{\vee}) = \mathcal{O}_{\overline{X}_w}.$$

(iii) *The image of the irreducible algebraic representation $L(\xi)$ is the line bundle on*

$$(\mathrm{Res}_{F/\mathbf{Q}_p} \mathrm{GL}_d/P_{\underline{\lambda}})_L \times (\mathrm{Res}_{F/\mathbf{Q}_p} \mathrm{GL}_d/P_{\underline{\lambda}})_L = H/B_H \times H/B_H \subseteq \overline{X}_{w_0,0}$$

associated to the character $((-\rho')_{\sigma}, (-\rho')_{\sigma})$ (compare Remark 6.2.7 for the definition of ρ').

We expect that, using the local description of the stack $\widehat{\mathfrak{X}}_{d,(\underline{\lambda}, \delta_{\mathrm{sm}}) - \mathrm{tri}}$ via the map

$$f_{\underline{\lambda}, \delta_{\mathrm{sm}}} : \widehat{\mathfrak{X}}_{d,(\underline{\lambda}, \delta_{\mathrm{sm}}) - \mathrm{tri}} \longrightarrow (\tilde{\mathfrak{h}} \times_{\mathfrak{h}} \tilde{\mathfrak{h}})_{\widehat{0}}$$

of Theorem 5.3.38, the pullback of Bezrukavnikov's functor gives rise to the restriction of the conjectural functor $\mathfrak{A}_G^{\mathrm{rig}}$ via the functor of Orlik–Strauch:

EXPECTATION 6.2.27. *Let $\varphi_1, \dots, \varphi_d \in L^{\times}$ such that $\varphi_i/\varphi_j \neq 1, q$ for $i \neq j$ and let δ_{sm} be as in 5.3.16. Then the functor*

$$\mathfrak{A}_G^{\mathrm{rig}} \circ \mathcal{F}_B^G(-, \delta_{\mathrm{sm}}) : \mathcal{O}(U(\mathfrak{g}_L))_{(\xi)} \longrightarrow D_{\mathrm{coh}}^b(\mathfrak{X}_d)$$

factors through $\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}$ and is given by the composition

$$(f_{\underline{\lambda},\delta_{\text{sm}}}^* \circ F_{\xi}(-))([F : \mathbf{Q}_p]\rho').$$

REMARK 6.2.28. Here the twist by ρ' makes sense on $\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}$, compare the discussion following Theorem 5.3.38.

Note that using Lemma 6.2.18 and the fact that $f_{\underline{\lambda},\delta_{\text{sm}}}$ is formally smooth with relative dualizing sheaf the line bundle associated to -2ρ this expectation is compatible with the twist by $(-\rho')_{\sigma}$ in (6.2.5) and with the twists in Expectation 6.2.14 and 6.2.23.

6.2.29. *The trianguline Breuil–Mézard conjecture.* We describe a shadow of the compatibility with parabolic induction and of the compatibility with smooth categorical Langlands in Conjecture 6.2.4. This shadow can be regarded as a version of the Breuil–Mézard conjecture (Conjecture 6.1.38) for locally analytic principal series representations that involves Grothendieck groups rather than groups of cycles (compare the discussion in Section 6.1.37). In certain regular cases, a version of this conjecture involving cycles was proven in [BHS17b], see Theorem 6.2.35 below.

As above we write $G = \text{GL}_d(F)$ and $T \subseteq B \subset G$ for the diagonal torus respectively the upper triangular Borel. We keep the assumption that L is large enough. Note that we may regard G as the group of \mathbf{Q}_p -valued points of $H = \text{Res}_{F/\mathbf{Q}_p} \text{GL}_d$ and similarly T and B are the \mathbf{Q}_p -valued points of T_H and T_B the Weil restrictions of the diagonal torus respectively of the upper triangular Borel. We fix a regular Hodge–Tate weight $\underline{\lambda}$ and as usual write ξ for the corresponding shifted character (which is a character of T_H defined over \bar{F}). We regard ξ as an (algebraic) character $z^{\xi} : T \rightarrow L^{\times}$ (compare (5.3.17)) and fix a smooth character $\delta_{\text{sm}} : T \rightarrow L^{\times}$.

Let us write $W \cong \prod_{K \hookrightarrow L} \mathcal{S}_n$ for the absolute Weyl group of H and $W_0 \cong \mathcal{S}_n \subseteq W$ for the relative Weyl group of H . Then for every $w \in W$ and $w' \in W_0$ we obtain well defined characters $z^{w \cdot \xi}$ and $w' \delta_{\text{sm}}$.

We write $K_0(\underline{\lambda}, \delta_{\text{sm}})$ for the free abelian group generated by the irreducible constituents of the locally analytically induced representations $\text{Ind}_B^G(z^{w \cdot \xi} w' \delta_{\text{sm}})$, for $w \in W$ and $w' \in W_0$.

REMARK 6.2.30. The notation K_0 should remind the reader of the fact that this free abelian group is the Grothendieck group of the category of finite length locally analytic representations whose irreducible subquotients occur in the above locally analytic principal series representations.

On the Galois side, recall the definition of the stack $\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}$ defined just before Theorem 5.3.38. We consider the union

$$\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}} \bigcup_{w' \in W_0} \widehat{\mathfrak{X}}_{d,(\underline{\lambda},w' \delta_{\text{sm}})\text{-tri}}$$

and the Grothendieck group $K_0(\text{Coh}(\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}))$ of coherent sheaves on this stack. Note that $\widehat{\mathfrak{X}}_{d,(\underline{\lambda},\delta_{\text{sm}})\text{-tri}}$ is precisely the set of (φ, Γ) -modules that admit a filtration with graded pieces $\mathcal{R}_F(z^{w \cdot \xi} w' \delta_{\text{sm}})$ for some $w \in W$ and $w' \in W_0$. The smooth character δ_{sm} defines a Bernstein component $\Omega = [T, \delta_{\text{sm}}]$ of the category of smooth representations of G and we write τ for the inertial type such that $\Omega = \Omega_{\tau}$.

As recalled in Section 5.2.9, the component Ω_τ defines connected components

$$\begin{aligned} \mathrm{WD}_{d,F,\tau} &\subset \mathrm{WD}_{d,F}, \\ \mathrm{Fil}_\lambda \mathrm{WD}_{d,F,\tau} &\subset \mathrm{Fil}_\lambda \mathrm{WD}_{d,F} \end{aligned}$$

of the stacks of rank d Weil–Deligne representations, respectively of rank d filtered Weil–Deligne representations of Hodge–Tate weight λ . The connected component $\mathrm{WD}_{d,F,\tau}$ admits a canonical map

$$(6.2.31) \quad \mathrm{WD}_{d,F,\tau} \longrightarrow \mathrm{Spec} \mathcal{Z}_{\Omega_\tau},$$

see (5.2.13), and δ_{sm} defines a point $[\delta_{\mathrm{sm}}] \in \mathrm{Spec} \mathcal{Z}_{\Omega_\tau}$. We write

$$\mathrm{WD}_{d,F}(\delta_{\mathrm{sm}}) \subseteq \mathrm{WD}_{d,F,\tau}$$

for the preimage of $[\delta_{\mathrm{sm}}]$ under (6.2.31).

REMARK 6.2.32. In terms of (smooth) categorical local Langlands this closed substack can be characterized as the union of the supports of the coherent sheaves $\mathfrak{A}_G^{\mathrm{sm}}(\pi)$ attached to the irreducible smooth representations π that occur in $\iota_B^G(\delta_{\mathrm{sm}})$, where ι_B^G denotes normalized parabolic induction. Note that we do not need to take all the possible $\iota_B^G(w'\delta_{\mathrm{sm}})$ for $w' \in W_0$ into account as $\iota_B^G(\delta)_{\mathrm{sm}}$ and $\iota_B^G(w'\delta)_{\mathrm{sm}}$ have the same Jordan–Hölder factors.

LEMMA 6.2.33. *The isomorphism*

$$\mathrm{Fil}_\lambda \mathrm{WD}_{d,F,\tau} \cong \mathrm{Fil}_\lambda \mathrm{Mod}_{d,\varphi,N,\tau}$$

of Lemma 5.2.3 induces a closed embedding

$$\mathrm{Fil}_\lambda \mathrm{WD}_{d,F,\tau}(\delta_{\mathrm{sm}})^{\mathrm{an}} \subset \mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}}.$$

PROOF. This is straightforward from the fact that the right hand side is the stack of all (φ, Γ) -modules that are trianguline with parameters $z^{w\lambda} w'\delta_{\mathrm{sm}}$. \square

The map in the following conjecture should be thought of as the map induced by $\mathfrak{A}_G^{\mathrm{rig}}$ on Grothendieck groups.

CONJECTURE 6.2.34. *There is a unique injective group homomorphism*

$$\mathfrak{a}_{\lambda,\delta_{\mathrm{sm}}} : K_0(\lambda, \delta_{\mathrm{sm}}) \longrightarrow K_0(\mathrm{Coh}(\mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}}))$$

such that

(i) *For $w \in W$ and $w' \in W_0$ one has*

$$\mathfrak{a}_{\lambda,\delta_{\mathrm{sm}}}([\mathrm{Ind}_B^G(z^{ww_0 \cdot \xi} w'\delta_{\mathrm{sm}})^{\mathrm{an}}]) = [\tilde{\beta}_{B,*} \circ \tilde{\alpha}_B^*(\mathfrak{A}_M^{\mathrm{rig}}(\iota(-)([F : \mathbf{Q}_p]\rho')))].$$

(ii) *For an irreducible smooth representation π in the block $\mathrm{Rep}_{\Omega_\tau}^{\mathrm{sm}} G$ that is a constituent of $\iota_B^G \delta_{\mathrm{sm}}$ we have*

$$\mathfrak{a}_{\lambda,\delta_{\mathrm{sm}}}([\pi \otimes L(\xi)]) = [(\mathrm{pr}_{\mathrm{WD}})^*(\mathfrak{A}_G^{\mathrm{sm}}(\pi))(-\rho')_\sigma].$$

Of course the formulas (i) and (ii) should be compared to Properties (1) and (2) in Conjecture 6.2.4. A weaker version of this conjecture is proved in [BHS17b, Theorem 4.3.8] under more restrictive hypothesis. In fact the proof of loc. cit. globalizes from spectra of deformation rings to the stack $\mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}}$. More precisely, let $\mathrm{CH}^0(\mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}})$ denote the free abelian group on the irreducible components of $\mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}}$. Mapping a coherent sheaf to its support (and forgetting components of non-maximal dimension) defines a morphism

$$\mathrm{supp} : K_0(\mathrm{Coh}(\mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}})) \longrightarrow \mathrm{CH}^0(\mathfrak{X}_{d,([\lambda,\delta_{\mathrm{sm}}])\text{-tri}}).$$

Then the methods of proof of [BHS17b, Theorem 4.3.8] can be used to prove the following theorem.

THEOREM 6.2.35. *Let $\delta_{\text{sm}} = (\text{unr}_{\varphi_1}, \dots, \text{unr}_{\varphi_n})$ be an unramified character with $\varphi_i/\varphi_j \notin \{1, q\}$ for $i \neq j$. Then there is a unique homomorphism*

$$\mathfrak{a}'_{\lambda, \delta_{\text{sm}}} : K_0(\lambda, \delta_{\text{sm}}) \longrightarrow \text{CH}^0(\mathfrak{X}_{d, [\lambda, \delta_{\text{sm}}] - \text{tri}})$$

such that the conditions in (i) and (ii) hold after replacing $\mathfrak{a}_{\lambda, \delta_{\text{sm}}}$ by $\mathfrak{a}'_{\lambda, \delta_{\text{sm}}}$ and after composing the right hand side of the equalities with the support map supp .

REMARK 6.2.36.

(a) In fact, in Conjecture 6.2.34, we could ask for more compatibilities, using parabolic induction from more general Levi subgroups than T . However, this would make the formulation of the conjecture a bit more involved.

(b) While Conjecture 6.2.4 is rather hard to attack as there is no candidate for a functor satisfying the constraints, Conjecture 6.2.34 is much more explicit: as in the Breuil–Mézard conjecture, the group homomorphism $\mathfrak{a}_{\lambda, \delta_{\text{sm}}}$ is determined by the conditions (i) and (ii) in the conjecture and the conjecture is rather about checking relations than about the construction of the map (though in this case, contrary to the Breuil–Mézard conjecture, there are only finitely many relations to check). For example, under the additional assumptions in Theorem 6.2.35, Bezrukavnikov’s functor 6.2.26 and Expectation 6.2.27 already determine $\mathfrak{a}_{\lambda, \delta_{\text{sm}}}$ (once one can verify the corresponding relations).

(c) The version of Conjecture 6.2.34 involving Grothendieck groups is much finer than its version involving cycles. For example, in the case $d = 2$, if π is the (parabolically induced) smooth representation that is the non-split extension of the trivial representation $\mathbf{1}$ and the (smooth) Steinberg representation St , then the cycle underlying $\mathfrak{a}_{\lambda, \delta_{\text{sm}}}(L(\xi) \otimes \pi)$ is trivial, whereas the class $\mathfrak{a}_{\lambda, \delta_{\text{sm}}}(L(\xi) \otimes \pi)$ in the Grothendieck group of coherent sheaves is non-trivial (compare e.g. [Hel20, Remark 4.43]). More precisely, the functor $\mathfrak{A}_G^{\text{sm}}$ maps St to the structure sheaf \mathcal{O}_{X_1} (in degree 0) of the irreducible component $X_1 \subset \text{WD}_{2, F, 1}$ (the stack of (φ, N) -modules of rank 2) where the monodromy N is generically non-trivial, whereas $\mathbf{1}$ is mapped to a line bundle \mathcal{L} on X_1 shifted to cohomological degree -1 . The line bundle \mathcal{L} and the structure sheaf \mathcal{O}_{X_1} are non-isomorphic as coherent sheaves on the stack X_1 , but they have of course the same underlying cycle. In the formula for the cycle defined by $\mathfrak{A}_G^{\text{sm}}(\pi)$ these two cycles cancel, as they live in cohomological degrees with distinct parities. On the other hand the class of $\mathfrak{A}_G^{\text{sm}}(\pi)$ in the Grothendieck group of coherent sheaves is still non-trivial.

(d) The construction of $\mathfrak{a}'_{\lambda, \delta_{\text{sm}}}$ lies at the heart of the construction of companion points on eigenvarieties in Section 9.6.29. In a sloppy way this can be compared to the Breuil–Mézard conjecture as follows: proving the Breuil–Mézard conjecture implies that one can (in some nice situations) construct automorphic forms with a prescribed Galois representation. Similarly, the locally analytic Breuil–Mézard multiplicity formula of Theorem 6.2.35 implies that one can construct (overconvergent finite slope) p -adic automorphic forms with prescribed Galois representation.

7. Some known cases of categorical p -adic Langlands

7.1. GL_1 . Unsurprisingly, the GL_1 case of Conjecture 6.1.14 is a consequence of local class field theory. Furthermore in this case we can describe the essential

image of the functor, and we do not need to pass to the derived level, because the functor comes from an equivalence of abelian categories. The key point is that by local class field theory we have an isomorphism $W_F^{\text{ab}} \xrightarrow{\sim} F^\times = \text{GL}_1(F)$, while by [EG22, Rem. 7.2.18], the stack \mathcal{X}_1 admits a description as a moduli stack of 1-dimensional continuous representations of W_F , and thus of representations of $W_F^{\text{ab}} \xrightarrow{\sim} \text{GL}_1(F)$.

More precisely, if we choose a uniformizer $\varpi_F \in F$ then we can write $F^\times = \mathcal{O}_F^\times \times \varpi_F^{\mathbb{Z}}$, and so write $\mathcal{O}[[F^\times]] = \mathcal{O}[[\mathcal{O}_F^\times]] \widehat{\otimes}_{\mathcal{O}} \mathcal{O}[X, X^{-1}]^\wedge$, where $\mathcal{O}[X, X^{-1}]^\wedge$ is the p -adic completion of $\mathcal{O}[X, X^{-1}]$. Then by [EG22, Prop. 7.2.17] (see also [Pha22, Cor. 1.2] for an alternative proof by Dat Pham, which is similar to the approach of Kedlaya–Pottharst–Xiao mentioned below), there is an isomorphism

$$(7.1.1) \quad \left[\left(\text{Spf } \mathcal{O}[[F^\times]] \right) / \widehat{\mathbf{G}}_m \right] \xrightarrow{\sim} \mathcal{X}_1$$

(where $\widehat{\mathbf{G}}_m = \text{Spf } \mathcal{O}[X, X^{-1}]^\wedge$ is the p -adic completion of \mathbf{G}_m , and in the formation of the quotient stack, the $\widehat{\mathbf{G}}_m$ -action is taken to be trivial).

It follows immediately from (7.1.1) that there is an equivalence of categories between the (abelian) category $\mathcal{O}[[F^\times]] - \text{Mod}$ and the category of quasicoherent sheaves on \mathcal{X}_1 whose $\widehat{\mathbf{G}}_m$ -action is trivial. Another way to formulate this equivalence is that the structure sheaf $\mathcal{O}_{\mathcal{X}_1}$ is naturally an $\mathcal{O}[[F^\times]]$ -module, and that our functor is given by

$$(7.1.2) \quad \pi \mapsto \mathcal{O}_{\mathcal{X}_1} \otimes_{\mathcal{O}[[F^\times]]} \pi.$$

In other words in the case $d = 1$, the sheaf L_∞ of Remark 6.1.23 is $\mathcal{O}_{\mathcal{X}_1}$.

There is a parallel description of the stack \mathfrak{X}_1 and an interpretation of coherent sheaves thereon in terms of locally analytic representations of F^\times . By a result of Kedlaya–Pottharst–Xiao [KPX14, Thm. 6.2.14] every rank one (φ, Γ) -module over $\mathcal{R}_{F,A}$, for an affinoid algebra A , is of the form $\mathcal{R}_{F,A}(\delta) \otimes \mathcal{L}$ for a continuous character $\delta : F^\times \rightarrow A^\times$ and a line bundle \mathcal{L} on $\text{Sp } A$. It follows that

$$(7.1.2) \quad \mathfrak{X}_1 = \mathcal{T} / \mathbf{G}_m = \left((\text{Spf } \mathcal{O}[[\mathcal{O}_F^\times]])^{\text{rig}} \times \mathbf{G}_m \right) / \mathbf{G}_m,$$

where $\mathcal{T} = \text{Hom}_{\text{cont}}(F^\times, \mathbf{G}_m(-))$ is the rigid analytic space of continuous characters of F^\times equipped with the trivial action of the rigid analytic multiplicative group \mathbf{G}_m . As above the choice of a uniformizer allows us to write $\mathcal{T} = \mathcal{W} \times \mathbf{G}_m$, where

$$\mathcal{W} = \text{Hom}_{\text{cont}}(\mathcal{O}_F^\times, \mathbf{G}_m(-)) = (\text{Spf } \mathcal{O}[[\mathcal{O}_F^\times]])^{\text{rig}}$$

is the space of continuous characters of \mathcal{O}_F^\times . In particular we see that we have a canonical map

$$\mathcal{X}_{1,\eta}^{\text{rig}} = (\mathcal{W} \times \widehat{\mathbf{G}}_m^{\text{rig}}) / \widehat{\mathbf{G}}_m^{\text{rig}} \longrightarrow (\mathcal{W} \times \mathbf{G}_m) / \mathbf{G}_m = \mathfrak{X}_1.$$

The image of this map is an open substack of \mathfrak{X}_1 , and we note that this map is not representable. Its non-trivial fibers are zero-dimensional and given by the stack quotient $\mathbf{G}_m / \widehat{\mathbf{G}}_m^{\text{rig}}$.

There is an isomorphism

$$\mathcal{D}(\mathcal{O}_F^\times) \cong \Gamma(\mathcal{W}, \mathcal{O}_{\mathcal{W}})$$

of the locally analytic distribution algebra of \mathcal{O}_F^\times with the global sections of the structure sheaf on the space \mathcal{W} of characters of \mathcal{O}_F^\times . We note that $\mathcal{D}(F^\times) =$

$\mathcal{D}(\mathcal{O}_F^\times)[T^{\pm 1}]$ is a Laurent polynomial ring over $\mathcal{D}(\mathcal{O}_F^\times)$, i.e. the ring of functions of the algebraic¹⁵ \mathbf{G}_m over $\mathcal{D}(\mathcal{O}_F^\times)$ whereas

$$\Gamma(\mathfrak{X}_1, \mathcal{O}_{\mathfrak{X}_1}) = \Gamma(\mathcal{T}, \mathcal{O}_{\mathcal{T}}) \cong \Gamma(\mathcal{W}, \mathcal{O}_{\mathcal{W}}) \widehat{\otimes}_L \Gamma(\mathbf{G}_m, \mathcal{O}_{\mathbf{G}_m})$$

is the space of functions on the rigid analytic \mathbf{G}_m over $\mathcal{D}(\mathcal{O}_F^\times)$.

As in (7.1.2) we hence want to define a functor

$$(7.1.3) \quad \pi \mapsto \mathcal{O}_{\mathfrak{X}_1} \widehat{\otimes}_{\mathcal{D}(F^\times)}^L \pi,$$

where $\mathcal{O}_{\mathfrak{X}_1} \widehat{\otimes}_{\mathcal{D}(F^\times)}^L \pi$ should be a derived completed tensor product. Unlike in the case of p -power torsion representations it turns out that the functor (7.1.2) is not defined on the level of abelian categories, and hence the categorical p -adic Langlands correspondence for locally analytic representations is only a derived statement already for GL_1 ! In order to make this more precise it is better to consider the adjoint functor of (7.1.2). To make the discussion easier we consider the case $F = \mathbf{Q}_p$ and even restrict to the case of \mathbf{Z}_p^\times instead of \mathbf{Q}_p^\times . The discussion of this case can again easily be reduced to the case of $(\mathbf{Z}_p, +) \cong (1 + p\mathbf{Z}_p, \cdot)$.

Recall that the distribution algebra $\mathcal{D}(\mathbf{Z}_p)$ is isomorphic to the ring of analytic functions $\Gamma(\mathbf{U}, \mathcal{O}_{\mathbf{U}})$ on the open unit disc \mathbf{U} over \mathbf{Q}_p . Here we identify $(\mathrm{Spf} \mathbf{Z}_p[[\mathbf{Z}_p]])^{\mathrm{rig}}$, the space of continuous characters of \mathbf{Z}_p , with the open unit disc $\mathbf{U} = (\mathrm{Spf} \mathbf{Z}_p[[T]])^{\mathrm{rig}}$ using an identification $\mathbf{Z}_p[[\mathbf{Z}_p]] \cong \mathbf{Z}_p[[T]]$.

The rigid analytic space \mathbf{U} is a smooth Stein space. For coherent sheaves on such spaces Chiarelotto [Chi90] has constructed cohomology with compact support $R\Gamma_c(\mathbf{U}, -)$ as well as a Serre duality

$$(7.1.4) \quad H^0(\mathbf{U}, \mathcal{F})' \cong H_c^1(\mathbf{U}, \mathbf{D}(\mathcal{F})),$$

where $(-)'$ denotes the strong dual (as above) and $\mathbf{D}(\mathcal{F}) = \mathbf{D}_{\mathcal{O}_{\mathbf{U}}}(\mathcal{F})$ denotes the Serre dual of the coherent sheaf \mathcal{F} . By functoriality $H_c^1(\mathbf{U}, \mathbf{D}(\mathcal{F}))$ (or more generally $R\Gamma_c(\mathbf{U}, \mathcal{F}^\bullet)$) is a module over $\mathcal{D}(\mathbf{Z}_p)$ (or a complex of such modules). In fact the induced \mathbf{Z}_p -action on the H_c^1 is locally analytic (and admissible), as by Serre duality, it is the strong dual of a coadmissible module over $\mathcal{D}(\mathbf{Z}_p)$. The functor (7.1.2) now should be defined as the right adjoint functor to the functor

$$\mathcal{F} \mapsto R\Gamma_c(\mathbf{U}, \mathcal{F})$$

from the derived category of coherent sheaves $\mathbf{D}_{\mathrm{coh}}^b(\mathbf{U})$ to a certain derived category of $\mathcal{D}(\mathbf{Z}_p)$ -modules in solid \mathbf{Q}_p -vector spaces (in the sense of Clausen–Scholze). We note that the functor $R\Gamma_c(\mathbf{U}, -)$ does not preserve the cohomological degree of coherent sheaves: if \mathcal{F} is a coherent sheaf concentrated in degree 0 that has finite support, then $R\Gamma_c(\mathbf{U}, \mathcal{F})$ agrees with the global sections $R\Gamma(\mathbf{U}, \mathcal{F})$ and is concentrated in degree 0 (and is a finite dimensional \mathbf{Q}_p vector space). On the other hand $R\Gamma_c(\mathbf{U}, \mathcal{O}_{\mathbf{U}})$ is concentrated in degree 1 where it is given by the space of locally analytic functions $\mathcal{C}^{\mathrm{an}}(\mathbf{Z}_p, \mathbf{Q}_p)$ on \mathbf{Z}_p (as $\mathcal{C}^{\mathrm{an}}(\mathbf{Z}_p, \mathbf{Q}_p)$ is identified with the strong dual of $\mathcal{D}(\mathbf{Z}_p) = R\Gamma(\mathbf{U}, \mathcal{O}_{\mathbf{U}})$).

Conversely the adjoint functor $-\widehat{\otimes}_{\mathcal{D}(\mathbf{Z}_p)}^L \mathcal{O}_{\mathbf{U}}$ takes finite dimensional locally analytic representations to coherent sheaves in degree 0 while it takes $\mathcal{C}^{\mathrm{an}}(\mathbf{Z}_p, \mathbf{Q}_p)$ to $\mathcal{O}_{\mathbf{U}}[1]$, the structure sheaf shifted to degree -1 .

¹⁵Note that we usually write \mathbf{G}_m for the rigid analytic space associated with the scheme \mathbf{G}_m . This should not cause confusion as we will always explicitly mention when \mathbf{G}_m should be seen as a scheme.

REMARK 7.1.5. Using solid condensed modules and analytic rings as introduced by Clausen–Scholze [Sch] (and in the context of rigid analytic spaces studied further by Andreychev [And21]) it is indeed possible to make the above precise. In order not to get too involved with condensed structures we will not spell this out here. However we note that one should compare the above situation to the compactification of the affine line

$$\iota : \mathrm{Spa}(\mathbf{Z}[x], \mathbf{Z}[x]) \rightarrow \mathrm{Spa}(\mathbf{Z}[x], \mathbf{Z})$$

in the context of discrete adic spaces, see [Sch, Theorem 8.1]. In this context the derived completion

$$\iota^* = - \otimes_{(\mathbf{Z}[x], \mathbf{Z})_{\blacksquare}}^L (\mathbf{Z}[x], \mathbf{Z}[x])_{\blacksquare} : \mathbf{D}_{\blacksquare}(\mathbf{Z}[x], \mathbf{Z}) \rightarrow \mathbf{D}_{\blacksquare}(\mathbf{Z}[x], \mathbf{Z}[x])$$

has a fully faithful left adjoint $\iota_!$ (the cohomology with compact support) such that $\iota^* \circ \iota_!$ is the identity, see [Sch, Observation 8.11]. Moreover $\iota_!(\mathbf{Z}[x]) = (\mathbf{Z}((x^{-1}))/\mathbf{Z}[x])[-1]$ lives in cohomological degree 1 (which implies that the derived completion ι^* maps $\mathbf{Z}((x^{-1}))/\mathbf{Z}[x]$ to the structure sheaf shifted to degree -1), see [Sch, Observation 8.12].

REMARK 7.1.6. The above remark suggests that a possible definition of the derived category of locally analytic \mathbf{Z}_p -representations could be the essential image of the functor $R\Gamma_c(\mathbf{U}, -)$ from a category of solid quasicohherent $\mathcal{O}_{\mathbf{U}}$ -modules to a category of $\mathcal{D}(\mathbf{Z}_p)$ -modules in the derived category of solid \mathbf{Q}_p -vector spaces.

It may seem surprising that in the Banach case the functor (7.1.2) is (the derived functor of) a functor between abelian categories, while in the locally analytic setup the functor (7.1.3) only makes sense on the derived level. In particular one of these functors cannot shift the cohomological degree, while the other one can¹⁶. Still, these two functors are compatible in the sense of Conjecture 6.2.4, as we now explain.

Consider the pro-system of ind-coherent sheaves

$$\mathcal{F}_m^{(n)} = F_m^{(n)} \otimes_{\mathbf{Z}_p[[T]]} \mathcal{O}_{\mathrm{Spf}\mathbf{Z}_p[[T]]}$$

on $\mathrm{Spf}(\mathbf{Z}/p^n[[T]]) = \mathrm{colim}_m \mathrm{Spec}((\mathbf{Z}/p^n)[T]/T^{p^m})$, where $F_m^{(n)} = (\mathbf{Z}/p^n)[\mathbf{Z}_p/p^m\mathbf{Z}_p]$ denotes the representation of \mathbf{Z}_p on the space of \mathbf{Z}/p^n -valued functions on \mathbf{Z}_p that are constant on $p^m\mathbf{Z}_p$ -cosets (here we of course use the isomorphism $\mathbf{Z}_p[[\mathbf{Z}_p]] \cong \mathbf{Z}_p[[T]]$). On the level of representations this pro-ind system gives rise (after taking the limits $\lim_{\leftarrow, n} \lim_{\rightarrow, m}$) to the representation of \mathbf{Z}_p on the space of continuous maps $\mathcal{C}^{\mathrm{cont}}(\mathbf{Z}_p, \mathbf{Z}_p)$ and hence, after inverting p and passing to locally analytic vectors we obtain $\mathcal{C}^{\mathrm{an}}(\mathbf{Z}_p, \mathbf{Q}_p)$. The compatibility of the Banach and the analytic context is hence settled by the following lemma.

LEMMA 7.1.7. *The generic fiber $(\mathcal{F}_m^{(n)})_{\eta}$ of the pro-ind-system $(\mathcal{F}_m^{(n)})$ (concentrated in degree 0) is the structure sheaf on the open unit disc $\mathcal{O}_{\mathbf{U}}[1]$ shifted to degree -1 .*

¹⁶One heuristic explanation for this phenomenon is the following: in the torsion case we consider smooth representations, and hence F^{\times} is regarded as a locally pro-finite group, in particular it should be considered as a 0-dimensional object. In the analytic context F^{\times} is regarded as a p -adic Lie group which has positive dimension.

Before sketching the proof, we have to define what we mean by the generic fiber $(\mathcal{F}_m^{(n)})_\eta$ of a pro-ind system $(\mathcal{F}_m^{(n)})$. Recall that the open unit disc has an admissible cover

$$\mathbf{U} = \bigcup_{k \geq 1} U_k$$

by the closed discs U_k of radius $p^{-1/k}$ and each of the closed discs has a nice affine formal model \mathcal{U}_k . Let $j_k : U_k \hookrightarrow \mathbf{U}$ denote the open embedding and let us write $\pi_k : \mathcal{U}_k \rightarrow \mathrm{Spf} \mathbf{Z}_p[[T]]$ for the canonical map. Then $j_k^*((\mathcal{F}_m^{(n)})_\eta)$ is by definition the generic fiber of the coherent sheaf on \mathcal{U}_k given by

$$\lim_{\leftarrow, n} \lim_{\rightarrow, m} \pi_k^* \mathcal{F}_m^{(n)},$$

and we define $(\mathcal{F}_m^{(n)})_\eta$ by gluing these sheaves on the cover $\mathbf{U} = \bigcup_{k \geq 1} U_k$.

SKETCH OF PROOF OF LEMMA 7.1.7. We sketch the case $k = 1$ in which case $\mathcal{U}_1 = \mathrm{Spf} \mathbf{Z}_p\langle X \rangle$ with $X = p^{-1}T$. For simplicity we only prove that

$$\lim_{\rightarrow, m} \bar{\pi}_1^* \mathcal{F}_m^{(1)}$$

is the structure sheaf on $\bar{\mathcal{U}}_1 = \mathrm{Spec} \mathbf{F}_p[X]$ shifted to degree -1 and leave the general case as an exercise. Here

$$\bar{\pi}_1 : \mathrm{Spec} \mathbf{F}_p[X] \rightarrow \mathrm{Spf} \mathbf{F}_p[[T]]$$

is the mod p fiber of π_1 , i.e. the morphism given by $T \mapsto 0$. One can compute that $\mathcal{F}_m^{(1)}$ is the coherent sheaf corresponding to the module $\mathbf{F}_p[T]/T^{p^m}$ with transition maps induced by $T \mapsto T^p$ (a quick way to see this is to note that $\lim_{\rightarrow, m} \mathcal{F}_m^{(1)}$ is the Pontryagin-dual of $\mathbf{F}_p[[\mathbf{Z}_p]]$ and $\lim_{\rightarrow, m} \mathbf{F}_p[T]/T^{p^m}$ is the Pontryagin dual of $\mathbf{F}_p[[T]]$, but we have fixed an identification $\mathbf{F}_p[[\mathbf{Z}_p]] \cong \mathbf{F}_p[[T]]$).

Computing the derived pullback $\bar{\pi}_1^*$ using the resolution

$$\mathbf{F}_p[[T]][X] \xrightarrow{T} \mathbf{F}_p[[T]][X]$$

of $\mathbf{F}_p[X]$, we find that $\bar{\pi}_1^* \mathcal{F}_m^{(1)}$ is a two-term complex with cohomology $\mathbf{F}_p[X] = T^{p^m-1}(\mathbf{F}_p[T]/T^{p^m})[X]$ in degree -1 and $\mathbf{F}_p[X] = (\mathbf{F}_p[T]/T)[X]$ in degree 0 . However the transition maps $T \mapsto T^p$ induce isomorphisms in degree -1 whereas they induce the zero maps in degree 0 . It follows that $\lim_{\rightarrow, m} \bar{\pi}_1^* \mathcal{F}_m^{(1)}$ is $\mathbf{F}_p[X]$ concentrated in degree -1 as claimed. \square

REMARK 7.1.8. We note however that the functor (7.1.3) cannot be fully faithful on the category $D_{\mathrm{f.p.}}^b(\mathrm{an.} F^\times)$: Roughly the functor involves analytification from the algebraic \mathbf{G}_m to the analytic \mathbf{G}_m . To give an explicit example, let π be the universal unramified representation

$$F^\times \rightarrow F^\times / \mathcal{O}_F^\times \cong \mathbf{Z} \rightarrow \Gamma(\mathbf{G}_m^{\mathrm{alg}}, \mathcal{O}_{\mathbf{G}_m^{\mathrm{alg}}}^\times) = L[T, T^{-1}]^\times$$

mapping 1 to T , where we write $\mathbf{G}_m^{\mathrm{alg}}$ for the scheme \mathbf{G}_m . Then (7.1.3) maps π to the structure sheaf on the quotient of the rigid analytic space $\mathbf{G}_m = \{1\} \times \mathbf{G}_m \subset \mathcal{W} \times \mathbf{G}_m = \mathcal{T}$, where $1 \in \mathcal{W}$ denotes the trivial character of \mathcal{O}_F^\times , by the trivial \mathbf{G}_m -action. It follows that $\mathrm{End}_{F^\times}(\pi, \pi) = L[T^{\pm 1}]$, whereas the endomorphism ring of $\mathcal{O}_{\mathfrak{X}_1} \widehat{\otimes}_{\mathcal{D}(F^\times)}^L \pi$ is the ring of functions on the rigid analytic \mathbf{G}_m . In fact the functor (7.1.3) can be seen as some kind of localization of a representation π on \mathfrak{X}_1 .

We expect that the functor is fully faithful on those representations π such that the canonical map

$$\pi \rightarrow R\Gamma_c(\mathfrak{X}_1, \mathcal{O}_{\mathfrak{X}_1} \widehat{\otimes}_{\mathcal{D}(F^\times)}^L \pi)$$

is an isomorphism, or equivalently those representations π on which the $\mathcal{D}(F^\times)$ -action extends to $\Gamma(\mathcal{T}, \mathcal{O}_{\mathcal{T}})$. This is the case if π has bounded slope, for example if π is an admissible representation.

7.2. The Banach case for $\mathrm{GL}_2(\mathbf{Q}_p)$ — I. The structure of \mathcal{X} . In this section and the following two, we summarise the results of the work in progress [DEGa] of Andrea Dotto with M.E. and T.G., which establishes a version of Conjecture 6.1.14 for $\mathrm{GL}_2(\mathbf{Q}_p)$. We caution that since this is work in progress, there may be imprecisions in what follows, and the sketched proofs may not reflect those in the final version of [DEGa].

For simplicity of exposition, we work throughout these two sections with representations having trivial central character. (The papers [DEGa, DEGb] allow an arbitrary fixed central character, and the arguments are essentially the same, with some notational overhead.) Accordingly we set $G := \mathrm{PGL}_2(\mathbf{Q}_p)$, and $K = \mathrm{PGL}_2(\mathbf{Z}_p)$, so that $\mathrm{sm}. G$ is the abelian category of smooth $\mathrm{GL}_2(\mathbf{Q}_p)$ -representations with trivial central character.

On the geometric side, we work with the moduli stack \mathcal{X} of rank 2 (φ, Γ) -modules for $\mathrm{GL}_2/\mathbf{Q}_p$ with determinant ε^{-1} . As explained following Remark 4.1.6, this is a formal algebraic stack over $\mathrm{Spf} \mathcal{O}$, whose underlying reduced substack $\mathcal{X}_{\mathrm{red}}$ is equidimensional of dimension 1. We make the geometry of \mathcal{X} much more explicit in the discussion below.

The following expected theorem verifies Conjecture 6.1.14 in the case of $\mathrm{GL}_2(\mathbf{Q}_p)$.

EXPECTED THEOREM 7.2.1 (Dotto–Emerton–Gee). *Suppose that $p \geq 5$, and let $G = \mathrm{PGL}_2(\mathbf{Q}_p)$. Then there is a pro-coherent sheaf (concentrated in degree 0) of $\mathcal{O}[[G]]$ -modules L_∞ over \mathcal{X} , such that*

$$\mathfrak{A}(\pi) := L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi$$

defines an \mathcal{O} -linear exact fully faithful functor $\mathfrak{A} : D_{\mathrm{f.p.}}^b(\mathrm{sm}. G) \rightarrow D_{\mathrm{coh}}^b(\mathcal{X})$, which then extends to a cocontinuous fully faithful functor $\mathfrak{A} : \mathrm{Ind} D_{\mathrm{f.p.}}^b(\mathrm{sm}. G) \rightarrow \mathrm{Ind} \mathrm{Coh}(\mathcal{X})$.

REMARK 7.2.2. We also expect to prove the various properties conjectured in Conjecture 6.1.14. In particular, the statement about the support of $\mathfrak{A}\left(\left(c\text{-Ind}_K^G W_\lambda \otimes_{\mathcal{O}} \sigma^{\mathrm{crys}, \circ}(\tau)\right)^{\widehat{p}}\right)$ should follow easily from the construction and the classification of locally algebraic vectors in the existing p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$.

The expected compatibility with duality is discussed in Section 7.3.22.

The following theorem is an ingredient in the proof of Expected Theorem 7.2.1. It is essentially the main result of [DEGb] (together with the construction of the morphism $\mathcal{X}_{\mathrm{red}} \rightarrow X$, which will appear in [DEGa] and is explained below). It is special to the case of $\mathrm{GL}_2(\mathbf{Q}_p)$, in the strong sense that we do not expect it to generalise in any obvious way, even to $\mathrm{GL}_2(\mathbf{Q}_{p^2})$.

THEOREM 7.2.3 (Dotto–Emerton–Gee). *Assume that $p \geq 5$ and $G = \mathrm{PGL}_2(\mathbf{Q}_p)$. There is a reduced scheme X given by an explicit chain of \mathbf{P}^1 s over k with ordinary*

double points, and a morphism $\mathcal{X}_{\text{red}} \rightarrow X$ inducing a bijection on closed $\overline{\mathbf{F}}_p$ -points. For each open subset U of X , there is a certain localization $(\text{sm. } G)_U$ of $\text{sm. } G$ such that $(\text{sm. } G)_X = \text{sm. } G$, and the collection $\{(\text{sm. } G)_U\}$ forms a stack (of abelian categories) over the Zariski site of X .

REMARK 7.2.4. We anticipate that Theorem 7.2.3 can be refined to give much more precise information about the compatibility of localization with the functor \mathfrak{A} , and about the Bernstein centres of the categories $(\text{sm. } G)_U$. Indeed, we expect that there is a formal scheme \widehat{X} with underlying reduced scheme X , which is equipped with a morphism $\mathcal{X} \rightarrow \widehat{X}$, and which simultaneously realises \widehat{X} as a moduli space for \mathcal{X} (in the sense that $\mathcal{X} \rightarrow \widehat{X}$ should be initial for maps from \mathcal{X} to formal algebraic spaces; indeed we expect that the map $\mathcal{X} \rightarrow \widehat{X}$ will satisfy the (formal algebraic analogues) of the conditions of [Alp14, Prop. 7.1.1]), and whose ring of functions over each U is identified with the Bernstein centre of $(\text{sm. } G)_U$. We hope to prove such results in a sequel to [DEGa].

REMARK 7.2.5. We expect that the statements of Expected Theorem 7.2.1 and Theorem 7.2.3 continue to hold without the assumption that $p \geq 5$, and that with some effort the proofs sketched below can be extended to this case. The difficulty when $p = 2$ or 3 is that there are many more special cases to consider (or rather, the existing special cases overlap in more complicated ways); on the Galois side, this is because the mod p cyclotomic character is quadratic (if $p = 3$) and trivial (if $p = 2$), and on the representation theory side, there are no very generic Serre weights in the sense of Definition 7.2.10 below (and indeed no generic Serre weights if $p = 2$).

REMARK 7.2.6. Under mild hypotheses, the expected compatibility with Taylor–Wiles–Kisin patching explained in Remark 6.1.30 is a consequence of the construction of our functor below, and local-global compatibility for p -adic local Langlands for $\text{GL}_2(\mathbf{Q}_p)$ of [Eme11a] (see also [CEG⁺18] for a more direct connection between the patching method for modular curves and p -adic local Langlands for $\text{GL}_2(\mathbf{Q}_p)$, under slightly different hypotheses). See also Expected Theorem 9.4.2.

7.2.7. *Motivation.* One reason that we are able to prove Conjecture 6.1.14 for $\text{GL}_2(\mathbf{Q}_p)$ (but not in any other cases) is that we are able to build upon the results of Colmez [Col10c] and Paškūnas [Paš13]. Roughly speaking, these results can be reinterpreted as proving Conjecture 6.1.14 “pointwise”; that is, after completing \mathcal{X} at a closed $\overline{\mathbf{F}}_p$ -point on the spectral side, and restricting to a block of locally admissible representations on the automorphic side.

There does not seem to be any hope of directly passing from these results to Expected Theorem 7.2.1, and morally our strategy is to carry out constructions over \mathcal{X} which are analogues of some of the constructions of Colmez and Paškūnas, and where possible to check properties of these constructions by passing to completions and using their results. In particular, \mathcal{X}_{red} generically consists of reducible representations, which (for $\text{GL}_2(\mathbf{Q}_p)$, but not more generally) correspond via the mod p local Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$ to (extensions of) principal series representations, and we are able to interpolate various explicit calculations in [Paš13] over the reducible locus.

In order to carry out this strategy we need to have an analogue on the automorphic side of familiar operations on sheaves such as restriction to open subsets

and passage to completion at closed points. In [DEGb] we develop such a localization theory for sm. G , proving (for the most part by purely representation-theoretic arguments, together with some use of the p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$) that it localizes over the projective scheme X of Theorem 7.2.3.

We then generalize Colmez's construction in [Col10c] of the $\mathrm{GL}_2(\mathbf{Q}_p)$ -representation he calls $D^{\natural} \boxtimes \mathbf{P}^1$, to give a direct construction of our candidate for L_{∞} . We therefore have an explicit functor \mathfrak{A} , whose behaviour after completion at closed points can be understood using the results of [Col10c, Paš13].

Our arguments rely on having a good understanding of the geometry of the stack \mathcal{X} , as well as an interpretation of this geometry in Galois representation-theoretic terms, and in the remainder of this section we describe this geometry. In the following section we explain the construction of L_{∞} and sketch the main points in the proof of Expected Theorem 7.2.1.

7.2.8. *The geometry of $\mathcal{X}_{\mathrm{red}}$.* We begin by describing the irreducible components of $\mathcal{X}_{\mathrm{red}}$ and their intersections. These are described by the weight part of Serre's conjecture, and we begin by recalling this.

For any $x \in \overline{\mathbf{F}}_p^{\times}$, we let $\lambda_x : \mathrm{Gal}_{\mathbf{Q}_p} \rightarrow \overline{\mathbf{F}}_p^{\times}$ be the unramified character taking a geometric Frobenius element to x . We write $\omega = \bar{\varepsilon}$ for the mod p cyclotomic character, so that every character $\mathrm{Gal}_{\mathbf{Q}_p} \rightarrow \overline{\mathbf{F}}_p^{\times}$ is of the form $\lambda_x \omega^i$ for some uniquely determined $x \in \overline{\mathbf{F}}_p^{\times}$ and $0 \leq i < p - 1$.

We write $\omega_2 : I_{\mathbf{Q}_p} \rightarrow \overline{\mathbf{F}}_p^{\times}$ for a choice of fundamental character of niveau 2; writing \mathbf{Q}_{p^2} for the quadratic unramified extension of \mathbf{Q}_p , we can extend ω_2 to a character $\mathrm{Gal}_{\mathbf{Q}_{p^2}} \rightarrow \overline{\mathbf{F}}_p^{\times}$ in such a way that $\omega_2^{p+1} = \omega|_{\mathrm{Gal}_{\mathbf{Q}_{p^2}}}$. Then the 2-dimensional absolutely irreducible representations $\mathrm{Gal}_{\mathbf{Q}_p} \rightarrow \mathrm{GL}_2(k)$ are precisely those of the form $\lambda_x \otimes \mathrm{Ind}_{\mathrm{Gal}_{\mathbf{Q}_{p^2}}}^{\mathrm{Gal}_{\mathbf{Q}_p}} \omega_2^n$ for some $x \in \overline{\mathbf{F}}_p^{\times}$ and $n \in \mathbf{Z}/(p^2 - 1)\mathbf{Z}$ with $(p + 1) \nmid n$. In this case x is uniquely determined up to multiplication by -1 , and n is uniquely determined up to multiplication by p . The determinant of this representation is an unramified twist of ω^n .

Serre weights for $\mathrm{GL}_2(\mathbf{F}_p)$ are the (isomorphism classes of) irreducible \mathbf{F}_p -representations of $\mathrm{GL}_2(\mathbf{F}_p)$, which are explicitly given by $\sigma_{a,b} = \det^a \mathrm{Sym}^b \mathbf{F}_p^2$, where $0 \leq a < p - 1$ and $0 \leq b \leq p - 1$. It is sometimes convenient to view a as an element of $\mathbf{Z}/p\mathbf{Z}$, and we will do so without further comment.

Following Serre [Ser87] (and as reinterpreted by Buzzard–Diamond–Jarvis [BDJ10]) one can associate to each representation $\bar{\rho} : \mathrm{Gal}_{\mathbf{Q}_p} \rightarrow \mathrm{GL}_2(\overline{\mathbf{F}}_p)$ a set of Serre weights σ . One way of defining this is to say that $\sigma_{a,b}$ is a Serre weight for $\bar{\rho}$ if and only if $\bar{\rho}$ admits a crystalline lift with Hodge–Tate weights $1 - a, -(a + b)$. This description can be made completely explicit as follows.

If $\bar{\rho} = \lambda_x \otimes \mathrm{Ind}_{\mathrm{Gal}_{\mathbf{Q}_{p^2}}}^{\mathrm{Gal}_{\mathbf{Q}_p}} \omega_2^n$ is irreducible, then $\sigma_{a,b}$ is a Serre weight for $\bar{\rho}$ if and only if $n \equiv (p + 1)a + b - p \pmod{p^2 - 1}$ or $n \equiv (p + 1)a + pb - 1 \pmod{p^2 - 1}$.

If $b \neq 0$, then $\sigma_{a,b}$ is a Serre weight for a reducible representation $\bar{\rho}$ if and only if if and only if

$$(7.2.9) \quad \bar{\rho} \cong \begin{pmatrix} \lambda_x \omega^{a+b} & * \\ 0 & \lambda_y \omega^{a-1} \end{pmatrix}$$

for some $x, y \in \overline{\mathbf{F}}_p^{\times}$. If $b = 0$, then we further demand that $\bar{\rho}$ is finite flat. If $x \neq y$ then this is automatic, while if $x = y$ it is equivalent to asking that the extension

be peu ramifiée (which by definition is equivalent to the corresponding Kummer class being integral).

Note that if $\det \bar{\rho} = \omega^{-1}$ then any Serre weight for $\bar{\rho}$ has trivial central character. We assume that all Serre weights have trivial central character from now on; that is, we only consider Serre weights $\sigma_{a,b}$ with $2a + b \equiv 0 \pmod{p-1}$. In particular, b is always even.

Then the irreducible components of \mathcal{X}_{red} are as follows: for each Serre weight $\sigma = \sigma_{a,b}$ with $b \neq p-1$, there is a unique irreducible component $\mathcal{X}(\sigma)$, whose $\bar{\mathbf{F}}_p$ -points are precisely those $\bar{\rho} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(\bar{\mathbf{F}}_p)$ admitting σ as a Serre weight. Furthermore if $\sigma = \sigma_{a,p-1}$ for $a = 0, (p-1)/2$, then there are two corresponding irreducible components $\mathcal{X}(\sigma)^\pm$, which differ by a twist by the unramified quadratic character λ_{-1} . In this case, the representations which admit $\sigma_{a,p-1}$ as a Serre weight are precisely the union of the $\bar{\mathbf{F}}_p$ -points of the stacks $\mathcal{X}(\sigma_{a,p-1})^+, \mathcal{X}(\sigma_{a,p-1})^-$, and $\mathcal{X}(\sigma_{a,0})$.

We now describe the stacks $\mathcal{X}(\sigma)$ more explicitly.

DEFINITION 7.2.10. A Serre weight $\sigma_{a,b}$ is *generic* if $0 \leq b \leq p-3$ (equivalently, if $b \neq p-1$). It is *very generic* if furthermore $b \neq 0$.

The following result is essentially a consequence of Fontaine–Laffaille theory (see Section 7.6.4 for slightly more explanation).


PROPOSITION 7.2.11. *If $\sigma = \sigma_{a,b}$ is generic, then we have an isomorphism $\mathcal{X}(\sigma) \xrightarrow{\sim} [(\mathbf{A}^2 \setminus \{0\})/\mathbf{G}_m]$, where if we write (t, x) for the coordinates on \mathbf{A}^2 , then the \mathbf{G}_m -action is given by*

$$u \cdot (t, x) = (t, u^2x).$$

The reducible locus is the locus $t \neq 0$, and the split locus is $x = 0$. More precisely, the locus $t \neq 0$ is the universal extension

$$\begin{pmatrix} \lambda_t \omega^{a+b} & * \\ 0 & \lambda_{t^{-1}} \omega^{a-1} \end{pmatrix}.$$

The variable x parameterizes the (unique non-zero) extension class.



$\mathcal{X}(\sigma)$ for generic σ

In the non-generic case, we have the following description of the stacks $\mathcal{X}(\sigma_{a,p-1})^\pm$, which can be proved by constructing universal families of extensions of rank 1 (φ, Γ) -modules.

PROPOSITION 7.2.12. *If $\sigma = \sigma_{a,p-1}$, then we have isomorphisms $[\mathbf{A}^2/\mathbf{G}_m] \xrightarrow{\sim} \mathcal{X}(\sigma)^\pm$, where if we write (x, y) for the coordinates on \mathbf{A}^2 , then the \mathbf{G}_m -action is given by*

$$u \cdot (x, y) = (ux, uy).$$

More precisely, $\mathcal{X}(\sigma)^\pm$ is the universal extension

$$\begin{pmatrix} \lambda_{\pm 1} \omega^{-a} & * \\ 0 & \lambda_{\pm 1} \omega^{-(a+1)} \end{pmatrix}.$$

The variables x, y parameterize the extension class in the 2-dimensional k -vector space $H^1(\text{Gal}_{\mathbf{Q}_p}, \omega)$.

DEFINITION 7.2.13. The *companion* of a generic Serre weight $\sigma = \sigma_{a,b}$ is the Serre weight $\sigma^{\text{co}} := \sigma_{a+b+1, p-3-b}$ (which is again generic). We refer to the unordered pair of $\sigma, \sigma^{\text{co}}$ as a *companion pair*.

(We do not define a companion for a Serre weight of the form $\sigma_{a,p-1}$.)

We say that a companion pair $\sigma, \sigma^{\text{co}}$ is very generic if σ and σ^{co} are very generic they not of the form $\sigma_{a,0}, \sigma_{a+1,p-3}$.

REMARK 7.2.14. The companion pairs $\sigma, \sigma^{\text{co}}$ are precisely the sets of Jordan–Hölder constituents of the reductions modulo p of tame cuspidal types (with trivial central character).

The following decomposition of \mathcal{X}_{red} will be useful below.

DEFINITION 7.2.15. If $\sigma = \sigma_{a,b}$ is a generic Serre weight, then we define a closed substack $\mathcal{X}(\sigma/\sigma^{\text{co}})$ of \mathcal{X}_{red} as follows: if $b \neq 0, p-3$ then we set

$$\mathcal{X}(\sigma/\sigma^{\text{co}}) = \mathcal{X}(\sigma) \cup \mathcal{X}(\sigma^{\text{co}}),$$

if $b = 0$ then we set

$$\mathcal{X}(\sigma/\sigma^{\text{co}}) = \mathcal{X}(\sigma_{a,0}) \cup \mathcal{X}(\sigma_{a,p-1})^+ \cup \mathcal{X}(\sigma_{a,p-1})^- \cup \mathcal{X}(\sigma_{a+1,p-3}),$$

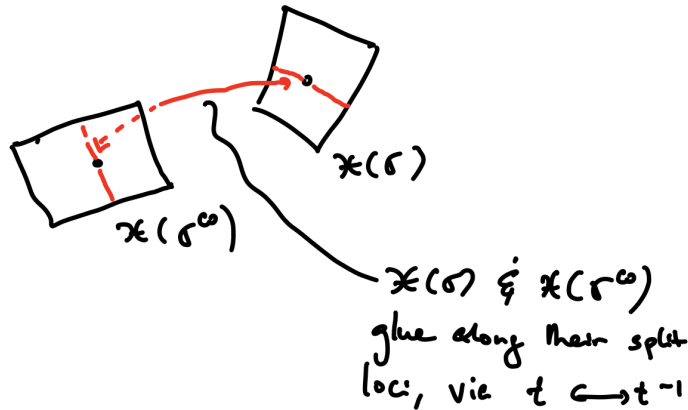
and if $b = p - 3$ then we set

$$\mathcal{X}(\sigma/\sigma^{\text{co}}) = \mathcal{X}(\sigma_{a-1,0}) \cup \mathcal{X}(\sigma_{a-1,p-1})^+ \cup \mathcal{X}(\sigma_{a-1,p-1})^- \cup \mathcal{X}(\sigma_{a,p-3}).$$

Note in particular that $\mathcal{X}(\sigma/\sigma^{\text{co}})$ depends only on the unordered pair $\{\sigma, \sigma^{\text{co}}\}$, and that

$$\mathcal{X}_{\text{red}} = \bigcup_{\sigma/\sigma^{\text{co}}} \mathcal{X}(\sigma/\sigma^{\text{co}}),$$

where $\sigma/\sigma^{\text{co}}$ runs over all companion pairs of Serre weights.



We let X denote a chain of \mathbf{P}^1 's over k with ordinary double points, of length $(p - 1)/2$. We choose a coordinate T on each irreducible component of X in such a way that each singular point corresponds to 0 on one intersecting component and ∞ on the other. We will refer to the points 0 and ∞ as *marked points*.

We label the components of X by pairs $\sigma, \sigma^{\text{co}}$ of companion Serre weights, and denote the corresponding \mathbf{P}^1 by $X(\sigma/\sigma^{\text{co}})$. We order the components by demanding that if $0 < b < p - 1$, then the component $X(\sigma_{a,b}/\sigma_{a+b+1,p-3-b})$ meets the component $X(\sigma_{a+b,p-1-b}/\sigma_{a+1,b-2})$. It is easy to check that there are precisely two such labellings of the components (the choice corresponding to a choice of end of the chain).



$$X(\sigma/\sigma^{\text{co}})$$

The point of this construction is that using our explicit descriptions of the stacks $\mathcal{X}(\sigma)$ we can define a morphism $\pi : \mathcal{X}_{\text{red}} \rightarrow X$. Again we content ourselves with this description in the case of a very generic pair $\sigma, \sigma^{\text{co}}$. We define a morphism $\pi_\sigma : \mathcal{X}(\sigma/\sigma^{\text{co}}) \rightarrow X(\sigma/\sigma^{\text{co}})$ as follows. We identify $\mathcal{X}(\sigma)$ with $[(\mathbf{A}^2 \setminus \{0\})/\mathbf{G}_m]$ as in Proposition 7.2.11, and with the coordinates (t, x) there, we send $(t, x) \mapsto t$. We also identify $\mathcal{X}(\sigma^{\text{co}})$ with $[(\mathbf{A}^2 \setminus \{0\})/\mathbf{G}_m]$, and send $(t, x) \mapsto t^{-1}$. Note that by definition we have $\pi_{\sigma^{\text{co}}}(x, t) = \pi_\sigma(x, t)^{-1}$.

For each companion pair $\sigma, \sigma^{\text{co}}$ of Serre weights, we make a choice of one of $\pi_\sigma : \mathcal{X}(\sigma/\sigma^{\text{co}}) \rightarrow X(\sigma/\sigma^{\text{co}})$ or $\pi_{\sigma^{\text{co}}} : \mathcal{X}(\sigma/\sigma^{\text{co}}) \rightarrow X(\sigma/\sigma^{\text{co}})$, in such a way that the two maps agree at the intersection points of adjacent components of X . (There are two possible such choices of maps, corresponding to the two choices of labelling of the components of X ; one can convert between the two possible labellings by twisting by $\det^{(p-1)/2}$.)

REMARK 7.2.16. If $\sigma/\sigma^{\text{co}}$ is very generic, then

$$\mathcal{X}(\sigma/\sigma^{\text{co}}) \cong (\mathbf{P}^1_t \times [(\text{Spec } k[x, y]/(xy))/(\mathbf{G}_m)_u]) \setminus (0 \times [\{x = 0\}/\mathbf{G}_m] \cup \infty \times [\{y = 0\}/\mathbf{G}_m]),$$

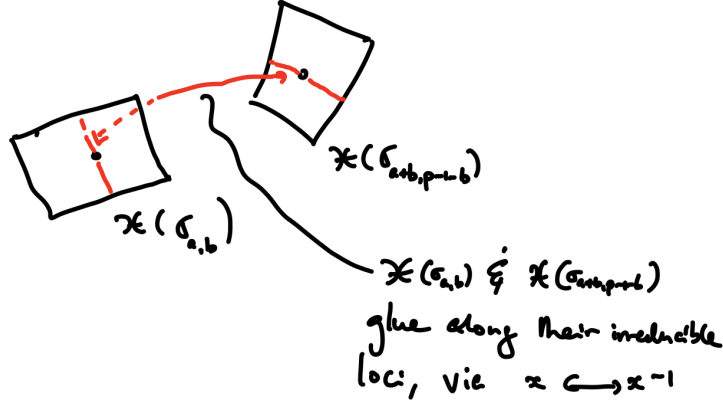
where \mathbf{G}_m acts via $u \cdot (x, y) = (u^2x, u^{-2}y)$. The closed substack $\mathcal{X}(\sigma)$ is cut out by the equation $y = 0$ (with x, t , and u then corresponding to the variables with the same names in the description of $\mathcal{X}(\sigma)$ given by Proposition 7.2.11), while $\mathcal{X}(\sigma^{\text{co}})$ is cut out by $x = 0$ (with y, t^{-1} , and u^{-1} then corresponding to the variables x, t , and u appearing in the description of $\mathcal{X}(\sigma^{\text{co}})$ given by Proposition 7.2.11). In terms of this explicit description, the map $\mathcal{X}(\sigma/\sigma^{\text{co}}) \rightarrow X(\sigma/\sigma^{\text{co}})$ is simply projection to the \mathbf{P}^1 factor.

REMARK 7.2.17. It can be checked that if $\sigma_{a,b}$ is a Serre weight with $b \neq p-1$, then we have an isomorphism

$$\mathcal{X}(\sigma_{a,b}) \cup \mathcal{X}(\sigma_{a+b,p-1-b}) \cong (\text{Spec } k[t_1, t_2]/(t_1t_2) \times [\mathbf{P}^1_x/(\mathbf{G}_m)_u]) \setminus (\{t_1 = 0\} \times [0/\mathbf{G}_m] \cup \{t_2 = 0\} \times [\infty/\mathbf{G}_m]),$$

where \mathbf{G}_m acts on \mathbf{P}^1 via $u \cdot x = u^2x$.

Here we identify the locus $t_2 = 0$ (resp. $t_1 = 0$) with $\mathcal{X}(\sigma_{a,b})$ (resp. $\mathcal{X}(\sigma_{a+b,p-1-b})$) via Proposition 7.2.11, with t_1 corresponding to t and x corresponding to x (resp. t_2 corresponding to t and x corresponding to x^{-1}).



From this, it is easy to see that the above prescriptions for the morphisms π_σ and $\pi_{\sigma^{co}}$ glue to give a morphism $\mathcal{X}_{\text{red}} \rightarrow X$.

REMARK 7.2.18. The motivation for the definition of X and of the morphism $\pi : \mathcal{X}_{\text{red}} \rightarrow X$ is as follows. Either by general principles, or by our explicit descriptions of the stacks $\mathcal{X}(\sigma)$, one easily sees that the *closed* $\overline{\mathbf{F}}_p$ -points of \mathcal{X} are precisely the semisimple $\overline{\rho} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(\overline{\mathbf{F}}_p)$ (with determinant $\overline{\varepsilon}^{-1}$). It follows from the definition of π that it induces a bijection on closed $\overline{\mathbf{F}}_p$ -points, and we think of X as a “coarse” moduli space¹⁷ for \mathcal{X}_{red} . On $\overline{\mathbf{F}}_p$ -points, we can think of the morphism π as sending a representation to its semisimplification; more explicitly, the marked points of the component $X(\sigma/\sigma^{co})$ correspond to irreducible $\overline{\rho}$, and the \mathbf{G}_m obtained by deleting the marked points parameterizes sums of characters with fixed product, and with the Frobenius eigenvalue determined by the \mathbf{G}_m .

REMARK 7.2.19. For any choice of F , it is easy to construct associated moduli space morphisms

$$(7.2.20) \quad (\mathcal{X}_{F,1})_{\text{red}} \rightarrow Y$$

satisfying the conditions of [Alp14, Prop. 7.1.1] Indeed, since $(\mathcal{X}_{F,1})_{\text{red}}$ is a global quotient stack of the form $[\text{Spec } A/\mathbf{G}_m]$, where \mathbf{G}_m acts trivially, one sets $Y = \text{Spec } A$, and the resulting morphism (7.2.20) is even a good moduli space morphism, in the sense of [Alp13].

Furthermore, since $\mathcal{X}_{F,1}$ can be written in the form $[\text{Spf } \widehat{A}/\widehat{\mathbf{G}}_m]$ for some adic affine formal scheme $\text{Spf } \widehat{A}$ which thickens $\text{Spec } A$, endowed with the trivial $\widehat{\mathbf{G}}_m$ -action, if we then set $\widehat{Y} := \text{Spf } \widehat{A}^{\widehat{\mathbf{G}}_m}$, we obtain a morphism $\mathcal{X}_{F,1} \rightarrow \widehat{Y}$ which satisfies the formal algebraic analogue of the conditions of [Alp14, Prop. 7.1.1], and realizes \widehat{Y} as the formal moduli space associated to $\mathcal{X}_{F,1}$ itself.


¹⁷We put “coarse” in quotes because the morphism $\mathcal{X}_{\text{red}} \rightarrow X$ is not a coarse moduli space morphism in the technical sense of e.g. [Alp14], or even an adequate moduli space morphism (for example, because it is not a closed map on underlying topological spaces). It does, however, satisfy the conditions of [Alp14, Prop. 7.1.1], and in particular is initial for morphisms from \mathcal{X}_{red} to locally separated algebraic spaces.

Returning to the case of $\mathrm{GL}_2(\mathbf{Q}_p)$, as already mentioned, we anticipate that the morphism $\mathcal{X}_{\mathrm{red}} \rightarrow X$ constructed above also admits a formal thickening $\mathcal{X} \rightarrow \widehat{X}$ satisfying the formal algebraic analogue of the conditions of [Alp14, Prop. 7.1.1]. On the other hand, we do not expect there to exist interesting moduli space morphisms $\mathcal{X}_{F,d} \rightarrow \widehat{Y}$, or $(\mathcal{X}_{F,d})_{\mathrm{red}} \rightarrow Y$, if $F \neq \mathbf{Q}_p$ and $d \geq 2$, or even if $F = \mathbf{Q}_p$ when $d > 2$. More precisely, in these cases we expect that any morphism $(\mathcal{X}_{F,d})_{\mathrm{red}} \rightarrow Y$, with Y a locally separated algebraic space, maps each connected component of its domain to a single point.

7.2.21. *Some open substacks.* Fix a very generic companion pair of Serre weights $\sigma, \sigma^{\mathrm{co}}$, and write $U(\sigma/\sigma^{\mathrm{co}})$ to denote the open subset of $|X(\sigma/\sigma^{\mathrm{co}})|$ obtained by deleting the marked points $0, \infty$. Then $U(\sigma/\sigma^{\mathrm{co}})$ is also open in $|X|$, and we write $\mathcal{U}(\sigma/\sigma^{\mathrm{co}})$, or sometimes just \mathcal{U} , to denote the preimage of $U(\sigma/\sigma^{\mathrm{co}})$, thought of as an open substack of \mathcal{X} . Then $\mathcal{U}_{\mathrm{red}} = \mathcal{U}(\sigma/\sigma^{\mathrm{co}})_{\mathrm{red}}$ is the (dense) open substack of $\mathcal{X}(\sigma/\sigma^{\mathrm{co}})$ obtained by deleting the locus (consisting of two closed points) corresponding to irreducible $\bar{\rho}$. More precisely, $\mathcal{U}_{\mathrm{red}}$ is the union of its two closed substacks $\mathcal{U} \cap \mathcal{X}(\sigma)$ and $\mathcal{U} \cap \mathcal{X}(\sigma^{\mathrm{co}})$.

In terms of the description of $\mathcal{X}(\sigma/\sigma^{\mathrm{co}})$ given in Remark 7.2.16, $\mathcal{U}(\sigma/\sigma^{\mathrm{co}})_{\mathrm{red}}$ is given by $[(\mathbf{G}_m \times \mathrm{Spec} k[x, y]/(xy))/\mathbf{G}_m] = [\mathrm{Spec} k[t^{\pm 1}, x, y]/(xy)/\mathbf{G}_m]$ with \mathbf{G}_m -action given by

$$u \cdot (x, y) = (u^2x, u^{-2}y).$$




By a deformation theory argument building on this description of $\mathcal{U}_{\mathrm{red}}$, we can explicitly describe \mathcal{U} as

$$(7.2.22) \quad \mathcal{U}(\sigma/\sigma^{\mathrm{co}}) = [\mathrm{Spf} \mathcal{O}[s, t^{\pm 1}, x, y]^{\wedge} / \widehat{\mathbf{G}}_m],$$

where the hat denotes (p, s, xy) -adic completion, and the $\widehat{\mathbf{G}}_m$ -action is given by

$$u \cdot (s, t, x, y) = (s, t, u^2x, u^{-2}y).$$

Here $\mathcal{U} \cap \mathcal{X}(\sigma)$ is the locus $\varpi = s = y = 0$, and $\mathcal{U} \cap \mathcal{X}(\sigma^{\mathrm{co}})$ is given by $\varpi = s = x = 0$ (with the identification with the description in Proposition 7.2.11 being given by $y \mapsto x^{-1}, t \mapsto t^{-1}$).



$\mathcal{U}(\sigma/\sigma^\omega)_{\text{red}}$ sits inside the formally smooth stack $\mathcal{U}(\sigma/\sigma^\omega)$

7.2.23. *More open substacks.* As in Remark 7.2.17, consider a Serre weight $\sigma_{a,b}$ with $b \neq p - 1$. The union $\mathcal{X}(\sigma_{a,b}) \cup \mathcal{X}(\sigma_{a+b,p-1-b})$ then contains a unique closed point whose associated Galois representation $\bar{\rho}$ is irreducible. If we delete the split locus from this union (which, in terms of the explicit description of Remark 7.2.17, means deleting the loci $x = 0$ and $x = \infty$), we obtain an open subset \mathcal{V}_{red} of \mathcal{X}_{red} containing $\bar{\rho}$, which admits the explicit description $\mathcal{V}_{\text{red}} = [\text{Spec } k[t_1, t_2]/(t_1 t_2)/\mu_2]$, where the μ_2 -action is trivial.

If we let \mathcal{V} denote the open substack of \mathcal{X} corresponding to the open subset \mathcal{V}_{red} of \mathcal{X}_{red} , then a deformation theory argument shows that \mathcal{V} admits the description

$$\mathcal{V} \cong [\text{Spf } \mathcal{O}[a, t_1, t_2]^\wedge / \mu_2],$$

where the hat indicates $(p, a, t_1 t_2)$ -adic completion, and where the μ_2 -action is again trivial.

7.3. The Banach case for $\text{GL}_2(\mathbf{Q}_p)$ — II. Proof sketches. In this section we outline the proof of Expected Theorem 7.2.1.

7.3.1. *The construction of L_∞ .* We begin by defining our pro-coherent sheaf L_∞ of $\text{GL}_2(\mathbf{Q}_p)$ -representations. In essence this is a natural generalization of Colmez’s $\text{GL}_2(\mathbf{Q}_p)$ -representation $D^\natural \boxtimes \mathbf{P}^1$. Some of the arguments of [Col10c] go through essentially unchanged in our setting, but others need more serious adaptation. (The main difficulties that arise are that Colmez’s D^\natural and D^{ur} are not compatible with flat base change, so that all arguments involving them have to be reworked; and that, when working with (φ, Γ) -modules with general coefficients, quotient of lattices need not be literally finite modules, so that arguments invoking such finiteness statements must also be reworked.)

To begin with, we note that by definition there is a universal rank 2 (φ, Γ) -module D over \mathcal{X} . We find it convenient to introduce a twist into its definition, in the following way. If A is a finite type \mathcal{O}/ϖ^a -algebra, then a morphism $\text{Spec } A \rightarrow \mathcal{X}$ corresponds to a rank 2 (φ, Γ) -module with A -coefficients and determinant the

inverse cyclotomic character, and the formation of this (φ, Γ) -module is compatible with base change in A . We write D_A for the twist of this (φ, Γ) -module by the cyclotomic character, so that D_A has determinant equal to the cyclotomic character (rather than its inverse).

We work with (φ, Γ) -modules for the full cyclotomic extension $\mathbf{Q}_p(\zeta_{p^\infty})/\mathbf{Q}_p$, so that D_A is a rank 2 projective module over the ring

$$\mathbf{A}_A = A((T)),$$

equipped with commuting (A -linear) semilinear actions of φ and Γ , where

$$\begin{aligned}\varphi(1+T) &= (1+T)^p, \\ \gamma(1+T) &= (1+T)^{\varepsilon(\gamma)}\end{aligned}$$

for any $\gamma \in \Gamma$ (where ε is as usual the p -adic cyclotomic character). There is also a left inverse $\psi : D_A \rightarrow D_A$ to φ ; this is an A -linear surjection which commutes with Γ , and satisfies

$$\begin{aligned}\psi(\varphi(a)m) &= a\psi(m), \\ \psi(a\varphi(m)) &= \psi(a)m\end{aligned}$$

for any $a \in \mathbf{A}_A$, $m \in D_A$.

Following Colmez, we use the actions of φ , ψ and Γ on D_A to explicitly build a $\mathrm{GL}_2(\mathbf{Q}_p)$ -representation. Write P^+ for the monoid $\begin{pmatrix} \mathbf{Z}_p \setminus \{0\} & \mathbf{Z}_p \\ 0 & 1 \end{pmatrix}$. Then there is an action of P^+ on D_A given by the continuous maps

$$\begin{pmatrix} p^k a & b \\ 0 & 1 \end{pmatrix} z = (1+T)^b \varphi^k \circ \sigma_a(z)$$

for $a \in \mathbf{Z}_p^\times$, $b \in \mathbf{Z}_p$, $k \in \mathbf{Z}_{\geq 0}$, where we write $\sigma_a \in \Gamma$ for the element with $\varepsilon(\sigma_a) = a$.

There is a natural action of P^+ on \mathbf{Z}_p (respectively of P on \mathbf{Q}_p) via $\begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} x = ax + b$, and as explained in [Col10b, §3.1], this action can be used to define a P^+ -equivariant sheaf of A -modules $U \mapsto D_A \boxtimes U$ on \mathbf{Z}_p . This is arranged by letting $D_A \boxtimes \mathbf{Z}_p$ be D_A itself, and for any $i \in \mathbf{Z}_p$ and $k \geq 0$, the sections of $D_A \boxtimes (i + p^k \mathbf{Z}_p)$ on $i + p^k \mathbf{Z}_p$ are defined to be $\begin{pmatrix} p^k & i \\ 0 & 1 \end{pmatrix} D_A \subseteq D_A$. The restriction map $\mathrm{Res}_{i+p^k \mathbf{Z}_p} : D_A \boxtimes \mathbf{Z}_p \rightarrow D_A \boxtimes (i + p^k \mathbf{Z}_p)$ is given by

$$\begin{pmatrix} 1 & i \\ 0 & 1 \end{pmatrix} \circ \varphi^k \circ \psi^k \circ \begin{pmatrix} 1 & -i \\ 0 & 1 \end{pmatrix}.$$

We thus have in particular $D_A \boxtimes \mathbf{Z}_p^\times = D_A^{\psi=0}$.

Given an affine linear map $f : U \rightarrow V$ between compact open subsets of \mathbf{Z}_p we have an induced map $f_* : M \boxtimes U \rightarrow M \boxtimes V$. This definition extends immediately to the case of continuous maps which are locally affine linear, and by taking limits, to functions which can be uniformly approximated by affine linear maps. In particular this applies to the map $z \rightarrow 1/z$ on \mathbf{Z}_p^\times , and we can then glue $D_A \boxtimes \mathbf{Z}_p := D_A$ to itself over $D_A \boxtimes \mathbf{Z}_p^\times$ via this map, giving a sheaf over $\mathbf{P}^1(\mathbf{Z}_p) (= \mathbf{P}^1(\mathbf{Q}_p) = \mathbf{Q}_p \cup \{\infty\})$. We denote the global sections of this sheaf by $D_A \boxtimes \mathbf{P}^1$; by functoriality, they have an action of $\mathrm{GL}_2(\mathbf{Q}_p)$ with trivial central character.

Since \mathbf{Q}_p is an open subset of $\mathbf{P}^1(\mathbf{Z}_p)$, we can also take the sections of $D_A \boxtimes \mathbf{P}^1$ over \mathbf{Q}_p , to obtain a B -representation that is denoted $D_A \boxtimes \mathbf{Q}_p$. Unwinding the

definitions, one obtains an explicit description of $D_A \boxtimes \mathbf{Q}_p$ as the inverse limit $\varprojlim_{\psi} D_A$.

The action of $\mathrm{GL}_2(\mathbf{Q}_p)$ on $D_A \boxtimes \mathbf{P}^1$ is compatible with flat base change in A , but $D_A \boxtimes \mathbf{P}^1$ is not the sheaf L_∞ that we are seeking. This is already apparent in the original setting of [Col10c], where $D_A \boxtimes \mathbf{P}^1$ does not realise the local Langlands correspondence, but rather an extension of the correspondence by its (twisted) dual. (See Section 7.3.22 below for an extension of this phenomenon to our setting.) We will therefore construct (again following Colmez) L_∞ as a subrepresentation of $D_A \boxtimes \mathbf{P}^1$.

In order to do so, we set $\mathbf{A}_A^+ = A[[T]]$, which inherits actions of φ , ψ and Γ from \mathbf{A}_A . It can be shown that D_A contains a minimal ψ -stable \mathbf{A}_A^+ -lattice D_A^{\natural} (it is a lattice in the sense that it is finitely generated over \mathbf{A}_A^+ and spans D_A over \mathbf{A}_A), whose formation is compatible with flat base change. There is a natural restriction map

$$\mathrm{Res}_{\mathbf{Q}_p} : D_A \boxtimes \mathbf{P}^1 \rightarrow D_A \boxtimes \mathbf{Q}_p,$$

and we set

$$D_A^{\natural} \boxtimes \mathbf{P}^1 := \{z \in D_A \boxtimes \mathbf{P}^1, \mathrm{Res}_{\mathbf{Q}_p}(z) \in D_A^{\natural} \boxtimes \mathbf{Q}_p\},$$

where $D_A^{\natural} \boxtimes \mathbf{Q}_p := \varprojlim_{\psi} D_A^{\natural}$. It turns out that $D_A^{\natural} \boxtimes \mathbf{P}^1$ is a $\mathrm{GL}_2(\mathbf{Q}_p)$ -stable lattice in $D_A \boxtimes \mathbf{P}^1$, but this is not at all obvious; indeed it is not at all obvious (even in the original setting of [Col10c]) that $D_A^{\natural} \boxtimes \mathbf{P}^1$ is open in $D_A \boxtimes \mathbf{P}^1$. (Here ‘‘lattice’’ is in the sense of [Dri06]: a lattice in a Tate module over A is an open A -submodule L with the property that if $U \subseteq L$ is open, then L/U is a finite A -module.)

The proof that $D_A^{\natural} \boxtimes \mathbf{P}^1$ is a $\mathrm{GL}_2(\mathbf{Q}_p)$ -stable lattice is complicated; it involves a reduction (via consideration of complete local rings) to the case that A is Artinian, and then as in [Col10c] it is ultimately proved via a further reduction to the explicit description of the p -adic local Langlands correspondence for 2-dimensional crystalline representations of $\mathrm{Gal}_{\mathbf{Q}_p}$ proved by Berger–Breuil [BB10]. Having established this, we have the following useful corollary.

COROLLARY 7.3.2 (Dotto–E.–G.). *Suppose that $p \geq 3$. Assume that A is a finite type \mathcal{O}/ϖ^a -algebra for some $a \geq 1$, and that D_A is a projective étale (φ, Γ) -module of rank 2 with A -coefficients and determinant ε . Then $D_A^{\natural} \boxtimes \mathbf{P}^1$ is a finite $A[[K]]$ -module.*

PROOF. This is deduced from the stability of $D_A^{\natural} \boxtimes \mathbf{P}^1$ under the action of $\mathrm{GL}_2(\mathbf{Q}_p)$, together with the Noetherianness of $A[[K]]$, and the fact that D_A^{\natural} , being a lattice, is a finite \mathbf{A}_A^+ -module. \square

It can then be shown that the formation of $D_A^{\natural} \boxtimes \mathbf{P}^1$ is compatible with flat base change in A , and we use this to construct a pro-coherent sheaf L_∞ over \mathcal{X} , equipped with a continuous $\mathrm{GL}_2(\mathbf{Q}_p)$ -action.

7.3.3. A reinterpretation of some results of Paškūnas. We now recall some of the main results of [Paš13], and reinterpret them in terms of our functor \mathfrak{A} . Let $(\mathrm{sm}. G)^{\mathrm{adm}}$ denote the full subcategory of $\mathrm{sm}. G$ consisting of *locally admissible* representations. (By definition, a representation $\pi \in \mathrm{sm}. G$ is locally admissible if every vector $v \in \pi$ is smooth and generates an admissible representation; in fact, for $\mathrm{GL}_2(\mathbf{Q}_p)$, locally admissible representations are automatically locally finite, so that for every v , the representation generated by v is of finite length.)

Unlike $\text{sm. } G$, the category $(\text{sm. } G)^{\text{l.adm}}$ admits a block decomposition, as we now explain. By definition, a block of $(\text{sm. } G)^{\text{l.adm}}$ is an equivalence class of (isomorphism classes of) irreducible objects under the equivalence relation generated by

$$\pi_1 \sim \pi_2 \text{ if } \text{Ext}_{\text{sm. } G}^1(\pi_1, \pi_2) \neq 0 \text{ or } \text{Ext}_{\text{sm. } G}^1(\pi_2, \pi_1) \neq 0.$$

Strictly speaking, we should have written “irreducible admissible objects”, but in fact the irreducible objects of $\text{sm. } G$ are all contained in $(\text{sm. } G)^{\text{l.adm}}$. Indeed, the absolutely irreducible smooth representations of $\text{GL}_2(\mathbf{Q}_p)$ in characteristic p admitting a central character¹⁸ have been classified by Barthel–Livné and Breuil [BL94, Bre03a], and in particular have been shown to be admissible. The non-absolutely irreducible case is rather easily reduced to the absolutely irreducible case (see e.g. [Ber12, Prop. 1.2] and [Paš13, §5.13]).

The blocks \mathfrak{B} are as follows:

- (1) $\mathfrak{B} = \{\pi\}$ for a supersingular irreducible representation π having trivial central character,
- (2) $\mathfrak{B} = \{\text{Ind}_B^G(\chi \otimes \chi^{-1}), \text{Ind}_B^G(\omega\chi^{-1} \otimes \omega^{-1}\chi)\}$, for some character $\chi : \mathbf{Q}_p^\times \rightarrow k^\times$ such that $\chi^2 \neq 1, \omega^2$,
- (3) $\mathfrak{B} = \{\chi, \chi \otimes \text{St}_G, \text{Ind}_B^G(\omega\chi \otimes \omega^{-1}\chi)\}$ for a quadratic character $\chi : \mathbf{Q}_p^\times \rightarrow k^\times$, and
- (4) blocks that do not contain absolutely irreducible objects.

A key observation for us is that there is a bijection between the blocks of types (1)–(3) and $X(k)$. (For the sake of exposition we ignore the blocks of type (4); they are easily handled by allowing finite extensions of our coefficient field k , and indeed correspond to points of X over extensions of k .) To describe this bijection, recall that for any Serre weight σ , the Hecke algebra

$$\mathcal{H}(\sigma) := \text{End}_G(c\text{-Ind}_K^G \sigma)$$

is equal to $k[T]$ for an explicit Hecke operator T , and the absolutely irreducible representations are all subquotients of the representations

$$(c\text{-Ind}_K^G \sigma)/(T - \lambda)$$

for $\lambda \in k$. Furthermore these representations are (absolutely) irreducible unless $\sigma = \sigma_{a,0}$ or $\sigma = \sigma_{a,p-1}$ for some a and $\lambda = \pm 1$, in which case there are two Jordan–Hölder factors, which are of the form $\chi, \chi \otimes \text{St}_G$ for some χ .

The bijection between the blocks and $X(k)$ relies on regarding the Hecke operator T as a coordinate on an \mathbf{A}^1 in the component $X(\sigma/\sigma^{\text{co}})$ of X . We do this as follows, restricting as above to the generic case for the sake of exposition.

DEFINITION 7.3.4. Let σ be a generic Serre weight. Then we defined above a morphism $\mathcal{X}(\sigma/\sigma^{\text{co}}) \rightarrow X(\sigma/\sigma^{\text{co}}) \subset X$ to be given by one of π_σ or $\pi_{\sigma^{\text{co}}}$, and we fixed a coordinate T on $X(\sigma/\sigma^{\text{co}})$ (a copy of \mathbf{P}^1). In the case that this morphism is given by π_σ , we define $f_\sigma : \text{Spec } \mathcal{H}(\sigma) \rightarrow X$ by sending $T \rightarrow T$; and in the case that this morphism is given by $\pi_{\sigma^{\text{co}}}$, we send $T \mapsto T^{-1}$.

¹⁸Subsequently, Berger showed [Ber12] that irreducible smooth representations of $\text{GL}_2(\mathbf{Q}_p)$ over $\overline{\mathbf{F}}_p$ necessarily admit central characters. In our particular context, we are interested in representations of $\text{PGL}_2(\mathbf{Q}_p)$, which is to say of $\text{GL}_2(\mathbf{Q}_p)$ having trivial central character, and so this issue with central characters does not arise.

Then if x is a k -point of X , the corresponding block is given by

$$\mathfrak{B}_x = \bigcup_{\sigma \text{ s.t. } x \in f_\sigma(\mathbf{A}^1)} \text{JH} \left(c\text{-Ind}_K^G(\sigma) \otimes_{\mathcal{H}(\sigma)} x \right).$$

More precisely, \mathfrak{B}_x is a block of $(\text{sm. } G)_k$, and the map $x \mapsto \mathfrak{B}_x$ (suitably extended to the case of non-generic Serre weights) is a bijection from $X(k)$ to the set of blocks of $\text{sm. } G$ containing absolutely irreducible representations.

If \mathfrak{B} is a block of $(\text{sm. } G)^{\text{l.adm}}$, then we write $(\text{sm. } G)_{\mathfrak{B}}^{\text{l.adm}}$ to denote the full subcategory of $(\text{sm. } G)^{\text{l.adm}}$ whose objects are those representations all of whose irreducible subquotients lie in the given block \mathfrak{B} . We then have the block decomposition

$$(\text{sm. } G)^{\text{l.adm}} \xrightarrow{\sim} \prod_{\mathfrak{B}} (\text{sm. } G)_{\mathfrak{B}}^{\text{l.adm}}.$$

In [Paš13] Paškūnas establishes equivalences of categories between the various categories $(\text{sm. } G)_{\mathfrak{B}}^{\text{l.adm}}$, and categories of modules over certain pseudodeformation rings. We use Paškūnas's results (and their proofs) as an input to the proof of Theorem 7.2.1, but in order to do so it is necessary to reformulate these results (and in some cases slightly extend them) as statements about fully faithful functors to categories of sheaves on stacks.

Let $\bar{\theta} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(k)$ be a semisimple representation (equivalently, a pseudorepresentation), corresponding to a point $x \in X(k)$. Assume (for simplicity of exposition) that the associated block \mathfrak{B}_x lies in one of cases (1) or (2) above; that is, we assume that x is not a point $T = \pm 1$ of $X(\sigma/\sigma^{\text{co}})$ for some $\{\sigma, \sigma^{\text{co}}\}$ of the form $\{\sigma_{a,0}, \sigma_{a+1,p-3}\}$. Write $(\text{sm. } G)_{\bar{\theta}}^{\text{l.adm}} := (\text{sm. } G)_{\mathfrak{B}_x}^{\text{l.adm}}$.

We let $\mathcal{X}_{\bar{\theta}}$ denote Carl Wang-Erickson's stack of Galois representations associated to $\bar{\theta}$. By definition, $\mathcal{X}_{\bar{\theta}}$ is characterised by the following property: if A is an \mathcal{O} -algebra in which p is nilpotent, then $\mathcal{X}^{\text{Gal}}(A)$ is the groupoid of continuous morphisms

$$\rho : \text{Gal}_{\overline{\mathbf{Q}}_p} \rightarrow \text{GL}_2(A)$$

with $\det \rho = \varepsilon^{-1}$ and are such that the pseudorepresentation associated to ρ modulo ϖ is equal to $\bar{\theta}$ (here A has the discrete topology, and $\text{Gal}_{\mathbf{Q}_p}$ its natural profinite topology).

Wang-Erickson showed in [WE18] that $\mathcal{X}_{\bar{\theta}}$ is a formal algebraic stack, and by [EG22, Thm. 6.7.2] and [EG20, Thm. 10.2.2], there is a natural monomorphism $\mathcal{X}_{\bar{\theta}} \hookrightarrow \mathcal{X}$ which induces a closed immersion on underlying topological spaces (with image corresponding to those representations $\bar{\rho} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(\overline{\mathbf{F}}_p)$ with semisimplification isomorphic to $\bar{\theta}$). This morphism is versal at the finite type points of $\mathcal{X}_{\bar{\theta}}$, so we naturally expect that it is some kind of completion of the formal algebraic stack \mathcal{X} at the closed substack $\mathcal{Z}_{\bar{\theta}}$ of \mathcal{X}_{red} whose $\overline{\mathbf{F}}_p$ -points are those $\bar{\rho}$ with semisimplification $\bar{\theta}$.

This expectation can be made precise by using the language of coherent completeness [AHR20, Defn. 2.1] in the following way. Write $\mathcal{I}_{\bar{\theta}}$ for the ideal sheaf of $\mathcal{Z}_{\bar{\theta}}$ in \mathcal{X} , and $\mathcal{X}_{\bar{\theta}}^{[n]}$ for the closed substack of \mathcal{X} cut out by $\mathcal{I}_{\bar{\theta}}^{n+1}$. Then it can be shown that for each n , $\mathcal{X}_{\bar{\theta}}^{[n]}$ is a closed substack of $\mathcal{X}_{\bar{\theta}}$, and that the natural functor (given by restriction)

$$\text{Coh}(\mathcal{X}_{\bar{\theta}}) \rightarrow \varprojlim_n \text{Coh}(\mathcal{X}_{\bar{\theta}}^{[n]})$$

is an equivalence of categories; that is, a coherent sheaf on $\mathcal{X}_{\bar{\theta}}$ is the same thing as a compatible system of coherent sheaves on the infinitesimal neighbourhoods of $\mathcal{Z}_{\bar{\theta}}$ in \mathcal{X} .

Given the above constructions, the following theorem is essentially a matter of unwinding some of the proofs of the main results of [Paš13]. Let $(\text{sm. } G)_{\bar{\theta}}^{\text{fg.adm}}$ denote the full subcategory of $(\text{sm. } G)_{\bar{\theta}}^{\text{L.adm}}$ consisting of finitely generated representations.

THEOREM 7.3.5 (Paškūnas+ε). *Suppose as above that $\bar{\theta} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(k)$ is a pseudorepresentation, corresponding to a block which lies in one of cases (1) or (2). Then the functor*

$$(7.3.6) \quad (L_{\infty})_{|\mathcal{X}_{\bar{\theta}}} \otimes_{\mathcal{O}[[G]]} -$$

is an exact fully faithful functor from $(\text{sm. } G)_{\bar{\theta}}^{\text{fg.adm}}$ to the category of coherent $\mathcal{O}_{\mathcal{X}_{\bar{\theta}}}$ -modules.

PROOF. Let $R_{\bar{\theta}}^{\text{ps}}$ denote the complete Noetherian local \mathcal{O} -algebra parameterizing deformations of $\bar{\theta}$ of fixed determinant ε^{-1} over finite type Artinian local \mathcal{O} -algebras, and let θ^u be the universal deformation of $\bar{\theta}$ over $R_{\bar{\theta}}^{\text{ps}}$. We write $\tilde{R}_{\bar{\theta}}$ for the Cayley–Hamilton algebra defined as follows: let J denote the closure of the two-sided ideal of $R_{\bar{\theta}}^{\text{ps}}[[\text{Gal}_{\mathbf{Q}_p}]]$ generated by the elements $g^2 - \theta^u(g)g + (\varepsilon^{-1})(g)$, for $g \in \text{Gal}_{\mathbf{Q}_p}$, and then set

$$\tilde{R}_{\bar{\theta}} := R_{\bar{\theta}}^{\text{ps}}[[\text{Gal}_{\mathbf{Q}_p}]]/J.$$

As in Section 3, we let V denote Colmez’s Montreal functor. Paškūnas shows in [Paš13] that the ring $\tilde{R}_{\bar{\theta}}$ is free of rank 4 as an $R_{\bar{\theta}}^{\text{ps}}$ -module, and the results of [Paš13] can be interpreted as showing that the functor $\pi \mapsto V(\pi)$ induces an equivalence between $(\text{sm. } G)_{\bar{\theta}}^{\text{fg.adm}}$ and the category of finite length $\tilde{R}_{\bar{\theta}}$ -modules.

There is a universal two-dimensional $\text{Gal}_{\mathbf{Q}_p}$ -representation \mathcal{V} lying over $\mathcal{X}_{\bar{\theta}}$, which can be regarded as a rank 2 locally free sheaf of $\mathcal{O}_{\mathcal{X}_{\bar{\theta}}}$ -modules equipped with an action of $\tilde{R}_{\bar{\theta}}$. Then by using the relationships between $D^{\natural} \boxtimes \mathbf{P}^1$ and the Montreal functor V established in [Col10c], one can show that the functor (7.3.6) is naturally equivalent to the functor

$$\mathcal{V} \otimes_{\tilde{R}_{\bar{\theta}}} V(-),$$

and it is now easy to deduce the result from the theorems proved in [Paš13]. \square

The case of a block of type (3) (i.e. the case that x is a point $T = \pm 1$ of $X(\sigma/\sigma^{\text{co}})$ for $\{\sigma, \sigma^{\text{co}}\}$ of the form $\{\sigma_{a,0}, \sigma_{a+1,p-3}\}$) is more subtle, as in contrast to the other blocks we see derived phenomena. This case has recently been studied by Christian Johansson, James Newton and Carl Wang-Erickson, and we expect that their results can be used to prove the following theorem.

THEOREM 7.3.7 (Johansson–Newton–Wang-Erickson+ε). *Suppose as above that $\bar{\theta} : \text{Gal}_{\mathbf{Q}_p} \rightarrow \text{GL}_2(k)$ is a pseudorepresentation, corresponding to a block $\mathfrak{B}_{\bar{\theta}} = \mathfrak{B}_x$ which lies in case (3). Then the functor*

$$(7.3.8) \quad (L_{\infty})_{|\mathcal{X}_{\bar{\theta}}} \otimes_{\mathcal{O}[[G]]}^L -$$

is a fully faithful functor with amplitude $[-1, 0]$ from $D(\mathfrak{B}_{\bar{\theta}}^{\text{fg}})$ to the bounded derived category of coherent $\mathcal{O}_{\mathcal{X}_{\bar{\theta}}}$ -modules.

7.3.9. *A useful resolution.* Let N be the normalizer of the usual (upper triangular) Iwahori subgroup I , and let δ be the nontrivial quadratic character of N/ZI . If π is any smooth G -representation, then we have a complex of $\mathcal{O}[[G]]$ -modules

$$(7.3.10) \quad 0 \rightarrow c\text{-Ind}_N^G(\delta\pi) \rightarrow c\text{-Ind}_K^G(\pi) \rightarrow \pi \rightarrow 0$$

where the third arrow is the tautological surjection $c\text{-Ind}_K^G \pi \rightarrow \pi$. The second arrow is defined as follows. Write $\Pi = \begin{pmatrix} 0 & 1 \\ p & 0 \end{pmatrix}$. Then the morphism

$$v \mapsto \Pi \otimes \Pi^{-1}v - 1 \otimes v \in \mathcal{O}[G] \otimes_{\mathcal{O}[K]} \pi = c\text{-Ind}_K^G \pi$$

is an N -equivariant morphism $\delta\pi \rightarrow c\text{-Ind}_K^G \pi$ which induces the second arrow $c\text{-Ind}_N^G \delta\pi \rightarrow c\text{-Ind}_K^G \pi$. It is easy to show that the complex (7.3.10) is acyclic.

REMARK 7.3.11. The resolution (7.3.10) is valid for $\text{PGL}_2(F)$ for any F/\mathbf{Q}_p , and can be used to reduce certain assertions about general π to the case $\pi = c\text{-Ind}_K^G \sigma$ for a Serre weight σ . In particular, we can reduce the evaluation of cocontinuous functors (such as our hypothetical functor \mathfrak{A}) on general π to the evaluation on $c\text{-Ind}_K^G \sigma$ in the following way. Since we are assuming that $p > 2$, $c\text{-Ind}_N^G \delta\pi$ is a direct summand of $c\text{-Ind}_I^G \pi = c\text{-Ind}_K^G(c\text{-Ind}_I^K \pi)$, so we can reduce to studying representations of the form $c\text{-Ind}_K^G V$.

Since compact induction is compatible with passing to filtered colimits, we then reduce to the case when V is finitely generated. Since compact induction is exact, we furthermore reduce to the case when π is of the form $c\text{-Ind}_K^G \sigma$ for some Serre weight σ , as claimed.

7.3.12. *The definition of the functor.* We are now in a position to define our functor. Since we have already defined the pro-coherent sheaf L_∞ of $\mathcal{O}[[G]]$ -representations over \mathcal{X} , we would simply like to define this via

$$(7.3.13) \quad \pi \mapsto L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi.$$

However, it is not a priori obvious that this definition is well-behaved, and we proceed in stages. Firstly, if π is of the form $c\text{-Ind}_K^G V$ with V finitely generated, then we note that

$$L_\infty \otimes_{\mathcal{O}[[G]]}^L c\text{-Ind}_K^G V = L_\infty \otimes_{\mathcal{O}[[K]]}^L V,$$

which is easily seen to be a bounded-above complex of coherent $\mathcal{O}_{\mathcal{X}}$ -modules by Corollary 7.3.2. We claim that it is in fact a sheaf, concentrated in degree zero. To see this, it is enough (by descending induction on the cohomological degree, and using the coherence of the cohomology groups) to check after restriction to each substack \mathcal{X}_θ , where the result follows from Theorem 7.3.5.

Passing to colimits, we see that for any V (finitely generated or otherwise), $L_\infty \otimes_{\mathcal{O}[[G]]}^L c\text{-Ind}_K^G V$ is a quasicohherent sheaf in degree zero. By Remark 7.3.11, we deduce that for any smooth representation π , $L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi$ is a complex of quasicohherent sheaves concentrated in degrees $[-1, 0]$, and that if π is finitely presented (equivalently, finitely generated), then this complex has coherent cohomology sheaves. We conclude that $L_\infty \otimes_{\mathcal{O}[[G]]}^L -$ defines a functor with amplitude in $[-1, 0]$ and takes bounded complexes of finitely presented representations to bounded coherent complexes of $\mathcal{O}_{\mathcal{X}}$ -modules.

7.3.14. *Localization and gluing.* If σ is a Serre weight, we heuristically regard $c\text{-Ind}_K^G \sigma$ as “living over” a copy of $\mathbf{A}^1 = \text{Spec } k[T] = \text{Spec } \mathcal{H}(\sigma)$ in X via Definition 7.3.4. This intuition is the basis of the localization theory for $\text{sm. } G$ developed in [DEGb], and we now explain how to make it precise. The results described in this section, and their proofs, are for the most part purely representation-theoretic (with some occasional uses of the classification of blocks, which is proved in [Paš13] using the p -adic local Langlands correspondence), although they are of course motivated by the geometric picture.

If Y is a closed subset of X , we let $(\text{sm. } G)_Y$ denote the subcategory of $\text{sm. } G$ consisting of those representations all of whose irreducible subquotients lie in a block corresponding to a closed point of Y . If U is an open subset of X , then we write

$$(\text{sm. } G)_U := \text{sm. } G / (\text{sm. } G)_Y,$$

where $Y := X \setminus U$. It is not hard to check that the subcategory $(\text{sm. } G)_Y$ is localizing in the usual sense: namely, there is a fully faithful right adjoint

$$(j_U)_* : (\text{sm. } G)_U \rightarrow \text{sm. } G$$

to the canonical quotient functor $(j_U)^* : \text{sm. } G \rightarrow (\text{sm. } G)_U$. Where U is understood, we write j_* for $(j_U)_*$, and j^* for $(j_U)^*$.

The following result, which is [DEGb, Prop. 3.1.7], confirms some basic intuitions about how the functors j_*, j^* should behave.

PROPOSITION 7.3.15. *Let Y be a closed subset of X , and write $U := X \setminus Y$.*

- (1) *If σ is a Serre weight, and if $f_\sigma^{-1}(Y) = V(g)$ (a closed subset of $\text{Spec } \mathcal{H}(\sigma)$) for some $g \in \mathcal{H}(\sigma)$, then the natural map*

$$c\text{-Ind}_K^G \sigma \rightarrow (c\text{-Ind}_K^G \sigma)[1/g]$$

can be identified with the morphism

$$c\text{-Ind}_K^G \sigma \rightarrow j_* j^* c\text{-Ind}_K^G \sigma.$$

- (2) *The functor j_* is exact and commutes with filtered colimits.*

If $\{U_0, \dots, U_n\}$ is any finite open cover of X , then for any object π of $\text{sm. } G$, we obtain a functorial Čech resolution

$$(7.3.16) \quad 0 \rightarrow \pi \rightarrow \prod_i (j_i)_* (j_i)^* \pi \rightarrow \cdots \rightarrow (j_{0, \dots, n})_* (j_{0, \dots, n})^* \pi \rightarrow 0$$

where as usual we write $U_{i \dots k}$ for $U_i \cap \cdots \cap U_k$, we have written $j_{i \dots k}$ for $j_{U_{i \dots k}}$, and the differentials are given by the usual formulas. The following is [DEGb, Prop. 3.2.3].

PROPOSITION 7.3.17. *For any object π of $\text{sm. } G$, the resolution (7.3.16) is acyclic.*

PROOF. Using the resolution (7.3.10), we reduce as in Remark 7.3.11 to the case that $\pi = c\text{-Ind}_K^G \sigma$ for a Serre weight σ . In this case Proposition 7.3.15 allows us to reduce to the standard assertion that the Čech complex associated to a finite open cover of an affine scheme by distinguished opens is acyclic. \square

From Proposition 7.3.17, we deduce that $\text{sm. } G$ is a stack in abelian categories over X , as claimed in Theorem 7.2.3.

We need a gluing statement for categories of representations that corresponds to Beauville–Laszlo gluing of sheaves. Write $(\text{sm. } G)^{\text{fg}}$ for the full subcategory of $\text{sm. } G$ consisting of finitely generated objects. Similarly, for any closed subset Y of X , we write $(\text{sm. } G)_Y^{\text{fg}}$ to denote the full subcategory of $(\text{sm. } G)_Y$ consisting of finitely generated objects. We may form the category of formal pro-objects $\text{pro}((\text{sm. } G)_Y^{\text{fg}})$.

DEFINITION 7.3.18. If π is an object of $(\text{sm. } G)^{\text{fg}}$, and if Y is a closed subset of X , then we write $\widehat{\pi}_Y$ for the object “ $\varprojlim_I \pi'$ ” of $\text{pro}((\text{sm. } G)_Y^{\text{fg}})$, where π' runs over all quotients of π which lie in $(\text{sm. } G)_Y^{\text{fg}}$.

The following is [DEGb, Thm. 3.8.1].

THEOREM 7.3.19. *If Y is finite, then the functor $\pi \mapsto \widehat{\pi}_Y$ is exact, and the canonical functor*

$$(\text{sm. } G)^{\text{fg}} \longrightarrow \text{pro}((\text{sm. } G)_Y^{\text{fg}}) \times_{(\text{pro}((\text{sm. } G)_Y^{\text{fg}})_U)} (\text{sm. } G)_U^{\text{fg}},$$

induced by completion along Y and by localization over U , is an equivalence of categories.

7.3.20. *A description of \mathfrak{A} generically.* Given an open subset U of $|X|$, we have the localized category $(\text{sm. } G)_U$, and an open substack $\mathcal{U} := \pi^{-1}(U)$ of \mathcal{X} . We now take U to be the complement of the finite set Y consisting of points corresponding to blocks of types (1) or (3); in particular we have deleted the marked points. Then U is dense in X , and as we now explain, it is in fact possible to explicitly describe the restriction of \mathfrak{A} to $(\text{sm. } G)_U$, and to show that it is a fully faithful functor to $D(\mathcal{U})$. The rough idea of our argument is to follow the strategy of [Paš13] for blocks of type (2), but to “algebraize in the unramified direction”.

The open set U is a disjoint union of its intersections with the various irreducible components of $|X|$, and we handle each of these intersections in turn. In the present discussion we restrict ourselves to describing the case of $U \cap X(\sigma/\sigma^{\text{co}})$ when $\sigma, \sigma^{\text{co}}$ is a companion pair of very generic Serre weights. This intersection was denoted $U(\sigma/\sigma^{\text{co}})$ in Section 7.2.21, but to simplify notation, we now redefine U as $U := U(\sigma/\sigma^{\text{co}})$. The preimage of U in \mathcal{X} is the open substack $\mathcal{U} := \mathcal{U}(\sigma/\sigma^{\text{co}})$.

Our aim, then, is to describe the composition of \mathfrak{A} with restriction to \mathcal{U} , which turns out to factor through $(\text{sm. } G)_U$. Our description of this composition will be Morita-theoretic, in terms of a projective pro-generator for the category $(\text{sm. } G)_U$, and the simplest way to construct this pro-generator in terms of the sheaf L_∞ that we have used to define \mathfrak{A} . Actually, we find it convenient to use a dual version L_∞^\vee of L_∞ , which can be constructed in a similar way to L_∞ itself (see also Section 7.3.22 below).

We now write

$$\mathcal{U} \setminus \mathcal{X}(\sigma^{\text{co}}) = [(\text{Spf } \mathcal{O}[s, t^{\pm 1}, x, y^{\pm 1}]^\wedge / \widehat{\mathbf{G}}_m) = [\text{Spf } A/\mu_2],$$

where $A = \mathcal{O}[s, t^{\pm 1}, x]^\wedge$ (the hat denoting (p, s, x) -adic completion), and the action of μ_2 is trivial. We can regard $L_\infty^\vee|_{\mathcal{U} \setminus \mathcal{X}(\sigma^{\text{co}})}$ as topological $A[[G]]$ -module P_α with an action of μ_2 . With some work, and a comparison to the results of [Paš13], we can show that P_α is a “pro-projective” object of $\text{pro}((\text{sm. } G)_U)$, meaning that if $\pi \rightarrow \pi'$ is an epimorphism in $(\text{sm. } G)_U$, then $\text{Hom}(P_\alpha, \pi) \rightarrow \text{Hom}(P_\alpha, \pi')$ is surjective. (Because projective limits of surjections of abelian groups need not be surjective, this does not imply that P_α is projective in $\text{pro}((\text{sm. } G)_U)$.)

Reversing the roles of σ and σ^{co} , we obtain another pro-projective object P_β . Together, P_α and P_β generate $(\text{sm. } G)_U$. By computing their completions at closed points using Theorem 7.3.5, we can show that $\mathfrak{A}(P_\alpha)$ and $\mathfrak{A}(P_\beta)$ are certain explicit line bundles on \mathcal{U} . In fact, with an appropriate interpretation of the notation, these are the line bundles corresponding to the graded modules $B(1)$ and $B(-1)$ considered in Example C.4.7, and the full faithfulness of the functor $\mathfrak{A} : (\text{sm. } G)_U \rightarrow D(\mathcal{U})$ is a straightforward consequence of Proposition C.4.9.

7.3.21. *Full faithfulness of the functor.* As at the beginning of Section 7.3.20, let U denote the complement of the finite set Y consisting of points corresponding to blocks of types (1) or (3). Ignoring as usual the issue of blocks of type (3) (which require some additional arguments), the full faithfulness of \mathfrak{A} follows easily from the full faithfulness of $\mathfrak{A} : (\text{sm. } G)_U \rightarrow D(\mathcal{U})$, together with Theorem 7.3.19 and Theorem 7.3.5.

7.3.22. *Compatibility with duality.* For any finite type \mathcal{O}/ϖ^a -algebra A , there is an obvious short exact sequence of $\text{GL}_2(\mathbf{Q}_p)$ -representations

$$(7.3.23) \quad 0 \rightarrow D_A^\natural \boxtimes \mathbf{P}^1 \rightarrow D_A \boxtimes \mathbf{P}^1 \rightarrow (D_A \boxtimes \mathbf{P}^1)/(D_A^\natural \boxtimes \mathbf{P}^1) \rightarrow 0.$$

This short exact sequence plays an important role in [Col10c], and using Colmez's results (via a reduction to the case that A is Artinian), we can prove that there is a perfect duality pairing (ultimately coming from the residue pairing in Tate local duality) between $D_A^\natural \boxtimes \mathbf{P}^1$ and $(D_A \boxtimes \mathbf{P}^1)/(D_A^\natural \boxtimes \mathbf{P}^1)$.

We write $\mathbf{D} : D_{\text{f.p.}}^b(\text{sm. } G) \rightarrow D_{\text{f.p.}}^b(\text{sm. } G)$ for the contravariant functor defined via

$$\mathbf{D}(-) = \text{RHom}(-, \mathcal{O}[[G]])[4].$$

Using the resolution (7.3.10), it is easy to check that \mathbf{D} has amplitude $[0, 1]$. We also write $\mathbf{D}_{\mathcal{X}}$ for the antiequivalence of $D_{\text{coh}}(\mathcal{X})$ given by Grothendieck–Serre duality, i.e. $\text{RHom}_{\mathcal{O}_{\mathcal{X}}}(-, \omega_{\mathcal{X}})$ where $\omega_{\mathcal{X}}$ is the dualizing sheaf of \mathcal{X} , placed in degree 0.

By the construction of L_∞ , we can use (7.3.23) to construct an extension

$$0 \rightarrow L_\infty \rightarrow ? \rightarrow \mathbf{D}_{\mathcal{X}}L_\infty \rightarrow 0.$$

This extension class is an element of

$$\text{Hom}_{\mathcal{O}[[G]]}(\mathbf{D}_{\mathcal{X}}L_\infty, L_\infty[1]) = \text{Ext}_{\mathcal{O}[[G]]}^1(\mathbf{D}_{\mathcal{X}}L_\infty, L_\infty)[1],$$

and composition with this extension class gives us a natural morphism of functors on $D_{\text{f.p.}}^b(\text{sm. } G)$

$$\text{RHom}_{\mathcal{O}[[G]]}(-, \mathbf{D}_{\mathcal{X}}L_\infty) \rightarrow \text{RHom}_{\mathcal{O}[[G]]}(-, L_\infty)[1]$$

which we expect can be shown to be an isomorphism, so that we have the expected compatibility with duality

$$\mathfrak{A} \circ \mathbf{D} \xrightarrow{\sim} (\mathbf{D}_{\mathcal{X}} \circ \mathfrak{A})[3].$$

As discussed in Remark 6.1.34, this implies that the various sheaves $\mathfrak{A}((c\text{-Ind}_K^G W_\lambda \otimes \sigma^{\text{crys}, \circ}(\tau))^\wedge)$ are Cohen–Macaulay. Since this is a property that can be checked by pulling back over versal rings, it also follows more directly from the construction of the functor \mathfrak{A} , its relationship to Paškūnas's results discussed in Section 7.3.3, and the results of [Paš15].

7.4. The Banach case for $\mathrm{GL}_2(\mathbf{Q}_p)$ — III. Examples. We now describe some explicit values of the functor \mathfrak{A} , beginning with its values on compact inductions of Serre weights and on (absolutely) irreducible representations, and then turning to some cases that are related to tamely potentially Barsotti–Tate Galois representations. These latter examples also illustrate the compatibility between the Banach and the analytic case posited in Section 6.2.10.

7.4.1. *Compact inductions of Serre weights.* We consider a Serre weight $\sigma_{a,b}$. Assume firstly that $b \neq p-1$, i.e. that $\sigma_{a,b}$ is generic. Then $\mathfrak{A}(c\text{-Ind}_K^G \sigma_{a,b})$ is a line bundle on $\mathcal{X}(\sigma_{a,b})$. This is to be expected, for example by comparison to the corresponding Taylor–Wiles patched modules, which are free over their support as a consequence of Fontaine–Laffaille theory, but it is less obvious precisely which line bundle on $\mathcal{X}(\sigma_{a,b})$ we obtain. To describe this, we note that by Proposition 7.2.11 (and the triviality of the Picard group of $\mathbf{A}^2 \setminus \{0\}$) we have $\mathrm{Pic}(\mathcal{X}(\sigma_{a,b})) = \mathbf{Z}$, and it can be checked that $\mathfrak{A}(c\text{-Ind}_K^G \sigma_{a,b})$ is the line bundle corresponding to $-1 \in \mathbf{Z}$. (This is in fact forced by the compatibility with duality; the discussion of Remark 6.1.32 shows that $\mathfrak{A}(c\text{-Ind}_K^G \sigma_{a,b})$ must be the unique self-dual line bundle on $\mathcal{X}(\sigma_{a,b})$. See also Section 7.7 for a related computation for $\mathrm{GL}_2(\mathbf{Q}_{p^2})$.)

If $b = p-1$ then $\mathfrak{A}(c\text{-Ind}_K^G \sigma_{a,p-1})$ is a line bundle on the union

$$(7.4.2) \quad \mathcal{X}(\sigma_{a,0}) \cup \mathcal{X}(\sigma_{a,p-1})^+ \cup \mathcal{X}(\sigma_{a,p-1})^-.$$

We anticipate that this line bundle can be described explicitly as follows: By Proposition 7.2.12, the stacks $\mathcal{X}(\sigma_{a,p-1})^\pm$ also have Picard group isomorphic to \mathbf{Z} , and we expect that it can be checked that the restriction of $\mathfrak{A}(c\text{-Ind}_K^G \sigma_{a,p-1})$ to each irreducible component in (7.4.2) corresponds to $-1 \in \mathbf{Z}$.

7.4.3. *Absolutely irreducible representations.* The results of Section 7.4.1 allow us to compute the values of \mathfrak{A} on absolutely irreducible representations. As discussed in Section 7.3.3, these are subquotients of the representations

$$(c\text{-Ind}_K^G \sigma)/(T - \lambda)$$

for $\lambda \in k$; and these representations are in fact absolutely irreducible unless $\sigma = \sigma_{a,0}$ or $\sigma = \sigma_{a,p-1}$ for some a and $\lambda = \pm 1$. Supposing that we are not in this case, we can deduce from the injectivity of $T - \lambda$ on the line bundle $\mathfrak{A}(c\text{-Ind}_K^G \sigma)$ that $\mathfrak{A}((c\text{-Ind}_K^G \sigma)/(T - \lambda))$ is a rank 1 sheaf supported on the closed locus $\pi_\sigma^{-1}(\lambda) \cap \mathcal{X}(\sigma)$. (In the case $\lambda \neq 0$ we describe this rank 1 sheaf more precisely in Section 7.4.4 below.)

Unsurprisingly, the trivial and Steinberg representations exhibit more complicated behaviour, which mirrors the corresponding picture for l -adic representations in characteristic zero (see Remark C.4.5). There are short exact sequences of G -representations

$$0 \rightarrow c\text{-Ind}_K^G \sigma_{0,0} \rightarrow c\text{-Ind}_K^G \sigma_{0,p-1} \rightarrow \mathrm{St} \oplus (\mathrm{nr}_{-1} \circ \det) \otimes \mathrm{St} \rightarrow 0$$

$$0 \rightarrow c\text{-Ind}_K^G \sigma_{0,p-1} \rightarrow c\text{-Ind}_K^G \sigma_{0,0} \rightarrow 1 \oplus \mathrm{nr}_{-1} \circ \det \rightarrow 0$$

which in combination with the descriptions of the $\mathfrak{A}(c\text{-Ind}_K^G \sigma_{a,b})$ above can be used to show that $\mathfrak{A}(\mathrm{St})$ is a line bundle on $\mathcal{X}(\mathrm{Sym}^{p-1})^+$, while $\mathfrak{A}(1)$ is a shift to cohomological degree -1 of another line bundle on $\mathcal{X}(\mathrm{Sym}^{p-1})^+$.

7.4.4. *Type (2) blocks.* We consider the case of a block of type (2) in more detail. In this case $\mathfrak{B} = \{\pi_1, \pi_2\}$, with $\pi_1 = \text{Ind}_B^G(\chi \otimes \chi^{-1})$ and $\pi_2 = \text{Ind}_B^G(\omega\chi^{-1} \otimes \omega^{-1}\chi)$ for some suitably generic characters $\chi : \mathbf{Q}_p^\times \rightarrow k^\times$. In this context there are non-split extensions $0 \rightarrow \pi_2 \rightarrow \kappa \rightarrow \pi_1 \rightarrow 0$ and $0 \rightarrow \pi_1 \rightarrow \kappa' \rightarrow \pi_2 \rightarrow 0$ (each unique up to isomorphism), and we will describe the various sheaves $\mathfrak{A}(\pi_1)$, $\mathfrak{A}(\pi_2)$, $\mathfrak{A}(\kappa)$, and $\mathfrak{A}(\kappa')$.

To begin with, we note that we may write $\pi_1 = (c\text{-Ind}_K^G \sigma)/(T - \lambda)$ and $\pi_2 = (c\text{-Ind}_K^G \sigma^{\text{co}})/(T - \lambda^{-1})$ for some very generic σ and its companion σ^{co} , and some appropriately chosen λ . All the sheaves under consideration will then have support contained in $\pi_\sigma^{-1}(\lambda)$.

The locus $\pi_\sigma^{-1}(\lambda)$ is a closed substack of $\mathcal{U}(\sigma/\sigma^{\text{co}})_{\text{red}}$. Following the notation of (7.2.22), we have

$$\mathcal{U}(\sigma/\sigma^{\text{co}})_{\text{red}} := [\text{Spec}(k[t^{\pm 1}, x, y]/(xy)) / \mathbf{G}_m],$$

with the \mathbf{G}_m -action being given by $u \cdot (x, y) = (u^2x, u^{-2}y)$. The map $\pi_\sigma : \mathcal{U}(\sigma/\sigma^{\text{co}})_{\text{red}} \rightarrow \mathbf{A}^1$ is then simply projection to the t -coordinate, and so

$$\pi_\sigma^{-1}(\lambda) = [(\text{Spec } k[x, y]/(xy)) / \mathbf{G}_m]$$

(the locus where $t = \lambda$).

We may give $k[x, y]/(xy)$ its grading induced by the \mathbf{G}_m -action (so x has degree 2 and y has degree -2), and may then describe coherent sheaves on $\pi_\sigma^{-1}(\lambda)$ as graded $k[x, y]/(xy)$ -modules. We then have the following descriptions of the sheaves under consideration: $\mathfrak{A}(\pi_1) = (k[x, y]/(y))(-1)$; $\mathfrak{A}(\pi_2) = (k[x, y]/(x))(1)$; $\mathfrak{A}(\kappa) = (k[x, y]/(xy))(-1)$; $\mathfrak{A}(\kappa') = (k[x, y]/(xy))(1)$. (Here the twists indicate shifts in grading, as usual.) Note the evident short exact sequences

$$0 \rightarrow (k[x, y]/(x))(1) \xrightarrow{y} (k[x, y]/(xy))(-1) \rightarrow (k[x, y]/(y))(-1) \rightarrow 0$$

and

$$0 \rightarrow (k[x, y]/(y))(-1) \xrightarrow{x} (k[x, y]/(xy))(1) \rightarrow (k[x, y]/(x))(1) \rightarrow 0;$$

these can be interpreted as arising from applying \mathfrak{A} to the short exact sequences that defined κ and κ' .

7.4.5. *Tame principal series types.* In this example and the next, we treat some p -adically completed compactly induced representations that are related to tamely potentially Barsotti–Tate representations. We content ourselves with merely sketching the various computations required to verify all our claims.

We begin with the tame principal series case, and to this end, let $\sigma/\sigma^{\text{co}}$ be very generic, and let τ be a tame principal series type such that σ is a Jordan–Hölder factor of $\overline{\sigma}^{\text{crys}, \circ}(\tau)$. Note that this implies that σ^{co} is *not* a Jordan–Hölder factor of $\overline{\sigma}^{\text{crys}, \circ}(\tau)$; indeed, if $\sigma = \sigma_{a,b}$, then the Jordan–Hölder factors of $\overline{\sigma}^{\text{crys}, \circ}(\tau)$ are σ itself, together with $\sigma_{a+b, p-1-b}$. (Explicitly, the type $\overline{\sigma}^{\text{crys}, \circ}(\tau)$ is inflated from a principal series representation of $\text{GL}_2(\mathbf{F}_p)$, as in Section 7.7.7.)

We may then consider the stack $\mathcal{X}^{\text{crys}, (1,0), \tau}$ of potentially Barsotti–Tate representations of type τ . We can describe this closed substack of \mathcal{X} quite explicitly, and we will proceed to do so. The first point to note is that the mod ϖ fibre $\mathcal{X}^{\text{crys}, (1,0), \tau}$ is equal to $\mathcal{X}(\sigma_{a,b}) \cup \mathcal{X}(\sigma_{a+b, p-1-b})$. (This is a precise form of the geometric Breuil–Mézard conjecture discussed in Section 6.1.37 for the particular type τ .) This union admits the explicit description given in Remark 7.2.17, which embeds it as an open substack of $\text{Spec } k[t_1, t_2]/(t_1 t_2) \times [\mathbf{P}^1 / \mathbf{G}_m]$. This latter stack

may be identified with the underlying reduced substack of the p -adic formal algebraic stack $(\mathrm{Spf}(\mathcal{O}[t_1, t_2]/(t_1 t_2 - p)))^\wedge \times [\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m]$, and $\mathcal{X}^{\mathrm{crys}, (1,0), \tau}$ is isomorphic to the open substack of this formal algebraic stack having $\mathcal{X}(\sigma_{a,b}) \cup \mathcal{X}(\sigma_{a+b, p-1-b})$ as its underlying reduced substack.

We now turn to describing the sheaf $\mathfrak{A}((c\text{-Ind}_K^G \sigma^{\mathrm{crys}, \circ}(\tau))^\wedge)$. Of course, this depends on the particular choice of lattice $\sigma^{\mathrm{crys}, \circ}(\tau)$. Among all such, there are two particular lattices Λ_1 and Λ_2 , determined uniquely (up to scaling) by the condition that the cosocle of Λ_1 (resp. Λ_2) is equal to σ (resp. $\sigma_{a+b, p-1-b}$). We may and do normalize these lattices so that $p\Lambda_1 \subset \Lambda_2 \subset \Lambda_1$. Then (up to scaling) any lattice Λ is of the form $\varpi^a \Lambda_1 + \Lambda_2$, with $0 \leq a \leq e$ (where e denotes the absolute ramification degree of our coefficient field L). (This description of lattices in tame principal series types for GL_2 is due to Breuil [Bre14, §2], and is a special case of [EGS15, Prop. 4.1.4].)

If $\chi : \mathbf{F}_p^\times \times \mathbf{F}_p^\times \rightarrow \mathbf{Z}_p^\times \subseteq \mathcal{O}^\times$ denotes the tame lift of the highest weight of σ , which we inflate to a character of the (usual upper triangular) Iwahori I , then $\Lambda_1 = \mathrm{Ind}_I^K \chi$, while if χ^s denotes the conjugate of χ by the non-trivial Weyl group element s , then $\Lambda_2 = \mathrm{Ind}_I^K \chi^s$. Conjugation by the element $\begin{pmatrix} 0 & 1 \\ p & 0 \end{pmatrix}$ in the normalizer (in G) of I interchanges χ and χ^s , and so there is an isomorphism

$$c\text{-Ind}_K^G \Lambda_1 = c\text{-Ind}_K^G \chi \cong c\text{-Ind}_I^G \chi^s = c\text{-Ind}_K^G \Lambda_2.$$

Following this isomorphism with the inclusion $c\text{-Ind}_K^G \Lambda_2 \hookrightarrow c\text{-Ind}_K^G \Lambda_1$ induced by the inclusion of Λ_2 in Λ_1 gives rise to an endomorphism U_p of $c\text{-Ind}_K^G \Lambda_1$. Interchanging the roles of Λ_1 and Λ_2 gives rise to another endomorphism U'_p of $c\text{-Ind}_K^G \Lambda_1$, with the property that $U_p U'_p = p$; and indeed

$$\mathrm{End}_G(c\text{-Ind}_K^G \Lambda_1) = \mathcal{O}[U_p, U'_p]/(U_p U'_p - p).$$

Passing to p -adic completions, we find that

$$\mathrm{End}_G((c\text{-Ind}_K^G \Lambda_1)^\wedge) = \mathcal{O}[U_p, U'_p]^\wedge / (U_p U'_p - p).$$

If $\Lambda = \varpi^a \Lambda_1 + \Lambda_2$, then the inclusion of Λ in Λ_1 induces an identification $c\text{-Ind}_K^G \Lambda = (\varpi^a, U_p) c\text{-Ind}_K^G \Lambda_1$, and hence a corresponding identification of p -adic completions $(c\text{-Ind}_K^G \Lambda)^\wedge = (\varpi^a, U_p)(c\text{-Ind}_K^G \Lambda_1)^\wedge$. This reduces the computation of $\mathfrak{A}((c\text{-Ind}_K^G \Lambda)^\wedge)$ for any lattice Λ to the computation of $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_1)^\wedge)$.

Now we choose our coefficient field L to be unramified over \mathbf{Q}_p (indeed the character χ , and hence the lattice Λ_1 , are defined over \mathbf{Z}_p), and with this hypothesis we see that $\mathcal{X}^{\mathrm{crys}, (1,0), \tau}$ is regular. Thus, since (as noted in Section 7.3.22) the sheaf $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_i)^\wedge)$ is Cohen–Macaulay, the Auslander–Buchsbaum formula shows that it is locally free, and in fact it must be invertible (e.g. for the reason sketched in Remark 6.1.34). Since $\Lambda_1^\vee = \Lambda_2$, and since $(c\text{-Ind}_K^G \Lambda_1)^\wedge \cong (c\text{-Ind}_K^G \Lambda_2)^\wedge$, we further see that this line bundle is Grothendieck–Serre self-dual on $\mathcal{X}^{\mathrm{crys}, (1,0), \tau}$.

Now one can check that $\mathrm{Pic}([\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m]) = \mathbf{Z} \times \mathbf{Z}$, with the first copy of \mathbf{Z} coming from $\mathrm{Pic}(\widehat{\mathbf{P}^1}) = \mathbf{Z}$, and the second copy coming from twisting by characters of $\widehat{\mathbf{G}}_m$. We let $\mathcal{O}(m)_{\widehat{\mathbf{P}^1}}(n)$ denote the invertible sheaf corresponding to a pair $(m, n) \in \mathbf{Z} \times \mathbf{Z}$.

Pulling back along the composite

$$f : \mathcal{X}^{\mathrm{crys}, (1,0), \tau} \xrightarrow{\mathrm{open}} (\mathrm{Spf}(\mathcal{O}[t_1, t_2]/(t_1 t_2 - p)))^\wedge \times [\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m] \xrightarrow{\mathrm{proj.}} [\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m]$$

induces an isomorphism

$$\mathrm{Pic}([\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m]) \xrightarrow{\sim} \mathrm{Pic}(\mathcal{X}^{\mathrm{crys},(1,0),\tau}).$$

Furthermore, the dualizing sheaf of $[\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m]$, namely $\mathcal{O}(-2)_{\widehat{\mathbf{P}^1}}$, pulls back to the dualizing sheaf of $\mathcal{X}^{\mathrm{crys},(1,0),\tau}$. Thus (since it is self-dual), we find that

$$\mathfrak{A}((c\text{-Ind}_K^G \Lambda_1)^\wedge) \cong f^* \mathcal{O}(-1)_{\widehat{\mathbf{P}^1}}.$$

Since \mathfrak{A} is fully faithful, we also obtain an isomorphism

$$\mathrm{End}_G((c\text{-Ind}_K^G \Lambda_1)^\wedge) \xrightarrow{\sim} \Gamma(\mathcal{X}^{\mathrm{crys},(1,0),\tau}, \mathcal{O}_{\mathcal{X}^{\mathrm{crys},(1,0),\tau}}) = \mathcal{O}[t_1, t_2]^\wedge / (t_1 t_2 - p),$$

which is given (after choosing the t_i suitably) by mapping U_p to t_1 and U'_p to t_2 . For any lattice $\Lambda = \varpi^a \Lambda_1 + \Lambda_2$, we then find that

$$\mathfrak{A}((c\text{-Ind}_K^G \Lambda)^\wedge) = (\varpi^a, t_1) f^* \mathcal{O}_{\widehat{\mathbf{P}^1}}(-1).$$

In order to compare these values of \mathfrak{A} with those we computed above for the compact inductions of Serre weights, it helps to note first that pullback to the special fibre induces an isomorphism

$$\mathrm{Pic}([\widehat{\mathbf{P}^1}/\widehat{\mathbf{G}}_m]) \xrightarrow{\sim} \mathrm{Pic}([\mathbf{P}^1/\mathbf{G}_m]).$$

Then, if we let x_0 and x_1 denote the homogeneous coordinates on \mathbf{P}^1 (so that the inhomogeneous coordinate x of Remark 7.2.17 is equal to x_1/x_0), and let u denote the coordinate on \mathbf{G}_m , each of $[\mathbf{A}_{x_1/x_0}^1/(\mathbf{G}_m)_u]$ and $[\mathbf{A}_{x_0/x_1}^1/(\mathbf{G}_m)_{u^{-1}}]$ embeds as an open substack of $[\mathbf{P}^1/\mathbf{G}_m]$, and restriction to these substacks induces an embedding

$$\mathrm{Pic}([\mathbf{P}^1/\mathbf{G}_m]) \hookrightarrow \mathrm{Pic}([\mathbf{A}_{x_1/x_0}^1/(\mathbf{G}_m)_u]) \times \mathrm{Pic}([\mathbf{A}_{x_0/x_1}^1/(\mathbf{G}_m)_{u^{-1}}]) = \mathbf{Z} \times \mathbf{Z},$$

given by

$$(m, n) \mapsto (m + n, m - n).$$

From this, we compute the restrictions

$$(f^* \mathcal{O}(-1)_{\widehat{\mathbf{P}^1}})|_{\mathcal{X}(\sigma_{a,b})} = \mathcal{O}_{\mathcal{X}(\sigma_{a,b})}(-1)$$

and

$$(f^* \mathcal{O}(-1)_{\widehat{\mathbf{P}^1}})|_{\mathcal{X}(\sigma_{a+b,p-1-b})} = \mathcal{O}_{\mathcal{X}(\sigma_{a+b,p-1-b})}(-1).$$

Thus, if we apply \mathfrak{A} to the short exact sequence

$$0 \rightarrow (c\text{-Ind}_K^G \Lambda_1)^\wedge \xrightarrow{U_p} (c\text{-Ind}_K^G \Lambda_1)^\wedge \rightarrow c\text{-Ind}_K^G \sigma_{a,b} \rightarrow 0,$$

respectively

$$0 \rightarrow (c\text{-Ind}_K^G \Lambda_1)^\wedge \xrightarrow{U'_p} (c\text{-Ind}_K^G \Lambda_1)^\wedge \rightarrow c\text{-Ind}_K^G \sigma_{a+b,p-1-b} \rightarrow 0,$$

we obtain the short exact sequence of sheaves

$$0 \rightarrow f^* \mathcal{O}_{\widehat{\mathbf{P}^1}}(-1) \xrightarrow{t_1} f^* \mathcal{O}_{\widehat{\mathbf{P}^1}}(-1) \rightarrow \mathcal{O}_{\mathcal{X}(\sigma_{a,b})}(-1) \rightarrow 0,$$

respectively

$$0 \rightarrow f^* \mathcal{O}_{\widehat{\mathbf{P}^1}}(-1) \xrightarrow{t_2} f^* \mathcal{O}_{\widehat{\mathbf{P}^1}}(-1) \rightarrow \mathcal{O}_{\mathcal{X}(\sigma_{a+b,p-1-b})}(-1) \rightarrow 0.$$

7.4.6. *Tame cuspidal types.* One can make a similar analysis to that of Section 7.4.5 in the case of a tame cuspidal type. Namely, there is a tame cuspidal type τ uniquely determined by the condition that the Jordan–Hölder factors of $\sigma^{\text{crys},\circ}(\tau)$ are σ and σ^{co} . The potentially Barsotti–Tate substack $\mathcal{X}^{\text{crys},(1,0),\tau}$ is then a p -adic formal lift of $\mathcal{X}(\sigma/\sigma^{\text{co}}) := \mathcal{X}(\sigma) \cup \mathcal{X}(\sigma^{\text{co}})$. As noted in Remark 7.2.16, $\mathcal{X}(\sigma/\sigma^{\text{co}})$ embeds as an open substack of $\mathbf{P}^1 \times [(\text{Spec } k[x, y]/(xy))/\mathbf{G}_m]$, and a deformation theory argument shows that $\mathcal{X}^{\text{crys},(1,0),\tau}$ is isomorphic to the open formal algebraic substack of the p -adic formal algebraic stack $\widehat{\mathbf{P}}^1 \times [(\text{Spf}(\mathcal{O}[x, y]/(xy - p))^\wedge/\widehat{\mathbf{G}}_m]$ lifting $\mathcal{X}(\sigma/\sigma^{\text{co}})$.

Just as in the principal series case, there are distinguished lattices Λ_1 and Λ_2 in $\sigma^{\text{crys}}(\tau)$, uniquely determined up to scaling, whose cosocles are σ and σ^{co} respectively. Fixing a choice of Λ_1 , we can then scale Λ_2 so that $p\Lambda_1 \subset \Lambda_2 \subset \Lambda_1$, and any lattice is (up to scaling) of the form $\Lambda = \varpi^a \Lambda_1 + \Lambda_2$, with $0 \leq a \leq e$. Thus, to describe the various sheaves $\mathfrak{A}((c\text{-Ind}_K^G \Lambda)^\wedge)$, it suffices to describe the sheaves $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_1)^\wedge)$ and $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_2)^\wedge)$, as well as the morphism

$$(7.4.7) \quad \mathfrak{A}((c\text{-Ind}_K^G \Lambda_2)^\wedge) \rightarrow \mathfrak{A}((c\text{-Ind}_K^G \Lambda_1)^\wedge)$$

induced by the inclusion $\Lambda_2 \subset \Lambda_1$.

As in the principal series we choose our coefficient field L to be unramified over \mathbf{Q}_p (tame types may always be defined over unramified extensions), and with this hypothesis we see that $\mathcal{X}^{\text{crys},(1,0),\tau}$ is again regular, so that the sheaves $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_i)^\wedge)$ are again invertible. Further, since $\Lambda_1^\vee = \Lambda_2$, we find that $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_1)^\wedge)$ and $\mathfrak{A}((c\text{-Ind}_K^G \Lambda_2)^\wedge)$ must be interchanged under Grothendieck–Serre duality.

Now pullback induces an isomorphism

$$\text{Pic}(\widehat{\mathbf{P}}^1 \times [(\text{Spf}(\mathcal{O}[x, y]/(xy - p))^\wedge/\widehat{\mathbf{G}}_m]) \xrightarrow{\sim} \text{Pic}(\mathcal{X}^{\text{crys},(1,0),\tau}),$$

and so any invertible sheaf on $\mathcal{X}^{\text{crys},(1,0),\tau}$ is necessarily of the form $\mathcal{L}(n)$, where \mathcal{L} is pulled back from $\widehat{\mathbf{P}}^1$, and n is an integer indicating the weight of the $\widehat{\mathbf{G}}_m$ -action. Furthermore, the dualizing sheaf on $\mathcal{X}^{\text{crys},(1,0),\tau}$ is obtained as the pull back of the dualizing sheaf $\mathcal{O}(-2)_{\widehat{\mathbf{P}}^1}$ on $\widehat{\mathbf{P}}^1 \times [(\text{Spf}(\mathcal{O}[x, y]/(xy - p))^\wedge/\widehat{\mathbf{G}}_m]$.

The short exact sequences

$$0 \rightarrow (c\text{-Ind}_K^G \Lambda_2)^\wedge \rightarrow (c\text{-Ind}_K^G \Lambda_1)^\wedge \rightarrow c\text{-Ind}_K^G \sigma \rightarrow 0$$

and

$$0 \rightarrow (c\text{-Ind}_K^G \Lambda_1)^\wedge \xrightarrow{p} (c\text{-Ind}_K^G \Lambda_2)^\wedge \rightarrow c\text{-Ind}_K^G \sigma^{\text{co}} \rightarrow 0$$

induce corresponding short exact sequences of sheaves after applying \mathfrak{A} , and from the known values of \mathfrak{A} on $c\text{-Ind}_K^G \sigma$ and $c\text{-Ind}_K^G \sigma^{\text{co}}$, we find that

$$\mathfrak{A}((c\text{-Ind}_K^G \Lambda_1)^\wedge) \cong \mathcal{L}(-1),$$

that

$$\mathfrak{A}((c\text{-Ind}_K^G \Lambda_2)^\wedge) \cong \mathcal{L}(1),$$

and that the morphism (7.4.7) is induced by multiplication by y . This analysis doesn't determine the invertible sheaf \mathcal{L} (since the restriction of any \mathcal{L} pulled back from \mathbf{P}^1 to either $\mathcal{X}(\sigma)$ or $\mathcal{X}(\sigma^{\text{co}})$ is necessarily trivial). However, the mutual duality of $\mathcal{L}(\pm 1)$ then shows that \mathcal{L} must be the pullback of $\mathcal{O}(-1)_{\widehat{\mathbf{P}}^1}$.

REMARK 7.4.8. One difference between the cuspidal and principal series cases is that in the cuspidal case, we have $\Gamma(\mathcal{X}^{\text{crys},(1,0),\tau}, \mathcal{O}_{\mathcal{X}^{\text{crys},(1,0),\tau}}) = \mathcal{O}$ (just as we have $\Gamma(\mathcal{X}(\sigma/\sigma^{\text{co}}), \mathcal{O}_{\mathcal{X}(\sigma/\sigma^{\text{co}})}) = k$). This reflects (via consideration of the functor \mathfrak{A}) the fact that the endomorphism ring of the compactly induced representation $(c\text{-Ind}_K^G \sigma^{\text{crys},\circ}(\tau))^\wedge$ (for any choice of lattice $\sigma^{\text{crys},\circ}(\tau)$) is trivial in this case.

7.4.9. *Illustrating the compatibility with the analytic case.* Allowing, for a moment, τ to denote either a tame principal series or tame cuspidal type, then, as indicated in (5.2.12) in the semi-stable case, there is a morphism

$$\pi : (\mathcal{X}^{\text{crys},(1,0),\tau})_\eta^{\text{rig}} \rightarrow (\text{Fil}_{(1,0)} \text{WD}_\tau)^{\text{an}},$$

whose target is the (rigid analytic) moduli stack of two-dimensional Weil–Deligne representations having inertial type τ , of fixed inverse cyclotomic determinant, equipped with a filtration by a one-dimensional subspace.

Suppose now that τ is a principal series type. Then

$$(\mathcal{X}^{\text{crys},(1,0),\tau})_\eta^{\text{rig}} \cong [(\{(t, x) \in (\mathbf{G}_m)_t \times \mathbf{P}_x^1 \mid 1 \geq |t| \geq |p|\} \setminus (\{(x, t) \mid 1 > |t|, |x|\} \cup \{(x, t) \mid |t| > |p|, |x| > 1\})) / (\widehat{\mathbf{G}}_m)_\eta^{\text{rig}}],$$

where $(\widehat{\mathbf{G}}_m)_\eta^{\text{rig}} = \{u \mid |u| = 1\} \subset \mathbf{G}_m$, acting via $u \cdot (t, x) = (t, u^2x)$, while

$$(\text{Fil}_{(1,0)} \text{WD}_\tau)^{\text{an}} \cong [(\mathbf{G}_m)_t \times \mathbf{P}_x^1 / (\mathbf{G}_m)_u],$$

where¹⁹ again the action is given by $u \cdot (t, x) = (t, u^2x)$. The map π is then the obvious one.

In terms of this description, the “forget the filtration” map pr_{WD} is the morphism

$$[(\mathbf{G}_m)_t \times \mathbf{P}_x^1 / (\mathbf{G}_m)_u] \rightarrow [(\mathbf{G}_m)_t / (\mathbf{G}_m)_u]$$

(the right hand stack being formed with respect to the trivial action of $(\mathbf{G}_m)_u$ on $(\mathbf{G}_m)_t$). The twist $((-\rho')_\sigma)$ which appears in (6.2.5) is, in this case, the twist by the pullback of the sheaf $\mathcal{O}(-1)_{\mathbf{P}^1}$ along the “forget the Weil–Deligne representation” map

$$\text{pr}_{\text{Fil}} : [(\mathbf{G}_m)_t \times \mathbf{P}_x^1 / (\mathbf{G}_m)_u] \rightarrow \mathbf{P}_x^1 / \text{SL}_2,$$

that is induced by projection to the second factor together with the inclusion of $(\mathbf{G}_m)_u$ as the diagonal torus in SL_2 .

Now the smooth categorical Langlands correspondence associates the trivial invertible sheaf, with trivial $(\mathbf{G}_m)_u$ -equivariant structure, to $c\text{-Ind}_K^G \sigma^{\text{crys}}(\tau)$. The expected compatibility (6.2.12) in Section 6.2.10 (taking into account also Conjecture 6.2.4 (1)) then predicts that

$$\mathfrak{A}((c\text{-Ind}_K^G \sigma^{\text{crys},\circ}(\tau))^\wedge)_\eta^{\text{rig}} \cong \text{pr}_{\text{Fil}}^*(\mathcal{O}(-1)_{\mathbf{P}^1}).$$

The line bundle $\mathcal{O}(-1)_{\mathbf{P}^1}$ on \mathbf{P}_x^1 pulls back to the trivial line bundle on the open subset

$$\{(t, x) \mid 1 \geq |t| \geq |p|\} \setminus (\{(x, t) \mid 1 > |t|, |x|\} \cup \{(x, t) \mid |t| > |p|, |x| > 1\})$$

¹⁹Here we write just \mathbf{G}_m and \mathbf{P}^1 for the multiplicative group and the projective line viewed as rigid analytic spaces.

of $(\mathbf{G}_m)_t \times \mathbf{P}_x^1$ under the projection to \mathbf{P}_x^1 , but it has non-trivial $(\mathbf{G}_m)_u$ -action: in fact, $(\mathbf{G}_m)_u$ acts on it with weight -1 . Thus (6.2.12) predicts an isomorphism

$$\mathfrak{A}((c\text{-Ind}_K^G \sigma^{\text{crys}, \circ}(\tau))^\wedge)_\eta^{\text{rig}} \cong \mathcal{O}_{(\mathcal{X}^{\text{crys}, (1,0), \tau})_\eta^{\text{rig}}}(-1),$$

in accordance with the computation of Section 7.4.5.

Now suppose that τ is cuspidal. In this case

$$(\mathcal{X}^{\text{crys}, (1,0), \tau})_\eta^{\text{rig}} \cong [(\{(t, x) \in \mathbf{P}^1 \times \mathbf{G}_m \mid 1 \geq |x| \geq |p|\} \setminus (\{(t, x) \mid 1 > |t|, |x| \geq |p|\} \cup \{(t, x) \mid |t| > 1, |x| > |p|\})) / (\widehat{\mathbf{G}}_m)_\eta^{\text{rig}}],$$

where $(\widehat{\mathbf{G}}_m)_\eta^{\text{rig}}$ acts via $u \cdot (t, x) = (t, u^2x)$, while

$$(\text{Fil}_{(1,0)} \text{WD}_\tau)^{\text{an}} \cong [\mathbf{P}_t^1 / \mu_2],$$

where μ_2 acts trivially. The map π is then the obvious one, given by forgetting x .

The map pr_{WD} is the projection

$$[\mathbf{P}_t^1 / \mu_2] \rightarrow [*/\mu_2],$$

while we again have the map

$$\text{pr}_{\text{Fil}} : [\mathbf{P}_t^1 / \mu_2] \rightarrow [\mathbf{P}_t^1 / \text{SL}_2]$$

that forgets the Weil–Deligne representations and which is induced by the central embedding of μ_2 into SL_2 . The expected compatibility (6.2.12) then predicts an isomorphism

$$\mathfrak{A}((c\text{-Ind}_K^G \sigma^{\text{crys}, \circ}(\tau))^\wedge)_\eta^{\text{rig}} \cong (\pi^* \text{pr}_{\text{Fil}}^* \mathcal{O}(-1)_{\mathbf{P}^1})(-1)$$

(the outer twist by -1 indicating the equivariant action of $(\widehat{\mathbf{G}}_m)_\eta^{\text{rig}}$), which is in accordance with the computations of Section 7.4.6. (Note that the twists by 1 and by -1 give non-isomorphic invertible sheaves on $\mathcal{X}^{\text{crys}, (1,0), \tau}$, but that these line bundles become isomorphic upon being pulled back to $(\mathcal{X}^{\text{crys}, (1,0), \tau})_\eta^{\text{rig}}$, since the invertible sheaf $\mathcal{O}_{(\mathcal{X}^{\text{crys}, (1,0), \tau})_\eta^{\text{rig}}}(2)$ is trivial, admitting as it does the coordinate x as a nowhere zero global section.)

7.5. Semiorthogonal decompositions. We now explain an explicit semiorthogonal decomposition of $\text{IndCoh}(\mathcal{X})$ (where \mathcal{X} is as in Section 7.2), which allows us to describe the essential image of the functor \mathfrak{A} . We also discuss the essential image of the analogous expected functor for representations of D^\times , where D is the nonsplit quaternion algebra with centre \mathbf{Q}_p , as well as the case of $\text{GL}_2(\mathbf{Q}_{p^2})$.

We begin by noting that we have a morphism $f : \mu_2 \times \mathcal{X} \rightarrow \mathcal{I}_\mathcal{X}$ (the inertia stack of \mathcal{X}) given by the action of $\zeta \in \mu_2$ on a (φ, Γ) -module by multiplication by ζ (note that since $\zeta^2 = 1$, this acts trivially on $\wedge^2 D$). This morphism is furthermore *central* (as one sees from its construction) with respect to the groupoid structure on $\mathcal{I}_\mathcal{X}$ (which is a groupoid over \mathcal{X}).

The morphism f induces a canonical action of μ_2 on $\text{Coh}(\mathcal{X})$, and thus on $\text{IndCoh}(\mathcal{X})$, as follows: We may find a formal scheme U with a smooth surjective morphism $U \rightarrow \mathcal{X}$ (in other words, a smooth chart for \mathcal{X}), and if $R := U \times_{\mathcal{X}} U$, then we may write $\mathcal{X} = [U/R]$. The morphism f above then induces a morphism $U \times \mu_2 \rightarrow R$ of groupoids over U , where μ_2 acts trivially on U . A coherent sheaf on \mathcal{X} may be regarded as a coherent sheaf on U with an R -equivariant structure (encoding the descent data back down to \mathcal{X}); in particular any such sheaf is equipped with

a μ_2 -equivariant structure, which (by the centrality of f) commutes with the R -equivariant structure. Thus coherent sheaves on \mathcal{X} are indeed equipped with a canonical μ_2 -equivariant structure.

We let $\mathrm{IndCoh}(\mathcal{X})^\pm$ be the subcategory of $\mathrm{IndCoh}(\mathcal{X})$, on which $-1 \in \mu_2$ acts as \pm . We sometimes refer to $\mathrm{IndCoh}(\mathcal{X})^+$ as the *even* subcategory, and $\mathrm{IndCoh}(\mathcal{X})^-$ as the *odd* subcategory.

As a first step towards a semiorthogonal decomposition of $\mathrm{IndCoh}(\mathcal{X})$, it is easy to see that we have an orthogonal decomposition into the two pieces $\mathrm{IndCoh}(\mathcal{X})^+$ and $\mathrm{IndCoh}(\mathcal{X})^-$. It follows from the construction of L_∞ that L_∞ is odd, so the essential image of \mathfrak{A} is contained in $\mathrm{IndCoh}(\mathcal{X})^-$.

7.5.1. Motivation. We refer to Appendix A.8 for some generalities on semiorthogonal decompositions. Our semiorthogonal decomposition is motivated by the possibility of a “ p -adic Fargues–Scholze conjecture” as in Remark 1.4.6, and by the structure of the stack Bun_G of G -bundles on the Fargues–Fontaine curve in the case $G = \mathrm{GL}_2(F)$. This stack admits a stratification with strata corresponding to particular isocrystals, and as in the ℓ -adic case, it is reasonable to expect that whatever categories of p -adic sheaves on Bun_G are considered, there will be a corresponding semiorthogonal decomposition into pieces which admit a description in terms of the representation theory of the automorphism groups of isocrystals. We warn the reader that we do not, however, expect that the semiorthogonal decomposition on $\mathrm{IndCoh}(\mathcal{X})$ that we consider in this section to match precisely with the (anticipated) semiorthogonal decomposition coming from Bun_G ; rather we expect that the two are related “by an upper triangular unipotent change of basis matrix” which may be quite complicated.²⁰

REMARK 7.5.2. It is expected that the formulation of p -adic local Langlands as an equivalence categories should involve a nilpotent singular support condition on $\mathrm{IndCoh}(\mathcal{X})$. We ignore this throughout this section; note that for the most part we confine ourselves to the generic part of the stack \mathcal{X} , where we expect this condition to be automatic.

7.5.3. Considerations from Bun_G . As already indicated, the structure of the semiorthogonal decomposition of $\mathrm{IndCoh}(\mathcal{X})$ that we aim to exhibit is motivated by a parallel structure on the category of constructible sheaves on Bun_G (the case $G = \mathrm{GL}_2(F)$ being the one of interest to us), arising from a natural stratification on Bun_G that we now recall.

Namely, following [FS21, §I.4] and [FS22, §2.6], the stack $\mathrm{Bun}_{\mathrm{GL}_2}$ admits a stratification into locally closed strata as follows. The connected components of $\mathrm{Bun}_{\mathrm{GL}_2}$ are indexed by elements $\alpha \in \frac{1}{2}\mathbf{Z}$, and each connected component contains a unique closed stratum. If $\alpha \in \mathbf{Z}$, this closed stratum is isomorphic to $[*/\mathrm{GL}_2(F)]$, while if $\alpha \notin \mathbf{Z}$, it is isomorphic to $[*/D^\times]$, where D is the non-split quaternion algebra with centre F . (Here and below we informally confuse the p -adic points of group schemes with the corresponding sheaves.) The other strata are indexed by the pairs of integers (m, n) with $m > n$ and $m + n = 2\alpha$, and are identified with the quotient stacks $[*/G_{m,n}]$, where

$$G_{m,n} = \begin{pmatrix} F^\times & \mathcal{BC}(\mathcal{O}(m-n)) \\ 0 & F^\times \end{pmatrix}$$

²⁰Xinwen Zhu has remarked to the authors that the situation is somewhat analogous to that of a sheaf-theoretic Radon transform as considered in [Yun09, §4].

is the semidirect product of a unipotent group diamond (in fact a Banach–Colmez space) $\mathcal{BC}(\mathcal{O}(m-n))$ by the torus $F^\times \times F^\times$. (The stack $\mathrm{Bun}_{\mathrm{PGL}_2}$ can be described by similar data, by identifying those isocrystals which differ by a twist by $\mathcal{O}(1)$; for example, the connected components are now indexed by $\alpha \in \frac{1}{2}\mathbf{Z}/\mathbf{Z}$.)

In the ℓ -adic case there is a corresponding (infinite) semiorthogonal decomposition of the category of sheaves on $\mathrm{Bun}_{\mathrm{GL}_2}$, by [FS21, Thm. I.5.1], and we expect that the same will be true in the p -adic case. (Indeed, the semiorthogonal decomposition of $\mathrm{IndCoh}(\mathcal{X})$ that we exhibit below was motivated by this expectation, and retrospectively justifies it.) In the ℓ -adic case one can furthermore ignore the Banach–Colmez spaces $\mathcal{BC}(\mathcal{O}(m-n))$, but it is *a priori* unclear that this should be the case for p -adic coefficients. As we explain below, we expect that they affect our semiorthogonal decomposition if and only if $[F : \mathbf{Q}_p] \geq m-n$. (In particular, for $\mathrm{GL}_2(\mathbf{Q}_p)$, we find that they give rise to a single constituent of the semiorthogonal decomposition of $\mathrm{IndCoh}(\mathcal{X})^+$, and do not intervene at all in the decomposition of $\mathrm{IndCoh}(\mathcal{X})^-$ — and so in particular do not intervene in the description of the essential image of \mathfrak{A} .)

As we have already indicated, in the case $F = \mathbf{Q}_p$, we construct below a semiorthogonal decomposition of $\mathrm{IndCoh}(\mathcal{X})$, and in particular we describe the image of \mathfrak{A} as the left orthogonal of a category of sheaves pushed forward from the reducible locus in \mathcal{X} . We furthermore give a semiorthogonal decomposition of the sheaves on the residually reducible locus. As indicated above, we anticipate that this latter decomposition should ultimately be related to the representation theory of \mathbf{Q}_p^\times . Furthermore, it should have a description in terms of a “spectral Eisenstein functor” involving the analogues in the Banach setting of the stacks of Borel (φ, Γ) -modules as in Remark 5.3.15. The geometry of these Borel stacks for $\mathrm{GL}_d(F)$, and the construction of a corresponding spectral Eisenstein functor, is the subject of work in progress of T.G. and Jack Sempliner (although in the case at hand, they can easily be constructed and studied directly, particularly on the generic components of \mathcal{X}).

7.5.4. Restriction to the reducible locus. We return to the setting $F = \mathbf{Q}_p$. The residually reducible locus $\mathcal{U} \subset \mathcal{X}$ is a union of connected components $\mathcal{U}(\sigma/\sigma^{\mathrm{co}})$, where as always for the sake of exposition we assume that σ is very generic. As we saw in (7.2.22), we have $\mathcal{U}(\sigma/\sigma^{\mathrm{co}}) = [\mathrm{Spf} \mathcal{O}[s, t^{\pm 1}, x, y]^\wedge / \widehat{\mathbf{G}}_m]$, where the wedge denotes (p, s, xy) -adic completion, and x and y have weights 2 and -2 respectively.

Write $A := \mathcal{O}[s, t^{\pm 1}]$, and write $B = A[x, y]$, regarded as a graded A -algebra by giving x weight 2 and y weight -2 . Let $I := (p, s) \subset A$ and $J := (I, xy) \subset B$ denote the indicated ideals, and let \mathcal{B}_J denote the full subcategory of the category of graded B -modules consisting of those modules each element of which is annihilated by a power of J . There is then a canonical equivalence

$$\mathrm{IndCoh}(\mathcal{U}(\sigma/\sigma^{\mathrm{co}})) \xrightarrow{\sim} \mathrm{IndCoh}(\mathcal{B}_J),$$

which respects the \pm -decomposition on its source and target (with respect to the μ_2 -action on the source, and with respect to parity of grading on the target).

It follows straightforwardly from Proposition C.4.17 that each \pm -component $\mathrm{IndCoh}(\mathcal{B}_J^\pm)$ admits a semiorthogonal decomposition into semiorthogonal pieces \mathcal{D}_0^\pm and $\mathcal{E}_0^\pm = (\mathcal{D}_0^\pm)^\perp$, where

$$\mathcal{D}_0^+ = \langle B/J \rangle,$$

$$\mathcal{E}_0^+ = \langle (A/I)[x](-2), (A/I)[x](-4), \dots, (A/I)[y](2), (A/I)[y](4), \dots \rangle,$$

$$\mathcal{D}_0^- = \langle B/J(1), B/J(-1) \rangle,$$

and

$$\mathcal{E}_0^- = \langle (A/I)[x](-3), (A/I)[x](-5), \dots, (A/I)[y](3), (A/I)[y](5), \dots \rangle.$$

By our Morita-theoretic construction of the restriction to $\mathcal{U}(\sigma/\sigma^{\text{co}})$ of the functor \mathfrak{A} given in Section 7.3.20, we see that the image of this restriction is precisely equal to \mathcal{D}_0^- .

7.5.5. *The essential image of \mathfrak{A} in the case of $\text{PGL}_2(\mathbf{Q}_p)$.* We now consider the pushforward along $j : \mathcal{U} \rightarrow \mathcal{X}$ of our semiorthogonal decomposition of $\text{IndCoh}(\mathcal{U})$ (ignoring as ever the non-generic components, and therefore the issue of nilpotent singular support), as in Appendix A.8.13. By (A.8.14), we have a semiorthogonal decomposition $\mathcal{A}_1, \mathcal{A}_2$ of $\text{IndCoh}(\mathcal{X})^-$, where $\mathcal{A}_2 = j_* \mathcal{E}_0^-$, and $\mathcal{A}_1 = (j^*)^{-1}(\mathcal{D}_0^-)$.

We claim that \mathcal{A}_1 is the essential image of \mathfrak{A} . To see this, note that for any x in \mathcal{A}_1 , we have the morphism

$$x \rightarrow j_* j^* x$$

(from the unit of the adjunction between j_* and j^*), whose cofiber is supported on the complement of \mathcal{U} . This cofiber is therefore in the essential image of \mathfrak{A} , because this essential image contains the (odd parity) skyscraper sheaves at the supersingular points (these are the images of the supersingular irreducible representations, via Theorem 7.3.5). It therefore suffices to show that the restriction to \mathcal{U} of the essential image of \mathfrak{A} is \mathcal{D}_0^- , which, as we already observed, follows from the explicit description of the restriction to \mathcal{U} in Section 7.3.20.

7.5.6. *Representations of D^\times .* We now consider the sheaves of even parity. We continue to assume that $p \geq 5$ for simplicity.

We begin by describing some features of a conjectural functor \mathfrak{A}_{D^\times} taking smooth representations of D^\times with trivial central character to the even category $\text{IndCoh}(\mathcal{X})^+$. We anticipate that such a functor exists, and enjoys similar properties to those of our functor \mathfrak{A} for $\text{PGL}_2(\mathbf{Q}_p)$. In particular, we expect it to be cocontinuous and fully faithful, and we expect it to satisfy a compatibility with Taylor–Wiles patching.

Using this latter expectation, we can predict the values of \mathfrak{A}_{D^\times} on a set of generators. In order to do so, we briefly recall some of the basic representation theory of D^\times . To avoid having to discuss rationality issues, we work with $\overline{\mathbf{F}}_p$ -coefficients. Writing Π for a uniformizer of D , and \mathcal{O}_D for the usual maximal order, we see that \mathcal{O}_D/Π is a quadratic extension of \mathbf{F}_p , which we identify with \mathbf{F}_{p^2} . Thus the irreducible $\overline{\mathbf{F}}_p$ -representations of \mathcal{O}_D^\times are given by the characters of $\mathbf{F}_{p^2}^\times$. Fixing an embedding $\mathbf{F}_{p^2} \hookrightarrow \overline{\mathbf{F}}_p$ defines a character $\overline{\eta} : \mathcal{O}_D^\times \rightarrow \overline{\mathbf{F}}_p^\times$, and every other character is of the form $\overline{\eta}^i$ for some $i \in \mathbf{Z}/(p^2 - 1)\mathbf{Z}$; these representations have trivial central character if and only if $(p - 1)$ divides i , which we assume from now on; then we can extend them to representations of $\Pi^{2\mathbf{Z}}\mathcal{O}_D^\times$ with trivial central character, which we also denote by $\overline{\eta}^i$.

If $(p + 1) \nmid i$ then $\sigma_i := \text{Ind}_{\Pi^{2\mathbf{Z}}\mathcal{O}_D^\times}^{D^\times} \overline{\eta}^i$ is an irreducible 2-dimensional representation of D^\times (with trivial central character); we have $\sigma_i \cong \sigma_{-i}$, and $\sigma_i \not\cong \sigma_j$ if $i^2 \neq j^2$. If $i = 0$ or $i = (p^2 - 1)/2$, then $\text{Ind}_{\Pi^{2\mathbf{Z}}\mathcal{O}_D^\times}^{D^\times} \overline{\eta}^i$ is the direct sum of two representations $\overline{\eta}_i^\pm$ with $\overline{\eta}_i^\pm(\Pi) = \pm 1$. These representations give a set of generators of the category of smooth D^\times representations with trivial central character.

We anticipate that the values of \mathfrak{A}_{D^\times} on these representations take a particularly simple form. The mod p Jacquet–Langlands correspondence of T.G. and Geraghty [GG15] associates to $\bar{\eta}_i$ a companion pair of Serre weights $\sigma, \sigma^{\text{co}}$ (in the sense of Definition 7.2.13). Indeed, by Remark 7.2.14, the pair $\sigma, \sigma^{\text{co}}$ is the set of Jordan–Hölder constituents of a cuspidal tame type, and this cuspidal tame type is the one corresponding to (the tame lift to characteristic zero of the representation of $I_{\mathbf{Q}_p}$ corresponding via local class field theory to) $\bar{\eta}_i \oplus \bar{\eta}_{-i}$.

Write $\mathcal{X}(\sigma_i) := \mathcal{X}(\sigma) \cup \mathcal{X}(\sigma^{\text{co}})$. (Note that this does not always agree with $\mathcal{X}(\sigma/\sigma^{\text{co}})$ as defined in Definition 7.2.15, because we are omitting components of the form $\mathcal{X}(\sigma_{a,p-1})^\pm$.)

Then we expect that (where $j = 0$ or 1)

$$(7.5.7) \quad \mathfrak{A}_{D^\times}(\bar{\eta}_{(p^2-1)j/2}^\pm) = \mathcal{O}_{\mathcal{X}(\sigma_{(p-1)j/2,p-1})^\pm},$$

while if $(p+1) \nmid i$ we have

$$(7.5.8) \quad \mathfrak{A}_{D^\times}(\sigma_i) = \mathcal{O}_{\mathcal{X}(\sigma_i)}.$$

Our justification for this expectation is as follows. By multiplicity one results for mod p quaternion modular forms, i.e. freeness results for the corresponding patched modules (see Section 7.5.13), and the expected compatibility of \mathfrak{A}_{D^\times} with Taylor–Wiles patching, we expect that all of these sheaves are invertible sheaves over their supports (which are the indicated closed substacks). (As far as we know, these multiplicity one results are not in the literature, so we sketch a proof of them below.) Since they are also in $\text{IndCoh}(\mathcal{X})^+$, it is natural to guess that they are given by the structure sheaves of their supports.

REMARK 7.5.9. It seems plausible that the compatibility of the functor \mathfrak{A}_{D^\times} with Conjecture 9.3.2 in Eisenstein situations (for a quaternion algebra over \mathbf{Q} which is ramified at p) would imply that these invertible sheaves are indeed just the structure sheaves of their support, adding further evidence for our expectation.

REMARK 7.5.10. If we drop the assumption that our representations have trivial central character, then the above statements extend in the obvious fashion, once we decree that a Serre weight $\sigma_{a,p-2}$ is its own companion (so the “companion pair” is a singleton in this case).

We can now repeat the analysis of Section 7.5.5. We again have a semiorthogonal decomposition $\mathcal{A}_1, \mathcal{A}_2$ of $\text{IndCoh}(\mathcal{X})^+$, where $\mathcal{A}_2 = j_*\mathcal{E}_0^+$, and $\mathcal{A}_1 = (j^*)^{-1}(\mathcal{D}_0^+)$. Using expectations (7.5.7) and (7.5.8), we see that the essential image of \mathfrak{A}_{D^\times} is contained in \mathcal{A}_1 , and that \mathcal{A}_1 is generated by this essential image, together with the (even parity) skyscraper sheaves at the irreducible points.

However, in contrast to the case of $\text{GL}_2(\mathbf{Q}_p)$, the essential image of \mathfrak{A}_{D^\times} does not contain these skyscraper sheaves! Instead, we expect that these skyscraper accounts are accounted for by the first non-open stratum of $\text{Bun}_{\text{PGL}_2}$, i.e. by the representation theory of $\mathcal{BC}(\mathcal{O}(1))$. We are therefore suggesting that $\mathcal{BC}(\mathcal{O}(1))$ behaves differently from $\mathcal{BC}(\mathcal{O}(n))$ for $n > 1$; as some justification for this, we note that $\mathcal{BC}(\mathcal{O}(n))$ is representable by a perfectoid space if $n = 1$, but not if $n > 1$.

7.5.11. *The case of $\text{GL}_2(\mathbf{Q}_{p^2})$.* It is possible to make a similar analysis in the case of $\text{GL}_2(\mathbf{Q}_{p^f})$. We can give generators for the essential image of the hypothetical functor \mathfrak{A} in a similar fashion to our predictions for D^\times above, and we can describe the irreducible components of the corresponding stack $\mathcal{X}_{\mathbf{Q}_{p^f},2}$ (for generic

Serre weights, we do this in Section 7.6). One can construct a semiorthogonal decomposition on the reducible locus by an analysis similar to the one given above.

The case $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ is very similar to the case of D^\times considered above. We again find that the essential image of (the hypothetical) \mathfrak{A} is contained in the even part of the corresponding category of sheaves, and is contained in the left orthogonal of an explicit category of sheaves pushed forward from the reducible locus. We again find that the skyscraper sheaves at irreducible points are contained in this left orthogonal, but seem not to be in the essential image of \mathfrak{A} .

We again anticipate that these skyscraper sheaves are accounted for by the representation theory of a Banach–Colmez space, in this case $\mathcal{BC}(\mathcal{O}(2))$; again, this Banach–Colmez space is represented by a perfectoid space. (For a general F , $\mathcal{BC}(\mathcal{O}(n))$ is represented by a perfectoid space if and only if $0 < n \leq [F : \mathbf{Q}_p]$, [FS21, Prop. II.2.5(iv)].)

REMARK 7.5.12. The failure of the skyscraper sheaves at irreducible points to be contained in the essential image of \mathfrak{A} is related to the failure of supersingular representations of $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ to be finitely presented; see Section 7.7.59 below. We do not know if there is a purely representation-theoretic interpretation of this failure of finite presentation in terms of $\mathcal{BC}(\mathcal{O}(2))$.

7.5.13. Multiplicity one. We end this section by sketching a proof of the multiplicity one result for D^\times mentioned above. For the sake of completeness we drop the assumption that our representations have trivial central character. Let $D_{\mathbf{Q}}$ be a quaternion algebra over \mathbf{Q} which is ramified precisely at p, ∞ , and let $\bar{r} : \mathrm{Gal}_{\mathbf{Q}} \rightarrow \mathrm{GL}_2(\bar{\mathbf{F}}_p)$ be such that $\bar{r}|_{\mathrm{Gal}_{\mathbf{Q}(\zeta_p)}}$ is irreducible. Let \mathfrak{m} be the corresponding maximal ideal of an appropriate prime-to- p Hecke algebra. (For simplicity the reader might imagine that \bar{r} is unramified outside p .) For each character $\bar{\eta}^i$ of \mathcal{O}_D^\times we have a corresponding space of quaternionic modular forms $M(\bar{\eta}^i)$ for the group $D_{\mathbf{Q}}^\times$, and a patched module $M_\infty(\bar{\eta}^i)$, which is the reduction modulo p of a patched module $M_\infty(\eta^i)$, where η is the lift of $\bar{\eta}$.

These patched modules have been studied in the papers [GS11, GG15, Roz15], and by modularity lifting theorems, one can show that $M_\infty(\eta^i)$ is maximal Cohen–Macaulay over a ring $R_\infty(\eta^i)$. The ring $R_\infty(\eta^i)$ is formally smooth over a framed deformation ring $R_{\bar{\rho}}(\eta^i)$ for $\bar{\rho} := \bar{r}|_{\mathrm{Gal}_{\mathbf{Q}_p}}$, which can be described as follows.

Firstly, if $(p+1)|i$ then after twisting we may suppose that $i = 0$, in which case $R_{\bar{\rho}}(\eta^i)$ is a framed deformation ring for non-crystalline semistable representations of Hodge–Tate weights $0, 1$. Secondly, if $(p+1) \nmid i$, then $R_{\bar{\rho}}(\eta^i)$ is a framed deformation ring for potentially crystalline representations of Hodge–Tate weights $0, 1$ and inertial type $\eta_i \oplus \eta_{-i}$.

With this description at hand, one can easily check that the supports of the patched modules agree with the prescriptions (7.5.8), (7.5.7). (This is the weight part of Serre’s conjecture for quaternionic modular forms, originally proved by T.G. and Savitt in [GS11]; the refinement to distinguish between the $\bar{\eta}_i^\pm$ was proved by Rozenzajn in [Roz15, Lem. 6.1.1], by considering “extended types”, i.e. by keeping track of an action of a Frobenius.)

Now, if $R_{\bar{\rho}}(\eta^i)$ is regular, $M_\infty(\eta^i)$ is automatically free (this is the usual argument for deducing multiplicity one results from Taylor–Wiles patching, due independently to Diamond and Fujiwara). In particular, this holds if $R_{\bar{\rho}}(\eta^i)$ is either formally smooth or is formally smooth over $\mathcal{O}[[x, y]]/(xy - p)$ where \mathcal{O} is the ring

of integers in an appropriate unramified extension of \mathbf{Q}_p . The deformation rings $R_{\bar{\rho}}(\eta^i)$ have been computed in many cases. In the case $(p+1)|i$, they are formally smooth (and can be described very explicitly, as all the lifts are ordinary) while if $(p+1) \nmid i$ then from [Sav05, Thm. 1.4] and [EGS15, Thm. 7.2.1], we know that they are regular except possibly for the cases that $\bar{\rho}$ is unramified, or $\bar{\rho}|_{I_p}$ a twist of $1 \oplus \bar{\varepsilon}$; we refer to these below as the exceptional cases.

We do not know of a reference which computes the rings $R_{\bar{\rho}}(\eta^i)$ in these exceptional cases, but it can be proved by “pure thought” (via a comparison to patched modules for crystalline representations of Hodge–Tate weights $0, p-1$) that in the case that $\bar{\rho}$ is a direct sum of distinct unramified characters, $R_{\bar{\rho}}(\eta^i)$ is not regular. In fact unpublished calculations of Bao Le Hung which will appear in [HMM22] show that it is formally smooth over $\mathcal{O}[[x, y]]/(xy - p^2)$ (at least if $p > 7$). In this case it seems unlikely that one can prove the freeness of $M_{\infty}(\eta^i)$ without using some geometric input. In the remaining unramified cases the deformation rings are again computed by Le Hung, and they are also not regular; note that the underlying cycles of their special fibres can also be read off from [San14] via the geometric Breuil–Mézard conjecture. In the case that $\bar{\rho}$ is a twist of $\lambda \oplus \bar{\varepsilon}$ where λ is a non-trivial unramified character, Le Hung’s calculations show that $R_{\bar{\rho}}(\eta^i)$ is formally smooth over $\mathcal{O}[[x, y]]/(xy - p)$, but we do not know what happens for $\lambda = 1$.

The following argument for freeness in these exceptional cases is taken from unpublished work of George Boxer, Frank Calegari and T.G. (It was independently discovered by Chengyang Bao, Andrea Dotto and Yulia Kotelnikova, and will appear in a forthcoming paper by them.) By Nakayama’s lemma, it is enough to prove that the corresponding eigenspace of quaternionic modular forms $M(\bar{\eta}^i)[\mathfrak{m}]$ is one-dimensional over $\bar{\mathbf{F}}_p$. Now, we interpret $M(\bar{\eta}^i)$ as a space of functions on the supersingular points SS of a modular curve X . More precisely, after twisting, we can identify $M(\bar{\eta}^i)_{\mathfrak{m}}$ with the term $H^0(SS, \omega^{n+p-1})_{\mathfrak{m}}$ in the short exact sequence

$$0 \rightarrow H^0(X, \omega^n)_{\mathfrak{m}} \rightarrow H^0(X, \omega^{n+p-1})_{\mathfrak{m}} \rightarrow H^0(SS, \omega^{n+p-1})_{\mathfrak{m}} \rightarrow 0$$

where either $n = 3$ and $\bar{\rho}|_{I_{\mathbf{Q}_p}} \cong \bar{\varepsilon} \oplus \bar{\varepsilon}$, or $n = 4$ and $\bar{\rho}|_{I_{\mathbf{Q}_p}} \cong \bar{\varepsilon} \oplus \bar{\varepsilon}^2$. In either case it follows from the weight part of Serre’s conjecture (and the assumption that $\bar{\rho}$ is reducible) that $H^0(X, \omega^n)_{\mathfrak{m}} = 0$, so $H^0(X, \omega^{n+p-1})_{\mathfrak{m}} \xrightarrow{\sim} H^0(SS, \omega^{n+p-1})_{\mathfrak{m}}$, and we need only show that $H^0(X, \omega^{n+p-1})[\mathfrak{m}]$ is 1-dimensional. Suppose that this is not the case; then there is an eigenform with $a_1 = 0$, and by the q -expansion principle, such an eigenform is in the kernel of the θ operator. But this kernel is trivial by a theorem of Katz [Kat77], so we are done.

7.6. Preliminaries in the Banach case for $\mathrm{GL}_2(\mathbf{Q}_{p^f})$. Following the case of $\mathrm{GL}_2(\mathbf{Q}_p)$, a natural next case to consider is that of $G := \mathrm{GL}_2(\mathbf{Q}_{p^f})$, for $f > 1$; here, as usual, \mathbf{Q}_{p^f} denotes the unramified extension of \mathbf{Q}_p with residue field $\mathbf{F}_q := \mathbf{F}_{p^f}$, and we write $K := \mathrm{GL}_2(\mathbf{Z}_{p^f})$ and $Z = Z(G)$. For example, the mod p local Langlands correspondence has been studied quite intensively in this case, as has the problem of mod p local-global compatibility. (We recall some of the history of these investigations as we go along; the reader who is unfamiliar with it may wish to consult [Bre10a] and the introduction to [BHH⁺21].) On the other hand, the structure of the category of mod p (or, more generally, mod p^n) smooth G -representations remains rather mysterious, the classification of supersingular irreducible admissible representations remains unknown, and as of yet there is no

precise formulation of a mod p or p -adic Langlands correspondence in the “traditional” mode as exemplified by [Col10c] and [Paš13] in the $\mathrm{GL}_2(\mathbf{Q}_p)$ -case.

We believe that the categorical perspective sheds quite a bit of light on this case, helping both to illuminate known results and to indicate why a traditional mode of formulation for the p -adic Langlands correspondence has proved so elusive. In the following discussion we elaborate on these points, by describing some of our²¹ preliminary results and expectations for the categorical Langlands program for the groups $\mathrm{GL}_2(\mathbf{Q}_{p^f})$. To ease notation and fix ideas, we primarily focus on the case when $f = 2$, but we begin in this section with some considerations that are valid for any f .

REMARK 7.6.1. The reader who is familiar with the various difficulties and mysteries in the p -adic local Langlands program for $\mathrm{GL}_2(\mathbf{Q}_{p^2})$, and who is anxious to see what light we can shed on them, may wish to jump to Section 7.7.59 and Remark 7.8.5.

REMARK 7.6.2. The similarities and differences between the case of $\mathrm{GL}_2(\mathbf{Q}_p)$ and $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ can be illustrated by the following example. In either case, for a generic Serre weight σ , the functor \mathfrak{A} takes $c\text{-Ind}_{KZ}^G \sigma$ to an explicit line bundle supported on the corresponding component $\mathcal{X}(\sigma)$ of $\mathcal{X}_{\mathrm{red}}$. In the case of $\mathrm{GL}_2(\mathbf{Q}_p)$, the correspondence is “not very derived”; in particular, a simple calculation shows that the supersingular irreducible admissible representation $(c\text{-Ind}_{KZ}^G \sigma)/T$ is taken to a (twisted) skyscraper sheaf supported at the corresponding irreducible Galois representation.

For $\mathrm{GL}_2(\mathbf{Q}_{p^2})$, the same calculation shows that $(c\text{-Ind}_{KZ}^G \sigma)/T$ is taken to a sheaf supported on the closed locus on $\mathcal{X}(\sigma)$ corresponding to $T = 0$. However, this support is now one-dimensional, and the representation $(c\text{-Ind}_{KZ}^G \sigma)/T$ has infinite length. In fact, we know (see (7.7.53)) that it admits as subrepresentations the full compact inductions $c\text{-Ind}_{KZ}^G \sigma'$ for certain weights $\sigma' \neq \sigma$, and it follows easily that the supersingular irreducible admissible representations of $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ cannot possibly be taken to actual sheaves (as opposed to complexes of sheaves), in complete contrast to the case of $\mathrm{GL}_2(\mathbf{Q}_p)$ (or the case $\ell \neq p$).

7.6.3. *The generic structure of $\mathcal{X}_{\mathrm{red}}$.* As always we fix our coefficients \mathcal{O} to be the ring of integers in some finite extension L of \mathbf{Q}_p , which we furthermore assume admits an embedding of \mathbf{Q}_{p^f} . (In other words, we may regard L as an extension of \mathbf{Q}_{p^f} .) We let k denote the residue field of \mathcal{O} . We also fix a character $\zeta : \mathbf{Q}_{p^f}^\times \rightarrow \mathcal{O}^\times$, and let \mathcal{X} denote the Noetherian formal algebraic stack classifying rank 2 projective étale (φ, Γ) -modules for $\mathrm{Gal}_{\mathbf{Q}_{p^f}}$ over p -adically complete \mathcal{O} -algebras, having their determinant fixed to equal $\zeta \varepsilon^{-1}$.

7.6.4. *The components $\mathcal{X}(\sigma)$.* Just as in the case of $\mathrm{GL}_2(\mathbf{Q}_p)$, the underlying reduced algebraic substack $\mathcal{X}_{\mathrm{red}}$ is of finite presentation over $\mathrm{Spec} k$, and is a union of finitely many irreducible components $\mathcal{X}(\sigma)$ labelled by Serre weights σ having central character $\zeta|_{\mathcal{O}^\times}$, with the exception that there is a pair of components $X(\sigma)^\pm$ when σ is a twist of the Steinberg weight. And again, as in that case, for a *generic*

²¹Many of these ideas have been worked out in joint work of M.E. and T.G. with Ana Caraiani, Michael Harris, Bao Le Hung, Brandon Levin, and David Savitt, under the auspices of the NSF grant 1952556 “Collaborative Research: Geometric Structures in the p -adic Langlands Program”, and we thank them for their permission to describe these results here.

σ (one of the form

$$\sigma = (\mathrm{Sym}^{a_0} \otimes \det^{b_0}) \otimes \cdots \otimes (\mathrm{Sym}^{a_i} \otimes \det^{b_i})^{\mathrm{Fr}^i} \otimes \cdots \otimes (\mathrm{Sym}^{a_{f-1}} \otimes \det^{b_{f-1}})^{\mathrm{Fr}^{f-1}}$$

with all the a_i satisfying $0 \leq a_i \leq p-3$) we have that $\mathcal{X}(\sigma)$ can be described as a moduli stack of Fontaine–Laffaille modules with coefficients in k -algebras.

In general, for a reductive group G over \mathbf{F}_p , the corresponding Fontaine–Laffaille moduli space can be described as the quotient $(U \backslash G)/T$, where U is the unipotent radical of some fixed Borel in G , acting by left translation, and T is the maximal torus of the the same Borel, acting by Frobenius twisted conjugation. In our case, the group G is the restriction of scalars of SL_2 from \mathbf{F}_{p^f} to \mathbf{F}_p . Since k is assumed to contain \mathbf{F}_{p^f} , we have

$$G = G/k = \mathrm{SL}_2^f,$$

$$U = U/k = \left(\begin{pmatrix} 1 & \mathbf{G}_a \\ 0 & 1 \end{pmatrix} \right)^f,$$

and

$$T = T/k = \mathbf{G}_m^f$$

with Frobenius cyclically permuting the factors.

Now $\left(\begin{pmatrix} 1 & \mathbf{G}_a \\ 0 & 1 \end{pmatrix} \right) \backslash \mathrm{SL}_2 = \mathbf{A}^2 \setminus \{0\}$ (via $\begin{pmatrix} * & * \\ x & y \end{pmatrix} \mapsto (x, y)$). Thus we have

$$U \backslash G := \prod_{i=0}^{f-1} (\mathbf{A}^2 \setminus \{0\})_{(x_i, y_i)}$$

(where the subscripts indicate our notation for the coordinates). The torus T then acts via

$$(t_0, \dots, t_{f-1}) \cdot (x_i, y_i) = (t_{i-1} t_i x_i, t_{i-1} t_i^{-1} y_i).$$

So (for generic σ) we have find that $\mathcal{X}(\sigma)$ is equal to the quotient stack

$$\mathcal{X}(\sigma) = [(\prod_{i=0}^{f-1} (\mathbf{A}^2 \setminus \{0\})_{(x_i, y_i)}) / \mathbf{G}_m^f]$$

with the torus action being given by the preceding formula.

The coordinates (x_i, y_i) on $\mathcal{X}(\sigma)$ are evidently not completely canonical, but they are not completely without meaning either; in particular, we will shortly see that their zero loci describe the loci of intersection of $\mathcal{X}(\sigma)$ with other components. A special role is played by the product

$$T := y_0 \cdots y_{f-1}.$$

Evidently, T is \mathbf{G}_m^f invariant, and so defines a regular function on $\mathcal{X}(\sigma)$. In fact, it generates the ring of functions on $\mathcal{X}(\sigma)$, as we record in the following lemma.

LEMMA 7.6.5.

- (1) $\Gamma(\mathcal{X}(\sigma), \mathcal{O}_{\mathcal{X}(\sigma)}) = k[T]$.
- (2) The distinguished open subset $D(T)$ (i.e. the non-vanishing locus of T) in $\mathcal{X}(\sigma)$ is precisely the niveau 1 ordinary locus in $\mathcal{X}(\sigma)$.

Concretely, we see that

$$D(T) = [\mathbf{A}_{(x_0, \dots, x_{f-1})}^f / (\mathbf{G}_m)_t] \times (\mathbf{G}_m)_T,$$

where t acts via

$$t \cdot (x_0, \dots, x_{f-1}) = (t^2 x_0, \dots, t^2 x_{f-1});$$

so $(\mathbf{G}_m)_t$ can be identified with the diagonal copy of \mathbf{G}_m in \mathbf{G}_m^f . There is a character $\chi : \text{Gal}_{\mathbf{Q}_{p^f}} \rightarrow k^\times$, depending on σ , with the property that the point (T, x_0, \dots, x_{f-1}) of $D(T)$ corresponds to a two-dimensional representation $\bar{\rho}$ of $\text{Gal}_{\mathbf{Q}_{p^f}}$ of the form

$$\begin{pmatrix} \text{ur}_T \chi & * \\ 0 & \text{ur}_{T^{-1}} \zeta \varepsilon^{-1} \chi^{-1} \end{pmatrix},$$

where $*$ is an element of $\text{Ext}_{\text{Gal}_{\mathbf{Q}_{p^f}}}^1(\text{ur}_{T^{-1}} \zeta \varepsilon^{-1} \chi^{-1}, \text{ur}_T \chi)$, an f -dimensional vector space with coordinates x_0, \dots, x_{f-1} .

The zero locus $Z(T) := \mathcal{X}(\sigma) \setminus D(T)$ is more complicated in its geometry; we describe it in detail in the case $f = 2$ in Section 7.7.

Suppose now that σ is *very generic* in the sense that each a_i satisfies $0 < a_i < p - 3$. Then there are $2f$ Serre weights σ' such $\text{Ext}_{\text{GL}_2(\mathbf{F}_q)}^1(\sigma, \sigma') \neq 0$, and each such σ' is generic. One then finds that $\mathcal{X}(\sigma)$ and $\mathcal{X}(\sigma')$ intersect in one of the codimension 1 closed substacks of $\mathcal{X}(\sigma)$ cut out by either $x_i = 0$ or $y_i = 0$ (for some choice of i).

7.6.6. *Picard groups.* Continue to assume that σ is generic. From the description of $\mathcal{X}(\sigma)$ as a quotient stack, it is easy to see that $\text{Pic}(\mathcal{X}(\sigma)) \cong \mathbf{Z}^f$ (the character lattice of \mathbf{G}_m^f). This isomorphism is not completely canonical, but it *is* canonical up to permuting the coordinates of \mathbf{Z}^f , as we now explain. For this, we assume that σ is in fact very generic.

To begin with, we consider the open substack $D(T)$ of $\mathcal{X}(\sigma)$. As already noted, this admits the description $D(T) = (\mathbf{A}^f / \mathbf{G}_m) \times \mathbf{G}_m$, from which one deduces that $\text{Pic}(D(T)) = \mathbf{Z}$ (the character lattice of \mathbf{G}_m). Furthermore, this identification *is canonical*, since the \mathbf{G}_m that appears is naturally oriented (corresponding to the fact that $D(T)$ parameterizes extensions of $\text{ur}_{T^{-1}} \zeta \varepsilon^{-1} \chi^{-1}$ by $\text{ur}_T \chi$ in a fixed direction).

For each $i = 0, \dots, f - 1$, write σ_i for the weight such that $\mathcal{X}(\sigma) \cap \mathcal{X}(\sigma_i) = \mathcal{X}(\sigma)^{y_i=0}$; since σ is very generic each σ_i is again generic. We write T_i for the analogue of the function T on $\mathcal{X}(\sigma_i)$; then

$$(T_i)_{\mathcal{X}(\sigma_i)^{y_i=0}} = \frac{x_{i-1}}{x_i} \prod_{j \neq i-1, i} y_j$$

Thus

$$\mathcal{X}(\sigma)^{y_i=0} \cap D(T_i) = \mathcal{X}(\sigma)^{x_{i-1} \prod_{j \neq i-1, i} y_j \neq 0, y_i=0}.$$

Now

$$D(x_{i-1} x_i \prod_{j \neq i-1, i} y_j) = [\mathbf{A}_{(x_0, \dots, x_{i-2}, y_{i-1}, y_i, x_{i+1}, \dots, x_{f-1})}^f / (\mathbf{G}_m)_t],$$

where t acts via

$$\begin{aligned} t \cdot (x_0, \dots, x_{i-2}, y_{i-1}, y_i, x_{i+1}, \dots, x_{f-1}) \\ = (t^2 x_0, \dots, t^2 x_{i-2}, t^2 y_{i-1}, t^{-2} y_i, t^2 x_{i+1}, \dots, t^2 x_{f-1}); \end{aligned}$$

the torus $(\mathbf{G}_m)_t$ can be regarded as the copy of \mathbf{G}_m embedded into \mathbf{G}_m^f via

$$t \mapsto (t, \dots, t, t^{-1} \text{ in the } i-1 \text{ slot}, t, \dots, t).$$

The closed substack $\mathcal{X}(\sigma)^{y_i=0} \cap D(T_i)$ is then the copy of $[\mathbf{A}^{f-1}/\mathbf{G}_m]$ cut out by $y_i = 0$. Of course the embedding

$$\mathcal{X}(\sigma)^{y_i=0} \cap D(T_i) \hookrightarrow D(T_i)$$

embeds $[\mathbf{A}^{f-1}/\mathbf{G}_m]$ into a different copy of $[\mathbf{A}^f/\mathbf{G}_m]$.

We've already observed that $\text{Pic}(D(T_i)) = \text{Pic}([\mathbf{A}^f/\mathbf{G}_m]) = \mathbf{Z}$, and the restriction map $\text{Pic}(D(T_i)) \rightarrow \text{Pic}([\mathbf{A}^{f-1}/\mathbf{G}_m]) = \mathbf{Z}$ is an isomorphism. Thus $\text{Pic}(\mathcal{X}(\sigma)^{y_i=0} \cap D(T_i)) = \mathbf{Z}$ (canonically). We now consider the product of the restriction morphisms

$$(7.6.7) \quad \mathbf{Z}^f \cong \text{Pic}(X(\sigma)) \rightarrow \text{Pic}(D(T)) \times \prod_{i=0}^{f-1} \text{Pic}(\mathcal{X}(\sigma)^{y_i=0} \cap D(T_i)) \\ = \mathbf{Z} \times \mathbf{Z}^f = \mathbf{Z}^{f+1}.$$

Recalling the descriptions of the $D(T)$ and the various $\mathcal{X}(\sigma)^{y_i=0} \cap D(T_i)$ as quotients $[\mathbf{A}^f/(\mathbf{G}_m)_t]$ and $[\mathbf{A}^{f-1}/(\mathbf{G}_m)_t]$, and of the embeddings of these various copies of $(\mathbf{G}_m)_t$ into $(\mathbf{G}_m)^f$, we find that this map is given by

$$(a_0, \dots, a_{f-1}) \mapsto \left(\sum_j a_j, \left(\sum_{j \neq i-1} a_j - a_i \right)_{i=0}^{f-1} \right) = \left(\sum_j a_j, \left(\sum_{j=0}^{f-1} a_j - 2a_i \right)_{i=0}^{f-1} \right).$$

So we see that (7.6.7) is an embedding, and that we can recover the various a_i from the coordinates in \mathbf{Z}^{f+1} ; explicitly, the (left) inverse is given by

$$(b, b_0, \dots, b_{f-1}) \mapsto \frac{1}{2}(b - b_0, \dots, b - b_{f-1}).$$

Now the coordinates (b, b_0, \dots, b_{f-1}) are canonical, up to permuting the b_i (corresponding to a relabelling of the σ_i), and so we see that the isomorphism $\text{Pic}(X(\sigma)) \cong \mathbf{Z}^f$ is canonical, up to permuting the coordinates a_i .

7.6.8. Tame types. If τ is a tame inertial type for \mathbf{Q}_{p^f} , then the associated K -representation $\sigma^{\text{crys}}(\tau)$ factors through $\text{GL}_2(k_F) = \text{GL}_2(\mathbf{F}_{p^f})$. If $\sigma^{\text{crys}}(\tau)^\circ$ is a K -invariant lattice in $\sigma^{\text{crys}}(\tau)$, then the multi-set of Jordan–Hölder factors of its reduction is in fact a set of distinct Serre weights which is independent of the particular choice of lattice, and which we denote by $\text{JH}(\overline{\sigma^{\text{crys}}(\tau)})$. If τ is furthermore sufficiently generic, then $\text{JH}(\overline{\sigma^{\text{crys}}(\tau)})$ consists of exactly 2^f weights.

We write $\mathcal{X}(\tau) = \bigcup_{\sigma \in \text{JH}(\overline{\sigma^{\text{crys}}(\tau)})} \mathcal{X}(\sigma)$; this is a closed substack of \mathcal{X}_{red} .

THEOREM 7.6.9. *If τ is a sufficiently generic tame type, then $\text{Pic}(\mathcal{X}(\tau)) \cong \mathbf{Z}^{2^f}$.*

In fact, the product of the restriction maps

$$\text{Pic}(\mathcal{X}(\tau)) \hookrightarrow \prod_{\sigma \in \text{JH}(\overline{\sigma^{\text{crys}}(\tau)})} \text{Pic}(\mathcal{X}(\sigma))$$

is injective, and one can identify the image of $\text{Pic}(\mathcal{X}(\tau))$ as a certain specific subgroup of the target.

7.6.10. *Values of the functor.* The results of [EGS15] show that if σ is a generic Serre weight, or an appropriately chosen lattice $\sigma^{\text{crys}}(\tau)^\circ$ for a sufficiently generic tame type τ (having determinant equal to $\zeta|_{\mathcal{O}_F^\times}$), then the patched modules $M_\infty(\sigma)$ and $M_\infty(\overline{\sigma^{\text{crys}}(\tau)^\circ})$ (constructed from an appropriate Shimura curve context) are free of rank one over their support. The weight part of Serre’s conjecture and the (geometric) Breuil–Mézard conjecture for potentially Barsotti–Tate representations imply that the support of $M_\infty(\sigma)$ is precisely the pullback to the corresponding versal ring of $\mathcal{X}(\sigma)$, and similarly the support of $M_\infty(\overline{\sigma^{\text{crys}}(\tau)^\circ})$ is the pullback of $\mathcal{X}(\tau)$. (See [EG22, Thm. 8.6.2].) Thus we expect that the value of the functor \mathfrak{A} on $c\text{-Ind}_{KZ}^G \sigma$ or $c\text{-Ind}_{KZ}^G \overline{\sigma^{\text{crys}}(\tau)^\circ}$ will be a line bundle on $\mathcal{X}(\sigma)$ or $\mathcal{X}(\tau)$ respectively. This explains our interest in computing the Picard groups of these stacks.

More generally, if V is a representation of $\text{GL}_2(k_F)$ over F which is multiplicity free (i.e. each of whose Jordan–Hölder factors occurs with multiplicity exactly one), and which has central character equal to $\bar{\zeta}$, then the value of \mathfrak{A} on $c\text{-Ind}_{KZ}^G V$ should be scheme-theoretically supported on $\bigcup_{\sigma \in \text{JH}(V)} \mathcal{X}(\sigma)$, and in particular on \mathcal{X}_{red} . (This follows from the expected compatibility with the Breuil–Mézard conjecture, and the expected maximal Cohen–Macaulay property of $\mathfrak{A}(c\text{-Ind}_{KZ}^G V)$.)

EXPECTED THEOREM 7.6.11 (Caraiani–E.–G.–Harris–Le Hung–Levin–Savitt). *There is a fully faithful functor \mathfrak{A} from the additive category of finitely presented smooth representations generated by the objects $c\text{-Ind}_{KZ}^G V$ (with V multiplicity free, and having all Jordan–Hölder factors being sufficiently generic) to the abelian category $\text{Coh}(\mathcal{X}_{\text{red}})$.*

The functor \mathfrak{A} of this theorem, which we certainly believe to be the restriction of the functor \mathfrak{A} of Conjecture 6.1.14 for $\text{GL}_2(\mathbf{Q}_p^f)$, is produced by an explicit construction, building on the description of the various Picard groups given above. In particular, we have the following expectation.

EXPECTATION 7.6.12. *For $G = \text{GL}_2(\mathbf{Q}_p^f)$, and for a very generic Serre weight σ , the sheaf $\mathfrak{A}(c\text{-Ind}_{KZ}^G \sigma)$ is the line bundle $\mathcal{O}(-1, \dots, -1)$ on $\mathcal{X}(\sigma)$, i.e. the line bundle corresponding to the element $(-1, \dots, -1) \in \text{Pic}(\mathcal{X}(\sigma)) \cong \mathbf{Z}^f$.*

(Note that $(-1, \dots, -1)$ is invariant under permutation, and so does indeed give a canonically defined element of $\text{Pic}(\mathcal{X}(\sigma))$.)

The functor \mathfrak{A} of Expected Theorem 7.6.11 is defined by imposing Expectation 7.6.12 as a definition; it turns out that the extension of \mathfrak{A} to the additive subcategory in its statement is then essentially forced.

REMARK 7.6.13. In the case $f = 1$ Expectation 7.6.12 agrees with the results of [DEGa], as explained in Section 7.4.1. In general the expected compatibility with patching, and the weight part of Serre’s conjecture, imply that $\mathfrak{A}(c\text{-Ind}_{KZ}^G \sigma)$ will be a line bundle on $\mathcal{X}(\sigma)$. One can then check that the full faithfulness of \mathfrak{A} , and the possible extensions between the various $c\text{-Ind}_{KZ}^G \sigma$, imply that if the corresponding element of the Picard group — identified with \mathbf{Z}^f as above — is independent of σ , then it is necessarily $(-1, \dots, -1)$. The interested reader is invited to check this in the case $f = 2$ using the calculations of Section 7.7 below. (Alternatively, this can be seen from the expected compatibility with duality, as mentioned in Section 7.4.1 in the case $F = \mathbf{Q}_p$.)

We also note that the discussion of Remark 7.8.4 below shows that Expectation 7.6.12 is compatible with the philosophy introduced by Breuil in [Bre11].

7.7. Partial results and examples in the Banach case for $\mathrm{GL}_2(\mathbf{Q}_{p^2})$. In order to continue, we restrict to the case $f = 2$, in which case we can make much of the preceding discussion quite explicit. To begin with, we simplify the notation by writing $(u, v), (x, y)$ (rather than $(x_0, y_0), (x_1, y_1)$) for the coordinates on G/U , and (s, t) (rather than t_0, t_1) for the coordinates on T .

Then, if σ is a generic Serre weight, we have

$$(7.7.1) \quad \mathcal{X}(\sigma) = [(\mathbf{A}^2 \setminus \{0\})_{(u,v)} \times (\mathbf{A}^2 \setminus \{0\})_{(x,y)} / (\mathbf{G}_m \times \mathbf{G}_m)_{(s,t)}],$$

where

$$(s, t) \cdot ((u, v), (x, y)) = ((stu, s^{-1}tv), (stx, st^{-1}y)).$$

Because the diagonal copy of μ_2 in T acts trivially in this formula, and because (as it furthermore turns out) these copies of μ_2 (for the various generic Serre weights σ) act trivially on all the sheaves that we will consider, it is notationally convenient to work instead on stacks that we will denote \mathcal{C}_σ , defined as follows.

We set $a = st, b = s^{-1}t$, so that

$$(7.7.2) \quad (s, t) \mapsto (a, b)$$

is a self-isogeny of $\mathbf{G}_m \times \mathbf{G}_m$ having the diagonal copy of μ_2 as its kernel. As already noted, the action of $\mathbf{G}_m \times \mathbf{G}_m$ on $(\mathbf{A}^2 \setminus \{0\})_{(u,v)} \times (\mathbf{A}^2 \setminus \{0\})_{(x,y)}$ occurring in the definition of $\mathcal{X}(\sigma)$ then factors through this isogeny, and we write \mathcal{C}_σ for the quotient of $(\mathbf{A}^2 \setminus \{0\})_{(u,v)} \times (\mathbf{A}^2 \setminus \{0\})_{(x,y)}$ by the corresponding isogenous copy of $\mathbf{G}_m \times \mathbf{G}_m$. Concretely, we write

$$(7.7.3) \quad \mathcal{C}_\sigma = [(\mathbf{A}^2 \setminus \{0\})_{(u,v)} \times (\mathbf{A}^2 \setminus \{0\})_{(x,y)} / (\mathbf{G}_m \times \mathbf{G}_m)_{(a,b)}],$$

where

$$(a, b) \cdot ((u, v), (x, y)) = ((au, bv), (ax, b^{-1}y)).$$


There is then a natural map $\mathcal{X}(\sigma) \rightarrow \mathcal{C}_\sigma$, and the various stacks \mathcal{C}_σ glue together, in a manner that emulates the gluing together of the $\mathcal{X}(\sigma)$ in $\mathcal{X}_{\mathrm{red}}$, to produce a reduced algebraic stack that we will denote $\mathcal{C}_{\mathrm{red}}$.

There are two open covers of \mathcal{C}_σ that we wish to introduce. To describe them, we first remark that, as usual, if f is a function on a stack, or more generally, a section of a line bundle, then we let $D(f)$ denote the non-vanishing locus of f ; this is an open subset of the stack in question. We then set $\mathcal{U}_\sigma := D(vy)$, and $\mathcal{V}_\sigma := D(u) \cup D(x)$, and note that $\mathcal{C}_\sigma = \mathcal{U}_\sigma \cup \mathcal{V}_\sigma$. We also write $\mathcal{Z}_\sigma := \mathcal{C}_\sigma \setminus \mathcal{U}_\sigma$.

There is an evident isomorphism

$$\mathcal{U}_\sigma = [(\mathbf{A}_{(u,x)}^2 / (\mathbf{G}_m)_a] \times (\mathbf{G}_m)_{vy}$$

(here the subscripts denote the coordinates), where of course $a \cdot (u, x) = (au, ax)$. The Galois-theoretic interpretation of this locus is straightforward: the points correspond to extensions of a pair of characters of some fixed inertial weights; the quantity vy corresponds to a certain Frobenius eigenvalue appearing in these characters, and the \mathbf{A}^2 is the 2-dimensional Ext^1 between these characters.



$$\mathcal{U}_\sigma = [\mathbb{A}^2 / \mathbb{G}_m] \times \mathbb{G}_m$$


There is an isomorphism

$$\mathcal{Z}_\sigma = \mathbf{P}_{[u:x]}^1 \times [\{(v, y) \mid vy = 0\} / (\mathbf{G}_m)_b] \setminus (\{u = v = 0\} \cup \{x = y = 0\}),$$

where b acts via $b \cdot (v, y) = (bv, b^{-1}y)$. If we consider the open substack of \mathcal{Z}_σ where $ux \neq 0$, we obtain the stack

$$(\mathbf{G}_m)_{x/u} \times [\{(v, y) \mid vy = 0\} / (\mathbf{G}_m)_b],$$

which is a family of extensions of characters of certain fixed inertial weights, the two lines corresponding to extensions in either possible directions, but which are “split along one of the two possible embeddings”. The variable x/u again parameterizes a Frobenius eigenvalue appearing in these characters. The points $\{u = y = 0\}$ and $\{v = x = 0\}$ each correspond to an irreducible Galois representation.



The open cover $\{\mathcal{U}_\sigma, \mathcal{V}_\sigma\}$ of \mathcal{C}_σ is useful when we want to think about points of $\mathcal{X}(\sigma)$ in terms of the Galois representations that they correspond to. However, there is another open cover we will work with, which at first seems less conceptually meaningful, but which leads to simple computations (especially when we consider how the various components $\mathcal{X}(\sigma)$ intersect one another); this is the open cover $\{D(u), D(v)\}$.

Concretely, we see that

$$(7.7.4) \quad D(u) = [(\mathbf{A}_v^1 \times (\mathbf{A}^2 \setminus \{0\})_{x/u, y}) / (\mathbf{G}_m)_b],$$

where $(\mathbf{G}_m)_b$ acts via $b \cdot (v, x/u, y) = (bv, x/u, b^{-1}y)$, while

$$(7.7.5) \quad D(v) = [(\mathbf{A}_u^1 \times (\mathbf{A}^2 \setminus \{0\})_{x, vy}) / (\mathbf{G}_m)_a],$$

where $(\mathbf{G}_m)_a$ acts via $a \cdot (u, x, vy) = (au, ax, vy)$. We can also describe their intersection:

$$D(u) \cap D(v) = D(uv) = (\mathbf{A}^2 \setminus \{0\})_{x/u, vy}.$$

Since $(\mathbf{A}^2 \setminus \{0\}) \times (\mathbf{A}^2 \setminus \{0\})$ has trivial Picard group, we see from (7.7.1) that

$$\text{Pic}(\mathcal{X}(\sigma)) = \mathfrak{X}^\bullet((\mathbf{G}_m \times \mathbf{G}_m)_{s,t}) = \mathbf{Z} \times \mathbf{Z},$$

and from (7.7.3) that

$$\text{Pic}(\mathcal{C}_\sigma) = \mathfrak{X}^\bullet((\mathbf{G}_m \times \mathbf{G}_m)_{a,b}) = \mathbf{Z} \times \mathbf{Z}.$$

The pull back map $\text{Pic}(\mathcal{C}_\sigma) \rightarrow \text{Pic}(\mathcal{X}(\sigma))$ is just the map of character lattices induced by the isogeny (7.7.2), and so is given by the matrix $\begin{pmatrix} -1 & 1 \\ -1 & -1 \end{pmatrix}$.

Similarly, we see that

$$\text{Pic}(D(u)) = X^\bullet((\mathbf{G}_m)_b) = \mathbf{Z},$$

while

$$\text{Pic}(D(v)) = X^\bullet((\mathbf{G}_m)_a) = \mathbf{Z}.$$

Thus the product of the restriction maps

$$\mathrm{Pic}(\mathcal{C}_\sigma) \rightarrow \mathrm{Pic}(D(v)) \times \mathrm{Pic}(D(u))$$

is an isomorphism, indeed is just the identity map from $\mathbf{Z} \times \mathbf{Z}$ to itself. This is one reason that the open cover $\{D(u), D(v)\}$ is useful: we can describe a line bundle on \mathcal{C}_σ simply by describing its restrictions to each of $D(u)$ and $D(v)$, and line bundles on each of these open substacks are given by prescribing a single twist (i.e. a single integer).

Following Expectation 7.6.12 (and taking into account our change of coordinates (7.7.2)), for a generic Serre weight σ we define

$$\mathfrak{A}(c\text{-Ind}_{KZ}^G \sigma) := \mathcal{O}_{\mathcal{C}_\sigma}(-1, 0).$$

The endomorphism T of $c\text{-Ind}_{KZ}^G \sigma$ will then correspond to the endomorphism of $\mathcal{O}_{\mathcal{C}_\sigma}(-1, 0)$ given by multiplication by vy .

7.7.6. Intersections and unions of generic components. Two generic components $\mathcal{X}(\sigma)$ and $\mathcal{X}(\sigma')$ intersect in a codimension one locus precisely if σ and σ' admit a non-trivial extension as representations of $\mathrm{GL}_2(\mathbf{F}_{p^2})$. For a sufficiently generic σ , there are four such σ' . The corresponding loci of intersection on $\mathcal{X}(\sigma)$ are respectively $\{u = 0\}$, $\{v = 0\}$, $\{x = 0\}$, and $\{y = 0\}$. On each of the corresponding components $\mathcal{X}_{\sigma'}$, these same loci are given by $\{v' = 0\}$, $\{u' = 0\}$, $\{y' = 0\}$, and $\{x' = 0\}$. (Here we have used primes to denote the variables on $\mathcal{X}_{\sigma'}$.) These loci are compatible with passing from the $\mathcal{X}(\sigma)$ to the \mathcal{C}_σ , and indeed the whole picture of $\mathcal{X}_{\mathrm{red}}$ that we will describe descends to the \mathcal{C}_σ level, and this is where we work from now on.

7.7.7. A brief interlude on representation theory. If ξ is a character of the torus $T(\mathbf{F}_q) := \mathbf{F}_q^\times \times \mathbf{F}_q^\times$ in $\mathrm{GL}_2(\mathbf{F}_q)$ (with values in k^\times , say, although the following discussion is independent of the particular field or ring of coefficients), then we can define the principal series $\mathrm{PS}(\xi) := \mathrm{Ind}_{B(\mathbf{F}_q)}^{\mathrm{GL}_2(\mathbf{F}_q)} \xi$ (where B denotes the subgroup scheme of upper triangular matrices in GL_2). We may inflate the $\mathrm{GL}_2(\mathbf{F}_q)$ -action on $\mathrm{PS}(\xi)$ to a K -action, and, assuming that ξ is suitably compatible with our central character ζ , we can then extend this K -action to a KZ -action by having Z act via ζ . We may then form the compact induction $c\text{-Ind}_{KZ}^G \mathrm{PS}(\xi)$. If we let s denote the non-trivial Weyl group involution of $T(\mathbf{F}_q)$, then we may similarly consider $\mathrm{PS}(\xi^s)$ and $c\text{-Ind}_{KZ}^G \mathrm{PS}(\xi^s)$.

Now if $\xi \neq \xi^s$, then $\mathrm{PS}(\xi)$ and $\mathrm{PS}(\xi^s)$ are not isomorphic as $\mathrm{GL}_2(\mathbf{F}_q)$ -representations (and hence are not isomorphic as KZ -representations). However, there *is* an isomorphism

$$(7.7.8) \quad c\text{-Ind}_{KZ}^G \mathrm{PS}(\xi) \xrightarrow{\sim} c\text{-Ind}_{KZ}^G \mathrm{PS}(\xi^s).$$

To see this, write

$$c\text{-Ind}_{KZ}^G \mathrm{PS}(\xi) = c\text{-Ind}_{IZ}^G \xi$$

and

$$c\text{-Ind}_{KZ}^G \mathrm{PS}(\xi^s) = c\text{-Ind}_{IZ}^G \xi^s,$$

and note that the matrix $\begin{pmatrix} 0 & 1 \\ p & 0 \end{pmatrix} \in G$ normalizes I and induces the action of s on $T(\mathbf{F}_q)$.

The isomorphism (7.7.8) is perhaps the most fundamental example of a morphism in the category of smooth G -representations which is not obtained by inducing a morphism of KZ -representations, and plays an important (if sometimes only implicit) role in our analysis below.

7.7.9. *A union of two components, and extensions of Serre weights.* We begin by describing $\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}$ in the case when these two components meet in codimension 1 (i.e. in the case that σ, σ' admit a non-trivial extension). Up to switching σ and σ' and swapping the variables u and v or x and y , we may assume that the intersection is along the locus $u = 0$ (in \mathcal{C}_σ), which coincides with the locus $v' = 0$ (in $\mathcal{C}_{\sigma'}$). This locus is evidently disjoint from (indeed, equal to the complement in \mathcal{C}_σ of) $D(u)$, and so is contained in $D(v)$. Similarly, this locus is disjoint from $D(v')$ in $\mathcal{C}_{\sigma'}$, and so is contained in $D(u')$. Thus

$$(7.7.10) \quad \{D(u), D(v) \cup D(u'), D(v')\}$$

is an open cover of $\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}$.

Recall from (7.7.5) that

$$D(v) = [(\mathbf{A}_u^1 \times (\mathbf{A}^2 \setminus \{0\})_{x,vy})/(\mathbf{G}_m)_a],$$

while applying (7.7.4) to $\mathcal{C}_{\sigma'}$ shows that

$$D(u') = [(\mathbf{A}_{v'}^1 \times (\mathbf{A}^2 \setminus \{0\})_{x'/u',y'})/(\mathbf{G}_m)_{b'}].$$

The union $D(v) \cup D(u')$ is then equal to

$$[\{(u, x, vy, v') \mid (x, vy) \neq (0, 0), uv' = 0\}/(\mathbf{G}_m)_a],$$

where the locus $\{v' = 0\}$ is identified with $D(v)$ in the evident way, while the locus $\{u = 0\}$ is identified with $D(u')$ via $b' = a^{-1}, x'/u' = vy, y' = x$.

Since $\mathbf{A}^1 \times (\mathbf{A}^2 \setminus \{0\})$ has trivial Picard, we see that restriction to either $D(v)$ or $D(u')$ induces isomorphisms

$$\mathrm{Pic}(D(v) \cup D(u')) \xrightarrow{\sim} \mathrm{Pic}(D(v)), \mathrm{Pic}(D(u')) \xrightarrow{\sim} \mathbf{Z},$$

the two different identifications with \mathbf{Z} differing by a sign (since $b' = a^{-1}$). A consideration of the open cover (7.7.10) then shows that the product of the restriction maps induces an isomorphism

$$\begin{aligned} \mathrm{Pic}(\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}) &\xrightarrow{\sim} \mathrm{Pic}(D(v')) \times \mathrm{Pic}(D(v) \cup D(u')) \times \mathrm{Pic}(D(u)) \\ &\xrightarrow{\sim} \mathrm{Pic}(D(v')) \times \mathrm{Pic}(D(v)) \times \mathrm{Pic}(D(u)) = \mathbf{Z} \times \mathbf{Z} \times \mathbf{Z}. \end{aligned}$$

In terms of these coordinates, we find that the restriction

$$\mathrm{Pic}(\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}) \rightarrow \mathrm{Pic}(\mathcal{C}_\sigma) = \mathrm{Pic}(D(v) \times D(u)) = \mathbf{Z} \times \mathbf{Z}$$

is given by

$$(l, m, n) \mapsto (m, n),$$

while

$$\mathrm{Pic}(\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}) \rightarrow \mathrm{Pic}(\mathcal{C}_{\sigma'}) = \mathrm{Pic}(D(v') \times D(u')) = \mathbf{Z} \times \mathbf{Z}$$

is given by

$$(l, m, n) \mapsto (l, -m).$$

We now present some examples of morphisms of coherent sheaves on $\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}$ which incorporate some aspects of the preceding discussion. These examples will play an important role in our discussion of the functor \mathfrak{A} .

EXAMPLE 7.7.11. Consider the short exact sequence

$$0 \rightarrow \mathcal{I}_{\mathcal{C}_\sigma} \rightarrow \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}} \rightarrow \mathcal{O}_{\mathcal{C}_\sigma} \rightarrow 0;$$

here, of course, $\mathcal{I}_{\mathcal{C}_\sigma}$ denotes the ideal sheaf cutting out \mathcal{C}_σ as a closed substack of the union $\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}$. The middle term is the trivial element of $\text{Pic}(\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'})$, the right-hand term is the trivial element of $\text{Pic}(\mathcal{C}_\sigma)$, and the ideal sheaf is the element $(0, -1)$ in $\text{Pic}(\mathcal{C}_{\sigma'})$ (as one sees by considering its restriction to the members of the open cover (7.7.10): its restriction to $D(u)$ is the zero ideal, its restriction to $D(v) \cup D(u')$ is the ideal sheaf spanned by v' , which is an invertible sheaf on $D(u')$ on which the $(\mathbf{G}_m)_{b'}$ -action has been twisted by weight -1 , and its restriction to $D(v')$ is the structure sheaf). Thus we may rewrite this short exact sequence as

$$(7.7.12) \quad 0 \rightarrow \mathcal{O}_{\mathcal{C}_{\sigma'}}(0, -1) \xrightarrow{v'} \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}} \rightarrow \mathcal{O}_{\mathcal{C}_\sigma} \rightarrow 0.$$

(Note that v' is a well-defined global section of $\mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(0, -1, n)$ for any n — and that $\mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(0, -1, n)$ restricts to $\mathcal{O}_{\mathcal{C}_{\sigma'}}(0, 1)$ on $\mathcal{C}_{\sigma'}$ — so that multiplication by v' is a well-defined morphism $\mathcal{O}_{\mathcal{C}_{\sigma'}}(0, -1) \rightarrow \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}$.)

EXAMPLE 7.7.13. We also have the short exact sequence

$$0 \rightarrow \mathcal{I}_{\mathcal{C}_{\sigma'}} \rightarrow \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}} \rightarrow \mathcal{O}_{\mathcal{C}_{\sigma'}} \rightarrow 0,$$

which can be rewritten as the short exact sequence

$$(7.7.14) \quad 0 \rightarrow \mathcal{O}_{\mathcal{C}_\sigma}(-1, 0) \xrightarrow{u} \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}} \rightarrow \mathcal{O}_{\mathcal{C}_{\sigma'}} \rightarrow 0.$$

EXAMPLE 7.7.15. The element $x = y'$ is a well-defined global section of $\mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(0, 1, 0)$. Thus multiplication by y' , followed by restriction to $\mathcal{C}_{\sigma'}$, induces a morphism

$$(7.7.16) \quad \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}} \xrightarrow{y'} \mathcal{O}_{\mathcal{C}_{\sigma'}}(0, -1).$$

The composite of this morphism with the first non-trivial arrow of (7.7.12) (i.e. with the “multiplication by v' ” map) gives a morphism

$$(7.7.17) \quad \mathcal{O}_{\mathcal{C}_{\sigma'}}(0, -1) \xrightarrow{(v'y')} \mathcal{O}_{\mathcal{C}_{\sigma'}}(0, -1).$$

By assumption, σ and σ' admit a nonzero extension as $\text{GL}_2(\mathbf{F}_q)$ -representations, and in fact the corresponding Ext^1 is one-dimensional. We let U denote the corresponding non-split extension of σ by σ' , and we define $\mathfrak{A}(c\text{-Ind}_{KZ}^G U) := \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(-1, -1, 0)$. The short exact sequence

$$0 \rightarrow c\text{-Ind}_{KZ}^G \sigma' \rightarrow c\text{-Ind}_{KZ}^G U \rightarrow c\text{-Ind}_{KZ}^G \sigma \rightarrow 0$$

corresponds to the short exact sequence

$$(7.7.18) \quad 0 \rightarrow \mathcal{O}_{\mathcal{C}_{\sigma'}}(-1, 0) \xrightarrow{v'} \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(-1, -1, 0) \rightarrow \mathcal{O}_{\mathcal{C}_\sigma}(-1, 0) \rightarrow 0,$$

which is a twist of (7.7.12).

There is a well-known morphism²² $c\text{-Ind}_{KZ}^G U \rightarrow c\text{-Ind}_{KZ}^G \sigma'$ (which is *not* the compact induction of a morphism $U \rightarrow \sigma'$; indeed, there is no non-zero such morphism) with the property that the composite of this morphism with the obvious

²²We may choose a character ξ so that σ' is the socle of $\text{PS}(\xi)$ and such that the extension U embeds as a $\text{GL}_2(\mathbf{F}_q)$ -subrepresentation of $\text{PS}(\xi)$. Then σ' is the cosocle of $\text{PS}(\xi^s)$, and so we may form the composite morphism $c\text{-Ind}_{KZ}^G U \hookrightarrow c\text{-Ind}_{KZ}^G \text{PS}(\xi) \xrightarrow{(7.7.8)} c\text{-Ind}_{KZ}^G \text{PS}(\xi^s) \rightarrow c\text{-Ind}_{KZ}^G \sigma'$.

inclusion $c\text{-Ind}_{KZ}^G \sigma' \hookrightarrow c\text{-Ind}_{KZ}^G U$ is equal to the endomorphism T of $c\text{-Ind}_{KZ}^G \sigma'$. This well-known morphism corresponds under \mathfrak{A} to the twist

$$\mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(-1, -1, 0) \xrightarrow{y'} \mathcal{O}_{\mathcal{C}_{\sigma'}}(-1, 0)$$

of (7.7.16). Composing this with the first non-trivial arrow in (7.7.18) then gives the twist

$$\mathcal{O}_{\mathcal{C}_{\sigma'}}(-1, 0) \xrightarrow{(v'y')} \mathcal{O}_{\mathcal{C}_{\sigma'}}(-1, 0)$$

of (7.7.17). This does indeed correspond to the endomorphism T of $c\text{-Ind}_{KZ}^G \sigma'$, and so we have verified an instance of the functoriality of \mathfrak{A} .

Similarly, if we let V be the non-split extension of σ' by σ , then we define $\mathfrak{A}(c\text{-Ind}_{KZ}^G V) := \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(-1, 0, 0)$. Then the short exact sequence


$$0 \rightarrow c\text{-Ind}_{KZ}^G \sigma \rightarrow c\text{-Ind}_{KZ}^G V \rightarrow c\text{-Ind}_{KZ}^G \sigma' \rightarrow 0$$

corresponds to the short exact sequence

$$0 \rightarrow \mathcal{O}_{\mathcal{C}_\sigma}(-1, 0) \xrightarrow{u} \mathcal{O}_{\mathcal{C}_\sigma \cup \mathcal{C}_{\sigma'}}(-1, 0, 0) \rightarrow \mathcal{O}_{\mathcal{C}_{\sigma'}}(-1, 0) \rightarrow 0$$

obtained by twisting (7.7.14).

7.7.19. *Union of four components — case one (principal series)*. Typically (i.e. for the induction of a suitably generic k^\times -valued character), the parabolic induction $\text{PS}(\xi)$ of a character ξ of the torus in $\text{GL}_2(\mathbf{F}_q)$ can be denoted schematically as follows:



Here the various σ_i are the Jordan–Hölder factors of $\text{PS}(\xi)$, the rows of the diagram are the layers of its socle filtration (starting from the bottom of the diagram) or equivalently of the cosocle filtration (starting from the top of the diagram), and the lines indicate non-trivial extensions.

We write $\mathcal{C}_i := \mathcal{C}_{\sigma_i}$; our aim now is to describe $\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$, and then to describe its Picard group. We begin by noting that this union admits the open cover

$$(7.7.20) \quad \{D(u_0x_0), D(v_1y_1), D(u_2x_2), D(v_3y_3), D(v_0x_0) \cup D(u_1y_1), \\ D(v_1x_1) \cup D(u_2y_2), D(v_2x_2) \cup D(u_3y_3), D(v_3x_3) \cup D(u_0y_0), \\ D(v_0y_0) \cup D(u_1x_1) \cup D(v_2y_2) \cup D(u_3x_3)\}.$$

(Here and below we use the notation u, v, x, y, a, b as before, but add a subscript $i \in \{0, 1, 2, 3\}$ to indicate the component to which the variable pertains.) Here are explicit descriptions of these various open subsets:

$$\begin{aligned} D(u_ix_i) &= (\mathbf{G}_m)_{u_i/x_i} \times [\mathbf{A}_{v_i, y_i}^2 / (\mathbf{G}_m)_{b_i}]; \\ D(v_iy_i) &= [\mathbf{A}_{u_i, x_i}^2 / (\mathbf{G}_m)_{a_i}] \times (\mathbf{G}_m)_{v_i y_i}; \\ D(v_ix_i) &= \mathbf{A}_{u_i/x_i}^1 \times \mathbf{A}_{v_i y_i}^1, \\ D(u_{i+1}y_{i+1}) &= \mathbf{A}_{x_{i+1}/u_{i+1}}^1 \times \mathbf{A}_{v_{i+1}y_{i+1}}^1, \end{aligned}$$

and for $i = 0, 2$, we have

$$D(v_i x_i) \cup D(u_{i+1} y_{i+1}) = \{u_i/x_i, v_i y_i, v_{i+1} y_{i+1} \mid (u_i/x_i)(v_{i+1} y_{i+1}) = 0\},$$

with the identification $v_i y_i = x_{i+1}/u_{i+1}$, while for $i = 1, 3$ we have

$$D(v_i x_i) \cup D(u_{i+1} y_{i+1}) = \{u_i/x_i, v_i y_i, x_{i+1}/u_{i+1} \mid (v_i y_i)(x_{i+1}/u_{i+1}) = 0\},$$

with the identification $u_i/x_i = v_{i+1} y_{i+1}$; and finally

$$\begin{aligned} D(v_0 y_0) \cup D(u_1 x_1) \cup D(v_2 y_2) \cup D(u_3 x_3) \\ = \{ \{ (u_0, x_0, v_1, y_3, v_0 y_0) \mid u_0 v_1 = x_0 y_3 = 0, v_0 y_0 \neq 0 \} / (\mathbf{G}_m)_{a_0} \}, \end{aligned}$$

where we have the identifications

$$\begin{aligned} x_0 = y_1, x_2 = y_3, u_0 = v_3, u_2 = v_1, \\ v_0 y_0 = x_1/u_1 = (v_2 y_2)^{-1} = u_3/x_3, \\ a_0 = b_1^{-1} = a_2^{-1} = b_3. \end{aligned}$$

Using the open cover (7.7.20), we compute

$$(7.7.21) \quad \Gamma(\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}) = k[v_1 y_1, v_3 y_3] / (v_1 y_1 v_3 y_3).$$

(To give slightly more detail: the function $v_1 y_1$ is well-defined on \mathcal{C}_1 , and extends by zero over the union of all four components. Similarly, the function $v_3 y_3$ is well-defined on \mathcal{C}_3 , and again extends by zero.)

One can show that there is an embedding

$$(7.7.22) \quad \text{Pic}(\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3) \hookrightarrow \mathbf{Z} \times \mathbf{Z} \times \mathbf{Z} \times \mathbf{Z} \times \mathbf{Z},$$

where the factors are given by $\text{Pic}(D(v_1 y_1))$, $\text{Pic}(D(v_3 y_3))$, $\text{Pic}(D(v_0 y_0) \cup D(u_1 x_1) \cup D(v_2 y_2) \cup D(u_3 x_3))$, $\text{Pic}(D(u_0 x_0))$, and $\text{Pic}(D(u_2 x_2))$, respectively, with the morphisms being defined by restriction. The image of this embedding is equal to

$$(7.7.23) \quad \{(j, k, l, m, n) \mid j - k + m - n = 0\};$$

in particular, $\text{Pic}(\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3) \cong \mathbf{Z}^4$.


The simplest way to fix our choice of coordinates precisely is to describe the corresponding embedding

$$(7.7.24) \quad \text{Pic}(\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3) \hookrightarrow \text{Pic}(\mathcal{C}_0) \times \text{Pic}(\mathcal{C}_1) \times \text{Pic}(\mathcal{C}_2) \times \text{Pic}(\mathcal{C}_3);$$

it is given by

$$(7.7.25) \quad (j, k, l, m, n) \mapsto ((l, m), (j, -l), (-l, n), (k, l)).$$

EXAMPLE 7.7.26. We have the following diagram of ideal subsheaves of $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$:



where the sheaf labels on the arrows are indicate the corresponding cokernel. Thus we may describe $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$ via the following extension diagram:

$$(7.7.27) \quad \begin{array}{ccc} & \mathcal{O}_{\mathcal{C}_3} & \\ & \swarrow \quad \searrow & \\ \mathcal{O}_{\mathcal{C}_2}(-1,0) & & \mathcal{O}_{\mathcal{C}_0}(-1,0) \\ & \swarrow \quad \searrow & \\ & \mathcal{O}_{\mathcal{C}_1} & \end{array}$$

Similarly, we have the diagram

$$\begin{array}{c} \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3} \\ \uparrow \mathcal{O}_{\mathcal{C}_1} \\ \mathcal{I}_{\mathcal{C}_1} \xlongequal{\quad} \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1} + \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2} \\ \swarrow \mathcal{O}_{\mathcal{C}_0}(-1,0) \quad \searrow \mathcal{O}_{\mathcal{C}_2}(-1,0) \\ \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1} \quad \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2} \\ \swarrow \quad \searrow \\ \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2} \xlongequal{\quad} \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1} \cap \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2} \xlongequal{\quad} v_3 y_3 \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3} \xleftarrow{\sim} \mathcal{O}_{\mathcal{C}_3} \end{array}$$

which allows us to describe $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$ via the extension diagram

$$(7.7.28) \quad \begin{array}{ccc} & \mathcal{O}_{\mathcal{C}_1} & \\ & \swarrow \quad \searrow & \\ \mathcal{O}_{\mathcal{C}_0}(-1,0) & & \mathcal{O}_{\mathcal{C}_2}(-1,0) \\ & \swarrow \quad \searrow & \\ & \mathcal{O}_{\mathcal{C}_3} & \end{array}$$

EXAMPLE 7.7.29. We also have the following diagrams of ideal subsheaves of $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$, namely

$$\begin{array}{c} \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3} \\ \uparrow \mathcal{O}_{\mathcal{C}_0} \\ \mathcal{I}_{\mathcal{C}_0} \xlongequal{\quad} \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_3} + \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1} \\ \swarrow \mathcal{O}_{\mathcal{C}_3}(0,1) \quad \searrow \mathcal{O}_{\mathcal{C}_1}(0,-1) \\ \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_3} \quad \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1} \\ \swarrow \quad \searrow \\ \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_3} \xlongequal{\quad} \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_3} \cap \mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_1} \xleftarrow{\sim} \mathcal{O}_{\mathcal{C}_2}(-2,0) \end{array}$$

and

$$\begin{array}{c}
 \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3} \\
 \uparrow \mathcal{O}_{\mathcal{C}_2} \\
 \mathcal{I}_{\mathcal{C}_2} \xleftarrow{\mathcal{O}_{\mathcal{C}_1(0,1)}} \mathcal{I}_{\mathcal{C}_2} \xrightarrow{\mathcal{O}_{\mathcal{C}_3(0,-1)}} \mathcal{I}_{\mathcal{C}_2} \\
 \text{-----} \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2} + \mathcal{I}_{\mathcal{C}_2 \cup \mathcal{C}_3} \\
 \swarrow \quad \searrow \\
 \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2} \quad \mathcal{I}_{\mathcal{C}_2 \cup \mathcal{C}_3} \\
 \swarrow \quad \searrow \\
 \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3} \text{-----} \mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2} \cap \mathcal{I}_{\mathcal{C}_2 \cup \mathcal{C}_3} \xleftarrow{\sim} \mathcal{O}_{\mathcal{C}_0}(-2, 0)
 \end{array}$$

which allows us to describe $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$ via the extension diagrams

(7.7.30)

$$\begin{array}{ccc}
 & \mathcal{O}_{\mathcal{C}_0} & \\
 & \swarrow \quad \searrow & \\
 \mathcal{O}_{\mathcal{C}_3}(0, 1) & & \mathcal{O}_{\mathcal{C}_1}(0, -1) \\
 & \swarrow \quad \searrow & \\
 & \mathcal{O}_{\mathcal{C}_2}(-2, 0) &
 \end{array}$$

and

(7.7.31)

$$\begin{array}{ccc}
 & \mathcal{O}_{\mathcal{C}_2} & \\
 & \swarrow \quad \searrow & \\
 \mathcal{O}_{\mathcal{C}_1}(0, 1) & & \mathcal{O}_{\mathcal{C}_3}(0, -1) \\
 & \swarrow \quad \searrow & \\
 & \mathcal{O}_{\mathcal{C}_0}(-2, 0) &
 \end{array}$$

We now explain the representation-theoretic interpretation of the above calculations. We write $U := \text{PS}(\xi)$, and recall the description

(7.7.32)

$$\begin{array}{ccc}
 & \sigma_3 & \\
 & \swarrow \quad \searrow & \\
 \sigma_2 & & \sigma_0 \\
 & \swarrow \quad \searrow & \\
 & \sigma_1 &
 \end{array}$$

of U in terms of its Jordan–Hölder factors. We then define

$$\mathfrak{A}(c\text{-Ind}_{KZ}^G U) := \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, -1, 0, 0, 0),$$

where the twisting is understood to be in terms of the embedding (7.7.22); note that $-1 - (-1) + 0 = 0 = 0$, so that the constraint of (7.7.23) is satisfied. If we form

the corresponding twist of the diagram (7.7.27), and take into account the formula of (7.7.25), we obtain the diagram

$$(7.7.33) \quad \begin{array}{ccc} & \mathcal{O}_{\mathcal{C}_3}(-1, 0) & \\ & \swarrow \quad \searrow & \\ \mathcal{O}_{\mathcal{C}_2}(-1, 0) & & \mathcal{O}_{\mathcal{C}_0}(-1, 0) \\ & \swarrow \quad \searrow & \\ & \mathcal{O}_{\mathcal{C}_1}(-1, 0) & \end{array}$$

which shows that our definition is compatible with the functoriality of \mathfrak{A} and diagram (7.7.32).

If we induce the character ξ^s (the non-trivial Weyl group conjugate of ξ) instead, we obtain a representation $U^s := \text{PS}(\xi^s)$, which can be described by the extension diagram

$$(7.7.34) \quad \begin{array}{ccc} & \sigma_1 & \\ & \swarrow \quad \searrow & \\ \sigma_0 & & \sigma_2 \\ & \swarrow \quad \searrow & \\ & \sigma_3 & \end{array}$$

Although U and U^s are *not* isomorphic representations of $\text{GL}_2(\mathbf{F}_{p^2})$, their compact inductions to $\text{GL}_2(\mathbf{Q}_{p^2})$ are isomorphic (via the isomorphism (7.7.8)), and so we also have that $\mathfrak{A}(c\text{-Ind}_{KZ}^G U^s) = \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, -1, 0, 0, 0)$. Note that if we twist diagram (7.7.28) by $(-1, -1, 0, 0, 0)$, we obtain the diagram


$$(7.7.35) \quad \begin{array}{ccc} & \mathcal{O}_{\mathcal{C}_1}(-1, 0) & \\ & \swarrow \quad \searrow & \\ \mathcal{O}_{\mathcal{C}_0}(-1, 0) & & \mathcal{O}_{\mathcal{C}_2}(-1, 0) \\ & \swarrow \quad \searrow & \\ & \mathcal{O}_{\mathcal{C}_3}(-1, 0) & \end{array}$$

This verifies the functoriality of \mathfrak{A} when applied to the diagram (7.7.34).

The composite $c\text{-Ind}_{KZ}^G \sigma_1 \hookrightarrow c\text{-Ind}_{KZ}^G U \xrightarrow{\sim} c\text{-Ind}_{KZ}^G U^s \rightarrow c\text{-Ind}_{KZ}^G \sigma_1$ is known to coincide with the Hecke operator T . A consideration of the diagrams in Example 7.7.26 shows that this map is given by multiplication by $v_1 y_1$, which is indeed $\mathfrak{A}(T)$. Related to this computation, we note that (7.7.21) is compatible with the full faithfulness of \mathfrak{A} and the known structure of the endomorphism ring of $c\text{-Ind}_{KZ}^G U$.


As is well-known, by reducing lattices in integral principal series, we can also obtain representations V and W of $\mathrm{GL}_2(\mathbf{F}_q)$, which can be described by the extension diagrams

(7.7.36)



and

(7.7.37)



respectively.

We define


$$\mathfrak{A}(c\text{-Ind}_{KZ}^G V) := \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, -1, -1, 0, 0)$$

and

$$\mathfrak{A}(c\text{-Ind}_{KZ}^G W) := \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, -1, 1, 0, 0).$$


Then twisting the diagrams (7.7.30) and (7.7.31) by $(-1, -1, -1, 0, 0)$ and $(-1, -1, 1, 0, 0)$ respectively, we obtain diagrams

(7.7.38)




and

(7.7.39)



confirming the functoriality of \mathfrak{A} when applied to (7.7.36) and (7.7.37).

7.7.40. *Union of four components — case two (cuspidal)*. The computations in this case follow similar lines to those in the principal series case. For a tame cuspidal type, there is no preferred lattice, though. Rather, the reduction (generically) contains four weights σ_0 , σ_1 , σ_2 , and σ_3 , each with multiplicity one, and one can arrange for them to appear in the extension diagram



or more generally in any cyclic permutation of this diagram.

For an appropriate choice of labelling of the weights, the union $\mathcal{C}_{\sigma_0} \cup \mathcal{C}_{\sigma_1} \cup \mathcal{C}_{\sigma_2} \cup \mathcal{C}_{\sigma_3}$ admits the open cover

$$\begin{aligned}
 (7.7.41) \quad & \{D(v_0x_0), D(u_1y_1), D(v_2x_2), D(u_3y_3), D(v_0y_0) \cup D(u_1x_1), \\
 & D(v_1y_1) \cup D(u_2x_2), D(v_2y_2) \cup D(u_3x_3), D(v_3y_3) \cup D(u_0x_0), \\
 & D(u_0y_0) \cup D(v_1x_1) \cup D(u_2y_2) \cup D(v_3x_3)\}.
 \end{aligned}$$

We have

$$\begin{aligned}
 D(u_iy_i) &= \mathbf{A}_{x_i/u_i}^1 \times \mathbf{A}_{v_iy_i}^1; \\
 D(v_ix_i) &= \mathbf{A}_{u_i/x_i}^1 \times \mathbf{A}_{v_iy_i}^1; \\
 D(u_ix_i) &= (\mathbf{G}_m)_{x_i/u_i} \times [\mathbf{A}_{v_iy_i}^2 / (\mathbf{G}_m)_{b_i}], \\
 D(v_iy_i) &= [\mathbf{A}_{u_i/x_i}^2 / (\mathbf{G}_m)_{a_i}] \times (\mathbf{G}_m)_{v_iy_i},
 \end{aligned}$$

and for $i = 0, 2$, we have

$$D(v_iy_i) \cup D(u_{i+1}x_{i+1}) = [\{(u_i, x_i, v_iy_i, y_{i+1}) \mid v_iy_i \neq 0, x_iy_{i+1} = 0\} / (\mathbf{G}_m)_{a_i}],$$

with the identifications $u_i = v_{i+1}$, $v_iy_i = u_{i+1}/x_{i+1}$, and $a_i = b_{i+1}$, while for $i = 1, 3$ we have

$$D(v_iy_i) \cup D(u_{i+1}x_{i+1}) = [\{(u_i, x_i, v_iy_i, v_{i+1}) \mid v_iy_i \neq 0, u_iv_{i+1} = 0\} / (\mathbf{G}_m)_{a_i}],$$

with the identifications $x_i = y_{i+1}$, $v_iy_i = x_{i+1}/u_{i+1}$, and $a_i = b_{i+1}^{-1}$; and finally

$$D(u_0y_0) \cup D(v_1x_1) \cup D(u_2y_2) \cup D(v_3x_3) = \{(v_0y_0, v_1y_1, v_2y_2, v_3y_3) \mid (v_0y_0)(v_2y_2) = (v_1y_1)(v_3y_3) = 0\},$$

with the identifications $x_0/u_0 = v_3y_3$, $u_1/x_1 = v_0y_0$, $x_2/u_2 = v_1y_1$, and $u_3/x_3 = v_2y_2$.

Using the open cover (7.7.41), we compute

$$(7.7.42) \quad \Gamma(\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}) = k.$$

We also compute that

$$\text{Pic}(\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3) = \mathbf{Z} \times \mathbf{Z} \times \mathbf{Z} \times \mathbf{Z},$$

where the factors are given by $\text{Pic}(D(v_2y_2) \cup D(u_3x_3))$, $\text{Pic}(D(v_1y_1) \cup D(u_2x_2))$, $\text{Pic}(D(v_0y_0) \cup D(u_1x_1))$, and $\text{Pic}(D(v_3y_3) \cup D(u_0x_0))$, respectively. We fix our choice of coordinates by describing the embedding

$$(7.7.43) \quad \text{Pic}(\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3) \hookrightarrow \text{Pic}(\mathcal{C}_0) \times \text{Pic}(\mathcal{C}_1) \times \text{Pic}(\mathcal{C}_2) \times \text{Pic}(\mathcal{C}_3);$$

it is given by

$$(7.7.44) \quad (k, l, m, n) \mapsto ((m, n), (-l, m), (k, l), (-n, k)).$$

EXAMPLE 7.7.45. We have the following diagrams of ideal subsheaves of $\mathcal{O}_{C_0 \cup C_1 \cup C_2 \cup C_3}$:

$$\begin{array}{c}
 \mathcal{O}_{C_0 \cup C_1 \cup C_2 \cup C_3} \\
 \uparrow \mathcal{O}_{C_0} \\
 \mathcal{I}_{C_0} \xleftarrow{\mathcal{O}_{C_3}(-1,0)} \mathcal{I}_{C_0} \xrightarrow{\mathcal{O}_{C_1}(0,1)} \mathcal{I}_{C_0 \cup C_3} \\
 \mathcal{I}_{C_0 \cup C_3} \xleftarrow{\quad} \mathcal{I}_{C_0} \xrightarrow{\quad} \mathcal{I}_{C_0 \cup C_3} \\
 \mathcal{I}_{C_0 \cup C_1 \cup C_3} \xleftarrow{\quad} \mathcal{I}_{C_0 \cup C_3} \xrightarrow{\quad} \mathcal{I}_{C_0 \cup C_1} \xleftarrow{\sim} \mathcal{O}_{C_2}(-1, -1)
 \end{array}$$


$\mathcal{I}_{C_0} \xrightarrow{\quad} \mathcal{I}_{C_0 \cup C_3} + \mathcal{I}_{C_0 \cup C_1}$
 $\mathcal{I}_{C_0 \cup C_1 \cup C_3} \xrightarrow{\quad} \mathcal{I}_{C_0 \cup C_3} \cap \mathcal{I}_{C_0 \cup C_1} \xleftarrow{\sim} \mathcal{O}_{C_2}(-1, -1)$

$$\begin{array}{c}
 \mathcal{O}_{C_0 \cup C_1 \cup C_2 \cup C_3} \\
 \uparrow \mathcal{O}_{C_1} \\
 \mathcal{I}_{C_1} \xleftarrow{\mathcal{O}_{C_0}(-1,0)} \mathcal{I}_{C_1} \xrightarrow{\mathcal{O}_{C_2}(0,-1)} \mathcal{I}_{C_1 \cup C_2} \\
 \mathcal{I}_{C_1 \cup C_2} \xleftarrow{\quad} \mathcal{I}_{C_1} \xrightarrow{\quad} \mathcal{I}_{C_1 \cup C_2} \\
 \mathcal{I}_{C_0 \cup C_1 \cup C_2} \xleftarrow{\quad} \mathcal{I}_{C_1 \cup C_2} \xrightarrow{\quad} \mathcal{I}_{C_0 \cup C_1} \xleftarrow{\sim} \mathcal{O}_{C_3}(-1, 1)
 \end{array}$$

$\mathcal{I}_{C_1} \xrightarrow{\quad} \mathcal{I}_{C_1 \cup C_2} + \mathcal{I}_{C_0 \cup C_1}$
 $\mathcal{I}_{C_0 \cup C_1 \cup C_2} \xrightarrow{\quad} \mathcal{I}_{C_0 \cup C_1} \cap \mathcal{I}_{C_1 \cup C_2} \xleftarrow{\sim} \mathcal{O}_{C_3}(-1, 1)$


$$\begin{array}{c}
 \mathcal{O}_{C_0 \cup C_1 \cup C_2 \cup C_3} \\
 \uparrow \mathcal{O}_{C_2} \\
 \mathcal{I}_{C_2} \xleftarrow{\mathcal{O}_{C_1}(-1,0)} \mathcal{I}_{C_2} \xrightarrow{\mathcal{O}_{C_3}(0,1)} \mathcal{I}_{C_2 \cup C_3} \\
 \mathcal{I}_{C_2 \cup C_3} \xleftarrow{\quad} \mathcal{I}_{C_2} \xrightarrow{\quad} \mathcal{I}_{C_2 \cup C_3} \\
 \mathcal{I}_{C_1 \cup C_2 \cup C_3} \xleftarrow{\quad} \mathcal{I}_{C_2 \cup C_3} \xrightarrow{\quad} \mathcal{I}_{C_1 \cup C_2} \xleftarrow{\sim} \mathcal{O}_{C_0}(-1, -1)
 \end{array}$$

$\mathcal{I}_{C_2} \xrightarrow{\quad} \mathcal{I}_{C_2 \cup C_3} + \mathcal{I}_{C_1 \cup C_2}$
 $\mathcal{I}_{C_1 \cup C_2 \cup C_3} \xrightarrow{\quad} \mathcal{I}_{C_1 \cup C_2} \cap \mathcal{I}_{C_2 \cup C_3} \xleftarrow{\sim} \mathcal{O}_{C_0}(-1, -1)$




The following representations (described via their extension diagrams) arise as the reduction of lattices in cuspidal types:


(7.7.46)



(7.7.47)




(7.7.48)



and

(7.7.49)



The corresponding values of \mathfrak{A} are $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(0, 1, -1, 0)$, $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, 1, 0, 0)$, $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, 0, 0, 1)$, and $\mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(0, 0, -1, 1)$. If we twist each of the diagrams in Example 7.7.45 by the corresponding twist (and compute using the formula (7.7.44)), we obtain an extension diagram constituted from the $\mathfrak{A}(\sigma_i)$ in the appropriate order; this confirms the functoriality of \mathfrak{A} on the various diagrams (7.7.46), (7.7.47), (7.7.48), and (7.7.49) in turn.

7.7.50. *A non-compactly induced example.* The isomorphism $c\text{-Ind}_{KZ}^G \overline{\text{PS}(\xi)} \xrightarrow{\sim} c\text{-Ind}_{KZ}^G \overline{\text{PS}(\xi^{(s)})}$ allows us to construct an interesting representation. (Here $\text{PS}(\xi)$ is the representation denoted U in Section 7.7.19, and $\overline{\text{PS}(\xi)}$ is the representation denoted U^s .) If we let σ_ξ denote the socle of $\overline{\text{PS}(\xi)}$, and $\sigma_{\xi^{(s)}}$ denote the socle of $\overline{\text{PS}(\xi^{(s)})}$, then there are embeddings

$$c\text{-Ind}_{KZ}^G \sigma_\xi \hookrightarrow c\text{-Ind}_{KZ}^G \overline{\text{PS}(\xi)}$$

and

$$c\text{-Ind}_{KZ}^G \sigma_{\xi^{(s)}} \hookrightarrow c\text{-Ind}_{KZ}^G \overline{\text{PS}(\xi^{(s)})} \xrightarrow{\sim} c\text{-Ind}_{KZ}^G \overline{\text{PS}(\xi)},$$

inducing

$$c\text{-Ind}_{KZ}^G \sigma_\xi \oplus c\text{-Ind}_{KZ}^G \sigma_{\xi^{(s)}} \hookrightarrow c\text{-Ind}_{KZ}^G \overline{\text{PS}(\xi)}.$$

Define Π to be the cokernel of this embedding.

We can compute $\mathfrak{A}(\Pi)$ using the calculations of Section 7.7.19. Indeed, noting that σ_ξ and $\sigma_{\xi^{(s)}}$ can be taken to be σ_1 and σ_3 in the notation of that discussion, we find that $\mathfrak{A}(\Pi)$ is the cokernel of the embedding

$$\mathcal{O}_{\mathcal{C}_1}(-1, 0) \oplus \mathcal{O}_{\mathcal{C}_3}(-1, 0) \hookrightarrow \mathcal{O}_{\mathcal{C}_0 \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}(-1, -1, 0, 0, 0).$$

The image of this embedding is identified with the twist by $(-1, -1, 0, 0, 0)$ of the ideal sheaf $\mathcal{I}_{\mathcal{C}_0 \cup \mathcal{C}_2}$, and one finds that $\mathfrak{A}(\Pi)$ is a certain invertible sheaf on $\mathcal{C}_0 \cup \mathcal{C}_2$. (In particular, it is *not* Cohen–Macaulay.)

7.7.51. *An example — weight cycling.* In the 2006 AIM meeting, the participants found an argument which has come to be known as “weight cycling”. Here it is in its original form. (For more details, the interested reader could consult the contemporary reports <https://aimath.org/WWN/padicmodularity/emertonreport.pdf> and https://www.ma.ic.ac.uk/~buzzard/maths/research/notes/notes_on_mod_p_local_langlands.pdf.)

PROPOSITION 7.7.52. *Suppose that π' is a representation of $\text{GL}_2(\mathbf{Q}_p)$ with the following properties:*

- (1) *The constituents of the K -socle of π' are among the constituents of the reduction of some sufficiently generic cuspidal type τ .*
- (2) *Each constituent of the K -socle of π' appears with multiplicity one.*
- (3) *Each constituent of the K -socle of π' is supersingular.*

Then there is a (necessarily irreducible) subrepresentation π of π' with the following properties:

- (a) *π admits a central character.*
- (b) *Every G -subrepresentation of π' contains π .*
- (c) *The K -socle of π coincides with the K -socle of π' .*
- (d) *All four constituents of the reduction of τ actually appear in the common K -socle of π' and π .*

PROOF. Any G -subrepresentation of π' contains at least one of the constituents of the K -socle of π' , and thus contains the G -subrepresentation that this constituent generates. Thus if we show that the G -subrepresentation π generated by any one of these constituents contains *all* of these constituents — that is, if we show that π satisfies (c) — then we will also have shown that π is independent of the initial choice of generating constituent, and thus that π also satisfies condition (b). (And the fact that π satisfies (b) implies that it is irreducible.)

So we turn to proving that π does indeed satisfy (c). Along the way, we will also show that it satisfies (a) and (d).

Let σ_0 denote the chosen constituent of the reduction of τ that appears in the K -socle of π' and generates π . Because σ_0 appears in the K -socle of π with multiplicity one, the centre Z of G acts on σ_0 , and hence also on the representation π that it generates, through a character. This proves (a).

Recalling that σ_0 is one of the weights appearing in the reduction of τ , we then label the remaining weights in the reduction of τ as $\sigma_1, \sigma_2, \sigma_3$. Since τ has a central character as a representation of K , each σ_i has a central character as a representation of K , and these characters all coincide (they are just the reduction of the central character of τ). In particular the central character of π is compatible with the central character of the σ_i (i.e. the restriction of the former to $K \cap Z$ coincides with the latter), and so we can and do extend each σ_i to a representation of KZ by having Z act through the central character of π .

We can then choose our labels so that for each $i = 0, \dots, 3$, we have an embedding (originally constructed by Paškūnas)

$$(7.7.53) \quad (c\text{-Ind}_{KZ}^G \sigma_{i+1} \oplus c\text{-Ind}_{KZ}^G \sigma'_{i+1}) \hookrightarrow c\text{-Ind}_{KZ}^G \sigma_i/T,$$

where $i+1$ is to be computed mod 4, and σ'_{i+1} is another Serre weight which is not isomorphic to any of the σ_j ($j = 0, \dots, 3$). (See Remark 7.7.54 below for a slight elaboration on this.)

Let $\pi'(\sigma_i)$ denote the cokernel of (7.7.53). Then a basic fact is that $\pi'(\sigma_i) = \pi'(\sigma_i^{(s)})$, where $\sigma_i^{(s)}$ denotes the Serre weight whose highest weight is obtained from that of σ_i by applying the Weyl group involution s . In particular, $\pi'(\sigma_i)$ is not only a quotient of $c\text{-Ind}_{KZ}^G \sigma_i$, but is also a quotient of $c\text{-Ind}_{KZ}^G \sigma_i^{(s)}$. Another basic fact is that none of the $\sigma_i^{(s)}$ (for $i = 0, \dots, 3$) is isomorphic to any of the σ_j ; i.e. none of the former representations are constituents of the reduction of τ . (Again, see Remark 7.7.54 below for an elaboration on this.) By our assumption on the K -socle of π' , this means that any morphism $c\text{-Ind}_{KZ}^G \sigma_i^{(s)} \rightarrow \pi'$ (for any value of i) must vanish, and consequently any morphism $\pi'(\sigma_i) \rightarrow \pi'$ must vanish. Since $\pi \subseteq \pi'$, the same applies with π in place of π' .

Now consider the morphism $c\text{-Ind}_{KZ}^G \sigma_0 \rightarrow \pi$ induced by the inclusion $\sigma_0 \subseteq \pi$. The supersingularity assumption implies (indeed, more-or-less means) that this factors through $c\text{-Ind}_{KZ}^G \sigma_0/T$. Since it is non-zero (it contains σ_0 in its image) it must not factor through $\pi'(\sigma_0)$, and so it must not vanish identically on the domain of (7.7.53). However, it *must* vanish on the second such summand (by our assumption on the K -socle of π'), and hence it is non-zero on the first summand. Thus we obtain a non-zero morphism $c\text{-Ind}_{KZ}^G \sigma_{i+1} \rightarrow \pi$. Equivalently, we obtain an embedding $\sigma_{i+1} \hookrightarrow \pi$.

Continuing in this manner, we cycle through all the weights σ_i , and find that the π' contains all of them in its K -socle, and that in fact they are all contained in π . This completes the proof of the proposition. \square

REMARK 7.7.54. Just to say slightly more about the facts regarding Serre weights that we used above: if χ_i denotes the highest weight of σ_i , then the weights in the reduction of $\text{PS}(\chi_i)$ are $\sigma_i, \sigma_{i+1}, \sigma'_{i+1}$, and $\sigma_i^{(s)}$. And the reduction of a (sufficiently generic) cuspidal type and of a (sufficiently generic) principal series type have either no or two weights in common. So, since τ contains σ_i and σ_{i+1} in its reduction, it does not contain σ'_{i+1} or $\sigma_i^{(s)}$.

REMARK 7.7.55. The original aim of the participants at AIM in 2006 had been to apply Proposition 7.7.52 with π' equal to the \mathfrak{m} -torsion in the completed cohomology of a Shimura curves (where \mathfrak{m} denotes a maximal ideal of an appropriate Hecke algebra). In this case the hypotheses of the proposition were known to hold under appropriate genericity hypotheses by [Gee11]. Much more recently it has been proved that in this case we even have $\pi' = \pi$ (see [BHH⁺21, Cor. 1.3.10]).

Shortly after the 2006 AIM meeting, Breuil and Paškūnas [BP12] gave a direct construction of an irreducible admissible representation π as in Proposition 7.7.52. In fact, they exhibited infinitely many pairwise non-isomorphic such representations.

The proof of Proposition 7.7.52 shows that there is actually a scalar invariant attached to π' — or, really, to its irreducible subrepresentation π . Namely, if we begin with the inclusion $\sigma_0 \subseteq \pi$ and cycle four times, we return to an embedding $\sigma_0 \hookrightarrow \pi$, which must then (by our multiplicity one assumption) be equal to a non-zero scalar multiple of the original embedding. In order to make this well-defined, one has to pin down the embeddings (7.7.53), but with a bit of care one can do that (we explain below how to do this after applying \mathfrak{A}), and one obtains a scalar invariant of π .

More precisely, we choose a non-zero scalar $\lambda = \lambda_0 \in k^\times$, write $\lambda_i = 1$ if $i = 1, 2, 3$, and consider the two step complex

$$(7.7.56) \quad \bigoplus_{i=0}^3 (c\text{-Ind}_{KZ}^G \sigma_i \oplus c\text{-Ind}_{KZ}^G \sigma'_i) \rightarrow \bigoplus_{i=0}^3 c\text{-Ind}_{KZ}^G \sigma_i/T,$$

where the arrow is defined on the summand $c\text{-Ind}_{KZ}^G \sigma_i \oplus c\text{-Ind}_{KZ}^G \sigma'_i$ via

$$\begin{aligned} c\text{-Ind}_{KZ}^G \sigma_i \oplus c\text{-Ind}_{KZ}^G \sigma'_i &\ni (v, w) \\ &\mapsto (v + w, -\lambda_i v \bmod T) \in c\text{-Ind}_{KZ}^G \sigma_{i-1}/T \oplus c\text{-Ind}_{KZ}^G \sigma_i/T; \end{aligned}$$

here we regard v and w as elements of $c\text{-Ind}_{KZ}^G \sigma_{i-1}/T$ via (7.7.53).

LEMMA 7.7.57. *The arrow of (7.7.56) is injective.*

PROOF. If the kernel is non-zero, then it has a non-zero K -socle, and so there is a non-zero morphism from $c\text{-Ind}_{KZ}^G \sigma'$ to this kernel, for some weight σ' . But if $c\text{-Ind}_{KZ}^G \sigma' \rightarrow \bigoplus_{i=0}^3 (c\text{-Ind}_{KZ}^G \sigma_i \oplus c\text{-Ind}_{KZ}^G \sigma'_i)$ is non-zero, then σ' is equal to one of the σ_i or σ'_i (we are assuming that these weights are sufficiently generic), and the morphism is just an embedding into the corresponding summand. The lemma then follows from the fact that the arrow of (7.7.56) is injective on each summand in its domain. \square

Let Π_λ denote the cokernel of (7.7.56). Then giving a morphism $\Pi_\lambda \rightarrow \pi$, for some other representation π , is equivalent to giving a copy of σ_0 (or indeed, any one of the σ_i) in π , such that the weight cycling of Proposition 7.7.52 “works” in π , and such that the corresponding scalar is equal to λ .

We can compute the coherent sheaf (or rather, complex of coherent sheaves) associated to Π_λ . The answer depends on λ . Indeed, the geometry allows us to normalize the embeddings (7.7.53). Namely, each of the arrows $c\text{-Ind}_{KZ}^G \sigma_{i+1} \rightarrow c\text{-Ind}_{KZ}^G \sigma_i/T$ and $c\text{-Ind}_{KZ}^G \sigma'_{i+1} \rightarrow c\text{-Ind}_{KZ}^G \sigma_i/T$ corresponds, on the coherent sheaf side, to an arrow $\mathcal{O}_{\mathcal{C}_{\sigma_{i+1}}}(-1, 0) \rightarrow \mathcal{O}_{\mathcal{C}_\sigma}(-1)/(vy)$, respectively $\mathcal{O}_{\mathcal{C}_{\sigma'_{i+1}}}(-1, 0) \rightarrow \mathcal{O}_{\mathcal{C}_\sigma}(-1)/(vy)$, and we can take these to simply be the restriction maps induced by the inclusions $\mathcal{C}_\sigma^{vy=0} \hookrightarrow \mathcal{C}_{\sigma_{i+1}}, \mathcal{C}_{\sigma'_{i+1}}$. Applying \mathfrak{A} to the morphism (7.7.56), with the morphisms being interpreted in the manner just described, one can (with a little effort) explicitly compute the complexes $\mathfrak{A}(\Pi_\lambda)$.

We summarize the results of this computation in the following statement, which shows in particular that the value $\lambda = 1$ is distinguished.

CONCLUSION 7.7.58.

- (1) $H^0(\mathfrak{A}(\Pi_1)) = \text{skyscraper supported at } x$, where x is the closed point obtained as the intersection of the components \mathcal{C}_{σ_i} ($i = 0, \dots, 3$), while $H^{-1}(\mathfrak{A}(\Pi_1))$ is a line bundle supported on $\bigcup_{i=0}^3 \mathcal{C}_{\sigma_i} \cup \mathcal{C}_{\sigma'_i}$.
- (2) If $\lambda \neq 1$, then $H^0(\mathfrak{A}(\Pi_\lambda)) = 0$, while $H^{-1}(\mathfrak{A}(\Pi_\lambda))$ is a torsion-free sheaf supported on $\mathcal{C} := \bigcup_{i=0}^3 \mathcal{C}_{\sigma_i} \cup \mathcal{C}_{\sigma'_i}$. More precisely, it is the pushforward to \mathcal{C} of a line bundle on $\mathcal{C} \setminus \{x\}$ which depends on the parameter λ .

7.7.59. *Irreducible quotients.* Consider the short exact sequence

$$0 \rightarrow \mathcal{K} \rightarrow \Pi_\lambda \rightarrow \pi \rightarrow 0$$

(where π is an irreducible admissible representation as above, and \mathcal{K} is defined to be the kernel of the surjection). Passing to the (now merely hypothetical) sheaves, and bearing in mind that \mathfrak{A} has amplitude $[-1, 0]$ (see Remark 6.1.27), we obtain a long exact sequence

$$\begin{aligned} 0 \rightarrow H^{-1}(\mathfrak{A}(\mathcal{K})) \rightarrow H^{-1}(\mathfrak{A}(\Pi_\lambda)) \rightarrow H^{-1}(\mathfrak{A}(\pi)) \\ \rightarrow H^0(\mathfrak{A}(\mathcal{K})) \rightarrow H^0(\mathfrak{A}(\Pi_\lambda)) \rightarrow H^0(\mathfrak{A}(\pi)) \rightarrow 0. \end{aligned}$$

If $\lambda \neq 1$, then we find that $H^0(\mathfrak{A}(\pi)) = 0$ (since $H^0(\mathfrak{A}(\Pi_\lambda)) = 0$), and so $\mathfrak{A}(\pi)$ is supported in degree -1 . If $\lambda = 1$, then since $H^0(\mathfrak{A}(\Pi_1))$ is the skyscraper supported at x , we see either that $H^0(\mathfrak{A}(\pi))$ is again the skyscraper supported at x , or else that $H^0(\mathfrak{A}(\pi)) = 0$.

On the other hand, since π is not finitely presented [Sch15c], Lemma A.3.2 (3) implies that $H^{-1}(\mathfrak{A}(\pi))$ is *not* coherent, no matter the value of λ .

What can we say about the sheaf $H^{-1}(\mathfrak{A}(\pi))$? Here is one property it has: since $\text{Hom}(c\text{-Ind}_{KZ}^G \sigma', \pi) = 0$ if $\sigma' \neq \sigma_i$ for some $i = 0, \dots, 3$, we find that

$$\text{Ext}^1(\mathcal{O}_{\mathcal{C}_{\sigma'}}, H^{-1}(\mathfrak{A}(\pi))) = 0$$

for any such σ' .

If $\lambda = 1$, a representation π as above (that is, an irreducible quotient of Π_λ whose K -socle consists precisely of the σ_i , each with multiplicity one) can also be constructed globally, via the completed cohomology of definite quaternion algebras

and/or Shimura curves. (That λ necessarily equals 1 in this case was proved by Dotto–Le [DL21a]. The representations in completed cohomology have been studied by many authors, using Taylor–Wiles patching as a key ingredient, and we refer to the introduction to [BHH⁺21] for a brief overview of what is known.)

In the case when π is constructed via completed cohomology, the expected relationship between the functor \mathfrak{A} and completed cohomology suggests that the map $H^0(\mathfrak{A}(\Pi_1)) \rightarrow H^0(\mathfrak{A}(\pi))$ should be non-zero, and so at least in this case we anticipate that $H^0(\mathfrak{A}(\pi))$ is equal to the skyscraper at x . We explore the sheaf $\mathfrak{A}(\pi)$ further in Section 7.8, where we see that our conjectures imply that the irreducible representation π should be uniquely determined by the requirement that $H^0(\mathfrak{A}(\pi)) \neq 0$ (see Remark 7.8.5).

7.8. The adjoint functor. Return to the situation of Conjecture 6.1.14, and the cocontinuous functor $\mathfrak{A} : \text{Ind } D_{\text{f.p.}}^b(\text{sm. } G) \rightarrow \text{Ind Coh}(\mathcal{X}_d)$. By Lemma A.3.2 (1) (i.e. the adjoint functor theorem), \mathfrak{A} admits a right adjoint $\mathfrak{B} : \text{Ind Coh}(\mathcal{X}_d) \rightarrow \text{Ind } D_{\text{f.p.}}^b(\text{sm. } G)$. By part (2) of this same lemma (which applies, because \mathfrak{A} preserves compact objects by assumption; see Remark 6.1.28), we see that \mathfrak{B} is again cocontinuous.

REMARK 7.8.1. In the setting of Remark 1.4.6 (i.e. of a p -adic version of the Fargues–Scholze conjecture, probably formulated using the results of Mann [Man22]), the functors \mathfrak{A} and \mathfrak{B} will presumably be induced by the functors $i_!$ and $i^!$, where i is the inclusion of the open stratum $[*/G(F)] \hookrightarrow \text{Bun}_G$.

We can use the calculations of the previous section to investigate the functor \mathfrak{B} in the cases of $\text{PGL}_2(\mathbf{Q}_p)$ and $\text{PGL}_2(\mathbf{Q}_{p^2})$. (Of course, this is not precisely the setting of Conjecture 6.1.14, but the above remarks go over unchanged to this case.) This will allow us to connect the conjectures of this paper to the more traditional approach to the p -adic and mod p local Langlands correspondences via the cohomology of Shimura curves, and in particular to the hope that it is possible to associate to each representation $\bar{\rho} : \text{Gal}_{\mathbf{Q}_{p^2}} \rightarrow \text{GL}_2(\overline{\mathbf{F}}_p)$ a specific admissible representation of $\text{GL}_2(\mathbf{Q}_{p^2})$, which is determined by its occurrence in any (that is to say, every!) global context related to a global lift of $\bar{\rho}$ (see for example [Bre10a, Q. 3.11]). We return to this theme in Remark 7.8.5 below.

We can study the expected properties of \mathfrak{B} by means of the expected compatibility of \mathfrak{A} with Taylor–Wiles patching. Let F/\mathbf{Q}_p be arbitrary, and let \mathcal{X} be the stack of rank 2 (φ, Γ) -modules with determinant ε^{-1} ; and let L_∞ be the (hypothetical, for $F \neq \mathbf{Q}_p$) pro-coherent sheaf which is a kernel for \mathfrak{A} , in the sense that

$$\mathfrak{A}(\pi) := L_\infty \otimes_{\mathcal{O}[[G]]}^L \pi.$$

Then the adjoint functor is given by

$$(7.8.2) \quad \mathfrak{B}(\mathcal{F}) := \text{RHom}_{\text{pro-}\mathcal{O}_{\mathcal{X}}}(L_\infty, \mathcal{F}).$$

Suppose now that $x : \text{Spec } k \rightarrow \mathcal{X}$ is a finite type point, corresponding to a continuous representation $\bar{\rho} : G_F \rightarrow \text{GL}_2(k)$, and let $f : \text{Spf } R_\infty \rightarrow \mathcal{X}$ be the versal morphism arising from patching at some globalization of $\bar{\rho}$. By Remark 6.1.30 we expect that f^*L_∞ is equal to the patched module M_∞ . For generic choices of $\bar{\rho}$, the ring R_∞ is formally smooth, and it follows from [GN22, Prop. 4.3.1] and the results of [BHH⁺20, HW20, Wan22] that M_∞ is pro-flat over R_∞ . We expect

that it is even pro-free. Accordingly, we expect that we can compute $\mathfrak{B}(\mathcal{F})$ via the formula

$$\begin{aligned}\mathfrak{B}(\mathcal{F}) &= \mathrm{RHom}_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(L_{\infty}, \mathcal{F}) \\ &= R\Gamma \circ R\mathcal{H}om_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(L_{\infty}, \mathcal{F}) \\ &= R\Gamma(\mathcal{X}, \mathcal{H}om_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(L_{\infty}, \mathcal{F})).\end{aligned}$$

Write $G_x = \mathrm{Aut}_{\mathrm{Gal}_F}(\bar{\rho})$. Assume further that $\bar{\rho}^{\mathrm{ss}}$ is not scalar (this is automatically the case if as above $\bar{\rho}$ is assumed to be generic); so $G_x = \mu_2$ if $\bar{\rho}$ is irreducible or non-split reducible, and equals \mathbf{G}_m if $\bar{\rho}$ is a sum of two distinct characters. Then the residual gerbe at x is an affine immersion²³

$$i : [\mathrm{Spec} k/G_x] \hookrightarrow \mathcal{X},$$

and we can form the quasicoherent sheaf $\delta_x := i_* \mathcal{O}_{[\mathrm{Spec} k/G_x]}$. If, as above, $f : \mathrm{Spf} R_{\infty} \rightarrow \mathcal{X}$ is a versal morphism at x arising from Taylor–Wiles patching, and \mathfrak{m}_{∞} denotes the maximal ideal of R_{∞} , then we see that

$$\mathcal{H}om_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(L_{\infty}, \delta_x) = M_{\infty}^{\vee}[\mathfrak{m}_{\infty}]$$

equipped with some G_x -action. (To see this, note that $L_{\infty} \otimes_{\mathcal{O}_{\mathcal{X}}} \delta_x = M_{\infty}/\mathfrak{m}_{\infty} M_{\infty}$ equipped with some G_x -action, and that the continuous k -dual of $M_{\infty}/\mathfrak{m}_{\infty} M_{\infty}$ equals $M_{\infty}^{\vee}[\mathfrak{m}_{\infty}]$.) Note that by the construction of M_{∞} by Taylor–Wiles patching, $M_{\infty}^{\vee}[\mathfrak{m}_{\infty}]$ is equal to the \mathfrak{m} -torsion in completed cohomology, where \mathfrak{m} is a certain maximal ideal of an appropriate Hecke algebra. Since G_x is linearly reductive (being equal to either μ_2 or \mathbf{G}_m), we then see that

$$\mathfrak{B}(\delta_x) = (M_{\infty}^{\vee}[\mathfrak{m}_{\infty}])^{G_x}.$$

We can recover all of $M_{\infty}^{\vee}[\mathfrak{m}_{\infty}]$ by considering the values

$$\mathfrak{B}(\delta_x(n)) = (M_{\infty}^{\vee}[\mathfrak{m}_{\infty}](n))^{G_x}$$

as n varies. Or, rather than considering the various twists $\delta_x(n)$ individually, if we let $g : \mathrm{Spec} k \rightarrow \mathcal{X}$ denote the composite

$$\mathrm{Spec} k \rightarrow [\mathrm{Spec} k/G_x] \xrightarrow{i} \mathcal{X},$$

then we can apply \mathfrak{B} to the quasicoherent sheaf $\Delta_x := g_* \mathcal{O}_{\mathrm{Spec} k}$. (This is the quasicoherent sheaf over the residual gerbe of x which corresponds to the regular representation of G_x .) Indeed, we compute that

$$\begin{aligned}\mathfrak{B}(\Delta_x) &= \mathrm{RHom}_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(L_{\infty}, g_* \mathcal{O}_{\mathrm{Spec} k}) \\ &= \mathrm{RHom}_{\mathrm{pro}\text{-}\mathcal{O}_{\mathrm{Spec} k}}(g^* L_{\infty}, \mathcal{O}_{\mathrm{Spec} k}) = \mathrm{Hom}_{\mathrm{cont}}(M_{\infty}/\mathfrak{m}_{\infty}, k) = M_{\infty}^{\vee}[\mathfrak{m}_{\infty}].\end{aligned}$$

Thus $x \mapsto \mathfrak{B}(\Delta_x)$ should recover the hoped-for “traditional” mod p local Langlands correspondence alluded to above. The G_x -action on $M_{\infty}^{\vee}[\mathfrak{m}_{\infty}]$ that arises by considering the sheaf-Hom $\mathcal{H}om_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(L_{\infty}, \delta_x)$ as above is some additional structure in the mod p Langlands correspondence that (as far as we are aware) has not been previously considered (although one could say that it is implicit in the ideas of Breuil and his collaborators expressed in papers such as [Bre11] and [BH15]).

²³If $\bar{\rho}$ is semisimple then i is a closed immersion; otherwise i factors as $[\mathrm{Spec} k/\mu_2] \xrightarrow{\mathrm{open}} [\mathbf{A}_k^1/\mathbf{G}_m] \xrightarrow{\mathrm{closed}} \mathcal{X}$, which exhibits it as a composite of affine immersions.

7.8.3. *Some examples.* In the cases that $F = \mathbf{Q}_p$ or $F = \mathbf{Q}_{p^2}$, we know something about $M_\infty^\vee[\mathfrak{m}]$ (in the former case by M.E.'s local-global compatibility theorem [Eme11a], and in the latter case by [BHH⁺21]), and we either know or can guess something about how G_x acts. We have the following examples.

- (1) If $F = \mathbf{Q}_p$ and $\bar{\rho} = \bar{\chi}_1 \oplus \bar{\chi}_2$ with $\bar{\chi}_1 \bar{\chi}_2^{-1} \neq \text{triv or } \varepsilon^{\pm 1}$, then $M_\infty^\vee[\mathfrak{m}]$ is a direct sum $\pi_1 \oplus \pi_2$ of two irreducible principal series representations, and in fact for an appropriate ordering of π_1, π_2 we have

$$\mathcal{H}om_{\text{pro-}\mathcal{O}_X}(L_\infty, \delta_x) = \pi_1(1) \oplus \pi_2(-1)$$

(where the twists denote the weight of the \mathbf{G}_m -action; note that the apparent asymmetry is due to the need to choose an ordering of $\bar{\chi}_1, \bar{\chi}_2$ in order to identify G_x with \mathbf{G}_m).

In this case $\Delta_x := \bigoplus_{n=-\infty}^{\infty} \delta_x(n)$, and so

$$\mathfrak{B}(\Delta_x) = \mathfrak{B}(\delta_x(-1)) \oplus \mathfrak{B}(\delta_x(1)) = \pi_1 \oplus \pi_2.$$

If σ_1, σ_2 are (a chosen ordering of) the two Serre weights of $\bar{\rho}$, we find that $(\mathfrak{A} \circ \mathfrak{B})(\delta_x(-1)) = \mathfrak{A}(\pi_1)$ is supported on $\mathcal{X}(\sigma_1)$ while $(\mathfrak{A} \circ \mathfrak{B})(\delta_x(1)) = \mathfrak{A}(\pi_2)$ is supported on $\mathcal{X}(\sigma_2)$. As a side remark, we note that although δ_x is supported on the intersection $\mathcal{X}(\sigma_1) \cap \mathcal{X}(\sigma_2)$, these representations are not; the representations $\mathfrak{A}(\pi_i)$ are supported on the closed locus in the respective component $\mathcal{X}(\sigma_i)$ corresponding to Galois representations with semisimplification $\bar{\rho}$. Thus $(\mathfrak{A} \circ \mathfrak{B})(\Delta_x) = \mathfrak{A}(\pi_1) \oplus \mathfrak{A}(\pi_2)$ is supported on the union of these two loci, which is the locus of all Galois representations having $\bar{\rho}$ as their semisimplification.

- (2) If $F = \mathbf{Q}_p$ and $\bar{\rho}$ is a non-split extension of characters $\bar{\chi}_1$ and $\bar{\chi}_2$, again with $\bar{\chi}_1 \bar{\chi}_2^{-1} \neq \text{triv or } \varepsilon^{\pm 1}$, then $M_\infty^\vee[\mathfrak{m}]$ is a non-split extension of the two irreducible principal series π_1 and π_2 appearing in case (1), and

$$\mathcal{H}om_{\text{pro-}\mathcal{O}_X}(L_\infty, \delta_x) = M_\infty^\vee[\mathfrak{m}](1),$$

the twist indicating that μ_2 acts non-trivially.

In this case $\Delta_x := \delta_x \oplus \delta_x(1)$, and so

$$\mathfrak{B}(\Delta_x) = \mathfrak{B}(\delta_x(1)) = M_\infty^\vee[\mathfrak{m}]$$

(the non-split extension of π_1 and π_2). Furthermore, $(\mathfrak{A} \circ \mathfrak{B})(\Delta_x) = (\mathfrak{A} \circ \mathfrak{B})(\delta_x(1))$ is a non-split extension of $\mathfrak{A}(\pi_1)$ and $\mathfrak{A}(\pi_2)$, and in particular has support equal to the closed locus corresponding to Galois representations with semisimplification equal to $\chi_1 \oplus \chi_2$.

- (3) If $F = \mathbf{Q}_p$ and $\bar{\rho}$ is irreducible, then $M_\infty^\vee[\mathfrak{m}] = \pi$ is a supersingular irreducible representation, and

$$\mathcal{H}om_{\text{pro-}\mathcal{O}_X}(L_\infty, \delta_x) = \pi_1(1)$$

(the twist indicating that the action of μ_2 is nontrivial).

As in case (2), we have $\Delta_x := \delta_x \oplus \delta_x(1)$, and so

$$\mathfrak{B}(\Delta_x) = \mathfrak{B}(\delta_x(1)) = \pi,$$

and $(\mathfrak{A} \circ \mathfrak{B})(\Delta_x) = (\mathfrak{A} \circ \mathfrak{B})(\delta_x(1)) = \delta_x(1)$.

- (4) If $F = \mathbf{Q}_{p^2}$ and $\bar{\rho} = \bar{\chi}_1 \oplus \bar{\chi}_2$ with $\bar{\chi}_1, \bar{\chi}_2$ suitably generic, then by [BHH⁺21, Cor. 1.3.10], we have $M_\infty^\vee[\mathfrak{m}] = \pi_1 \oplus \pi \oplus \pi_2$, where π_1, π_2 are explicit

irreducible principal series representations, and π is a (non-explicit!) supersingular irreducible representation. We expect that

$$\mathcal{H}om_{\text{pro-}\mathcal{O}_X}(L_\infty, \delta_x) = \pi_1(2) \oplus \pi \oplus \pi_2(-2).$$

In this case, as in case (1), we have $\Delta_x := \bigoplus_{n=-\infty}^{\infty} \delta_x(n)$, and so

$$\mathfrak{B}(\Delta_x) = \mathfrak{B}(\delta_x(-2)) \oplus \mathfrak{B}(\delta_x) \oplus \mathfrak{B}(\delta_x(2)) = \pi_1 \oplus \pi \oplus \pi_2.$$

In particular, $\mathfrak{B}(\delta_x) = \pi$ is irreducible and supersingular. (A similar analysis could be performed in the case of a nonsplit extension of characters, using the results of [HW22].)

- (5) If $F = \mathbf{Q}_{p^2}$ and $\bar{\rho}$ is irreducible and suitably generic, then by [BHH⁺21, Thm. 1.3.11], we have $M_\infty^\vee[\mathfrak{m}] = \pi$, where π is a supersingular irreducible representation. We expect that

$$\mathcal{H}om_{\text{pro-}\mathcal{O}_X}(L_\infty, \delta_x) = \pi$$

(with trivial μ_2 -action).

As in case (3), we have $\Delta_x := \delta_x \oplus \delta_x(1)$, and so

$$\mathfrak{B}(\Delta_x) = \mathfrak{B}(\delta_x) = \pi$$

is irreducible and supersingular.

REMARK 7.8.4. The twists in examples (4) and (5) were, at first, arrived at heuristically, based on examples (1), (2), and (3) and the philosophy suggested by Breuil's paper [Bre11]. However, one can confirm them (within our conjectural framework, i.e. assuming that \mathfrak{A} exists with its conjectured properties, and that the values of $\mathfrak{A}(c\text{-Ind}_{KZ}^G \sigma)$ for sufficiently generic Serre weights σ are as conjectured in Expectation 7.6.12) in the following manner: The Serre weights σ of $\bar{\rho}$ are known, and are precisely the irreducible constituents of the K -socle of $M_\infty^\vee[\mathfrak{m}_\infty]$, and are known to appear with multiplicity one. It is also known how they distribute themselves among the various direct summands of $M_\infty^\vee[\mathfrak{m}_\infty]$. (See [GLS12, GK14, EGS15, Bre14] for these results.) Thus, as σ ranges over the Serre weights of $\bar{\rho}$, the multiplicity spaces $\text{Hom}_G(c\text{-Ind}_{KZ}^G \sigma, M_\infty^\vee[\mathfrak{m}_\infty])$ are one-dimensional, and it is known which direct summand of $M_\infty^\vee[\mathfrak{m}_\infty]$ any non-zero morphism factors through.

On the other hand, we can write this Hom space as

$$\begin{aligned} \text{Hom}_G(c\text{-Ind}_{KZ}^G \sigma, M_\infty^\vee[\mathfrak{m}_\infty]) &= \text{Hom}_G(c\text{-Ind}_{KZ}^G \sigma, \mathfrak{B}(\Delta_x)) \\ &= \text{Hom}_G\left(c\text{-Ind}_{KZ}^G \sigma, \mathfrak{B}\left(\bigoplus_n \delta_x(n)\right)\right) = \text{Hom}_G\left(c\text{-Ind}_{KZ}^G \sigma, \bigoplus_n \mathfrak{B}(\delta_x(n))\right) \\ &= \bigoplus_n \text{Hom}_G\left(c\text{-Ind}_{KZ}^G \sigma, \mathfrak{B}(\delta_x(n))\right), \end{aligned}$$

the sum ranging over either $n \in \mathbf{Z}$ or $n \in \mathbf{Z}/2\mathbf{Z}$, depending on whether $\bar{\rho}$ is a direct sum of characters or is irreducible. Now the adjunction between \mathfrak{A} and \mathfrak{B} shows that

$$\begin{aligned} \text{Hom}_G\left(c\text{-Ind}_{KZ}^G \sigma, \mathfrak{B}(\delta_x(n))\right) &= \text{Hom}_G(\mathfrak{A}(c\text{-Ind}_{KZ}^G \sigma), \delta_x(n)) \\ &= \text{Hom}_G(\mathcal{O}_{\mathcal{X}(\sigma)}(-1, -1), \delta_x(n)) = i^*(\mathcal{O}_{\mathcal{X}(\sigma)}(-1, -1))(n)^{G_x}, \end{aligned}$$

where, as above, $i : [\text{Spec } k/G_x] \hookrightarrow \mathcal{X}$ is the closed immersion of the residual gerbe at x into \mathcal{X} .

Since $\mathcal{O}_{\mathcal{X}(\sigma)}(-1, -1)$ is an invertible sheaf, we see that $i^*(\mathcal{O}_{\mathcal{X}(\sigma)}(-1, -1))$ is a character of G_x , and so has the form $k(m)$ for some integer m (which we see is determined by σ and x). Thus $i^*(\mathcal{O}_{\mathcal{X}(\sigma)}(-1, -1))(n)^{G_x} = k(m+n)^{G_x}$ is non-zero for exactly one value of n (namely $n = -m$), and this determines the twist of the corresponding constituent of $M_\infty^\vee[\mathfrak{m}_\infty]$; namely, it appears with a twist by m . Computing the values of m for the various possible choices of x (which is to say, of $\bar{\rho}$) and of σ then yields the twists appearing in examples (4) and (5).

REMARK 7.8.5. Since the results of Breuil and Paškūnas [BP12] it has been clear that the formulation of the p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ would have to be very different from the (classical) correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$, for the reason that there appear to be many more representations of $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ than there are 2-dimensional representations of $\mathrm{Gal}_{\mathbf{Q}_{p^2}}$. In particular, as already discussed in Section 7.7.59, there are infinite families of supersingular irreducible admissible mod p representations of $\mathrm{GL}_2(\mathbf{Q}_{p^2})$, all of which appear to correspond to the same irreducible representation $\bar{\rho}$ of $\mathrm{Gal}_{\mathbf{Q}_{p^2}}$.

Our conjectural functor \mathfrak{A} gives one explanation of this phenomenon: namely, the apparent lack of Galois representations is made up for by an abundance of coherent sheaves on \mathcal{X} , and in particular, these supersingular representations all correspond to Ind-coherent sheaves on \mathcal{X} , which are distinguished by their (quasi-coherent, but not coherent) H^{-1} sheaves.

There is, however, an earlier approach to reconciling the results of [BP12] with the expectation that there is indeed a p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_{p^2})$, which is to hope that there is a uniquely determined canonical element $\pi_{\bar{\rho}}$ of the infinite family of supersingular representations constructed in [BP12]. Furthermore, the representation $\pi_{\bar{\rho}}$ is determined by local-global compatibility, in the sense that it occurs in the cohomology of Shimura curves, for any globalisation of $\bar{\rho}$; so the existence of $\pi_{\bar{\rho}}$ can be rephrased as showing that the representations occurring in Shimura curves depends only on $\bar{\rho}$ (and that it is indeed an supersingular irreducible representation of the type constructed by [BP12]). This question has been investigated by many authors, and there is a considerable amount of evidence in favour of this hope; we refer again to the introduction to [BHH⁺21] for a brief account of these investigations. (This hope motivated the proposed construction of a p -adic local Langlands correspondence via Taylor–Wiles patching in [CEG⁺16], which in turn motivated the conjectures of these notes.)

The computations in this section are also consistent with this expectation, and indeed we see from examples (4) and (5) above that if x is the closed point corresponding to a suitably generic $\bar{\rho}$, then we should have $\pi_{\bar{\rho}} = \mathfrak{B}(\delta_x)$. We claim that $\pi_{\bar{\rho}}$ can be characterised among all irreducible representations π of $\mathrm{GL}_2(\mathbf{Q}_{p^2})$ by the following property: it is the unique irreducible π for which the support of $H^0(\mathfrak{A}(\pi))$ contains x .

Indeed, suppose that π is irreducible and that the support of $H^0(\mathfrak{A}(\pi))$ contains x . Since π and $\pi_{\bar{\rho}} = \mathfrak{B}(\delta_x)$ are irreducible, it is enough to show that

$$\mathrm{Hom}_{\mathrm{GL}_2(\mathbf{Q}_{p^2})}(\pi, \mathfrak{B}(\delta_x)) \neq 0.$$

By adjunction, this is equivalent to showing that

$$\mathrm{Hom}_{\mathrm{pro}\text{-}\mathcal{O}_{\mathcal{X}}}(\mathfrak{A}(\pi), \delta_x) \neq 0.$$

To see that this holds, we consider the canonical morphism

$$\mathfrak{A}(\pi) \rightarrow \tau^{\geq 0}\mathfrak{A}(\pi),$$

and note that $\tau^{\geq 0}\mathfrak{A}(\pi)$ is coherent (by Remark 6.1.28), and thus admits a morphism to δ_x by the assumptions that the support of $H^0(\mathfrak{A}(\pi))$ contains x (here we use that $\mathfrak{A}(\pi)$ is contained in the even part of $\text{Ind Coh}(\mathcal{X})$, see Section 7.5.11).

Note that since the supersingular representation $\pi_{\bar{\rho}}$ admits a surjection from a representation $c\text{-Ind}_{KZ}^G \sigma$ (where σ is some Serre weight), the coherent sheaf $\tau^{\geq 0}\mathfrak{A}(\pi)$ admits a surjection from a line bundle. In particular in case (5) we see that $H^0(\mathfrak{A}(\pi))$ should equal δ_x .

REMARK 7.8.6. In cases (4) and (5) above we do not have an explicit description of $\mathfrak{A}(\pi)$, except that in case (5) its H^0 is equal to δ_x (as discussed in Remark 7.8.5). (Similarly, we expect that in case (4) the H^0 is the structure sheaf of the closed locus of all points whose semisimplification is x .) However, $\mathfrak{A}(\pi)$ itself could in principle be computed via the machinery of Section A.8.12, because it is equal to $(\mathfrak{A} \circ \mathfrak{B})(\delta_x)$, and we expect to be able to write down a collection of generators for the image of the hypothetical functor \mathfrak{A} (namely, the various line bundles corresponding to the $\mathfrak{A}(c\text{-Ind}_{KZ}^G \sigma)$, which in the generic case we made explicit in Section 7.7).

REMARK 7.8.7. One anticipated feature of the the traditional mod p local Langlands correspondence for $\text{GL}_2(F)$ is that, for reducible $\bar{\rho}$, and writing (in terms of the notation introduced above) $\pi(\bar{\rho}) := \mathfrak{B}(\Delta_x)$, one should have

$$(7.8.8) \quad \pi(\bar{\rho}^{\text{ss}}) \cong \pi(\bar{\rho})^{\text{ss}}.$$

It is natural to ask if this can be seen from our present perspective.

To this end, we note that if $\bar{\rho}^{\text{ss}}$ is a direct sum of two suitably generic characters χ_1 and χ_2 of $G_{\mathbf{Q}_p^f}$, then the locus of $\bar{\rho}$ which are extensions of χ_1 by χ_2 is equal to a copy of $[\mathbf{A}^f/\mathbf{G}_m]$. If we let \mathcal{F} denote the structure sheaf of $[\mathbf{A}^f/\mathbf{G}_m]$, pushed forward to a coherent sheaf on \mathcal{X} , then $\mathcal{H}om(L_\infty, \mathcal{F})$ can be interpreted as a \mathbf{G}_m -equivariant flat family of smooth $\text{GL}_2(F)$ -representations over \mathbf{A}^f . If we forget the \mathbf{G}_m -equivariant structure, this family interpolates the various $\pi(\bar{\rho})$ as $\bar{\rho}$ runs over the Galois representations (i.e. extensions of characters) parameterized by \mathbf{A}^f . It seems plausible that the existence of this flat family (with its \mathbf{G}_m -equivariant structure!) could imply the anticipated isomorphism (7.8.8).

7.9. Global functions and Hecke algebras. The full faithfulness of the putative functor \mathfrak{A} has some concrete consequences for the geometry of the stacks \mathcal{X}_d . For example, we saw above that if $G = \text{GL}_2(F)$ with F/\mathbf{Q}_p unramified, then for a generic Serre weight σ we expect $\mathfrak{A}(c\text{-Ind}_K^G \sigma)$ to be a line bundle on the component $\mathcal{X}(\sigma)$; so that in particular the (underived) Hecke algebra $\text{End}_G(c\text{-Ind}_K^G \sigma)$ should be identified with the global functions on $\mathcal{X}(\sigma)$.

This expectation follows immediately from the explicit description of $\mathcal{X}(\sigma)$ in terms of Fontaine–Laffaille theory, and similar results can be proved for tame types by explicit integral p -adic Hodge theory (again in the generic case for $\text{GL}_2(F)$ with F/\mathbf{Q}_p unramified).

More generally, the following result is work in progress of Eivind Hjellev, Louis Jaburi, Rachel Knak, Yulia Kotelnikova, Billy Lee, and Shenrong Wang, under the remit of the NSF grant 1952556 “Collaborative Research: Geometric Structures in the p -adic Langlands Program”: if σ is a suitably generic Serre weight for $\text{GL}_d(F)$,

with F/\mathbf{Q}_p unramified, and $\mathcal{X}(\sigma)$ is the irreducible component of the stack $\mathcal{X}_{F,d}$ labelled by σ , then the global functions on $\mathcal{X}(\sigma)$ are given by a ring of the form $k[T_1, \dots, T_{d-1}, T_d, T_d^{-1}]$. Note that this ring is isomorphic to the Hecke algebra $\text{End}_G(c\text{-Ind}_K^G \sigma)$, but there is no obviously defined direct map between the Hecke algebra and the ring of global functions; rather, it is possible to characterize T_i as a renormalized Hecke operator on the ordinary locus in $\mathcal{X}(\sigma)$.

In the non-generic case, the paper [GKK⁺22] of Anthony Guzman, Kalyani Kansal, Jason Kountouridis, Ben Savoie, and Xiyuan Wang (also under NSF grant 1952556) proves under mild hypotheses that for F/\mathbf{Q}_p unramified and σ a not-necessarily-generic Serre weight of $\text{GL}_2(F)$, the global functions on the normalization of $\mathcal{X}(\sigma)$ can be identified with the corresponding Hecke algebra.

8. Categorical local Langlands in the case $\ell \neq p$

Although our focus so far has been on the categorical formulation of the p -adic local Langlands conjecture for the group GL_d over a p -adic local field, we also wish to address the global context; indeed, there is a case to be made that the main point of developing the local aspects of number theory is to provide tools for investigating global problems. We will explain the global theory in Section 9; this naturally involves the local $\ell \neq p$ case, which we study in this section.

The main global problem we will focus on is the conjectural description of the cohomology of Shimura varieties (and, even more speculatively, of general congruence quotients) in terms of the sheaves arising from the categorical local conjecture. On the one hand, this is related to the completed cohomological approach to the Fontaine–Mazur conjecture (as in [Eme11a, Pan22a]). On the other, it is related to Taylor–Wiles patching (and in particular, provides a conjectural explanation for the expected relationship between patched modules and the coherent sheaves of the categorical conjecture).

In order to study global problems, especially those involving Shimura varieties, we have to be able to consider reductive groups that are not necessarily split. Although currently our understanding of p -adic local Langlands in the non- GL_d case is even more fragmentary than our understanding for GL_d , the local situation at primes away from p is better understood; indeed, there are precise conjectures in this case, on which significant progress has been made, in particular in the papers [FS21, Zhu20, Hel20, BZCHN20, HZ]. In this section we briefly recall some of these local conjectures and results, referring to these papers and to the notes [FS22] for further details.

We begin by briefly recalling the formalism of C -groups; these will be the targets of the Galois representations arising in the local context when non-split groups are involved. We then describe the moduli stacks of local Langlands parameters in the $\ell \neq p$ case, and very briefly recall the categorical local Langlands conjectures that we need for our global discussion.

8.1. C -groups. As always, we have fixed a finite extension L of \mathbf{Q}_p , whose ring of integers \mathcal{O} serves as our coefficients.

If F is a finite extension of either \mathbf{Q}_ℓ (for some prime ℓ) or of \mathbf{Q} , and G is a connected reductive linear algebraic group over F , then we can form the C -group cG of G as in [BG14]; see also [Zhu20, §3.0]. This is a variant of the L -group ${}^L G$ which is better adapted to issues relating to fields of definition. By construction, it

is a semi-direct product

$${}^cG := \widehat{G} \rtimes (\mathbf{G}_m \times \text{Gal}(\widetilde{F}/F)),$$

where \widetilde{F} is a certain finite Galois extension of F , chosen “large enough” for the purposes at hand, and \widehat{G} is the usual Langlands dual group. (In particular, \widetilde{F} should contain a splitting field for the quasi-split inner form of G ; the definitions involving cG that appear below are easily seen to be independent of the particular choice of \widetilde{F} .) Here \mathbf{G}_m acts via the cocharacter $\rho_{\text{ad}} : \mathbf{G}_m \rightarrow \widehat{G}_{\text{ad}}$ defined by one-half the sum of the positive coroots of \widehat{G} . We regard cG as a group scheme over the ring of integers \mathcal{O} .

REMARK 8.1.1. The C -group can also be defined as the L -group of a certain \mathbf{G}_m -extension \widetilde{G} of G , see [BG14, Defn. 5.3.2].

8.2. Langlands parameters in the local case — $\ell \neq p$. If F is an ℓ -adic field with $\ell \neq p$, then one can define a stack $\mathcal{X}_{G,F}$ which parameterizes representations of a Weil–Deligne-type group of F (suitably interpreted) into cG . There are in fact various ways of defining $\mathcal{X}_{G,F}$, depending on the base over which one wishes to have it live. We will consider $\mathcal{X}_{G,F}$ as a stack over $\text{Spf } \mathcal{O}$, and then we can make the following definition: we first choose elements $\sigma, \tau \in \text{Gal}_F := \text{Gal}(\widetilde{F}/F)$ that respectively lift Frobenius and a topological generator of tame inertia, and let Γ_F denote the subgroup of Gal_F generated by P_F (the wild inertia subgroup of Gal_F), σ , and τ . We have a short exact sequence of groups

$$1 \rightarrow P_F \rightarrow \Gamma_F \rightarrow \Gamma_F^{\text{tame}} := \langle \sigma, \tau \ \sigma\tau\sigma^{-1} = \tau^q \rangle \rightarrow 1$$

(where q denotes the order of the residue field of F). Note that P_F has a neighbourhood basis of the identity consisting of open subgroups Q which are normal in Γ_F , so that we may write $\Gamma_F = \varprojlim_Q \Gamma_F/Q$ as a projective limit of discrete quotients, with each Q being a subgroup of P_F (and hence pro- ℓ).

For any \mathcal{O} -algebra A in which p is nilpotent, and for any of the precedingly mentioned subgroups Q , we define $\mathcal{X}_{G,F}^Q(A)$ to be the groupoid of homomorphisms $\Gamma_F/Q \rightarrow {}^cG(A)$ which are C -homomorphisms, i.e. which satisfy the C -group version of the usual “ L -homomorphism” condition: the composite

$$(8.2.1) \quad \Gamma_F \rightarrow \Gamma_F/Q \rightarrow {}^cG(A) \rightarrow A^\times \times \text{Gal}(\widetilde{F}/F)$$

(the second arrow being the tautological projection) should equal the product of the inverse cyclotomic character $\Gamma_F \rightarrow \text{Gal}_F \rightarrow \mathbf{Z}_p^\times$ and the canonical surjection $\Gamma_F \rightarrow \text{Gal}_F \rightarrow \text{Gal}(\widetilde{F}/F)$. We then set $\mathcal{X}_{G,F} := \varinjlim_Q \mathcal{X}_{G,F}^Q$; each transition map is an open and closed embedding (because each Q is pro- ℓ and $\ell \neq p$), and in this way $\mathcal{X}_{G,F}$ is given the structure of an Ind-algebraic stack.

Actually, while this discussion defines $\mathcal{X}_{G,F}$ as a classical stack, we really want to define it as a *derived* stack, which we can do by following the ideas explained in Appendix F. A key point is that, even when $\mathcal{X}_{G,F}$ is defined as a derived stack, it turns out to be classical, and even a local complete intersection ([Zhu20], Prop. 3.1.6). Since local Galois cohomology is supported in cohomological degrees $[0, 2]$, this amounts to the classical stack underlying $\mathcal{X}_{G,F}$ being of the expected dimension, which can be proved by an explicit computation, which we now briefly recall.

The first point is that, for an A -valued point $\rho : \Gamma_F \rightarrow {}^cG(A)$ of $\mathcal{X}_{G,F}$ (here A is a (classical!) \mathcal{O} -algebra in which p is nilpotent), the tangent complex to $\mathcal{X}_{G,F}$ at ρ is computed by the shifted continuous cochain complex $C^\bullet(\Gamma_F, \mathfrak{g} \otimes_{\mathcal{O}} A)[1]$, where Γ_F acts on $\mathfrak{g} \otimes_{\mathcal{O}} A$ via the composite of ρ with the adjoint action of cG on its Lie algebra. This complex has cohomology supported in degrees $[-1, 1]$, and so we find that the cotangent complex of $\mathcal{X}_{G,F}$ at any point has cohomology vanishing in degrees ≤ -2 . This implies that $\mathcal{X}_{G,F}$ is *quasi-smooth*, i.e. a (possibly derived) local complete intersection. Furthermore, $\mathcal{X}_{G,F}$ is of its expected dimension (the Euler characteristic of $C^\bullet(\Gamma_F, \mathfrak{g} \otimes_{\mathcal{O}} A)[1]$ equals zero, essentially by the local Euler characteristic formula, and $\mathcal{X}_{G,F}$ is flat over $\mathrm{Spf} \mathcal{O}$ of relative dimension zero — see [Zhu20, Prop. 3.2.7]). This implies that $\mathcal{X}_{G,F}$ is actually a classical local complete intersection, as claimed (see the proof of [Zhu20, Prop. 2.2.13], on which the proof of [Zhu20, Prop. 3.1.6] ultimately depends).

8.3. Categorical local Langlands. Let F be an ℓ -adic local field, for some prime $\ell \neq p$, and let G be a connected reductive group over F . We let $D(G(F))$ denote the stable ∞ -category of complexes of smooth representations of $G(F)$ on locally p -torsion \mathcal{O} -modules (*locally p -torsion* meaning that each element of the module is annihilated by some power of p). We let $D(G(F))'$ denote a “renormalization” of $D(G(F))$ analogous to the one described in [Zhu20, §4.1] (in which $\mathrm{Rep}(G(F), \mathcal{O})$ is replaced by $\mathrm{Rep}^{\mathrm{ren}}(G(F), \mathcal{O})$, the key difference between that context and ours being that no locally p -torsion hypothesis is applied there); in $D(G(F))'$, all compactly induced representations are compact objects (by definition).

We then have the following categorical form of the local Langlands conjecture in the $\ell \neq p$ context.

CONJECTURE 8.3.1. *There is a fully faithful cocontinuous functor of stable ∞ -categories*

$$\mathfrak{A}_G : D(G(F))' \rightarrow \mathrm{Ind} \mathrm{Coh}(\mathcal{X}_{G,F}).$$

REMARK 8.3.2. As in Remark 6.1.23, we expect that the functor \mathfrak{A}_G of Conjecture 8.3.1 should be given by the (derived) tensor product with a kernel. In the case $G = \mathrm{GL}_n$ this expectation was first made explicit by E.H. in [Hel20]; here for the Iwahori block of L -representations of $\mathrm{GL}_n(F)$, the functor is expected to be given by the derived tensor product over the Iwahori Hecke algebra with the Iwahori invariants in the Emerton–Helm–Moss local Langlands correspondence in families.

Conjecture 8.3.1, as stated, is rather imprecise. A much more precise version is stated by Zhu as [Zhu20, Conj. 4.5.1]. (In that statement there is also no restriction to locally p -torsion representations, and the target category is the Ind-coherent sheaves on an algebraic \mathcal{O} -stack, of which our stack $\mathcal{X}_{G,F}$ is the p -adic completion.) The analogous conjecture in the case of L -coefficients has been made independently by E.H. [Hel20]. When G is split, and continuing to work with L -coefficients, an analogous conjecture was made, yet again independently, by Ben-Zvi–Chen–Helm–Nadler [BZCHN20], who furthermore construct the restriction of the conjectural functor to the Iwahori block. This restriction (and an extension of it to the general case of unipotent parameters) was also constructed independently by Hemo–Zhu [HZ]. We remark that when G is not of the form GL_d (for some d), the

functor \mathfrak{A}_G depends on the choice of auxiliary data (e.g. a pinning and associated Whittaker data, in the case that G is quasi-split). There should be line bundles on the stack $\mathcal{X}_{G,F}$ that allow one to twist between the functors arising from different choices of auxiliary data (see [Zhu20, Rem. 4.5.7]).

In the case $G = \mathrm{GL}_d$, the construction of [BZCHN20] extends to the entire category of smooth $\mathrm{GL}_d(F)$ -representations (essentially because all the Hecke algebras appearing in the theory of the Bernstein centre for GL_d are Iwahori Hecke algebras). As already explained in Remark 8.3.2, E.H. gives another conjectural description of this functor (still in the case of L -coefficients) in terms of the “local Langlands in families” whose existence was conjectured in [EH14], and proved by Helm and Moss [HM18]. It seems likely that E.H.’s conjectural description should apply also in the case of \mathcal{O} -coefficients; see e.g. [Zhu20, Rem. 4.4.4], where Zhu also gives yet another conjectural description of the functor in the case of GL_d , inspired by a conjecture made by Braverman–Finkelberg [BF19] in the context of geometric Langlands. (As far as we know, even for L -coefficients, the construction of [BZCHN20] has not been shown to match with either E.H.’s conjectural description via local Langlands in families, or Zhu’s conjectural description à la Braverman–Finkelberg.)

In fact, the fully faithful functor of Conjecture 8.3.1 should actually arise as the restriction of an equivalence of categories coming from an arithmetic analogue of the geometric Langlands correspondence in its categorical form. Such a conjecture has been made by Fargues and Scholze [FS21]. Another such conjecture has also been made by Zhu [Zhu20, §4.6, esp. Conj. 4.6.4 and Rem. 4.6.7].

REMARK 8.3.3. Let $G = \mathrm{GL}_d(F)$. In characteristic zero (over a fixed p -adic base field L) there should be a version of Conjecture 8.3.1 involving smooth G -representations on L -vector spaces and (ind-)coherent sheaves on the stack

$$\mathcal{Y}_{G,F} = \bigcup_{F'} \mathrm{Mod}_{d,\varphi,N,F'/F}$$

of Section 5.2 instead of the stack of L -parameters $\mathcal{X}_{G,F}$. Note that the stack of L -parameters for GL_d agrees with the moduli stack of d -dimensional Weil–Deligne representations $\mathrm{WD}_{d,F}$. The point is that we expect local Langlands in families to make sense over $\mathcal{Y}_{G,F}$. Indeed, following [Hel16] local Langlands in families is a family of representation \mathcal{V} over $\mathcal{X}_{G,F}$ constructed using a choice of a Whittaker datum (N, ψ) and an isomorphism from the ring of invariant functions on $\mathcal{X}_{G,F}$ to the Bernstein center \mathcal{Z} of the category of smooth G -representations on L -vector spaces. Then \mathcal{V} is the unique quotient of

$$(c\text{-Ind}_N^G \psi) \otimes_{\mathcal{Z}} \mathcal{O}_{\mathcal{X}_{G,F}}$$

with a certain prescribed quotient at the generic points of $\mathcal{X}_{G,F}$. Using Proposition 5.2.8 and Proposition 5.2.10 the same construction should carry over to $\mathcal{Y}_{G,F}$. There are two difficulties in this strategy: we need to fix a Whittaker datum (N, ψ) which is only possible if L contains all the p -power roots of unity (and hence in particular L can’t be finite over \mathbf{Q}_p). And the stacks in question are not geometrically irreducible (and their connected components are not geometrically connected). Hence one needs to prove some rationality results in order to define the family \mathcal{V} over a given base field L .

9. Moduli stacks of global Langlands parameters and the cohomology of Shimura varieties

In this section we briefly describe the global aspects of the theory. We sketch the construction of moduli stacks of global Galois representations, and of the associated global-to-local restriction maps, and then explain how these maps can be used, along with the conjectural categorical local Langlands correspondences, to give a conjectural formula for the cohomology of Shimura varieties. We next explain some results in the case of GL_2/\mathbf{Q} obtained in [EGZ], which explain the connection between these conjectures, the Taylor–Wiles method, and M.E.’s results on local-global compatibility in the p -adic Langlands program [Eme11a].

9.1. Global moduli stacks. We begin with some basic facts regarding moduli stacks of global Langlands parameters. Although one can usually reduce to the case of considering reductive groups defined over \mathbf{Q} (via an application of restriction of scalars), it is no more difficult to consider the case of general number fields, and we do this. Thus we fix a number field F and a connected reductive group G defined over F . We also fix a finite set S of places of F containing all places above p and ∞ .

We can define a derived stack $\mathcal{X}_{G,F,S}$ which lies over $\mathrm{Spf} \mathcal{O}$, and whose groupoid $\mathcal{X}_{G,F,S}(A)$ of A -valued points, for a classical \mathcal{O} -algebra A in which p is nilpotent, coincides with the groupoid of continuous representations $\rho : \mathrm{Gal}_{F,S} \rightarrow {}^cG(A)$ (the target being given its discrete topology) which are C -homomorphisms (i.e. those which satisfy (8.2.1)). (This defines the points of $\mathcal{X}_{G,F,S}$ on classical rings, i.e. it characterizes the underlying classical stack $\mathcal{X}_{G,F,S}^{cl}$; the definition of the points of $\mathcal{X}_{G,F,S}$ on derived rings is more involved — see [Zhu20, §2.4] for the general formalism, the discussion at the end of [Zhu20, §3.4] for a summary of the specifics in this context, and forthcoming paper [EZ] for the precise details.)

In the case when $G = \mathrm{GL}_d$ for some $d \geq 1$, the underlying classical stack $\mathcal{X}_{\mathrm{GL}_d,F,S}^{cl}$ has already been introduced by Wang–Erickson [WE18]. He shows that it is a disjoint union of formal algebraic stacks, each of finite type over the formal Spec of complete Noetherian local ring (in fact, an appropriate *pseudodeformation ring*). In general that $\mathcal{X}_{G,F,S}$ has an analogous structure, and in particular is a formal derived algebraic stack. More precisely, if $\bar{\theta}$ is a residual pseudorepresentation of $\mathrm{Gal}_{F,S}$, then we let $X_{\bar{\theta}}$ denote the associated formal deformation space of $\bar{\theta}$. Passing from a Galois representation to its associated pseudorepresentation gives rise to a morphism $\mathcal{X}_{G,F,S} \rightarrow \coprod_{\bar{\theta}} X_{\bar{\theta}}$, and we let $\mathcal{X}_{\bar{\theta}}$ denote the preimage of $X_{\bar{\theta}}$. We thus obtain a decomposition

$$(9.1.1) \quad \mathcal{X}_{G,F,S} := \coprod_{\bar{\theta}} \mathcal{X}_{\bar{\theta}},$$

which is precisely the decomposition of $\mathcal{X}_{G,F,S}$ into its connected components.

Just as in the local case, if $\rho : \mathrm{Gal}_{F,S} \rightarrow {}^cG(A)$ is an A -valued point of $\mathcal{X}_{G,F,S}$, for some classical \mathcal{O} -algebra A in which p is nilpotent, then the tangent complex to $\mathcal{X}_{G,F,S}$ at ρ is computed via group cohomology, in this case the complex of continuous cochains $C^\bullet(\mathrm{Gal}_{F,S}, \mathfrak{g})[1]$. From now on we assume that $p > 2$. Then, just like in the local case, the cohomology of $\mathrm{Gal}_{F,S}$ vanishes in degrees > 2 , and so we find that $\mathcal{X}_{G,F,S}$ is formally quasi-smooth over $\mathrm{Spf} \mathcal{O}$. However, unlike in the local case, $\mathcal{X}_{G,F,S}$ need not be of its expected dimension (the expected dimension at

a particular point being computable as the Euler characteristic of $C^\bullet(\mathrm{Gal}_{F,S}, \mathfrak{g})[1]$, which is given by the global Euler characteristic formula, and so depends on the local behaviour at the real places of F of the Galois representation corresponding to the point under consideration); more precisely, it can have components that are of greater than the expected dimension. Because it is formally quasi-smooth, though, it *is* a derived local complete intersection, i.e. can be written as the quotient by an appropriately defined conjugation action of ${}^cG^\wedge$ (the p -adic completion of cG) on a derived formal affine scheme of the form

$$\mathrm{Spf} \mathcal{O}[[X_1, \dots, X_d]]/(f_1, \dots, f_r),$$

where $1 + d - r - \dim {}^cG$ is equal to the expected dimension of $\mathcal{X}_{G,F,S}$. (Here the first summand “1” corresponds to the dimension of \mathcal{O} .) The relations f_1, \dots, f_r need not form a regular sequence, though, and so the quotient ring must be understood in a derived sense. The derived structure of $\mathcal{X}_{G,F,S}$ is thus supported on those components of $\mathcal{X}_{G,F,S}^{cl}$ which are of greater-than-expected dimension.

9.2. The global-to-local map. The restriction of Galois representations is expected to give rise to a morphism

$$(9.2.1) \quad f : \mathcal{X}_{G,F,S} \rightarrow \prod_{v \in S} \mathcal{X}_{G,F_v}.$$

We can only define this morphism when $G = \mathrm{GL}_d$ locally at places $v|p$, since it is only in this case that we have a definition of the local stacks \mathcal{X}_{G,F_v} . Assuming that this is the case, the definition of this morphism is non-trivial, since the local stack \mathcal{X}_{G,F_v} is defined in terms of (φ, Γ) -modules rather than directly in terms of local Galois representations, and even in these cases, is defined only at the classical level. However, there is a tautological restriction morphism to a derived version of Wang-Erickson’s p -adic local stacks [WE18], and the results of [BIP21] show that these stacks are actually classical (*cf.* Remark 4.1.7); and these classical stacks are substacks of our local stacks \mathcal{X}_{G,F_v} by Remark 4.1.2.

In principle, we should also consider local stacks \mathcal{X}_{G,F_v} for the real archimedean primes v of F . However, since we are assuming that $p > 2$, these will simply be classical stacks [Zhu20, Prop. 2.3.2], and classify conjugacy classes of involutions lying above the complex conjugation c_v under the projection of cG onto its Galois factor. In fact, we only want to consider a particular conjugacy class, namely the “odd” class in cG — see [BV13, Prop. 6.1]. We let $\mathcal{X}_{G,F,S}^{\mathrm{odd}}$ denote the union of connected components of $\mathcal{X}_{G,F,S}$ classifying Galois representations which are odd at every real prime of F .

REMARK 9.2.2. As will be explained in [EZ], standard Fontaine–Mazur-type conjectures imply that the morphism f is representable by algebraic (derived) stacks.

REMARK 9.2.3. The tangent complex of f (i.e. the relative tangent complex of $\mathcal{X}_{G,F,S}^{\mathrm{odd}}$ over $\prod_{v \in S} \mathcal{X}_{G,F_v}$) at a point is computed (up to a shift) by the cone of the restriction morphism from global-to-local Galois cohomology; namely we have a distinguished triangle

$$\cdots \rightarrow \mathfrak{t}_f \rightarrow C^\bullet(\mathrm{Gal}_{F,S}, \mathfrak{g})[1] \rightarrow \prod_{v \in S} C^\bullet(\Gamma_{F_v}, \mathfrak{g})[1] \rightarrow \cdots,$$

so that the relative tangent complex \mathbf{t}_f is a Selmer-type complex of the sort considered in [GV18, App. B], which can possibly admit a non-zero H^3 term, which in turn, by global duality, is dual to another Galois H^0 , namely $H^0(\mathrm{Gal}_{F,S}, \widehat{\mathfrak{g}}(1))$. (A similar analysis, in the more traditional language of lifting rings, is made in [Kis07b, §4]. Indeed, this local analysis of f is a more geometric version of the traditional method of presenting a global deformation ring or lifting ring over a tensor product of local lifting rings. Such presentations are important in modularity lifting theorems (and in their applications to potential modularity); an example in our earlier presentation is the morphism 3.1.3.

The condition of f being relatively quasi-smooth at $\bar{\rho}$ is an interesting one. For example, if $G = \mathrm{GL}_d$, so that $\widehat{\mathfrak{g}} = \mathrm{End}(\bar{\rho})$, then $H^0(\mathrm{Gal}_{F,S}, \widehat{\mathfrak{g}}(1)) = \mathrm{End}_{\mathrm{Gal}_{F,S}}(\bar{\rho}, \bar{\rho}(1))$, and so we are considering whether or not $\bar{\rho}$ admits a non-trivial $\mathrm{Gal}_{F,S}$ -equivariant map to its Tate twist $\bar{\rho}(1)$. If $F = \mathbf{Q}$ and $G = \mathrm{GL}_2$, then (recalling $p > 2$) this is impossible unless either $\bar{\rho}^{\mathrm{ss}}$ is a twist of the direct sum $1 \oplus \bar{\chi}$ (where $\bar{\chi}$ is the mod p cyclotomic character) or $p = 3$ and $\bar{\rho}$ is induced from $\mathrm{Gal}_{\mathbf{Q}(\sqrt{-3})}$. These are precisely the cases excluded by the local-to-global principle for liftings of $\bar{\rho}$ of Diamond–Taylor [DT94, Thm. 1], and it is known that such a local-to-global principle fails in this case. In the case that $\bar{\rho}$ is induced from $\mathrm{Gal}_{\mathbf{Q}(\sqrt{-3})}$, this goes back to work of Carayol and Serre, and an explicit example is given in [All19, §6.4]. In the reducible case the question of the possible levels has been investigated by Ribet and Yoo, see in particular [Yoo19]. (See also [LHL16] for analogous results in the case $p = 2$.)

More generally, the condition that $H^0(\mathrm{Gal}_{F,S}, \widehat{\mathfrak{g}}(1)) = 0$ is closely related to the “big image” conditions in the Taylor–Wiles method, and to the *genericity* condition of [CS17].

9.3. Computing the cohomology of Shimura varieties via categorical Langlands. We state a conjectural formula for the cohomology of Shimura varieties, which is related to the various conjectures of [Zhu20], §4.7. (This conjecture will be discussed in more detail in [EZ].) While the general conjecture seems far out of reach (since, for example, it implies all kinds of automorphy theorems!), we discuss the case of the modular curve in Section 9.4.

9.3.1. *Cohomology of Shimura varieties.* Suppose now that $F = \mathbf{Q}$, and that the reductive group G admits a Shimura datum. We let F (which is now freed up as a piece of notation!) denote the reflex field of this Shimura datum, and let $\mu : (\mathbf{G}_m)_{/F} \rightarrow G_{/F}$ denote the associated Hodge cocharacter (well defined up to conjugacy). In the definition of cG , we choose the extension \tilde{F} of \mathbf{Q} large enough that it contains F , and then write

$${}^cG_F := \widehat{G} \rtimes (\mathbf{G}_m \times \mathrm{Gal}(\tilde{F}/F)) \subseteq \widehat{G} \rtimes (\mathbf{G}_m \times \mathrm{Gal}(\tilde{F}/\mathbf{Q})) =: {}^cG.$$

The Hodge cocharacter μ , thought of as a weight of \widehat{G} , gives rise to an irreducible (highest weight) representation V_μ of \widehat{G} , which naturally extends to a representation of cG_F . For any set S of primes (containing p and ∞ as usual), the moduli stack $\mathcal{X}_{G,\mathbf{Q},S}$ parameterizes a universal family of (derived) C -homomorphisms $\mathrm{Gal}_{\mathbf{Q}} \rightarrow {}^cG$, which we can restrict to Gal_F so as to obtain a family of representations $\mathrm{Gal}_F \rightarrow {}^cG_F$. Composing this family with the representation V_μ of cG_F , we obtain a vector bundle \mathcal{V}_μ over $\mathcal{X}_{G,\mathbf{Q},S}$, which is endowed with an action of Gal_F .

Choose a level subgroup K_f , i.e. a compact open subgroup of $G(\mathbf{A})$, which we suppose may be factored as $K_f = \prod_v K_v$, with K_v compact open in $G(\mathbf{Q}_v)$ for

each finite place v . We choose (as we may) S large enough that G is unramified outside S , and so that, furthermore, there is an extension \mathcal{G} of G to a reductive group scheme over $\mathbf{Z}[1/S]$ such that $K_v = \mathcal{G}(\mathbf{Z}_v)$ for $v \notin S$. The locally symmetric quotient $Y(K_f) := G(\mathbf{Q}) \backslash (X \times G(\mathbf{A}_f)) / K_f$ is a Shimura variety.

For each finite prime $v \in S$, we let W_v be a smooth representation of K_v on a finite torsion \mathcal{O} -module. We let $W := \bigotimes_{v \in S} W_v$ denote the corresponding representation of K_f (which acts through its projection onto $\prod_{v \in S} K_v$; the tensor product is taken over \mathcal{O} , of course), and we let \mathcal{W} denote the associated local system $\mathcal{W} := G(\mathbf{Q}) \backslash (X \times G(\mathbf{A}_f) \times W) / K_f$ over $Y(K_f)$. For each v , the compact induction $c\text{-Ind}_{K_v}^{G(\mathbf{Q}_v)} W_v$ is a smooth $G(\mathbf{Q}_v)$ -representation. Applying the functors of Conjectures 6.1.14 and 8.3.1, we obtain (conjecturally!) a coherent sheaf on $\mathcal{X}_{G, \mathbf{Q}_v}$, which we denote simply by \mathfrak{A}_v . We may form the exterior tensor product $\mathfrak{A} := \boxtimes_{v \in S_f} \mathfrak{A}_v$ on $\prod_{v \in S_f} \mathcal{X}_{G, \mathbf{Q}_v}$.

As above, let $f : \mathcal{X}_{G, \mathbf{Q}, S}^{\text{odd}} \rightarrow \prod_{v \in S} \mathcal{X}_{G, \mathbf{Q}_v}$ denote the global-to-local restriction morphism. We then have the following conjectural formula for the cohomology of the Shimura variety.

CONJECTURE 9.3.2. *Let d be the complex dimension of $Y(K_f)$. Then there is an isomorphism*

$$(9.3.3) \quad R\Gamma_c(Y(K_f), \mathcal{W})[d] \xrightarrow{\sim} R\Gamma(\mathcal{X}_{G, \mathbf{Q}, S}^{\text{odd}}, f^! \mathfrak{A} \otimes \mathcal{V}_\mu)$$

compatible with the Gal_F and (derived) Hecke actions on source and target.

By taking the inverse limit of the mod p^n cases, the conjecture naturally extends to the case of cohomology with p -adic coefficients. We can also take a limit over the tame levels at places in S , to study completed cohomology, as in Expected Theorem 9.4.2 below.

REMARK 9.3.4. As briefly mentioned in our discussion of Conjecture 8.3.1 in Subsection 8.3, the local functors appearing in the statement of Conjecture 9.3.2 depend on auxiliary choices of Whittaker-type data. In the statement of the conjecture, we should make a global choice of this data, and then compute the local functors \mathfrak{A}_v in terms of the induced local data. The resulting sheaf $f^! \mathfrak{A}$ should then be well-determined independent of the global choice. (A different choice of global data should cause each value of \mathfrak{A}_v to be twisted by a certain line bundle, but the tensor product of the pullbacks to $\mathcal{X}_{G, \mathbf{Q}, S}^{\text{odd}}$ of these various line bundles should then be trivial; see the discussion at the top of p. 63 of [Zhu20] for the analogous discussion in the function field context.)

REMARK 9.3.5. Conjecture 9.3.2 implies in particular that various representations of Gal_F occur in the cohomology of Shimura varieties, so in particular it implies cases of the Fontaine–Mazur conjecture. To make this implication precise, one needs to know something about the supports of the sheaves \mathfrak{A}_v . For $v|p$, Conjecture 6.1.14 (4) is the relevant statement: it shows that the \mathfrak{A}_v will detect all points of $\mathcal{X}_{G, \mathbf{Q}, S}^{\text{odd}}$ whose restrictions to \mathcal{X}_v satisfy an appropriate p -adic Hodge theoretic condition. Analogous statements in the case $v \nmid p$, describing the support of \mathfrak{A}_v in terms of ramification conditions, should also hold. We don't spell these out here, but refer to the discussions of the $v \nmid p$ case in [BZCHN20], [FS21], [Hel20], and [Zhu20]; the essential point in this case is that the categorical theory should be compatible, in a suitable sense, with the classical local Langlands correspondence.

REMARK 9.3.6. The Hecke algebras that act naturally on the source and target of (9.3.3) are the derived Hecke algebras $\mathrm{RHom}_G(c\text{-Ind}_{K_v}^{G(\mathbf{Q}_v)} W_v, c\text{-Ind}_{K_v}^{G(\mathbf{Q}_v)} W_v)$ for finite places v (where we interpret W_v simply as the trivial representation of K_v in the case $v \notin S$; to see that the derived Hecke algebras do indeed act at places $v \notin S$, one compares the situations at S and $S \cup \{v\}$ just after [Zhu20, Conj. 4.7.5]). We refer the reader to [Zhu20] for a discussion of these derived Hecke algebras, and several conjectures related to them, in the case when $v \neq p$.

REMARK 9.3.7. While the factor \mathcal{V}_μ has a global aspect, since it carries an action of the global absolute Galois group Gal_F , if we ignore this then it can be regarded as the contribution from ∞ to the formula of Conjecture 9.3.2. We believe that an analogous isomorphism should hold for *any* connected reductive group G (i.e. for the cohomology of all the locally symmetric spaces $Y(K_f)$, whether or not they admit the structure of Shimura variety), with an appropriate choice of factor coming from ∞ (which is now presumably just a vector bundle on $\mathcal{X}_{G, \mathbf{Q}, S}^{\mathrm{odd}}$, pulled back from an appropriately chosen vector bundle on the odd moduli stack at ∞ , and not admitting the additional structure of a Galois representation). In particular, we expect to recover the derived phenomena observed and conjectured by Galatius and Venkatesh in [GV18, Ven19].

9.4. An example in the global case: $(\mathrm{PGL}_2)_{/\mathbf{Q}}$. In [EGZ], M.E., T.G. and Xinwen Zhu use the Taylor–Wiles–Kisin patching method to study Conjecture 9.3.2 in the case of the modular curve. We briefly explain some of this work in progress in a special case, which is adapted to our discussion of the patching method in Section 3.

9.4.1. *Relationship to patching.* To fix ideas, set $G = \mathrm{PGL}_2$, so we are in the modular curve case. Choose $\bar{r} : \mathrm{Gal}_{\mathbf{Q}, \{p, \infty\}} \rightarrow \mathrm{GL}_2(k)$ as in Section 3.1, so that in particular $\bar{r}|_{\mathrm{Gal}_{\mathbf{Q}(c_p)}}$ is absolutely irreducible, and \bar{r} has determinant $\bar{\varepsilon}^{-1}$. Write $\bar{\rho} := \bar{r}|_{\mathrm{Gal}_{\mathbf{Q}_p}}$. We consider (9.2.1) in the special case that $S = \{p, \infty\}$, so that we have the morphism

$$f : \mathcal{X}_{\mathrm{GL}_2, \mathbf{Q}, \{p, \infty\}}^{\varepsilon^{-1}, \mathrm{odd}} \rightarrow \mathcal{X}_{\mathrm{GL}_2, \mathbf{Q}_p}^{\varepsilon^{-1}},$$

where we are only considering representations with determinant ε^{-1} , so that the target is the stack denoted \mathcal{X} in Section 7.2. Write \mathcal{V} for the rank 2 vector bundle on $\mathcal{X}_{\mathrm{GL}_2, \mathbf{Q}, \{p, \infty\}}^{\varepsilon^{-1}, \mathrm{odd}}$ given by the universal $\mathrm{Gal}_{\mathbf{Q}}$ -representation.

As in Section 3.1, we write $R(\bar{\rho})$ for the universal deformation ring of $\bar{\rho}$ with fixed determinant ε^{-1} . By definition there is a versal morphism $g : \mathrm{Spf} R(\bar{r}_p) \rightarrow \mathcal{X}_p$ that maps the closed point in its domain to the local restriction $f(\bar{r})$. The patching construction gives an $R(\bar{r}|_{\mathrm{Gal}_{\mathbf{Q}_p}})$ -module M_∞ with an action of $\mathrm{PGL}_2(\mathbf{Q}_p)$.

We expect to prove the following result in [EGZ], describing the localization at the maximal ideal \mathfrak{m} of the prime-to- p Hecke algebra corresponding to \bar{r} of the completed homology group \tilde{H}_1 . As in (9.1.1) we write $\mathcal{X}_{\bar{r}}$ for the irreducible component of $\mathcal{X}_{\mathrm{GL}_2, \mathbf{Q}, \{p, \infty\}}^{\varepsilon^{-1}, \mathrm{odd}}$ corresponding to \bar{r} .

EXPECTED THEOREM 9.4.2 (M.E.–T.G.–Zhu). *Let L_∞ be the sheaf of $\mathrm{PGL}_2(\mathbf{Q}_p)$ -representations of Expected Theorem 7.2.1. Then there is a $\mathrm{PGL}_2(\mathbf{Q}_p)$ -equivariant isomorphism*

$$(9.4.3) \quad \tilde{H}_{1, \mathfrak{m}} \xrightarrow{\sim} R\Gamma(\mathcal{X}_{\bar{r}}, f^! L_\infty \otimes \mathcal{V})[-1]$$

compatible with the actions of $\text{Gal}_{\mathbf{Q}}$ and the (derived) Hecke algebras at primes $\ell \neq p$ on each side.

In addition, there is a $\text{PGL}_2(\mathbf{Q}_p)$ -equivariant isomorphism $M_\infty \xrightarrow{\sim} g^* L_\infty$.

REMARK 9.4.4. We anticipate that the hypotheses in Expected Theorem 9.4.2 can be considerably relaxed. In particular, we can allow \bar{r} to be at ramified at places not dividing p , and we hope to relax the other assumptions on \bar{r} as much as possible in the final version of [EGZ].

We expect to prove Expected Theorem 9.4.2 as follows. The basic idea is that M_∞ is “patched” out of the cohomology of modular curves, and we want to show that the patching process “undoes” the $f^!$ appearing in the conjectural formula for cohomology. We show (using existing modularity lifting theorems) that $\mathcal{X}_{\bar{r}}$ is equal to $[\text{Spf } R_{\{p,\infty\}}(\bar{r})/\mu_2]$. We then show that $R_{\{p,\infty\}}(\bar{r})$ is the quotient of $R(\bar{\rho})$ by a certain regular sequence which is also a regular sequence on M_∞ , in order to compute the pullback $f^!$ via a Koszul complex, and ultimately reduce (9.4.3) to the local-global compatibility results originally proved by M.E. in [Eme11a], or more precisely to their interpretation and reproof in [CEG⁺18].

9.5. The structure of $\mathcal{X}_{\text{GL}_2, \mathbf{Q}, S}^{\text{odd}}$. Another aim of [EGZ] is to use Taylor–Wiles patching and Iwasawa theory to describe the structure of the stacks $\mathcal{X}_{\text{GL}_2, \mathbf{Q}, S}^{\text{odd}}$. We briefly describe some of our expected results. For simplicity we again restrict to the stacks $\mathcal{X}_{\text{GL}_2, \mathbf{Q}, S}^{\varepsilon^{-1}, \text{odd}}$ with determinant ε^{-1} .

Similar to (9.1.1) we have the decomposition

$$\mathcal{X}_{\text{GL}_2, \mathbf{Q}, S}^{\text{odd}} := \coprod_{\bar{\theta}} \mathcal{X}_{\bar{\theta}}$$

into connected components, where $\bar{\theta}$ runs over the odd residual pseudorepresentation of $\text{Gal}_{\mathbf{Q}, S}$ with determinant $\bar{\varepsilon}^{-1}$. The simplest case is when $\bar{\theta}$ arises from an irreducible odd representation $\bar{r} : \text{Gal}_{\mathbf{Q}, S} \rightarrow \text{GL}_2(\overline{\mathbf{F}}_p)$. We then write $X_{\bar{\theta}} = X_{\bar{r}} = \text{Spf } R_{\bar{r}}$, where $R_{\bar{r}}$ is the deformation ring of \bar{r} , and we find that $\mathcal{X}_{\bar{\theta}} = \mathcal{X}_{\bar{r}} = [\text{Spf } R_{\bar{r}}/\widehat{\mathbf{G}}_m]$, where $\widehat{\mathbf{G}}_m$ denotes the $\mathfrak{m}_{R_{\bar{r}}}$ -adic completion of \mathbf{G}_m , acting trivially on $\text{Spf } R_{\bar{r}}$. In this case, then, there is no essential difference between the theory of the moduli stack $\mathcal{X}_{\bar{r}}$ and usual Galois deformation theory. Serre’s conjecture tells us that \bar{r} is modular, and then known “big $R = \text{big } \mathbf{T}$ ” theorems imply (at least under the usual Taylor–Wiles hypotheses on \bar{r}) that $R_{\bar{r}}$ may be identified with an appropriately defined Hecke algebra acting on a space of p -adic modular forms.

The case when $\bar{\theta}$ is reducible is more subtle, and our goal in this case is again to describe $\mathcal{X}_{\bar{\theta}}$ as much as possible in terms of modular forms. To give a sense of what this means, set $\bar{\theta} = 1 + \bar{\varepsilon}^{-1}$, and let $\mathcal{Y}_{\bar{\theta}}$ be the (derived) substack obtained by imposing the local conditions of being unipotently tame at the primes $\ell \in S \setminus \{p, \infty\}$, as well as the local condition of being finite flat at p . Let $N = \prod_{\ell \in S \setminus \{p, \infty\}} \ell$, and choose \mathcal{O} large enough that each cuspidal eigenform f of weight two and level N has coefficients in \mathcal{O} . Then each such cuspform f whose associated mod p pseudorepresentation $\bar{\theta}_f$ is equal to $1 + \bar{\varepsilon}^{-1}$ contributes an irreducible component to $\mathcal{Y}_{\bar{\theta}}$ of the form $[\text{Spf } \mathcal{O}\langle X, Y \rangle / (XY - c) / \widehat{\mathbf{G}}_m]$, where c is a certain modulus of reducibility, $\mathcal{O}\langle X, Y \rangle$ denotes the p -adic completion of $\mathcal{O}[X, Y]$, and $\widehat{\mathbf{G}}_m$ denotes the p -adic completion of \mathbf{G}_m , acting via $t \cdot (X, Y) = (t^2 X, t^{-2} Y)$. This component parameterizes lattices in the p -adic Galois representation attached to f .

In addition to such components, there is an “Eisenstein” component, whose underlying classical stack is of the form $[\mathrm{Spf} \mathcal{O}\langle X_1, \dots, X_n \rangle / \widehat{\mathbf{G}}_m]$, where n denotes the order of $S \setminus \{p, \infty\}$ (i.e. the number of factors of N), and $\widehat{\mathbf{G}}_m$ acts via $t \cdot X_i = t^2 X_i$. This component parameterizes extensions of ε^{-1} by 1; the space of such extensions is spanned by the Kummer classes associated to the n factors of N . If $n > 1$, then this component is of larger-than-expected dimension, and so (following the discussion above) will carry non-trivial derived structure.

There can be yet another irreducible component of \mathcal{X} , parameterizing extensions of 1 by ε^{-1} (i.e. extensions “the wrong way”). Assuming $p \geq 5$ for simplicity, these are governed by those $\ell \in S$ for which $\ell \equiv -1 \pmod{p}$. If there are m such ℓ , then the family of such extensions is $(m-1)$ -dimensional (there is an m -dimensional space of extensions in characteristic p , which is then quotiented by a $\widehat{\mathbf{G}}_m$ -action) and is supported (set-theoretically) purely in characteristic p . If $m > 1$ it is a component of $\mathcal{Y}_{\overline{\rho}}$. If $m > 2$ it is also of greater-than-expected dimension, and so admits derived structure.

We expect that the preceding discussion enumerates all possible components of $\mathcal{Y}_{\overline{\rho}}$. We hope that this (and more general statements) can be proved by suitable applications of (known cases of the) the Fontaine–Mazur conjecture (especially in the residually reducible case, where it has been proved by Skinner–Wiles [SW97] and Pan [Pan22a]), together with the main conjecture of Iwasawa theory [MW84].

9.6. Eigenvarieties and sheaves of overconvergent p -adic automorphic forms. We now state an analogue of Conjecture 9.3.2, giving a conjectural description of the cohomology of some (overconvergent) p -adic coefficient systems. This gives a conjectural description of the sheaves of overconvergent p -adic automorphic forms that appear on eigenvarieties. We maintain the notation from the preceding section.

There are basically two approaches to p -adic automorphic forms and eigenvarieties: one can use the cohomology of locally symmetric spaces with coefficients in p -adic coefficient systems (see e.g. [Che04] or [Han17]), or one can use the completed cohomology groups introduced by M.E. and further developed by M.E. and Frank Calegari (see [Eme06b, Eme14, CE12]) and apply the locally analytic Jacquet functor introduced by M.E. in [Eme06a] (see [Eme06b] for a construction of eigenvarieties using this approach). We briefly discuss the relation between the two approaches and their relation with the functor $\mathfrak{A}_G^{\mathrm{rig}}$ of Conjecture 6.2.4.

REMARK 9.6.1. In general an eigenvariety can be associated with a tame level K^p , Bernstein blocks $\Omega_v = [M_v, \sigma_v]$ of $\mathrm{Rep}_L^{\mathrm{sm}} G(F_v)$ at places $v|p$, and algebraic representations W_v (defined up to twist with an algebraic character of M_v) of the supercuspidal support M_v of Ω_v . We will focus on the classical case where all the Bernstein blocks are principal (and hence W_v is not an extra datum) and refer to [BD21] for the treatment of arbitrary Bernstein components. We point out that in this more general context it is also possible to state similar conjectures to Conjectures 9.6.8 and 9.6.27 below.

9.6.2. The case of Shimura sets. Let us first discuss the special case where the reductive group G is compact at infinite places, and hence the locally symmetric spaces $Y(K_f)$ are just point sets. We refer to these sets as Shimura sets even though strictly speaking they are not Shimura varieties. In order to fix the setup we assume from now on that F is totally real and the group G is

- either a definite unitary group associated to a hermitian space over a CM extension E of F ,
- or the group associated to a definite quaternion algebra over F ,

and that G is split at places dividing p .

Let us fix for each place $v|p$ a split maximal torus and a Borel $T_v \subset B_v \subset G_v = G(F_v) = \mathrm{GL}_d(F_v)$. We also fix the choice of an Iwahori subgroup I_v adapted to B_v and write $T_v^0 = T_v \cap I_v$ for the maximal compact subgroup of the torus T_v , and $B_{v,0} = B_v \cap I_v$. Moreover, we fix a *tame level* K^p which is a compact open subgroup of $G(\mathbf{A}_F^{\infty,p})$. Let $\kappa_v : T_v^0 \rightarrow L^\times$ be a locally analytic character. Then we can define the following space of *module valued overconvergent p -adic automorphic forms*

(9.6.3)

$$M^\dagger(G, K^p, (\kappa_v)_{v|p}) = \left\{ \begin{array}{l} f : G(F) \backslash G(\mathbf{A}_F^\infty) / K^p \rightarrow \widehat{\bigotimes}_{v|p} (\mathrm{Ind}_{B_{v,0}}^{I_v} \kappa_v)^{\mathrm{an}} \\ \text{s. th. } f(gu) = u^{-1} f(g) \text{ for all } g \in G(\mathbf{A}_F^\infty), u \in \prod_{v|p} I_v \end{array} \right\}$$

compare e.g. [Loe11, 3.3]. This space comes equipped with an action of the so-called Atkin–Lehner ring $\mathcal{A}^+(G_v)$ whose definition (and action) we now briefly sketch.

For each place $v|p$ the choice of the Borel subgroup B_v and its subgroup $N_{v,0} = N_v \cap B_{v,0}$ cuts out a submonoid $T_v^+ \subset T_v$ consisting of all $t \in T_v$ such that $tN_{v,0}t^{-1} \subset N_{v,0}$. Clearly T_v^+ contains the maximal compact subgroup T_v^0 . We fix the choice of a free abelian subgroup $\Sigma_v \subset T_v$ such that $\Sigma_v \cong T_v/T_v^0$ (in the case of the diagonal torus and after choosing a uniformizer ϖ_v we might for example choose the diagonal matrices all of whose entries are powers of ϖ_v) and set $\Sigma_v^+ = \Sigma_v \cap T_v^+$, which is a free abelian monoid. Then we define

$$\begin{aligned} \mathcal{A}^+(G_v) &= L[\Sigma_v^+] \\ \mathcal{A}^+(G) &= \bigotimes_{v|p} \mathcal{A}^+(G_v). \end{aligned}$$

It turns out [Loe11, 3.4] that $\mathcal{A}^+(G_v)$ is canonically isomorphic to the subalgebra of the Hecke algebra $\mathcal{H}(G_v)$ consisting of functions that are I_v -biinvariant and that are supported on the submonoid $\Sigma_v^+ I_v$ of G_v generated by Σ_v^+ and I_v . The action of I_v on $\mathrm{Ind}_{B_{v,0}}^{I_v} \kappa_v$ canonically extends to $\Sigma_v^+ I_v$ [Loe11, Theorem 2.4.7] and hence this monoid acts on $M^\dagger(G, K^p, (\kappa_v)_{v|p})$ via $(\gamma.f) = \gamma f(g\gamma)$. Viewing elements of $\mathcal{A}^+(G)$ as locally constant functions on $G(\prod_{v|p} F_v)$ we hence obtain an action of $\mathcal{A}^+(G)$ on $M^\dagger(G, K^p, (\kappa_v)_{v|p})$.

We now explain a different characterization of the above space of overconvergent p -adic automorphic forms together with its $\mathcal{A}^+(G)$ -action that involves completed cohomology and is in fact less involved to write down (in particular it is easier to write down the action of the Atkin–Lehner ring in this setup). We however preferred to give the above definition first as it is directly related to the cohomology with coefficients in a p -adic coefficient system defined by the representations $\mathrm{Ind}_{B_{v,0}}^{I_v} \kappa_v$ (which makes it easier to formulate Conjecture 9.6.8 below).

Consider the space

$$\hat{S}(K^p, L) = \{f : G(F) \backslash G(\mathbf{A}_F^\infty) / K^p \rightarrow L \text{ continuous}\}$$

which is an admissible Banach space representation of $G(F_p)$ and coincides with the completed cohomology (base changed to L) of $Y(K^p)$ as our group is compact

at infinity. The locally analytic vectors in this Banach space representations are given by

$$(9.6.4) \quad \hat{S}(K^p, L)^{\text{an}} = \{f : G(F) \backslash G(\mathbf{A}_F^\infty) / K^p \rightarrow L \text{ locally analytic}\}.$$

As above we write $N_v \subset B_v$ for the unipotent radical of B_v and $N_{v,0} = N_v \cap B_{v,0}$. If we set $N_0 = \prod_{v|p} N_{v,0}$ we can then define an action of T^+ on the N_0 -invariants $\hat{S}(K^p, L)^{N_0}$ by

$$(9.6.4) \quad t.f = \sum_{n \in N_0 / tN_0t^{-1}} ntf.$$

The following proposition compares the modules (9.6.3) and (9.6.4) equipped with their $\mathcal{A}^+(G)$ -action (where $\mathcal{A}^+(G_v)$ acts on the N_0 -invariants of (9.6.4) via the inclusion $L[\Sigma_v^+] \subset L[T_v^+]$).

PROPOSITION 9.6.5. *Let $N_0 = \prod_v N_{v,0}$. Then there is an isomorphism*

$$M^\dagger(G, K^p, (\kappa_v)_{v|p}) \xrightarrow{\sim} (\hat{S}(K^p, L)^{\text{an}})^{N_0} [(\kappa_v)_{v|p}]$$

equivariant for the action of the Atkin–Lehner ring $\mathcal{A}^+(G) = \otimes_{v|p} \mathcal{A}^+(G_v)$ and for the Hecke-action away from p . Here the $(-)[(\kappa_v)_{v|p}]$ on the right hand side means the subspace of the N_0 -invariants on which $T^0 = \prod_v T_v^0$ acts via $\prod_v \kappa_v$.

PROOF. This is for example proven in [Loe11, Prop. 3.10.1, Prop. 3.10.2]. \square

Instead of a fixed locally analytic character one can apply the same construction to a family of characters of T^0 or even to the universal character

$$\kappa^{\text{univ}} : T^0 \rightarrow \Gamma(\hat{T}^0, \mathcal{O}_{\hat{T}^0})^\times,$$

where \hat{T}^0 denotes the space of continuous characters of T^0 . This means that in the definition of $M^\dagger(G, K^p, \kappa^{\text{univ}})$ we replace the representation $\text{Ind}_{B_{v,0}}^{I_v} \kappa_v$ by the I_v -representation $\text{Ind}_{N_{v,0}}^{I_v} \mathbf{1}$ (which contains all the representations $\text{Ind}_{B_{v,0}}^{I_v} \kappa_v$ for the various characters κ_v).

PROPOSITION 9.6.6. *There is an isomorphism*

$$M^\dagger(G, K^p, \kappa^{\text{univ}}) \xrightarrow{\sim} (\hat{S}(K^p, L)^{\text{an}})^{N_0},$$

equivariant for the action of $L[T^+]$ and for the action of the Hecke-operators away from p .

PROOF. This is a special case of [Fu21, Theorem 1.2]; see also the discussion of [Eme06b, §3.2]. \square

We now conjecturally relate these spaces of overconvergent p -adic automorphic forms to the functor $\mathfrak{A}_G^{\text{rig}}$. On the Galois side, let us write $\mathfrak{X}_{G,F,S} = \mathcal{X}_{G,F,S}^{\text{rig}}$ for the rigid analytic generic fiber of the stack $\mathcal{X}_{G,F,S}$. Note that $\mathcal{X}_{G,F,S}$ in general is a derived stack (not a classical stack), and we hence have to adapt the definition of its generic fiber (which will be a derived rigid analytic Artin stack), generalizing (5.1.11) to the derived setting. As we do not want to discuss derived rigid analytic geometry here, we do not spell out the precise definition of $\mathfrak{X}_{G,F,S}$.

We note that similarly to (9.2.1), given $v|p$ there is a morphism

$$f_{\text{rig},v} : \mathfrak{X}_{G,F,S} = (\mathcal{X}_{G,F,S})_\eta^{\text{rig}} \rightarrow \mathfrak{X}_{G,F_v},$$

that is given by composing the corresponding map to $(\mathcal{X}_{G,F_v})_\eta^{\text{rig}}$ with the morphism (5.1.27). We write

$$f_{\text{rig}} : \mathfrak{X}_{G,F,S} \rightarrow \prod_v \mathfrak{X}_{G,F_v}$$

for the product of the maps $f_{\text{rig},v}$. By abuse of notation we also write f_{rig} for the restriction of f_{rig} to $\mathfrak{X}_{G,F,S}^{\text{odd}} = (\mathcal{X}_{G,F,S}^{\text{odd}})_\eta^{\text{rig}} \subset \mathfrak{X}_{G,F,S}$.

Let us write $\sigma_{\kappa_v} = \text{Ind}_{B_{v,0}}^{I_v} \kappa_v$ and $\sigma_v = \text{Ind}_{N_{v,0}}^{I_v} \mathbf{1}$. We can use these representation and the functors $\mathfrak{A}_{G(F_v)}^{\text{rig}}$ to define the coherent sheaves

$$(9.6.7) \quad \begin{aligned} \mathfrak{A}_{\text{rig}}((\kappa_v)_{v|p}) &= \left(\bigotimes_{v \nmid p} \mathfrak{A}_{G(F_v)}(c\text{-Ind}_{K_v}^{G(F_v)} \mathbf{1}) \right)^{\text{rig}} \otimes \left(\bigotimes \mathfrak{A}_{G(F_v)}^{\text{rig}}(c\text{-Ind}_{I_v}^{G(F_v)} \sigma_{\kappa_v}) \right), \\ \mathfrak{A}_{\text{rig}}^{\text{univ}} &= \left(\bigotimes_{v \nmid p} \mathfrak{A}_{G(F_v)}(c\text{-Ind}_{K_v}^{G(F_v)} \mathbf{1}) \right)^{\text{rig}} \otimes \left(\bigotimes \mathfrak{A}_{G(F_v)}^{\text{rig}}(c\text{-Ind}_{I_v}^{G(F_v)} \sigma_v) \right). \end{aligned}$$

CONJECTURE 9.6.8. *There are isomorphisms*

$$\begin{aligned} M^\dagger(G, K^p, (\kappa_v)_{v|p}) &\cong R\Gamma_c(\mathfrak{X}_{G,F,S}^{\text{odd}}, f_{\text{rig}}^* \mathfrak{A}_{\text{rig}}((\kappa_v)_{v|p})), \\ M^\dagger(G, K^p, \kappa^{\text{univ}}) &\cong R\Gamma_c(\mathfrak{X}_{G,F,S}^{\text{odd}}, f_{\text{rig}}^* \mathfrak{A}_{\text{rig}}^{\text{univ}}) \end{aligned}$$

compatible with the Hecke action and the action of the Atkin–Lehner ring.

REMARK 9.6.9. The Atkin–Lehner ring acts on the right hand side via the canonical action of $L[T^+]$ on $c\text{-Ind}_{I_v}^{G_v} \sigma_{\kappa_v}$, see for example [KS12, Lemma 2.2, Lemma 2.3].

We point out that this conjecture is somehow dual in nature to Conjecture 9.3.2: we use cohomology on the side of the symmetric space (which is just a point set in this case) and cohomology with compact support on the side of coherent sheaves (note that the stack $\mathfrak{X}_{G,F,S}^{\text{odd}}$ is usually the quotient of a Stein space by the action of a reductive group and hence we can define cohomology with compact support as in [Chi90], similar to its use in Section 7.1). This duality also explains why in this conjecture we use the $*$ -pullback along f_{rig} instead of the $!$ -pullback of Conjecture 9.3.2.

REMARK 9.6.10. In the definition of $M^\dagger(G, K^p, (\kappa_v)_{v|p})$ we could also have inserted non-trivial K_v -representations σ_v for some places $v \nmid p$ as coefficients. We then expect the same isomorphism as in Conjecture 9.6.8, but then of course the trivial K_v -representation $\mathbf{1}$ in (9.6.7) has to be replaced by σ_v .

9.6.11. *Finite slope spaces and Emerton’s locally analytic Jacquet functor.* We recall the definition of finite slope spaces and of the locally analytic Jacquet module [Eme06a].

The spaces $M^\dagger(G, K^p, (\kappa_v)_{v|p})$ and $(\hat{S}(K^p, L)^{\text{an}})^{N_0}$ are locally convex topological L -vector spaces (of compact type) that come equipped with an action of the monoid $L[T^+]$. The construction of a finite slope space V_{fs} attached to a locally convex L -vector space V equipped with a topological T^+ -action (or a T_v^+ -action for a place $v|p$) is a canonical procedure to extend the $L[T^+]$ -action to a locally analytic representation of $T \supset T^+$. Its formal definition is given by

$$V_{\text{fs}} = \text{Hom}_{L[T^+], b}(\Gamma(\hat{T}, \mathcal{O}_{\hat{T}}), V)$$

compare [Eme06a, Definition 3.2.1]. Here \hat{T} is the rigid analytic space of characters of T (note that $\Gamma(\hat{T}, \mathcal{O}_{\hat{T}})$ contains the distribution algebra $\mathcal{D}(T)$ but is strictly larger) and the subscript b indicates that we take the space of continuous homomorphisms and equip it with the strong topology. Then the $\Gamma(\hat{T}, \mathcal{O}_{\hat{T}})$ -module structure on V_{fs} restricts to an action of $\mathcal{D}(T)$ and hence to an action of T which in fact is locally analytic [Eme06a, Proposition 3.2.4]. If this representation is (essentially) admissible, the strong dual of the finite slope space is a coadmissible module over the Frechet–Stein algebra $\Gamma(\hat{T}, \mathcal{O}_{\hat{T}})$. As \hat{T} is a Stein space this implies that there exists a (necessarily unique) coherent sheaf \mathcal{M}_V^* on \hat{T} such that

$$(9.6.11) \quad (V_{\text{fs}})' \cong \Gamma(\hat{T}, \mathcal{M}_V^*)$$

as $\Gamma(\hat{T}, \mathcal{O}_{\hat{T}})$ -modules.

REMARK 9.6.12.

(i) If $V = \Pi^{N_{v,0}}$ for a locally analytic G_v -representation Π , then V is equipped with a T^+ -action defined by (9.6.4). Then essential admissibility of V_{fs} (as a locally analytic T_v -representation) can for example be assured by assuming that $\Pi|_H = \mathcal{C}^{\text{an}}(H, L)^m$ for some open compact subgroup $H \subset G_v$ and some $m > 0$, see [Eme06a, Proposition 3.2.24] and [BHS17b, Proposition 5.3] (which relies on [Eme06a]). In this situation [BHS17b, Proposition 5.3] moreover shows that locally on \hat{T}_v the sheaf \mathcal{M}_V^* is finite projective over \hat{T}_v^0 .

(ii) Assume that the T -representation on V_{fs} is essentially admissible. Then using Serre duality on the Stein space \hat{T} (as done for the open unit disc \mathbf{U} in (7.1.4)) we can rewrite (9.6.11) as

$$V_{\text{fs}} = R\Gamma_c(\hat{T}, \mathcal{M}_V),$$

where $\mathcal{M}_V = \mathbf{D}(\mathcal{M}_V^*)$ is the Serre dual (on \hat{T}) to \mathcal{M}_V^* . A priori this is only a complex of coherent sheaves. However, in good situations (e.g. in the situation of (i)) the sheaf \mathcal{M}_V^* turns out to be a Cohen–Macaulay module and hence \mathcal{M}_V is concentrated in one degree (though this degree might be strictly negative, depending on the dimension of the support of \mathcal{M}_V^*).

DEFINITION 9.6.13. The locally analytic Jaquet module $J_{B_v}(\Pi_v)$ of a locally analytic G_v -representation Π_v is the T_v -representation on $(\Pi_v^{N_{v,0}})_{\text{fs}}$.

With this definition at hand Proposition 9.6.6 directly implies the following:

PROPOSITION 9.6.14. *There is an isomorphism*

$$M^\dagger(G, K^p, \kappa^{\text{univ}})_{\text{fs}} \xrightarrow{\cong} J_B(\hat{S}(K^p, L)^{\text{an}})$$

of locally analytic T -representations.

If we restrict attention to the finite slope subspace of $M^\dagger(G, K^p, \kappa^{\text{univ}})$, then we can formulate a conjecture that gives a precise description of the coherent sheaf whose compactly supported cohomology conjecturally computes $M^\dagger(G, K^p, \kappa^{\text{univ}})_{\text{fs}}$ and that does not involve the conjectural functors $\mathfrak{A}_G^{\text{rig}}$ (though of course the conjecture is inspired by these conjectural functors and their expected properties with respect to parabolic induction).

As above, for $v \nmid p$, let us write $\mathfrak{A}_{\text{rig},v} = (\mathfrak{A}_{G(F_v)}(c\text{-Ind}_{K_v^G} \mathbf{1}))^{\text{rig}}$. For $v|p$ consider the morphism

$$f_{\text{rig},v} \times \iota_v : \mathfrak{X}_{G,F,S}^{\text{odd}} \times \hat{T} \longrightarrow \mathfrak{X}_{G,F_v} \times \hat{T}_v,$$

where ι_v is given by projection to the v -th factor followed by the automorphism

$$(\delta_1, \dots, \delta_n) \mapsto \delta_{B_v} \cdot (\delta_1, \delta_2(\varepsilon \circ \text{rec}_{F_v}), \dots, \delta_n(\varepsilon \circ \text{rec}_{F_v})^{n-1})$$

of \hat{T}_v , where δ_{B_v} denotes the modulus character. Moreover consider the diagram

$$\begin{array}{ccc} & \bar{\mathfrak{X}}_{B_v} & \\ \pi_v \swarrow & & \searrow \alpha_v \\ \mathfrak{X}_{G, F_v} \times \hat{T}_v & & \mathfrak{X}_{T_v}. \end{array}$$

and define

$$\mathfrak{A}_{\text{rig}, v} = \pi_{v,*} \alpha_v^*(\mathcal{O}_{\mathfrak{X}_{T_v}}([F_v : \mathbf{Q}_p](0, -1, \dots, -n + 1))[\text{rk}_{\mathbf{Q}_p} T_v]).$$

Here, as in Conjecture 6.2.4 (2), the twist $(0, -1, \dots, -n + 1)$ means that we take a sheaf on \hat{T}_v and consider it as a sheaf on $\mathfrak{X}_{T_v} = \hat{T}_v / \mathbf{G}_m^n$ by equipping it with the \mathbf{G}_m^n -equivariant structure given by the algebraic character $(0, -1, \dots, -n + 1)$.

REMARK 9.6.15. We note that it would be more canonical to talk about $\pi_{v,!}$ instead of $\pi_{v,*}$. But as π_v is proper by Theorem 5.3.11 these two functors agree and we do not need to worry about the definition of $\pi_{v,!}$ (that most likely would involve condensed structures).

Writing

$$g_{\text{rig}} = \prod_{v \nmid p} f_{\text{rig}, v} \times \prod_{v \mid p} g_{\text{rig}, v} : \mathfrak{X}_{G, F, S}^{\text{odd}} \times \hat{T} \rightarrow \prod_v \mathfrak{X}_{G, F_v} \times \hat{T}$$

we can formulate the following conjecture:

CONJECTURE 9.6.16. *There is an isomorphism*

$$\mathcal{M}^\dagger(G, K^p, \kappa^{\text{univ}})_{\text{fs}} \cong R\Gamma_c(\mathfrak{X}_{G, F, S} \times \hat{T}, g_{\text{rig}}^*(\bigotimes_v \mathfrak{A}_{\text{rig}, v})).$$

equivariant for the action of the Hecke-operators away from p and for the action of the Atkin–Lehner ring at the places dividing p .

REMARK 9.6.17. Note that the degree shift in the definition of $\mathfrak{A}_{\text{rig}, v}$ is necessary: it is expected that the coherent sheaf in Conjecture 9.6.16 has support of dimension $\sum_{v \mid p} \text{rk}_{\mathbf{Q}_p} T_v$. Hence with this degree shift we should expect that the compactly supported cohomology sits in degree 0. This should also explain why we actually need some degree shift in the analytic case of Section 7.1.

9.6.18. *Description of Hecke eigenspaces of p -adic automorphic forms.* We assume that Π_v is a G_v -representation satisfying the assumptions in Remark 9.6.12, that is we assume that $\Pi_v|_{H_v}$ is isomorphic to $\mathcal{C}^{\text{an}}(H_v, L)^m$ for some $m > 0$ and some compact open subgroup $H_v \subset G_v$. Given a locally analytic character δ of T we can compute the eigenspace

$$\text{Hom}_T(\delta, J_{B_v}(\Pi_v)) = \text{Hom}_{T^+}(\delta, \Pi_v^{N_v, 0})$$

using the coherent sheaf $\mathcal{M}_{\Pi_v}^* := \mathcal{M}_{\Pi_v^{N_v, 0}}^*$ defined as in (9.6.11). Indeed we can consider δ as a point in \hat{T} and compute that

$$(9.6.18) \quad \text{Hom}_T(\delta, J_{B_v}(\Pi_v)) = (\mathcal{M}_{\Pi_v}^* \otimes_{\mathcal{O}_{\hat{T}}} k(\delta))'$$

is the dual space to the fiber of \mathcal{M}^* at the point δ . We remark that there is also a derived version of this formula given by

$$R\mathrm{Hom}_T(\delta, J_{B_v}(\Pi_v)) = (R\Gamma(\hat{T}, \mathcal{M}_{\Pi_v}^* \otimes_{\mathcal{O}_{\hat{T}}}^L k(\delta)))'.$$

REMARK 9.6.19.

(i) At least if δ is a locally algebraic character the eigenspace $\mathrm{Hom}_T(\delta, J_{B_v}(\Pi_v))$ can be computed using Breuil’s adjunction formula [Bre15, Remarque 4.4]

$$\mathrm{Hom}_T(\delta, J_{B_v}(\Pi_v)) = \mathrm{Hom}_{G_v}(\mathcal{F}_{B_v}^{G_v}(\delta), \Pi_v),$$

where $\mathcal{F}_{B_v}^{G_v}(-)$ is the functor defined by Orlik and Strauch, see Section 6.2.13 above. Roughly this adjunction says that the finite slope part

$$M^\dagger(G, K^p, \kappa^{\mathrm{univ}})_{\mathrm{fs}} = J_B(\hat{S}(K^p, L)^{\mathrm{an}})$$

consists of those forms $f \in \hat{S}(K^p, L)^{\mathrm{an}}$ that generate a locally analytic subrepresentation all of whose subquotients occur in parabolically induced representations.

(ii) If δ_v is a locally analytic character of T_v it induces by restriction a character κ_v of T_v^+ and a character $\chi : \mathcal{A}^+(G_v) = L[\Sigma_v^+] \rightarrow L$. Conversely any such pair (κ_v, χ) canonically defines a character δ_v of T_v . With these notations and using (9.6.18) we can rewrite

$$\begin{aligned} \mathrm{Hom}_T(\delta, J_B(\hat{S}(K^p, L)^{\mathrm{an}})) &= \mathrm{Hom}_{T_v^+}(\delta, (\hat{S}(K^p, L)^{\mathrm{an}})^{N_0}) \\ &= M^\dagger(G, K^p, (\kappa_v)_v)^{\mathcal{A}^+(G)=\chi}. \end{aligned}$$

9.6.20. *Eigenvarieties.* We indicate (slightly informally) the relation with the construction of eigenvarieties as e.g. in [Eme06b]. The Hecke action in Conjecture 9.6.8 is in particular the action of a big Hecke algebra \mathbf{T} that is a tensor product of spherical Hecke algebras at good places away from p . Let us write $\hat{\mathbf{T}}(K^p)$ for the inverse limit of the images of \mathbf{T} in the endomorphism rings of

$$\hat{S}(K^p, \mathcal{O}/\varpi^m) = \{f : G(F)\backslash G(\mathbf{A}_F^\infty)/K^p \rightarrow \mathcal{O}/\varpi^m \text{ locally constant}\}.$$

Then $\hat{\mathbf{T}}(K^p)$ is a noetherian semi-local ring and we may form its attached formal scheme $\mathcal{Y}(K^p) = \mathrm{Spf} \hat{\mathbf{T}}(K^p)$ as well as its rigid analytic generic fiber $\mathfrak{Y}(K^p) = \mathcal{Y}(K^p)_\eta^{\mathrm{rig}}$. Following [Eme06b] (or our Remark 9.6.12 above and the discussion preceding it) the dual of the Jacquet module $J_B(\hat{S}(K^p, L)^{\mathrm{an}})'$ becomes in a natural way a coherent sheaf $\mathcal{M}^*(G, K^p)$ on $\mathfrak{Y}(K^p) \times \hat{T}$. The support of this coherent sheaf is called the eigenvariety $\mathcal{E}(G, K^p)$ associated with K^p (and the choice of the principal Bernstein blocks that we omit from the notation, compare Remark 9.6.1).

REMARK 9.6.21. Note that of course one could as well use $M^\dagger(G, K^p, \kappa^{\mathrm{univ}})_{\mathrm{fs}}$ to construct the eigenvariety. This is the point of view of [Che04] or [Loe11].

Conjecturally, as indicated for example in Section 1.3, the completed big Hecke algebra $\hat{\mathbf{T}}(K^p)$ should coincide with a product of pseudodeformation rings, and these pseudodeformation rings should either be defined as or proven to be isomorphic to the invariant global sections of the classical stack $\mathcal{X}_{G,F,S}^{\mathrm{cl}}$ underlying $\mathcal{X}_{G,F,S}$; that is, conjecturally one has an isomorphism

$$\hat{\mathbf{T}}(K^p) \cong \Gamma(\mathcal{X}_{G,F,S}^{\mathrm{cl}}, \mathcal{O}_{\mathcal{X}_{G,F,S}^{\mathrm{cl}}}).$$

Let us write $\mathcal{X}_{G,F,S}^{\text{ps}}$ for the disjoint union of the formal spectra of pseudodeformation rings corresponding to (G, F, S) . Then, using a map $\mathcal{Y}(K^p) \rightarrow \mathcal{X}_{G,F,S}^{\text{ps}}$ (constructed using the construction of Galois representations associated to automorphic forms) we may regard $\mathcal{M}^*(G, K^p)$ as a coherent sheaf on the Stein space $\mathfrak{X}_{G,F,S}^{\text{ps}} \times \hat{T}$, with $\mathfrak{X}_{G,F,S}^{\text{ps}} = (\mathcal{X}_{G,F,S}^{\text{ps}})_{\eta}^{\text{rig}}$. Conjecture 9.6.16 now roughly says that the coherent sheaf $\mathcal{M}^*(G, K^p)$ in some sense localizes over $\mathfrak{X}_{G,F,S} \times \hat{T}$, namely that there is a coherent sheaf on this stack (the Serre dual of the explicitly described coherent sheaf in the conjecture) whose (derived) global sections are concentrated in degree zero and agree with the global sections of the sheaf $\mathcal{M}^*(G, K^p)$ on the Stein space $\mathfrak{X}_{G,F,S}^{\text{ps}} \times \hat{T}$.

REMARK 9.6.22.

(i) We remark that we can in fact strengthen the computation of eigenspaces (9.6.18) in order to not just compute the eigenspace for the Hecke action at p (i.e. the action of the Atkin-Lehner ring $\mathcal{A}^+(G)$) but also for the Hecke action away from p , i.e. for the action of the Hecke operators in $\hat{\mathbf{T}}(K^p)$. Let $\mathfrak{m}_{\rho} \subset \mathbf{T}(K^p)$ be the ideal attached to a Hecke eigensystem with associated Galois representation $\rho : \text{Gal}_{F,S} \rightarrow \text{GL}_n(L)$. Then \mathfrak{m}_{ρ} defines a point $x_{\rho} \in \mathcal{Y}(K^p)_{\eta}^{\text{rig}}$ and we find

$$(9.6.22) \quad \begin{aligned} \text{Hom}_G(\widehat{\otimes}_v \mathcal{F}_B^G(\delta_v), \hat{S}(K^p, L)^{\text{an}}[\mathfrak{m}_{\rho}]) &= \text{Hom}_T\left(\prod_v \delta_v, J_B(\hat{S}(K^p, L)^{\text{an}})[\mathfrak{m}_{\rho}]\right) \\ &= (\mathcal{M}^*(G, K^p) \otimes k(x))', \end{aligned}$$

where $x = (x_{\rho}, (\delta_v)_v) \in \mathcal{Y}(K^p)_{\eta}^{\text{rig}} \times \hat{T}$.

(ii) In good situations the spaces $\mathcal{Y}(K^p)_{\eta}^{\text{rig}}$ and $\mathfrak{X}_{G,F,S}$ agree (up to the trivial action of a copy of \mathbf{G}_m) after passing to connected components where for example the residual Galois representation is a fixed irreducible and odd representation. In these cases (up to using Serre duality) Conjecture 9.6.16 gives a very precise description of the coherent sheaf $\mathcal{M}^*(G, K^p)$ that only involves stacks of Galois representations (or (φ, Γ) -modules). Hence we also obtain a (conjectural) purely Galois-theoretic description of the eigenspaces (9.6.22).

9.6.23. *More general Shimura varieties.* We return to the more general situation, i.e. we drop the assumption that G is compact at infinity. In particular the locally symmetric spaces (Shimura varieties) $Y(K_f)$ can have cohomology in several degrees, and it is better to work with complexes rather than individual cohomology groups.

Again we fix split maximal tori T_v and Borel subgroups B_v at the places dividing p . Let $\kappa = \prod_{v|p} \kappa_v : T \rightarrow L$ denote a locally analytic character. We fix K^p and I_v as above and let $K_f = K^p \times \prod_{v|p} I_v$. Then the representation

$$W(\kappa) = \widehat{\otimes}_{v|p} \text{Ind}_{B_{v,0}}^{I_v} \kappa_v$$

defines a “locally analytic” coefficient system again denoted by $W(\kappa)$ over $Y(K_f)$, see [Han17, 2.2, 3] (the coefficient system $W(\kappa)$ is denoted \mathcal{A}_{κ} in loc. cit.). The cohomology complex $R\Gamma(Y(K_f), W(\kappa))$ is then equipped with an action of the Atkin–Lehner ring $\mathcal{A}^+(G)$. Of course there is also a variant $W(\kappa^{\text{univ}})$ of this construction involving the universal character κ^{univ} .

REMARK 9.6.24. The complex $R\Gamma(Y(K_f), W(\kappa))$ should be seen as an object in the derived category of solid L -vector spaces, in particular it is only well defined

up to quasi-isomorphism. Its cohomology groups are L -vector spaces of compact type. The action of $\mathcal{A}^+(G)$ should also be regarded as an action in the derived category, i.e. the endomorphism defined by a given element is only well defined up to homotopy, etc.

Similarly we can consider the completed cohomology, compare [Fu21, 5] for the discussion of this in the derived sense. Again we consider complexes instead of individual cohomology groups, that is we consider the complex

$$R\tilde{\Gamma}(Y(K^p), L) = \left(\lim_{\leftarrow, m} \lim_{\rightarrow, K_p} R\Gamma(Y(K^p K_p), \mathcal{O}/\varpi^m) \right) \otimes_{\mathcal{O}} L.$$

Again we regard this as an object in the derived category of solid L -vector spaces. In this derived category it is equipped with an action of G such that the cohomology groups are admissible G -representations. Even better, the dual of $R\tilde{\Gamma}(Y(K^p), L)$, i.e. the complex computing completed homology, is a perfect complex, when we only consider it as a complex of $\mathcal{O}[[K]]$ -modules. Similar to (9.6.4) we may pass to the locally analytic vectors $R\tilde{\Gamma}(Y(K^p), L)^{\text{an}}$. We then have a comparison similar to Proposition 9.6.5.

PROPOSITION 9.6.25. *There is a quasi-isomorphism*

$$R\Gamma(Y(K_f), W(\kappa^{\text{univ}})) \cong (R\tilde{\Gamma}(Y(K^p), L)^{\text{an}})^{N_0}$$

equivariant for the action of the Atkin–Lehner ring.

PROOF. This is basically [Fu21, Theorem 1.2]. □

REMARK 9.6.26. There is of course a comparison of the individual complexes $R\Gamma(Y(K_f), W(\kappa))$ with a fixed κ instead of the universal one. For that comparison one has to take a derived eigenspace on the right hand side. We omit the technical details here.

Fu [Fu21, 6] continues the parallels with Section 9.6.2 and gives a derived construction of eigenvarieties building upon the complex $(R\tilde{\Gamma}(Y(K^p), L)^{\text{an}})^{N_0}$, respectively its dual. We will not discuss this here, but we point out that (at least on a heuristic level) the dual of the finite slope part $(R\tilde{\Gamma}(Y(K^p), L)^{\text{an}})^{N_0}_{\text{fs}}$ conjecturally should localize to an object of the derived category of coherent sheaves on $\mathfrak{X}_{G,F,S} \times \hat{T}$. Instead of making the construction of eigenvarieties more precise we state the following analogue of Conjecture 9.6.8 (and hence of Conjecture 9.3.2) in this context. Similarly to Section 9.6.2 we have a stack $\mathfrak{X}_{G,F,S}$ and a global-to-local map f_{rig} .

CONJECTURE 9.6.27. *Let d denote the dimension of $Y(K_f)$. Then there is an isomorphism*

$$R\Gamma(Y(K_f), W(\kappa))[d] \cong R\Gamma_c(\mathfrak{X}_{G,F,S}^{\text{odd}}, f_{\text{rig}}^* \mathfrak{A}_{\text{rig}}((\kappa_v)_{v|p}) \otimes V_{\mu})$$

compatible with the (derived) Hecke action and the action of the Atkin–Lehner ring.

REMARK 9.6.28. In fact, as in Section 9.3.1, one should actually replace G by its Weil restriction to \mathbf{Q} and obtain an additional action of Gal_F (F now being the reflex field of the Shimura datum) on the cohomology complexes. As in Conjecture 9.3.2 the isomorphism in Conjecture 9.6.27 should then also be equivariant for the Gal_F -action.

9.6.29. *Companion points and Breuil’s socle conjecture.* We return to the setup of Section 9.6.2 and assume in addition that G is a definite unitary group (of course similar results are expected in a more general context). In [Bre15, Conj. 5.3, Conj. 6.1] Breuil has made a precise conjecture about the socles (i.e. the largest semisimple subrepresentations) of the (admissible) locally analytic representations $\hat{S}(G, K^p)^{\text{an}}[\mathfrak{m}_\rho]$ for Galois representations ρ (with coefficients in L) that are crystalline at places dividing p and are such that the eigenvalues $\varphi_{1,v}, \dots, \varphi_{n,v} \in L$ of the crystalline Frobenius on the Weil–Deligne representation $\text{WD}(\rho|_{\text{Gal}_{F_v}})$ satisfy $\varphi_i/\varphi_j \notin \{1, q_v\}$ for $i \neq j$. Breuil’s conjecture describes which (tensor products indexed by the places dividing p of) irreducible subquotients of locally analytic principal series representations $(\text{Ind}_{B_v}^{G(F_v)} \delta_v)^{\text{an}}$ embed into the space $\hat{S}(G, K^p)^{\text{an}}[\mathfrak{m}_\rho]$. A weaker version of this conjecture is the precise description of the characters $(\delta_v)_v \in \prod_v \hat{T}_v$ such that

$$(\rho, (\delta_v)_v) \in \mathcal{E}(G, K^p) \subset \mathfrak{X}_{G,F,S}^{\text{ps}} \times \prod_v \hat{T}_v.$$

We give a precise version of the weaker conjecture and refer to [Bre15] for the precise formulation of Breuil’s socle conjecture. Using Remark 9.6.22 one can easily show that Breuil’s conjecture implies the weaker version presented here.

We continue to assume that $\rho_v = \rho|_{\text{Gal}_{F_v}}$ is crystalline with pairwise distinct Frobenius eigenvalues on $\text{WD}(\rho_v)$ that satisfy the above regularity assumption. Assume that ρ_v has regular Hodge–Tate weights λ_v and write ξ_v for the corresponding highest weight. To each Frobenius stable flag $\mathcal{F}_v \subset \text{WD}(\rho_v)$ (defined over $\bar{\mathbf{Q}}_p$) we may associate

- an unramified character

$$\delta_{\mathcal{F}_v} = \text{unr}_{\varphi_{1,v}} \otimes \dots \otimes \text{unr}_{\varphi_{n,v}} : T_v \rightarrow L^\times$$

where $\varphi_{1,v}, \dots, \varphi_{n,v}$ is the ordering of the Frobenius eigenvalues given by \mathcal{F}_v .

- for each embedding $\tau : F_v \hookrightarrow \bar{\mathbf{Q}}_p$ an element $w_\tau \in \mathcal{S}_n$ giving the relative position of the (τ -part of the) Hodge-Filtration on $D_{\text{dR}}(\rho_v) \otimes_{F,\tau} \bar{\mathbf{Q}}_p$ with respect to \mathcal{F}_v . We write $w(\mathcal{F}_v)$ for the tuple $(w_\tau)_\tau \in \prod_{\tau:F_v \hookrightarrow \bar{\mathbf{Q}}_p} \mathcal{S}_n$.

CONJECTURE 9.6.30. *Let ρ be a representation as above and let $(\delta_v)_{v|p} \in \hat{T}(L)$. Then $(\rho, (\delta_v)_v) \in \mathcal{E}(G, K^p)$ if and only if each δ_v is of the form*

$$\delta_v = z^{w_v w_0 \cdot \xi_v} \delta_{\text{sm}}$$

for some choice of a collection of Frobenius stable flags $(\mathcal{F}_v)_{v|p}$ and Weyl group elements $w(\mathcal{F}_v) \preceq w_v = (w_\tau)_{\tau:F_v \hookrightarrow L}$. Here w_0 denotes the longest Weyl group element and \preceq denotes the usual Bruhat order.

We point out that this conjecture would be a direct consequence of Conjecture 9.6.16: we only need to compute the support of the coherent sheaf occurring on the right hand side of Conjecture 9.6.16. However, this support is basically computed using Theorem 5.3.18.

REMARK 9.6.31. The points

$$x_{w_v} = (\rho, (\delta_{\mathcal{F}_v} z^{w_v w_0 \cdot \xi_v})_v)$$

for w_v as in the conjecture are often referred to as *companion points* of the point $x_{w_0} = (\rho, (\delta_{\mathcal{F}_v} z^{\xi_v})_v)$. Note that the latter point conjecturally corresponds to (the

Hecke eigensystem of) a classical automorphic form f . The p -adic automorphic forms in the Hecke eigensystem defined by x_{w_v} are then often referred to as *companion forms* of f .

Under additional assumptions, so-called Taylor–Wiles assumptions related to the patching construction in 3, this conjecture is proven in [BHS19, Theorem 5.3.3, Remark 5.3.2].

THEOREM 9.6.32. *Assume that*

- $p > 2$
- the CM extension E of F is unramified and does not contain a p -th root of unity.
- the group G is quasi-split at finite places and K_v is hyperspecial if v is a place of F inert in E .
- the residual Galois representation $\bar{\rho}$ of ρ is adequate.

Then Conjecture 9.6.30 holds true.

REMARK 9.6.33.

(i) In fact Conjecture 6.2.4 (iii) and Conjecture 9.6.8 imply the full conjecture in the locally analytic socle and [BHS19, Theorem 5.3.3] proves the full socle conjecture under the assumptions in the above theorem. The main input into this is the (proof of) the locally analytic Breuil–Mézard conjecture, Theorem 6.2.35.

(ii) There are generalizations of the above Theorem in the case of non-regular Hodge–Tate weights by Zhixiang Wu [Wu21a] and to the case of eigenvarieties associated to more general Bernstein components by Breuil–Ding [BD21].

9.6.34. *Classical and non-classical p -adic automorphic forms.* As in Section 9.6.29 we continue to work in the fixed global setup of a definite unitary group G . Let $x = (\rho, (\delta_v)_v) \in \mathfrak{X}_{G,F,S}^{\text{ps}} \times \prod_v \hat{T}_v$ be a point. We assume that the characters $\delta_v = z^{\xi_v} \delta_{v,\text{sm}}$ are locally algebraic with unramified smooth part $\delta_{v,\text{sm}}$. Recall from Remark 9.6.19 and Remark 9.6.22 that we may view δ as a tuple $((\kappa_v)_v, \chi)$ consisting of a character $\kappa = (\kappa_v)_v : T^0 = \prod_{v|p} T_v^0 \rightarrow L^\times$ and a character

$$\chi = \bigotimes_v \chi_{\delta_{v,\text{sm}}} : \mathcal{A}^+(G) \rightarrow L$$

of the Atkin–Lehner ring. Moreover, we have the identification of eigenspaces

$$(9.6.35) \quad M^\dagger(G, K^p, (\xi_v)_{v|p})[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi} \cong \text{Hom}_G(\widehat{\otimes}_v \mathcal{F}_B^G(\delta_v), \hat{S}(K^p, L)^{\text{an}}[\mathfrak{m}_\rho]).$$

Assume that ξ_v is the algebraic character associated with a regular Hodge–Tate weight $\underline{\lambda}_v$. Then the locally analytic representation $\text{Ind}_{B_{v,0}}^{I_v} z^{\xi_v}$ contains the algebraic representation $V_{\underline{\lambda}_v}$ of highest weight ξ_v . In particular we find that the space $M^\dagger(G, K^p, (z^{\xi_v})_{v|p})$ of overconvergent p -adic automorphic forms of weight $\xi = (z^{\xi_v})_{v|p}$ contains the space

$$S(K^p I, (\kappa_v)_v) = \left\{ \begin{array}{l} f : G(F) \backslash G(\mathbf{A}_F^\infty) / K^p \rightarrow \bigotimes_{v|p} V_{\underline{\lambda}_v} \\ \text{s. th. } f(gu) = u^{-1} f(g) \text{ for all } g \in G(\mathbf{A}_F^\infty), u \in \prod_{v|p} I_v \end{array} \right\}$$

of classical automorphic forms of weight ξ .

LEMMA 9.6.36. *Under the identification (9.6.35) an overconvergent p -adic automorphic form*

$$f \in M^\dagger(G, K^p, (z^{\xi_v})_{v|p})[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi}$$

is classical (i.e. lies in $S(K^p I, (z^{\xi_v})_v)$) if and only if the corresponding morphism

$$\widehat{\otimes}_v \mathcal{F}_B^G(\delta_v) \longrightarrow \widehat{S}(K^p, L)^{\text{an}}[\mathfrak{m}_\rho]$$

*factors through the locally algebraic quotient*²⁴

$$\widehat{\otimes}_v \mathcal{F}_B^G(\delta_v) \longrightarrow \bigotimes_{v|p} V_{\underline{\lambda}_v} \otimes (\text{Ind}_{B_v}^{G(F_v)} \delta_{\text{sm}})^{\text{sm}}.$$

PROOF. The proof is similar to the proof of [BHS17a, Proposition 3.4]. \square

The following conjecture is basically implied by standard Fontaine–Mazur type conjectures.

CONJECTURE 9.6.37. *Let $\rho \in \mathfrak{X}_{G,F,S}$ be a crystalline representation with regular Hodge–Tate weights $\underline{\lambda}$ and let χ be a character of the Atkin–Lehner ring such that the eigenspace*

$$M^\dagger(G, K^p, (z^{\xi_v})_{v|p})[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi}$$

is non-zero, i.e. ρ is associated to a p -adic automorphic form f . Then the eigenspace

$$S(K^p I, (z^{\xi_v})_v)[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi}$$

is non-zero, i.e. ρ is associated to a classical automorphic form.

In view of our conjectures on functors from locally analytic representations to coherent sheaves on stacks of Galois representations, one can approach this conjecture by studying the morphism of coherent sheaves obtained by applying the conjectural functor $\mathfrak{A}_G^{\text{fig}}$ to the map

$$\widehat{\otimes}_v \mathcal{F}_B^G(\delta_v) \longrightarrow \bigotimes_{v|p} (V_{\underline{\lambda}_v} \otimes \text{Ind}_{B_v}^{G(F_v)} \delta_{\text{sm}})^{\text{sm}}.$$

of locally analytic representations. In fact this strategy led to a proof of Conjecture 9.6.37 under the additional Taylor–Wiles assumptions as in Theorem 9.6.32, [BHS19, Theorem 5.1.3].

It is an interesting question whether the p -adic automorphic form f in Conjecture 9.6.37 is automatically classical. Note that this is *not* covered by the phenomenon of companion points or companion forms, as companion forms are p -adic forms of non-classical weight (that hence can't be classical!), but here we have fixed the weight to be classical. Still our conjectures imply that one cannot expect f to be classical. Conjecture 9.6.16 and Remark 9.6.22 give a precise conjectural description of the eigenspaces

$$S(K^p I, (z^{\xi_v})_v)[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi} \text{ and } M^\dagger(G, K^p, (z^{\xi_v})_{v|p})[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi}$$

in terms of a coherent sheaf that can be computed using the Drinfeld compactification \mathfrak{X}_B . It turns out that this compactification is not Gorenstein in general and these singularities have the consequence that the fiber dimension of the sheaf $\mathcal{M}^*(G, K^p)$ can jump. This basically implies that

$$M^\dagger(G, K^p, (z^{\xi_v})_{v|p})[\mathfrak{m}_\rho]^{\mathcal{A}^+(G)=\chi}.$$

²⁴See Remark 6.2.20 for the description of that quotient map.

should sometimes have larger dimension than its classical subspace. This expectation can be made into a concrete example in the case of a definite unitary group in three variables (and will appear in [HHS]).

Appendix A. Background on category theory

We recall various results in (∞) -category theory that will be applied in the main body of the notes.

A.1. Stable ∞ -categories. The functors we consider in the categorical local Langlands program will necessarily be derived, in so far as they won't necessarily take objects from the abelian category of smooth representations of the p -adic group G under consideration to an abelian category of coherent or quasicoherent sheaves on the corresponding moduli stack of Langlands parameters; rather, certain representations will be taken to complexes of sheaves. (This is already forced by the desired compatibility with Kisin–Taylor–Wiles patching, and so is an inevitable part of the story.) For this reason, it is necessary to work with appropriate categories of complexes of sheaves.

If we were writing these notes in the twentieth century, we would then proceed to explain that this obliges us to work with various derived or triangulated categories, and we would go on to explain which ones. But we will take advantage of contemporary advances in homological algebra by working with stable ∞ -categories, which avoid the well-known technical drawbacks of triangulated categories, and are especially useful to work with when one wants to apply various gluing or limiting processes (both of which we *will* be applying, e.g. since we are working with sheaves on formal algebraic stacks, which have both a topological structure — necessitating gluing arguments — and an Ind-stack structure, necessitating limiting arguments).

For the basics of stable ∞ -categories, and of t -structures on them, we refer to [Lur17, Ch. 1]. We note, though, that we will use *cohomological* indexing, while this reference uses homological indexing. We very briefly recall some of these basics, mainly for the purpose of orienting those readers (and those authors!) who are more familiar with the traditional theory of derived and triangulated categories than with the theory of stable ∞ -categories.

If \mathcal{C} is a stable ∞ -category, then its underlying homotopy category has a canonical triangulated category structure. In practice one can often just imagine that one is working in this underlying triangulated category, but with the improvement that cones (and thus homotopy limits and colimits) are canonical.

One important example of a stable ∞ -category is the derived ∞ -category $D(\mathcal{A})$ associated to a Grothendieck abelian category \mathcal{A} , whose corresponding homotopy category is the usual (unbounded) derived category of \mathcal{A} . (Recall that \mathcal{A} is a Grothendieck abelian category if it is cocomplete, the formation of filtered colimits in \mathcal{A} is exact, and \mathcal{A} admits a set of generators.) It is defined in [Lur17, Def. 1.3.5.8] as the *differential graded nerve* of the full subcategory of the category of chain complexes in \mathcal{A} consisting of objects which are *fibrant* for a certain model structure on the category of chain complexes. (See [Lur17, Prop. 1.3.5.3] for the construction of this model structure, and [Lur17, Prop. 1.3.5.6] for a partial description of the fibrant objects.) We can also describe $D(\mathcal{A})$ as the ∞ -category localization of the category of chain complexes with respect to quasi-isomorphisms. (See [Lur17, Prop. 1.3.5.15] and its proof.)

If X and Y are objects of an ∞ -category \mathcal{C} , then we have a mapping anima²⁵ $\mathrm{Maps}_{\mathcal{C}}(X, Y)$. When \mathcal{C} is furthermore stable, this mapping anima will in fact be a connective spectrum. In our contexts our categories will be furthermore \mathcal{O} -linear (for some ring of coefficients \mathcal{O}), and so $\mathrm{Maps}_{\mathcal{C}}(X, Y)$ will be a connective \mathcal{O} -module spectrum, or equivalently an animated \mathcal{O} -module, which (by the Dold–Kan correspondence) can be thought of as an object of $D^{\leq 0}(\mathcal{O} - \mathrm{Mod})$, the stable ∞ -category overlying the derived category of non-positively graded cochain complexes of \mathcal{O} -modules.

Again using stability, one can furthermore define a (typically non-connective) spectrum $\mathrm{RHom}_{\mathcal{C}}(X, Y)$, from which the mapping space can be recovered via the formula

$$\mathrm{Maps}_{\mathcal{C}}(X, Y) = \tau^{\leq 0} \mathrm{RHom}_{\mathcal{C}}(X, Y).$$

Again, in the contexts we consider, our categories will be \mathcal{O} -linear, and so the various $\mathrm{RHom}_{\mathcal{C}}(X, Y)$ will be \mathcal{O} -module spectra, which we can regard as objects of $D(\mathcal{O} - \mathrm{Mod})$ (unbounded in both directions).

Placing a t -structure on \mathcal{C} amounts (by definition) to placing a t -structure on its underlying triangulated category. The full subcategories²⁶ $\mathcal{C}^{\leq n}$, $\mathcal{C}^{\geq n}$ (for $n \in \mathbf{Z}$) are then defined as pullbacks of the corresponding subcategories of the homotopy category. The key facts are that:

- (1) $\mathrm{Maps}_{\mathcal{C}}(X, Y) = 0$ (i.e. is contractible) if $X \in \mathcal{C}^{\leq n}$ and $Y \in \mathcal{C}^{\geq n+1}$;
- (2) The inclusion $\mathcal{C}^{\leq n} \hookrightarrow \mathcal{C}$ admits a right adjoint $\tau^{\leq n}$.
- (3) The inclusion $\mathcal{C}^{\geq n} \hookrightarrow \mathcal{C}$ admits a left adjoint $\tau^{\geq n}$.

DEFINITION A.1.1. If \mathcal{C} is a stable ∞ -category endowed with a t -structure, then we define the *heart* of \mathcal{C} to be the intersection $\mathcal{C}^{\heartsuit} := \mathcal{C}^{\geq 0} \cap \mathcal{C}^{\leq 0}$.

The heart \mathcal{C}^{\heartsuit} is an abelian category [Lur17, Rem. 1.2.1.12]. (More precisely, it is equivalent to the nerve of the heart of the associated t -structure on the homotopy category of \mathcal{C} , and this latter heart is literally an abelian category. Recall that forming nerves is the technical device for converting usual categories into ∞ -categories, in the framework of [Lur09a].)

If \mathcal{A} is a Grothendieck abelian category, then the stable ∞ -category $D(\mathcal{A})$ is endowed with a canonical t -structure, arising from the usual t -structure on the derived category of \mathcal{A} [Lur17, Prop. 1.3.5.21]. There is a fully faithful functor $\mathcal{A} \hookrightarrow D(\mathcal{A})$ which induces an equivalence between \mathcal{A} and the heart of the t -structure on $D(\mathcal{A})$. (By definition the heart of $D(\mathcal{A})$ coincides with the heart of its full subcategory $D^+(\mathcal{A})$ discussed below, and then this claim follows from the dual version of [Lur17, Prop. 1.3.2.19].) We sometimes refer to \mathcal{A} as the subcategory of $D(\mathcal{A})$ consisting of *static* objects (sometimes known as “discrete” or “classical” objects).

We may form the full subcategory $D^+(\mathcal{A}) = \bigcup_{n \in \mathbf{Z}} D^{\geq n}(\mathcal{A})$. This category may also be constructed as the differential graded nerve of the full subcategory of bounded below chain complexes of injective objects of \mathcal{A} . (This is because the bounded below fibrant complexes are precisely the bounded below complexes of injectives; see [Lur17, Prop. 1.3.5.6]. This construction of $D^+(\mathcal{A})$ is then dual

²⁵In [Lur09b] and [Lur17], the terminology *spaces* is used for the ∞ -category arising from simplicial sets under the usual Quillen model structure, and for its objects, but we follow [CS21] in using the terminology *anima* instead.

²⁶Meaning sub- ∞ -categories, of course!

to the construction of $D^-(\mathcal{A})$ for abelian categories with enough projectives given in [Lur17, §1.3.2]. Note also that this construction of $D^+(\mathcal{A})$ only requires that \mathcal{A} has enough injectives, and doesn't require \mathcal{A} to be Grothendieck.)

A.2. Associative algebra objects. The notion of E_1 -rings (sometimes known as E_1 -ring spectra or A_∞ -ring spectra) is defined in [Lur17, Defn. 7.1.0.1]. By [Lur17, Rem. 7.1.0.3], static E_1 -rings are associative rings in the usual sense. There is a theory of left and right modules over an E_1 -ring A , and we write LMod_A , RMod_A respectively for these stable ∞ -categories. In the case that A is static, these identify with the derived ∞ -categories of the usual abelian categories of modules, [Lur17, Rem. 7.1.1.16].

If x is an object of a stable ∞ -category \mathcal{C} , then by [Lur17, Rem. 7.1.2.2] the spectrum $\mathrm{RHom}_{\mathcal{C}}(x, x)$ acquires the structure of an E_1 -ring which we denote $\mathrm{End}_{\mathcal{C}}(x)$, so that in particular $\pi_n \mathrm{End}_{\mathcal{C}}(x) = \mathrm{Ext}_{\mathcal{C}}^{-n}(x, x)$ for all $n \in \mathbf{Z}$.

A.3. Cocontinuous functors.

DEFINITION A.3.1. We say that a functor between ∞ -categories is *cocontinuous* if it preserves small colimits. (Note that in some parts of the literature, this is the definition of a *continuous* functor.)

The notes [Cam] provide a useful overview of categorical notions related to compact objects. The following lemma is partly based on some of the results discussed in these notes.

LEMMA A.3.2. *Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a functor between cocomplete and compactly generated stable ∞ -categories.*

- (1) *F is cocontinuous if and only if it admits a right adjoint G .*
- (2) *If F is cocontinuous, then F preserves compact objects if and only if its adjoint G is cocontinuous.*
- (3) *If F is cocontinuous, fully faithful, and preserves compact objects, then F also reflects compact objects, i.e. $F(X)$ compact implies X itself is compact.*

PROOF. (1) This is the adjoint functor theorem [Lur09a, Cor. 5.5.2.9].

(2) Suppose firstly that F preserves compact objects. In order to show that G is cocontinuous, it is enough to show that it preserves finite colimits and filtered colimits. For the former, recall that a functor between stable ∞ -categories preserves finite colimits if and only if it preserves finite limits; and since G is a right adjoint, it in fact preserves all limits.

Thus for any filtered colimit $\mathrm{colim}_j Y_j \xrightarrow{\sim} Y$ in \mathcal{D} , we have to show that the induced morphism $\mathrm{colim}_j G(Y_j) \rightarrow G(Y)$ is an isomorphism. By Yoneda's lemma, it's enough to do this after applying $\mathrm{Mor}(X, -)$ for any object X of \mathcal{C} , and since \mathcal{C} is compactly generated, we can further assume that X is compact. We then find that

$$\begin{aligned} \mathrm{Mor}(X, \mathrm{colim}_j G(Y_j)) &= \mathrm{colim}_j \mathrm{Mor}(X, G(Y_j)) \xrightarrow{\sim} \mathrm{colim}_j \mathrm{Mor}(F(X), Y_j) \\ &\xrightarrow{\sim} \mathrm{Mor}(F(X), \mathrm{colim}_j Y_j) \xrightarrow{\sim} \mathrm{Mor}(F(X), Y) \xrightarrow{\sim} \mathrm{Mor}(X, G(Y)), \end{aligned}$$

as required.

Conversely, if G is cocontinuous, it in particular preserves filtered colimits. Accordingly, if X is compact, and $\mathrm{colim}_j Y_j$ is a filtered colimit in \mathcal{D} , then we may

write

$$\begin{aligned} \mathrm{Mor}(F(X), \mathrm{colim}_j Y_j) &\xrightarrow{\sim} \mathrm{Mor}(X, G(\mathrm{colim}_j Y_j)) \xrightarrow{\sim} \mathrm{Mor}(X, \mathrm{colim}_j G(Y_j)) \\ &= \mathrm{colim}_j \mathrm{Mor}(X, G(Y_j)) \xrightarrow{\sim} \mathrm{colim}_j \mathrm{Mor}(F(X), Y_j), \end{aligned}$$

so that $F(X)$ is compact, as required.

(3) We must show that $F(X)$ compact implies that X itself is compact. Here are two arguments:

(i) Write $X \xrightarrow{\sim} \varinjlim_i X_i$, with X_i compact. Applying the functor F , we obtain $F(X) \xrightarrow{\sim} \varinjlim_i F(X_i)$. If $F(X)$ is compact, then this latter isomorphism factors through one of the $F(X_i)$. Since F is fully faithful, we find that our original isomorphism similarly so factors, and so the identity automorphism of X factors as $X \rightarrow X_i \rightarrow X$. Thus X is a retract of the compact object X_i , and so is itself compact.

(ii) Since F is fully faithful, its right adjoint is a retract, i.e. the unit of adjunction is a natural isomorphism $\mathrm{id}_{\mathcal{C}} \xrightarrow{\sim} G \circ F$; in particular, any object of \mathcal{C} is in the essential image of G . Thus it suffices to show that the natural map

$$\bigoplus_j \mathrm{Mor}(X, G(Y_j)) \rightarrow \mathrm{Mor}(X, \bigoplus_j G(Y_j))$$

is an isomorphism, for any collection $\{Y_j\}$ of objects of \mathcal{D} . (Here we use the fact that in a stable ∞ -category, to test compactness of an object it suffices to test Homs against coproducts.) By (2), the adjoint G preserves filtered colimits, and so we find that

$$\begin{aligned} \bigoplus_j \mathrm{Mor}(X, G(Y_j)) &\xrightarrow{\sim} \bigoplus_j \mathrm{Mor}(F(X), Y_j) \xrightarrow{\sim} \mathrm{Mor}(F(X), \bigoplus_j Y_j) \\ &\xrightarrow{\sim} \mathrm{Mor}(X, G(\bigoplus_j Y_j)) \xrightarrow{\sim} \mathrm{Mor}(X, \bigoplus_j G(Y_j)), \end{aligned}$$

as required. \square

A.4. Pro-categories (and Ind-categories). We very briefly recall the definitions of Pro- and Ind- ∞ -categories; see [Lur09a, §5.3.5] for more details. (More precisely, this reference treats the case of Ind- ∞ -categories, but for any ∞ -category \mathcal{C} , there is an equivalence $\mathrm{Ind}(\mathcal{C}) \cong \mathrm{Pro}(\mathcal{C}^{\mathrm{op}})^{\mathrm{op}}$.)

If \mathcal{C} is a (small) ∞ -category, one can define its associated pro-category $\mathrm{Pro}(\mathcal{C})$ to be the category whose objects are the diagrams $F : I \rightarrow \mathcal{C}$ indexed by a cofiltered small ∞ -category I , with the morphisms between two diagrams $F : I \rightarrow \mathcal{C}, G : J \rightarrow \mathcal{C}$ being defined by the formula

$$\mathrm{Mor}_{\mathrm{Pro}(\mathcal{C})}(F, G) = \varprojlim_J \varinjlim_I \mathrm{Mor}_{\mathcal{C}}(F(i), G(j)).$$

(In fact, as explained in the introduction to [Lur09a, §5.3], this is not the most convenient way to set up the theory of Ind- and Pro- ∞ -categories, but we ignore this point.)

In the case that \mathcal{C} is an ordinary category, we will sometimes denote the diagram $F : I \rightarrow \mathcal{C}$ via “ $\varprojlim_I F(i)$ ”, and refer to it as a pro-object of \mathcal{C} . The point of this notation is to distinguish the pro-object from the limit $\varprojlim_I F(i)$ in \mathcal{C} itself, if this

limit happens to exist. Often, when employing this notation, we write X_i rather than $F(i)$, and so denote the object F of $\text{Pro}(\mathcal{C})$ as “ $\varprojlim_I X_i$ ”.

A.5. Complete t -structures. References for this material are [BZNP17, §6] and [Lur17, §§1.2, 1.3]. (Important caution: these references use homological indexing, rather than the cohomological indexing that we use!)

DEFINITION A.5.1. If \mathcal{C} is a stable ∞ -category equipped with a t -structure, then it admits a *left completion*

$$\widehat{\mathcal{C}} := \lim_{\mathbf{Z}} \mathcal{C}^{\geq n}.$$

The limit is indexed by the totally ordered set \mathbf{Z} , and the transition functors are given by the truncations $\tau^{\geq n} : \mathcal{C}^{\geq n-1} \rightarrow \mathcal{C}^{\geq n}$.

There is a canonical functor $\mathcal{C} \rightarrow \widehat{\mathcal{C}}$, defined by mapping an object X to the sequence $(\tau^{\geq n} X)$, which has a natural interpretation as an object of $\widehat{\mathcal{C}}$. We say that \mathcal{C} is *left complete* if this functor is an equivalence.

There are dual notions of right completion and right completeness. We don't introduce notation for the right completion, but to be explicit: \mathcal{C} is *right complete* if the canonical functor $\mathcal{C} \rightarrow \lim_{\mathbf{Z}} \mathcal{C}^{\leq n}$ defined by $X \mapsto (\tau^{\leq n} X)$ is an equivalence.

Related to these notions of completeness are corresponding notions of separatedness.

DEFINITION A.5.2. If \mathcal{C} is a stable ∞ -category equipped with a t -structure, then we say it is *left separated* if

$$\mathcal{C}^{-\infty} := \bigcap_{n \in \mathbf{Z}} \mathcal{C}^{\leq n}$$

consists only of zero objects. Dually, we say it is *right separated* if

$$\mathcal{C}^{\infty} := \bigcap_{n \in \mathbf{Z}} \mathcal{C}^{\geq n}$$

consists only of zero objects.

If \mathcal{A} is a Grothendieck abelian category, the stable ∞ -category $D(\mathcal{A})$ is left and right separated, and also right complete [Lur17, Prop. 1.3.5.21]. It need not be left complete in general [Nee11].

If \mathcal{A} admits countable products, and if these are furthermore exact, then $D(\mathcal{A})$ admits countable products (take products level wise on complexes), and it follows from [Lur17, Prop. 1.2.1.19] that $D(\mathcal{A})$ is left complete (since Condition (2) there is satisfied). For example, if \mathcal{A} is the abelian category of modules over a ring, then $D(\mathcal{A})$ is left complete as well as right complete. We will be interested in contexts in which products in \mathcal{A} , although they exist, may not be exact, and so we note the following mild generalization of [Lur17, Prop. 1.2.1.19]. (The possibility of such a generalization is noted in [BZNP17, Rmk. 6.1.5], where the hypothesis of right t -exactness of products is required to hold only “up to a finite shift”.)

PROPOSITION A.5.3. *Let \mathcal{C} be a stable ∞ -category equipped with a t -structure. Suppose that \mathcal{C} admits countable products, and that the formation of countable products is of bounded amplitude (i.e. there is some $a \geq 0$ such that if $\{X_n\}$ is any sequence of objects of $\mathcal{C}^{\leq 0}$, then $\prod_n X_n$ lies in $\mathcal{C}^{\leq a}$.) Then \mathcal{C} is left complete if and only if \mathcal{C} is left separated.*

PROOF. One follows the proof of [Lur17, Prop. 1.2.1.19], noting now that if $\cdots \rightarrow X_n \rightarrow X_{n-1} \rightarrow \cdots$ is a tower of objects in $\mathcal{C}^{\leq 0}$, then $\lim X_n$ lies in $\mathcal{C}^{\leq a+1}$.

Now the proof works just as written, except that rather than considering the factorization $\lim f \rightarrow f(n-1) \rightarrow f(n)$, we consider the factorization $\lim f \rightarrow f(n-a-1) \rightarrow f(n)$. \square

REMARK A.5.4. We elaborate slightly on the construction of [Nee11] alluded to above: it gives examples of Grothendieck abelian categories \mathcal{A} for which both the hypothesis of Proposition A.5.3 (that countable products in $D(\mathcal{A})$ have bounded cohomological amplitude), and the conclusion (that $D(\mathcal{A})$ is left complete) fail. Example E.2.3 below gives an explicit example of this construction, which arises naturally in the context of smooth representation theory of (non-analytic) pro- p -groups.

One of the advantages of the framework of stable ∞ -categories, compared to that of triangulated categories, is the following simple and very general result (which is the dual version to [Lur17, Thm. 1.3.3.2]) on the existence of right derived functors. Recall that we say that a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ between stable ∞ -categories endowed with t -structures is *left t -exact* if $F(\mathcal{C}^{\geq 0}) \subseteq \mathcal{D}^{\geq 0}$, that it is *right t -exact* if $F(\mathcal{C}^{\leq 0}) \subseteq \mathcal{D}^{\leq 0}$, and *t -exact* if it is both left and right t -exact.

THEOREM A.5.5. *If \mathcal{A} is an abelian category with enough injectives, and if \mathcal{C} is a stable ∞ -category equipped with a right complete t -structure, then $F \mapsto \tau^{\leq 0} F|_{\mathcal{A}}$ (the restriction being taken by identifying \mathcal{A} with the heart of $D^+(\mathcal{A})$) is an equivalence between the ∞ -category of left t -exact functors $D^+(\mathcal{A}) \rightarrow \mathcal{C}$ which carry injective objects of \mathcal{A} into \mathcal{C}^{\heartsuit} , and the ordinary category of left exact functors from \mathcal{A} to \mathcal{C}^{\heartsuit} .*

As a particular consequence we have the following result [Lur17, Prop. 1.3.3.7] (which notoriously does not hold in the same level of generality in the triangulated category context, because of the non-functoriality of the formation of cones in that context).

PROPOSITION A.5.6. *Suppose that \mathcal{C} is a stable ∞ -category equipped with a right complete t -structure, and that its heart \mathcal{C}^{\heartsuit} has enough injectives. Then the inclusion $\mathcal{C}^{\heartsuit} \hookrightarrow \mathcal{C}$ extends essentially uniquely²⁷ to a t -exact functor $F : D^+(\mathcal{C}^{\heartsuit}) \rightarrow \mathcal{C}$.*

Furthermore, the following properties are equivalent:

- (1) *This functor is fully faithful.*
- (2) *For any objects X and Y of \mathcal{C}^{\heartsuit} the natural map $\mathrm{Ext}_{\mathcal{C}^{\heartsuit}}^i(X, Y) \rightarrow \mathrm{Ext}_{\mathcal{C}}^i(X, Y)$ is an isomorphism for $i > 0$.*
- (3) *For any objects X and Y of \mathcal{C}^{\heartsuit} with Y injective, $\mathrm{Ext}_{\mathcal{C}}^i(X, Y) = 0$ for $i > 0$.*
- (4) *For any objects X and Y of \mathcal{C}^{\heartsuit} with Y injective, there is an epimorphism $Z \rightarrow X$ such that $\mathrm{Ext}_{\mathcal{C}}^i(Z, Y) = 0$ for $i > 0$.*

Finally, if these equivalent properties hold, then the essential image of F is equal to $\mathcal{C}^+ := \cup_n \mathcal{C}^{\geq n}$.

²⁷Meaning that the sub- ∞ -category of the ∞ -category of functors $\mathrm{Fun}(D^+(\mathcal{C}^{\heartsuit}), \mathcal{C})$ whose objects consist of such extensions, whose 1-morphisms consist of natural transformations that restrict to the identity on the restriction of F to \mathcal{C}^{\heartsuit} , and whose higher morphisms consist of arbitrary higher morphisms between these 1-morphisms, is contractible. This is the ∞ -categorical analogue of being unique up to a unique natural isomorphism that restricts to the identity on the restriction of F to \mathcal{C}^{\heartsuit} .

PROOF. Except for the insertion of item (2) (which we add just for expositional clarity) this is simply the dual version of [Lur17, Prop. 1.3.3.7]. To see the equivalence of (2) with the other conditions, note that (1) certainly implies (2) (even for $i = 0$; however this case is in any event automatic, since \mathcal{C}^\heartsuit is full in \mathcal{C} by construction), while (2) clearly implies (3). \square

Below we will use this proposition both as stated, and in the dual form for projective objects and categories bounded above (which is the case literally treated in [Lur17, Prop. 1.3.3.7]).

A.6. Compact and coherent objects and regular t -structures. A reference for this material is [BZNP17, §6]. (Reminder: this reference uses homological indexing, rather than the cohomological indexing that we use.)

DEFINITION A.6.1. Let \mathcal{C} be an ∞ -category admitting filtered colimits. We say that an object $x \in \mathcal{C}$ is *compact* if $\mathrm{Maps}_{\mathcal{C}}(x, -)$ preserves filtered colimits (see [Lur09a, Defn. 5.3.4.5]).

REMARK A.6.2. If \mathcal{C} is a stable ∞ -category admitting filtered colimits, then it is cocomplete. Then $x \in \mathcal{C}$ is compact if and only if $\mathrm{RHom}_{\mathcal{C}}(x, -)$ preserves filtered colimits, if and only if $\mathrm{RHom}_{\mathcal{C}}(x, -)$ preserves all (small) colimits. (This is because $\mathrm{RHom}_{\mathcal{C}}(x, -)$ automatically preserves all limits and is therefore exact.)

For stable ∞ -categories equipped with a t -structure, we have the related notion of coherent objects. (See [BZNP17, §6.2].) The following is [BZNP17, Def. 6.2.2].

DEFINITION A.6.3. Let \mathcal{C} be a stable ∞ -category that admits filtered colimits, and is equipped with a t -structure that is compatible with filtered colimits (which is to say that $\mathcal{C}^{\geq 0}$, or equivalently every $\mathcal{C}^{\geq n}$, is closed under the formation of filtered colimits in \mathcal{C}). We say that an object $x \in \mathcal{C}$ is *coherent* if x is bounded below, i.e. x lies in $\mathcal{C}^{\geq n}$ for some n , and if x is furthermore compact in $\mathcal{C}^{\geq m}$ for all $m \leq n$.

DEFINITION A.6.4. In the context of Definition A.6.3, we write $\mathrm{Coh}(\mathcal{C})$ for the full sub- ∞ -category of coherent objects of \mathcal{C} . This is a stable sub- ∞ -category by [BZNP17, Rem. 6.2.3].

REMARK A.6.5. In the context of Definition A.6.3, if $m \geq n$, then maps from x to objects of $\mathcal{C}^{\geq m}$ factor through $\tau^{\geq m}x$, and so if x is compact in $\mathcal{C}^{\geq n}$ then we see that $\tau^{\geq m}x$ is compact as an object of $\mathcal{C}^{\geq m}$ for $m \geq n$.

REMARK A.6.6. If \mathcal{C} is as in Definition A.6.3 and is furthermore right complete, and if $x \in \mathcal{C}$ is coherent, then x is bounded above as well as below. (Apply the compactness assumption to the identity map $x \rightarrow x = \mathrm{colim} \tau^{\leq n}x$.)

The following result combines parts of [BZNP17, Lems. 6.2.4, 6.2.5].

LEMMA A.6.7. *Let \mathcal{C} be a stable ∞ -category endowed with a right complete t -structure that is also compatible with filtered colimits. Then the following are equivalent:*

- (1) *The inclusion $\mathcal{C}^{\geq 0} \hookrightarrow \mathcal{C}^{\geq -1}$ preserves compact objects.*
- (2) *The full subcategory of compact objects of \mathcal{C}^\heartsuit (the heart of the t -structure on \mathcal{C}) is an abelian subcategory of \mathcal{C}^\heartsuit .*

If these conditions hold, then the coherent objects in \mathcal{C} are precisely those objects x which are bounded both above and below and such that each $H^n(x)$ is a compact object of \mathcal{C}^\heartsuit .

DEFINITION A.6.8. A Grothendieck abelian category is *locally coherent* if it is compactly generated, and if its full subcategory of compact objects is abelian. (See for example [Her97, Thm. 1.6] for some equivalent characterisations of this notion.)

The following definition is [BZNP17, Def. 6.2.7].

DEFINITION A.6.9. We say that a t -structure on a stable ∞ -category admitting filtered colimits is *coherent* if it is right complete and is compatible with filtered colimits, and if furthermore \mathcal{C}^\heartsuit is locally coherent.

Since the compact objects of a locally coherent abelian category form an abelian subcategory (by definition), a coherent t -structure in particular satisfies the equivalent conditions of Lemma A.6.7.

We now define what it means for a coherent t -structure to be regular (following [BZNP17, §6.3]). We first note that if \mathcal{C} is equipped with a coherent t -structure, then since by assumption \mathcal{C} admits filtered colimits, there is a canonical functor

$$(A.6.10) \quad \mathrm{Ind\,Coh}(\mathcal{C}) \rightarrow \mathcal{C}.$$

This functor is cocontinuous (essentially by construction).

DEFINITION A.6.11. We say that a coherent t -structure on \mathcal{C} is *regular* if \mathcal{C} is compactly generated by its coherent objects (i.e. its coherent objects are compact, and generate), or if (equivalently) the functor (A.6.10) is an equivalence.

REMARK A.6.12. By [BZNP17, Prop. 6.3.2], if \mathcal{C} has a regular coherent t -structure, then the compact objects of \mathcal{C} are exactly the coherent objects.

If \mathcal{C} is equipped with a coherent t -structure, then the pullback of this t -structure along (A.6.10) endows $\mathrm{Ind\,Coh}(\mathcal{C})$ with a t -structure, which is again coherent and is in fact regular [BZNP17, Prop. 6.3.2]. The functor (A.6.10) is t -exact by construction. For any n , it restricts to an *equivalence*

$$\mathrm{Ind\,Coh}(\mathcal{C})^{\geq n} \xrightarrow{\sim} \mathcal{C}^{\geq n}.$$

If \mathcal{C} is equipped with a coherent t -structure, then the t -structure on its left completion $\widehat{\mathcal{C}}$ is again coherent [BZNP17, Prop. 6.3.4]. Again, the canonical functor

$$\mathcal{C} \rightarrow \widehat{\mathcal{C}}$$

induces equivalences

$$\mathcal{C}^{\geq n} \xrightarrow{\sim} \widehat{\mathcal{C}}^{\geq n}.$$

Thus the formation of $\mathrm{Ind\,Coh}(\mathcal{C})$ and of $\widehat{\mathcal{C}}$ are two different ways of “filling out” $\mathcal{C}^+ = \cup_n \mathcal{C}^{\geq n}$ in the leftward direction.

REMARK A.6.13. If the t -structure on \mathcal{C} is left separated (e.g. because it is left complete), then the kernel of the functor (A.6.10) consists precisely of $\mathrm{Ind\,Coh}(\mathcal{C})^{-\infty}$. In examples when (A.6.10) is not an equivalence, i.e. when the t -structure on \mathcal{C} is not regular, it is often the case that $\mathrm{Ind\,Coh}(\mathcal{C})^{-\infty}$ is non-zero, i.e. that the t -structure on $\mathrm{Ind\,Coh}(\mathcal{C})$ is not left separated.

For later reference, we note the following general fact about t -structures and Ind-categories [AGH19, Prop. 2.13].

PROPOSITION A.6.14. *If \mathcal{C} is a stable ∞ -category endowed with a t -structure, then $\mathrm{Ind}\mathcal{C}$ inherits a t -structure, characterized by the requirements that $\mathcal{C} \hookrightarrow \mathrm{Ind}\mathcal{C}$ be t -exact, and that the inclusion $\mathcal{C}^{\geq 0} \rightarrow \mathcal{C}$ induce an equivalence $\mathrm{Ind}(\mathcal{C}^{\geq 0}) \xrightarrow{\sim} \mathrm{Ind}(\mathcal{C})^{\geq 0}$. If the t -structure on \mathcal{C} is furthermore right bounded, then the t -structure on $\mathrm{Ind}\mathcal{C}$ is right complete.*

A.7. Derived functors and adjoints. We recall the following result.

PROPOSITION A.7.1. *Suppose that $F : \mathcal{A} \rightarrow \mathcal{B}$ is an exact functor between Grothendieck abelian categories which is compatible with colimits, and let $G : \mathcal{B} \rightarrow \mathcal{A}$ denote the right adjoint to F (which exists by the adjoint functor theorem). Then*

- (1) *F induces a t -exact cocontinuous functor $D(\mathcal{A}) \rightarrow D(\mathcal{B})$, which we again denote by F .*
- (2) *The right derived functor $RG : D(\mathcal{B}) \rightarrow D(\mathcal{A})$ exists, and is right adjoint to F .*

EXAMPLE A.7.2. A typical context in which one might apply Proposition A.7.1 is when \mathcal{B} is a Grothendieck category, and \mathcal{A} is a localizing subcategory, i.e. a Serre subcategory which is furthermore closed under the formation of arbitrary direct sums (and hence arbitrary colimits) in \mathcal{B} . Then \mathcal{A} is also a Grothendieck abelian category, and we may take $F : \mathcal{A} \rightarrow \mathcal{B}$ to be the inclusion.

Although F is fully faithful by construction, its extension to a functor $D(\mathcal{A}) \rightarrow D(\mathcal{B})$ need not be fully faithful. (Example E.2.3 below gives an illustration of this.) However, we have the following positive result in this direction,

PROPOSITION A.7.3. *As in the preceding discussion, let $F : \mathcal{A} \hookrightarrow \mathcal{B}$ be the inclusion of a localizing subcategory into a Grothendieck abelian category, and consider its t -exact extensions $D^+(\mathcal{A}) \rightarrow D^+(\mathcal{B})$ and $D(\mathcal{A}) \rightarrow D(\mathcal{B})$.*

- (1) *If F preserves injectives, then the extension $D^+(\mathcal{A}) \rightarrow D^+(\mathcal{B})$ is again fully faithful, and its essential image consists of the full subcategory $D_{\mathcal{A}}^+(\mathcal{B})$ consisting of objects whose cohomologies lie in \mathcal{A} .*
- (2) *Suppose that the extension $D^+(\mathcal{A}) \rightarrow D^+(\mathcal{B})$ is fully faithful. Suppose, furthermore, that the formation of products in \mathcal{B} is exact, and that the derived right adjoint RG of Proposition A.7.1 has finite cohomological dimension. Then the extension $D(\mathcal{A}) \rightarrow D(\mathcal{B})$ is again fully faithful, with essential image equal to $D_{\mathcal{A}}(\mathcal{B})$ (the subscript \mathcal{A} again denoting the full subcategory consisting of objects whose cohomologies lie in \mathcal{A}).*

PROOF. The t -structure on $D(\mathcal{B})$ induces a t -structure on $D_{\mathcal{A}}^+(\mathcal{B})$, whose heart is precisely \mathcal{A} . If we assume that F preserves injectives, then Proposition A.5.6 gives rise to a t -exact equivalence

$$D^+(\mathcal{A}) \xrightarrow{\sim} D_{\mathcal{A}}^+(\mathcal{B})$$

which extends the identity functor on hearts. The essential uniqueness statement of that proposition shows that this equivalence must essentially coincide with the canonical extension of F . This proves (1).

Suppose now that the extension of F to $D^+(\mathcal{A})$ is fully faithful; as above, we continue to denote this extension, as well as the extension to $D(\mathcal{A})$, by F . For any object X of $D(\mathcal{A})$, the unit of adjunction gives rise to a natural morphism $X \rightarrow RG(F(X))$. Thus if n is any integer, applying this to the exact triangle

$\tau^{<n}X \rightarrow X \rightarrow \tau^{\geq n}X$ (and recalling that F is t -exact), we obtain a morphism of exact triangles

$$\begin{array}{ccccc} \tau^{<n}X & \longrightarrow & X & \longrightarrow & \tau^{\geq n}X \\ \downarrow & & \downarrow & & \downarrow \\ RG(\tau^{<n}F(X)) & \longrightarrow & RG(F(X)) & \longrightarrow & RG(\tau^{\geq n}F(X)) \end{array}$$

Since the functor F is assumed to be fully faithful on $D^+(\mathcal{A})$, we see that the right hand vertical arrow is an isomorphism. Of course $\tau^{<n}X$ has vanishing cohomology in degrees $\geq n$, and if $a \geq 0$ denote the cohomological amplitude of RG , then $RG(\tau^{<n}F(X))$ has vanishing cohomology in degrees $\geq n+a$. Thus if $i \geq n+a$, we obtain a commutative square

$$\begin{array}{ccc} H^i(X) & \xrightarrow{\sim} & H^i(\tau^{\geq n}X) \\ \downarrow & & \downarrow \sim \\ H^i(RG(F(X))) & \xrightarrow{\sim} & H^i(RG(\tau^{\geq n}F(X))) \end{array}$$

in which (as indicated) all but the left hand vertical arrow are known to be isomorphisms. Thus this arrow is an isomorphism as well.

Since n was arbitrary, we find that the unit of adjunction $X \rightarrow RG(F(X))$ induces an isomorphism on cohomology in all degrees, and hence is an isomorphism in $D(\mathcal{A})$. Thus RG is a left quasi-inverse to F , and so the latter is fully faithful, as claimed.

Thus $D(\mathcal{A})$ is equivalent to a full subcategory of $D(\mathcal{B})$, which is evidently contained in $D_{\mathcal{A}}(\mathcal{B})$. At this point there may well be a direct argument to show that $D_{\mathcal{A}}(\mathcal{B})$ is actually equal to the essential image of F , but we finish the argument in a slightly roundabout way.

First, we note that since the formation of products in \mathcal{B} is exact, the category $D(\mathcal{B})$ admits products, and their formation is t -exact. Proposition A.5.3 shows that $D(\mathcal{B})$ is furthermore left complete. One then sees that $D_{\mathcal{A}}(\mathcal{B})$ is also left complete. Now, if $\{X_n\}_{n=0}^{\infty}$ is a sequence of objects of $D(\mathcal{A})^{\leq 0}$, then $\prod_{n=0}^{\infty} F(X_n)$ is an object of $D(\mathcal{B})^{\leq 0}$, and so $RG(\prod_{n=0}^{\infty} F(X_n))$ is an object of $D(\mathcal{A})^{\leq a}$ (where, as above, a denotes the cohomological amplitude of RG). Furthermore, one checks using the adjointness of F and RG and the full faithfulness of F that $RG(\prod_{n=0}^{\infty} F(X_n))$ is the product of the X_n in $D(\mathcal{A})$. Thus countable products exist in $D(\mathcal{A})$, and their formation has bounded cohomological amplitude. Applying Proposition A.5.3 again, we find that that $D(\mathcal{A})$ is left complete as well.

Thus we may consider the square

$$\begin{array}{ccc} \widehat{D(\mathcal{A})} & \xrightarrow{\widehat{F}} & \widehat{D_{\mathcal{A}}(\mathcal{B})} \\ \sim \uparrow & & \sim \uparrow \\ D(\mathcal{A}) & \xrightarrow{F} & D_{\text{sm}}(\mathcal{B}) \end{array}$$

in which the hats denote left completions, the vertical arrows, which are the canonical functors from the indicated categories to their left completions, are (as we have just noted) equivalences, the bottom horizontal arrow is the functor F , and the top

horizontal arrow is the functor \widehat{F} induced by F on left completions. We claim that this diagram commutes (up to natural isomorphism). Since \widehat{F} is an equivalence (as F restricts to an equivalence $D(\mathcal{A})^{\geq n} \rightarrow D_{\text{sm}}(\mathcal{B})$ for each n), we conclude that the bottom horizontal arrow is an equivalence, as required.

To see the claimed commutativity, note that if X is an object of $D(\mathcal{A})$, mapping to the element $\lim \tau^{\geq n} X$ in $\widehat{D(\mathcal{A})}$, then we have a natural isomorphism

$$\widehat{F}(\lim \tau^{\geq n} X) := \lim F(\tau^{\geq n} X) \xrightarrow{\sim} \lim \tau^{\geq n} F(X),$$

using the t -exactness of F , and $\lim \tau^{\geq n} F(X)$ is precisely the image of the object $F(X)$ of $D_{\mathcal{A}}(\mathcal{B})$ in $\widehat{D_{\mathcal{A}}(\mathcal{B})}$. \square

REMARK A.7.4. In the context of Proposition A.7.3 (2), the proof shows that both $D(\mathcal{A})$ and $D(\mathcal{B})$ are left complete, and that one may compute the extension of F from $D^+(\mathcal{A})$ to $D(\mathcal{A})$ via the formula $F(X) := \lim F(\tau_{\geq n} X)$.

A.8. Semiorthogonal decompositions. In this section we outline the basic concepts related to semiorthogonal decompositions of stable ∞ -categories. The stable ∞ -categories under consideration will always be cocomplete. Furthermore, when we speak of a cocomplete subcategory, we will always mean this in the strongest sense, i.e. that the subcategory in question is closed under the formation of colimits in the ambient category.

We begin by proving some generalities about generators in cocomplete stable ∞ -categories.

DEFINITION A.8.1. Let \mathcal{C} be a stable ∞ -category, and X a set of objects of \mathcal{C} . Then we define the left and right orthogonals

$$\begin{aligned} {}^\perp X &:= \{y \in \mathcal{C} \mid \text{RHom}(y, x) = 0 \text{ for all } x \in X\}, \\ X^\perp &:= \{y \in \mathcal{C} \mid \text{RHom}(x, y) = 0 \text{ for all } x \in X\}. \end{aligned}$$

LEMMA A.8.2. *If \mathcal{C} is a cocomplete stable ∞ -category, and X is a set of objects of \mathcal{C} , then ${}^\perp X$ is a cocomplete stable subcategory of \mathcal{C} .*

PROOF. The cocompleteness of ${}^\perp X$ is immediate from

$$\text{RHom}(\text{colim}_i y_i, x) = \lim_i \text{RHom}(y_i, x),$$

so we only need to check stability. By [Lur17, Lem. 1.1.3.3], it is enough to check that \mathcal{C} is stable under the formation of cofibers and translations. The claim for cofibers is immediate from the definitions, and that for translations from the equality

$$\text{RHom}(y[1], x) = \Omega \text{RHom}(y, x)$$

which shows that $y \in {}^\perp X$ if and only if $y[1] \in {}^\perp X$. \square

LEMMA A.8.3. *If \mathcal{C} is a cocomplete stable ∞ -category, and X is a set of compact objects of \mathcal{C} , then X^\perp is a cocomplete stable subcategory of \mathcal{C} .*

PROOF. Each $x \in X$ is compact, so by Remark A.6.2 we have

$$\text{RHom}(x, \text{colim}_i y_i) = \text{colim}_i \text{RHom}(x, y_i),$$

showing that X^\perp is cocomplete. It is immediate from the definitions that X^\perp is stable under the formation of fibers, and the result follows as in the proof of Lemma A.8.2 (using a variant of [Lur17, Lem. 1.1.3.3] for fibers rather than cofibers, which follows from *loc. cit.* by passage to the opposite category). \square

We will use Lemma A.8.2 to define the subcategory $\langle X \rangle$ of \mathcal{C} generated by X . Before doing so, we find it useful to consider the following more general situation. Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a fully faithful and cocontinuous functor between cocomplete stable ∞ -categories. By Lemma A.3.2 (1), the functor F admits a right adjoint $G : \mathcal{D} \rightarrow \mathcal{C}$. Let $\mathcal{A} \subseteq \mathcal{D}$ denote the kernel of G , and let $\mathcal{B} \subseteq \mathcal{D}$ denote the essential image of F . These are both thick full subcategories of \mathcal{D} (obviously in the case of \mathcal{A} , and as a consequence of Lemma A.8.4 (6) in the case of \mathcal{B}).

LEMMA A.8.4. *In the preceding situation we have the following statements.*

- (1) F induces an equivalence $\mathcal{C} \xrightarrow{\sim} \mathcal{B}$.
- (2) G induces an equivalence $\mathcal{D}/\mathcal{A} \xrightarrow{\sim} \mathcal{C}$.
- (3) The composite $\mathcal{B} \subseteq \mathcal{D} \rightarrow \mathcal{D}/\mathcal{A}$ is an equivalence.
- (4) \mathcal{B} is closed under the formation of colimits in \mathcal{D} .
- (5) $\mathcal{A} = \mathcal{B}^\perp$.
- (6) $\mathcal{B} = {}^\perp\mathcal{A}$.

PROOF. (1) follows from the very definition of \mathcal{B} as the essential image of F , together with the assumption that F is fully faithful. To prove (2), note first that by the definition of \mathcal{A} as the kernel of G , the functor G induces a functor $\mathcal{D}/\mathcal{A} \rightarrow \mathcal{C}$. Since the unit of adjunction $\text{id}_{\mathcal{C}} \xrightarrow{\sim} GF$ is an equivalence (by full faithfulness of F), and the counit is an equivalence modulo \mathcal{A} (i.e. gives an equivalence $GFG \xrightarrow{\sim} G$) we deduce that this induced functor is indeed an equivalence. (3) is a reformulation, using (1) and (2), of the fact that $GF \xrightarrow{\sim} \text{id}_{\mathcal{C}}$. (4) follows from the fact that F is

To prove (5), note that $Y \in \mathcal{A}$ iff $G(Y) = 0$ iff $\text{RHom}(X, G(Y)) = 0$ for all $X \in \mathcal{C}$ iff $\text{RHom}(F(X), Y) = 0$ for all $X \in \mathcal{C}$ (by the adjunction between F and G) iff $Y \in \mathcal{B}^\perp$. To prove (6), note first that (5) gives the inclusion $\mathcal{B} \subseteq {}^\perp\mathcal{A}$. Now let $Y \in \mathcal{D}$, and let Z denote the cofiber of the counit of adjunction $FG(Y) \rightarrow Y$. Since $GF \xrightarrow{\sim} \text{id}_{\mathcal{C}}$, the morphism $GFG(Y) \rightarrow G(Y)$ is an equivalence, so that $Z \in \mathcal{A}$. In particular, if $Y \in {}^\perp\mathcal{A}$, then as $FG(Y) \in \mathcal{B} \subseteq {}^\perp\mathcal{A}$, we must also have $Z \in {}^\perp\mathcal{A}$; and since $Z \in \mathcal{A}$, we see that $\text{id}_Z = 0$, so $Z = 0$, as required. \square

DEFINITION A.8.5. If \mathcal{C} is a cocomplete stable ∞ -category, and X is a set of objects of \mathcal{C} , then we define the *subcategory generated by X* to be $\langle X \rangle := {}^\perp(X^\perp)$. If $\langle X \rangle = \mathcal{C}$ then we say that X is a *set of generators* of \mathcal{C} .

This may look like a rather strange definition, but we note that by Lemma A.8.2, $\langle X \rangle$ is a cocomplete stable subcategory of \mathcal{C} . The following lemma shows that it is the smallest cocomplete stable subcategory containing X , justifying the definition.

LEMMA A.8.6. *Suppose that \mathcal{C} is a cocomplete stable ∞ -category, that X is a set of objects of \mathcal{C} , and that \mathcal{E} is a cocomplete stable subcategory of \mathcal{C} which contains every object in X . Then \mathcal{E} contains $\langle X \rangle$.*

PROOF. Applying Lemma A.8.4 with F given by the inclusion $\mathcal{E} \subseteq \mathcal{C}$, we see that $\mathcal{E} = {}^\perp(\mathcal{E}^\perp)$. Since $\langle X \rangle = {}^\perp(X^\perp)$ (by definition), we need to show that ${}^\perp(X^\perp)$ is contained in ${}^\perp(\mathcal{E}^\perp)$. By definition, it suffices to prove that X^\perp contains \mathcal{E}^\perp , which is immediate from our assumption that \mathcal{E} contains X . \square

The following corollary is immediate.

COROLLARY A.8.7. *Let \mathcal{D} be a cocomplete stable ∞ -category, and let X be a set of objects of \mathcal{D} . Then the following are equivalent.*

- (1) X is a set of generators of \mathcal{D} .
- (2) The only cocomplete stable subcategory of \mathcal{D} containing every element of X is \mathcal{D} itself.
- (3) If y is an object of \mathcal{D} satisfying $\mathrm{RHom}(x, y) = 0$ for every $x \in X$, then $y = 0$.

We now recall what we mean by a semiorthogonal decomposition (which may differ slightly from other definitions in the literature, in that we allow the decomposition to be infinite). Let \mathcal{C} be a cocomplete stable ∞ -category, and let $\mathcal{A}_1, \mathcal{A}_2, \dots$ be cocomplete stable subcategories (either finite or infinite in number).

DEFINITION A.8.8. We say that $\mathcal{A}_1, \mathcal{A}_2, \dots$ is a *semiorthogonal decomposition* of \mathcal{C} if

- (1) The subcategories $\mathcal{A}_1, \mathcal{A}_2, \dots$ generate \mathcal{C} , and
- (2) if $x \in \mathcal{A}_i, y \in \mathcal{A}_j$ with $i < j$, then $\mathrm{RHom}(x, y) = 0$.

REMARK A.8.9. By definition, condition (2) in Definition A.8.8 is equivalent to asking that if $i < j$ then $\mathcal{A}_i \subseteq {}^\perp \mathcal{A}_j$, and is also equivalent to asking that $\mathcal{A}_j \subseteq \mathcal{A}_i^\perp$.

We can define semiorthogonal decompositions in terms of generators, as in the following definition and lemma.

DEFINITION A.8.10. Let \mathcal{C} be a cocomplete stable ∞ -category. We say that a (possibly infinite) sequence x_1, x_2, \dots of objects of \mathcal{C} is *weakly exceptional* if:

- (1) The x_i are all compact objects of \mathcal{C} .
- (2) The x_i generate \mathcal{C} .
- (3) $\mathrm{RHom}(x_i, x_j) = 0$ for $i < j$.

LEMMA A.8.11. *Let \mathcal{C} be a cocomplete stable ∞ -category, let x_1, x_2, \dots be a weakly exceptional sequence of objects of \mathcal{C} , and let $\mathcal{A}_i := \langle x_i \rangle$ be the subcategory generated by x_i . Then $\mathcal{A}_1, \mathcal{A}_2, \dots$ is a semiorthogonal decomposition of \mathcal{C} .*

PROOF. By assumption the sequence x_1, x_2, \dots generates \mathcal{C} . Since x_i is compact, we see from Lemma A.8.3 that x_i^\perp is cocomplete and stable, and since it contains x_j by assumption, we see that it must contain $\langle x_j \rangle = \mathcal{A}_j$.

Now let $y \in \mathcal{A}_j$ be arbitrary. We have just shown that $\mathrm{RHom}(x_i, y) = 0$, so that $x_i \in {}^\perp y$, which is cocomplete and stable by Lemma A.8.2. Thus ${}^\perp y$ contains $\mathcal{A}_i = \langle x_i \rangle$, and since y is arbitrary, we have shown that $\mathrm{RHom}(x, y) = 0$ for any $x \in \mathcal{A}_i, y \in \mathcal{A}_j$, as required. \square

A.8.12. *Generators and adjunctions.* We now return to the setting of Lemma A.8.4. Let $F' : \mathcal{C} \xrightarrow{\sim} \mathcal{B}$ denote the equivalence of part (1), and let $G' : \mathcal{B} \xrightarrow{\sim} \mathcal{C}$ denote its quasi-inverse. Let $\iota : \mathcal{B} \subseteq \mathcal{D}$ denote the inclusion, and $\pi : \mathcal{D} \rightarrow \mathcal{D}/\mathcal{A}$ denote the projection. Then $F \xrightarrow{\sim} \iota F'$, while $G \xrightarrow{\sim} G'(\pi\iota)^{-1}\pi$ (noting that $\pi\iota$ is an equivalence, by (3)). Now, by (5) and (6), either of the subcategories \mathcal{A} and \mathcal{B} determines the other, and the functors ι and π are defined canonically in terms of \mathcal{A} and \mathcal{B} . Thus the original data of the functor F is determined by giving either of the subcategories \mathcal{A} or \mathcal{B} , together with the equivalence F' .

Suppose that $\{X_i\}_{i \in I}$ is a collection of generators of \mathcal{C} . Then we see that

$$\mathcal{B}^\perp = \cap_{i \in I} F(X_i)^\perp = \{Y \in \mathcal{D} \mid \mathrm{RHom}(F(X_i), Y) = 0 \text{ for all } i \in I\}.$$

By Lemma A.8.4 (5), we then see that $\mathcal{A} = \cap_{i \in I} F(X_i)^\perp$.

Thus, suppose that we know the values of the functor F on the generating collection $\{X_i\}_{i \in I}$, say $Y_i := F(X_i)$. Then we can characterize \mathcal{B} as the stable full sub- ∞ -category generated by the Y_i . Given \mathcal{B} , we can then determine \mathcal{A} , and although from this data we can't determine F itself, we can determine the composite FG . Namely, we have $FG \xrightarrow{\sim} \iota(\pi\iota)^{-1}\pi$.

Another way to describe this is as follows: any Y in \mathcal{D} fits into a cofiber sequence (equivalently, fiber sequence, or exact triangle) — the *semiorthogonal decomposition* of Y — which is unique up to equivalence, of the form

$$X \rightarrow Y \rightarrow Z$$

with $X \in \mathcal{B}$ and $Z \in \mathcal{A}$. Namely, $X = FG(Y)$, and Z is the cofiber of the counit of the adjunction $FG(Y) \rightarrow Y$. We then recover $FG(Y)$ as the term X in this distinguished triangle.

We can apply this formalism in applications where the functor F is conjectural. More precisely, we can suppose given a collection of objects $\{Y_i\}_{i \in I}$ of \mathcal{D} (which we imagine arise as the image of a generating set of objects $X_i \in \mathcal{C}$ under a functor F ; but we won't actually be given F). Then we can define \mathcal{B} to be the subcategory generated by the Y_i , just as in the preceding paragraph. By construction the inclusion $\mathcal{B} \subset \mathcal{D}$ is fully faithful and preserves colimits, and so by Lemma A.3.2 (1) admits a right adjoint. Lemma A.8.4 then applies: we may form $\mathcal{A} := \mathcal{B}^\perp$, every object Y of \mathcal{D} admits a functorial semiorthogonal decomposition given by a cofiber sequence

$$X \rightarrow Y \rightarrow Z$$

with $X \in \mathcal{B}$ and $Z \in \mathcal{A}$, and the functor $Y \mapsto X$ is the right adjoint to the inclusion of \mathcal{B} .

A.8.13. *Restriction to open substacks.* Suppose further that we have a cocomplete stable ∞ -category \mathcal{D}' with a functor $j^* : \mathcal{D} \rightarrow \mathcal{D}'$, together with a fully faithful right adjoint $j_* : \mathcal{D}' \rightarrow \mathcal{D}$. (In applications, \mathcal{D} will be a category of sheaves on a stack \mathcal{X} , \mathcal{D}' will be the category of sheaves on an open substack $j : \mathcal{U} \rightarrow \mathcal{X}$, and j^*, j_* will have their usual meanings.) Let $\mathcal{B}', \mathcal{A}'$ be sub-stable ∞ -categories of \mathcal{D}' which give a semiorthogonal decomposition as above: so $\mathcal{A}', \mathcal{B}'$ together generate \mathcal{D}' , and we have $\mathcal{B}' = {}^\perp \mathcal{A}'$.

Set $\mathcal{A} := j_* \mathcal{A}'$. Then (since the full faithfulness of j_* implies that $j^* j_*$ is naturally equivalent to the identity) we have $\mathcal{A}' = j^* \mathcal{A}$, and by adjointness we have

$$(A.8.14) \quad {}^\perp \mathcal{A} = (j^*)^{-1}({}^\perp \mathcal{A}') = (j^*)^{-1}(\mathcal{B}').$$

Appendix B. Coherent, Ind-Coherent, and Pro-Coherent sheaves on Ind-algebraic stacks

B.1. Modules over I -adically complete rings. We will establish some results relating various stable ∞ -categories of modules. These provide the commutative algebra background for the theory of (Ind-)coherent sheaves on formal algebraic stacks. In fact our results apply in certain non-commutative contexts as well, and in this case we will use them below to deduce certain facts about stable ∞ -categories of smooth representations of p -adic analytic groups.

The setup is as follows: we let A be a (not necessarily commutative) Noetherian ring, and I a two-sided ideal in A with respect to which A is I -adically complete. We furthermore assume that I satisfies the *Artin-Rees property*: namely, that any A -submodule M of $A^{\oplus n}$ (for any $n \geq 0$) is closed with respect to the I -adic topology

on $A^{\oplus n}$, and the topology on M induced by the I -adic topology on $A^{\oplus n}$ coincides with the I -adic topology on M . It follows from this hypothesis that any finitely generated A -module is complete and separated with respect to its I -adic topology, and that morphisms between finitely generated A -modules have I -adically closed closed image and are strict with respect to the I -adic topologies on their source and target.

If A is commutative, then the Artin–Rees property holds automatically. If A is non-commutative, then it need not hold (as far as we know), but it does hold in one particular case that we are interested in.

EXAMPLE B.1.1. If H is a compact p -adic analytic group, if P is a normal open pro- p subgroup of H (the identity has a neighbourhood basis consisting of such P), and if I is the kernel of the projection $\mathcal{O}[[H]] \rightarrow k[H/P]$, then the I -adic topology on $\mathcal{O}[[H]]$, and more generally on any finitely generated $\mathcal{O}[[H]]$ -module, coincides with the canonical topology (as discussed e.g. in Section D.2 below). Thus I satisfies the Artin–Rees property.

DEFINITION B.1.2. We write $D(A) := D(A - \text{Mod})$, and similarly for any n we write $D(A/I^n) := D(A/I^n - \text{Mod})$.

Since A is Noetherian, the t -structure on $D(A)$ is coherent. Indeed the compact objects in its heart, which is the abelian category $A - \text{Mod}$, are precisely the finitely presented A -modules, and they form a generating abelian subcategory of $A - \text{Mod}$. The full subcategory of $D(A)$ consisting of coherent objects is thus precisely the category $D_{\text{f.p.}}^b(A)$ consisting of (cohomologically) bounded complexes whose cohomologies are finitely presented (or, equivalently, finitely generated, since A is Noetherian). Similar remarks apply with A/I^n in place of A .

There are obvious exact inclusions

$$(B.1.3) \quad A/I^m - \text{Mod} \hookrightarrow A/I^n - \text{Mod} \hookrightarrow A - \text{Mod},$$

for $m \leq n$, which induce t -exact functors

$$D_{\text{f.p.}}^b(A/I^m) \rightarrow D_{\text{f.p.}}^b(A/I^n) \rightarrow D_{\text{f.p.}}^b(A).$$

Since these functors are compatible in an evident way as m and n vary, we obtain a functor

$$(B.1.4) \quad \text{colim}_n D_{\text{f.p.}}^b(A/I^n) \rightarrow D_{\text{f.p.}}^b(A)$$

(the colimit being taken in the ∞ -category of stable ∞ -categories). We will show below that (B.1.4) is fully faithful, and describe its essential image. To begin with, we consider an analogous construction on the level of abelian categories.

DEFINITION B.1.5. We denote by $(A - \text{Mod})_I$ the full abelian subcategory (indeed Serre subcategory) of $A - \text{Mod}$ consisting of modules each element of which is annihilated by some power of I . Equivalently, $(A - \text{Mod})_I$ is the closure under filtered colimits of the union $\bigcup_n A/I^n - \text{Mod}$, the union taking place in $A - \text{Mod}$.

REMARK B.1.6. Since $(A - \text{Mod})_I$ is a Serre subcategory of the Grothendieck abelian category $A - \text{Mod}$ which is furthermore closed under the formation of arbitrary direct sums, it is itself a Grothendieck abelian category. In particular $D((A - \text{Mod})_I)$ is defined, and is equipped with its canonical t -structure, having $(A - \text{Mod})_I$ as its heart.

The compact objects in the category $(A - \text{Mod})_I$ are precisely the finitely presented modules (i.e. those that are compact simply as A -modules), and these form a generating abelian subcategory of $(A - \text{Mod})_I$. Thus the t -structure on $D((A - \text{Mod})_I)$ is coherent. Its full subcategory of coherent objects is equal to $D_{\text{f.p.}}^b((A - \text{Mod})_I)$, the subcategory consisting of complexes whose cohomologies are finitely presented.

There is one more collection of categories we need to introduce.

DEFINITION B.1.7. We let $D_I(A)$ denote the full subcategory of $D(A)$ consisting of objects whose cohomologies lie in $(A - \text{Mod})_I$, and set $D_{\text{f.p.,}I}^b(A) = D_{\text{f.p.}}^b(A) \cap D_I(A)$.

The heart of the t -structure on $D_I(A)$ is precisely $(A - \text{Mod})_I$, which is locally coherent, and hence $D_{\text{f.p.,}I}^b(A)$ is precisely the full subcategory of coherent objects of $D_I(A)$.

The inclusions $A/I^n - \text{Mod} \hookrightarrow A - \text{Mod}$ of (B.1.3) evidently factor through $(A - \text{Mod})_I$, and so (B.1.4) factors as

$$(B.1.8) \quad \text{colim}_n D_{\text{f.p.}}^b(A/I^n) \rightarrow D_{\text{f.p.}}^b((A - \text{Mod})_I) \rightarrow D_{\text{f.p.,}I}^b(A) \hookrightarrow D_{\text{f.p.}}^b(A).$$

We will show that the first two functors in this factorization are equivalences.

PROPOSITION B.1.9. *The functor*

$$(B.1.10) \quad \text{colim}_n D_{\text{f.p.}}^b(A/I^n) \rightarrow D_{\text{f.p.}}^b((A - \text{Mod})_I)$$

(the first arrow of (B.1.8)) is an equivalence.

PROOF. We will prove this by constructing a quasi-inverse to (B.1.10). But we begin by noting that (B.1.10) is essentially surjective, since we will need to use this in our construction of the quasi-inverse. This essential surjectivity amounts to the fact that if X^\bullet is a bounded complex of objects in $(A - \text{Mod})_I$, each of whose cohomologies is finitely generated, then we can find a quasi-isomorphism $Y^\bullet \rightarrow X^\bullet$ of bounded complexes with each Y^\bullet being a finitely presented object of $(A - \text{Mod})_I$ (so that Y^\bullet is in fact a complex of A/I^n -modules for some n). (This is standard, using that since A is Noetherian, all submodules of finitely generated A -modules are finitely presented, but for the convenience of the reader we sketch a proof. Assume that X^m is finitely presented for $m > n$, and let $(X^n)'$ be a finitely generated A -submodule of X^n which surjects onto both $H^n(X^\bullet)$ and $\text{im}(d^n : X^n \rightarrow X^{n+1})$. Let Y^\bullet be the subcomplex of X^\bullet with $Y^r = X^r$ for $r \neq n-1, n$, $Y^n = (X^n)'$, $Y^{n-1} = (d^{n-1})^{-1}(Y^n)$. Then the inclusion $Y^\bullet \rightarrow X^\bullet$ is a quasi-isomorphism, and Y^m is finitely presented for all $m \geq n$, and we may proceed by downward induction on n .)

We now turn to constructing the required quasi-inverse. Proposition A.6.14 endows $\text{Ind colim}_n D_{\text{f.p.}}^b(A/I^n)$ with a right complete t -structure, whose heart is precisely $(A - \text{Mod})_I$, so that Proposition A.5.6 gives rise to a t -exact functor

$$(B.1.11) \quad D^+((A - \text{Mod})_I) \rightarrow \text{Ind colim}_n D_{\text{f.p.}}^b(A/I^n).$$

We claim that (B.1.11) restricts to a functor

$$(B.1.12) \quad D_{\text{f.p.}}^b((A - \text{Mod})_I) \rightarrow \text{colim}_n D_{\text{f.p.}}^b(A/I^n),$$

which is furthermore quasi-inverse to (B.1.10).

The composite of (B.1.11) with the functor $D^+(A/I^m) \rightarrow D^+((A - \text{Mod})_I)$ (for some fixed m) is a t -exact functor

$$(B.1.13) \quad D^+(A/I^m) \rightarrow \text{Ind colim}_n D_{f.p.}^b(A/I^n)$$

whose restriction to $A/I^m - \text{Mod}$ coincides with the canonical functor

$$A/I^m - \text{Mod} = \text{Ind } D_{f.p.}^b(A/I^m)^\heartsuit \hookrightarrow \text{Ind } D_{f.p.}^b(A/I^m) \rightarrow \text{Ind colim}_n D_{f.p.}^b(A/I^n).$$

The same is true of the canonical functor

$$(B.1.14) \quad D^+(A/I^m) \xrightarrow{\sim} (\text{Ind } D_{f.p.}^b(A/I^m))^+ \rightarrow \text{Ind colim}_n D_{f.p.}^b(A/I^n).$$

Thus Proposition A.5.6 shows that the two functors (B.1.13) and (B.1.14) essentially coincide.

Consequently, the composite of (B.1.11) with the canonical functor $D_{f.p.}^b(A/I^m) \rightarrow D^+((A - \text{Mod})_I)$, which we can factor as

$$D_{f.p.}^b(A/I^m) \rightarrow D^+(A/I^m) \xrightarrow{(B.1.13)} \text{Ind colim}_n D_{f.p.}^b(A/I^n),$$

essentially coincides with the composite

$$D_{f.p.}^b(A/I^m) \rightarrow D^+(A/I^m) \xrightarrow{(B.1.14)} \text{Ind colim}_n D_{f.p.}^b(A/I^n),$$

which in turn evidently essentially coincides with the composite

$$D_{f.p.}^b(A/I^m) \rightarrow \text{colim}_n D_{f.p.}^b(A/I^m) \hookrightarrow \text{Ind colim}_n D_{f.p.}^b(A/I^n).$$

Varying m , and relabelling it as (the dummy variable) n , we find that the composite of (B.1.11) with the functor

$$\text{colim}_n D_{f.p.}^b(A/I^n) \xrightarrow{(B.1.10)} D_{f.p.}^b((A - \text{Mod})_I) \hookrightarrow D^+((A - \text{Mod})_I)$$

coincides (up to natural equivalence) with the inclusion

$$\text{colim}_n D_{f.p.}^b(A/I^n) \hookrightarrow \text{Ind colim}_n D_{f.p.}^b(A/I^n).$$

Since (B.1.10) is essentially surjective, we also find that the restriction of (B.1.11) to $D_{f.p.}^b((A - \text{Mod})_I)$ does have image lying in $\text{colim}_n D_{f.p.}^b(A/I^n)$, and thus does induce a functor (B.1.12). Furthermore, what we have shown so far is that the composite of (B.1.12) with (B.1.10) essentially coincides with the identity on $\text{colim}_n D_{f.p.}^b(A/I^n)$.

It remains to consider the composite of (B.1.10) with (B.1.12). This composite is a t -exact functor from $D_{f.p.}^b((A - \text{Mod})_I)$ to itself, which restricts to the identity on $D_{f.p.}^b((A - \text{Mod})_I)^\heartsuit$. Theorem A.5.5 shows that it essentially coincides with the identity functor. (Technically, we should be applying Theorem A.5.5 to $D^+((A - \text{Mod})_I)$ rather than its full subcategory $D_{f.p.}^b((A - \text{Mod})_I)$. But we can do this completely formally, by forming the canonical extension of our functor to a t -exact endofunctor of $\text{Ind}(D_{f.p.}^b((A - \text{Mod})_I))$, restricting this to $(\text{Ind}(D_{f.p.}^b((A - \text{Mod})_I)))^+$, and then using the canonical equivalence $D^+((A - \text{Mod})_I) \xrightarrow{\sim} \text{Ind}(D_{f.p.}^b((A - \text{Mod})_I))^+$.) \square

Before showing that the second functor in (B.1.8) is an equivalence, we recall what is presumably a standard result comparing injective objects in $A - \text{Mod}$ and $(A - \text{Mod})_I$.

The inclusion $(A - \text{Mod})_I \hookrightarrow A$ has a right adjoint, namely $M \mapsto M[I^\infty]$. Since the right adjoint to an exact functor preserves injectives, we see that if M is an

injective A -module, then $M[I^\infty]$ is an injective object of $(A - \text{Mod})_I$. But in fact more is true, as the following lemma shows.

LEMMA B.1.15. *If M is an injective A -module, then $M[I^\infty]$ is again an injective A -module.*

PROOF. By Baer’s criterion, we must show that if J is an ideal of A , then any morphism $f : J \rightarrow M[I^\infty]$ extends to a morphism $A \rightarrow M[I^\infty]$. Since J is finitely generated (A being Noetherian), the morphism f has image lying in $M[I^m]$ for some m , and so factors through the quotient $J/I^m J$ of J . Now our Artin–Rees hypothesis implies that $I^n \cap J \subseteq I^m J$ for some $n \geq m$. Thus f factors through a morphism $g : J/I^n \cap J \rightarrow M[I^\infty]$. If we rewrite the domain of g as $I^n + J/I^n$, and regard this as a submodule of R/I^n , then the hypothesis that M is injective ensures that g extends to a morphism $R/I^n \rightarrow M$, which obviously has image lying in $M[I^n] \subseteq M[I^\infty]$. Composing this extension of g with the quotient map $R \rightarrow R/I^n$ gives the required extension of f . \square

COROLLARY B.1.16. *The inclusion $(A - \text{Mod})_I \hookrightarrow A - \text{Mod}$ preserves injectives.*

PROOF. Let M be an injective object of $(A - \text{Mod})_I$, and let $M \hookrightarrow N$ be an inclusion of M into an injective A -module. This inclusion evidently factors through an inclusion $M \hookrightarrow N[I^\infty]$. This is a morphism of injective objects of $(A - \text{Mod})_I$, and hence is a split inclusion. On the other hand, Lemma B.1.15 shows that $N[I^\infty]$ is again an injective A -module, and hence so is its direct summand M . \square

PROPOSITION B.1.17. *The canonical functor $D^+((A - \text{Mod})_I) \rightarrow D_I^+(A)$ is an equivalence, which restricts to an equivalence $D_{\text{f.p.}}^b((A - \text{Mod})_I) \xrightarrow{\sim} D_{\text{f.p.,}I}^b(A)$ (and this is the second arrow of (B.1.8)).*

PROOF. The first claimed equivalence follows by combining Corollary B.1.16 with Proposition A.7.3 (1). Restricting this equivalence to the compact objects in the truncations $\tau^{\geq n}$ for each n , we see that the equivalence then restricts to an equivalence $D_{\text{f.p.}}^b((A - \text{Mod})_I) \xrightarrow{\sim} D_{\text{f.p.,}I}^b(A)$. \square

B.2. Ind-coherent sheaves on Ind-algebraic stacks. In this section we introduce (or, really, recall from the literature) some basic notions related to derived categories of sheaves of \mathcal{O} -modules on algebraic and formal algebraic stacks. Our discussion is slightly unsatisfactory, for the reason that we wish to work with stable ∞ -categories, rather than merely the underlying triangulated categories, but our basic reference for the theory of stacks is [Sta], which does not use ∞ -categorical notions.

For example, in the ∞ -categorical world, equivariance is a perfectly good notion for objects in the (stable ∞ -categorical version of the) derived category of quasicoherent sheaves on a scheme, and so by writing an algebraic stack \mathcal{X} as a quotient $[U/R]$ for some scheme U and groupoid R , one can define the stable ∞ -category of quasicoherent complexes on \mathcal{X} as the R -equivariant objects of the corresponding category for U . (Cf. Remark B.2.4 below.) In this way one can even define the *derived category* of quasicoherent complexes on \mathcal{X} *prior* to having defined the abelian category of quasicoherent sheaves on \mathcal{X} ; the latter can then be defined as the heart of a t -structure on the former. In this approach, the more general concept of $\mathcal{O}_{\mathcal{X}}$ -module, or the derived category of complexes of such, does not even arise. (And

the question of, e.g., the relationship between $D(\mathrm{QCoh}(\mathcal{X}))$ and $D_{\mathrm{qc}}(\mathcal{O}_{\mathcal{X}})$ doesn't come up.) However, in the framework of **[Sta]**, such a definition is not feasible.

We don't really try to reconcile these conflicting approaches here. We simply state various possible definitions, with pointers to the literature, and do our best to indicate the relationships between them. Ultimately, the ∞ -categorical perspective will be the one best suited to our goals.

B.2.1. The case of algebraic stacks. If \mathcal{X} is an algebraic stack, then we will let $D(\mathcal{X})$ denote the stable ∞ -category of quasicoherent complexes on \mathcal{X} . Let us briefly indicate what this means. (Note that the citations we are making typically work in the context of derived/triangulated categories, rather than stable ∞ -categories, but their results are readily adapted to this latter setting.)

There are actually (at least!) two possible approaches to defining this category (for any algebraic stack \mathcal{X}): we can first consider the abelian category $\mathrm{QCoh}(\mathcal{X})$ of quasicoherent sheaves on \mathcal{X} (and even this category admits multiple definitions: via small lisse-étale sites, small *fppf* sites (called the *flat-fppf* site in **[Sta, Tag 0786]**) or large *fppf* sites, which however lead to equivalent notions **[Sta, Tag 07B1]**), which is a Grothendieck abelian category, and then form $D(\mathrm{QCoh}(\mathcal{X}))$.

Alternatively, we can consider the larger Grothendieck abelian category $\mathcal{O}_{\mathcal{X}} - \mathrm{Mod}$ of *all* sheaves of $\mathcal{O}_{\mathcal{X}}$ -modules (again on some appropriate site of \mathcal{X}), form its associated stable ∞ -category $D(\mathcal{O}_{\mathcal{X}})$, and then (attempt to) consider the full subcategory $D_{\mathrm{qc}}(\mathcal{O}_{\mathcal{X}})$ consisting of objects whose cohomologies are quasicoherent, i.e. lie in $\mathrm{QCoh}(\mathcal{X})$. If we work with small sites, this definition is valid and reasonable, because the inclusion $\mathrm{QCoh}(\mathcal{X}) \rightarrow \mathcal{O}_{\mathcal{X}} - \mathrm{Mod}$ is exact in that context. If we work with big sites, then the inclusion $\mathrm{QCoh}(\mathcal{X}) \rightarrow \mathcal{O}_{\mathcal{X}} - \mathrm{Mod}$ is generally not exact (see point (6) of **[Sta, Tag 06VE]**), and so this construction is *not* valid; but a more elaborate version of it can be carried out instead **[Sta, Tag 07B5]**). Again, the resulting category is independent, up to canonical equivalence, of the method of construction.

There is a canonical t -exact functor $D(\mathrm{QCoh}(\mathcal{X})) \rightarrow D_{\mathrm{qc}}(\mathcal{O}_{\mathcal{X}})$. We don't know if this is an equivalence, even under reasonable finiteness hypotheses on \mathcal{X} . However, if \mathcal{X} is quasi-compact and *geometric* (meaning that it has affine diagonal), then this functor does restrict to an equivalence

$$(B.2.2) \quad D^+(\mathrm{QCoh}(\mathcal{X})) \xrightarrow{\sim} D_{\mathrm{qc}}^+(\mathcal{O}_{\mathcal{X}});$$

this is **[ATJLPRVG18, Prop. 1.6]**.

The heart of the t -structure on either $D(\mathrm{QCoh}(\mathcal{X}))$ or $D_{\mathrm{qc}}(\mathcal{O}_{\mathcal{X}})$ is the category $\mathrm{QCoh}(\mathcal{O}_{\mathcal{X}})$. If \mathcal{X} is a Noetherian algebraic stack, then this category is locally coherent, and its full subcategory of compact objects is precisely the subcategory $\mathrm{Coh}(\mathcal{X})$ of coherent sheaves on \mathcal{X} . In particular, the t -structures on both $D(\mathrm{QCoh}(\mathcal{X}))$ and $D_{\mathrm{qc}}(\mathcal{O}_{\mathcal{X}})$ are then coherent. Suppose that \mathcal{X} is furthermore geometric; then the equivalence (B.2.2) induces an equivalence of subcategories of compact objects (which are precisely the coherent objects in $D(\mathrm{QCoh}(\mathcal{X}))$ and $D_{\mathrm{qc}}(\mathcal{O}_{\mathcal{X}})$ respectively). We identify these stable ∞ -categories of coherent objects via this equivalence, and denote either one by $D_{\mathrm{coh}}^b(\mathcal{O}_{\mathcal{X}})$. This is reasonable, since these objects (being coherent objects for a coherent t -structure) are bounded, and the heart of $D_{\mathrm{coh}}^b(\mathcal{O}_{\mathcal{X}})$ is precisely $\mathrm{Coh}(\mathcal{X})$ (again by the general properties of coherent t -structures).

DEFINITION B.2.3. Suppose that \mathcal{X} is a Noetherian and geometric algebraic stack (so that, by the preceding discussion, $D_{\text{coh}}^b(\mathcal{O}_{\mathcal{X}})$ is unambiguously defined). We then define $\text{Ind Coh}(\mathcal{X}) := \text{Ind } D_{\text{coh}}^b(\mathcal{O}_{\mathcal{X}})$, and sometimes refer to it as the stable ∞ -category of Ind-coherent complexes on \mathcal{X} .

REMARK B.2.4. If one works in the ∞ (rather than triangulated) categorical context from the beginning, then it is possible to give yet another definition of (a version of) $D(\text{QCoh}(\mathcal{X}))$, via descent from the affine case; see e.g. [Gai, §1.1.3]. (Working ∞ -categorically is crucial for this definition to work, since the coherent higher homotopies that are implicit in that context are required for descent to be a valid process on the derived level.) This stable ∞ -category is usually denoted $\text{QCoh}(\mathcal{X})$, and in fact is defined in much greater generality — namely for derived prestacks. Unwinding the definition of [Gai, §1.1.3] for an algebraic stack that can be written in the form $\mathcal{X} \xrightarrow{\sim} [U/R]$ where $U = \text{Spec } A$ is an affine scheme and R is an *fppf* groupoid on M , we find that $D_{\text{qc}}(\mathcal{O}_{\mathcal{X}})$ is equal to the stable ∞ -category of R -equivariant objects in $D_{\text{qc}}(U) = D(A)$.

Note that the definition given in Section 5.1.7 of the stable ∞ -category $\mathbf{D}_{\text{coh}}(\mathfrak{X})$ on a rigid analytic Artin stack \mathfrak{X} is analogous to this definition.

One can similarly define $\text{Ind Coh}(\mathcal{X})$ in the same style, and so also in much greater generality, see e.g. [Gai12] and [GR11, §2.3.1].

B.2.5. *The case of Ind-algebraic stacks.* Suppose that \mathcal{X} is an Ind-algebraic stack which admits a description

$$(B.2.6) \quad \mathcal{X} \xrightarrow{\sim} \text{colim}_n \mathcal{X}_n,$$

with the \mathcal{X}_n being geometric Noetherian algebraic stacks, and the transition morphisms $\mathcal{X}_n \rightarrow \mathcal{X}_{n+1}$ being closed immersions. These closed immersions then induce t -exact functors

$$D_{\text{coh}}^b(\mathcal{X}_n) \rightarrow D_{\text{coh}}^b(\mathcal{X}_{n+1}),$$

and thus (by passing to Ind categories)

$$\text{Ind Coh}(\mathcal{X}_n) \hookrightarrow \text{Ind Coh}(\mathcal{X}_{n+1}).$$

DEFINITION B.2.7. We define $\text{Ind Coh}(\mathcal{X}) := \text{colim}_n \text{Ind Coh}(\mathcal{X}_n)$, the colimit being taken in the ∞ -category whose objects are cocomplete stable ∞ -categories and whose morphisms are cocontinuous functors.

REMARK B.2.8. The pushforward functors $\text{Ind Coh}(\mathcal{X}_n) \rightarrow \text{Ind Coh}(\mathcal{X}_{n+1})$ used to define the colimit in Definition B.2.7 are t -exact, and so $\text{Ind Coh}(\mathcal{X})$ is naturally equipped with a t -structure, which is furthermore coherent, and with respect to which $\text{Ind Coh}(\mathcal{X})$ is left regular.

We define $\text{Coh}(\mathcal{X})$ to be the abelian category of compact objects in $\text{Ind Coh}(\mathcal{X})^\heartsuit$. Equivalently, $\text{Coh}(\mathcal{X}) := \text{colim}_n \text{Coh}(\mathcal{X}_n)$.

REMARK B.2.9. Any two descriptions of \mathcal{X} of the form (B.2.6) are mutually cofinal in an evident sense, and thus $\text{Ind Coh}(\mathcal{X})$ and $\text{Coh}(\mathcal{X})$ are well-defined up to canonical equivalence independent of the choice of such a description.

REMARK B.2.10. The definition of $\text{Ind Coh}(\mathcal{X})$ given in [Gai12] and [GR11, §2.3.1], which we recalled in Remark B.2.4 above, can also be applied directly to the Ind-algebraic stack \mathcal{X} . We've given the more concrete definition above simply because we find it easier to think about.

B.3. Pro-coherent sheaves on formal algebraic stacks. Suppose that \mathcal{X} is a geometric Noetherian formal algebraic stack, described as an Ind-algebraic stack via $\text{colim}_n \mathcal{X}_n \xrightarrow{\sim} \mathcal{X}$, where the \mathcal{X}_n are geometric Noetherian algebraic stacks, and the transition maps $\mathcal{X}_n \rightarrow \mathcal{X}_{n+1}$ are thickenings, i.e. closed immersions inducing isomorphisms on underlying reduced substacks. We may then define $\text{Ind Coh}(\mathcal{X})$ and $\text{Coh}(\mathcal{X})$ following Definition B.2.7 and Remark B.2.8.

There are natural “sheaves” on \mathcal{X} that one might want to consider, which do not arise as objects of $\text{Coh}(\mathcal{X})$ according to our definition. For example, if A is an I -adically complete Noetherian ring, then the usual definition of coherent sheaves on $\text{Spf } A$ yields a category equivalent (via passage to global sections) to the category of finitely generated A -modules, whereas according to our definition, $\text{Coh}(\text{Spf } A)$ is equivalent (via passage to global sections) to the subcategory of finitely generated A -modules annihilated by I^n for some n . In particular, the *structure sheaf* on $\text{Spf } A$, corresponding to the ring A itself, is *not* a coherent sheaf on $\text{Spf } A$. Instead, we identify it with the formal pro-object “ $\varprojlim_n A/I^n$ ” and regard it as a *pro-coherent sheaf* on $\text{Spf } A$.

DEFINITION B.3.1. The abelian category of *pro-coherent sheaves* on \mathcal{X} is defined to be the formal pro-category of $\text{Pro Coh}(\mathcal{X})$. This category has enough projectives, and so we may form $D^-(\text{Pro Coh}(\mathcal{X}))$, following [Lur17, §1.3.2]. This is the (bounded above) *stable ∞ -category of pro-coherent complexes* on \mathcal{X} .

EXAMPLE B.3.2. If A is a Cohen–Macaulay I -adically complete Noetherian ring, with dualizing module D , then D also gives rise to a pro-coherent sheaf on $\text{Spf } A$. More generally, if a Noetherian formal algebraic stack admits a dualizing complex, this will naturally be an object of $D^-(\text{Pro Coh}(\mathcal{X}))$. We adopt this point of view on dualizing sheaves/complexes in the body of our notes, since we will want to apply Grothendieck–Serre duality in the context of pro-coherent sheaves.

We note, though, that if we restrict our duality theory to objects of $\text{Ind Coh}(\mathcal{X})$, then we can construct an alternative version of the dualizing complex, which is itself an object of $\text{Ind Coh}(\mathcal{X})$. Namely, if we write $\mathcal{X} \xrightarrow{\sim} \text{colim}_n \mathcal{X}_n$, and let $i_{n,n+1} : \mathcal{X}_n \hookrightarrow \mathcal{X}_{n+1}$ denote the transition maps (assumed to be closed immersions) and $i_n : \mathcal{X}_n \hookrightarrow \mathcal{X}$ denote the canonical closed immersions, then if ω_n denotes the dualizing complex on \mathcal{X}_n , the natural isomorphism $\omega_n \xrightarrow{\sim} i_{n,n+1}^! \omega_{n+1}$ induces a morphism $(i_{n,n+1})_* \omega_n \rightarrow \omega_{n+1}$, and hence a morphism $(i_n)_* \omega_n \rightarrow (i_{n+1})_* \omega_{n+1}$. If we set $\omega := \text{colim}_n (i_n)_* \omega_n$, then ω is an object of $\text{Ind Coh}(\mathcal{X})$ which serves as a dualizing sheaf (for objects of $\text{Ind Coh}(\mathcal{X})$).

As a concrete example, consider $A = k[[T]]$ (with its T -adic topology). Then $k[[T]]$ is a dualizing module for A . On the other hand, if we write $A = \varprojlim_n k[T]/(T^n)$, then the construction of the preceding paragraph gives rise to the A -module $\text{colim}_n k[T]/(T^n)$, the transition maps being given by multiplication by T , which is isomorphic to $k((T))/k[[T]]$. This module corresponds to an object of $\text{Ind Coh}(\text{Spf } A)$, and a consideration of the short exact sequence

$$0 \rightarrow k[[T]] \rightarrow k((T)) \rightarrow k((T))/k[[T]] \rightarrow 0$$

shows that $\text{RHom}_A(-, k[[T]])$ and $\text{RHom}_A(-, k((T))/k[[T]])$ coincide (up to a shift) on objects of $\text{Ind Coh}(\text{Spf } A)$ (since $\text{RHom}_A(-, k((T)))$ vanishes on T -power torsion A -modules).

Appendix C. Localization and Morita theory

“Localization” refers to the idea of interpreting modules M over some ring A as sheaves (of some sort; for us they will always be quasicohherent sheaves, or some variant thereof, such as complexes of quasicohherent sheaves or ind-coherent sheaves) on an appropriate sort of space (for us, a scheme or variant thereof, such as an algebraic stack or a formal algebraic stack). The basic intuition comes from spectral theory: if A is a k -algebra, then an A -module M is a k -vector space endowed with an action of A as a ring of operators. We can then try to decompose M in terms of its various simultaneous eigenspaces with respect to the action of A (if A is commutative) or more generally into a direct sum or direct integral of simple A -modules (if A is non-commutative). We can hope that the collection of simple A -modules has the structure of a space of some sort, and that this decomposition is effected by spreading M out over the space of simple A -modules, with the fibre of the associated sheaf at a given point corresponding to the multiplicity space that the measures the contribution of that particular simple module to M .

C.1. Localizing A -modules over $\text{Spec } A$ (the paradigmatic example).

The most basic example of localization in algebraic geometry is the theory of quasicohherent sheaves over affine varieties: if $X = \text{Spec } A$ is an affine variety, then the functor of global sections

$$(C.1.1) \quad \Gamma(X, -) : \text{QCoh}(X) \rightarrow \Gamma(X, \mathcal{O}_X) - \text{Mod}$$

is an equivalence. The quasi-inverse is, of course, given by the localization functor $M \mapsto \mathcal{O}_X \otimes_A M$.

Another way to describe the global sections functor, when applied to sheaves of \mathcal{O}_X -modules, is as $\text{Hom}_{\mathcal{O}_X}(\mathcal{O}_X, -)$. Thus this fundamental example of localization illustrates a general paradigm: the sheaf \mathcal{O}_X is naturally an (\mathcal{O}_X, A) -bimodule (in a somewhat trivial manner, just via the identification of A with the ring of global sections of \mathcal{O}_X), and so we have the pair of adjoint functors $\mathcal{O}_X \otimes_A -$ and $\text{Hom}_{\mathcal{O}_X}(\mathcal{O}_X, -)$, which together induce an equivalence between the category of A -modules and the category of quasicohherent A -modules. This is an instance of a more general structure provided by *Morita theory*, which we discuss more formally below.

To elaborate even further on this example, we remark that we can regard any commutative ring A as an “abstract” ring of commuting operators, and an A -module M as a concrete realization of A as a ring of operators (namely, via its operation on M). Note then that $X = \text{Spec } A$ can be thought of as parameterizing all the possible systems of simultaneous eigenvalues that the elements of A can assume (since the points of $\text{Spec } A$ correspond, via taking the kernel, to equivalence classes of morphisms $A \rightarrow F$ with F a field, two morphisms $A \rightarrow E$ and $A \rightarrow F$ being equivalent if we can find a common field extension K of E and F such that the induced morphisms $A \rightarrow K$ coincide); hence the designation of X as the “spectrum” of A . Roughly, the support of the quasicohherent sheaf \mathcal{M} associated to M consists of those systems of simultaneous eigenvalues that appear in M . (More precisely, the fibre of \mathcal{M} at a point of X is precisely the co-eigenspace of M for the corresponding system of eigenvalues.) Thus localization provides an algebraic interpretation of spectral theory.

EXAMPLE C.1.2. If we set $A = k[x]$ (for some field k) and apply the preceding discussion, we recover the usual derivation of the spectral theory of a single operator via an analysis of the structure of $k[x]$ -modules.

We are interested in other, less obvious, examples of localization. But in each case, just as in the paradigmatic example of localizing A -modules, we hope that the structure of the sheaf associated to a module (its fibres, its support, and so on) will illuminate, and separate out, various pieces of internal structure that are more obscure when considered purely in terms of modules.

C.2. Morita theory. Morita theory, at least as we intend it here, is a slightly catch-all term for the theory that describes functors, especially equivalences, between categories of modules over various rings. We recall the classical (i.e. abelian, i.e. module-theoretic) case of the theory first, before turning to the derived context (where the ideas of the theory easily extend to provide a flexible and convenient context for studying functors that are not necessarily equivalences).

C.2.1. *The abelian case.* One way to phrase the classical (i.e. module theoretic) version of Morita theory is that if A and B are (not necessarily commutative) associative rings, with respective categories of left modules $A\text{-Mod}$ and $B\text{-Mod}$, then equivalences of categories $A\text{-Mod} \rightarrow B\text{-Mod}$ are all of the form

$$M \mapsto P \otimes_A M,$$

where P is a finitely generated projective B -module which generates $B\text{-Mod}$ equipped with an isomorphism $A \xrightarrow{\sim} \text{End}_B(P)$. A quasi-inverse is given by

$$N \mapsto \text{Hom}_B(P, N).$$

(Of course, Morita theory applies to this adjoint functor as well; we find that $Q := P^\vee := \text{Hom}_B(P, B)$ — which is an (A, B) -bimodule — is necessarily projective as an A -module, and the quasi-inverse can be rewritten as $N \mapsto N \otimes_B Q$.)

More generally, any functor $F : A\text{-Mod} \rightarrow B\text{-Mod}$ which commutes with colimits is of the form

$$(C.2.2) \quad F : M \mapsto P \otimes_A M$$

for some (B, A) -bimodule P . (The bimodule P is given by evaluating the functor on A . The formula (C.2.2) is then obtained by evaluating F on some presentation of M .) The right adjoint of F (which exists since F commutes with colimits) is evidently given by $\text{Hom}_B(P, -)$.

If F is furthermore fully faithful, then we see that $A \rightarrow \text{End}_B(P)$ must be an isomorphism. If it is in fact an equivalence, then P must inherit all the purely categorical properties that A has in $A\text{-Mod}$ (compactness, projectivity, being a generator), which gives the statement of classical Morita theory.

EXAMPLE C.2.3. A standard example of classical Morita theory is given by the case $B = M_n(A)$ and $P = A^{\oplus n}$ (thought of as length n column vectors) with its obvious $(M_n(A), A)$ -bimodule structure. Since $B \xrightarrow{\sim} P^{\oplus n}$ as a module over itself, we see that P is finitely generated and projective, and so $M \mapsto P \otimes_A M = M^{\oplus n}$ gives an equivalence between the categories $A\text{-Mod}$ and $M_n(A)\text{-Mod}$.

Going the other way, if we set $Q = A^{\oplus n}$ (now thought of as length n row vectors), with its obvious $(A, M_n(A))$ -bimodule structure, then $N \mapsto Q \otimes_{M_n(A)} N$ gives the quasi-inverse equivalence between $M_n(A)\text{-Mod}$ and $A\text{-Mod}$.

EXAMPLE C.2.4. When A is commutative, the preceding example can be turned into an example of localization, by writing $X = \text{Spec } A$, and considering the functor

$$N \mapsto \mathcal{O}_X^{\oplus n} \otimes_{M_n(A)} N,$$

which induces an equivalence of categories between $M_n(A) - \text{Mod}$ and the category of quasicoherent sheaves on X .

Of course we are interested in studying modules over non-commutative rings that are more complicated than matrix rings over commutative rings. In such contexts, it is generally not reasonable to expect to obtain an equivalence of categories with the category of modules over a commutative ring, or with the category of quasicoherent sheaves on a scheme. As we will see, it is more plausible to hope for a fully faithful embedding than an equivalence, and it also makes sense to expand the possibilities for the “space” X which will carry the sheaves, e.g. by allowing it to be an algebraic stack.

It is also more natural to pass to a derived setting. Indeed, if one considers the problem of determining when a functor of the form (C.2.2) is fully faithful (but not necessarily an equivalence), one immediately runs into questions about the homological nature of F which are more naturally considered in the derived world. We will thus briefly discuss derived Morita theory next, with an emphasis on the extra flexibility it provides (in comparison to the classical abelian setting) for investigating functors which are not necessarily equivalences.

C.2.5. *The case of stable ∞ -categories.* If A, B are E_1 -rings (see Appendix A.2 for the terminology of E_1 -rings), then by [Lur17, Prop. 7.1.2.4], the cocontinuous functors $\text{LMod}_A \rightarrow \text{LMod}_B$ are given by tensoring with (A, B) -bimodules. In particular if A and B are static (i.e. are associative rings in the usual sense), then equivalences of stable ∞ -categories $D(A - \text{Mod}) \rightarrow D(B - \text{Mod})$ are all of the form

$$(C.2.6) \quad M \mapsto T \otimes_A^L M,$$

where T is a perfect complex of B -modules which generates $D(B - \text{Mod})$, and is equipped with an isomorphism of E_1 -rings $A \xrightarrow{\sim} \text{REnd}_B(T)$, where we write $\text{REnd}_B(T)$ for the E_1 -ring $\text{End}_{D(B - \text{Mod})}(T, T)$. (See [Lur17, Prop. 7.1.2.1, Rem. 7.1.2.3].)

If we drop the assumption that T generates $D(B - \text{Mod})$, but continue to assume that $A \xrightarrow{\sim} \text{REnd}_B(T)$, then (C.2.6) still defines a cocontinuous fully faithful functor $F : D(A - \text{Mod}) \rightarrow D(B - \text{Mod})$, with a right adjoint $G : D(B - \text{Mod}) \rightarrow D(A - \text{Mod})$ given by

$$(C.2.7) \quad N \mapsto \text{RHom}_B(T, N).$$

More generally, if T is a perfect complex of B -modules and S is a perfect complex of A -modules which generates $D(A - \text{Mod})$ and is equipped with an isomorphism $\text{REnd}_A(S) \xrightarrow{\sim} \text{REnd}_B(T) := \mathcal{H}$, then combining the equivalence of $D(A - \text{Mod})$ with $\mathcal{H} - \text{Mod}$ given by S and the cocontinuous fully faithful functor from $\mathcal{H} - \text{Mod}$ to $D(B - \text{Mod})$ given by T , we have a cocontinuous fully faithful functor $D(A - \text{Mod}) \rightarrow D(B - \text{Mod})$ given by

$$(C.2.8) \quad M \mapsto T \otimes_{\mathcal{H}}^L \text{RHom}_A(S, M),$$

with a right adjoint $D(B - \text{Mod}) \rightarrow D(A - \text{Mod})$ given by

$$N \mapsto S \otimes_{\mathcal{H}}^L \text{RHom}_B(T, N).$$

There are, of course, many generalisations of this framework. In particular, in the body of the text we consider examples where $D(B - \text{Mod})$ is replaced by an appropriate derived category of sheaves on a stack, as in the examples below (see also Remark 6.1.23).

REMARK C.2.9. The notation \mathcal{H} for the derived endomorphism ring is intended to be suggestive of a (derived) Hecke algebra. See for example [Hel20, (1.1)] for a functor from a category of representations to a category of sheaves defined in this fashion.

REMARK C.2.10. In the literature, derived Hecke algebras are often thought of as dgas, rather than E_1 -rings. The correspondence between these notions is as follows: by [Lur17, Prop. 7.1.4.6], if A is a (static) commutative ring, then the ∞ -category of E_1 -algebras over A can be identified with the ∞ -category of differential graded A -algebras. For example, if we take $A = \mathbf{F}_p$, Schneider’s derived Hecke algebra [Sch15a] can equally well be regarded as an \mathbf{F}_p -dga or an \mathbf{F}_p - E_1 -algebra.

C.3. Non-commutative phenomena. There are some phenomena that can occur in the theory of modules over non-commutative rings that don’t occur in the commutative case, which have to be taken into account when trying to localize in the non-commutative setting. We briefly recall some of these, and explain why allowing localization over algebraic stacks (rather than just schemes) helps to ameliorate the situation.

There is one basic phenomenon that modules over a non-commutative ring A can exhibit which does not occur in the commutative context, namely, if S_1 and S_2 are non-isomorphic simple A -modules, then although $\text{Hom}_A(S_1, S_2)$ evidently vanishes, the higher Exts, i.e. $\text{Ext}_A^i(S_1, S_2)$ for $i > 0$, need not vanish. This means that if $A - \text{Mod}$ is mapped functorially into a category of sheaves in such a manner that induces an isomorphism on higher Exts as well as just Homs, then the images of S_1 and S_2 cannot have disjoint support.

Remembering that the simple objects in the category of quasicohherent sheaves on a Noetherian²⁸ scheme are precisely the skyscraper sheaves at closed points, so that non-isomorphic simple objects have disjoint support, we see for example that if A is a non-commutative ring admitting simple modules S_1 and S_2 with this property (i.e. such that $\text{Hom}_A(S_1, S_2) = 0$ but $\text{RHom}_A(S_1, S_2) \neq 0$), then the category $A - \text{Mod}$ certainly cannot be *equivalent* to the category of quasicohherent sheaves on such a scheme.

REMARK C.3.1. The preceding discussion helps to explain why in Beilinson–Bernstein localization theory, which does give an equivalence (between certain categories of Lie algebra representations and the category of algebraic \mathcal{D} -modules on a flag variety), it is quasicohherent \mathcal{D} -modules that appear, rather than merely quasicohherent \mathcal{O} -modules. The point is that the Riemann–Hilbert correspondence shows that \mathcal{D} -modules behave like perverse sheaves, in which (to first approximation) the simple objects are described by their support (a stratum in some stratification of the ambient scheme X), and there are interesting extension classes between sheaves supported on different strata. (The most basic example is the sequence

$$0 \rightarrow j_! \mathbf{C} \rightarrow \mathbf{C} \rightarrow i_* \mathcal{C} \rightarrow 0,$$

²⁸We make this assumption just to ensure that the discussion is not overburdened with technical complications.

with $j : \mathbf{A}^1 \setminus \{0\} \rightarrow \mathbf{A}^1$, resp. $i : \{0\} \rightarrow \mathbf{A}^1$ the natural open, resp. closed, immersion.)

If we only want fully faithful functors, the various (im)possibilities are less clear, since then a simple A -module needn't map to a simple quasicoherent sheaf (it might just be that the proper quasicoherent subsheaves of its image are not themselves in the image of the functor). If we pass to the derived context, the situation becomes yet murkier, because then the notions of *simple module* and *submodule* aren't defined. Nevertheless, the difference in behaviour between the non-commutative and commutative situations creates a tension which makes it difficult to construct interesting localization functors for modules over non-commutative rings.

One way to lower this tension is by allowing the target category to be a category of quasicoherent sheaves on an algebraic stack \mathcal{X} , as we now explain.

If $x \in |\mathcal{X}|$ is a closed point of (the underlying topological space of) an algebraic stack \mathcal{X} , then it typically does not correspond to a closed immersion $\mathrm{Spec} k \hookrightarrow \mathcal{X}$ for some field k , but rather (under mild finiteness hypotheses, such as having quasi-compact diagonal; see [Sta, Tag 06RD], as well as the discussion of [EG22, App. E]) to a closed immersion $\mathcal{Z} \hookrightarrow \mathcal{X}$ with \mathcal{Z} being a one-point *gerbe*. If, for example, our stack \mathcal{X} is of finite type over an algebraically closed field k , then $\mathcal{Z} = [\mathrm{Spec} k/G]$ for some algebraic group over G .

Suppose then, to fix ideas, that we have a closed embedding $[\mathrm{Spec} k/G] \hookrightarrow \mathcal{X}$ for some field k and some algebraic group G over k , giving rise to the closed point $x \in |\mathcal{X}|$. Representations of G then correspond to quasicoherent sheaves on \mathcal{X} that are supported at x ; irreducible representations in particular correspond to simple objects in the category of quasicoherent sheaves on \mathcal{X} . In examples, we can often produce non-split extensions of these sheaves by considering the higher infinitesimal neighbourhoods of x in \mathcal{X} .

EXAMPLE C.3.2. Consider $\mathcal{X} = [\mathbf{A}^1/\mathbf{G}_m]$, with $\mathbf{A}^1 = \mathrm{Spec} k[x]$, $\mathbf{G}_m = \mathrm{Spec} k[t^{\pm 1}]$, and \mathbf{G}_m acting on \mathbf{A}^1 via $t \cdot x = tx$. Let $\mathcal{I} \subset \mathcal{O}_{\mathcal{X}}$ denote the ideal sheaf cutting out the closed point $\{x = 0\}$. Then we have a non-split short exact sequence of coherent sheaves

$$0 \rightarrow k(-1) \cong \mathcal{I}/\mathcal{I}^2 \rightarrow \mathcal{O}_{\mathcal{X}}/\mathcal{I}^2 \rightarrow \mathcal{O}_{\mathcal{X}}/\mathcal{I} \cong k(0) \rightarrow 0$$

(where as usual we identify \mathbf{G}_m representations with graded k -modules).

In short, in the category of coherent sheaves on algebraic stacks, it is possible to find non-isomorphic simple objects that admit non-split extensions. This makes categories of coherent sheaves on algebraic stacks amenable targets for localization functors defined on categories of modules over non-commutative rings.

C.4. Examples. We now give some examples of Morita theory and localization, building up from some very simple ones to some which occur in the representation theory of $\mathrm{GL}_2(\mathbf{Q}_p)$. We work throughout over some commutative ring k .

EXAMPLE C.4.1. As we saw in Example C.1.2 above, the category $k[x] - \mathrm{Mod}$ is equivalent to $\mathrm{QCoh}(\mathbf{A}^1)$; this Morita equivalence is mediated by the structure sheaf $\mathcal{O}_{\mathbf{A}^1}$ (which is of course a module over its ring of global sections, namely $k[x]$). Slightly less trivially, as a special case of Example C.2.3, these categories are equivalent to the categories of left modules for the matrix algebra

$$\begin{pmatrix} k[x] & k[x] \\ k[x] & k[x] \end{pmatrix},$$

via the bimodule $\mathcal{O}_{\mathbf{A}^1} \oplus \mathcal{O}_{\mathbf{A}^1}$.

EXAMPLE C.4.2. Consider instead the algebra

$$A := \begin{pmatrix} k[x] & k[x] \\ xk[x] & k[x] \end{pmatrix}$$

Its category of left modules is not Morita equivalent to the category of modules over any commutative ring. However, there is a fully faithful functor from $D(A - \text{Mod})$ to $D(\text{QCoh}(\mathcal{X}))$, where

$$\mathcal{X} := [(\text{Spec } k[x, y]/(xy))/\mathbf{G}_m],$$

with the \mathbf{G}_m -action given by²⁹

$$t \cdot (x, y) = (x, t^2y),$$

and the functor being given by (derived!) tensoring with

$$P := \mathcal{O}_{\mathcal{X}} \oplus (\mathcal{O}_{\mathcal{X}}/y).$$

In order for this to make sense, we need P to be a right A -module. We claim that in fact

$$\text{REnd}_{\mathcal{O}_{\mathcal{X}}}(P) = \text{End}_{\mathcal{O}_{\mathcal{X}}}(P) = A^{\text{op}},$$

so that the functor $P \otimes_A -$ is fully faithful.

In general, for an object $X \oplus Y$ in an abelian category, we have

$$\text{REnd}(X \oplus Y) = \begin{pmatrix} \text{REnd}(X) & \text{RHom}(Y, X) \\ \text{RHom}(Y, X) & \text{REnd}(Y) \end{pmatrix},$$

acting on the left on $X \oplus Y$, thought of as column vectors. Then

$$\text{REnd}(X \oplus Y)^{\text{op}} = \begin{pmatrix} \text{REnd}(X) & \text{RHom}(X, Y) \\ \text{RHom}(Y, X) & \text{REnd}(Y) \end{pmatrix},$$

acting on the right on $X \oplus Y$ (now thought of as row vectors).

In our example, \mathcal{X} is an affine scheme modulo a linearly reductive group, in fact the torus \mathbf{G}_m , and so the global sections functor on \mathcal{X} is exact; equivalently, $\mathcal{O}_{\mathcal{X}}$ is projective, and $\text{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}, -) = \text{Hom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}, -)$ is simply the functor of global sections. In particular

$$\text{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}, \mathcal{O}_{\mathcal{X}}) = \text{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}, \mathcal{O}_{\mathcal{X}}/y) = k[x].$$

To compute $\text{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, -)$, we use the projective resolution

$$\dots \xrightarrow{y} \mathcal{O}_{\mathcal{X}}(-2) \xrightarrow{x} \mathcal{O}_{\mathcal{X}}(-2) \xrightarrow{y} \mathcal{O}_{\mathcal{X}} \rightarrow \mathcal{O}_{\mathcal{X}}/y \rightarrow 0.$$

(Here as usual our notation for invertible sheaves on the quotient stack \mathcal{X} corresponds to our notation for \mathbf{Z} -graded modules, with the convention that $\text{Gr}^i M(n) = \text{Gr}^{i+n} M$.)

We have $\text{Hom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}(-n), \mathcal{O}_{\mathcal{X}}) = k[x]$ if $n = 0$ and $= k \cdot y^n$ if $n > 0$, and similarly $\text{Hom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}(-n), \mathcal{O}_{\mathcal{X}}/y) = k[x]$ if $n = 0$ and $= 0$ if $n > 0$. It follows that $\text{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, \mathcal{O}_{\mathcal{X}})$ is computed by the complex

$$0 \rightarrow k[x] \rightarrow k \xrightarrow{0} k \xrightarrow{\sim} k \xrightarrow{0} k \rightarrow \dots$$

²⁹The appearance of t^2 rather than t in the formula for the action is of no intrinsic importance for this example; its only significance is that it yields an example that connects directly with an instance of categorical Langlands, as we note below.

(where the morphism $k[x] \rightarrow k$ is given by evaluation at 0) and that $\mathrm{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, \mathcal{O}_{\mathcal{X}}/y)$ is computed by the complex

$$0 \rightarrow k[x] \rightarrow 0 \rightarrow 0 \rightarrow \dots,$$

so that

$$\mathrm{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, \mathcal{O}_{\mathcal{X}}) = \mathrm{Hom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, \mathcal{O}_{\mathcal{X}}) = xk[x]$$

and that

$$\mathrm{RHom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, \mathcal{O}_{\mathcal{X}}/y) = \mathrm{Hom}_{\mathcal{O}_{\mathcal{X}}}(\mathcal{O}_{\mathcal{X}}/y, \mathcal{O}_{\mathcal{X}}/y) = k[x].$$

Thus indeed $\mathrm{REnd}_{\mathcal{O}_{\mathcal{X}}}(P)^{\mathrm{op}} = A$.

Continuing with this example a little more, writing $M = k[x] \oplus xk[x]$ (thought of as column vectors, with the evident left action of A) and $N = k[x] \oplus k[x]$ (again thought of as column vectors, with the evident left action of A) then $A = M \oplus N$ as left A -modules, so that M and N are both projective as A -modules, and $P \otimes_A M = \mathcal{O}_{\mathcal{X}}$, while $P \otimes_A N = \mathcal{O}_{\mathcal{X}}/y$.

The evident inclusion

$$(C.4.3) \quad M \hookrightarrow N$$

has cokernel $S_1 := 0 \oplus k$ (whose A -module structure comes by thinking of elements as column vectors and having x act via 0), while the the morphism

$$(C.4.4) \quad N \xrightarrow{x} M$$

has cokernel $S_2 := k \oplus 0$ (which again has the A -module structure given by regarding its elements as column vectors and having x act via 0). The two A -modules S_1, S_2 are simple and not isomorphic (and, up to isomorphism, are precisely the two simple A -modules on which the the element x in the centre $k[x]$ of A acts via 0). The quotient M/xM is a non-split extension of S_2 by S_1 , while N/xN is a non-split extension of S_1 by S_2 .

If we pass to sheaves on \mathcal{X} , the inclusion (C.4.3) corresponds to the surjection $\mathcal{O}_{\mathcal{X}} \rightarrow \mathcal{O}_{\mathcal{X}}/y$, and thus S_1 corresponds to the cone of this morphism, which is $(\mathcal{O}_{\mathcal{X}}/x)(-2)$. The morphism (C.4.4) corresponds to the morphism $\mathcal{O}_{\mathcal{X}}/y \xrightarrow{x} \mathcal{O}_{\mathcal{X}}$, and thus S_2 corresponds to the cone of this morphism, which is $\mathcal{O}_{\mathcal{X}}/x$. Thus we see an example of the possibility discussed in Section C.3, i.e. of two non-isomorphic simple modules, admitting a non-split extension between them, corresponding to coherent sheaves which are non-isomorphic but have the same support.

REMARK C.4.5. Example C.4.2 is closely connected to the ℓ -adic ($\ell \neq p$) representation theory of $\mathrm{PGL}_2(\mathbf{Q}_p)$, see for example [Hel20, Rem. 4.43]. Indeed the Iwahori Hecke algebra \mathcal{H}_I for $\mathrm{PGL}_2(\mathbf{Q}_p)$ is a matrix algebra

$$\mathcal{H}_I = \begin{pmatrix} \mathbf{Q}_{\ell}[T] & \mathbf{Q}_{\ell}[T] \\ (T^2 - (p+1)^2)\mathbf{Q}_{\ell}[T] & \mathbf{Q}_{\ell}[T] \end{pmatrix}$$

and there is a fully faithful functor (determined by the coherent Springer sheaf) from the derived category of left \mathcal{H} -modules to the category of ind-coherent sheaves on the moduli stack $\mathcal{X}_2^{\mathrm{unip}}$ of 2-dimensional unipotently ramified Weil–Deligne representations with inverse cyclotomic determinant.

This stack has three irreducible components; the unramified component, and two ‘‘Steinberg’’ components which are related by a quadratic twist. The two Steinberg components do not intersect, but each intersects the unramified component, at respectively the representation $1 \oplus \varepsilon^{-1}$ and its quadratic twist. In a neighbourhood

of either one of these intersection points we can identify $\mathcal{X}_2^{\text{unip}}$ with (a corresponding neighbourhood of the singular point on) the stack \mathcal{X} above. The locus $y = 0$ then corresponds to the Steinberg component, and P corresponds to the coherent Springer sheaf (see e.g. [Zhu20, Ex. 4.4.4]).

The module M corresponds to the Steinberg representation (or its quadratic twist), whose associated sheaf is $\mathcal{O}_{\mathcal{X}}/y$, the structure sheaf of the Steinberg component; while N corresponds to the trivial representation (or its quadratic twist), whose associated sheaf is $(\mathcal{O}_{\mathcal{X}}/y)(-2)[1]$, a shift of a twist of the structure sheaf of the Steinberg component.

REMARK C.4.6. Note that the object P of Example C.4.2 is not a perfect complex, but it is coherent; thus already in this example we see that it is natural to consider ind-coherent sheaves on stacks of Langlands parameters, rather than quasicohherent sheaves.

The following variant on Example C.4.2 plays a similar role in the p -adic Langlands correspondence for $\text{GL}_2(\mathbf{Q}_p)$, as described in Section 7.3.

EXAMPLE C.4.7. Let A be a (commutative) ring, and consider the graded ring

$$B := A[x, y],$$

where the grading is defined by giving x and y weights 2 and -2 respectively. Let \mathcal{B} denote the category of graded B -modules (equivalently, quasicohherent sheaves on the quotient stack $[\text{Spec } B/\mathbf{G}_m]$). If M is a graded module, and $n \in \mathbf{Z}$, then as usual we let $M(n)$ denote the graded module whose underlying module coincides with that of M , and for which $\text{Gr}^i M(n) = \text{Gr}^{i+n} M$.

We consider the module $P := B(1) \oplus B(-1)$, and define $R := \text{End}_{\mathcal{B}}(P)^{\text{op}}$; then R acts on the right on P . Concretely, we find that

$$\begin{aligned} R &= \begin{pmatrix} \text{End}(B(1)) & \text{Hom}(B(1), B(-1)) \\ \text{Hom}(B(-1), B(1)) & \text{End}(B(-1)) \end{pmatrix} = \begin{pmatrix} \text{Gr}^0 B & \text{Gr}^0 B(-2) \\ \text{Gr}^0 B(2) & \text{Gr}^0 B \end{pmatrix} \\ &= \begin{pmatrix} \text{Gr}^0 B & \text{Gr}^{-2} B \\ \text{Gr}^2 B & \text{Gr}^0 B \end{pmatrix} = \begin{pmatrix} A[xy] & yA[xy] \\ xA[xy] & A[xy] \end{pmatrix} \end{aligned}$$

acting on P via the usual right multiplication of matrices on row vectors.

Certainly P is a projective object of \mathcal{B} . Slightly less obviously, we have the following lemma.

LEMMA C.4.8. P is projective as a right R -module.

PROOF. We first note that as a right module over itself, we have $R = R_1 \oplus R_2$, where $R_1 = A[xy] \oplus yA[xy]$ (the first row of R regarded as above as a matrix order) and $R_2 = xA[xy] \oplus A[xy]$ (the second row of R). Each of R_1 and R_2 is projective (being a direct summand of R).

We next note that $P = \bigoplus_{m=-\infty}^{\infty} \text{Gr}^{1+2m} P$; since R acts as graded endomorphisms of P , this is an isomorphism of R -modules. If $1 + 2m \geq 1$ (i.e. $m \geq 0$), then $\text{Gr}^{1+2m} P = x^m A[xy] \oplus x^{m+1} A[xy] \xrightarrow{\sim} A[xy] \oplus xA[xy] = R_2$, while if $1 + 2m \leq -1$ (i.e. $m \leq -1$), then $\text{Gr}^{1+2m} P = y^{-m-1} A[xy] \oplus y^{-m} A[xy] \xrightarrow{\sim} A[xy] \oplus yA[xy] = R_1$. Thus P is a direct sum of projective R -modules, and so is itself a projective R -module. \square

We have the following immediate consequence.

PROPOSITION C.4.9. *The functor $N \mapsto P \otimes_R N$ induces a fully faithful and exact embedding from the abelian category of R -modules into \mathcal{B} , whose essential image is equal to the full subcategory \mathcal{C} of \mathcal{B} generated by P .*

In fact, since P is projective over B , so that $\mathrm{RHom}_B(P, P) = R$, the preceding proposition has a derived analogue (following the discussion of C.2.5).

PROPOSITION C.4.10. *The functor $N \mapsto P \otimes_R N$ induces a fully faithful and t -exact embedding of stable ∞ -categories with t -structures $D(R) \rightarrow D(\mathcal{B})$, whose essential image is equal to the full subcategory of $D(\mathcal{B})$ generated by P .*

C.4.11. *Characterizing an image.* As already remarked, Example C.4.7 has an application to the categorical p -adic Langlands correspondence in the $\mathrm{GL}_2(\mathbf{Q}_p)$ -case. In fact, in this application we will study a variant of this functor, in which we work with stable ∞ -categories and also impose a support condition.

We now describe our setup precisely. To begin with, we maintain the notation of Example C.4.7, but furthermore assume that the ring A is Noetherian. We also introduce additional notation. Namely, we let \mathcal{R} denote the abelian category of R -modules, and let \mathcal{R}_0 denote the full subcategory of \mathcal{R} consisting of those modules, each element of which is annihilated by some power of xy . Similarly, we let \mathcal{B}_0 denote the subcategory of the category \mathcal{B} of graded B -modules consisting of those modules satisfying the same condition, i.e. each element is annihilated by some power of xy .

Proposition B.1.17 (and its evident variant for graded rings) shows that the embeddings $\mathcal{R}_0 \hookrightarrow R\text{-Mod}$ and $\mathcal{B}_0 \hookrightarrow \mathcal{B}$ induce equivalences

$$(C.4.12) \quad D^+(\mathcal{R}_0) \xrightarrow{\sim} D_{(xy)}^+(R)$$

and

$$(C.4.13) \quad D^+(\mathcal{B}_0) \xrightarrow{\sim} D_{(xy)}^+(\mathcal{B}).$$

(The subscript (xy) here denotes the full subcategory consisting of those objects each of whose cohomologies is annihilated elementwise by some power of xy .)

LEMMA C.4.14. *The equivalences (C.4.12) and (C.4.13) extend to equivalences $D(\mathcal{R}_0) \xrightarrow{\sim} D_{(xy)}(R)$ and $D(\mathcal{B}_0) \xrightarrow{\sim} D_{(xy)}(\mathcal{B})$.*

PROOF. This will follow from Proposition A.7.3 (2), once we show that the adjoint functors to the inclusions $\mathcal{R}_0 \hookrightarrow R\text{-Mod}$ and $\mathcal{B}_0 \hookrightarrow \mathcal{B}$ have bounded cohomological dimension. However, in either case, the derived functor of this adjoint is computed via $X \mapsto \mathrm{colim}_n(X \xrightarrow{(xy)^n} X)$, and so has a cohomological amplitude of 1. \square

The t -exact fully faithful functor $N \mapsto P \otimes_R N$ of Proposition C.4.10 evidently restricts to a t -exact functor fully faithful functor $D_{(xy)}(R) \rightarrow D_{(xy)}(\mathcal{B})$, which by Lemma C.4.14 we may also regard as a t -exact fully faithful functor $D(\mathcal{R}_0) \rightarrow D(\mathcal{B}_0)$. We will characterize the image of this functor in terms of a certain semi-orthogonal decomposition, which we now explain.

The adjoint functor $\mathrm{Hom}_B(P, -) : D_{(xy)}(\mathcal{B}) \rightarrow D_{(xy)}(R)$ admits the more explicit description (remembering that $P := B(1) \oplus B(-1)$)

$$M \mapsto \mathrm{Gr}^{-1} M \oplus \mathrm{Gr}^1 M.$$

(The right hand side being equipped with its evident R -structure: each summand is naturally an $A[xy]$ -module; multiplication by x takes the first summand to the second; and multiplication by y takes the second summand to the first.)

PROPOSITION C.4.15. *Let \mathcal{K} denote the full subcategory of $D_{(xy)}(\mathcal{B})$ consisting of those objects M such that $\mathrm{Gr}^{-1} M = \mathrm{Gr}^1 M = 0$. Then the image of $P \otimes_R -$ is equal to ${}^\perp\mathcal{K}$.*

PROOF. This follows from Lemma A.8.4 (6). □

For applications to categorical p -adic Langlands, we wish to have an alternative description of the kernel \mathcal{K} , in terms of a set of generators. In order to do this, we change the context slightly, replacing the unbounded stable infinity categories $D(\mathcal{R}_0)$ and $D(\mathcal{B}_0)$ by their regularizations $\mathrm{Ind\,Coh}(\mathcal{R}_0)$ and $\mathrm{Ind\,Coh}(\mathcal{B}_0)$, following the procedure described in Section A.6. The functor $P \otimes_R -$ and its adjoint $\mathrm{Hom}_{\mathcal{B}}(P, -)$ are both exact and cocontinuous (the latter because P is finitely generated and projective as a B -module), and preserve compact objects. Thus they induce an adjoint pair of functors between $\mathrm{Ind\,Coh}(\mathcal{R}_0)$ and $\mathrm{Ind\,Coh}(\mathcal{B}_0)$, and the results proved above in the context of $D(\mathcal{R}_0)$ and $D(\mathcal{B}_0)$ carry over to the Ind-coherent context.

The category \mathcal{B}_0 splits into a product $\mathcal{B}_0^+ \times \mathcal{B}_0^-$, where \mathcal{B}_0^+ (resp. \mathcal{B}_0^-) consists of objects whose graded pieces are supported purely in even (resp. odd) degrees. Correspondingly, $\mathrm{Ind\,Coh}(\mathcal{B}_0)$ decomposes as a product $\mathrm{Ind\,Coh}(\mathcal{B}_0^+) \times \mathrm{Ind\,Coh}(\mathcal{B}_0^-)$.

The functor $P \otimes_R -$ evidently factors through $\mathrm{Ind\,Coh}(\mathcal{B}_0^-)$, and we let \mathcal{D}_0^- denote its essential image in this category. We also consider the analogous functor $B \otimes_{A[xy]} - : \mathrm{Ind\,Coh}(\mathcal{A}_0) \rightarrow \mathrm{Ind\,Coh}(\mathcal{B}_0^+)$, where \mathcal{A}_0 denotes the subcategory of $A[xy]$ -Mod consisting of modules annihilated elementwise by some power of xy . We let \mathcal{D}_0^+ denote the essential image in $\mathrm{Ind\,Coh}(\mathcal{B}_0^+)$ of this latter functor.

We regard $A[x]$ as a graded B -module via the isomorphism $B/(y) \xrightarrow{\sim} A[x]$, and similarly for $A[y]$. Similarly we regard A as a graded B -module via the isomorphism $B_i/(x, y) \xrightarrow{\sim} A$. All three of $A[x]$, $A[y]$, and A are then objects of \mathcal{B}_0 .

Using the graded resolution

$$(C.4.16) \quad 0 \rightarrow B(2) \xrightarrow{y} B \rightarrow A[x] \rightarrow 0,$$

we see (for any choice of twist n) that $\mathrm{RHom}_{\mathcal{B}}(A[x], A[x](n))$ is computed by the complex (in cohomological degrees 0 and 1)

$$\mathrm{Gr}_0(A[x](n)) \xrightarrow{0} \mathrm{Gr}_0(A[x](n-2)),$$

while $\mathrm{RHom}_{\mathcal{B}}(A[x], A[y](n))$ is computed by the complex (in cohomological degrees 0 and 1)

$$0 \rightarrow \mathrm{Gr}^0 A(n-2).$$

Similarly, we find that $\mathrm{RHom}_{\mathcal{B}}(A[y], A[y](n))$ is computed by the complex (in cohomological degrees 0 and 1)

$$\mathrm{Gr}_0(A[y](n)) \xrightarrow{0} \mathrm{Gr}_0(A[y](n+2)),$$

while $\mathrm{RHom}_{\mathcal{B}}(A[y], A[x](n))$ is computed by the complex (in cohomological degrees 0 and 1)

$$0 \rightarrow \mathrm{Gr}^0 A(n+2).$$

Consequently, we find that the collections of objects $\{A[x](-2), A[x](-4), \dots\}$ and $\{A[y](2), A[y](4), \dots\}$ each form a weakly exceptional collection in $\mathrm{Ind\,Coh}(\mathcal{B}_0^+)$

in the sense of Definition A.8.10, and are furthermore mutually orthogonal. By Lemma A.8.11, they induce a semiorthogonal decomposition of the full subcategory \mathcal{E}_0^+ of $\text{Ind Coh}(\mathcal{B}_0^+)$ that they generate.

Similarly, the collections $\{A[x](-3), A[x](-5), \dots\}$ and $\{A[y](3), A[y](5), \dots\}$ each form a weakly exceptional collection in $\text{Ind Coh}(\mathcal{B}_0^-)$, and are mutually orthogonal. They similarly induce a semiorthogonal decomposition of the full subcategory \mathcal{E}_0^- of $\text{Ind Coh}(\mathcal{B}_0^-)$ that they generate.

PROPOSITION C.4.17.

- (1) \mathcal{D}_0^+ is the cocomplete full subcategory of $\text{Ind Coh}(\mathcal{B}_0^+)$ generated by the object $B/(xy)$.
- (2) $\mathcal{E}_0^+ = (\mathcal{D}_0^+)^\perp$.
- (3) \mathcal{D}_0^- is the cocomplete full subcategory of $\text{Ind Coh}(\mathcal{B}_0^-)$ generated by the objects $B/(xy)(1)$ and $B/(xy)(-1)$.
- (4) $\mathcal{E}_0^- = (\mathcal{D}_0^-)^\perp$.

PROOF. The cocomplete stable ∞ -category $\text{Ind Coh}(\mathcal{A}_0)$ is generated by $A := A[xy]/(xy)$. Thus \mathcal{D}_0^+ , which is defined to be the image of this category under the functor $B \otimes_{A[xy]} -$, is generated by the image of A , namely $B/(xy)$. This proves (1).

Lemma A.8.4, together with the Ind-coherent analogue of Proposition C.4.15, shows that $(\mathcal{D}_0^+)^\perp$ consists of those objects M for which $\text{Gr}^0 M = 0$. On the other hand, Lemma A.8.6 shows that \mathcal{E}_0^+ can be described as the smallest cocomplete stable subcategory of $D(\mathcal{B}_0^+)$ which contains the set of objects X . Since $\text{Gr} M_0 = 0$ for each object in X , we thus find that $\mathcal{E}_0^+ \subseteq (\mathcal{D}_0^+)^\perp$.

Now, a consideration of (the twist by -2 of) the presentation (C.4.16), along with the fact that $\text{Gr} M_0 = 0$ for objects of (\mathcal{D}_0^+) , shows that $\text{RHom}_{\mathcal{B}}(A[x](-2), M) = \text{Gr}^2 M$ for objects M of $(\mathcal{D}_0^+)^\perp$. In particular, $\text{RHom}_{\mathcal{B}}(A[x](-2), A[x](-2)) = A$. Thus $N \mapsto A[x](-2) \otimes_A N$ is a fully faithful functor $\text{Ind Coh}(A) \rightarrow (\mathcal{D}_0^+)^\perp$, with adjoint given by $M \mapsto \text{Gr}^2 M$. In particular, the cone of $A[x](-2) \otimes_A \text{Gr}^2 M \rightarrow M$ has vanishing Gr^2 . Continuing inductively with a consideration of the various $A[x](-n)$, and then also of the $A[y](n)$ (for positive even values of n), we find that if M is an object of $(\mathcal{E}_0^+)^\perp \cap (\mathcal{D}_0^+)^\perp$, then $\text{Gr}^i M = 0$ for all values of i , and thus that $M = 0$. Consequently, $\mathcal{E}_0^+ = (\mathcal{D}_0^+)^\perp$, proving (2).

The proofs of (3) and (4) are entirely analogous to those of (1) and (2), and so we omit the details. \square

Appendix D. Topological abelian groups, rings and modules

D.1. Functional analysis. Let $\{G_n\}$ be a sequence of linearly topologized³⁰ Hausdorff topological abelian groups, equipped with closed embeddings $G_n \hookrightarrow G_{n+1}$. We can then define the inductive limit $G := \varinjlim_n G_n$, endowed with the inductive limit linear topology (i.e. the topology for which a subset is open if and only if its inverse image in each G_n is open); this is again a linearly topologized topological abelian group. We recall some basic properties that G satisfies; these are analogues of standard properties of inductive limits of sequences of locally convex topological spaces in various contexts (LF -spaces, spaces of compact type, \dots), and our arguments are an adaptation of the arguments of [RR73, Ch. VII].

REMARK D.1.1.

³⁰Meaning that 0 admits a neighbourhood basis of subgroups.

- (1) To simplify the notation, we regard each G_n as a subset of G via the canonical embedding. We will see below that this embedding of groups is furthermore a topological embedding, so this should cause no confusion.
- (2) By definition, a subgroup U of G is open if each pullback $G_n \cap U$ is an open subgroup of G_n for each n . A typical way to construct such open subgroups is to be given an open subgroup V_n of G_n for each n , and to define $U_n := \sum_{i=1}^n V_i$. Then U_n is again an open subgroup of G_n , and evidently $U_n \subseteq U_{n+1}$, so that $U := \bigcup_n U_n = \sum_{i=1}^{\infty} V_i$ is a subgroup of G . Since $U_n \subseteq G_n \cap U$, we see that U is an open subgroup of G . If the V_n satisfy the additional condition $G_n \cap V_{n+1} \subseteq V_n$, then one sees that $U_n = G_n \cap U_{n+1}$, and thus $U_n = G_n \cap U$.

If we are given open subgroups V_n , then we may inductively construct open subgroups $V'_n \subseteq V_n$ such that furthermore $G_n \cap V'_{n+1} \subseteq V'_n$. Namely, assuming V'_i is constructed, for $i = 1, \dots, n$, and since $G_n \hookrightarrow G_{n+1}$ is a topological embedding, we may find an open subgroup W_{n+1} such that $G_n \cap W_{n+1} \subseteq V'_n$. We then set $V'_{n+1} := V_{n+1} \cap W_{n+1}$.

LEMMA D.1.2. *If $X \subseteq G$ has the property that $X \cap G_n$ is finite for each n , then X is closed and discrete as a subset of G .*

PROOF. We first show that X is discrete. To this end, choose $x \in X$; we must find a neighbourhood of x in G which is disjoint from $X \setminus \{x\}$. To ease notation, replace X by $X - x$, and x by 0. Since $G_n \cap (X \setminus \{0\})$ is finite for each n , we may inductively choose open subgroups V_n of G_n so that $U_n := V_1 + \dots + V_n$ is disjoint from $G_n \cap (X \setminus \{0\})$. Then (following Remark D.1.1 (2)), we see that $U := \bigcup U_n$ is an open subgroup of G which is disjoint from $X \setminus \{0\}$.

Now suppose that x is an element of $G \setminus X$. Then $X' := X \cup \{x\}$ again satisfies the hypothesis of the lemma, and thus, by what we have already proved, is discrete. Thus we may find a neighbourhood of x which is disjoint from X , and so x does not lie in the closure of X . This proves that X is closed, as claimed. \square

PROPOSITION D.1.3.

- (1) *Each of the canonical maps $G_n \hookrightarrow G$ is a closed embedding.*
- (2) *G is Hausdorff, and if each G_n is furthermore complete, then G is complete.*
- (3) *Any compact subset of G is contained in some G_n .*
- (4) *If H is a subgroup of G , and if we set $H_n := G_n \cap H$, then each of the natural maps $\varinjlim_n H_n \rightarrow H$ and $\varinjlim_n G_n/H_n \rightarrow G/H$ is a topological isomorphism. Furthermore, H is closed in G if and only if H_n is closed in G_n for each n .*

PROOF. Fix n_0 , and let U_{n_0} be any open subgroup of G_{n_0} . If we set $V_n := U_{n_0} \cap G_n$ for $n \leq n_0$, and choose V_n such that $V_n \cap G_{n-1} \subseteq V_{n-1}$ for $n > n_0$, then we may follow Remark D.1.1 (2) to find an open subgroup $U \subset G$ such that $U \cap G_{n_0} = U_{n_0}$. Thus the topology of G_{n_0} is induced by that of G , proving (1).

To see that G is Hausdorff, we have to show that if $g \in G \setminus \{0\}$, then we can find an open subgroup of G not containing g . Choose n_0 such that $g \in G_{n_0} \setminus \{0\}$, and let U_{n_0} be any open subgroup of G_{n_0} such that $g \notin U_{n_0}$. By the result of the preceding paragraph, we may find an open subgroup U of G such that $G_{n_0} \cap U = U_{n_0}$. Thus we see that $g \notin U$, as required. This proves the first claim of (2).

Suppose now that each G_n is complete, and let \widehat{G} denote the completion of G . The tautological morphism $G \rightarrow \widehat{G}$ is a topological embedding (since G is Hausdorff, as we've just seen). Since each G_n is complete, the composite embedding $G_n \hookrightarrow G \hookrightarrow \widehat{G}$ realizes G_n as a closed subgroup of \widehat{G} . If G is not complete, then we may find $x \in \widehat{G} \setminus G$, and then, since each G_n is closed in \widehat{G} (as we've just noted) and does not contain x , we may find an open subgroup W_n of \widehat{G} such that $x + W_n$ is disjoint from G_n ; we may furthermore choose that W_n so that $W_{n+1} \subseteq W_n$.

Now let $U := \sum_{n=1}^{\infty} W_n \cap G_n$; by Remark D.1.1, this is an open subgroup of G . By the construction of completions, its completion \widehat{U} (i.e. its closure in \widehat{G}) is an open subgroup of \widehat{G} . Since G is dense in \widehat{G} , we see that $(x + \widehat{U}) \cap G$ is non-empty, and thus that $(x + \widehat{U}) \cap G_n$ is non-empty for some value of n . On the other hand, we see that

$$\begin{aligned} \widehat{U} &\subseteq U + W_n = \sum_{i=1}^{\infty} (W_i \cap G_i) + W_n \\ &= \sum_{i \leq n} (W_i \cap G_i) + \sum_{i > n} (W_i \cap G_i) + W_n \subseteq \sum_{i \leq n} G_n + \sum_{i > n} W_i + W_n \\ &= G_n + W_{n+1} + W_n = G_n + W_n, \end{aligned}$$

Thus $(x + G_n + W_n) \cap G_n$ is non-empty, and so $(x + W_n) \cap G_n$ is non-empty, contradicting our choice of W_n . This shows that necessarily $G = \widehat{G}$, completing the proof of (2).

Let $X \subseteq G$ be a compact subset; since G is Hausdorff (by what we've just proved), X is closed. Suppose that $X \not\subseteq G_n$ for every n ; then we may find a sequence x_n of elements of X such that $x_n \notin G_n$. Then $\bigcup_{n=1}^{\infty} \{x_n\}$ is an infinite subset of X which is closed and discrete in G (and thus in X), by Lemma D.1.2. This contradicts the compactness of X , and so (3) is proved.

Now let H be a subgroup of G . By definition of the induced topology, the open subgroups of H are of the form $H \cap U$, with U an open subgroup of G . Write $U_n := G_n \cap U$, so that U_n is an open subgroup of G and $U = \bigcup_{n=1}^{\infty} U_n$. Then $H \cap U_n = H_n \cap U_n = H_n \cap U$ is an open subgroup of H_n , and Remark D.1.1 shows that

$$\bigcup_{n=1}^{\infty} H_n \cap U_n = H \cap \bigcup_{n=1}^{\infty} U_n = H \cap U$$

is also an open subgroup of $\varinjlim H_n$, proving the first claim of (4). The topological isomorphism $\varinjlim G_n/H_n \xrightarrow{\sim} G/H$ is immediately verified via a comparison of the universal mapping properties of the source and target.

Certainly, if H is closed in G then H_n is closed in G_n for each n . Conversely, suppose that H_n is closed in G_n for each n . Then each quotient G_n/H_n is again Hausdorff, and so part (2) shows that $\varinjlim G_n/H_n$ is Hausdorff. Thus H , which is the kernel of the canonical morphism $G = \varinjlim G_n \rightarrow \varinjlim G_n/H_n$, is closed in G , as claimed, completing the proof of (4). \square

D.2. Countably generated topological modules. We first note the following simple lemma.

LEMMA D.2.1. *If A is a left Noetherian (not necessarily commutative) ring, then the category of countably generated A -modules is an abelian subcategory of the category of all A -modules.*

PROOF. The only point that is not completely clear is closure under passage to subobjects. To see this, note that if N is countably generated, then we may write $N = \varinjlim_n N_n$, with each N_n finitely generated. Then if M is an A -submodule of N , we have $M = \varinjlim_n M_n$, where $M_n := M \cap N_n$. Since A is left Noetherian by assumption, we see that each M_n is again finitely generated, and so M is indeed countably generated. \square

Suppose now that A is a compact Hausdorff left Noetherian topological (possibly non-commutative) ring. Any finitely generated A -module then admits a unique compact Hausdorff topology — its *canonical topology* — with respect to which it becomes a topological A -module, and all morphisms between finitely generated A -modules are continuous and strict (i.e. induce topological quotient maps from their domains onto their images) with respect to the canonical topologies on their source and target. (See [Eme10, Prop. 2.1.3] for a proof in a particular case, which immediately generalizes. The key point is that, given a finite presentation $A^{\oplus m} \rightarrow A^{\oplus n} \rightarrow M \rightarrow 0$, the image of the first arrow is necessarily closed, being a compact subset of a Hausdorff space, and so M obtains a natural compact Hausdorff quotient topology, which is easily checked to be independent of the choice of finite presentation. Similar arguments with finite presentations establish the claim about morphisms.)

In the cases we are interested in, the topology on A will be defined by a neighbourhood basis of 0 consisting of open ideals, and then one sees that the canonical topology on any finitely generated A -module is similarly defined by a neighbourhood basis of 0 consisting of open A -submodules. In this case we say that A and M admit *linear topologies*, and we assume from now on that the topology on A is indeed linear.

We let $A^{\oplus \mathbf{N}}$ denote a countable direct sum of copies of A . Writing $A^{\oplus \mathbf{N}} := \varinjlim_n A^{\oplus n}$ (with respect to the evident transition maps), we endow $A^{\oplus \mathbf{N}}$ with its inductive limit linear topology.

PROPOSITION D.2.2.

- (1) *If M is a countably generated A -module, then M admits a unique topology — which we call its canonical topology — with respect to which some (equivalently any) surjection $A^{\oplus \mathbf{N}} \rightarrow M$ is a topological quotient map. Endowed with its canonical topology, M becomes a linear topological A -module.*
- (2) *Morphisms between countably generated A -modules are necessarily continuous and strict with respect to the canonical topology on source and target.*

PROOF. If M is a countably generated A -module, then we may write $M = \varinjlim_n M_n$, where $\{M_n\}$ is an increasing sequence of finitely generated A -submodules of M . If $\{M'_n\}$ is any other increasing sequence of finitely generated A -submodules of M , then evidently $\{M_n\}$ and $\{M'_n\}$ are mutually cofinal, in the sense that each M_n is contained in some $M'_{n'}$, and conversely. Thus, if we endow each M_n with

its canonical topology, the resulting inductive limit linear topology on M , coming from writing it as the colimit of the M_n , is independent of the choice of the sequence $\{M_n\}$. We declare this topology to be the canonical topology of M .

Since each M_n admits a neighbourhood basis of 0 consisting of open A -submodules, it follows from Remark D.1.1 (2) that the same is true of M . One then easily deduces that M is a topological A -module.

If $f : M \rightarrow N$ is a morphism of finitely generated A -modules, then one may choose the $\{M_n\}$ and $\{N_n\}$ such that f restricts to morphisms $M_n \rightarrow N_n$. These morphisms are necessarily continuous with respect to the canonical topology on source and target, and hence so is the morphism

$$f : M = \varinjlim_n M_n \rightarrow \varinjlim_n N_n = N.$$

If f is injective, then we may in fact first choose the N_n , and then set $M_n := M \cap N_n$. Then M_n is closed in N_n , and so Proposition D.1.3 (4) shows that $M \rightarrow N$ is a closed embedding.

If f is surjective, then we may in fact first choose the M_n , and then set $N_n := f(M_n)$. Again by Proposition D.1.3 (4), we see that $M \rightarrow N$ is a topological quotient map.

Altogether, we have now proved (2), and (1) is a special case of the strictness claim of (2). \square

Appendix E. Representations of p -adic analytic groups

E.1. Smooth G -representations and modules over the ring $\mathcal{O}[[G]]$. We suppose that G is a p -adic analytic group, and fix a compact open subgroup H of G . Then the usual completed group ring $\mathcal{O}[[H]] := \varprojlim_{J \triangleleft H \text{ open}} \mathcal{O}[H/J]$ is a compact (two-sided) Noetherian linear topological ring by a result of Lazard [Laz65].

DEFINITION E.1.1. We write

$$\mathcal{O}[[G]] := \mathcal{O}[G] \otimes_{\mathcal{O}[H]} \mathcal{O}[[H]] = c\text{-Ind}_H^G \mathcal{O}[[H]];$$

this is *a priori* a (left) representation of G with a compatible right module structure over $\mathcal{O}[[H]]$, but following [Koh17, §1] (see also [Sho20, §3]), it actually has a natural ring structure, uniquely determined by the requirement that each of $\mathcal{O}[G]$ and $\mathcal{O}[[H]]$ is a subring. The \mathcal{O} -algebra $\mathcal{O}[[G]]$ is furthermore independent of the choice of H up to canonical isomorphism.

Any $\mathcal{O}[[G]]$ -module is in particular an $\mathcal{O}[G]$ -module, and so can be thought of as an \mathcal{O} -linear G -representation. Not every \mathcal{O} -linear G -representation admits a structure of $\mathcal{O}[[G]]$ -module compatible with its $\mathcal{O}[G]$ -action (which is just to say that not every $\mathcal{O}[G]$ -module extends to a $\mathcal{O}[[G]]$ -module), but an important fact is that smooth G -representations (whose definition we now recall) do.

DEFINITION E.1.2.

- (1) If M is an \mathcal{O} -linear G -representation, we say that an element $m \in M$ is *smooth* if m is \mathcal{O} -torsion, and if M is fixed by some open subgroup of G . We let M_{sm} denote the subset of smooth elements of M ; it is a G -invariant \mathcal{O} -submodule of M .
- (2) A representation of G on a \mathcal{O} -module M is called *smooth* if $M_{\text{sm}} = M$.

If M is a smooth G -representation, then the $\mathcal{O}[G]$ -action on M extends canonically to an $\mathcal{O}[[G]]$ -action. Thus the abelian category of smooth \mathcal{O} -linear G -representations may be regarded as a full subcategory of the category of $\mathcal{O}[[G]]$ -modules. (See [Sho20, Lem. 3.5].)

REMARK E.1.3. By definition, an element of an \mathcal{O} -linear G -representation M is smooth if and only if it is smooth when M is regarded as an H -representation for some (or, equivalently, any) compact open subgroup H of G . By definition of the topology on $\mathcal{O}[[H]]$, we see that an element $m \in M$ is smooth if and only if the cyclic $\mathcal{O}[[H]]$ -module generated by m is discrete when regarded as a quotient of $\mathcal{O}[[H]]$.

REMARK E.1.4. Just from the definitions, we see that if M is an \mathcal{O} -linear G -representation, then M_{sm} is the maximal smooth sub- G -representation of M . Furthermore, if M itself is an $\mathcal{O}[[G]]$ -module, then M_{sm} (defined in terms of the underlying $\mathcal{O}[G]$ -module structure) becomes a sub- $\mathcal{O}[[G]]$ -module of M , when endowed with its canonical $\mathcal{O}[[G]]$ -module structure. (This follows from the fact that if H is any compact p -adic analytic group, e.g. a compact open subgroup of G , then the augmentation ideal of $\mathcal{O}[H]$ generates the augmentation ideal of $\mathcal{O}[[H]]$.)

REMARK E.1.5. Most of the preceding discussion (the definitions of $\mathcal{O}[[G]]$ and of smooth G -representations, and the realization of smooth representations as particular kinds of $\mathcal{O}[[G]]$ -modules) applies more generally to any topological group G which contains a profinite open subgroup. However, many of the results that we prove below (for p -adic analytic G) do not hold in this more general context. (See Example E.2.3 below for an illustration of this.)

We let $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$ denote the full subcategory of $\mathcal{O}[[G]] - \text{Mod}$ consisting of countably generated modules.

LEMMA E.1.6. *$(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$ is a Serre subcategory of $\mathcal{O}[[G]] - \text{Mod}$, which has enough projectives. Its projective objects are precisely the countably generated projective $\mathcal{O}[[G]]$ -modules.*

PROOF. The only non-trivial point to check regarding the Serre subcategory claim is that submodules of countably generated $\mathcal{O}[[G]]$ -modules are again countably generated. To see this, note first that, for any compact open subgroup H of G , the ring $\mathcal{O}[[G]]$ is countably generated as a $\mathcal{O}[[H]]$ -module. Thus any countably generated $\mathcal{O}[[G]]$ -module is also countably generated as an $\mathcal{O}[[H]]$ -module. Lemma D.2.1 shows that any $\mathcal{O}[[G]]$ -submodule is then countably generated over $\mathcal{O}[[H]]$, and thus also over $\mathcal{O}[[G]]$.

Evidently any countably generated $\mathcal{O}[[G]]$ -module is a quotient of a countably generated free $\mathcal{O}[[G]]$ -module. Since the latter are projective even in $\mathcal{O}[[G]] - \text{Mod}$, they are projective in $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$. This proves the claim about enough projectives. Any projective object of $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$ is then a retract of a countably generated free module, and so is a countably generated projective module. \square

As already noted in the proof of Lemma E.1.6, any countably generated $\mathcal{O}[[G]]$ -module is also countably generated as an $\mathcal{O}[[H]]$ -module, and so admits a canonical topology (independent of the choice of H), as in Proposition D.2.2.

We note the following elementary lemma.

LEMMA E.1.7. *A countably generated topological $\mathcal{O}[[H]]$ -module M is discrete in its canonical topology if and only if M is a smooth H -representation.*

PROOF. If M is discrete, and if $m \in M$, then the sub- $\mathcal{O}[[H]]$ -module of M generated by m is discrete. Remark E.1.3 then shows that m is smooth. Thus $M_{\text{sm}} = M$.

Conversely, suppose that M is smooth, and choose a surjection $\mathcal{O}[[H]]^{\oplus \mathbb{N}} \rightarrow M$. For any $n \geq 0$, the image of $\mathcal{O}[[H]]^{\oplus n}$ factors through a discrete, and so finite, quotient, and so has finite image $M_n \subseteq M$. We have $M = \varinjlim_n M_n$, and by construction, the canonical topology on M is the inductive limit linear topology on M . Lemma D.1.2 implies that M is discrete, as claimed. \square

We now define a functor from the category $(\mathcal{O}[[G]]\text{-Mod})^{\text{c.g.}}$ of countably generated $\mathcal{O}[[G]]$ -modules to the category $\text{Pro sm. } G$ of pro-smooth G -representations via

$$M \mapsto \text{Pro}(M) := \varprojlim M/U,$$

where U runs over the codirected set of $\mathcal{O}[[G]]$ -submodules of M that are open with respect to the canonical topology on M . It follows from Lemma E.1.7 that M/U is in fact a smooth G -representation. The functoriality of the formation of $\text{Pro}(M)$ follows from the fact that if $f : M \rightarrow N$ is a morphism of countably generated $\mathcal{O}[[G]]$ -modules, then f is continuous with respect to the canonical topologies on its source and target. Indeed, we then see that $f^{-1}(V)$ is an open $\mathcal{O}[[G]]$ -submodule of M whenever V is an open $\mathcal{O}[[G]]$ -submodule of N , so that we have an induced morphism

$$\text{Pro}(M) := \varprojlim M/U \rightarrow \varprojlim M/f^{-1}(V) \rightarrow \varprojlim N/V =: \text{Pro}(N).$$

PROPOSITION E.1.8.

- (1) *The functor $M \mapsto \text{Pro}(M)$ is right exact.*
- (2) *If H is a compact open subgroup of G , and if N is a countably generated $\mathcal{O}[[H]]$ -module, then $\mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} N$ is acyclic for the left-derived functors $L^i \text{Pro}$.*
- (3) *All countably generated smooth representations are acyclic for the left-derived functors $L^i \text{Pro}$.*

PROOF. Let $0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$ be a short exact sequence of countably generated $\mathcal{O}[[G]]$ -modules. If U is an open $\mathcal{O}[[G]]$ -submodule of M , we let $U' = M' \cap U$, and let U'' denote the image of U in M'' . Since $M \rightarrow M''$ is an open G -equivariant mapping, we see that, as U runs over all G -invariant open submodules of M , the images U'' run over all G -invariant open submodules of M'' . Thus we obtain a short exact sequence

$$0 \rightarrow \varprojlim M'/U' \rightarrow \text{Pro}(M) \rightarrow \text{Pro}(M'') \rightarrow 0.$$

There is an evident epimorphism $\text{Pro}(M') \rightarrow \varprojlim M'/U'$, so that

$$(E.1.9) \quad \text{Pro}(M') \rightarrow \text{Pro}(M) \rightarrow \text{Pro}(M'') \rightarrow 0$$

is exact, as claimed.

We now prove (2). Namely, if N is a countably generated $\mathcal{O}[[H]]$ -module, then we may find a resolution $P^\bullet \rightarrow N$ consisting of countably generated free $\mathcal{O}[[H]]$ -modules, and $\mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} P^\bullet$ is then a resolution of $M := \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} N$ by

countable generated free $\mathcal{O}[[G]]$ -modules. (Recall that $\mathcal{O}[[G]]$ is flat over $\mathcal{O}[[H]]$.) Using this resolution to compute the left derived functors $L^i \text{Pro}$, the claim of (2) follows by applying Lemma E.1.10 (2) below.

Finally, we turn to proving (3). Suppose to begin with that M'' is a countably generated smooth G -representation, and let M be a countably generated free $\mathcal{O}[[G]]$ -module admitting a surjection $M \rightarrow M''$. Lemma E.1.7 shows that the kernel M' of this surjection is open in M , and so Lemma E.1.10 (1) below shows that $L^1 \text{Pro}(M'') = 0$; thus $L^1 \text{Pro}$ vanishes on countably generated smooth representations.

Now let H be a compact open subgroup of G . We may write find a surjection $M := \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} N \rightarrow M''$ for some smooth representation N of H . The kernel N' of this surjection is again countably generated and smooth, and since (2) shows that M is $L^\bullet \text{Pro}$ -acyclic, we find that $L^{i+1} \text{Pro}(M'') \xrightarrow{\sim} L^i \text{Pro}(M')$ for $i \geq 1$. An obvious dimension shifting argument, taking into account the result of the preceding paragraph, gives (3). \square

It is not clear to us that (E.1.9) is left-exact in general; such left-exactness would be equivalent to the intersections U' being cofinal in the collection of open G -invariant submodules of M' . (Without the G -invariance condition, this would always hold, since $M' \rightarrow M$ is necessarily a topological embedding.) The following lemma gives some situations in which this *does* hold.

LEMMA E.1.10. (E.1.9) is left-exact if either

- (1) M' is open in M , or
- (2) the embedding $M' \hookrightarrow M$ is of the form

$$M' := \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} N' \hookrightarrow \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} N =: M,$$

where $N' \hookrightarrow N$ is an embedding of $\mathcal{O}[[H]]$ -modules (for some compact open subgroup H of G).

PROOF. (1) If M' is open in M , then any open G -invariant submodule U' of M' is itself an open G -invariant submodule of M , and so we can set $U = U'$. Thus in this case (E.1.9) is exact.

(2) If V' is an open G -invariant submodule of M' , then $W' := N' \cap V'$ is an open $\mathcal{O}[[H]]$ -submodule of N' , and so $U' := \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} W'$ is an open G -invariant submodule of M' contained in V' ; thus submodules of this latter form are cofinal among all open G -invariant submodules of M' . Now since $N' \hookrightarrow N$ is necessarily a topological embedding, we may find an open $\mathcal{O}[[H]]$ -submodule W of N such that $W' = N' \cap W$. Then $U := \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} W$ is an open G -invariant submodule of M for which $M' \cap U = U'$. Thus in this case (E.1.9) is again exact. \square

E.2. Derived categories of smooth representations. Let G be a p -adic analytic group, and let $\mathcal{O}[[G]]$ be the ring of Definition E.1.1. We let $D(\mathcal{O}[[G]])$ denote the stable ∞ -category associated to the abelian category of $\mathcal{O}[[G]]$ -modules. It admits its standard t -structure, with respect to which it is both left and right complete. (This is true for the stable ∞ -category associated to the abelian category of modules over any ring, as we noted in Section A.5.)

DEFINITION E.2.1. We let $D_{\text{sm}}(\mathcal{O}[[G]])$ denote the full subcategory of $D(\mathcal{O}[[G]])$ consisting of objects whose cohomologies in every degree are smooth G -representations.

The t -structure on $D(\mathcal{O}[[G]])$ induces a t -structure on $D_{\text{sm}}(\mathcal{O}[[G]])$, whose heart is the abelian category $\text{sm. } G$ of smooth G -representations. The full subcategory $\text{sm. } G$ of $\mathcal{O}[[G]]\text{-Mod}$ is closed under passage to subobjects, quotients, and colimits in $\mathcal{O}[[G]]$. Thus $\text{sm. } G$ is again a Grothendieck abelian category, and in particular admits enough injectives.

PROPOSITION E.2.2. *There is a canonical equivalence $D(\text{sm. } G) \xrightarrow{\sim} D_{\text{sm}}(\mathcal{O}[[G]])$.*

PROOF. The inclusion of smooth G -representations as a full subcategory of $\mathcal{O}[[G]]$ -modules (which is an exact functor!) induces a t -exact functor $F : D(\text{sm. } G) \rightarrow D(\mathcal{O}[[G]])$, which evidently factors through $D_{\text{sm}}(\mathcal{O}[[G]])$. We begin by showing that the restriction of F to $D^+(\text{sm. } G)$ is fully faithful.

Firstly, we note that $\text{sm. } G$ is precisely the heart of $D_{\text{sm}}(\mathcal{O}[[G]])$, and so the restriction of F to $D^+(\text{sm. } G)$ is exactly the functor of Proposition A.5.6. A consideration of that proposition shows that to show that this restriction of F is fully faithful, it suffices to show that if X is a smooth G -representation, and Y is an injective smooth G -representation, then we may find a smooth G -representation Z admitting a surjection $Z \rightarrow X$ such $\text{Ext}_{\mathcal{O}[[G]]}^i(Z, Y) = 0$ for $i > 0$. For this, choose a compact open subgroup H of G , and set $Z = c\text{-Ind}_H^G X = \mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} X$, which admits a natural surjection to X .

Now the forgetful functor $\mathcal{O}[[G]]\text{-Mod} \rightarrow \mathcal{O}[[H]]\text{-Mod}$ admits an exact left adjoint, namely $\mathcal{O}[[G]] \otimes_{\mathcal{O}[[H]]} -$. Thus it preserves injectives. This same functor, restricted to $\text{sm. } H$, where it can be interpreted as $c\text{-Ind}_H^G$, provides a left adjoint to the forgetful functor $\text{sm. } G \rightarrow \text{sm. } H$; thus this latter functor also preserves injectives. Consequently $\text{Ext}_{\mathcal{O}[[G]]}^i(c\text{-Ind}_H^G X, Y) = \text{Ext}_{\mathcal{O}[[H]]}^i(X, Y)$, with Y injective in $\text{sm. } H$, and so we reduce to verifying that if X and Y are objects of $\text{sm. } H$ with Y injective, then $\text{Ext}_{\mathcal{O}[[H]]}^i(X, Y) = 0$ for $i > 0$. In short, we have replaced G in the original problem by its compact open subgroup H . In this case, the required statement follows from Corollary B.1.16 (applied in the context of Example B.1.1).

We next note that the inclusion $\text{sm. } G \hookrightarrow \mathcal{O}[[G]]\text{-Mod}$ admits a right adjoint (as it must, by the adjoint functor theorem, since colimits of smooth representations are smooth), which admits a quite explicit description: it is

$$M \mapsto M_{\text{sm}} := \text{colim}_{H,n} M[\varpi^n]^H,$$

where in the colimit, n ranges over positive integers and H ranges over open subgroups of G . By Proposition A.7.1, we may then form the corresponding right derived functor $R\Gamma_{\text{sm}} : D(\mathcal{O}[[G]]) \rightarrow D(\text{sm. } G)$, which is right adjoint to F . Furthermore, $R\Gamma_{\text{sm}}$ has finite cohomological amplitude.³¹ The proposition now follows from Proposition A.7.3 (2). \square

EXAMPLE E.2.3. As the following (counter)example will show, Proposition E.2.2 genuinely depends upon the hypothesis that G is p -adic analytic, or at least on various properties of the completed group ring $\mathcal{O}[[G]]$ as a pro-Artinian ring which need not hold for more general locally pro- p -groups.

For our example, we take $G = \prod_{i=0}^{\infty} \mathbf{Z}/p\mathbf{Z}$, with its natural profinite topology. To ease the comparison with the results we cite, we work with coefficients in k rather than \mathcal{O} , but the reader may easily convince themselves that this alteration

³¹If G has dimension d , and if H is sufficiently small (so that it is a uniform pro- p -group), then the derived functor of $M \mapsto M[\pi^n]^H$ has cohomological amplitude $d + 1$, and so the same is true of $R\Gamma_{\text{sm}}$ (which is the filtered colimit of these various derived functors).

of the coefficients plays no essential role, and that the discussion readily extends to the case of \mathcal{O} -coefficients. The group ring of the product $G_n := \prod_{i=0}^n \mathbf{Z}/p\mathbf{Z}$ admits the description $k[G_n] = k[x_0, \dots, x_n]/(x_0^p, \dots, x_n^p)$, and so the completed group ring $k[[G]]$ admits the description

$$k[[G]] = k[[x_0, \dots, x_n, \dots]]/(x_0^p, \dots, x_n^p, \dots) := \varprojlim_n k[[x_0, \dots, x_n]]/(x_0^p, \dots, x_n^p).$$

If we let \mathfrak{m} denote the maximal ideal (i.e. the augmentation ideal) of $k[[G]]$, then although \mathfrak{m} is open in $k[[G]]$, the natural profinite topology on $k[[G]]$ is weaker than its \mathfrak{m} -adic topology. E.g. $\mathfrak{m}/\mathfrak{m}^2 = \varprojlim_n k\langle x_0, \dots, x_n \rangle$ is infinite dimensional, and the profinite topology on $k[[G]]$ endows $\mathfrak{m}/\mathfrak{m}^2$ with its natural profinite topology, rather than the discrete topology.

Let R denote the polynomial subring $k[x_0, \dots, x_n, \dots]$ of $k[[G]]$. The category of smooth G -representations may be identified with the category of R -modules in which each element is killed by one of the ideals $(x_i)_{i \geq n}$ for some $n \geq 0$. This is an example of one of the categories denoted \mathcal{A} that are introduced in [Nee11, Construction 1.1], and [Nee11, Thm. 1.1] then shows that $D(\text{sm. } G)$ is not left complete, while [Nee11, Rem. 1.2] shows that the formation of countable products in $D(\text{sm. } G)$ is of unbounded cohomological amplitude (illustrating the necessity of *some* bounded amplitude criterion in Proposition A.5.3). Thus the argument in the proof of Proposition E.2.2 which extends the result from $D^+(\text{sm. } G)$ to $D(\text{sm. } G)$ breaks down.

In fact even the canonical functor $F : D^+(\text{sm. } G) \rightarrow D^+(k[[G]])$ is not fully faithful. Indeed, the one-dimensional trivial representation of G corresponds to the $k[[G]]$ -module $k := k[[G]]/\mathfrak{m}$, and so $\text{Ext}_{k[[G]]}^1(k, k) = \text{Hom}_k(\mathfrak{m}/\mathfrak{m}^2, k)$ (the k -linear dual to k). On the other hand $\text{Ext}_{\text{sm. } G}^1(k, k)$ corresponds to the proper subspace of $\text{Hom}_k(\mathfrak{m}/\mathfrak{m}^2, k)$ consisting of functionals that are continuous when $\mathfrak{m}/\mathfrak{m}^2$ is endowed with its natural profinite topology.

E.3. Compact and coherent objects. We begin with some preliminaries based on the results of [Sho20, §§2,3].

LEMMA E.3.1. *If M is a smooth \mathcal{O} -linear G -representation, then the following are equivalent:*

- (1) *M is finitely generated as an $\mathcal{O}[G]$ -module.*
- (2) *M is finitely generated as an $\mathcal{O}[[G]]$ -module.*
- (3) *There is a compact open subgroup H of G , a smooth representation of H on a finitely generated torsion \mathcal{O} -module N , and a surjection $c\text{-Ind}_H^G N \rightarrow M$.*
- (4) *For every compact open subgroup H of G , there is a smooth representation of H on a finitely generated torsion \mathcal{O} -module N , and a surjection $c\text{-Ind}_H^G N \rightarrow M$.*

PROOF. Clearly (4) implies (3) implies (1) implies (2). Suppose then that (2) holds, and let $S \subset M$ be a finite generating set for M as an $\mathcal{O}[[G]]$ -module. If H is any compact open subgroup of G , and if N is the $\mathcal{O}[[H]]$ -submodule of M generated by S (which coincides with the $\mathcal{O}[H]$ -submodule generated by S , since the smoothness of M implies that an open subgroup of H acts trivially on S), then $c\text{-Ind}_H^G = \mathcal{O}[G] \otimes_{\mathcal{O}[H]} N$ surjects onto M , proving (4). \square

DEFINITION E.3.2. We say that a smooth G -representation is finitely generated if it satisfies the equivalent conditions of Lemma E.3.1.

DEFINITION E.3.3. A smooth representation of G on an \mathcal{O} -module M is of *finite presentation* if there is an exact sequence

$$M_1 \rightarrow M_2 \rightarrow M \rightarrow 0$$

of smooth G -representations with M_1 and M_2 finitely generated.

The following result is [Sho20, Prop. 3.8].

LEMMA E.3.4. *A smooth G -representation is of finite presentation (in the sense of Definition E.3.2) if and only if it is finitely presented as an $\mathcal{O}[[G]]$ -module.*

We now have the following result.

PROPOSITION E.3.5. *The category $\text{sm. } G$ is compactly generated. The compact objects of $\text{sm. } G$ are precisely the smooth representations of finite presentation.*

PROOF. We sketch the proof, which is standard. To begin with, note that if M is a smooth \mathcal{O} -linear G -representation, and $\{M_i\}$ is the (directed!) set of finitely generated subrepresentations of M , then evidently $M = \bigcup_i M_i = \text{colim } M_i$. For each value of i , choose a surjection $c\text{-Ind}_{H_i}^G N_i \rightarrow M_i$, with H_i compact open in G and N_i a smooth representation of H_i on a finitely generated torsion \mathcal{O} -module. If M'_i denotes the kernel of this surjection, then, as we have just seen, $M'_i = \text{colim}(M'_i)_j$, where $(M'_i)_j$ runs over the finitely generated subrepresentations of M'_i . Then $M = \text{colim}_{i,j} (c\text{-Ind}_{H_i}^G N_i) / (M'_i)_j$, which exhibits M as a colimit of representations of finite presentation. Thus, if we prove that representations of finite presentation are compact, we will have proved that $\text{sm. } G$ is compactly generated.

Now if M is compact, then applying its compactness property to its representation as a colimit of representations of finite presentation, we find that M is a direct summand of a representation of finite presentation, and so is easily seen itself to be a representation of finite presentation (either directly, or via an application of Lemma E.3.4). Conversely, if M is of finite presentation, then Lemma E.3.4 shows that M is compact (even in the category of $\mathcal{O}[[G]]$ -modules). \square

We don't know in general whether the smooth G -representations of finite presentation form an abelian subcategory of $\text{sm. } G$, although it seems plausible that they do; see Conjecture 6.1.4, as well as Remark 6.1.6 for a discussion of known results in this direction. If they do, then we find that $\text{sm. } G$ is a locally coherent abelian category, and thus that the t -structure on $D(\text{sm. } G)$ is coherent. In this case we find (applying Lemma A.6.7) that the subcategory of coherent objects in $D(\text{sm. } G)$ is precisely the full subcategory $D_{\text{f.p.}}^b(\text{sm. } G)$ of (cohomologically) bounded complexes whose cohomologies are finitely presented.

E.4. Complexes of $\mathcal{O}[[G]]$ -modules as pro-smooth complexes. The formal Pro-category $\text{Pro sm. } G$ admits enough projectives. Thus we may form its associated stable ∞ -category $D^-(\text{Pro sm. } G)$, with its natural t -structure.

We may also form the formal Pro- ∞ -category $\text{Pro } D^b(\text{sm. } G)$. The t -structure on $D^b(\text{sm. } G)$ induces a left complete t -structure on $\text{Pro } D^b(\text{sm. } G)$ (this is the dual statement to Proposition A.6.14), whose heart is equivalent to $\text{Pro sm. } G$. Thus (by the dual version of Proposition A.5.6, i.e. [Lur17, Prop. 1.3.3.7]) there is a canonical functor

$$(E.4.1) \quad D^-(\text{Pro sm. } G) \rightarrow \text{Pro } D^b(\text{sm. } G).$$

Recall that $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$ denotes the abelian category of countably generated $\mathcal{O}[[G]]$ -modules. By Lemma E.1.6, this is an abelian category with enough projectives (namely, the countably generated free $\mathcal{O}[[G]]$ -modules), and so we may consider its stable ∞ -category of bounded above complexes $D^-((\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}})$. Since in fact (by Lemma E.1.6) $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$ is a Serre subcategory of $\mathcal{O}[[G]] - \text{Mod}$, it also makes sense to consider $D_{\text{c.g.}}^-(\mathcal{O}[[G]])$, the full sub-stable ∞ -category of $D^-(\mathcal{O}[[G]])$ consisting of complexes with countably generated cohomologies. The canonical t -structure on $D^-(\mathcal{O}[[G]])$ induces a t -structure on $D_{\text{c.g.}}^-(\mathcal{O}[[G]])$, whose heart is precisely $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$.

LEMMA E.4.2. *There is a canonical equivalence*

$$(E.4.3) \quad D^-((\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}) \xrightarrow{\sim} D_{\text{c.g.}}^-(\mathcal{O}[[G]]).$$

PROOF. This follows from the dual version of Proposition A.5.6, i.e. [Lur17, Prop. 1.3.3.7], since the projective objects in $(\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}$ are precisely the countably generated projective $\mathcal{O}[[G]]$ -modules, which can be used to compute Exts in $\mathcal{O}[[G]] - \text{Mod}$. \square

Now left deriving the functor of Proposition E.1.8 gives a functor

$$L \text{Pro} : D^-((\mathcal{O}[[G]] - \text{Mod})^{\text{c.g.}}) \rightarrow D^-(\text{Pro sm. } G).$$

Precomposing with an inverse to the equivalence of (E.4.3), and postcomposing with (E.4.1), we obtain a functor

$$(E.4.4) \quad D_{\text{c.g.}}^-(\mathcal{O}[[G]]) \rightarrow \text{Pro } D^b(\text{sm. } G).$$

Let $D_{\text{c.g.}}^b(\text{sm. } G)$ denote the full subcategory of $D^b(\text{sm. } G)$ consisting of objects whose cohomologies are countably generated, and let $D_{\text{c.g.sm}}^b(\mathcal{O}[[G]])$ denote the full subcategory of $D^b(\mathcal{O}[[G]])$ consisting of complexes whose cohomologies are countably generated and smooth. The equivalence of Proposition E.2.2 evidently induces an equivalence

$$(E.4.5) \quad D_{\text{c.g.}}^b(\text{sm. } G) \xrightarrow{\sim} D_{\text{c.g.sm}}^b(\mathcal{O}[[G]]).$$

It follows from Proposition E.1.8 (3) that the composite of the equivalence (E.4.5) with (the restriction to $D_{\text{c.g.sm}}^b(\mathcal{O}[[G]])$ of) the functor (E.4.4) is equivalent to the composite of the canonical functors

$$D_{\text{c.g.}}^b(\text{sm. } G) \hookrightarrow D^b(\text{sm. } G) \hookrightarrow \text{Pro } D^b(\text{sm. } G).$$

In summary, we have the functor $D_{\text{c.g.}}^-(\mathcal{O}[[G]]) \rightarrow \text{Pro } D^b(\text{sm. } G)$ of (E.4.4), which when restricted to $D_{\text{c.g.}}^b(\text{sm. } G)^b$ (thought of as a subcategory of $D_{\text{c.g.}}^-(\mathcal{O}[[G]])$ via the equivalence (E.4.5)) is equivalent to the evident inclusion.

Appendix F. Derived moduli stacks of group representations

In this appendix we give a brief introduction to and motivation of derived moduli stacks of Galois representations. We refer to [Zhu20, §2] for the details.

Let

$$\Gamma := \langle g_1, \dots, g_m \mid r_1 = \dots = r_n = 1 \rangle$$

be a finitely presented group, with generators g_1, \dots, g_m and relations r_1, \dots, r_n . (The motivational example to bear in mind is the Galois group of a finite extension of local or global fields; in the body of the notes we will consider instead various topologically finitely presented profinite groups, given for example by the absolute

Galois groups of p -adic local fields, but for the sake of motivation the case of a literally finite group suffices.)

Then giving a representation $\Gamma \rightarrow \mathrm{GL}_d(A)$ (for some ring A) amounts to giving elements $x_1, \dots, x_m \in \mathrm{GL}_d(A)$ satisfying the relations r_j . Let V_Γ be the closed subscheme of $(\mathrm{GL}_d)^m$ classifying m -tuples satisfying the relations r_j . There is an action of GL_d on V_Γ via simultaneous conjugation, and the quotient stack $\mathcal{X}_\Gamma := [V_\Gamma/\mathrm{GL}_d]$ is then the moduli stack of d -dimensional representations of Γ (over arbitrary commutative rings, as we have written it here, although of course we could work over any given base ring \mathcal{O} , and then consider representations over \mathcal{O} -algebras). It is an algebraic stack (over \mathbf{Z} , or whichever other base ring \mathcal{O} we choose).

If we fix a particular representation $\rho : \Gamma \rightarrow \mathrm{GL}_d(l)$ (defined over some field l), then this representation gives rise to a l -valued point of \mathcal{X}_Γ , and the lifting ring R_ρ of ρ is a versal ring to \mathcal{X}_Γ at the point ρ . The tangent space to \mathcal{X}_Γ at ρ can be described in terms of group cohomology as $H^1(\Gamma, \mathrm{Ad}\rho)$, while $H^0(\Gamma, \mathrm{Ad}\rho)$ computes the infinitesimal automorphisms of ρ . Passing to higher degree cohomology, one knows that elements of $H^2(\Gamma, \mathrm{Ad}\rho)$ give rise to obstructions to higher order deformations of ρ . However, this obstruction theory is an extra piece of data — it is not determined intrinsically by the stack \mathcal{X}_Γ itself. For example, in general one can't “compute” $H^2(\Gamma, \mathrm{Ad}\rho)$ in terms of \mathcal{X}_Γ (unlike H^0 and H^1).

This comes up concretely when one uses group cohomology to describe the lifting ring R_ρ : if we consider liftings over l -algebras, and write $h^i := \dim_l H^i(\Gamma, \mathrm{Ad}\rho)$, then we can write R_ρ as a quotient $l[[x_1, \dots, x_{h^1-h_0+d^2}]]/(f_1, \dots, f_{h^2})$. When \mathcal{X}_Γ has the “expected dimension” $-h^0 + h^1 - h^2$ at ρ , the f_i must form a regular sequence in $l[[x_1, \dots, x_{h^1-h_0+d^2}]]$, and R_ρ is a local complete intersection. But if \mathcal{X}_Γ has larger-than-expected dimension, then the ring-theoretic structure of R_ρ is less clear, and less clearly related to the quantity h^2 .

One reason, then, for considering a *derived* analog of \mathcal{X}_Γ is that it tightens the relationship between group cohomology and the infinitesimal structure of \mathcal{X}_Γ at its points ρ . The precise definition of the derived representation stack is somewhat involved (see [Zhu20, §2]), but it proceeds roughly as follows: as is usual in a derived setting, rather than just taking into account the relations r_i appearing in Γ , we consider the “higher relations”, i.e. the relations between the relations and so on. In the non-abelian setting of not-necessarily-commutative groups, this can be expressed formally by finding a simplicial resolution Γ_\bullet of Γ by free groups (or, in fact, for technical reasons, by free monoids). The schemes V_{Γ_n} classifying homomorphisms $\Gamma_n \rightarrow \mathrm{GL}_d(A)$ (for coefficient rings A) are then of the form $\mathrm{Spec} R_{\Gamma_n}$ for certain rings R_{Γ_n} that are localizations of polynomial rings (since the Γ_n are free monoids), and which are themselves organized into a simplicial commutative ring R_{Γ_\bullet} . We may consider the corresponding derived affine scheme V_{Γ_\bullet} , and form the derived algebraic stack $\mathcal{X}_{\Gamma_\bullet} := [V_{\Gamma_\bullet}/\mathrm{GL}_d]$, whose underlying classical substack coincides with \mathcal{X}_Γ . In a suitable ∞ -categorical framework (see e.g. the discussion of [Zhu20, §2]) the derived stack $\mathcal{X}_{\Gamma_\bullet}$ is well-defined independently of the choice of resolution Γ_\bullet .

There is a standard useful analogy, in which one compares the passage from classical to derived rings (or, in the terminology of [CS21], which we employ, “animated rings”) to the passage from reduced rings to arbitrary rings that occurs in passing from varieties to schemes, and the passage from classical to derived schemes

or stacks to the passage from varieties to not-necessarily-reduced schemes. Even for a classical variety over a field l , it is interesting to consider its points over non-reduced rings: e.g. its tangent spaces are computed via $l[\varepsilon]/(\varepsilon^2)$ -valued points. Similarly, for a derived stack \mathcal{X} (even one which is actually classical), there are non-classical analogues of $l[\varepsilon]/(\varepsilon^2)$ whose points in a classical or derived stack \mathcal{X} determine the obstruction theory of \mathcal{X} , and using such points, we can extend the classical computation of tangent spaces to compute the cotangent complex $\mathbf{L}_{\mathcal{X},x}$ of \mathcal{X} at points x of \mathcal{X} . In particular, when $\mathcal{X} = \mathcal{X}_{\Gamma,\bullet}$, one finds that

$$\mathbf{L}_{\mathcal{X}_{\Gamma,\bullet},\rho} = C^\bullet(\Gamma, \mathrm{Ad}\rho)^*[-1]$$

(a shift of the dual to the cochain complex that computes the group cohomology $H^\bullet(\Gamma, \mathrm{Ad}\rho)$) — see e.g. [Zhu20, §2.2]. This extends the geometric interpretations of H^0 and H^1 in the classical context.

This is especially useful in cases when $H^i(\Gamma, \mathrm{Ad}\rho) = 0$ for $i > 2$, since then the cohomology of $\mathbf{L}_{\mathcal{X}_{\Gamma,\bullet},\rho}$ appears only in degrees $[-1, 1]$. This implies that $\mathcal{X}_{\Gamma,\bullet}$ is *quasi-smooth* (i.e. derived l.c.i.) in a neighbourhood of $\bar{\rho}$. Concretely, this means that there is a formally smooth morphism

$$\mathrm{Spec} \, l[[x_1, \dots, x_{h^1-h_0+d^2}]]/(f_1, \dots, f_{h^2}) \rightarrow \mathcal{X}_{\Gamma,\bullet}$$

taking the closed point of the domain to ρ , provided that the domain is understood in a derived sense: if the f_i do not form a regular sequence, then the ring $l[[x_1, \dots, x_{h^1-h_0+d^2}]]/(f_1, \dots, f_{h^2})$ will be a non-classical animated ring (its higher homotopy groups will be computed by the Koszul complex of the sequence (f_i)). We then see that $\mathcal{X}_{\Gamma,\bullet}$ is in fact classical in a neighbourhood of ρ if and only if the underlying classical stack \mathcal{X}_Γ is of the expected dimension $-h_0+h_1-h_2 = \chi(\mathbf{L}_{\mathcal{X}_{\Gamma,\bullet},\rho})$ in such a neighbourhood.

In Sections 8 and 9 we apply these ideas when Γ is (at least closely related to) a local or global absolute Galois group. (Their extension to the context of profinite groups is treated in [Zhu20, §2.4].) To this end, we recall that local Galois cohomology always vanishes in degrees > 2 , and that the same is true for global Galois cohomology, provided we work with odd residue characteristics (as we always will).

References

- [AC14] Patrick B. Allen and Frank Calegari, *Finiteness of unramified deformation rings*, Algebra Number Theory **8** (2014), no. 9, 2263–2272.
- [AGH19] Benjamin Antieau, David Gepner, and Jeremiah Heller, *K-theoretic obstructions to bounded t-structures*, Invent. Math. **216** (2019), no. 1, 241–300.
- [AHR20] Jarod Alper, Jack Hall, and David Rydh, *A Luma étale slice theorem for algebraic stacks*, Ann. of Math. (2) **191** (2020), no. 3, 675–738.
- [All19] Patrick B. Allen, *On automorphic points in polarized deformation rings*, American Journal of Mathematics **141** (2019), no. 1, 119–167.
- [Alp13] Jarod Alper, *Good moduli spaces for Artin stacks*, Ann. Inst. Fourier (Grenoble) **63** (2013), no. 6, 2349–2402.
- [Alp14] ———, *Adequate moduli spaces and geometrically reductive group schemes*, Algebr. Geom. **1** (2014), no. 4, 489–531.
- [And21] Grigory Andreychev, *Pseudocoherent and perfect complexes and vector bundles on analytic adic spaces*, 2021.
- [Art89] James Arthur, *Unipotent automorphic representations: conjectures*, no. 171–172, 1989, Orbites unipotentes et représentations, II, pp. 13–71.

- [ATJLPRVG18] Leovigildo Alonso Tarrío, Ana Jeremías López, Marta Pérez Rodríguez, and María J. Vale Gonsalves, *On the derived category of quasi-coherent sheaves on an Adams geometric stack*, J. Pure Appl. Algebra **222** (2018), no. 4, 828–845.
- [BB10] Laurent Berger and Christophe Breuil, *Sur quelques représentations potentiellement cristallines de $GL_2(\mathbf{Q}_p)$* , Astérisque (2010), no. 330, 155–211.
- [BC08] Laurent Berger and Pierre Colmez, *Familles de représentations de de Rham et monodromie p -adique*, no. 319, 2008, Représentations p -adiques de groupes p -adiques. I. Représentations galoisiennes et (ϕ, Γ) -modules, pp. 303–337.
- [BC09] Joël Bellaïche and Gaëtan Chenevier, *Families of Galois representations and Selmer groups*, Astérisque (2009), no. 324, xii+314.
- [BCDT01] Christophe Breuil, Brian Conrad, Fred Diamond, and Richard Taylor, *On the modularity of elliptic curves over \mathbf{Q}* , J.A.M.S. **14** (2001), 843–939.
- [BD21] Christophe Breuil and Yiwen Ding, *Bernstein eigenvarieties*, 2021.
- [BDJ10] Kevin Buzzard, Fred Diamond, and Frazer Jarvis, *On Serre’s conjecture for mod ℓ Galois representations over totally real fields*, Duke Math. J. **155** (2010), no. 1, 105–161.
- [Ber08a] Laurent Berger, *Construction de (ϕ, Γ) -modules: représentations p -adiques et B -paires*, Algebra Number Theory **2** (2008), no. 1, 91–120.
- [Ber08b] ———, *Équations différentielles p -adiques et (ϕ, N) -modules filtrés*, no. 319, 2008, Représentations p -adiques de groupes p -adiques. I. Représentations galoisiennes et (ϕ, Γ) -modules, pp. 13–38.
- [Ber10] ———, *Représentations modulaires de $GL_2(\mathbf{Q}_p)$ et représentations galoisiennes de dimension 2*, Astérisque (2010), no. 330, 263–279.
- [Ber12] ———, *Central characters for smooth irreducible modular representations of $GL_2(\mathbf{f}\mathbf{Q}_p)$* , Rend. Semin. Mat. Univ. Padova **128** (2012), 1–6 (2013).
- [Bez16] Roman Bezrukavnikov, *On two geometric realizations of an affine Hecke algebra*, Publ. Math. Inst. Hautes Études Sci. **123** (2016), 1–67.
- [BF19] Alexander Braverman and Michael Finkelberg, *A quasi-coherent description of the category $D\text{-mod}(gr_{GL(n)})$* , 2019.
- [BG02] A. Braverman and D. Gaitsgory, *Geometric Eisenstein series*, Invent. Math. **150** (2002), no. 2, 287–384.
- [BG14] Kevin Buzzard and Toby Gee, *The conjectural connections between automorphic representations and Galois representations*. Vol. 1, London Math. Soc. Lecture Note Ser., vol. 414, Cambridge Univ. Press, Cambridge, 2014, pp. 135–187.
- [BG19] Rebecca Bellovin and Toby Gee, *G -valued local deformation rings and global lifts*, Algebra Number Theory **13** (2019), no. 2, 333–378.
- [BH15] Christophe Breuil and Florian Herzig, *Ordinary representations of $G(\mathbf{Q}_p)$ and fundamental algebraic representations*, Duke Math. J. **164** (2015), no. 7, 1271–1352.
- [BHH⁺20] Christophe Breuil, Florian Herzig, Yongquan Hu, Stefano Morra, and Benjamin Schraen, *Gelfand–kirillov dimension and mod p cohomology for GL_2* , 2020.
- [BHH⁺21] ———, *Conjectures and results on modular representations of $GL_n(k)$ for a p -adic field k* , 2021.
- [BHS17a] Christophe Breuil, Eugen Hellmann, and Benjamin Schraen, *Smoothness and classicality on eigenvarieties*, Invent. Math. **209** (2017), no. 1, 197–274.
- [BHS17b] ———, *Une interprétation modulaire de la variété trianguline*, Math. Ann. **367** (2017), no. 3-4, 1587–1645.
- [BHS19] ———, *A local model for the trianguline variety and applications*, Publ. Math. Inst. Hautes Études Sci. **130** (2019), 299–412.
- [BIP21] Gebhard Böckle, Ashwin Iyengar, and Vytautas Paškūnas, *On local Galois deformation rings*, 2021.
- [BIP22] Gebhard Böckle, Ashwin Iyengar, and Vytautas Paškūnas, *Zariski density of crystalline points*, 2022.
- [BK] Roman Bezrukavnikov and Tina Kanstrup, *canonical bases and representation categories*.
- [BL94] L. Barthel and R. Livné, *Irreducible modular representations of \mathfrak{gl}_2 of a local field*, Duke Math J. **75** (1994), no. 2, 261–292.

- [BL95] Arnaud Beauville and Yves Laszlo, *Un lemme de descente*, C. R. Acad. Sci. Paris Sér. I Math. **320** (1995), no. 3, 335–340.
- [BM02] Christophe Breuil and Ariane Mézard, *Multiplicités modulaires et représentations de $GL_2(\mathbf{Z}_p)$ et de $\text{Gal}(\overline{\mathbf{Q}_p}/\mathbf{Q}_p)$ en $l = p$* , Duke Math. J. **115** (2002), no. 2, 205–310, With an appendix by Guy Henniart.
- [BM14] Christophe Breuil and Ariane Mézard, *Multiplicités modulaires raffinées*, Bull. Soc. Math. France **142** (2014), no. 1, 127–175.
- [Böc01] Gebhard Böckle, *On the density of modular points in universal deformation spaces*, Amer. J. Math. **123** (2001), no. 5, 985–1007.
- [BP12] Christophe Breuil and Vytautas Paškūnas, *Towards a modulo p Langlands correspondence for GL_2* , Mem. Amer. Math. Soc. **216** (2012), no. 1016, vi+114.
- [BP20] George Boxer and Vincent Pilloni, *Higher Hida and Coleman theories on the modular curve*, 2020.
- [BP21] ———, *Higher Coleman theory*, 2021.
- [Bre03a] Christophe Breuil, *Sur quelques représentations modulaires et p -adiques de $GL_2(\mathbf{Q}_p)$. I*, Compositio Math. **138** (2003), no. 2, 165–188.
- [Bre03b] ———, *Sur quelques représentations modulaires et p -adiques de $GL_2(\mathbf{Q}_p)$. II*, J. Inst. Math. Jussieu **2** (2003), no. 1, 23–58.
- [Bre08] ———, *Introduction générale*, Astérisque (2008), no. 319, 1–12, Représentations p -adiques de groupes p -adiques. I. Représentations galoisiennes et (ϕ, Γ) -modules.
- [Bre10a] ———, *The emerging p -adic Langlands programme*, Proceedings of the International Congress of Mathematicians. Volume II, Hindustan Book Agency, New Delhi, 2010, pp. 203–230.
- [Bre10b] ———, *Série spéciale p -adique et cohomologie étale complétée*, Astérisque (2010), no. 331, 65–115.
- [Bre11] ———, *Diagrammes de Diamond et (ϕ, Γ) -modules*, Israel J. Math. **182** (2011), 349–382.
- [Bre14] ———, *Sur un problème de compatibilité local-global modulo p pour GL_2* , J. Reine Angew. Math. **692** (2014), 1–76.
- [Bre15] ———, *Vers le socle localement analytique pour GL_n II*, Math. Ann. **361** (2015), no. 3–4, 741–785.
- [BS07] Christophe Breuil and Peter Schneider, *First steps towards p -adic Langlands functoriality*, J. Reine Angew. Math. **610** (2007), 149–180.
- [BS21] Bhargav Bhatt and Peter Scholze, *Prismatic F -crystals and crystalline Galois representations*, 2021.
- [BV13] Nicolas Bergeron and Akshay Venkatesh, *The asymptotic growth of torsion homology for arithmetic groups*, J. Inst. Math. Jussieu **12** (2013), no. 2, 391–447.
- [BZ77] I. N. Bernstein and A. V. Zelevinsky, *Induced representations of reductive p -adic groups. I*, Ann. Sci. École Norm. Sup. (4) **10** (1977), no. 4, 441–472.
- [BZCHN20] David Ben-Zvi, Harrison Chen, David Helm, and David Nadler, *Coherent Springer theory and the categorical Deligne–Langlands correspondence*, 2020.
- [BZN18] David Ben-Zvi and David Nadler, *Betti geometric Langlands*, Algebraic geometry: Salt Lake City 2015, Proc. Sympos. Pure Math., vol. 97, Amer. Math. Soc., Providence, RI, 2018, pp. 3–41.
- [BZNP17] David Ben-Zvi, David Nadler, and Anatoly Preygel, *Integral transforms for coherent sheaves*, J. Eur. Math. Soc. (JEMS) **19** (2017), no. 12, 3763–3812.
- [Cal21] Frank Calegari, *Reciprocity in the Langlands program since Fermat’s Last Theorem*, 2021.
- [Cam] Campbell.
- [Car83] Henri Carayol, *Sur les représentations l -adiques attachées aux formes modulaires de Hilbert*, C. R. Acad. Sci. Paris Sér. I Math. **296** (1983), no. 15, 629–632.
- [CC98] F. Cherbonnier and P. Colmez, *Représentations p -adiques surconvergentes*, Invent. Math. **133** (1998), no. 3, 581–611.
- [CDN20] Pierre Colmez, Gabriel Dospinescu, and Wiesław Nizioł, *Cohomologie p -adique de la tour de Drinfeld: le cas de la dimension 1*, J. Amer. Math. Soc. **33** (2020), no. 2, 311–362.

- [CDN21] ———, *Integral p -adic étale cohomology of Drinfeld symmetric spaces*, Duke Math. J. **170** (2021), no. 3, 575–613.
- [CDN22] Pierre Colmez, Gabriel Dospinescu, and Wiesława Nizioł, *Factorization de la cohomologie étale p -adique de la tour de Drinfeld*, 2022.
- [CDP14] Pierre Colmez, Gabriel Dospinescu, and Vytautas Paškūnas, *The p -adic local Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$* , Camb. J. Math. **2** (2014), no. 1, 1–47.
- [CDT99] Brian Conrad, Fred Diamond, and Richard Taylor, *Modularity of certain potentially Barsotti-Tate Galois representations*, J.A.M.S. **12** (1999), 521–567.
- [CE12] Frank Calegari and Matthew Emerton, *Completed cohomology—a survey*, Non-abelian fundamental groups and Iwasawa theory, London Math. Soc. Lecture Note Ser., vol. 393, Cambridge Univ. Press, Cambridge, 2012, pp. 239–257.
- [CEG⁺16] Ana Caraiani, Matthew Emerton, Toby Gee, David Geraghty, Vytautas Paškūnas, and Sug Woo Shin, *Patching and the p -adic local Langlands correspondence*, Camb. J. Math. **4** (2016), no. 2, 197–287.
- [CEG⁺18] ———, *Patching and the p -adic Langlands program for $\mathrm{GL}_2(\mathbf{Q}_p)$* , Compos. Math. **154** (2018), no. 3, 503–548.
- [Che04] Gaëtan Chenevier, *Familles p -adiques de formes automorphes pour GL_n* , J. Reine Angew. Math. **570** (2004), 143–217.
- [Che13] ———, *Sur la densité des représentations cristallines de $\mathrm{Gal}(\overline{\mathbf{Q}}_p/\mathbf{Q}_p)$* , Math. Ann. **355** (2013), no. 4, 1469–1525.
- [Che14] ———, *The p -adic analytic space of pseudocharacters of a profinite group and pseudorepresentations over arbitrary rings*, Automorphic forms and Galois representations. Vol. 1, London Math. Soc. Lecture Note Ser., vol. 414, Cambridge Univ. Press, Cambridge, 2014, pp. 221–285.
- [Chi90] Bruno Chiarellotto, *Duality in rigid analysis, p -adic analysis* (Trento, 1989), Lecture Notes in Math., vol. 1454, Springer, Berlin, 1990, pp. 142–172.
- [Col08] Pierre Colmez, *Représentations triangulines de dimension 2*, no. 319, 2008, Représentations p -adiques de groupes p -adiques. I. Représentations galoisiennes et (ϕ, Γ) -modules, pp. 213–258.
- [Col10a] ———, *La série principale unitaire de $\mathrm{GL}_2(\mathbf{Q}_p)$* , Astérisque (2010), no. 330, 213–262.
- [Col10b] ———, *(ϕ, Γ) -modules et représentations du mirabolique de $\mathrm{GL}_2(\mathbf{Q}_p)$* , Astérisque (2010), no. 330, 61–153.
- [Col10c] ———, *Représentations de $\mathrm{GL}_2(\mathbf{Q}_p)$ et (ϕ, Γ) -modules*, Astérisque (2010), no. 330, 281–509.
- [Col14] ———, *La série principale unitaire de $\mathrm{GL}_2(\mathbf{Q}_p)$: vecteurs localement analytiques*, Automorphic forms and Galois representations. Vol. 1, London Math. Soc. Lecture Note Ser., vol. 414, Cambridge Univ. Press, Cambridge, 2014, pp. 286–358.
- [Col16] ———, *Représentations localement analytiques de $\mathbf{fGL}_2(\mathbf{fQ}_p)$ et (φ, Γ) -modules*, Represent. Theory **20** (2016), 187–248.
- [Col19] ———, *Correspondance de Langlands locale p -adique et changement de poids*, J. Eur. Math. Soc. (JEMS) **21** (2019), no. 3, 797–838.
- [Con06] Brian Conrad, *Relative ampleness in rigid geometry*, Ann. Inst. Fourier (Grenoble) **56** (2006), no. 4, 1049–1126.
- [CS17] Ana Caraiani and Peter Scholze, *On the generic part of the cohomology of compact unitary Shimura varieties*, Ann. of Math. (2) **186** (2017), no. 3, 649–766.
- [CS21] Kestutis Cesnavicius and Peter Scholze, *Purity for flat cohomology*, 2021.
- [CT09] Brian Conrad and Michael Temkin, *Non-Archimedean analytification of algebraic spaces*, J. Algebraic Geom. **18** (2009), no. 4, 731–788.
- [DDT97] Henri Darmon, Fred Diamond, and Richard Taylor, *Fermat’s last theorem, Elliptic curves, modular forms & Fermat’s last theorem* (Hong Kong, 1993), Int. Press, Cambridge, MA, 1997, pp. 2–140.
- [DEG^a] Andrea Dotto, Matthew Emerton, and Toby Gee, *A categorical p -adic Langlands correspondence for $\mathrm{GL}_2(\mathbf{Q}_p)$* .
- [DEG^b] ———, *Localization of smooth p -power torsion representations of $\mathrm{GL}_2(\mathbf{Q}_p)$* .
- [Del73] P. Deligne, *Letter to Piatetski-Shapiro*, 1973.

- [DHKM20] Jean-François Dat, David Helm, Robert Kurinczuk, and Gilbert Moss, *Moduli of Langlands Parameters*, 2020.
- [Dia97] Fred Diamond, *The Taylor-Wiles construction and multiplicity one*, Invent. Math. **128** (1997), no. 2, 379–391.
- [DL21a] Andrea Dotto and Daniel Le, *Diagrams in the mod p cohomology of Shimura curves*, Compos. Math. **157** (2021), no. 8, 1653–1723.
- [DL21b] Heng Du and Tong Liu, *A prismatic approach to (φ, \hat{G}) -modules and F -crystals*, 2021.
- [DLB17] Gabriel Dospinescu and Arthur-César Le Bras, *Revêtements du demi-plan de Drinfeld et correspondance de Langlands p -adique*, Ann. of Math. (2) **186** (2017), no. 2, 321–411.
- [DPS20] Gabriel Dospinescu, Vytautas Paškūnas, and Benjamin Schraen, *Infinitesimal characters in arithmetic families*, 2020.
- [Dri06] Vladimir Drinfeld, *Infinite-dimensional vector bundles in algebraic geometry: an introduction*, The unity of mathematics, Progr. Math., vol. 244, Birkhäuser Boston, Boston, MA, 2006, pp. 263–304.
- [DS15] Della Dumbaugh and Joachim Schwermer, *Emil Artin and beyond—class field theory and L -functions*, Heritage of European Mathematics, European Mathematical Society (EMS), Zürich, 2015, With contributions by James Cogdell and Robert Langlands.
- [DS22] Fred Diamond and Shu Sasaki, *A serre weight conjecture for geometric Hilbert modular forms in characteristic $\neq p$* , Journal of the European Mathematical Society (2022).
- [DT94] Fred Diamond and Richard Taylor, *Lifting modular mod l representations*, Duke Math. J. **74** (1994), no. 2, 253–269.
- [EG14] Matthew Emerton and Toby Gee, *A geometric perspective on the Breuil-Mézard conjecture*, J. Inst. Math. Jussieu **13** (2014), no. 1, 183–223.
- [EG20] Matthew Emerton and Toby Gee, *Moduli stacks of (φ, Γ) -modules: a survey*, 2020.
- [EG22] Matthew Emerton and Toby Gee, *Moduli stacks of étale (φ, Γ) -modules and the existence of crystalline lifts*, Annals of Math Studies (to appear), 2022.
- [EGS15] Matthew Emerton, Toby Gee, and David Savitt, *Lattices in the cohomology of Shimura curves*, Invent. Math. **200** (2015), no. 1, 1–96.
- [EGZ] Matthew Emerton, Toby Gee, and Xinwen Zhu, In progress.
- [EH14] Matthew Emerton and David Helm, *The local Langlands correspondence for GL_n in families*, Ann. Sci. Éc. Norm. Supér. (4) **47** (2014), no. 4, 655–722.
- [Eme] Matthew Emerton, *Formal algebraic stacks*, In preparation.
- [Eme06a] ———, *Jacquet modules of locally analytic representations of p -adic reductive groups. I. Construction and first properties*, Ann. Sci. École Norm. Sup. (4) **39** (2006), no. 5, 775–839.
- [Eme06b] ———, *On the interpolation of systems of eigenvalues attached to automorphic Hecke eigenforms*, Invent. Math. **164** (2006), no. 1, 1–84.
- [Eme10] ———, *Ordinary parts of admissible representations of p -adic reductive groups I. Definition and first properties*, Astérisque (2010), no. 331, 355–402.
- [Eme11a] ———, *Local-global compatibility in the p -adic Langlands programme for GL_2/\mathbf{Q}* , 2011.
- [Eme11b] ———, *p -adic families of modular forms (after Hida, Coleman, and Mazur)*, no. 339, 2011, Séminaire Bourbaki. Vol. 2009/2010. Exposés 1012–1026, pp. Exp. No. 1013, vii, 31–61.
- [Eme14] ———, *Completed cohomology and the p -adic Langlands program*, Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. II, Kyung Moon Sa, Seoul, 2014, pp. 319–342.
- [EP20] Matthew Emerton and Vytautas Paškūnas, *On the density of supercuspidal points of fixed regular weight in local deformation rings and global Hecke algebras*, J. Éc. polytech. Math. **7** (2020), 337–371.
- [EZ] Matthew Emerton and Xinwen Zhu, In progress.
- [Fen19] Tony Feng, *The spectral Hecke algebra*, 2019.

- [FF18] Laurent Fargues and Jean-Marc Fontaine, *Courbes et fibrés vectoriels en théorie de Hodge p -adique*, Astérisque (2018), no. 406, xiii+382, With a preface by Pierre Colmez.
- [FM95] Jean-Marc Fontaine and Barry Mazur, *Geometric Galois representations*, Elliptic curves, modular forms, & Fermat’s last theorem (Hong Kong, 1993), Ser. Number Theory, I, Int. Press, Cambridge, MA, 1995, pp. 41–78.
- [Fon90] Jean-Marc Fontaine, *Représentations p -adiques des corps locaux. I*, The Grothendieck Festschrift, Vol. II, Progr. Math., vol. 87, Birkhäuser Boston, Boston, MA, 1990, pp. 249–309.
- [Fon94] ———, *Représentations l -adiques potentiellement semi-stables*, no. 223, 1994, Périodes p -adiques (Bures-sur-Yvette, 1988), pp. 321–347.
- [Fon04] ———, *Arithmétique des représentations galoisiennes p -adiques*, no. 295, 2004, Cohomologies p -adiques et applications arithmétiques. III, pp. xi, 1–115.
- [FS21] Laurent Fargues and Peter Scholze, *Geometrization of the local Langlands correspondence*.
- [FS22] ———, *The Langlands program and the moduli of bundles on the curve*.
- [Fu21] Weibo Fu, *A derived construction of eigenvarieties*, 2021.
- [Fuj06] K. Fujiwara, *Deformation rings and Hecke algebras in the totally real case*, ArXiv Mathematics e-prints (2006).
- [Gai] Dennis Gaitsgory, *Notes on Geometric Langlands: Quasi-coherent sheaves on stacks*.
- [Gai12] Dennis Gaitsgory, *Ind-coherent sheaves*, 2012.
- [Gai15] Dennis Gaitsgory, *Outline of the proof of the geometric Langlands conjecture for GL_2* , Astérisque (2015), no. 370, 1–112.
- [Gai17] ———, *Progrès récents dans la théorie de Langlands géométrique*, no. 390, 2017, Séminaire Bourbaki. Vol. 2015/2016. Exposés 1104–1119, pp. Exp. No. 1109, 139–168.
- [Gee11] Toby Gee, *On the weights of mod p Hilbert modular forms*, Invent. Math. **184** (2011), no. 1, 1–46.
- [GG15] Toby Gee and David Geraghty, *The Breuil–Mézard conjecture for quaternion algebras*, Ann. Inst. Fourier (Grenoble) **65** (2015), no. 4, 1557–1575.
- [GK14] Toby Gee and Mark Kisin, *The Breuil–Mézard conjecture for potentially Barsotti–Tate representations*, Forum Math. Pi **2** (2014), e1 (56 pages).
- [GKK⁺22] Anthony Guzman, Kalyani Kansal, Iason Kountouridis, Ben Savoie, and Xiyuan Wang, *Smoothness of components of the Emerton–Gee stack for GL_2* , 2022.
- [GLS12] Toby Gee, Tong Liu, and David Savitt, *The Buzzard–Diamond–Jarvis conjecture for unitary groups*, Preprint, 2012.
- [GM98] Fernando Q. Gouvêa and Barry Mazur, *On the density of modular representations*, Computational perspectives on number theory (Chicago, IL, 1995), AMS/IP Stud. Adv. Math., vol. 7, Amer. Math. Soc., Providence, RI, 1998, pp. 127–142.
- [GN22] Toby Gee and James Newton, *Patching and the completed homology of locally symmetric spaces*, J. Inst. Math. Jussieu **21** (2022), no. 2, 395–458.
- [GR11] Dennis Gaitsgory and Nick Rozenblyum, *Dg indschemes*, 2011.
- [GRJ18] Ildar Gaisin and Joaquín Rodríguez Jacinto, *Arithmetic families of (ϕ, Γ) -modules and locally analytic representations of $GL_2(\mathbb{Q}_p)$* , Doc. Math. **23** (2018), 1313–1404.
- [GS11] Toby Gee and David Savitt, *Serre weights for quaternion algebras*, Compos. Math. **147** (2011), no. 4, 1059–1086.
- [GV18] S. Galatius and A. Venkatesh, *Derived Galois deformation rings*, Adv. Math. **327** (2018), 470–623.
- [Han17] David Hansen, *Universal eigenvarieties, trianguline Galois representations, and p -adic Langlands functoriality*, J. Reine Angew. Math. **730** (2017), 1–64, With an appendix by James Newton.
- [Har16] Michael Harris, *Speculations on the mod p representation theory of p -adic groups*, Ann. Fac. Sci. Toulouse Math. (6) **25** (2016), no. 2-3, 403–418.
- [Hel13] Eugen Hellmann, *On arithmetic families of filtered ϕ -modules and crystalline representations*, J. Inst. Math. Jussieu **12** (2013), no. 4, 677–726.

- [Hel16] David Helm, *Whittaker models and the integral Bernstein center for GL_n* , Duke Math. J. **165** (2016), no. 9, 1597–1628.
- [Hel20] Eugen Hellmann, *On the derived category of the Iwahori–Hecke algebra*, 2020.
- [Hen00] Guy Henniart, *Une preuve simple des conjectures de Langlands pour $GL(n)$ sur un corps p -adique*, Invent. Math. **139** (2000), no. 2, 439–455.
- [Her97] Ivo Herzog, *The Ziegler spectrum of a locally coherent Grothendieck category*, Proc. London Math. Soc. (3) **74** (1997), no. 3, 503–558.
- [Her09] Florian Herzig, *The weight in a Serre-type conjecture for tame n -dimensional Galois representations*, Duke Math. J. **149** (2009), no. 1, 37–116.
- [HH20] Urs Hartl and Eugen Hellmann, *The universal family of semistable p -adic Galois representations*, Algebra Number Theory **14** (2020), no. 5, 1055–1121.
- [HHS] Eugen Hellmann, Valentin Hernandez, and Benjamin Schraen, In progress.
- [HM18] David Helm and Gilbert Moss, *Converse theorems and the local Langlands correspondence in families*, Invent. Math. **214** (2018), no. 2, 999–1022.
- [HMM22] Bao V. Le Hung, Ariane Mézard, and Stefano Morra, *Local model theory for non-generic tame potentially Barsotti–Tate deformation rings (in preparation)*.
- [HMS18] Eugen Hellmann, Christophe M. Margerin, and Benjamin Schraen, *Density of automorphic points in deformation rings of polarized global Galois representations*, 2018.
- [HT01] Michael Harris and Richard Taylor, *The geometry and cohomology of some simple Shimura varieties*, Annals of Mathematics Studies, vol. 151, Princeton University Press, Princeton, NJ, 2001, With an appendix by Vladimir G. Berkovich.
- [HW20] Yongquan Hu and Haoran Wang, *On the mod p cohomology for GL_2 : the non-semisimple case*, 2020.
- [HW22] ———, *On the mod p cohomology for GL_2 : the non-semisimple case*, Camb. J. Math. **10** (2022), no. 2, 261–431.
- [HZ] Tamir Hemo and Xinwen Zhu, *Unipotent categorical local Langlands correspondence*, In progress.
- [JC21] Joaquín Rodríguez Jacinto and Juan Esteban Rodríguez Camargo, *Solid locally analytic representations of p -adic Lie groups*, 2021.
- [Kat77] Nicholas M. Katz, *A result on modular forms in characteristic p* , Modular functions of one variable, V (Proc. Second Internat. Conf., Univ. Bonn, Bonn, 1976), 1977, pp. 53–61. Lecture Notes in Math., Vol. 601.
- [Ked08] Kiran S. Kedlaya, *Slope filtrations for relative Frobenius*, no. 319, 2008, Représentations p -adiques de groupes p -adiques. I. Représentations galoisiennes et (ϕ, Γ) -modules, pp. 259–301.
- [Ked16] ———, *Noetherian properties of Fargues-Fontaine curves*, Int. Math. Res. Not. IMRN (2016), no. 8, 2544–2567.
- [Kis03] Mark Kisin, *Overconvergent modular forms and the Fontaine-Mazur conjecture*, Invent. Math. **153** (2003), no. 2, 373–454.
- [Kis06] ———, *Crystalline representations and F -crystals*, Algebraic geometry and number theory, Progr. Math., vol. 253, Birkhäuser Boston, Boston, MA, 2006, pp. 459–496.
- [Kis07a] ———, *Modularity for some geometric Galois representations, L -functions and Galois representations*, London Math. Soc. Lecture Note Ser., vol. 320, Cambridge Univ. Press, Cambridge, 2007, With an appendix by Ofer Gabber, pp. 438–470.
- [Kis07b] ———, *Modularity of 2-dimensional Galois representations*, Current developments in mathematics, 2005, Int. Press, Somerville, MA, 2007, pp. 191–230.
- [Kis09a] ———, *The Fontaine-Mazur conjecture for GL_2* , J. Amer. Math. Soc. **22** (2009), no. 3, 641–690.
- [Kis09b] ———, *Modularity of 2-adic Barsotti-Tate representations*, Invent. Math. **178** (2009), no. 3, 587–634.
- [Kis09c] ———, *Moduli of finite flat group schemes, and modularity*, Ann. of Math. (2) **170** (2009), no. 3, 1085–1180.
- [Kis10a] ———, *Deformations of $G_{\mathbf{Q}_p}$ and $GL_2(\mathbf{Q}_p)$ representations*, Astérisque (2010), no. 330, 511–528.

- [Kis10b] ———, *The structure of potentially semi-stable deformation rings*, Proceedings of the International Congress of Mathematicians. Volume II, Hindustan Book Agency, New Delhi, 2010, pp. 294–311.
- [Koh17] Jan Kohlhaase, *Smooth duality in natural characteristic*, Adv. Math. **317** (2017), 1–49.
- [KPX14] Kiran S. Kedlaya, Jonathan Pottharst, and Liang Xiao, *Cohomology of arithmetic families of (φ, Γ) -modules*, J. Amer. Math. Soc. **27** (2014), no. 4, 1043–1115.
- [KS12] Jan Kohlhaase and Benjamin Schraen, *Homological vanishing theorems for locally analytic representations*, Math. Ann. **353** (2012), no. 1, 219–258.
- [KW09] Chandrashekhara Khare and Jean-Pierre Wintenberger, *Serre’s modularity conjecture. I*, Invent. Math. **178** (2009), no. 3, 485–504.
- [Laf18] Vincent Lafforgue, *Chtoucas pour les groupes réductifs et paramétrisation de Langlands globale*, J. Amer. Math. Soc. **31** (2018), no. 3, 719–891.
- [Lan70] R. P. Langlands, *Problems in the theory of automorphic forms*, Lectures in modern analysis and applications, III, 1970, pp. 18–61. Lecture Notes in Math., Vol. 170.
- [Lan71] Robert P. Langlands, *Euler products*, Yale Mathematical Monographs, vol. 1, Yale University Press, New Haven, Conn.-London, 1971, A James K. Whittemore Lecture in Mathematics given at Yale University, 1967.
- [Lan73] R. P. Langlands, *Modular forms and ℓ -adic representations*, Modular functions of one variable, II (Proc. Internat. Summer School, Univ. Antwerp, Antwerp, 1972), Springer, Berlin, 1973, pp. 361–500. Lecture Notes in Math., Vol. 349.
- [Lan79] ———, *Automorphic representations, Shimura varieties, and motives. Ein Märchen*, Automorphic forms, representations and L -functions (Proc. Sympos. Pure Math., Oregon State Univ., Corvallis, Ore., 1977), Part 2, Proc. Sympos. Pure Math., XXXIII, Amer. Math. Soc., Providence, R.I., 1979, pp. 205–246.
- [Lan80] Robert P. Langlands, *Base Change for $GL(2)$* , Annals of Math. Studies **96** (1980).
- [Laz65] Michel Lazard, *Groupes analytiques p -adiques*, Inst. Hautes Études Sci. Publ. Math. (1965), no. 26, 389–603.
- [Le19] Daniel Le, *On some nonadmissible smooth irreducible representations for GL_2* , Math. Res. Lett. **26** (2019), no. 6, 1747–1758.
- [LHL16] Bao V. Le Hung and Chao Li, *Level raising mod 2 and arbitrary 2-Selmer ranks*, Compos. Math. **152** (2016), no. 8, 1576–1608.
- [Liu15] R. Liu, *Triangulation of refined families*, Comment. Math. Helv. **90** (2015), no. 4, 831–904.
- [Loe11] David Loeffler, *Overconvergent algebraic automorphic forms*, Proc. Lond. Math. Soc. (3) **102** (2011), no. 2, 193–228.
- [Lur09a] Jacob Lurie, *Higher topos theory*, Annals of Mathematics Studies, vol. 170, Princeton University Press, Princeton, NJ, 2009.
- [Lur09b] Jacob Lurie, *Stable infinity categories*, 2009.
- [Lur17] Jacob Lurie, *Higher algebra (version of september 18, 2017)*, 2017.
- [Man22] Lucas Mann, *A p -adic 6-functor formalism in rigid-analytic geometry*, 2022.
- [Maz89] Barry Mazur, *Deforming Galois representations*, Galois groups over $\mathbf{f}Q$ (Berkeley, CA, 1987), Math. Sci. Res. Inst. Publ., vol. 16, Springer, New York, 1989, pp. 385–437.
- [MW84] B. Mazur and A. Wiles, *Class fields of abelian extensions of $\mathbf{f}Q$* , Invent. Math. **76** (1984), no. 2, 179–330.
- [Nak14] Kentaro Nakamura, *Zariski density of crystalline representations for any p -adic field*, J. Math. Sci. Univ. Tokyo **21** (2014), no. 1, 79–127.
- [Nee11] Amnon Neeman, *Non-left-complete derived categories*, Math. Res. Lett. **18** (2011), no. 5, 827–832.
- [NT19] James Newton and Jack A. Thorne, *Adjoint selmer groups of automorphic Galois representations of unitary type*, 2019.
- [OS15] Sascha Orlik and Matthias Strauch, *On Jordan-Hölder series of some locally analytic representations*, J. Amer. Math. Soc. **28** (2015), no. 1, 99–157.

- [Pan17] Lue Pan, *First covering of the Drinfel'd upper half-plane and Banach representations of $GL_2(\mathbf{Q}_p)$* , Algebra Number Theory **11** (2017), no. 2, 405–503.
- [Pan22a] ———, *The Fontaine-Mazur conjecture in the residually reducible case*, J. Amer. Math. Soc. **35** (2022), no. 4, 1031–1169.
- [Pan22b] ———, *On locally analytic vectors of the completed cohomology of modular curves*, Forum Math. Pi **10** (2022), Paper No. e7, 82.
- [Paš10] Vytautas Paškūnas, *Admissible unitary completions of locally \mathbf{Q}_p -rational representations of $GL_2(F)$* , Represent. Theory **14** (2010), 324–354.
- [Paš13] ———, *The image of Colmez's Montreal functor*, Publ. Math. Inst. Hautes Études Sci. **118** (2013), 1–191.
- [Paš15] Vytautas Paškūnas, *On the Breuil–Mézard conjecture*, Duke Math. J. **164** (2015), no. 2, 297–359.
- [Pha22] Dat Pham, *The Emerton–Gee stack of rank one (φ, Γ) -modules*, 2022.
- [Pil20] V. Pilloni, *Higher coherent cohomology and p -adic modular forms of singular weights*, Duke Math. J. **169** (2020), no. 9, 1647–1807.
- [Por22] Gal Porat, *Overconvergence of étale (ϕ, Γ) -modules in families*, 2022.
- [PR09] G. Pappas and M. Rapoport, *Φ -modules and coefficient spaces*, Mosc. Math. J. **9** (2009), no. 3, 625–663, back matter.
- [Roz15] Sandra Rozensztajn, *Potentially semi-stable deformation rings for discrete series extended types*, J. Éc. polytech. Math. **2** (2015), 179–211.
- [RR73] A. P. Robertson and Wendy Robertson, *Topological vector spaces*, second ed., Cambridge Tracts in Mathematics and Mathematical Physics, No. 53, Cambridge University Press, London-New York, 1973.
- [San14] Fabian Sander, *Hilbert–Samuel multiplicities of certain deformation rings*, Math. Res. Lett. **21** (2014), no. 3, 605–615.
- [Sav05] David Savitt, *On a conjecture of Conrad, Diamond, and Taylor*, Duke Math. J. **128** (2005), no. 1, 141–197.
- [Sch] Peter Scholze.
- [Sch13] ———, *The local Langlands correspondence for GL_n over p -adic fields*, Invent. Math. **192** (2013), no. 3, 663–715.
- [Sch15a] Peter Schneider, *Smooth representations and Hecke modules in characteristic p* , Pacific J. Math. **279** (2015), no. 1–2, 447–464.
- [Sch15b] Peter Scholze, *On the p -adic cohomology of the Lubin–Tate tower*, ArXiv e-prints (2015).
- [Sch15c] Benjamin Schraen, *Sur la présentation des représentations supersingulières de $GL_2(F)$* , J. Reine Angew. Math. **704** (2015), 187–208.
- [SD73] H. P. F. Swinnerton-Dyer, *On l -adic representations and congruences for coefficients of modular forms*, Modular functions of one variable, III (Proc. Internat. Summer School, Univ. Antwerp, 1972), 1973, pp. 1–55. Lecture Notes in Math., Vol. 350.
- [Ser73] Jean-Pierre Serre, *Congruences et formes modulaires [d'après H. P. F. Swinnerton-Dyer]*, Séminaire Bourbaki, 24e année (1971/1972), Exp. No. 416, 1973, pp. 319–338. Lecture Notes in Math., Vol. 317.
- [Ser87] ———, *Sur les représentations modulaires de degré 2 de $\text{Gal}(\overline{\mathbf{Q}}/\mathbf{Q})$* , Duke Math. J. **54** (1987), 179–230.
- [Sho20] Jack Shotton, *The category of finitely presented smooth mod p representations of $GL_2(F)$* , Doc. Math. **25** (2020), 143–157.
- [ST02] Peter Schneider and Jeremy Teitelbaum, *Locally analytic distributions and p -adic representation theory, with applications to GL_2* , J. Amer. Math. Soc. **15** (2002), no. 2, 443–468.
- [ST03] ———, *Algebras of p -adic distributions and admissible representations*, Invent. Math. **153** (2003), no. 1, 145–196.
- [Sta] The Stacks Project Authors, *Stacks Project*, <http://stacks.math.columbia.edu>.
- [SW97] Chris Skinner and Andrew Wiles, *Ordinary representations and modular forms*, Proc. Nat. Acad. Sci. U.S.A. **94** (1997), no. 20, 10520–10527.
- [SW99] C. M. Skinner and A. J. Wiles, *Residually reducible representations and modular forms*, Inst. Hautes Études Sci. Publ. Math. (1999), no. 89, 5–126 (2000).

- [SZ99] P. Schneider and E.-W. Zink, *K-types for the tempered components of a p -adic general linear group*, J. Reine Angew. Math. **517** (1999), 161–208, With an appendix by Schneider and U. Stuhler.
- [Tay08] Richard Taylor, *Automorphy for some l -adic lifts of automorphic mod l Galois representations. II*, Pub. Math. IHES **108** (2008), 183–239.
- [Tim21] James Timmins, *Coherence of augmented Iwasawa algebras*, 2021.
- [Tun81] Jerrold Tunnell, *Artin’s Conjecture for representations of octahedral type*, Bull. AMS **5** (1981), 173–175.
- [TW95] Richard Taylor and Andrew Wiles, *Ring-theoretic properties of certain Hecke algebras*, Annals of Mathematics **142** (1995), 553–572.
- [Ven19] Akshay Venkatesh, *Derived Hecke algebra and cohomology of arithmetic groups*, Forum Math. Pi **7** (2019), e7, 119.
- [Vig89] Marie-France Vignéras, *Représentations modulaires de $GL(2, F)$ en caractéristique l , F corps p -adique, $p \neq l$* , Compositio Math. **72** (1989), no. 1, 33–66.
- [Wan22] Yitong Wang, *On the mod p cohomology for GL_2* , 2022.
- [War17] Evan B. Warner, *Adic Moduli Spaces*, ProQuest LLC, Ann Arbor, MI, 2017, Thesis (Ph.D.)–Stanford University.
- [WE18] Carl Wang-Erickson, *Algebraic families of Galois representations and potentially semi-stable pseudodeformation rings*, Math. Ann. **371** (2018), no. 3-4, 1615–1681.
- [Wil88] A. Wiles, *On ordinary λ -adic representations associated to modular forms*, Invent. Math. **94** (1988), no. 3, 529–573.
- [Wil95] Andrew Wiles, *Modular elliptic curves and Fermat’s last theorem*, Annals of Mathematics **142** (1995), 443–551.
- [Wu21a] Zhixiang Wu, *Local models for the trianguline variety and partially classical families*, 2021.
- [Wu21b] ———, *A note on presentations of supersingular representations of $GL_2(F)$* , Manuscripta Math. **165** (2021), no. 3-4, 583–596.
- [Wu21c] Zhiyou Wu, *Galois representations, (φ, Γ) -modules and prismatic F -crystals*, 2021.
- [WWE17] Preston Wake and Carl Wang-Erickson, *Ordinary pseudorepresentations and modular forms*, Proc. Amer. Math. Soc. Ser. B **4** (2017), 53–71.
- [Yoo19] Hwajong Yoo, *Non-optimal levels of a reducible mod l modular representation*, Trans. Amer. Math. Soc. **371** (2019), no. 6, 3805–3830.
- [Yun09] Zhiwei Yun, *Weights of mixed tilting sheaves and geometric Ringel duality*, Selecta Math. (N.S.) **14** (2009), no. 2, 299–320.
- [Zel80] A. V. Zelevinsky, *Induced representations of reductive p -adic groups. II. On irreducible representations of $GL(n)$* , Ann. Sci. École Norm. Sup. (4) **13** (1980), no. 2, 165–210.
- [Zhu20] Xinwen Zhu, *Coherent sheaves on the stack of Langlands parameters*, 2020.

Email address: `emerton@math.uchicago.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, 5734 S. UNIVERSITY AVE., CHICAGO, IL 60637, USA

Email address: `toby.gee@imperial.ac.uk`

DEPARTMENT OF MATHEMATICS, IMPERIAL COLLEGE LONDON, LONDON SW7 2AZ, UK

Email address: `e.hellmann@uni-muenster.de`

MATHEMATISCHES INSTITUT, WESTFÄLISCHE WILHELMS UNIVERSITÄT MÜNSTER, EINSTEIN-STRASSE 62, 48149 MÜNSTER, GERMANY