Foundations of Computational Mathematics (2019) 19:991–1011
https://doi.org/10.1007/s10208-019-09416-0

FOUNDATIONS OF COMPUTATIONAL MATHEMATICS
The Journal of the Society for the Foundations of Computational Mathematics

Interpolation, the Rudimentary Geometry of Spaces of Lipschitz Functions, and Geometric Complexity

Shmuel Weinberger[1]

Received: 7 May 2018 / Revised: 7 June 2018 / Accepted: 12 December 2018 /
Published online: 9 May 2019

Abstract
We consider seriously the analogy between interpolation of nonlinear functions and manifold learning from samples, and examine the results of transferring ideas from each of these domains to the other. Illustrative examples are given in approximation theory, variational calculus (closed geodesics), and quantitative cobordism theory.

Keywords Interpolation · Function space · Persistent homology · Homotopy · Quantitative cobordism · Embedding

Mathematics Subject Classification Primary 68U05 · 57R75 · 41A10; Secondary 55Q05 · 68T10

1 Manifesto

There are fundamental connections and analogies between spaces and functions that are commonplace to the working mathematician—yet for some reason the study of the former is called analysis and the latter topology.[1] This division of intellectual labor is, of course, natural: a smooth function is a rather different kind of object than a smooth manifold. Nevertheless, we explore here some ideas that arise naturally when focusing our attention on similarities between these.

Perhaps the first example of this is the relation between a vector space V and its dual V^*, which consists of vectors and which of (linear) functions is entirely arbitrary in the finite-dimensional case. The use of Euclidean geometry uncritically in an infinite-

[1] Or geometry; I shall not distinguish between these.

Communicated by Martin Sombra.

Shmuel Weinberger: Partially supported by an NSF Grant.

✉ Shmuel Weinberger
shmuel@math.uchicago.edu

1 Department of Mathematics, University of Chicago, Chicago, IL 60637, USA

dimensional setting leads so naturally to one of the great genius ideas of the nineteenth century: Fourier series (and other orthogonal series).

A nonlinear version of this considers compact smooth hypersurfaces of, say, $\mathbf{R}^n$ and studies these via a smooth proper map $\phi\colon \mathbf{R}^n \to [-1, \infty)$, so that the hypersurface corresponds to $\phi^{-1}(0)$. (It is also common to dispense with the properness and replace the range by $[-1, 1]$; in that case one loses the compactness of the hypersurface.)

Doing this requires a basic theorem of topology guaranteeing the orientability of (closed) hypersurfaces in $\mathbf{R}^n$ (or more generally, in simply connected manifolds), or the Jordan separation theorem (which enables a concrete version of the orientation, by pointing outward away from a compact component, or toward infinity). A proof of this can also be given by a somewhat more sophisticated construction of either a map $\phi\colon \mathbf{R}^n \to \mathbf{RP}^k$ such that the transverse inverse image of $\mathbf{RP}^{k-1}$ is the submanifold (which captures near the submanifold the geometry of the twisting of the unit normal) and then extends to the complement by an essentially arbitrary map to the disk—and then letting k be large, and using elementary homotopy theory.[2]

We note the importance of considering functions with values in spaces other than $\mathbf{R}$ (or more general linear spaces). In other approaches this is perhaps less apparent: for learning a classification one will map a space X to say $\{-1, 1\}$ of labels, embed the latter in $\mathbf{R}$, use linear methods, and interpret the function values not in $\{-1, 1\}$ as being somewhat less certain. (For instance, one compresses f to $\text{sign}(f)$, and is more highly uncertain when f is near 0.)

Twentieth-century topology saw a similar exploitation of this connection. The Pontrjagin–Thom construction [57] is a generalization of the above to submanifolds of arbitrary codimension, but then the target is necessarily more complicated. This arises already in codimension two—the Klein bottle embeds in $\mathbf{R}^4$, and it is not a "smooth complete intersection." (See [29] for the application of this idea to manifold learning.)

Many of the successes of high-dimensional topology rest on the effectiveness of this methodology. One can call this the algebraicization of topology. This trend includes immersion theory and its subsequent developments [25,34,53] and surgery theory [9,60], our most powerful tool for the classification of manifolds. If one cares about the geometric nature of the solutions to these problems, barring a completely new approach to them,[3] one needs to understand the solutions of the problems of nonlinear algebraic topology.[4]

Our interest is in learning, describing, and understanding functions and manifolds. In order to do this, we need to have an understanding of their complexity, and their natures. Various notions of these will arise in this paper and will occasionally be compared. On the function side, we will typically use the Lipschitz constant as such a

[2] Alternatively, if a connected submanifold M didn't separate, one would construct a closed curve transverse to M lying in a neighborhood of M, and then argue that this would contradict the simple connectivity of the sphere (by homotopy invariance of intersection numbers).

[3] In some cases the connection between the homotopy theory and the geometric problem is sufficiently tight that *any* new geometric complexity statement would readily imply some, although perhaps not optimal, result about the complexity of homotopies.

[4] Many classical problems in analysis are also solved by homotopy methods. Typically, one starts, though, with an explicit homotopy to apply them to. I hope that the ideas discussed in later sections can be of use in that setting.

measure. On the (sub)manifold, we will need to control both size and local geometry. We begin the discussion in the next section. In all cases, for us a complexity notion should lead to a precompactness or a finiteness theorem.

In Sect. 2, we will consider linear interpolation and approximation of Lipschitz functions. This will lead to some simple ideas of approximation theory, manifold learning, and persistent homology, some of whose discussion will spill into Sect. 3. We also mention some subtleties that arise in the C^0 geometric theory that do not seem to have counterparts in the function setting. Section 4 is slightly more substantial and begins a study of PH (persistent homology) of the Lipschitz functional on the free loop space of a Riemannian manifold, with applications to closed geodesics and also some less familiar functionals. In the remaining sections, we consider a program of Gromov's to understand the complexity of geometric objects whose existence is predicted by topology and its relation to the geometry (as opposed to topology) of function spaces.

This paper is a survey of the ideas I discussed at the FOCM meeting in Barcelona 2017 and reflects collaborations I've had over the years with G. Chambers, S. Dranishnikov, D. Dotterer, S. Ferry, F. Manin, A. Nabutovsky, P. Niyogi, and S. Smale, and the conversations I've had with countless mathematicians during this time. Unattributed facts are either well known or appear in the work of Gromov, whose work has had too large an influence on my thinking for me to always be aware of it. As I write these lines, I realize that I had tried to compress too much material into a one hour lecture; I apologize to my hosts and the audience at the lecture on which this is based. Although there has been some substantial advances on topic related to the ones discussed here (notably [39], but also [43]) since the Barcelona meeting, I have chosen not to consider how they impact the picture.

And, lastly, I am happy to express my thanks to the referee whose comments have improved the quality of this paper.

2 Interpolation 101

We begin with the simplest problem. A function $f: X \to Y$ is "out there" in some sense, and we would like to "learn it" (see [19]). The simplest setting might be that we have access to function values for a set of points that we get to choose. Let's also assume that X and Y are "nice spaces"—later more specific hypotheses will be imposed.

Let's say that "learning" a function means being able to approximate it within ε (à la [59]; in particular, we'll assume that Y is metric, and we shall deal here with the sup norm on functions). With no information about X, Y, and f, this is impossible. Even if f is assumed continuous, it could oscillate at a frequency so rapid as to foil any such attempt. However, with control on the modulus of continuity, we can achieve this for any compact X: For definiteness, assume that f is K-Lipschitz, i.e., that

$$d(f(x_1), f(x_2)) \le K d(x_1, x_2)$$

(where d on the left is a measurement in Y and on the right is one in X).

In that case, to learn f one can merely choose an ε/K-dense set of points in X and record (accurately!) the function values at these points, and at any other point output the function value associated with the nearest point. Notice that the number of points that we need is the covering number of X at scale ε/K—which, if X is d-dimensional is $O((K/\varepsilon)^d)$.

Even the above simple observation has a useful implication.

Corollary 1 *If X is a compact d-dimensional space, and Y is compact, the covering number of the K-Lipschitz functions from X to Y at scale ε is $O(\exp(K/\varepsilon)^d)$.*

The base of the exponential has to do with the ε covering number of Y.

This solution is somewhat dissatisfactory, because it will never be right: we always[5] output a discontinuous function. It would be nice to be able to output a function f_n that is itself K'-Lipschitz for a constant K' that isn't much larger than K.

Interpolation is the simplest solution to this problem, but to do so inevitably requires a condition on Y. We shall assume, at first, the simplest possibility, that Y is a Riemannian manifold with positive convexity radius that is at least ε.

Recall that for a compact Riemannian manifold Y [22] there is an injectivity radius $\mathrm{inj}(Y)$, so that any two points of a distance at most inj have a unique geodesic connecting them. It also has a (typically smaller) convexity radius $\mathrm{conv}(Y)$, so that balls of size smaller than this are convex, i.e., contain the geodesic connecting any of these points.

Suppose now that the set in X that we choose forms the vertices of a triangulation (with none of the simplices terribly eccentric, so we can piece together estimates from one simplex to their union without much pain[6]). In that case, we can "connect the dots"; i.e., extend the map from the vertices of the simplex using iterated coning to define a map that is piecewise linear in barycentric coordinates.

In this situation, if two functions f and g are at most $\mathrm{conv}(Y)$ apart, then there is a natural homotopy between them that moves f to g pointwise along the geodesic that connects $f(x)$ to $g(x)$, and that stays within this class of functions, so the function space balls of this size are contractible. This means that $\mathrm{conv}(Y)$-close maps $X \to Y$ are homotopic.

Corollary 2 *The number of homotopy classes of maps $X \to Y$ with Lipschitz constant K is at most $O(\exp(K^d))$.*

For $d = 1$ this is the theory of the growth of the fundamental group of Y that is popular in geometric group theory (see, e.g., [33,39]). Growth is at most exponential (for finitely generated groups) but can be smaller in many interesting cases.

The fact that the function space balls are contractible, not merely connected, has a nice corollary.

Corollary 3 *The map from $K - \mathrm{Lip}(X : Y) \to C^0(X : Y)$ factors through a finite complex (with at most $O(\exp(K^d))$ vertices). In particular, the image in homology H_i of the former can be bounded by that dimension and vanishes for i high enough dimension.*

[5] Unless something extraordinary is happening—say f is locally constant.

[6] This can be achieved using either of the subdivision schemes of [23] or [30].

Frequently, however, the homology of the function space is nonzero in infinitely many dimensions. This produces interesting variational invariants, e.g., the first K for which one has i-dimensional homology—which can be thought of as a nonlinear analog of the spectrum of the Laplacian[7] (see [35]). In the next section, we will observe other "spectral invariants" within the homology of sublevel sets.

The above leads one to frequently adopt $\varepsilon = \text{conv}(Y)$ as a natural *scale* at which to study functions. At that scale, one can parametrize balls by picking a center f and then exponentiating vector fields along f whose uniform norms are at most conv. In other words, at this scale, we have a useful linear local coordinate chart.

We close this section with a few remarks:

Remark 1 One can give a similar analysis to any equicontinuous function space, e.g., the Hölder spaces $C^{1,\alpha}$ etc. The important thing is to have an explicit understanding of the relationship between ε and δ in the definition of continuity.

Remark 2 In order to do the connect the dots construction, one needs much less than a convexity radius, one needs a *local contractibility function* $\rho\colon [0, \varepsilon) \to [0, \infty)$ for Y, such that $\rho(r) \geq r$ (for $r < \varepsilon$) such that every ball of radius r is nullhomotopic in the concentric ball of radius $\rho(r)$. (Such a space is called $\text{LC}(\rho)$.) Note that one needs to apply this d-times if the domain is d-dimensional, so the scale at which one needs to work in this setting is much smaller than ε.

Remark 3 Some of these ideas arise naturally when thinking about geometric group theory and trying to approximate from a group Γ (or its Cayley graph) the universal cover of its classifying space $B\Gamma$, assuming that there is a compact model for the latter. (See [10,52].) In that case, the universal cover will be $\text{LC}(\rho)$ with ε infinite, but $\rho(0)$ positive. $\text{LC}(\rho)$ is a homotopical analogue of a, perhaps exotic, regularity condition on a function.

To approximate this universal cover, one covers the discrete metric space Γ by metric balls of size k and takes the nerve of this cover. Let's call this locally finite simplicial complex with cocompact Γ action $N_k(\Gamma)$. There are natural simplicial maps $N_k(\Gamma) \to N_l(\Gamma)$ if $l > k$; the sequence of locally finite homology groups of this sequence of spaces carries a lot of information about Γ: in particular, one can read off the cohomological dimension (when it's finite) by seeing in which dimensions the limit is nontrivial. However, there is a lot of interesting information in the sequence that is not in the limit, as we will comment in the next section.

Remark 4 These connect the dot ideas yield easily an analogue for metric spaces of the principle (used in Corollary 2) that close enough K-Lipschitz maps are homotopic. It is that in the Gromov–Hausdorff space of d-dimensional $\text{LC}(\rho)$ spaces, close enough spaces are homotopy equivalent. (See [28] for a deep analysis that goes considerably farther.)

Recall that Gromov–Hausdorff space (see [36]) is a metric space whose points are isometry classes of compact metric spaces. The distance between X and Y is the

[7] The first level for $E = \int\langle\nabla f, \nabla f\rangle / \int\langle f, f\rangle$ where i-dimensional mod 2 homology appears in the projective space of the Sobolev space $H^{1,2}(M)$ is in the ith positive eigenvalue of the Laplacian.

smallest d so that in some metric space containing X and Y isometrically, X and Y are in the d neighborhoods of each other.

Close enough spaces can be mapped to each other by matching a pair of sufficiently dense finite subsets of X and Y, and connecting the dots, and the composites can be checked to be homotopic to the identity by a similar induction. The closeness of the metric spaces that is sufficient to give the homotopy equivalence can be determined explicitly in terms of ρ, d, and ε.

In [21] we study the extent to which something similar holds for homeomorphism for manifolds. Note that if M is a manifold, like the sphere, for which any homotopy equivalent manifold is homeomorphic to it, then in any LC(ρ) manifold close enough to any such metric M will be homeomorphic to it. We show (that aside from dimension 4) this is true for all manifolds whose universal covers are the sphere or Euclidean space (although at least in the former case there are many manifolds of homotopy equivalent to M but not homeomorphic to it—see [60]). But that it is not true in general. One cannot predict from ρ, d, and ε how close general manifolds have to be to be homeomorphic.

Let me phrase this somewhat differently. If M is a compact Riemannian manifold, then there is an ε_M, so that any LC(ρ) manifold within ε_M of M is homeomorphic to M. *However,* this ε_M is *not* bounded below on a precompact set of manifolds in the Gromov–Hausdorff space of Riemannian LC(ρ) manifolds. (However, for some manifolds, the issue does *not* arise, for deep surgery theoretic reasons.)

This suggests an obstruction to C^0 sample reconstruction theorems in the LC(ρ) setting.[8] We remark that for some ρ (e.g., linear) this phenomenon does *not* occur, and one can hope to try to learn a manifold with such a metric from a sufficiently densely sampled point cloud.

In the C^2 setting, this is indeed possible and is the subject of a rather large literature that we can just indicate (see, e.g., [2,4,5,12,29,46]). For simplicity, let's work in the Euclidean setting and consider manifolds with *positive reach* τ. This means that the τ-tubular neighborhood of M is embedded, i.e., that any point of Euclidean space within τ of M is nearest to a unique point of M. This condition reflects both local (i.e., the curvature of the second fundamental form) and global properties of M (e.g., preventing nearby concentric spheres), see, e.g., [46]. It is a regularity condition on manifolds like LC(ρ) where ρ is the identity on $[0, \tau)$.

The paper most in the spirit of this section is [29] that uses deep interpolation ideas for smooth functions (the Whitney problem, see, e.g., [26,27]) to describe how to interpolate a manifold through a set of noisy samples of a manifold with positive reach (with high probability) when it's possible.

3 Persistent Homology

In the previous section, we described the idea of interpolation as connecting the dots and filling in high-dimensional simplicies of either a graph of a function (implicitly

[8] The reader could well want to know what can distinguish manifolds which mutually approach one another. The simplest invariants are odd primary characteristic classes of the topological tangent bundle. However, there are additional subtle secondary invariants that are torsion analogues of the invariants used by Atiyah and Bott [1] to distinguish lens spaces from one another.

defining an approximation to the function) or a manifold based on the proximity of points to one another.

However, there is a simpler method for getting homotopical information (e.g., the homology) as alluded to in Remark 3. We refer to [18,24,48,50,62] for more discussion.

Imagine a point cloud, i.e., an abstract discrete metric space. Suppose we know nothing about the scale at which the sampling took place, or, for example, if it's LC(ρ) which set of interlocking scales one should use for the purpose of inference. Surely we don't have to just do nothing?

For the infinite metric space, we formed a sequence of nerves of coverings (increasing our scales at each moment) and took a limit (feeling confident, since we were approximating a hypothetical contractible universal cover[9]). At the very foundations of (co)homology theory, the Čech theory is obtained by a similar process using limits with respect to refinements (i.e., resampling with a denser set of points and adjusting to use the nerve at a smaller scale) rather than with respect to coarsening.

So, we consider the sequence of nerves, *but don't take the limit*. We then have a sequence of spaces whose homotopy invariants (most simply the homology) we can study as a sequence, keeping track of when things are born and when they die.

This theory works best over a field. Consider a set that contains 3 points. According to the triangle inequality, this metric space is actually isometric to a (perhaps degenerate) triangle. Let's consider what happens when we take the nerves at various scales. On H_0 we start at $\varepsilon = 0$ with 3 generators, say $[p]$, $[q]$, and $[r]$, and this is the case until ε is the smallest distance between points. At that point, we decide that a better generating set would be $[p]$, $[p] - [q]$, $[r]$, where the components of p and q merge. Note the basis element $[p]$ isn't quite well defined, but $[p] - [q]$ is. When we hit the next larger distance, $[r] - [p]$ or $[r] - [q]$ is killed. At this point the nerve is connected and nothing else ever gives us (more or) fewer components. So we can think of this module as being a sum of 3 modules, all born at $\varepsilon = 0$, one of infinite length, and two of each of the shortest two legs of the triangle.

If there are only three points, then 1-dimensional homology never forms because by definition when the three edges are in the complex, so is the 2-simplex that bounds this. However, if we had, say 4 points isometric to the vertices of a square, then H^1 would be formed when ε is the shortest distance between points, but it would die when ε is the larger diagonal distance. (At that point a 2-sphere would form, but it bounds a 3-simplex instantaneously).

These data are encoded in a «bar code» or a «persistence diagram» (that is a collection of points in the first quadrant, above the diagonal, recording when homology generators are born and when they die.)

In quite great generality, if points are sampled densely enough at a scale r at which the space is locally contractible, then the homology that is born before $r/2$ that survives through scale r is isomorphic to the homology of the space.

An interesting situation happens if one samples *very* densely a torus $\mathbf{R}^2/(1000\mathbf{Z} \oplus .001\mathbf{Z})$. This is simply the rectangular torus that's a product of two circles—one of length 1000 and one of length $1/1000$. The 0-dimensional homology will have some

[9] The process described, though, is useful even when $B\Gamma$ does not exist as a finite complex. The homology described can be used as a substitute that occasionally has more useful properties. See [20] for an example of this.

very tiny bars (as we've assumed the sample is *very* dense), but will quickly stabilize to the «right answer» . However, H_1 will have 2 interesting large bars, (i.e., the bars that are topological features of this torus) being born around the density of the sample and surviving until around .001, but then one direction will collapse. This homology class was honestly part of the space, but it dies anyway. (It's genuineness was noted by the fact that it gave a quite long bar: at least the bar is quite long if one sampled at a .00001 scale!) The other class will die as well at scale 1000.

The Persistent Homology in this case is measuring interesting geometry and telling us that at a very fine scale we have a torus, but from a larger scale it is more like a circle.

The H_0 story is highly relevant to the problem of «clustering» . We have data and want to lump them together into various populations. There is no best way to do this (see [38]), but as argued in [15], PH_0 gives a useful tool to come up with a set of possible clusterings that can each be of value for different purposes.[10] This, too, can be indicative of the multiscale geometry that a (data) set might have.

PH formally speaking just requires a real-valued filtration of a space. Then associated with ε, one assigns the homology of the part of the space with filtration $< \varepsilon$. This enables many variations. For example, if we had a point set in Euclidean space, rather than just an abstract one, one could filter Euclidean space by the distance to the point set rather than considering the formal nerves. This choice is sometimes referred to in the literature as choosing the Čech complex over the Rips complex. Each has its advantages.

A good way to get a feeling for persistent homology is to consider a Morse function on a manifold (see [41]) whose basic object of study is the homology of sublevel sets of (generic) real-valued smooth functions on manifolds. The main result is that *each critical point of index k creates an endpoint of a persistence interval: either it begins a k-dimensional homology bar, or marks the end of a $(k-1)$-dimensional homology bar (and conversely, every beginning and end of a bar is due to a critical point of the required index)*.

Ordinarily, people tend to record Morse theoretical conclusions by a series of inequalities relating the ordinary homology of M (the infinite-length bars) to the critical points, but in PH the relation is as simple as possible. We shall soon see that sometimes there's something topological even in finite length bars. For the meantime, we will use the notation $\mathrm{PH}_i(f)$ for the ith Persistent Homology of a function $f: M \to R$. Usually our functions will be bounded below (a real condition when M is noncompact[11] as it will be in some later examples), so that none of the intervals start at $-\infty$, but this is not essential.

We note an important *stability* property of $\mathrm{PH}(f)$. Suppose that $|f-g| < \varepsilon$. Then for every C,

$$g^{-1}((-\infty, C-\varepsilon)) \subset f^{-1}((-\infty, C)) \subset g^{-1}((-\infty, C+\varepsilon)).$$

This interleaves the homology of the sublevel sets of f and g. As a result the barcode of f can only change in the following ways through a small C^0 perturbation:

[10] We shall not discuss here what to do about noise or when populations have overlapping images.

[11] Of course, in the noncompact case, Morse theory in its ordinary sense requires a properness condition, or a Palais–Smale condition in the infinite-dimensional setting.

1. Short bars of length at most 2ε can be created.
2. Such short bars can be removed.
3. The top and/or bottom of a long bar can be moved by up to ε.

The above is an inelegant formulation of the following theorem of [17]:

Theorem (Stability Theorem) *If* $f, g\colon M \to \mathbf{R}$ *then* $d_B(\mathrm{PH}(f), \mathrm{PH}(g)) \leq ||f - g||_{C^0}$

where the *bottleneck distance* used in the left hand side is the inf over all one-to-one correspondences between the bars in the persistence diagram of f and of g of the amount that an endpoint has to be moved, allowing for the insertion and deletion at will of bars of length 0.

For example, any perturbation of $\sin(x)$ on $\mathbf{R}$ of size < 1 must still have infinitely many local minima and maxima, or even any bounded size permutation of $(x \sin x)^2$ must have infinitely many local minima and maxima (and they will be spaced fairly densely in $\mathbf{R}$).

Returning now to another of our motivating examples, if one considers a smooth manifold X in Euclidean space, and Y a sufficiently dense sample (say a finite subset), and perhaps chosen in slightly noisy fashion, the functions $\mathrm{dist}(-, X)$ and $\mathrm{dist}(-, Y)$ will be C^0-close. Thus, the intervals longer than the density and the noise will correspond. If the reach is larger than this, then one can infer the homology of X by considering Y.

Implicit in this theorem is a remarkable C^0-*semicontinuity of critical values*. Defining the index of critical points requires C^2 topology, and Morse theory is C^2-stable. (The Morse condition is that the Hessian of mixed second order derivatives is nonsingular.) However, because of the connection to PH, any C^0-close function will necessarily have at least as many critical points of each index.

4 Simple Applications

We now shall try to apply the ideas of the previous section to function spaces. But first let's discuss some very classical examples.[12]

Suppose I ask you to approximate $\sin(nx)$ on $[0, 2\pi]$ by lower order trigonometric functions, then you'd ask me to pick a norm. If I pick L^2 the answer is trivial: since $\sin(nx)$ is orthogonal to that space, the closest point to it is 0 and that is the best approximation. But, if I ask you for the C^0-closest approximation, that's trickier, but the answer is the same. Why?

Consider what $\mathrm{PH}_0(f)$ looks like for a degree k trigonometric polynomial. It has a bunch of bars going from the local minima to various local maxima.[13] The number of these is $\frac{1}{2}$ the number of critical points.

The critical points are the zeroes of a degree k trigonometric polynomial (i.e., the zeroes), which are the roots of a polynomial of degree k: Think of $2\sin$ and $2\cos$ as $z - z^{-1}$ and $z + z^{-1}$ on the circle and z^k times the trigonometric polynomial will

[12] Other applications to approximation theory can be found in [49].

[13] Except the global maximum, which doesn't close any H_0 bar. Thought of as being the circle, it begins the nontrivial H_1 bar.

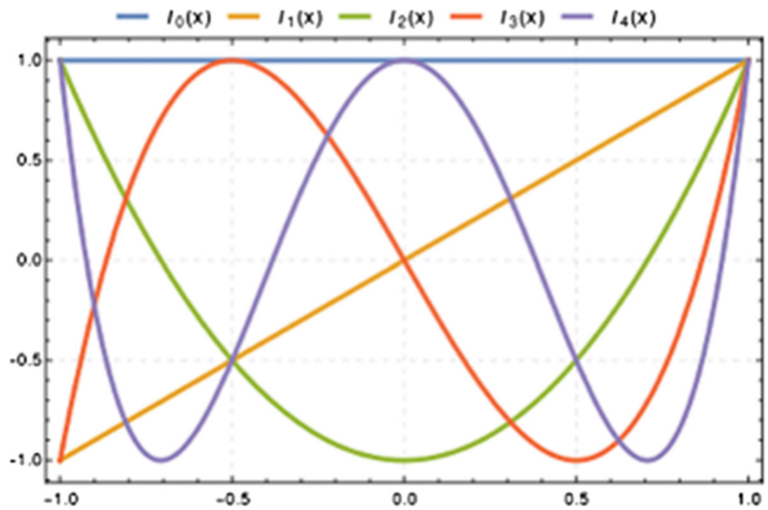

Fig. 1 The Chebyshev polynomials

be a complex degree $2k$ polynomial with at most the number of roots as the original trigonometric polynomial. So there are at most k bars.

Since $\sin(nx)$ has n bars, all of length 2, the smallest the bottleneck distance to any lower-degree trigonometric polynomial must be at least 1, giving the result. Note by the way that unlike standard approximation (or Fourier) ideas this also immediately gives the C^0 distance from $\sin(nx)$ to even «reparametrized» degree $(n-1)$ trigonometric polynomials: composing T_{n-1} with a homeomorphism of $[0, 2\pi]$ to itself will not change the barcode in the slightest.

This is a straightforward argument, but here's a consequence that is not as transparent. The distance of $5\sin(x) + \sin(2x)$ to the degree-one trigonometric polynomials is the same as that of $\sin(2x)$, but the barcode of this function has only one bar! The proof fails for this obviously equivalent problem.

The same argument now applies to estimate on $[-1, 1]$ the distance of x^n to the lower-degree polynomials. It has a boring barcode, but instead we realize that this problem should be done using orthogonal polynomials instead. The Chebyshev polynomials $P(x) = \cos(n\cos^{-1}(x))$ are a good choice. (See Fig. 1.) Note there are n bars for this all going from -1 to 1 (aside from the 1 infinite bar). So the distance one sees from the Chebyshev polynomial to the lower-degree polynomials is 1, and since the coefficient of x^n in $P_n(x)$ is clearly 2^{n-1} that the distance from x^n to the lower-degree polynomials is $2^{-(n-1)}$.

Let's turn our attention to functions on function spaces, following an example sketched in [62] explaining a theorem of Gromov. It transpires that PH is an excellent language for doing «large scale Morse theory» .

Let M be a compact Riemannian manifold, and let $\Lambda M = C^2(S^1 : M)$ be the space of smooth loops on M. We shall consider the function $E\colon \Lambda M \to \mathbf{R}$ given by $\log(\int\langle f', f'\rangle)$ where $\langle, \rangle$ is the Riemannian inner product on M. Notice that changing the inner product on a compact Riemannian manifold can only change E by a uniformly finite amount (log of the ratio of the two inner products). As a result, PH(E) is well defined as a set of bars up to finite bottleneck distance.

The bars of infinite length coming from $-\infty$ correspond to the homology of M, thought of as the constant loops. (These last forever because there is an evaluation map $\Lambda M \to M$.) The ones lasting to $+\infty$ correspond to the topological invariant $H_i(\Lambda M)$.

Much more interesting are finite length bars. What does a bar of length K in PH_0 signify? At the bottom of this bar will be a closed geodesic of some energy e^E. When we try to move this curve to lower the energy (i.e., to connect to a component of E^{-1} of a lower level), we ultimately succeed in doing so, but to do it, we need to multiply our energy by e^K to do so. By restricting to the component of nulhomotopic loops $\Lambda_e M$, we obtain the following:

Proposition *The property of whether all nullhomotopic geodesics on M can be nulhomotoped through curves of length that is only a linear multiple of their lengths, is independent of the metric: it is equivalent to* $\mathrm{PH}_0(E\colon \Lambda_e M \to R)$ *having only bounded bars (aside from the one infinite bar* $[-\infty, \infty]$*).*

All of this can be done with an arbitrary polyhedron as well, with slightly different details. Moreover, with a little thought, one can see that PH_0 (up to finite bottleneck distance) only depends on $\pi_1(M)$.

Proposition *If* $\pi_1(M)$ *has superexponential Dehn function,*[14] *then* PH_0 *has infinitely many bars of arbitrarily large finite length.*

If there were a bound on the bar length, then one could find a C so that every nulhomotopic curve of length L can be homotoped to one of length $L/2$ through curves of length at most CL. The entropy argument in Sect. 2 shows that the area swept out by this annulus (i.e., this path of curves) is at most $\exp(CL)$. As a result, ultimately all nullhomotopic curves of length L will bound a disk of at most exponential in L area proving the theorem.

Indeed, if $D_\pi(L)$ is a superexponential Dehn function of a group, then there will be infinitely many closed nullhomotopic geodesics on any closed Riemannian manifold with fundamental group π that will require multiplying their length by $\log D_\pi(L)/L$ to nullhomotope. Thus, groups with radically different Dehn functions will have quite different barcodes for their PH_0's.

Many other functionals on spaces of Riemannian metrics or function spaces, especially with nonsimply connected targets tend to have extremely long finite length bars. For example, the result of Nabutovsky [45] can be expressed as follows:

Proposition *Let* HS $=$ *the space of unparametrized smooth embedded* S^n*'s in the unit ball* B^{n+1}*, then although* HS *is connected,* $\mathrm{PH}_0(1/\tau\colon \mathrm{HS} \to \mathbf{R})$ *(recall that* τ *denotes the reach of a submanifold) has infinitely many bars of finite length. Given L, the function that gives the largest length bar going through level L cannot be bounded by any computable function.*

See, e.g., [45,61] for surveys of such ideas and our joint paper [47] for applications to Riemannian variational problems. It seems that this gives a method for detecting

[14] See the proof for a definition, or see, e.g., the Wikipedia page on Dehn functions of groups.

the existence of a rugged Morse landscape with more critical points being imposed by "texture" than forced by topology.[15]

In the case of simply connected function spaces, the kinds of phenomena indicated in the previous proposition cannot hold: because of the work of Brown [8] bars cannot be uncomputably long. The following seems just perhaps conceivable (as indicated in our discussions below). It is closely related to the ideas in [37].

Conjecture *Suppose X and Y are finite complexes, Y is simply connected, then*

$$\mathrm{PH}_i(\log(L)\colon \mathrm{Lip}(X:Y) \to \mathbf{R})$$

has only finite length bars.

All bars for PH_0 have at most linear length, when length is plotted as a function of their bottom value, as a consequence of the important recent paper [39].

5 Unoriented Cobordism

Having turned interpolation into a tool for understanding geometric sampling, and used PH as a way to use that idea to understand the homology of the underlying space, we turned around and considered the application of these methods to understanding functions.

In this section, we will review a basic example of the algebraicization of geometric problems, the cobordism problem, and give some recent results on the complexity of nullcobordism, and its connection to the conjecture made at the end of the previous section.

We recall Thom's theorem [55,57].

Theorem *A smooth closed manifold M^n is the boundary of some compact W^{n+1} iff the cycle represented by M in H_n(Grassmanian of $(n+1)$-planes in $\mathbf{R}^{N+1}$) is trivial.*

More precisely, smoothly embed M in $\mathbf{R}^N$ for N sufficiently large and then at each point m the tangent plane $T_m M$ is an element of the Grassmanian of n-planes in $\mathbf{R}^N$. Using this function, the image of M is a cycle in H_n(Grassmanian of n-planes in $\mathbf{R}^N$). We just add on a trivial last direction to consider it in H_n(Grassmanian of $(n+1)$-planes in $\mathbf{R}^{N+1}$).

Once N is larger than $2n+1$, this class is independent of the embedding by standard differential topological arguments. If M were a boundary then the cycle it defined would bound (in H_n(Grassmanian of $(n+1)$-planes in $\mathbf{R}^{N+1}$)) the chain represented by W.

That the converse holds is the result of deep geometric and algebraic topological arguments.

[15] Here I mean a large scale idea of texture that is analogous to the differences that one would notice on the small scale if one took the graphs of (Baire generic) Hölder functions for different exponents. If you consider the persistence homology, the Hölder exponent would be apparent in the "length spectrum of the bars" as will be explained in detail in a future paper with Yuliy Baryshnikov (and is not hard to see by wrinkling a map by putting in as many bumps at small scale as permitted by dimension and the Hölder condition).

Remark The same theorem holds (using oriented Grassmanians) for oriented manifolds and oriented as a result of subsequent deep work of Milnor, Averbuch, Novikov and Wall, see [55]. The geometric part is the same, but the algebraic topology is much more difficult. We shall see that the quantitative aspects of these theorems also have differences in their level of depth.

Gromov [36,37] pointed out that Thom's theorem raises a very natural quantitative problem. Let's define the *complexity* of M, $k(M)$ to be $\inf(\mathrm{Vol}(M, g))$ as g ranges over all Riemannian metrics with sectional curvatures pointwise bounded by 1 (in absolute value) and with inj ≥ 1 and, if there is a boundary, we assume that it has a neighborhood isometric to $\partial \times [0, 1]$, to avoid any local complication near the boundary either in the interior or near the boundary. Now, suppose that M is a boundary (as Thom tells us when that is), what is $\inf(k(W))$ as W ranges over manifolds with $\partial W = M$?

Thom's proof does not directly give any estimates at all; Gromov [37] sketched an argument to deduce from Thom's work a bound that is a tower of exponentials

$$k(W) \leq \exp(\exp(\exp(\dots(k(M))\dots))),$$

where the number of exponentials grows with n—but he speculated that the truth was linear[16]! This is still unknown, but the work of [11,30] goes a considerable distance in that direction.

The reduction to algebraic topology goes like this. Start by embedding M in S^N for N large. The normal bundle ν of M is pulled back from the universal bundle ξ of $N-n$ planes in $\mathbf{R}^{N+1}$. By the tubular neighborhood theorem, there is an ε-neighborhood of M that has a map to the unit disk bundle over the Grassmanian. This map sends the set of points of distance ε to the unit sphere bundle. The map

$$(N_\varepsilon(M), \partial N_\varepsilon(M)) \to (D(\xi), S(\xi))$$

gives rise to a map $\varphi_M : S^N \to D(\xi)/S(\xi)$, the space obtained by identifying the whole sphere bundle with a point, and which we shall denote by $\mathrm{Th}(\xi)$, simply by mapping all of the complement of $N_\varepsilon(M)$ to the point that $S(\xi)$ was identified with.

Lemma $\varphi_M : S^N \to \mathrm{Th}(\xi)$ *is nullhomotopic iff* M *is a boundary.*

To the modern mathematician, the proof is almost obvious: $\mathrm{Th}(\xi)$ is a manifold with a single singularity at the point $* = S(\xi)$. The Grassmanian in $\mathrm{Th}(\xi)$ is entirely in the manifold part. Ignoring this singularity, $\varphi_M : S^N \to \mathrm{Th}(\xi)$ is smooth wherever this makes sense and is transverse to the Grassmanian. (Thom had to define transversality in his paper, though, and establish its properties.) If M bounded W, we'd embed W in $S^N \times [0, 1]$ and repeat his construction making a homotopy of φ_M to the constant map $*$.

[16] This linearity conjecture implies an estimate for η-invariants of manifolds (spectral invariants of odd-dimensional manifolds defined by Atiyah–Patodi–Singer) in terms of volume; Cheeger and Gromov proved these directly [13]. While the quantitative cobordism work we discuss is not strong enough to give the Cheeger–Gromov inequality, a different approach—with very interesting complexity applications—was given by [14].

Conversely, if φ_M were nullhomotopic, we'd smoothly approximate the nullhomotopy (rel φ_M where things are good enough already) and then take the transverse inverse image of the Grassmanian. The inverse function theorem would tell us that this is a smooth manifold bounding M.

The rest of the proof of Thom's theorem goes by relating the condition in the lemma to the one in the statement of the theorem. Aside from a minor point relating Grassmanians associated with tangent and normal bundles, and the use of the Thom isomorphism theorem (relating the homology of $\mathrm{Th}(\xi)$ to that of a Grassmanian) this boils down to showing that in this situation a Hurewicz homomorphism (see, e.g., [54]) is essentially an injection.

Recall that for any X, the Hurewicz map is the map $\pi_n(X) \to H_n(X)$, sending a homotopy class to the image of $[S^n]$ in $H_n(X)$. For simply connected X (and $\mathrm{Th}(\xi)$ is simply connected), the Hurewicz theorem says that this map is an isomorphism for the first n for which either group is nonzero. The problem is that in the cobordism problem this hypothesis does not hold except for the case relevant to 0-manifolds.

In algebraic topology, the method around this is called the method of killing homotopy groups. If X is $(n-1)$-connected, then we can learn $\pi_n(X)$ by computing $H_n(X)$. Now one builds an infinite-dimensional space $K(\pi_n(X), n)$ which has just one nonzero homotopy group, and up to homotopy makes X a fiber bundle over this space. The fiber F of this map $X \to K(\pi_n(X), n)$ is n-connected, so to compute $\pi_{n+1}(F)$ one can compute $H_{n+1}(F)$—this is frequently done using spectral sequences or other algebraic tools. Note how all the geometry has flown away: only to an algebraic topologist is $K(\pi_n(X), n)$ a simple space. The map $X \to K(\pi_n(X), n)$ also entails significant geometric cost. General theory shows that $\pi_{n+1}(F)$ is isomorphic to $\pi_{n+1}(X)$ and one can try to repeat the process to learn $\pi_{n+2}(F) = \pi_{n+2}(X)$ etc.

There are a number of points to consider in trying to make this type of reasoning effective:

(1) We need to relate the geometric complexity of M to some kind of analytic complexity of φ_M.
(2) We need to learn something about the complexity of the nullcobordism from that of the map.
(3) We need to infer from the analytic complexity of the nullhomotopy, geometric information about the nullcobordism.

Regarding (1) it should not surprise the reader that the focus is on Lipschitz constants. Note that the proof requires a geometric act of violence: embedding M into a sphere.[17] This must cause distortion but we can try to avoid doing too badly, see [7,31,40] for a beginning. (Metric embedding is a major subject with deep applications in theoretical computer science.) Note that the Lipschitz constant of φ_M comes from the motion of the normal spaces around M, i.e., from curvature, and also the ε which is essentially the reach of the image of M. If we make M bigger to lower curvature, it will fold back to itself and make the reach small. We will not discuss the embedding aspects since they are complementary to the concerns of this paper.

[17] Actually, there is a beautiful proof of Thom's theorem [6] that avoids this. This method actually directly leads to a polynomial estimate (of degree 2^n) for $k(W)$. Unlike what we are about to discuss, this does not extend to the oriented case.

(3) in general feels like it should be difficult, and involves the issues considered by Yomdin [63]. However, that work deals with the worst case estimates for geometry of fibers. Happily in this situation, one can deal with average case (or even best case) scenarios and the complexity estimates for this last part are dealt with by ideas related to the simplicial approximation theorem.[18]

(2) is at the core of our concerns. Essentially Gromov's analysis was that each step in the method of killing homotopy groups costs an exponential. It leads to a proof of the conjecture made in the previous section with a tower of exponentials for the length of the bars. For our purposes, we just are interested in PH_0.

If we had a bounded bar for PH_0, then an L-Lipschitz map would be homotopic through CL-Lipschitz maps to a constant. However, we need the Lipschitz map of the nullhomotopy to be able to use it to make the Thom lemma effective.

The general methods about covering numbers from Sect. 2 could be applied to give a Lipschitz constant of $\exp(LN)$ for this nullhomotopy, which would then get rid of most of the tower of exponentials. *However,* part of the emerging story is that there is better information available than this. For now, let's formulate the following:

Question *Suppose Y is a finite complex, and so is X. Is there a C, so that if $f, g\colon X \to Y$ are homotopic L-Lipschitz maps, there is a CL-Lipschitz homotopy $F\colon X \times [0, 1] \to Y$ between them?*

Of course, we rarely expect a positive answer to this question when Y is not simply connected, although it **is** OK when Y has nonpositive curvature. It is not true in general even when Y is simply connected.

Theorem ([30]) *The answer to the above question is affirmative with a constant C that only depends on* $\dim(X)$ *(for X's with nice local structure) iff all the homotopy groups of Y are finite.*

The condition on Y can be checked. It requires that $\pi_1 Y$ is finite (which can be verified, but cannot be decided) and then that the reduced $\mathbf{Q}$-homology of its universal cover vanishes. In particular it holds for $\mathrm{Th}(\xi)$ and the above theorem implies a polynomial bound in the unoriented case, when combined with efficient treatments of (1) and (3).

Unfortunately the theorem contains its own obstacle for dealing with the oriented case. Since no multiple of $\mathbf{CP}^2$ is an oriented boundary, there is an element of infinite order. We have no way of producing Lipschitz nullhomotopies from the [30] method unless the geometry of X is almost irrelevant.

It is instructive to see what happens for $X = B^n(R)$ a ball of radius R, and $Y = S^n$. Of course, with this X all maps from X are nullhomotopic. However, the nullhomotopy this gives will have proportionality constant around R.

And that is the correct estimate. This is a straightforward Stokes' theorem argument about φ dVol and Φ dVol for a nullhomotopy, assuming CL-Lipschitz, where $\varphi\colon X \to Y$ is the composite of maps $B(R) \to B(1) \to T^n \to T^n \to S^n$ where $B(R) \to B(1)$ is multiplication by $1/R$, $B(1) \to T^n$ is a standard map which is onto via going onto

[18] To see the issue, the number of point inverses for a Lipschitz map from S^1 to itself can we be infinite, even with $L = 2$. However, the degree of a Lipschitz map is bounded by L and one can C^0-approximate such a map by one where the number inverse images grows linearly with L.

a fundamental domain, $T^n \to T^n$ is multiplication by $\lceil LR \rceil$ (greatest integer), and $T^n \to S^n$ is a standard degree-one map.

6 Rational Homotopy Theory

Let's reformulate the theorem adumbrated in the previous section: A finite complex Y has the property that all maps from finite-dimensional X into Y that are Lipschitz and homotopic are Lipschitz homotopic iff all the homotopy groups of Y are finite. (The independence of the Lipschitz constant on the geometry of X beyond dimension enables one to use an Arzela–Ascoli argument and see that theorem for finite complexes implies the one for infinite.)

However, much more critical is what happens if X is a fixed finite complex, and we allow the dependence of the Lipschitz constant for the homotopy to involve the geometry of X.

That this is necessary to make progress is clear. Indeed, the discussion at the end of the last section even suggests that a key issue is the nature of the kinds of isoperimetric inequalities that hold on X: the ball's boundary having $o(\mathrm{Vol(Ball)})$ was the key to the construction.

The following example from [11] though shows that even this weakening is impossible to accomplish in general. To explain the example, we need a little bit of notation.[19] Consider the manifold $S^n \times S^m$ with the Morse function $||u||^2 + ||v||^2$ (with obvious notation). It has 4 critical points, with Morse index 0, m, n, and $m+n$. As a result, the boundary of the tubular neighborhood maps to the next sublevel set, clearly homotopy equivalent to $S^n \vee S^m$. We call this map

$$\Phi_{m,n} \colon S^{n+m-1} \to S^n \vee S^m.$$

Let $L \colon S^n \to S^n$ be the L-Lipschitz map of degree L^n. The map $\Psi_L \colon S^{3n-2} \to S^n \vee S^n \vee S^n$ obtained by composing $\Phi_{n,2n-1}(L \vee \Phi_{n,n}(L \vee L))$ is L-Lipschitz and is $L^{3n}\Psi_1$ as an element of $\pi_{3n-2}(S^n \vee S^n \vee S^n)$. Consider the space $Y = (S^n \vee S^n \vee S^n) \vee (S^n \vee S^n \vee S^n) \cup e^{3n-1}$ where we attach a $(3n-1)$-cell along $\Psi_1 - \Psi_1'$ where the $'$ denotes the same map with values in the 2nd wedge of n-spheres.

Proposition *$\Psi_L - \Psi_L' \colon S^{3n-2} \to Y$ is L-Lipschitz and nullhomotopic, but the Lipschitz constant of any nullhomotopy is at least $L^{3n/(3n-1)}$.*

The idea is that one uses the relation $\Psi_1 - \Psi_1' = 0$ L^{3n} times and that is mediated over the space $S^{3n-2} \times [0, 1]$ of dimension $3n - 1$. The details are rather similar to the argument for the $B^n \to S^n$ example in the previous section.

These examples (and [32]) suggest an important role for rational homotopy theory in these problems. (We note that Sullivan [56] has given an algebraicization of rational homotopy theory using differential forms, and their algebraic analogue, commutative differential graded algebras.) As that theory is simplest for simply connected spaces,

[19] The construction we are about to explain is called the Whitehead product in homotopy theory. (See, e.g., [54].)

we'll henceforth restrict attention to the simply connected case. The result of [30] gives boundedness (independently of L) of the diameter of the L-Lipschitz maps in the CL-Lipschitz maps iff Y has the rational homotopy type of a point. On the other hand, the arguments using differential forms shows that the rational homotopy type of $Y = (S^n \vee S^n \vee S^n) \vee (S^n \vee S^n \vee S^n) \cup e^{3n-1}$ doesn't have linear size homotopies.

The following is an example of the type of reduction to rational homotopy theory that we currently have. (Of course, the difficult part of the [30] theorem would be a trivial consequence of rational homotopy invariance: when Y is a point, one has no trouble getting estimates in the function space.)

Theorem ([16]) *Rationally equivalent simply connected finite simplicial complexes admit nullhomotopies of the same shapes. That is, suppose we are given the following data:*

(1) *Rationally homotopy equivalent simply connected finite metric simplicial complexes Y and Y;*
(2) *A finite n-dimensional simplicial complex X;*
(3) *A simplicial pair $(K, X \times ([0, 1] \cup [2, 3]))$ which is homeomorphic to*

$$(X \times [0, 3], X \times ([0, 1] \cup [2, 3]))$$

and given the standard metric on simplicies. Here the product of X with each unit interval is given an arbitrary fixed simplicial structure which restricts at $t = 0$ and $t = 1$ to the simplicial structure on X.

Then there is a constant $C = C(X, Y, Z) > 0$ such that if for every nullhomotopic L-Lipschitz map $f : X \to Y$ there is an M-Lipschitz homotopy $F : K \to Y$, then for every L/C-Lipschitz map $g : X \to Z$ there is a $(CM + C)$-Lipschitz nullhomotopy $G : K \to Z$.

The awkwardness of the precise theorem is due to the fact that rationally equivalent spaces need not have maps to one another, but they will both have maps into a third space that induces the equivalence.

In [11,16] we study some special rational homotopy types (and [39] goes much deeper) and see that for a fixed X, linear (or occasionally higher-degree polynomial) homotopies can be constructed. In particular, we have the following theorem.

Theorem ([11]) *If Y is a finite simply connected complex whose rational cohomology is that of a product of odd-dimensional spheres, then for any finite complex X, there is a C, so that if $f, g : X \to Y$ are homotopic L-Lipschitz maps, then they are CL-Lipschitz homotopic.*

Remark This theorem is of a different nature than the [30] theorem. (1) In that theorem one can actually build a map that is CL-Lipschitz in the X direction and 1-Lipschitz in the [0, 1] direction. And (2), in the [30] situation one can homotope an arbitrary homotopy between f and g to one that has small Lipschitz constant, but in the situation of this theorem, that is not possible, as one quickly sees if $Y = S^n$ and $X = S^{n-1}$. In some sense, to prove such theorems it is important to first understand what the homotopy does at the level of the differential forms, and then try to produce a genuine geometric homotopy that resembles it.

As before, following from Thom's ideas this then gives (using only the simple embedding ideas), the following quantitative estimate:

Corollary *If M is a closed (oriented) manifold with complexity $k(M)$ that bounds a(n oriented) compact manifold, then there is such a compact manifold W with $k(W) = O(k(M)^c)$ for some c (depending on n).*

Remark In [43] this is improved to c any arbitrary number larger than 1. (Of course, the implicit constant in the «O» will depend on the dimension.) Here the main tool is to produce embeddings for which the Thom map has quite small Lipschitz constant. Interestingly, the best thing to do is *not* to make the dimension of the ambient sphere around $2n + 1$ (à la Whitney) but to pick a high multiple of n. Doing this lowers distortion. (Although raising dimension increases the complexity of inverse images, [43] deals with the tension between these effects.)

Given the ubiquity of the use of Johnson–Lindenstrauss ideas in geometric algorithms, the use of two scales in [43], taken from [31] could be of broader applicability. We just state the final Gromov–Guth–Whitney embedding theorem:

Theorem *Let M be a closed Riemannian m-manifold with complexity V. Then for every $n \geq 2m + 1$, there is a smooth embedding $g\colon M \to \mathbf{R}^n$ such that*

- *$g(M)$ is contained in a ball of radius $R = C(m, n)V^{1/(n-m)}(\log V)^{2m+2}$.*
- *For every unit vector $v \in TM$, $K_0(m, n)R \leq |Dg(v)| \leq K_1(m, n)R$.*
- *The reach of $g(M)$ is at least 1.*

7 Final Remark

The theme of this paper has been the geometricization of analysis, of understanding the complexity of maps and of spaces, and how doing so helps us understand each.

Although topology has some highly ineffective parts (as was evidenced by the theorems at the end of Sect. 4, and which have some very pleasant positive implications such as the existence of many solutions to certain variational problems), the complexity of the solution to topological problems is frequently, a posteriori, low.

It is almost tautologous that algebraic topology is really about the algebraic topology of function spaces. Its successes have told us, for example, about connectedness of such spaces. However, every driver knows how important it is to have a road map, rather than the mere knowledge that one is on a connected continent. The results I have sketched show that certain Lipschitz function spaces not only have boundable «volume » (as measured by covering numbers, the most basic invariant of approximation theory, or VC dimension, as the machine learning analogue) but frequently unusually small diameter. For targets with no rational homotopy Lipschitz function spaces had bounded diameter, and for simple rational homotopy types one had linear results. Determining exactly how much one needs to give up on Lipschitz constant and length of homotopy for simply connected targets[20] is an important challenge (see [39] for important recent progress).

[20] However, for $\mathrm{Lip}(S^1 : M)$ for M a 3-dimensional Sol manifold the diameter is $\exp(L)$, as suggested by a worst case analysis based on covering numbers.

This suggests that the geometry of function spaces is a rich object deserving of study. It also suggests that even in nonlinear situations (governed by either nonpositive curvature or simple connectivity) one should be able to navigate these spaces and (optimistically) produce fast computational algorithms.

References

1. M. Atiyah and R. Bott. A Lefschetz fixed point formula for elliptic complexes: II. Applications. *Ann. of Math.*, 88:451–491, 1968.
2. N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. *Discrete Comput. Geom.*, 22:481–504, 1999.
3. Y. Barzdin. On the realization of networks in three-dimensional space. In A. N. Shiryayev, editor, *Selected Works of A. N. Kolmogorov, Volume III*, volume 27 of *Mathematics and its Applications (Soviet Series)*, pages 194–202. Springer, Dordrecht, 1993.
4. J. Boissonnat, F. Chazal, and M. Yvinec. Computational geometry and topology for data analysis. To appear.
5. J. Boissonnat, L. Guibas, and S. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete Comput. Geom.*, 42(1):37–70, 2009.
6. S. Buoncristiano and D. Hacon. An elementary geometric proof of two theorems of Thom. *Topology*, 20(1):97–99, 1981.
7. J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52:46–52, 1985.
8. E. Brown. Finite computability of Postnikov complexes. *Ann. of Math.*, 65(1):1–20, 1957.
9. W. Browder. *Surgery on simply connected manifolds*. Springer Verlag, 1972.
10. J. Block and S. Weinberger. Large scale homology theories and geometry. In *Geometric Topology: 1993 Georgia International Topology Conference*, AMS/IP Stud. Adv. Math., pages 522–569, Providence, RI, 1997. Amer. Math. Soc.
11. G. Chambers, D. Dotterer, F. Manin, and S. Weinberger. Quantitative null-cobordism. *J. Amer. Math. Soc.*, 31(4):1165–1203, 2018.
12. S. Cheng, T. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA '05 Proceedings of the sixtieth annual ACM-SIAM symposium on discrete algorithms*, pages 1018–1027, Philadelphia, PA, 2005. Society for Industrial and Applied Mathematics.
13. J. Cheeger and M. Gromov. On the characteristic numbers of complete manifolds of bounded curvature and finite volume. In I. Chavel and H. M. Farkas, editors, *Differential Geometry and Complex Analysis*, pages 115–154. Springer, Berlin, Heidelberg, 1985.
14. J. Cha. A topological approach to Cheeger-Gromov universal bounds for von Neumann ρ-invariants. *Comm. Pure and Applied Math.*, 69:1154–1209, 2016.
15. G. Carlsson and F. Memoli. Classifying clustering schemes. *Found. Comput. Math.*, 13:221–252, 2013.
16. G. Chambers, F. Manin, and S. Weinberger. Quantitative null homotopy and rational homotopy type. *Geom. Funct. Anal.*, 28(3):563–588, 2018.
17. D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37(1):103–120, 2007.
18. G. Carlsson and A. Zomorodian. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.
19. F. Cucker and D. Zhou. *Learning Theory: An Approximation Theory Viewpoint*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2007.
20. A. Dranishnikov, S. Ferry, and S. Weinberger. Large Riemannian manifolds which are flexible. *Ann. of Math.*, 157:919–938, 2003.
21. A. Dranishnikov, S. Ferry, and S. Weinberger. An infinite dimensional phenomenon in finite dimensional topology. Preprint, 2017.
22. M. DoCarmo. *Riemannian Geometry*. Birkhäuser Verlag, 1992.
23. H. Edelsbrunner and D. Grayson. Edgewise subdivision of a simplex. *Discrete Comput. Geom.*, 24(4):707–719, 2000.
24. H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.

25. Y. Eliashberg and N. Mishachev. *Introduction to the h-Principle*, volume 48 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.
26. C. Fefferman. Whitney's extension problem for c^m. *Ann. of Math.*, 164:313–359, 2006.
27. C. Fefferman. c^m-extension by linear operators. *Ann. of Math.*, 166:779–835, 2007.
28. S. Ferry. Topological finiteness theorems for manifolds in Gromov-Hausdorff space. *Duke Math. J.*, 74(1):95–106, 1994.
29. C. Fefferman, S. Mitter, and H. Naryanan. Testing the manifold hypothesis. *J. Amer. Math. Soc.*, 29:983–1049, 2016.
30. S. Ferry and S. Weinberger. Quantitative algebraic topology and Lipschitz homotopy. *Proc. Natl. Acad. Sci. U.S.A.*, 110:19246–19250, 2013.
31. M. Gromov and L. Guth. Generalizations of the Kolmogorov-Barzdin embedding estimates. *Duke Math. J.*, 161:2549–2603, 2012.
32. M. Gromov. Homotopical effects of dilation. *J. Diff. Geo.*, 13:313–310, 1978.
33. M. Gromov. Groups of polynomial growth and expanding maps. *Publ. Math. IHÉS*, 53:51–78, 1981.
34. M. Gromov. *Partial Differential Relations*. Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. Springer Verlag, 1986.
35. M. Gromov. Dimension, non-linear spectra, and width. In J. Lindenstrauss and V. D. Milman, editors, *Geometric Aspects of Functional Analysis*, volume 1317 of *Lecture Notes in Mathematics*, pages 132–184. Springer, Berlin, Heidelberg, 1988.
36. M. Gromov. *Metric Structures for Riemannian and non-Riemannian Spaces*. Modern Birkhäuser Classics. Birkhäuser Verlag, 1999.
37. M. Gromov. Quantitative homotopy theory. In H. Rossi, editor, *Prospects in Mathematics (Princeton, NJ, 1996)*, pages 45–49. Amer. Math. Soc., Providence, RI, 1999.
38. J. Kleinberg. An impossibility theorm for clustering. In S. Becker, K. Obermayer, and S. Thrun, editors, *Advances in Neural Information Processing Systems 15 (NIPS 2002)*, pages 463–470. MIT Press, Cambridge, MA, 2002.
39. F. Manin. Plato's cave and differential forms. Preprint, 2018.
40. J. Matousek. Lecture notes on metric embeddings. Preprint, 2013.
41. J. Milnor. *Morse Theory*, volume 51 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1963.
42. J. Milnor. A note on curvature and fundamental group. *J. Differ. Geom.*, 2:1–7, 1968.
43. F. Manin and S. Weinberger. The Gromov-Guth embedding theorem. Appendix to [11].
44. A. Nabutovsky. Non-recursive functions, knots "with thick ropes" and self-clenching "thick" hyperspheres. *Commun. Pure Appl. Math.*, 48(4):1–50, 1995.
45. A. Nabutovsky. Morse landscapes of Riemannian functionals and related problems. In *Proceedings of the International Congress of Mathematicians: Hyderabad, India*, pages 862–881, 2010.
46. P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39:419–441, 2008.
47. A. Nabutovsky and S. Weinberger. Variational problems for Riemannian functionals and arithmetic groups. *Publ. Math. IHÉS*, 92:5–62, 2000.
48. S. Oudot. *Persistence Theory: From Quiver Representations to Data Analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.
49. I. Polterovich, L. Polterovich, and V. Stojisavljević. Persistence barcodes and Laplace eigenfunctions on surfaces. Preprint, 2017.
50. L. Polterovich, D. Rosen, K. Samvelyan, and J. Zhang. Persistent homology for symplectic topologists. Preprint, 2018.
51. V. Robins. Toward computing homology from finite approximations. *Topology Proceedings*, 24:503–532, 1999.
52. J. Roe. *Index theory, coarse geometry, and topology of manifolds*, volume 90 of *CMBS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, RI, 1996.
53. S. Smale. The classification of immersions of spheres in Euclidean spaces. *Ann. of Math.*, 69:327–344, 1959.
54. E. Spanier. *Algebraic Topology*. McGraw-Hill, 1956.
55. R. Strong. *Notes on cobordism theory*. Mathematical Notes. Princeton University Press, Princeton, NJ, 1968.
56. D. Sullivan. Infinitesimal computations in topology. *Publ. Math. IHÉS*, 47:269–331, 1977.

57. R. Thom. Quelques propriétés globales des variétés différentiables. *Comment. Math. Helv.*, 28:17–86, 1954.
58. J. Traub and A. Werschulz. *Complexity and information*. Lezioni Lincee. Cambridge University Press, Cambridge, United Kingdom, 1999.
59. L. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984.
60. C. T. C. Wall. *Surgery on Compact Manifolds*. London Mathematical Society Monographs. Academic Press, Cambridge, MA, 1969.
61. S. Weinberger. *Computers, Rigidity, and Moduli*. Princeton University Press, Princeton, NJ, 2004.
62. S. Weinberger. What is... persistent homology. *Notices Amer. Math. Soc.*, 58(1), 2011.
63. Y. Yomdin and G. Comte. *Tame Geometry with Applications in Smooth Analysis*, volume 1834 of *Lecture Notes in Mathematics*. Springer Verlag Berlin Heidelberg, 2004.