Geometry with Valuations

Lectures delivered by Nir Avni
Notes by Akhil Mathew

Fall 2011, Harvard

Contents

Introduction

Nir Avni taught a course (Math 254y) titled "Geometry with valuations" at Harvard in Fall 2011. These are my "live-TEXed" notes from the course.

Conventions are as follows: Each lecture gets its own "chapter," and appears in the table of contents with the date. Some lectures are marked "section," which means that they were taken at a recitation session. The recitation sessions were taught by .

Of course, these notes are not a faithful representation of the course, either in the mathematics itself or in the quotes, jokes, and philosophical musings; in particular, the errors are my fault. By the same token, any virtues in the notes are to be credited to the lecturer and not the scribe.

Please email corrections to `amathew@college.harvard.edu`.

Lecture 1
9/6

This is the real beginning of the course. (**N.B.** This was really the second lecture. The first lecture was not TEXed.) We have to start with a lot of boring things (after the fun things last lecture). However, we want to give them names. The course will assume that everyone has seen first-order logic. However, we'll set the notation. In order not to make this too boring, we'll apply this to algebraic geometry. In some ways this is foolish, because we know more about algebraic geometry than we know about logic. However, there are two reasons to use logic:

1. Like category theory, it directs you to interesting concepts.

2. It's more general. For instance, take the Nullstellensatz in alg. geo.; it's not clear in other contexts where to look for analogs of it. Model theory will give us ideas.

§1 The Ax-Grothendieck theorem

We want to start by aiming for:

1.1 Theorem. *Suppose $f : \mathbb{C}^n \to \mathbb{C}^n$ is an injective polynomial map. Then f is surjective.*

This is a theorem in algebraic geometry, which can be proved using algebraic geometry, but model theory gives a different perspective. We will start with a very loose idea.

Proof. In three steps:

1. This theorem holds with $\mathbb{C}$ replaced by a finite field $\mathbb{F}_q$. This is obvious: the pigeon-hole principle. An injective endomorphism of a finite set is a surjection.

2. The result holds for $\mathbb{C}$ replaced by $\overline{\mathbb{F}_p}$, the algebraic closure of a finite field. Indeed, if a polynomial map is defined over $\overline{\mathbb{F}_p}$, then it is defined over $\mathbb{F}_q$ for some large q. We can consider the map $f : \overline{\mathbb{F}_p}^n \to \overline{\mathbb{F}_p}^n$ as the limit of the maps over finite fields. Since all these are surjective, the map f is surjective.

3. Morally,
$$\lim_{p \to \infty} \overline{\mathbb{F}_p} = \mathbb{C}.$$
Therefore, since the claim holds for $\overline{\mathbb{F}_p}$, the result for $\mathbb{C}$ (which we want) follows.

▲

We have **not** made precise what the silly claim about limits really means. We will do this over time.

Remark. The converse is not true. A surjective polynomial map is not necessarily injective (e.g. squaring). It fails at the second step.

Now we need to make this precise. What is the topology on the set of all fields? It has a name; it is the **first-order topology.**

1.2 Definition. The **first-order topology** on the class of fields has as a basis sets of the form

$$U_\phi = \{F : F \text{ satisfies a fixed first–order formula } \phi\}.$$

Now the class of all fields is not a set. So you should really restrict to fields that aren't too large (e.g. we could use universes). Anyway, in the end we're not just going to use fields, but fields up to elementary equivalence.

It is easy to see that a finite intersection of a bunch of U_ϕ is still another set of the same form (i.e. $U_{\phi_1} \cap \cdots \cap U_{\phi_n} = U_{\phi_1 \wedge \cdots \wedge \phi_n}$). The empty set is in this basis, for instance we could take $\phi = \{\forall x, x = 0\}$ to get $U_\phi = \emptyset$. We could also get the full class of fields, e.g. $\phi = \forall x, x = x$.

Let us clarify what *first-order* really means. More formally: we use the following first-order language, with:

1. Two constants $0, 1$.
2. Two functions $+, \times$ which are binary operations.
3. One unary function $-$.
4. No relations.

1.3 Definition. A **first-order formula** is an expression (well, a "gramatically correct" expression) that can be built of these two constants, three operations, variables, logical connectors $(\wedge, \vee)$ and existential and universal quantifiers $\exists, \forall$.

To really define this formally, we would have to spend a lot more time. Let us just give some examples.

1.4 Example. For instance, the first-order formula $(\forall x)(x = 1 + 1)$ is grammatically kosher (though it is never satisfied). However, $(\forall x)(x+)$ is *not* OK.

We are really talking about something very elementary, and very trivial, but it's not fun to describe properly what a grammatically correct expression is. It will be easy to tell one.

In general, formulas that one writes will have *free* variables, i.e. variables that are not quantified by anything.

1.5 Example. $(\exists y)(y \times y = x)$ is a formula saying something about x. This is a formula saying something about x.

However, other formulas will not have free variables.

1.6 Definition. A formula without free variables is called a **sentence.**

These don't talk about elements. They talk about "truths" (that apply to a given structure or not). The most important notions that we need to have now are the first-order language above (called the **first-order language of rings**) and the notions of **formula** and **sentence**.

Given a sentence ϕ, we can (as before) talk about a field satisfying it or not. This is how we defined U_ϕ above: it is the set of all fields in which ϕ holds. This generates a topology, for which the $\{U_\phi\}$ form a basis on the class of all fields. It is *not* Hausdorff, and not even T_0. There are fields which are not isomorphic that satisfy *exactly* the same first-order formulas. For instance, every two algebraically closed fields of the same characteristic satisfy the same formulas. In algebraic geometry, this is called the **Lefschetz principle.**

Remark. Since the complement of U_ϕ is $U_{\not\phi}$, each U_ϕ is also closed. In particular, the topology is totally disconnected.

Here's the real theorem:

1.7 Theorem (Godel)**.** *This topology makes the class of all fields into a compact space.*

This is where the math happens.

In particular, there is an accumulation point to any infinite set of fields (e.g. $\left\{\overline{\mathbb{F}_p}\right\}$), which is going to be a field (not necessarily unique). We are going to see that this sequence converges in fact to $\mathbb{C}$, as it does to $\overline{\mathbb{Q}}$. Thus, we can elucidate further the proof of Ax-Grothendieck.

Well, so let K be an accumulation point of the $\left\{\overline{\mathbb{F}_p}\right\}$. Consider the sentence $\phi_n = (\forall a_0 \forall a_1 \dots \forall a_n)(a_n \neq 0 \implies \exists x)(a_n x^n + \cdots + a_0 = 0)$. This sentence states that all polynomials of degree n have roots. Since each $\overline{\mathbb{F}_p}$ lives inside the closed set U_{ϕ_n}, the accumulation point K does too. Thus K satisfies all the sentences $\{\phi_n\}$, i.e. is algebraically closed.

Now consider the sentence $\psi_3 = (1 + 1 + 1 \neq 0)$. Then almost all (all but one) $\overline{\mathbb{F}_p}$ belong to U_{ψ_3}, so the accumulation point must belong to U_{ϕ_3} (which is closed). So $1+1+1 \neq 0$ in K, which means K can't have characteristic three. Similarly, K can't have any nonzero characteristic, so K is an algebraically closed field of characteristic zero.

Each $\overline{\mathbb{F}_p}$ satisfies the condition that each injective polynomial map is surjective; this is a countable collection of first-order sentences. (We can't quantify over polynomial maps, but we can quantify over polynomial maps of degree $\leq N$ for each N, because these are described by a finite amount of data.) Reasoning similarly with the topology, we see that K satisfies this property: every injective polynomial endomorphism of K^n is surjective. Now we can use the Lefschetz principle to replace K with $\mathbb{C}$. So $\mathbb{C}$ has the Ax-Grothendieck property. This is the explanation for the "limiting process" used in the earlier proof.

§2 First-order logic

We will talk about more general first-order languages, formulas, and sentences. A **first-order language** will have a fixed set of (symbols of) constants, functions, and relations. For instance, we have:

1.8 Example. 1. The language of rings (with addition, multiplication, additive inversion, $1, 0$) as before. This will let us talk about residue fields of valued fields.

2. The language of ordered groups (which has a constant zero, a unary function $x \mapsto -x$, a binary addition function, and an order relation).

3. We want to add the language of valued fields, but this is slightly more complicated. We'll defer this for now.

This is purely *syntactic*. All we can say is that an expression is syntactically correct. To talk about truthfulness, we need structures. Previously, we thought of fields as special types of structures.

1.9 Definition. A **structure** for a first-order language is a set together with an *interpretation* of the language: constant symbols go to elements of the set, function symbols go to functions, relation symbols go to relations (of the appropriate arity).

Given a structure, we can ask whether it *satisfies* a structure. If a sentence ϕ (i.e. a formula without free variables, as before) holds in the structure M, we write

$$M \vdash \phi.$$

If $\phi(x_1, \dots, x_n)$ is a formula on the free variables $x_1, \dots, x_n$, and M is a structure, we can write

$$\phi(M) = \{(a_1, \dots, a_n) \in M^n : M \vdash \phi(a_1, \dots, a_n).\}$$

We think of ϕ as a "family" of sentences.

All the action starts when one starts giving axioms.

1.10 Definition. A set of sentences over a given first-order language is called a **theory**.

Not all theories are the same. If T is a theory, though, and M a structure, we say that $M \vdash T$ if $M \vdash \phi$ for each $\phi \in T$. Then we say that M is a **model** of T.

We are only interested in theories that have models. The great theorem of model theory is that this is actually a syntactic property of the set T; we don't have to try it on models to check whether a theory has models.

1.11 Theorem (Completeness)**.** *A theory T has models if and only if one cannot "derive a contradiction" from it using T as axioms and ordinary laws of logic.*

In other words, T has models if and only if one *cannot* derive both ϕ and $\not\phi$ from T (for some, or any, sentence ϕ). This is a nontrivial fact, and is actually equivalent to the axiom of choice.

Here is why completeness implies compactness, though, as well as the more general form of compactness:

1.12 Theorem (Compactness)**.** *Let T be a theory (over a fixed first-order language), and endow the class of all models of T with the first-order topology (defined as previously). Then the associated space is compact.*

Proof. We need to show that a decreasing collection of closed basic sets has the finite intersection property, then it has a nonempty intersection. Let $\{\phi_\alpha\}$ be a collection of sentences such that each finite subset of them has a model (i.e. each $U_{\phi_{\alpha_1}} \cap \cdots \cap U_{\phi_{\alpha_n}}$ is nonempty) which is also a T-model. Then we want to show that there is a model for the whole theory. This will imply the result.

However, the hypothesis means that one cannot derive a contradiction from $T \cup \{\phi_{\alpha_1} \cup \cdots \cup \phi_{\alpha_n}\}$ for each choice of $\alpha_1, \dots, \alpha_n$. This means that one cannot derive a contradiction from the big set $T \cup \{\phi_\alpha\}$, because a proof has finite length and can only involve a finite number. But by the completeness theorem, there is a model for $T \cup \{\phi_\alpha\}$, which is a model of T satisfying all the $\{\phi_\alpha\}$. ▲

Let us discuss a consequence of this compactness theorem. Fix a theory T with models over a given language. Suppose ϕ is a formula, and let's write $\phi(x)$ even if there are many free variables (so x can stand for many variables). Suppose for every model M of T, the set $\phi(M)$ is finite. Then, there is a constant C such that $\phi(M) \leq C$ for *all* models M of T. This is a direct consequence of compactness. Otherwise, we could consider the sentences $\{\phi(M) > n\}$; every finite subset would be satisfiable, and thus all of them would be simultaneously satisfiable.

In particular: **There is no way in first-order logic to get arbitrarily big but finite sets.**

§3 Further directions

Anyway, we can now start looking at limits now that this formalism is established. We might wonder:

1.13 Example. What is the limit (or rather, limits) of the set of finite fields (or just the set of prime fields)?

Suppose F is a limit point of the set $\{\mathbb{F}_q\}$ of all finite fields, in the theory of fields (and the language of rings). F is then infinite, since the fields in this set get arbitrarily big (and for any N, only finitely many have cardinality $< N$). However, F is not algebraically closed, because each finite field has a unique quadratic extension. This is a first-order property, and as a result F too has a unique quadratic extension. More generally, F has a unique extension of degree n for each n, because the same is true for a finite field. Thus the absolute Galois group of F is $\hat{\mathbb{Z}}$.

F is also pseudo-algebraically closed: any absolutely (i.e. geometrically) irreducible variety has a rational point. This is a first-order property, and follows because this is "asymptotically" true over finite fields.

Surprisingly, this is precisely what characterizes the limit points of the set of finite fields. Any pseudo-algebraically field with absolute Galois group $\hat{\mathbb{Z}}$ is a limit point of finite fields; these are called **pseudo-finite fields.** They satisfy the same sentences "asymptotically" as finite fields. However, not all pseudo-finite fields satisfy the same relations. For instance, half the prime fields are such that -1 is a square, and half are such that this fails. It follows that there are pseudo-finite fields where -1 is a square and pseudo-finite fields where -1 is not a square. (However, it turns out that this is the only thing that distinguishes pseudo-finite fields as far as model theory is concerned.)

Let $\mathfrak{A}$ be a non-principal ultrafilter on the set of prime fields $\{\mathbb{F}_p\}$. Let $K = \lim_{\mathfrak{A}} \mathbb{F}_p$. We can also do something similar: we can consider the language of valued fields (to be described later!) and consider the limit $\lim_{\mathfrak{A}} \mathbb{F}_p((t))$ of the *valued fields* $\mathbb{F}_p((t))$. It's going to be some valued field. Similarly, we can do the same with $\lim_{\mathfrak{A}} \mathbb{Q}_p$. We are going to show

$$\lim_{\mathfrak{A}} \mathbb{F}_p((t)) = \lim_{\mathfrak{A}} \mathbb{Q}_p.$$

This is somewhat surprising. Both will turn out to be valued fields with residue field K. From this we will show:

1.14 Corollary (Ax-Kochen)**.** *For every sentence ϕ, for almost all p, $\mathbb{Q}_p$ satisfies ϕ if and only if $\mathbb{F}_p((t))$ satisfies ϕ.*

Moreover, we will show that for any formula $\phi(x_1, \ldots, x_n)$, the sets of solutions $\phi(\mathbb{Q}_p)$ and $\phi(\mathbb{F}_p((t))$ will have *equal Haar measure* for $p \gg 0$.

Lecture 2
9/8

§1 Elementary equivalence and embeddings

Last time, we talked about "limits" of structures over a first-order language. These limits were *not* unique. Here is one reason they are not unique: one can have two models of a theory that satisfy *precisely* the same sentences. For instance, we are going to see that every sentence in the language of rings that is satisfied in $\overline{\mathbb{Q}}$ is satisfied in $\mathbb{C}$.

2.1 Example. The statement "F is not algebraic over the prime field" cannot be stated in first-order logic, as a result.

Here is another reason. Suppose ϕ were such a sentence that meant "the field is algebraic over the prime field." Then build a larger language with uncountably many constants c_α, and add the conditions $c_\alpha \neq c_\beta$. Every finite subset of this family of sentences in the larger language is satisfiable, so the whole family is satisfiable. The result will be an uncountable field algebraic over the finite field, a contradiction.

There is a sharp restriction between *finite* and *infinite* in model theory. Finite things can be bounded. If a sentence has infinite models, it has arbitrarily large infinite models.

2.2 Definition. Two structures M_1, M_2 (over some fixed first-order language) are **elementarily equivalent** if they satisfy the same sentences. That is, M_1 satisfies ϕ if and only if M_2 does.

Given two structures M_1, M_2 and a map $f : M_1 \hookrightarrow M_2$ (assumed injective), then f is called an **elementary embedding** if M_1, M_2 are elementarily equivalent and, moreover, for every *formula* (possibly with free variables) $\phi(x_1, \ldots, x_n)$ and any $a_1, \ldots, a_n \in M_1$, we have that $\phi(a_1, \ldots, a_n)$ satisfied in M_1 if and only if $\phi(f(a_1), \ldots, f(a_n))$ satisfied in M_2. (The second condition implies the first.)

§2 Definable sets

From this, we will give the definition of the main thing we care about, definable sets.

2.3 Definition. Let T be a theory. A T**-definable set** is an equivalence class of formulas under the equivalence relation of $\phi(x) \sim \psi(x)$ (where x may be several variables) if and only if for all models M of T, the sets $\phi(M), \psi(M)$ are equal.

This equivalence relation is a semantic condition, but it can be phrased syntactically too. Namely, by Gödel's theorem, $\phi(x) \sim \psi(x)$ if and only if T proves the sentence $\forall x(\phi(x) \iff \psi(x))$.

Alternatively, a definable set is a *functor* from the category of models of T with morphisms the elementary embeddings to the category of sets, which is of the form $M \mapsto \phi(M)$ for some ϕ.

These are the generalized "varieties" we can think about. The analogy is that one thinks of varieties as functors (from the opposite category of rings to sets) in the fancy modern language.

2.4 Example. An scheme of finite type over $\mathbb{Z}$ is a definable set on the theory of rings, or fields, or algebraically closed fields.

If we want to use the relative context (i.e. talk about varieties over a field), one has to enlarge the language to include the constants. So, for instance, if one is interested in $\mathbb{C}$-varieties, one has to enlarge the language of rings to add constants for each element of $\mathbb{C}$. Thus one gets a language of $\mathbb{C}$-algebras.

2.5 Example. Any affine $\mathbb{C}$-variety gives a definable set in the theory of $\mathbb{C}$-algebras.

The notion of a definable set is dependent on the theory T. Namely, two definable sets may be equal in one theory T, but not in another theory T'. For instance, consider the definable set $\mathbb{A}^1$ defined by the one-variable formula $\phi(x) = \{1 = 1\}$, and the definable set $\psi(x) = (\exists y)(y^2 = x)$. These definable sets are the *same* over the theory of algebraically closed fields. However, they are not the same over the theory of fields. This is nothing really new, but simply a warning about the notations.

2.6 Example. If X, Y are definable sets, then one can define definable sets $X \times Y$, $X \cup Y, X \cap Y$. (The latter require that the formulas have the same number of variables.) One can also talk about the notion $X \subset Y$, etc.

Now we want to make the collection of definable sets into a *category.*

2.7 Definition. A **definable function** between X, Y is a definable subset $Z \subset X \times Y$ such that for every model M (of the theory), $Z(M) \subset X(M) \times Y(M)$ is the graph of a function $X(M) \to Y(M)$.

§3 Chevalley's theorem and elimination of quantifiers

One of the most basic theorems in algebraic geometry is the theorem of Chevalley, which states that the image of a constructible set under a regular map of varieties is itself constructible. Recall that a *constructible set* (in some ambient variety) is a boolean combination of subvarieties. Thus, in particular, a constructible set is a definable set (because every subvariety is definable), and can be defined using formulas of the form $P(x_1, \ldots, x_n) = 0$ or $P(x_1, \ldots, x_n) \neq 0$. In particular, the constructible sets are precisely those definable sets expressible using formulas *without quantifiers.*

Chevalley's theorem can be stated in the following form. Let $\mathbb{A}_k^n$ denote the affine space over an algebraically closed field.

2.8 Theorem. *If $X \subset \mathbb{A}_k^{n+m}$ is constructible, then its projection to $\mathbb{A}_k^n$ is constructible.*

Here its projection to $\mathbb{A}_k^n$ is equal to the definable set given by the formula $\phi(x) = (\exists y)((x, y) \in X)$. So what Chevalley's theorem states is that this formula ϕ is *equivalent* to a formula without quantifiers.

Here is the model-theoretic version of Chevalley's theorem:

2.9 Theorem. *Any formula in the theory of algebraically closed fields (denoted ACF) is equivalent to a formula without quantifiers.*

This is called *elimination of quantifiers.* A theory is said to admit elimination of quantifiers if every formula is equivalent (with respect to this theory) to a formula without quantifiers. Or, every definable set is actually quantifier-free. We will prove it below after some work.

We want to give a condition for elimination of quantifiers. Let $\phi(x)$ be a formula, and it is equivalent to a quantifier-free formula. Suppose M, N are two models of the given theory, and they both contain a sub-structure $A \subset M, N$ (so closed under all sub-structures). The claim is that:

2.10 Proposition. *If for every $a \in A$, M satisfies $\phi(a)$ if and only if N satisfies $\phi(a)$.*

Proof. This is because $\phi(a)$ is equivalent to a formula asserting that some functions of a are equal or unequal. ▲

That is, the structure used to test a formula on an element is irrelevant, for a quantifier-free formula. We are going to prove the converse:

2.11 Theorem. *Let $\phi(x)$ be a formula. Suppose for pair of T-models M, N, and for every substructure*[1] *$A \subset M, N$, then an element $a \in A$ satisfies $\phi(x)$ in M if and only if it does in N. Then $\phi(x)$ is T-equivalent to a quantifier-free formula.*

Proof. Assume that $\phi(x)$ is a formula satisfying the hypotheses of the theorem. Suppose $\phi(x)$ is satisfiable, i.e. there is a model with $\phi(M) \neq \emptyset$ (or the result is trivial). We are going to use a trick: we will *enlarge* the language by adding a constant symbol c with no relations on it. A model for this extended language is just a model of the original language with a point. We can extend the theory T to a theory T' on the new language, and a model of T' is the same thing as a model of T with a distinguished point. (If x is actually an array of variables, we let c be an array of constants correspondingly.) We are going to show that T' implies that $\phi(c) \iff \psi(c)$ for some quantifier-free formula ψ. This will be enough, since c can be arbitrary.

[1] A is not necessarily a model. For instance, a non-algebraically-closed subfield of an algebraically closed fields in ACF.

We start with a weaker assertion. There is a satisfiable, quantifier-free formula $\psi(x)$ such that $\psi(c) \implies \phi(c)$. Indeed, suppose the contrary. Let M be a model and think of $\phi(M)$ as some subset of a power of M. We are going to show that there is a *smaller* quantifier-free definable set inside of ϕ. This is our claim. Pick a point $a \in M$ such that $M \implies \phi(a)$. Consider the collection Ω of all quantifier-free formulas satisfied by a. We want to show that Ω implies ϕ. By assumption, $T \cup \{\psi(c)\}_{\psi\in\Omega}$ together with $\neg\phi(c)$ has a model N (since ϕ is not implied by quantifier-free formulas by hypothesis), in view of the compactness theorem.

But we have the T-models M, N, and the substructures $A \subset M, B \subset N$ generated by a or c. This is because to check that substructures are isomorphic, we only need to check on quantifier-free formulas, and the quantifier-free formulas satisfied by a, c are the same by construction. There is thus an isomorphism $A \simeq B$ sending $a \mapsto c$, because these satisfy the *same relations* with respect to the functions in the language. As a result, $\phi(a)$ if and only if $\phi(c)$ by hypothesis (ϕ had a special property). This, however, is a contradiction by choice of c and a.

We thus have proved the claim that there is a satisfiable, quantifier-free formula $\psi(x)$ such that $\psi(c) \implies \phi(c)$.

Now let's do the same trick again. Let X be the set of all quantifier-free formulas ψ such that $\psi(c) \implies \phi(c)$. We have shown $X \neq \emptyset$. Suppose there is no $\psi \in X$ such that $\phi(c) \implies \psi(c)$. If there is, then we are done. Consider the theory $T \cup \{\phi(c)\} \cup \{\neg\psi(c)\}_{\psi\in X}$. This is either consistent or inconsistent.

1. If it is inconsistent (i.e. does not have a model), then $T \cup \{\phi(c)\} \cup F$ is inconsistent for F a finite subset of $\neg X$. Thus a finite number of $\neg\psi(c), \psi \in X$ proves $\neg\phi(c)$ (with T), so
$$T \implies (\phi(c) \implies \psi_1(c) \vee \cdots \vee \psi_n(c)).$$
We win.

2. If it is consistent, choose a model (M, c). This satisfies $\phi(c)$, but no $\psi(c)$ for $\psi \in X$.

   We saw earlier that there is a quantifier-free formula Ψ such that $\Psi(c) \implies \phi(c)$ and such that $\Psi(c)$ holds. (This is a little stronger than the first claim, but came out in the proof.) But then $\Psi \in X$, a contradiction, because c was taken so that $\Psi(c)$ was false.

To reiterate, we considered for a definable set and a point in it, the class of quantifier-free neighborhoods of this point. We showed that we could always find such a neighborhood, and then after that, we more or less took the union of all these neighborhoods of points. ▲

Finally, we want to show that ACF has elimination of quantifiers. From this, we can prove the Lefschetz principle. Any two algebraically closed fields of characteristic p contain a common substructure (the prime field). As a result, since any sentence is equivalent (by elimination of quantifiers) to a sentence involving only $1, 0, =, +, \times$, it follows that any sentence is true or false uniformly in all algebraically closed fields of a given characteristic.

Let us now prove that ACF admits elimination of quantifiers.

Proof. It's enough to show that any formula of the form $(\exists x)\phi(x, y_1, \dots, y_n)$ (here x really means one variable) with ϕ quantifier-free is equivalent to some quantifier-free condition on $y_1, \dots, y_n$. Then we can induct on the number of quantifiers.

That is, by the previous theorem, if M, N are algebraically closed fields with a common subfield A, we need to show that for $a_1, \dots, a_n \in A$, the one-dimensional formula $\phi(x, a_1, \dots, a_n)$ has a solution in M if and only if it has one in N. But in the language of rings, any quantifier-free set (such as $\{x : \phi(x, a_1, \dots, a_n)\}$) is either finite or cofinite, over an algebraically closed field. If finite, then there are finitely many solutions contained in the algebraic closure of the fraction field $K(A)$, which is contained in both M, N. If cofinite, then there are solutions in both M, N. (*This can also be re-done using elimination theory.* ▲

As stated earlier, we have:

2.12 Corollary. *A sentence in the language of rings is true in one algebraically closed field of char. p if and only if it is true in all algebraically closed fields of char. p.*

Here is another example of a theory with elimination of quantifiers.

2.13 Example. The theory of divisible ordered abelian groups admits elimination of quantifiers. To see this, we reduce to the case of one quantifier, as in the earlier proof for algebraically closed fields. So we need to analyze how quantifier-free one-dimensional sets look like. Let $\phi(x)$ be a one-dimensional, quantifier-free formula using coefficients in some abelian group A. We want to show that it has a solution in one divisible ordered abelian group M containing A iff it does in another N.

But ϕ is equivalent to a union of intervals whose endpoints are in $A \otimes \mathbb{Q}$, and for such sets, non-emptiness is equivalent in M or N (e.g. the midpoint of such an interval will be in both M, N).

Here is a much more interesting example.

2.14 Example. Consider the theory RCF' of $\mathbb{R}$ in the language of rings: that is, it consists of all sentences satisfied by $\mathbb{R}$. This doesn't have elimination of quantifiers. For instance, $(\exists y)(x = y^2)$ cuts out the positive real numbers, and is *not* equivalent to a quantifier-free formula. However, this the only obstruction to elimination of quantifiers. If we add an ordering, then we don't have a problem: let RCF be the theory of $\mathbb{R}$ in the language of *ordered* rings. Then RCF satisfies elimination of quantifiers. We'll do this next time.

Lecture 3
9/13

§1 Elimination of quantifiers II

So we want to recap the argument last time on elimination of quantifiers, but in a slightly different language. This will help pinpoint where the math actually is. For instance, we'll (re)prove that the image of a variety is constructible.

3.1 Proposition. *If $V \subset \mathbb{A}^{n+m}$ is a variety defined over $\mathbb{Q}$, then the image under projection $\pi : \mathbb{A}^{n+m} \to \mathbb{A}^n$ is a constructible set defined over $\mathbb{Q}$. In particular, for any algebraically closed field F of characteristic zero , $\pi V(F)$ is a boolean combination of varieties defined over $\mathbb{Q}$.*

Proof. Let $X_1, X_2, \ldots,$ be an enumeration of all constructible sets, defined over $\mathbb{Q}$, contained in in this projection $\pi(V)$. (Here we're using the fact that one is working over a countable field.) We want to show that the union $\bigcup X_i$ is $\pi(V)$.

Indeed, then $\pi(V)$ is a countable union of constructible sets, and we can thus find an *increasing* sequence of constructible sets Y_i such that $\pi(V) = \bigcup Y_i$. Suppose $\pi(V) \neq Y_i$, for each i. There is a large algebraically closed field F and $c \in \pi(V)(F)$ that is not in *any* of these Y_i: this follows because we can find some field F_i and a point in $\pi(V)(F_i) - Y_i(F_i)$ for each i, and then we can use compactness. Or more explicitly, to construct F from the $\{F_i\}$, one can take the quotient $\prod F_i/\mathfrak{m}$ for $\mathfrak{m}$ a non-principal maximal ideal.

So, we have a field F and a point $c \in \pi(V)(F) - \bigcup Y_i(F)$. For *any* algebraically closed field F, we must however have

$$\pi(V)(F) = \bigcup X_i(F).$$

This will be a contradiction. To prove this claim, suppose given a counterexample $a \in \pi(V)(F) - \bigcup X_i(F)$. Let $\{Z_i\}$ be an enumeration of all varieties defined over $\mathbb{Q}$ that contain a. None of these varieties is contained in $\pi(V)(F)$ by construction. There are elements $\{b_i\}$ in $\pi(V)(F) - Z_i(b)$. Now we will do the same trick again, of passing to a larger field. Let $R = \prod F$, $\mathfrak{m} \subset R$ a non-principal maximal ideal. Then R is a F-algebra via the diagonal, and the quotient $K = R/\mathfrak{m}$ is also an F-algebra.

Let $b = (b_i)$ be the product element of $\prod F$, regarded as an element of K. By construction, $b \notin \pi(V)(K)$. But a, imbedded via the diagonal, is in $\pi(V)(K)$. The map $\mathbb{Q}(a) \to \mathbb{Q}(b)$ is, however, an isomorphism, so a rationally defined variety contains a if and only if it contains b. This map $\mathbb{Q}(a) \to \mathbb{Q}(b)$ can be extended to an elementary embedding $F \hookrightarrow K$, which contradicts the model completeness of algebraically closed fields proved earlier. **I was a little confused here and should try to fix this later.** ▲

§2 Real closed fields

Our next goal is to prove that algebraically closed *valued* fields have elimination of quantifiers. But first we want to give another example of this phenomenon, where we study the topology of real algebraic varieties. We will later do this again for Berkovich spaces.

Recall that RCF is the theory of $\mathbb{R}$ (the set of all sentences $\mathbb{R}$ satisfies) in the language of $+, \cdot, \times$: for instance, every odd-degree polynomial has a root. This is the theory of *real closed fields.*

3.2 Proposition. *RCF has elimination of quantifiers.*

Proof. It all comes down to a theorem about quantifier-free one-dimensional problems. Here is the claim.

3.3 Proposition. *If X is a quantifier-free, one-dimensional (i.e. a subset of $\mathbb{R}$) set, defined over some real closed field K, then X is a boolean combination of intervals with endpoints in K.*

Proof. The fact that it is a boolean combination of intervals is quite easy. All you can really say is that a certain polynomial is less than zero. The real content of the theorem is that the endpoints are in K. It follows from the fact that if $K \subset F$ are real closed fields, and f is a polynomial in $K[x]$, then the F-roots of f are inside K. The reason is that K has no nontrivial algebraic extensions other than the algebraic closure $K(\sqrt{-1})$. Once you add $\sqrt{-1}$, however, you cannot imbed inside a real closed field. So the roots of f must lie in $K = \overline{K} \cap F$. ▲

Actually, a quantifier-free one-dimensional set is even a *disjoint union* of intervals with endpoints in K. In fact, these intervals must be themselves definable (i.e. whose endpoints are definable), over the given language (which possibly has constants added to it). Let's postpone that latter statement, though.

To show that RCF has elimination of quantifiers, we need to show that a one-dimensional, quantifier-free definable set is nonempty over one real closed field if and only if it is in a different, bigger model by the basic criterion. (Note that though RCF is a complete theory, that does not mean it admits elimination of quantifiers.) However, this follows by the description as a union of intervals. ▲

Consider now 2-dimensional definable sets over a real closed field. Let $X \subset \mathbb{A}^2$ be definable over RCF, possibly with additional constants. We treat this as fibered over $\mathbb{A}^1$, via projection $\pi : X \to \mathbb{A}^1$.

3.4 Proposition. *There is a definable partition of $\mathbb{A}^1 = X_1 \cup \cdots \cup X_n$ such that the number of intervals of $\pi^{-1}(x) \cap X \subset \mathbb{A}^1$ is constant for each $x \in X_i$, and such that the endpoints are definable functions.*

Proof. Enumerate all formulas of the form $\phi(x, y) : x \in \bigcup_i [f_i(y), g_i(y)]$ (or open intervals) for some definable functions f_i, g_i. Let us say that the sets X_ϕ are defined as $\{y \in \mathbb{A}^1 : X_y = (\forall x)(x \in X_y \iff \phi(x, y)\}$, i.e. those sets where the fibers can be described by a certain formula. For every model F, $\mathbb{A}^1(F)$ is the union of all X_ϕ: every one-dimensional set is a finite union of intervals whose endpoints are definable. By compactness, the relative version of this claim (the present proposition) follows. ▲

In particular, suppose X definable and *closed* and bounded. (Closedness can be defined via balls.) In this case, X can be fibered as unions of closed intervals with definably varying endpoints. There is thus a finite simplicial complex to which $X(\mathbb{R})$ is homeomorphic. This is because there are finitely many intervals over which X is a union of n intervals, etc.

Even more is true. If $X \subset \mathbb{A}^n \times \mathbb{A}^m$ is a closed, bounded definable family of definable subsets, then only finitely many different definable homeomorphism classes of $\mathbb{R}$-points among the fibers. We essentially saw this above for $n + m = 2$.

Proof. For every finite simplicial complex Δ and every definable function f from the geometric realization to $\mathbb{A}^n$, the definable set $X_{\Delta,f}$ consisting of "f is a homeomorphism of Δ and the fiber y," then $\mathbb{A}^m$ is the union of all these. This means it must be a finite union. ▲

§3 Valued fields

Now let's start the course. Next time, we'll prove that algebraically closed valued fields have elimination of quantifiers. Let us just set up the language. This language is a bit complicated.

3.5 Definition. The **language of valued fields** will have three types of variables ("sorts"):

1. The valued field (usually F).
2. The residue field (usually k).
3. The value group (usually Γ).

There is one relation $<$, and is a relation only between elements of Γ. There are many functions:

1. Binary functions $+, \cdot$ on F, k, Γ (i.e. addition, multiplication, etc.; $\cdot$ is not defined on Γ)
2. A valuation function $v : F^* \to \Gamma$ (or $v : F \to \Gamma \cup \{\infty\}$).
3. A function $\mathrm{ac} : F \to k$ (the "angular component").

We will define the **theory of algebraically closed, valued fields**, written $ACVF$, as having the following axioms.

1. $(F, +, \cdot)$ and $(k, +, \cdot)$ are algebraically closed fields.
2. $(\Gamma, +, <)$ is a divisible ordered abelian group (abbreviated **DOAG**).
3. v is a non-archimedean valuation.
4. ac (the angular component) is to be thought of as a generalization of reduction mod a prime. It is required to be a homomorphism $F^* \to k^*$, and if $v(a) < v(b)$, we should have $\mathrm{ac}(a + b) = \mathrm{ac}(a)$. If $v(a) = v(b) = v(a + b)$, then we require $\mathrm{ac}(a + b) = \mathrm{ac}(a) + \mathrm{ac}(b)$.
5. v, ac are surjective.
6. $\mathrm{ac}^{-1}(0) = 0$.

Remark. It might happen that $\gamma \in \Gamma$ is definable, but there are no definable elements of F whose valuation is γ. For example, one could add a random constant to Γ without doing anything else, and make that constant definable. The definable set $\{x : v(x) \geq \gamma\}$ is an invertible module over the ring of integers in K, which is *not* definably isomorphic to this ring of integers. (It will be isomorphic after some extension of scalars.)

There will not be a class this Thursday.

Lecture 4
9/20

It's about time to talk about valued fields in this course, which we haven't.

§1 Valuations

Let (F, val) be a nonarchimedean valued field. From this data, we can get a *ring of integers* $\mathcal{O}$ consisting of elements $x \in F$ with $\text{val}(x) \geq 0$. Then $\mathcal{O}$ is a proper subring of F, and we observe that for each $x \in F$, either $x \in \mathcal{O}$ or $x^{-1} \in \mathcal{O}$. That is, $\mathcal{O}$ is a *valuation subring* of the field F.

The opposite is also true. When one has a valuation subring $\mathcal{O} \subset F$, one gets a valuation of F such that $\mathcal{O}$ becomes the ring of integers. We start with:

4.1 Lemma. *If $\mathcal{O}$ is a valuation ring in F, then $\mathcal{O}$ is local.*

Proof. We need to show that if $x, y \in \mathcal{O}$ are noninvertible, then $x+y$ is non-invertible. Suppose x, y are not invertible. One of x/y and y/x is in $\mathcal{O}$, by definition. Without loss of generality, $x/y \in \mathcal{O}$. Consequently $\frac{x+y}{y} \in \mathcal{O}$, because it is $x/y+1$. So $x+y = \left(\frac{x+y}{y}\right) y$ is the product of a non-invertible element of $\mathcal{O}$ and an element of $\mathcal{O}$, so it cannot be invertible. Alternatively, if $x+y$ were invertible, then $\frac{1}{y}$ would have to be in $\mathcal{O}$. ▲

Now given a valuation ring $\mathcal{O} \subset F$, we want to get a valuation. Consider $F^*/\mathcal{O}^*$; this is a group. The claim is that this is an *ordered group.* To get this, we have to decide which elements are greater than the unit. We declare that the image of the maximal ideal $\mathfrak{m} \subset \mathcal{O}$ is greater than the identity, i.e. consists of the positive elements. We have a map

$$F^* \to F^*/\mathcal{O}^*,$$

which will be the valuation. Consequently, we can obtain a valuation from a valuation ring.

We find:

4.2 Proposition. *There is a correspondence between valuations on F and valuation rings in F.*

What's the point of this? Let F be a valued field, and let $K \supset F$ be a field extension. We want to *extend* the valuation of F to K. We can now translate this into the question: let $\mathcal{O} \subset F$ be a valuation ring. Can we extend $\mathcal{O}$ to a valuation ring in K? That is, can we find a valuation ring $\mathcal{O}' \subset K$ such that $\mathcal{O}' \cap F = \mathcal{O}$?

There is a simple way of making this extension. Consider the map $\mathcal{O} \to \mathcal{O}/\mathfrak{m}$ where $\mathfrak{m} \subset \mathcal{O}$ is the maximal ideal. Imbed $\mathcal{O}/\mathfrak{m}$ inside its algebraic closure $\overline{\mathcal{O}/\mathfrak{m}}$; call the map $\mathcal{O} \to \overline{\mathcal{O}/\mathfrak{m}}$ f_0. Consider the collection of pairs (R, f) where R is a subring of K containing $\mathcal{O}$, and $f : R \to \overline{\mathcal{O}/\mathfrak{m}}$ is a homomorphism extending f_0. By Zorn's lemma, we can find a *maximal* such pair (R, f). The claim is that R is a valuation ring in K.

Indeed, let $x \in K$. We want to show that $x \in R$, or $x^{-1} \in R$. In other words, we can either extend $f : R \to \overline{\mathcal{O}/\mathfrak{m}}$ to $R[x]$ or we can extend it to $R[x^{-1}]$. To see this, let $\mathfrak{P} \subset R$ be the kernel of f. Note first that we can extend f to $R_{\mathfrak{P}}$ in a natural way, and consequently R must be local (if R was chosen maximal) with maximal ideal $\mathfrak{P}$. One obstruction to extending $R \to \overline{\mathcal{O}/\mathfrak{m}}$ to $R[x]$ is if $\mathfrak{P}R[x] = R[x]$. In this event, an extension is clearly impossible.

Let us assume x algebraic over R, because if otherwise we can make the extension directly.

So, suppose $\mathfrak{P}R[x] = R[x]$. This means that

$$1 = \sum m_i x^i, \quad m_i \in \mathfrak{P},$$

which means that x^{-1} is *integral* over R. However, once we have this, the going-up theorem implies that we can extend a maximal ideal $\mathfrak{P} \subset R$ to a maximal ideal of $R[x^{-1}]$. This implies that we can extend the homomorphism, which means that to begin with $x^{-1} \in R$.

The other possibility is that $\mathfrak{P}R[x] \neq R[x]$, so $\mathfrak{P}R[x]$ is contained in a maximal ideal $\mathfrak{M} \subset R[x]$. Then the extension is to take $R[x]/\mathfrak{M}$ as a field and to take the natural projection $R[x] \to R[x]/\mathfrak{M}$. Since R was maximal, it must follow that $x \in R$.

We have proved:

4.3 Proposition. *Any valuation on F can be extended to a valuation on any field extension $K \supset F$.*

§2 Quantifier elimination

We want to prove quantifier elimination for algebraically closed valued fields. We will need a lemma:

4.4 Lemma. *Let $F \subset K$ be an extension of algebraically closed, valued fields. Let $\phi(x)$ be a one-dimensional, quantifier-free formula with coefficients in the smaller field F. Consider $a, b \in K$ in the same F-balls (i.e. for any ball whose center and radius[2] are in F resp. the value group, a and b are either both in it or neither in it). Then $\phi(a)$ holds if and only if $\phi(b)$ holds.*

So, once two elements are in the same balls, they are in the same definable sets.

4.5 Example. Suppose K has an element whose valuation is smaller than that of any nonzero element of F. Then that element is in the same balls as twice itself. One can construct such a field (which we can suppose algebraically closed) by considering the language of F, and adding a constant c with the axioms that the valuation of c is less than the valuation of each nonzero element of F. Using the compactness theorem, we can find a model of this new language which is also a valued field K, and where the constant c corresponds to an such an element.

This is like the construction of the hyperreals.

Proof of the lemma. Consider the valued fields $F(a)$, $F(b)$. The valuations on these fields is obtained by restricting that of K. We have residue fields $k(F(a)), k(F(b))$, and value groups $\Gamma(F(a)), \Gamma(F(b))$. There is a map $F(a) \to F(b)$ sending $a \mapsto b$, along with

[2] Possibly ∞.

maps of the residue groups and valued groups, which together are an *isomorphism* of valued fields under F. If we prove this, we will be done. Let us check this.

If $a \in F$, then $b \in F$ as well and $a = b$. Consequently a, b are transcendental over F, so we get a map $F(a) \to F(b)$ of fields under K. More interestingly, we have to show that this map actually leads to a map between $\Gamma(F(a))$ and $\Gamma(F(b))$. We need to show that for any polynomial (or rational function) $f(x) \in F(x)$ and $\gamma \in \Gamma(F)$, if the valuation of $f(a)$ is $< \gamma$, then that of $f(b) < \gamma$. Consequently, we can define a map $\Gamma(F(a)) \to \Gamma(F(b))$ by sending the valuation of $f(a)$ to that of $f(b)$. (Why is this well-defined? Suppose the valuation of $f(a)$ is that of $g(a)$; then that of $f(a)/g(a)$ is equal to zero. Thus that of $f(b)/g(b)$ is zero.)

So, F is algebraically closed. Consequently, $f = \prod(x - \alpha_i)/\prod(x - \beta_i)$. Let us pretend the β_i's don't exist for now, i.e. we prove things for a polynomial. Now the claim about $f(a)$ versus $f(b)$ follows from the hypothesis that a, b are in the same F-balls. For instance, if

$$\sum \mathrm{val}(a - \alpha_i) > \gamma,$$

then we can find[3] γ_i with $\sum \gamma_i > \gamma$ with $\mathrm{val}(a - \alpha_i) > \gamma_i$. (This follows by divisibility.) Now the same holds with a replaced by b because these things lie in the same F-balls.

The last thing to show is that the map on residue fields is well-defined. This too, will be gotten back to at some point. ▲

So, modulo the small IOU marked by a footnote, we can try to prove the main result.

We have shown something semantic: if two elements are in the same F-balls, then they cannot be distinguished by quantifier-free formulas. Here is the syntactic counterpart: every quantifier-free subset of F is a boolean combination of F-balls. This is very similar to the proof of elimination of quantifiers.

Proof. If $\phi(x)$ is nonempty and F is some model of ACVF, and $\alpha \in F$ is some element such that $\phi(\alpha)$ holds, there is a boolean combination of balls that contains α and is contained in $\phi(F) = \{x : \phi(x)\}$. Otherwise, one could enumerate all boolean combinations of F-balls that contain α, and assume that for each one of them there is a point outside $\phi(x)$. Then one would take the ultralimit and would find a point contained in all F-balls but does not satisfy ϕ, a contradiction.

So to each $\alpha \in F$, we can find a ball containing it contained in $\phi(F)$. Consider the set of boolean combinations SW of balls that are contained in $\phi(F)$. The claim is that a finite union is the definable set $\phi(F)$. This is because for every element, there is a boolean combination of balls containing it, so we can apply compactness. ▲

4.6 Corollary. *A nonempty quantifier-free subset of F has a point in every model of ACVF.*

"Nonempty" means that there is a point in some model.

Proof. So we need to show that every boolean combination of balls is nonempty in one algebraically closed valued field if and only if it is nonempty in another. I.e., if B, and

[3] **This needs to be clarified.**

$C_1, \ldots, C_n$ are balls, then $B - C_1 - \cdots - C_n$ is either empty everywhere (in all models) are nonempty everywhere. The point is that one is working with a nonarchimedean field, so any two balls are either contained in each other or they are disjoint. And this is independent of the model, because the value groups are always divisible (for an algebraically closed field). ▲

4.7 Theorem. *The theory of algebraically closed valued fields admits elimination of quantifiers.*

Proof. This follows from the previous result and the general criterion for when a theory has elimination of quantifiers. Indeed, it is enough to show that any formula with one quantifier is equivalent to a formula without quantifiers. I.e., the formula looks like $(\exists x)\phi(x, y_1, \ldots, y_n)$. We need to show that this formula's truthness is independent of which field you evaluate it at, but this follows because the quantifier-free definable set $\phi(x, y_1, \ldots, y_n)$ is a union of balls (as before). There are also the cases of x ranging over the valued field, but also one should check the case of the residue field and x ranging over the ordered group. But this is easier, for instance, because algebraically closed fields admit elimination of quantifiers. ▲