HOMOLOGY OF RANDOM GEOMETRIC COMPLEXES

SAMUEL SHEN

Abstract. Geometric complexes are a powerful tool of geometric decomposition; *random* geometric complexes add a probabilistic element that makes approximation methods much more flexible. We will demonstrate one such asymptotic method of studying homology, revealing the uniformly percolating behavior of homology groups of *every* dimension of the random Rips complex by an analysis only of its underlying geometric graph. We follow Kahle [3] while relying heavily on results about random geometric graphs from Penrose [5] and utilizing the Fundamental Theorem of Discrete Morse Theory [7] developed by Forman, both of which are at the intuitive heart of the subject of geometric complexes, which is geometric simplification.

Contents

1. Introduction

The study of geometric complexes is a field where we examine the geometric interactions between a set of points in a space, with ties to topology, geometry, algebra, and combinatorics. A precursor to geometric complexes is geometric graphs, where we attach edges between points that are within a certain distance of each other to lay a graph structure on our points, which are now vertices. The power of a geometric graph is immediately intuitive as a way to manage distance-based networks such as social circles if the vertices are individuals, digital systems if the vertices are nodes, molecular bonds if the vertices are atoms etc.

Our goal is to reveal that there is often a richer geometry behind a set of points embedded in a space than a geometric graph is capable of revealing. Indeed, oftentimes we wish not only to extract information from a "point cloud," but also to utilize those points to understand the ambient space in which they lie. Extending the notion of connected components and cycles in graphs, we will study the homotopy and homology of geometric complexes, which offers a far more comprehensive geometry on a vertex set and its underlying space, especially when the dimension of our space is great.

To geometrically understand a space, ideally we can cover it with the *body* of a geometric complex, which then "tessellates" the space. The geometric complex will "triangulate" the space into simpler parts, known as *simplices*, which possess *simple* combinatorial description and geometric properties. The idea is to make a complex space more understandable by analyzing its elementary components. Indeed, we will introduce and use results from Discrete Morse Theory, which is precisely the study of identifying which elementary components of a geometric complex comprise the *entire* homotopy of a triangulated space. With some classes of spaces, such as finite topological spaces [9], we can understand their homotopy completely by simple homotopy equivalence to geometric complexes.

However, sometimes we do not even possess enough information about a geometric object to triangulate it via a complex or to find a homotopy equivalent one. This is where we finally arrive at *random* geometric complexes. Even if we are unable to yield a specific "tessellating" complex, we can try to derive universal properties about complexes taken randomly on a space. A technique we will use is to take the asymptotic behavior of random complexes to become asymptotically more intricate, meaning more vertices and smaller *scale*, to capture the details of our space.

We motivate and examine one of the geometric complexes that is most prevalent in the field: the Rips Complex.

The Rips Complex is interesting to study because of its close connection to the more primitive geometric graph, from which we first motivated geometric complexes. Our goal will be to access the rich homological information contained in a geometric complex while limiting our analysis solely to the study of its *1-skeleton*, which is an underlying geometric graph for a Rips Complex. Thus, we will be able to use combinatorial methods to access topological information.

The study of the random Rips Complex can be seen as an extension of the study of *random geometric graphs* [5]. Instead of studying percolation, which can be viewed as the study of 0-dimensional homology, we examine when homology of higher dimensions appears and disappears. Furthermore, instead of directly varying the probability that two adjacent vertices are connected and finding the *percolation threshold*, we can instead vary the scale of the Rips Complex in relation to the number of vertices we have and find thresholds of vanishing and non-vanishing homology. We present the remarkable result by Kahle [3] from 2010 that: if the scale of our random Rips Complex is asymptotically within a specific *tight* interval (with respect to the number of points in the complex), homology of *every single* dimension exists as we asymptotically increase the number of points in our random complex and does not exist otherwise.

This kind of asymptotic analysis ensures that we are approximating our space as finely as possible as the number of vertices increases and scale decreases causing our Rips Complex to fit snuggly into the ambient space. Kahle's result is extremely powerful when we are able to use a random Rips complex to asymptotically approximate a space, as we can pinpoint the existence and non-existence of homology groups of the space by just the scale of our complex. The connections with percolation theory are revealing as they tell us that we could actually access far more information from just a geometric *graph* than might seem immediately obvious, though the venture into random geometric *complexes* was necessary to make this apparent.

2. Preliminaries

2.1. Geometric Complexes.

We first establish the basic construction and power of a geometric complex as a way to decompose spaces into simpler parts.

Definition 2.1. Given a set of vertices $V = \{v_0, v_1, ..., v_k\}$, we define its *convex hull* by the smallest *convex* set "wrapping" the vertices:

$$Conv(V) = \left\{ \sum_{i=0}^{k} \alpha_i \cdot v_i | \sum_{i=0}^{k} \alpha_i = 1, \alpha_i \geq 0 \right\}$$

Definition 2.2. A *geometric k-simplex* σ in X is the convex hull of an affinely independent set of vertices $V = \{v_0, v_1, ..., v_k\} \subseteq X$ (requiring $k+1$ vertices to construct a k-simplex).

Definition 2.3. If $U \subseteq V$, $\tau = Conv(U)$ is a *face* of $\sigma = Conv(V)$ and σ is a *coface* of τ.

Definition 2.4. A *geometric simplicial complex* or *geometric complex* $K \subseteq X$ is a finite collection of geometric simplices such that:

(1) If $\sigma \in K$ and τ is a face of σ, then $\tau \in K$
(2) If $\sigma, \tau \in K$, then $\sigma \cap \tau$ is either $\emptyset$ or a common face of σ and τ

Definition 2.5. The *body* of K is $\bigcup_{\sigma \in K} \sigma \subseteq X$, the conglomeration of all the simplices of the complex as a single set which is *triangulated* by the complex K.

Definition 2.6. The *n-skeleton* of K is its set of m-simplices with $m \leq n$

The vocabulary we've defined above including *face*, *body*, *triangulate*, and *skeleton* highlight that the utility of geometric complexes is to understand shapes by "tesselating" them: objects that can be disjointly decomposed in an organized fashion are then much easier to study.

Remark 2.7. It is interchangeable to denote a simplex either by just its vertices (abstract representation) or by the collection of all points in the simplex (geometric realization). We will opt for the abstract representation.

2.2. **The Rips Complexes.** We motivate the use of Rips Complexes by showing that it has hidden combinatorial structure which aids us in the study of its *homological* properties.

Definition 2.8. The *Vietoris-Rips Complex* on vertices $V \subseteq X$ with scale $r \geq 0$ is defined by (where each simplex is given by its abstract vertex set representation):

$$\mathit{Rips}(V,\ r) = \{\sigma \subset V | \forall x, y \in \sigma, d(x,y) \leq r\}$$

Definition 2.9. The *geometric graph* on n points $X_n \subseteq \mathbb{R}^d$ with scale $r > 0$ is the *undirected* graph $G(X_n, r) = (V, E)$ with vertex set $V = X_n$ and edge set $E = \{\{x, y\} \subseteq X_n | d(x, y) \leq r\}$.

Definition 2.10. A *k-clique* of a graph $G = (V, E)$ is a subset $W_k \subseteq V$ of k vertices that are all adjacent to each other i.e. $\forall x, y \in W_k$ where $x \neq y$, $\{x, y\} \in E$.

Definition 2.11. The *clique complex* of a graph $G = (V, E)$ inherits the same vertex set V where each (k+1)-*clique* of G forms a corresponding k-simplex.

Definition 2.12. Simplicial complexes that arise as the clique complex of a *geometric graph* are called *flag complexes*.

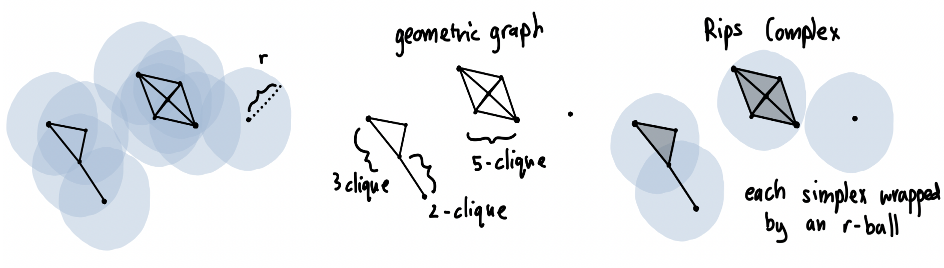

FIGURE 1. Geometric Graph and corresponding Rips Complex in $\mathbb{R}^2$

Remark 2.13. Observe that $\mathit{Rips}(X_n, r)$ is defined to be exactly the clique complex of the geometric graph $G(X_n, r)$. Consequently, our study of the Rips complex can be shifted to studying the underlying geometric graph.

2.3. **Simplicial Graph Theory.** We reveal the way in which the geometric graph of a Rips Complex is connected to its homology.

Definition 2.14. The *(k+1)-dimensional cross-polytope* is the convex hull of the $2k + 2$ mirrored $\mathbb{R}^{k+1}$ standard basis vectors $\{\pm e_i\}_{1 \leq i \leq k+1}$. The boundary of the cross-polytope is a k-dimensional simplicial complex denoted O_k, which we will refer to as a *"vertex-minimal" k-sphere.*

We now introduce the main combinatorial sleight of hand: the k'th homology of a flag complex can be identified by "vertex-minimal" spheres in the underlying geometric graph. This is geometrically intuitive up to 3-dimensions but we extend the result generally.

Lemma 2.15. *If $\triangle$ is a flag complex, then any nontrivial element of $H_k(\triangle)$ also exists on a subcomplex $S \subset \triangle$ with at least $2k+2$ vertices. If S has exactly $2k+2$ vertices, then S is combinatorially isomorphic to O_k.*

Proof. (Sketch) Any non-trivial element of $H_k(\triangle)$ indicates the existence of a k-cycle that is not the boundary of a (k+1)-simplex and so must contain at least 2k + 2 vertices. We construct a minimal subcomplex S by combining all and only all the simplices part of this k-cycle. By definition, since O_k is the simplex with the smallest number of vertices homeomorphic to a k-sphere and S has exactly 2k + 2 vertices, it must be combinatorially isomorphic to O_k by the uniqueness of the "vertex-minimal" k-sphere. □

Remark 2.16. We will thus transfer our study of k'th homology to the study of the subgraphs of the geometric graph isomorphic to O_k.

Definition 2.17. A graph $H = (V_H, E_H)$ is an *induced* subgraph of $G = (V_G, E_G)$ if for any $x, y \in V_H$, we have that $\{x, y\} \in E_H$ if and only if $\{x, y\} \in E_G$.

Remark 2.18. A connected graph is called *feasible* if it is geometrically realizable as an induced subgraph of a *geometric graph.* We harmlessly add a feasibility requirement as we are only interested in the geometric graph.

Example 2.19. The complete bipartite graph $K_{1,7}$ is not feasible since a regular heptahedron has longer distance from each vertex to its center than side lengths.

Definition 2.20. Induced Subgraph Counting Functions Denote the number of induced subgraphs of $G(X_n, r)$ isomorphic to H by $G_n(H)$. In particular, $G_n(O_k)$ will tell us how many k-dimensional "holes" there are in the Rips complex.

Definition 2.21. Normalized Cross-Polytope Counting Function For a feasible graph H of order k, define the indicator function $h_H : X_k \subset \mathbb{R}^d \to \mathbb{R}$ by $h_H(X_k) = 1$ if $G(X_k, 1)$ is isomorphic to H and 0 if not. If $f : \mathbb{R}^d \to \mathbb{R}$ is the density function from which the points X_k are drawn, we define:

$$\mu_H = k!^{-1} \int_{\mathbb{R}^d} f(x)^k dx \int_{(\mathbb{R}^d)^{k-1}} h_H(\{0, x_1, ..., x_{k-1}\}) d(x_1, ..., x_{k-1})$$

Theorem 2.22. ***Expectation of Subgraph Counts*** *Suppose that* $\lim_{n\to\infty} r = 0$ *and* H *is a feasible graph of order* $k \geq 2$*, then:*

$$\lim_{n\to\infty} r^{-d(k-1))} n^{-k} \mathbb{E}[G_n(H)] = \mu_H$$

Proof. The proof is purely measure-theoretic and thus rather unenlightening so we will omit it, but it can be found in Proposition 3.1 in Penrose [5] □

Remark 2.23. Theorem 2.22 is the main result from Random Geometric Graph Theory [5] which we rely on to understand the presence of "vertex-minimal" spheres. Its versatility comes largely from the fact that we only need our scale to asymptotically approach 0 with respect to the vertex count without any requirements to the speed at which we do so.

2.4. **Discrete Morse Theory.** We will build up to the fundamental theorem of discrete morse theory, which is the extremely powerful idea that we can analyze just a select few simplices in a geometric complex to understand its homotopy class completely.

Definition 2.24. CW Complexes A *d-cell* σ is homeomorphic to the d-ball. We use *attaching maps* $f : \partial\sigma \to X$ to build a *CW complex* X by attaching cells of increasing dimension (building first the 0-skeleton then 1-skeleton and so on) via $X \cup_f \sigma$ where $\partial\sigma$ is "glued" to $f(\partial\sigma)$ on the pre-existing structure of X.

Remark 2.25. A *d-simplex* is a *d-cell* so geometric complexes are CW Complexes

Definition 2.26. We now motivate Discrete Morse Theory with the notion of *homotopic simplification.* (σ_{p-1}, τ_p) is called a *free pair* if τ is the only coface of σ. Removing a free pair is a deformation retraction called an *elementary collapse*, preserving homotopy type while "simplifying the complex." Our goal is to simplify the complex maximally and encode this simplification while preserving homotopy.

Definition 2.27. A *discrete vector field* V on K is a pairing of the faces of K denoted by $V = \{(\sigma^{(p-1)}, \tau^{(p)}))|\sigma \subset \tau\}$ where each simplex is in at most one pair.

Definition 2.28. A *V-path* on a discrete vector field V is a sequence of simplices $\{\sigma_0^{(p)}, \tau_0^{(p+1)}, \sigma_1^{(p)}, \tau_1^{(p+1)}, ..., \tau_{k-1}^{(p+1)}, \sigma_k^{(p)}\}$ such that $(\sigma_i^{(p)}, \tau_i^{(p+1)}) \in V$ and $\sigma_i^{(p)}$ is a face of $\tau_{i-1}^{(p+1)}$. Crucially, if $\sigma_0^{(p)} = \sigma_k^{(p)}$, then the V-path is *closed*.

Remark 2.29. The definition of a V-path reveals that a discrete vector field implicitly assigns arrows pointing from a (p+1)-simplex to *all* of its p-faces while potentially swapping the direction of at most one arrow per face. These arrows direct us toward some chosen scalar hierarchy, which we hope is decreasing "homotopic importance." A closed V-loop is a sort of *escherian stairwell* preventing us from constructing a notion of gravity.

Definition 2.30. A discrete vector field without closed V-paths is called a *discrete gradient vector field*. A face that is not paired in a discrete gradient vector field is called *critical*, a "highest" or "lowest" point in the hierarchy.

Example 2.31. If we map the *complete* elementary collapse of a geometric complex by arrows between free pairs where we will "rip" out a random simplex when no feasible collapses remain, those ripped out simplices are precisely the *critical* ones that constitute the entire homotopy type of the complex. The best choice of a discrete gradient vector field will give us the best critical simplices to approximate the complex homotopically.

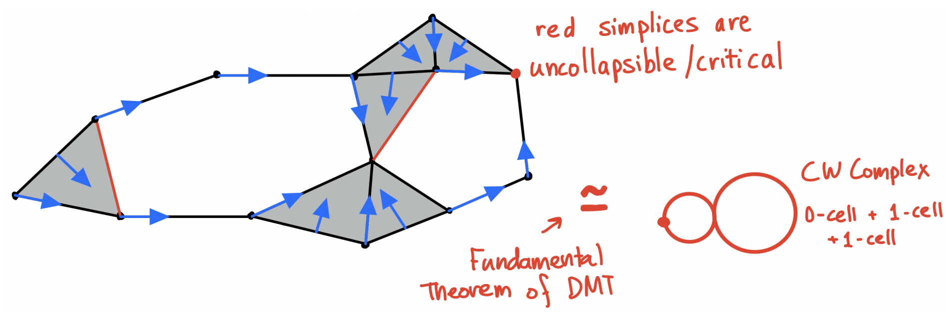

FIGURE 2. Complete elementary collapse with free pairs (depicted by blue arrows) forming discrete gradient vector field

Theorem 2.32. ***Fundamental Theorem of Discrete Morse Theory (Forman)*** *If $\triangle$ is a complex with discrete gradient vector field V, then $\triangle$ is homotopy equivalent to a CW complex with a k-cell attached for each critical k-simplex in V. This formally describes that critical simplices are responsible for the entire homotopy of the complex.*

3. Interval of Non-Vanishing Homology

We figure out for which scales the random Rips Complex has and does not have non-trivial homology as we asymptotically increase vertex count and decrease scale.

Definition 3.1. **Asymptotic Property** An object X_n has a property $\mathbf{P}$ *asymptotically almost surely (a.a.s)* if $\lim_{n\to\infty} \mathbb{P}\{X_n \in \mathbf{P}\} = 1$.

Remark 3.2. We will now prove the remarkable fact that as a random Rips Complex becomes arbitrarily large, we can predict for which scales $r > 0$ the k'th homology for *any* $k \in \mathbb{N}$ is non-trivial. Intuitively, there is no homology when the scale is too small for faces to form nor when the scale is too large and the complex is completely connected, which is also when the underlying geometric graph percolates.

Proposition 3.3. ***Interval Bounds*** *Let our random Rips complex be generated by a uniform distribution of points on a compact and convex set $K \subset \mathbb{R}^d$ with nonempty interior. Then given any $k \geq 0$, the following three statements describe when the k'th homology appears and disappears:*

$$(1) \textit{If } r = o\left(n^{-\frac{2k+2}{d(2k+1)}}\right), \textit{ then a.a.s. } H_k = 0$$

$$(2) \textit{If } r = w\left(n^{-\frac{2k+2}{d(2k+1)}}\right) \textit{ and } r = o\left(\left(\frac{log(n)}{n}\right)^{\frac{1}{d}}\right), \textit{ then a.a.s. } H_k \neq 0$$

$$(3) \textit{If } r = w\left(\left(\frac{log(n)}{n}\right)^{\frac{1}{d}}\right), \textit{ then a.a.s. } H_k = 0$$

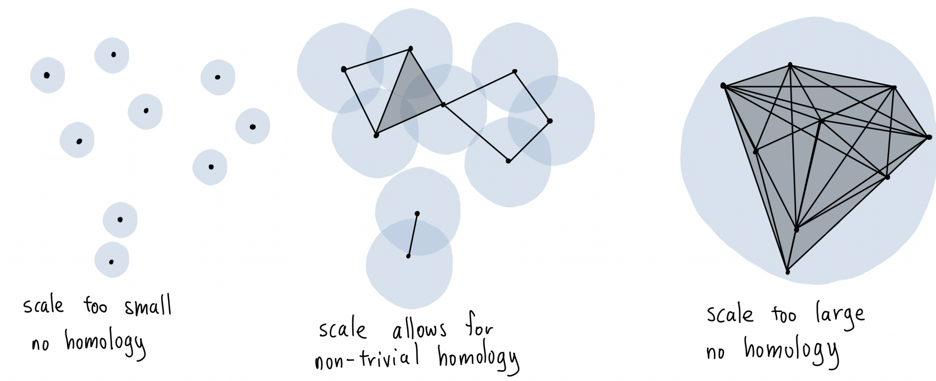

FIGURE 3. Depiction of homology for a single set of random vertices at varying scales for a Rips Complex in $\mathbb{R}^2$

Remark 3.4. The convexity requirement is so that our Rips complex stays within our ambient space. While it may seem troublesome to impose that our ambient space K be bounded, this is actually quite a weak stipulation as K can be arbitrarily large. Thus, we can analayze the topology in a local fashion and apply limiting behavior if we want global properties. The rest of this section will prove the 3 components of Proposition 3.3.

Theorem 3.5. ***Lower Vanishing*** *If* $r = o\left(n^{-\frac{2k+2}{d(2k+1)}}\right)$, *then a.a.s.* $H_k = 0$

Proof. If $r = o\left(n^{-\frac{2k+2}{d(2k+1)}}\right)$, this means that:

$$\lim_{n\to\infty} \frac{r}{n^{-\frac{2k+2}{d(2k+1)}}} = \lim_{n\to\infty} rn^{\frac{2k+2}{d(2k+1)}} = 0$$

which implies that $\lim_{n\to\infty} r^{d(2k+1)}n^{2k+2} = 0$ by continuity. Since we also observe that $\lim_{n\to\infty} r = 0$, we apply our main combinatorial lemma for induced subgraph counting (Theorem 2.22) by setting $H = O_k$ as intentioned where $|O_k| = 2k+2$ to get:

$$\mu_{O_k} = (2k+2)!^{-1} \int\limits_{\mathbb{R}^d} f(x)^{2k+2} dx \int\limits_{(\mathbb{R}^d)^{2k+1}} h_{O_k}(\{0, x_1, ..., x_{2k+1}\}) d(x_1, ..., x_{2k+1})$$

$$= (2k+2)!^{-1} \cdot C \cdot (2k+2)! \cdot Vol(S^k) = C \cdot Vol(S^k) > 0$$

since the density function is integrated to some positive constant $C > 0$ and there is a shell of radius 1 around 0 and we can place the vertices to be isomorphic to O_k where we have $(2k+2)!$ permutations of each placement of vertices. Then by Theorem 2.22, we have that:

$$\lim_{n\to\infty} \frac{\mathbb{E}[G_n(O_k)]}{r^{d(2k+1)}n^{2k+2}} = \mu_{O_k} > 0$$

which implies that, since the denominator converges to 0, $\lim_{n\to\infty} \mathbb{E}[G_n(O_k)] = 0$. As established in Lemma 2.15, if the geometric graph has no induced subgraphs forming "vertex-minimal spheres," the corresponding Rips complex has no k'th homology. □

Theorem 3.6. ***Interval of Non-Vanishing***
If $r = w\left(n^{-\frac{2k+2}{d(2k+1)}}\right)$ *and* $r = o\left(\left(\frac{log(n)}{n}\right)^{\frac{1}{d}}\right)$, *then a.a.s.* $H_k \neq 0$

Proof. The proof is identical to Theorem 3.5 as we still observe that $\lim_{r\to\infty} = 0$ still holds so we apply our combinatorial method (Theorem 2.22) but this time $\lim_{n\to\infty} r^{d(2k+1)}n^{2k+2} = \infty$ where since $\mu_{O_k} < \infty$, we have that $\lim_{n\to\infty} \mathbb{E}[G_n(O_k)] = \infty$ must hold so that homology is non-vanishing (and in fact becomes richer as we take more vertices and proportionally smaller scales from the unbounded expectation). □

Remark 3.7. We can observe that the previous upper bound for the scale $r > 0$ is as high as we can go before we lose $\lim_{n\to\infty} r = 0$ so can no longer apply the induced subgraph counting method (Theorem 2.22). In fact, the underlying geometric graph percolates precisely when homology becomes trivial. We will now formally prove that once we surpass this bound, we indeed lose all but 0'th and 1'st homology groups using discrete morse theoretic methods. First, we introduce a geometric lemma.

Lemma 3.8. ***Geometric Lemma*** *Given* $\{y_0, ..., y_l\} \subseteq \mathbb{R}^d$ *such that* $||y_0|| \leq ||y_1|| \leq ... \leq ||y_l||$ *and* $||y_1|| \geq \frac{1}{2}$, $||y_0 - y_1|| > 1$ *and* $||y_i - y_j|| \leq 1$ *for every other* $0 \leq i < j \leq l$, *then there exists* ϵ_d *such that* $\mu\left(\bigcap_{i=1}^{l} B(y_i, 1) \cap B(0, ||y_1||)\right) \geq \epsilon_d$ *(where* μ *is Lebesgue measure i.e. volume).*

Proof. Let $y_m = \frac{y_0+y_1}{2}$. Since $||y_0 - y_1|| > 1$, then $||y_m - y_0|| = ||y_m - y_1|| > \frac{1}{2}$. Now let $\theta > 0$ be the angle between $y_0 - y_2$ and $y_1 - y_2$. Since $||y_0 - y_2|| \leq 1$, $||y_1 - y_2|| \leq 1$, and $||y_0 - y_1|| > 1$, the law of cosines yields that:

$$(y_0 - y_2) \cdot (y_1 - y_2) = ||y_0 - y_2||||y_1 - y_2||cos\theta$$

$$= \frac{1}{2}\left(||y_0 - y_2||^2 + ||y_1 - y_2||^2 - ||y_0 - y_1||^2\right) < \frac{1}{2}$$

Then it follows that:

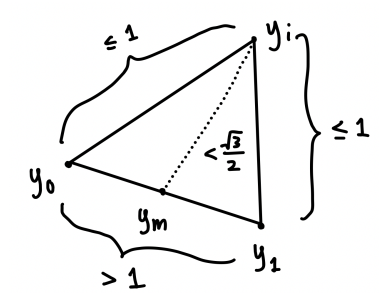

$$||y_m - y_2||^2 = (y_m - y_2) \cdot (y_m - y_2)$$

$$= \left(\frac{y_0 + y_1}{2} - y_2\right)\left(\frac{y_0 + y_1}{2} - y_2\right) = \left(\frac{y_0 - y_2}{2} + \frac{y_1 - y_2}{2}\right)\left(\frac{y_0 - y_2}{2} + \frac{y_1 - y_2}{2}\right)$$

$$= \frac{1}{4}\left(||y_0 - y_2||^2 + ||y_1 - y_2||^2 + 2(y_0 - y_2) \cdot (y_1 - y_2)\right)$$

$$< \frac{1}{4}\left(1 + 1 + 2\left(\frac{1}{2}\right)\right) = \frac{3}{4}$$

so we can conclude that $||y_m - y_2|| < \frac{\sqrt{3}}{2}$. Replace y_2 with any y_i for $3 \leq i \leq l$ and the same argument holds. Now let $\rho = 1 - \frac{\sqrt{3}}{2}$. By the triangle inequality, $B(y_m, \rho) \subset B(y_i, 1)$ for $1 \leq i \leq l$. Then $B(y_m, \rho) \cap B(0, ||y_1||) \subset \bigcap_{i=1}^{l} B(y_i, 1) \cap B(0, ||y_1||)$.

Since $||y_0|| \leq ||y_1||$, $||y_1|| \geq \frac{1}{2}$, $||y_0 - y_1|| > 1$, and $||y_m|| \leq ||y_1||$, it follows that $\mu\left(\bigcap_{i=1}^{l} B(y_i, 1) \cap B(0, ||y_1||)\right) \geq \mu\left(B(y_m, \rho) \cap B(0, ||y_1||)\right) \geq \mu\left(B(y_1, \rho) \cap B(0, ||y_1||)\right) \geq \epsilon_d$ where we can set ϵ_d to the smallest volume when $||y_1|| = \frac{1}{2}$ where the volume depends on the Euclidean dimension d only. □

Lemma 3.9. ***Scaled Geometric Lemma*** *Given arbitrary* $r > 0$ *and* $\{y_0, ..., y_l\} \subseteq \mathbb{R}^d$ *such that* $||y_i - y_j|| \leq r$ *and* $(\frac{1}{2})r \leq ||y_1||$, $||y_0 - y_1|| > r$, *and* $||y_i - y_j|| \leq r$ *for every other* $0 \leq i < j \leq l$ *(excluding* $i = 0, j = 1$*), there exists* ϵ_d *(for any* $r > 0$*) such that* $\mu\left(\bigcap_{i=1}^{l} B(y_i, r) \cap B(0, ||y_1||)\right) \geq \epsilon_d r^d$.

Proof. Mutatis mutandis as unscaled Lemma 3.8 with arbitrary r replacing 1. □

Theorem 3.10. ***Upper Vanishing*** *If* $r = w\left(\left(\frac{log(n)}{n}\right)^{\frac{1}{d}}\right)$, *then a.a.s.* $H_k = 0$

Proof. Without loss of generality by our nonempty interior assumption, let $B(0, 1) \subseteq K$ (the space on which the vertices are distributed). Since we know with probability 1 that no two of our random vertices will be the same distance to the origin, we can index $X_n = \{x_1, ..., x_n\}$ so that $||x_1|| < ||x_2|| < ... < ||x_n||$.

Now we construct a *discrete vector field* V on $Rips(X_n, r)$ as follows: whenever possible, pair a face $S = \{x_{i_1}, x_{i_2}, ..., x_{i_j}\}$ with coface $\{x_{i_0}\} \cup S$ such that $i_0 < i_1$ and i_0 is as small as possible (S will choose the coface with the smallest possible i_0). We verify the validity of this discrete vector field as follows:

- S cannot get paired with two cofaces $\{x_a\} \cup S$ and $\{x_b\} \cup S$ since it will prefer $\min\{a, b\}$
- S cannot get paired with a face and a coface since if S is paired with coface $\{x_a\} \cup S$, then $||x_a|| < ||s||$ for every $s \in S$ so any face $L \subset S$ would *also* prefer to be paired with $\{x_a\} \cup S$ instead of S.

Since each face is in at most one pair, V is a well-defined discrete vector field. Additionally, since the indices are decreasing along any V-path of V because the smallest indice of two adjacent faces and cofaces (separated by 1 face in the V-path) must be lower by the way we choose index-minimal pairings, V has no closed V-loops and is a discrete graident vector field.

Denote $W = nr^d$ for notational ease. Define p_k to be the probability that a random set of $k + 1$ vertices of X_n span a k-face of $Rips(X_n, r)$. If we are given a single vertex v of the set, we know that the remaining k vertices must lie in $B(v, r)$ so that:

$$p_k = O(r^d k) = O\left(\left(\frac{W}{n}\right)^k\right) \tag{3.11}$$

Suppose that $\{x_{i_1}, x_{i_2}, ..., x_{i_{k+1}}\}$ span a critical face F. Then we know that:

- There is no common neighbor x_a of *all* the vertices of F where $a < i_1$ or else F would pair up with $\{x_a\} \cup F$
- There is a common neighbor x_{i_0} of $x_{i_2}, ..., x_{i_{k+1}}$ (and not of x_{i_1}) such that $i_0 < i_1$ or else F would be paired up with face $F \setminus \{x_{i_1}\}$

All requirements of the *scaled geometric lemma* (Lemma 3.9) have been met, namely that $||x_{i_0}|| < ||x_{i_1}|| < ... < ||x_{i_{k+1}}||$, $||x_{i_0} - x_{i_1}|| > r$ since x_{i_0} and x_{i_1} are not neighbors, $||x_{i_m} - x_{i_n}|| > r$ for every other $0 \leq m < n \leq k + 1$ since they are neighbors,

and $||x_{i_1}|| \geq (\frac{1}{2})r$ because otherwise $||x_{i_0} - x_{i_1}|| < r$, which is a contradiction. The scaled geometric lemma allows us to conclude:
Let $I = \bigcap_{j=1}^{k+1} B(x_{i_j}, r) \cap B(0, ||x_{i_1}||)$. There exists $\epsilon_d > 0$ such that $\mu(I) \geq \epsilon_d r^d$. We observe that the probability that a single random point in K falls in I is $\frac{\mu(I)}{\mu(K)} \geq \frac{\epsilon_d r^d}{\mu(K)}$ by our uniformity assumption.
We now observe that any vertex x_a that falls in I will necessarily be neighbors with all the vertices in F and also satisfy $a < i_1$. Thus, the vertices of our critical face $\{x_{i_1}, x_{i_2}, ..., x_{i_{k+1}}\}$ and x_{i_0} must be the only vertices that live in I. Thus, if we let p_c be the probability that an arbitrary (not necessarily critical anymore here) k-face F is *critical*, we can bound p_c by the probability that the remaining $n - (k+2)$ points of X_n are not in the set I constructed from F.

$$p_c \leq \left(1 - \frac{\epsilon_d}{\mu(K)} r^d\right)^{n-k-2} \tag{3.12}$$

$$\leq exp\left(-\frac{\epsilon_d}{\mu(K)} r^d (n-k-2)\right) = O(exp(-cW))$$

where c is any constant satisfying $0 < c < \frac{\epsilon_d}{\mu(K)}$
Now let C_k denote the number of critical k-faces with respect to our discrete gradient vector field V where we conclude that:

$$\mathbb{E}[C_k] \leq \binom{n}{k+1} p_k p_c \tag{3.13}$$

as we multiply all the number of ways to choose a $k+1$ vertex subset by the probability that the $k+1$ set is a face by the probability that the $k+1$ face is critical. Extend (3.13) further by applying the bounds in (3.11) and (3.12)

$$\leq \binom{n}{k+1} \left(\frac{W}{n}\right)^k e^{-cW} = O\left(W^k e^{cW} n\right)$$

Since $r = w\left(\left(\frac{log(n)}{n}\right)^{\frac{1}{d}}\right)$, we have that eventually both $r \geq \left(\frac{log(n)}{n}\right)^{\frac{1}{d}}$ i.e. $r^d n \geq log(n)$ and $n \geq e$ so that we extend (3.13) further by

$$= O\left((nr^d)^k e^{-cnr^d} n\right) = O\left((nr^d)^k n^{1-cnr^d}\right) = O\left((log(n))^k n^{1-clog(n)}\right)$$

since the exponential term overpowers polynomial term asymptotically so we can replace both appearances of nr^d with $log(n)$ even though the polynomial replacement alone will lower our bound. Our extension of (3.13) allows us to conclude that $\mathbb{E}[C_k] \to 0$ as $n \to \infty$.
The only critical face that always exists is the 0-face that is the vertex closest to the origin, which cannot be paired with any other face since it has the smallest norm by our construction of our discrete gradient vector field V. This is the only vertex set that is guaranteed to be both a face and unpairable.

By the *Fundamental Theorem of Morse Theory* (Theorem 2.32), we have that $Rips(X_n, r)$ is a.a.s. homotopy equivalent to a CW-complex with just a 0-cell and no other cells attached. Hence, $Rips(X_n, r)$ is asymptotically contractible and has no homology excluding 0'th homology. □

Remark 3.14. Proposition 3.3 establishes a tight "goldilocks" interval in which the scale and number of vertices of the random Rips Complex is perfectly balanced such that an infinitely dense complex will have non-trivial homology groups of every dimension. We've narrowed our analysis primarily to the 1-skeleton of the Rips Complex as promised while utilizing the core ideas from discrete morse theory, namely extracting core simplices of a complex to crystallize our understanding of its topology, both of which bring out the intuitive essence of geometric complexes. The technique of using a random Rips Complex to fill out the "intricacies" of a space is highly robust as the asymptotically increasing density frees us from having to stipulate how our points are distributed, so that our results can be flexibly applied in analyzing the homology of nearly any space via the random Rips Complex.

Acknowledgments

I would like to express my gratitude to Professor May for organizing an enriching REU this Summer and for invaluable feedback on my paper, to Professor Rudenko and Professor Babai for amazing apprentice lectures and problem sessions, and to my mentor DeVon Ingram for working with me throughout the program.

References

[1] Žiga Virk. Introduction to Persistent Homology. http://zalozba.fri.uni-lj.si/virk2022.pdf

[2] Measure Theoretic Probability. P.J.C. Spreij. https://staff.fnwi.uva.nl/p.j.c.spreij/onderwijs/master/mtp.pdf

[3] M. Kahle. Random geometric complexes. Discrete Comput. Geom., 45(3):553–573, 2011.

[4] R. Meester and R. Roy. Continuum percolation, volume 119 of Cambridge Tracts in Mathematics. Cambridge University Press, Cambridge, 1996

[5] M. Penrose. Random geometric graphs, volume 5 of Oxford Studies in Probability. Oxford University Press, Oxford, 2003.

[6] O. Bobrowski and M. Kahle. Topology of random geometric complexes: a survey. J. Appl. Comput. Topol., 1(3):331–364, Jun 201

[7] Robin Forman. A user's guide to discrete Morse theory. S ́em. Lothar. Combin., 48:Art. B48c, 35 pp. (electronic), 2002.

[8] Allen Hatcher. Algebraic topology. Cambridge University Press, Cambridge, 2002.

[9] J.P. May and Elle Pishevar. Finite Spaces and Larger Contexts. http://www.math.uchicago.edu/ may/REU2023/FiniteSpacesNew