

A one-dimensional homologically persistent skeleton of an unstructured point cloud in any metric space

V. Kurlin¹¹ Microsoft Research Cambridge, 21 Station Road, Cambridge CB1 2FB, UK and
Department of Mathematical Sciences, Durham University, Durham DH1 3LE, UK

Abstract

Real data are often given as a noisy unstructured point cloud, which is hard to visualize. The important problem is to represent topological structures hidden in a cloud by using skeletons with cycles. All past skeletonization methods require extra parameters such as a scale or a noise bound. We define a homologically persistent skeleton, which depends only on a cloud of points and contains optimal subgraphs representing 1-dimensional cycles in the cloud across all scales. The full skeleton is a universal structure encoding topological persistence of cycles directly on the cloud. Hence a 1-dimensional shape of a cloud can be now easily predicted by visualizing our skeleton instead of guessing a scale for the original unstructured cloud. We derive more subgraphs to reconstruct provably close approximations to an unknown graph given only by a noisy sample in any metric space. For a cloud of n points in the plane, the full skeleton and all its important subgraphs can be computed in time $O(n \log n)$.

Categories and Subject Descriptors (according to ACM CCS): I.5.1 [Pattern Recognition]: Models—Structural

1. Introduction: motivations and our contributions

A point cloud C is a finite metric space, namely a finite set of points with only pairwise distances. The traditional way to visualize a shape of C is to select a scale α and join all points of C at a distance not more than 2α . If a cloud C is high-dimensional, choosing a suitable scale can be hard.

Topological Data Analysis resolves this difficulty by studying data across all scales. The output is a persistence diagram of pairs (birth, death) encoding life spans of topological features such as closed cycles from their birth to death. In the persistence diagram we can select pairs corresponding to closed cycles with a high persistence death – birth.

However, the pairs (birth, death) alone are insufficient to actually locate cycles in a given cloud C . That is why, to visualize the persistence of 1-dimensional cycles hidden in C across all scales, we introduce a Homologically Persistent Skeleton $\text{HoPeS}(C)$ whose vertices are all points of C .

Now all 1-dimensional cycles that persist in a given cloud C over many scales are visualized directly on C . The skeleton $\text{HoPeS}(C)$ depends only on C , but contains shortest subgraphs that provably represent the 1-dimensional topology of C at every continuous scale, see Optimality Theorem 9.

$\text{HoPeS}(C)$ is obtained from a Minimum Spanning Tree $\text{MST}(C)$ by adding all critical edges giving birth to (homology classes of) 1D cycles in the persistence diagram, see two red critical edges with labels (birth, death) in Fig. 1.




Figure 1: Left: cloud $C \subset \mathbb{R}^2$. Middle: $\text{MST}(C)$. Right: a Homologically Persistent Skeleton with 2 red critical edges.

Assuming that a cloud C is a noisy sample of an unknown graph G , we find another natural hierarchy of derived subgraphs $\text{HoPeS}_{k,l}(C)$ that provably approximate G , see Theorem 15. Starting from any unstructured cloud C , we can visualize the full skeleton and then study its subgraphs that contain most persistent cycles depending on integers $k, l \geq 1$, which are easier to choose than a continuous scale α .

Here is a summary of our motivations for $\text{HoPeS}(C)$.

- Visualize in one universal skeleton 1-dimensional cycles hidden in a high-dimensional cloud C across all scales α .
- Extend a classical Minimum Spanning Tree $\text{MST}(C)$ of a cloud C to an optimal graph $\text{HoPeS}(C)$ with cycles.
- Solve the skeletonization problem for unstructured clouds with guarantees and without using any extra parameters.

A high-level description of our contributions is below.

- Definition 6 introduces a Homologically Persistent Skeleton $\text{HoPeS}(C)$ of a point cloud C , which is the first universal structure that visualizes 1-dimensional cycles across all scales directly on the given cloud C in any metric space.
- For any $\alpha > 0$, Theorem 9 proves that $\text{HoPeS}(C)$ contains a reduced skeleton with a minimum length over all graphs G having the homology of ‘thickened’ C at the scale α .
- The new Gap Search method in Propositions 11 and 13 strengthens the seminal stability of persistence diagrams by providing bijections between natural finite subdiagrams.
- For any ε -sample C of an unknown graph G , Theorem 15 gives conditions on G when subskeletons $\text{HoPeS}_{k,l}(C)$ have the homotopy type of G and are within the 2ε -offset of G .
- Corollary 16 proves that $\text{HoPeS}_{k,l}(C)$ is globally stable for any small perturbations of noisy samples C of graphs G .

2. Comparison with related past skeletonization work

Morse-Smale complex $\text{MS}(f, M)$ is defined for a function f on a manifold M (or for a discrete gradient field on a complex). To compare $\text{MS}(f, M)$ with $\text{HoPeS}(C)$, which depends on a cloud C , we need more structure on C . In practice f is a density depending on close neighbors of a point in C . Let $d_C(p)$ be the distance from $p \in \mathbb{R}^d$ to a closest point of C . Then $\text{MS}(d_C, \mathbb{R}^d)$ is a subdivision of a Delaunay triangulation and has the 1D skeleton larger than $\text{HoPeS}(C)$.

Forman’s discrete Morse theory for a cell complex with a discrete gradient field builds a smaller homotopy equivalent complex whose number of critical cells is minimized by the algorithm in [LLT04]. Optimality Theorem 9 minimizes the total length of a skeleton, not the number of critical edges. Removing low persistent edges from the full skeleton $\text{HoPeS}(C)$ to get smaller skeletons $\text{HoPeS}_{k,l}(C)$ is similar to the simplification [EHZ03] of Morse-Smale complexes, where critical points with close heights are cancelled.

All known skeletonization algorithms for clouds seem to require extra input parameters such as a scale α or a noise bound ε . Hence all these algorithms can not accept our minimal input, which is only a cloud C . That is why any experimental comparison of solutions to 2 different problems will be unfair and we compare only theoretical aspects below.

Delaunay-based skeletons. R. Singh et al. [SCP00] approximated a skeleton of a shape by a subgraph of a Delaunay triangulation using 3 thresholds: K for the minimum number of edges in a cycle and $\delta_{min}, \delta_{max}$ for inserting/merging Voronoi regions. Similar parameters are used in [KK02].

Skeletonization via Reeb graphs. Starting from a noisy sample C of an unknown graph G with a scale parameter, X. Ge et al. [GSBW11] considered the Reeb graph of the Vietoris-Rips complex on a cloud C at a given scale α . The Reeb graph is not intrinsically embedded into any space even if $C \subset \mathbb{R}^2$. The reconstruction in [GSBW11, Theorem 3.1] outputs a graph with a correct homotopy type, while all our derived skeletons $\text{HoPeS}_{k,l}(C)$ also give close geometric approximations in the 2ε -offset of an unknown graph G .

Metric graph reconstruction. M. Aanjaneya et al. [ACC*12] studied a related problem approximating a metric on a large input graph Y by a metric on a small output graph \hat{X} . If Y is a good ε -approximation to an unknown graph X , then [ACC*12, Theorem 2] guarantees the existence of a homeomorphism $X \rightarrow \hat{X}$ that distorts the metrics on X and \hat{X} with a multiplicative factor $1 + c\varepsilon$ for $c > \frac{30}{b}$, where $b > 14.5\varepsilon$ is the length of a shortest edge of X . According to [GSBW11], the algorithm may not run on all inputs, but only for carefully chosen parameters. The skeletons $\text{HoPeS}_{k,l}(C)$ are well-defined for any cloud C and $k, l \geq 1$.

Filamentary structures using Reeb-type graphs. F. Chazal et al. [CHS15] defined the α -Reeb graph G for a metric space X at a user-defined scale α . If X is ε -close to an unknown graph with edges of minimum length 8ε , the output G is $34(\beta(G) + 1)\varepsilon$ -close to X in the Gromov-Hausdorff distance [CHS15, Theorem 4.9] between spaces, not within one space X . Here $\beta(G)$ is the first Betti number of G . The algorithm has the fast time $O(n \log n)$ for n points in X .

Graph Induced Complex GIC. T. Dey et al. [DFW13] built GIC depending on a scale α and a user-defined graph that spans a cloud C . If C is an ε -sample of a good manifold, then GIC has the same homology H_1 as the Vietoris-Rips complex on C at scales $\alpha \geq 4\varepsilon$. Theorem 15 describes graphs G that can be geometrically and topologically approximated from any ε -sample C without extra input parameters.

Skeleton for α -offsets in \mathbb{R}^2 . This work extends [CK13, Kur14b, Kur14a] from locating holes in 2D clouds to a full skeleton. The Gap Search method in section 6 vastly improves [Kur15a, Theorem 7], which was stated for one subskeleton under stronger assumptions on a graph $G \subset \mathbb{R}^2$.

The key advantage of a Homologically Persistent Skeleton $\text{HoPeS}(C)$ is its universal scale-independent structure. In comparison with the persistence diagram of isolated dots (homology classes), $\text{HoPeS}(C)$ shows all persistent cycles directly on a cloud C . In comparison with all algorithms that require a scale α , the skeleton $\text{HoPeS}(C)$ contains a hierarchy of derived skeletons $\text{HoPeS}_{k,l}(C)$ independent of α .

The derived skeletons are most persistent subgraphs of $\text{HoPeS}(C)$ depending on integer indices $k, l \geq 1$, which are easier to choose a posteriori rather than a continuous scale α a priori. We may start from the ‘simplest guess’ $k = l = 1$ and then try $k = 2, l = 1$ without re-running the algorithm, but only selecting a different subgraph of $\text{HoPeS}(C)$.

3. Filtrations of complexes and persistent homology

The shape of a cloud C in any metric space M is usually visualized by fixing a scale $\alpha > 0$ and considering the α -offset $C^\alpha \subset M$ that is the union of closed balls with the radius α and centers at all points of C . When the scale α is increasing, the α -offsets form the filtration $\{C^\alpha\}$ that is the nested sequence $C = C^0 \subset \dots \subset C^\alpha \subset \dots \subset C^{+\infty} = M$. A cloud C is an ϵ -sample of a set $G \subset M$ if $C \subset G^\epsilon$ and $G \subset C^\epsilon$.

Since C^α may have complicated shapes, but continuously deform to simpler Čech complexes [Kur15b, Appendix A]. A faster computable filtration on a cloud C is by the Vietoris-Rips complexes. For any scale α , the complex $\text{VR}(C; \alpha)$ has a k -dimensional simplex on points $v_0, \dots, v_k \in C$ whenever the distance $D(v_i, v_j) \leq 2\alpha$ for all $0 \leq i < j \leq k$.




Figure 2: The Vietoris-Rips complexes $\text{VR}(C; 1.5)$, $\text{VR}(C; 2)$ and $\text{VR}(C; 2.5)$ for the point cloud C in Fig. 1.

$\text{VR}(C; \alpha)$ can be high-dimensional even if $C \subset \mathbb{R}^2$. Any $\text{VR}(C; \alpha)$ is uniquely determined by its 1-dimensional skeleton $\text{VR}_1(C; \alpha)$ whose simplices are spanned by cliques (complete subgraphs) of $\text{VR}_1(C; \alpha)$. The last picture in Fig. 2 shows a plane projection of $\text{VR}(C; 2.5)$ with 2 symmetric tetrahedra having longest edges of length 5.

The reader may consider only the Vietoris-Rips filtration $\{\text{VR}(C; \alpha)\}$. However, our key results work for general complexes and require the concept of a length relative to an ascending filtration $\{Q(C; \alpha)\}$ of complexes on C , where $Q(C; 0) = C$ and $Q(C; \alpha) \subset Q(C; \alpha')$ for any $\alpha \leq \alpha'$.

The length of the edge e between any two points of C relative to $\{Q(C; \alpha)\}$ is the doubled scale α when e enters $Q(\alpha)$, so $|e| = 2 \min\{\alpha : e \in Q(C; \alpha)\}$. If e doesn't enter $Q(C; \alpha)$ for any α , we set $|e| = +\infty$. The condition $Q(C; 0) = C$ implies that the length of any edge between points of C is positive. For both filtrations $\{C^\alpha\}$ and $\{\text{VR}(C; \alpha)\}$, this length $|e|$ coincides with the original metric on the cloud C .

The key idea of Topological Data Analysis is to study topological features such as homology that persist across many scales in a filtration $\{Q(C; \alpha)\}$. The 0-dimensional homology $H_0(Q)$ of a complex Q is the vector space formally generated by connected components of Q . The 1-dimensional homology $H_1(Q)$ is similarly formed by linear combinations of non-trivial 1-dimensional cycles, say with coefficients in $\mathbb{Z}_2 = \{0, 1\}$, see [Kur15b, Appendix A].

For instance, $\text{VR}(1.5)$ in Fig. 2 is one cycle, so $H_1 = \mathbb{Z}_2$ has dimension 1. The complex $\text{VR}(C; 2)$ looks like the character θ with 2 ‘independent’ small cycles, whose ‘sum’ gives the big cycle, hence $H_1 = \mathbb{Z}_2 \oplus \mathbb{Z}_2$ has dimension 2. All cycles in the larger complex $\text{VR}(C; 2.5)$ are contractible, hence $H_1 = 0$. If a complex Q is disconnected, then its homology $H_1(Q)$ is considered as the direct sum of the 1-dimensional homology groups of all connected components of Q .

Increasing the scale α in a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C , we will watch how (homology classes of) 1-dimensional cycles are born and die in $H_1(Q(C; \alpha))$. Any inclusion $f : Q(C; \alpha_i) \subset Q(C; \alpha_j)$ induces a homomorphism $f_* : H_1(Q(C; \alpha_i)) \rightarrow H_1(Q(C; \alpha_j))$. These homomorphisms for $\alpha_i < \alpha_j$ are crucial for defining life intervals (from birth to death) of homological classes below.

For $C \subset \mathbb{R}^2$ in Fig. 1 and $\alpha = 1.5$, when 2 horizontal edges of length 3 enter $\text{VR}(C, 1.5)$, a first cycle appears, so a homology class γ_1 is born at $\text{birth}(\gamma_1) = 1.5$. For $\alpha = 2$, the horizontal edge of length 4 enters $\text{VR}(C, 2)$ and the big cycle splits into 2 smaller cycles. So a homology class γ_2 is born at $\text{birth}(\gamma_2) = 2$. Both γ_1, γ_2 die at $\alpha = 2.5$ when their representative cycles become contractible within $\text{VR}(C; 2.5)$.

Definition 1 (births and deaths) For any filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C in a metric space, a homology class $\gamma \in H_1(Q(C; \alpha_i))$ is born at $\alpha_i = \text{birth}(\gamma)$ if γ is not in the full image under the induced homomorphism $H_1(Q(C; \alpha)) \rightarrow H_1(Q(C; \alpha_i))$ for any $\alpha < \alpha_i$. The class γ dies at $\alpha_j = \text{death}(\gamma) \geq \alpha_i$ when the image of γ under $H_1(Q(C; \alpha_i)) \rightarrow H_1(Q(C; \alpha_j))$ merges into the full image under $H_1(Q(C; \alpha)) \rightarrow H_1(Q(C; \alpha_j))$ for some $\alpha < \alpha_i$.

The births and deaths from Definition 1 are all critical scales $\alpha_1, \dots, \alpha_m$ when the homology group $H_1(Q(C; \alpha))$ changes, so the induced homomorphisms $H_1(Q(C; \alpha_1)) \rightarrow H_1(Q(C; \alpha_2)) \rightarrow \dots \rightarrow H_1(Q(C; \alpha_m))$ are not isomorphisms. This sequence of homomorphisms for the complexes $\text{VR}(C; \alpha)$ in Fig. 2 is $0 \rightarrow \mathbb{Z}_2 \rightarrow \mathbb{Z}_2 \oplus \mathbb{Z}_2 \rightarrow 0$ corresponding to the scales $\alpha = 0, 1.5, 2, 2.5$. The persistence diagram below consists of the pairs (birth, death), which will be called dots to distinguish them from points of a cloud C .

Definition 2 (persistence diagram) Fix a filtration $\{Q(C; \alpha)\}$ of complexes on a set C . Let $\alpha_1, \dots, \alpha_m$ be all critical scales when the homology group $H_1(Q(C; \alpha))$ changes. Let μ_j be the number of independent classes in $H_1(Q(C; \alpha))$ that are born at α_i and die at α_j . The persistence diagram $\text{PD}\{Q(C; \alpha)\} \subset \mathbb{R}^2$ is the multi-set consisting

of all dots $(\alpha_i, \alpha_j) \in \mathbb{R}^2$ with the multiplicities $\mu_{ij} \geq 1$ and all diagonal dots (x, x) with the infinite multiplicity.

For the cloud C in Fig. 1 and the filtration $\{C^\alpha\}$, the homology H_1 has 2 classes that persist over $1.5 \leq \alpha < R$ and $2 \leq \alpha < R$, where $R = \frac{15}{8}\sqrt{17}$ is the circumradius of the largest Delaunay triangle with sides 4, $\sqrt{17}$, 5 in Fig. 2. We use the approximate value $R \approx 2.577$ for simplicity. Hence the persistence diagram $\text{PD}\{C^\alpha\}$ has 2 off-diagonal red dots $(1.5, 2.577)$ and $(2, 2.577)$ in the last picture of Fig. 3.




Figure 3: The α -offsets $C^{1.5}$, C^2 and the 1-dimensional persistence diagram $\text{PD}\{C^\alpha\}$ for the cloud C in Fig. 1.

Using persistence for noisy data is justified by Stability Theorem [EH10, section VIII.2] roughly saying that perturbing a cloud by ϵ perturbs the persistence diagram also by at most ϵ , see [Kur15b, Appendix A]. This stability is crucial for our Reconstruction Theorem 15, while the classical medial axis is unstable for noisy inputs according to [ABE09].

4. A Homologically Persistent Skeleton $\text{HoPeS}(C)$

To motivate the new concept of a Homologically Persistent Skeleton in key Definition 6, we explain below the simpler 0-dimensional case when a Minimum Spanning Tree $\text{MST}(C)$ nicely summarizes the single-edge clusters of a cloud C at all scales α . In sections 4 and 5, we always fix a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C in a metric space M . So MST and HoPeS depend on any filtration $\{Q(C; \alpha)\}$, but we use the simpler notations $\text{MST}(C)$ and $\text{HoPeS}(C)$.

Definition 3 (MST) Fix a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C . A *Minimum Spanning Tree* $\text{MST}(C)$ is a connected graph with the vertex set C and the minimum total length of edges relative to the filtration $\{Q(C; \alpha)\}$. For any scale $\alpha \geq 0$, a *forest* $\text{MST}(C; \alpha)$ is obtained from $\text{MST}(C)$ by removing all open edges that are longer than 2α .

A graph G spans a possibly disconnected complex Q on a cloud C if G has the vertex set C , every edge of G belongs to Q and the inclusion $G \subset Q$ induces a 1-1 correspondence between connected components. Lemma 4 says that $\text{MST}(C)$ for every $\alpha > 0$ contains the shortest forest that identifies all single-edge clusters of C . The *single-edge clustering* of C at a scale $\alpha > 0$ means that points $p, q \in C$ belong to the same cluster if and only if the distance $D(p, q) \leq \alpha$.

Lemma 4 Fix a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C in a metric space. For any fixed scale $\alpha \geq 0$, the forest

$\text{MST}(C; \alpha)$ has the minimum total length of edges among all graphs that span $Q(C; \alpha)$ at the same scale α .

Proof Let $e_1, \dots, e_m \subset \text{MST}(C)$ be all edges longer than 2α , so $\text{MST}(C) = e_1 \cup \dots \cup e_m \cup \text{MST}(C; \alpha)$. Assume that there is a graph G that spans $Q(C; \alpha)$ and is shorter than $\text{MST}(C; \alpha)$. Then the connected graph $G \cup e_1 \cup \dots \cup e_m$ has the vertex set C and is shorter than a Minimum Spanning Tree $\text{MST}(C)$, which contradicts Definition 3. \square

Our first main Theorem 9 extends Lemma 4 from $\text{MST}(C)$ to the skeleton $\text{HoPeS}(C)$ with cycles summarizing all 1-dimensional persistence of C instead of only clusters.

Definition 5 (critical edges) Fix a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C . Each off-diagonal dot $(\alpha_i, \alpha_j) \in \text{PD}\{Q(C; \alpha)\}$ corresponds to a homology class γ that persists in $H_1(Q(C; \alpha))$ over $\alpha_i \leq \alpha < \alpha_j$. The class γ was formed (born) after adding a last edge e to $Q(C; \alpha_i)$. This edge e is called *critical* and has the associated label (α_i, α_j) .

For a fixed scale $\alpha > 0$, the above inequalities $\alpha_i \leq \alpha < \alpha_j$ describe the axes-aligned rectangle in the persistence diagram $\text{PD}\{Q(C; \alpha)\}$ with the bottom right corner at the dot (α, α) . This rectangle contains (birth, death) of all classes that are ‘alive’ at the scale α . That is why at a fixed scale α we will ignore all critical edges e with $\text{death}(e) \leq \alpha$.

Definition 6 (HoPeS) A *Homologically Persistent Skeleton* $\text{HoPeS}(C)$ is the union of a minimum spanning tree $\text{MST}(C)$ and all critical edges with their labels (birth, death).

For any scale $\alpha \geq 0$, the *reduced skeleton* $\text{HoPeS}(C; \alpha)$ is obtained from $\text{HoPeS}(C)$ by removing all edges that are longer than 2α and all critical edges e with $\text{death}(e) \leq \alpha$.

If some distances between points of C are the same, then $\text{MST}(C)$, hence $\text{HoPeS}(C)$, is not unique. The complex $Q(C; 0)$ in any our filtration coincides with a cloud C , so $\text{HoPeS}(C; 0) = C$. By Definition 6 a critical edge e belongs to the reduced skeleton $\text{HoPeS}(C; \alpha)$ if and only if $\text{birth}(e) \leq \alpha < \text{death}(e)$. Hence the homology class born due to the addition of e is alive at the scale α . Namely, any critical edge e is added to $\text{HoPeS}(C; \alpha)$ at $\alpha = \text{birth}(e)$ and will be later removed at the larger scale $\alpha = \text{death}(e)$.

When α is increasing, the sequence of $\text{HoPeS}(C; \alpha)$ may not be monotone. But if $\text{HoPeS}(C; \alpha)$ has become connected, it will stay connected for all larger α . Indeed, removing a critical edge destroys only a cycle, not connectivity.

5. Optimality of reduced skeletons $\text{HoPeS}(C; \alpha)$

In this section we prove first properties of the skeleton $\text{HoPeS}(C)$ including Optimality Theorem 9. Proposition 7 says that the persistence diagram $\text{PD}\{Q(C; \alpha)\}$ of isolated dots can be turned into a graph structure on a cloud C . So the skeleton $\text{HoPeS}(C)$ naturally ‘structurizes’ isolated dots in the 1-dimensional persistence diagram $\text{PD}\{Q(C; \alpha)\}$.

Proposition 7 Fix a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C in a metric space. The 1-dimensional persistence diagram $\text{PD}\{Q(C; \alpha)\}$ can be canonically reconstructed from a full Homologically Persistent Skeleton $\text{HoPeS}(C)$.

Proof By Definitions 5 and 6 any off-diagonal dot $(\alpha_i, \alpha_j) \in \text{PD}\{Q(C; \alpha)\}$ corresponds to a critical edge of $\text{HoPeS}(C)$ with the label (α_i, α_j) . We can easily read the labels on all critical edges of $\text{HoPeS}(C)$ to get all off-diagonal dots in the persistence diagram $\text{PD}\{Q(C; \alpha)\}$. \square

Proposition 8 means that $\text{HoPeS}(C)$ can be used for comparing point clouds up to isometries and uniform scaling in a metric space. Indeed, the definition of the bottleneck distance d_B between persistence diagrams in [Kur15b, Appendix A] is easily re-stated for skeletons with labels (birth, death).

Proposition 8 The topology of a Homologically Persistent Skeleton $\text{HoPeS}(C)$ with all labels is preserved under any isometric transformation of a cloud C . If a cloud C is fixed and a filtration $\{Q(C; \alpha)\}$ is re-parameterized by $\alpha \mapsto \text{constant} \cdot \alpha$, then $\text{HoPeS}(C)$ has the same combinatorial structure, but all labels will be multiplied by the constant.

Proof The complexes $Q(C; \alpha)$ are unaffected by isometries of C . If α is multiplied by a constant, we multiply all births and deaths in $\text{PD}\{Q(C; \alpha)\}$ by the same constant. \square

Theorem 9 naturally extends the optimality of $\text{MST}(C)$ of a cloud C in Lemma 4 from dimension 0 (single-edge clusters of C) to dimension 1 (persistent cycles hidden in C).

Theorem 9 (optimality of reduced skeletons at all scales) Fix a filtration $\{Q(C; \alpha)\}$ of complexes on a cloud C in a metric space. For any fixed scale $\alpha > 0$, the reduced skeleton $\text{HoPeS}(C; \alpha)$ has the minimum total length of edges over all graphs $G \subset Q(C; \alpha)$ that span $Q(C; \alpha)$ and induce an isomorphism in homology $H_1(G) \rightarrow H_1(Q(C; \alpha))$.

Theorem 9 and all further results have proofs in [Kur15b, Appendices B,C]. In section 8 we discuss computations and approximations of $\text{HoPeS}(C)$ in any metric space.

6. The Gap Search strengthens stability of persistence

In sections 6, 7 we consider 1-dimensional persistence diagrams only for α -offsets of a subspace C in a metric space M , where C can be a cloud or a graph $G \subset M$. We strengthen Stability Theorem [EH10, section VIII.2] restricting a bijection between original infinite persistence diagrams to a bijection between finite subdiagrams in Propositions 11, 13.

We will use a Homologically Persistent Skeleton $\text{HoPeS}(C)$ to reconstruct an unknown graph G in a metric space M from a noisy sample C . However, a full skeleton $\text{HoPeS}(C)$ contains too many critical edges that are in a 1-1 correspondence with all off-diagonal dots in the persistence diagram $\text{PD}\{C^\alpha\}$. We will derive smaller skeletons $\text{HoPeS}_{k,l}(C)$ by removing critical edges with (1) a low persistence and (2) a high persistence and large birth times.

Critical edges of type (1) above correspond to homology classes that die shortly after birth. Critical edges of type (2) form cycles that live long, but are born only at large scales α . These large cycles are not present in a graph G given by a sample C , but are formed when distant edges of G start overlapping with each other after a substantial thickening of G . The diagonal gaps quantify our perception of the persistence diagram, because we naturally look for a widest gap to separate persistent features from noisy artefacts.

The persistence diagram $\text{PD}\{C^\alpha\}$ technically includes all diagonal dots $(x, x) \in \mathbb{R}^2$. Below we define smaller subdiagrams of $\text{PD}\{C^\alpha\}$ consisting of finitely many dots.

Definition 10 (gap dgap_k , subdiagram DS_k , scale ds_k) For any subspace C of a metric space, a *diagonal gap* is a strip $\{0 \leq a < y - x < b\}$ that has dots of $\text{PD}\{C^\alpha\}$ in both boundary lines $\{y - x = a\}$ and $\{y - x = b\}$, but not inside the strip. For any integer $k \geq 1$, the k -th *widest diagonal gap* $\text{dgap}_k(C)$ has the k -th largest vertical width $|\text{dgap}_k(C)| = b - a$.

The *diagonal subdiagram* $\text{DS}_k(C) \subset \text{PD}\{C^\alpha\}$ consists of only the dots above the lowest of the first k widest $\text{dgap}_i(C)$, $i = 1, \dots, k$. So each $\text{DS}_k(C)$ is bounded below by $y - x = a$ and has the *diagonal scale* $\text{ds}_k(C) = a$.

In Definition 10 if $\text{PD}\{C^\alpha\}$ has different diagonal gaps with the same width, we say that a lower diagonal gap has a larger width. If $\text{PD}\{C^\alpha\}$ has dots only in m different lines $\{y - x = a_i > 0\}$, $i = 1, \dots, m$, we have exactly m diagonal gaps $\text{dgap}_i(C)$. We ignore the highest gap $\{y - x > \max a_i\}$, so we set $\text{dgap}_i(C) = \emptyset$ and $|\text{dgap}_i(C)| = 0$ for $i > m$.




Figure 4: Subdiagrams $\text{DS}_k(C), \text{VS}_{k,l}(C) \subset \text{PD}\{C^\alpha\}$ from Definitions 10 and 12 for the cloud C in Fig. 1.

The cloud C in Fig. 1 has the persistence diagram $\text{PD}\{C^\alpha\}$ with only 2 off-diagonal dots $(1.5, 2.577)$ and $(2, 2.577)$. Then $\text{dgap}_1(C) = \{0 < y - x < 0.577\}$ is below $\text{dgap}_2(C) = \{0.577 < y - x < 1.577\}$. So $\text{DS}_1(C) = \text{DS}_2(C)$ consists of $(1.5, 2.577), (2, 2.577)$, $\text{ds}_1(C) = \text{ds}_2(C) = 0.577$ in Fig. 4.

Definition 10 makes sense in any persistence diagram, say for α -offsets G^α of a graph $G \subset M$. The graph θ in Fig. 5 has the 1-dimensional persistence diagram $\text{PD}\{\theta^\alpha\}$ containing only one off-diagonal dot $(0, 2.577)$ of multiplicity 2. Then $\text{dgap}_1(\theta) = \{0 < y - x < 2.577\}$, $|\text{dgap}_2(\theta)| = 0$ and $\text{DS}_1(\theta) = \{(0, 2.577), (0, 2.577)\}$ and $\text{ds}_1(\theta) = 2.577$.

Stability Theorem [EH10, section VIII.2] briefly says that, for any ε -sample C of G , there is a bijection between infinite

persistence diagrams $\psi : \text{PD}\{G^\alpha\} \rightarrow \text{PD}\{C^\alpha\}$ so that the L_∞ -distance $\|q - \psi(q)\|_\infty \leq \varepsilon$ for any $q \in \text{PD}\{G^\alpha\}$.

Proposition 11 below restricts such a bijection ψ to finite subdiagrams of only finitely many dots (with high persistence) above the k -th widest diagonal gap in persistence.

Proposition 11 Let C be any ε -sample of a graph G . If $|\text{dgap}_k(G)| - |\text{dgap}_{k+1}(G)| > 8\varepsilon$, there is a bijection $\psi : \text{DS}_k(G) \rightarrow \text{DS}_k(C)$ so that $\|q - \psi(q)\|_\infty \leq \varepsilon$, $q \in \text{DS}_k(G)$.

If we reconstruct cycles of G from a noisy sample C using a close diagram $\text{PD}\{C^\alpha\}$, we should look for dots (birth, death) having a high persistence death – birth and small birth. Hence we search for a vertical gap that separates dots (birth, death) having a small birth and all other dots.

Definition 12 (gap $\text{vgap}_{k,l}$, subdiagram $\text{VS}_{k,l}$, scale $\text{vs}_{k,l}$) In the subdiagram $\text{DS}_k(C)$ from Definition 10 a vertical gap is a widest vertical strip $\{a < x < b\}$ whose boundary contains a dot from $\text{DS}_k(C)$ in the line $\{x = a\}$, but not inside the strip. For $l \geq 1$, the l -th widest vertical gap $\text{vgap}_{k,l}(C)$ has the l -th largest horizontal width $|\text{vgap}_{k,l}(C)| = b - a$. The vertical subdiagram $\text{VS}_{k,l}(C) \subset \text{DS}_k(C)$ consists of only the dots to the left of the first l widest vertical gaps $\text{vgap}_{k,j}(C)$, $j = 1, \dots, l$. So each $\text{VS}_{k,l}(C)$ is bounded on the right by $x = a$ and has the vertical scale $\text{vs}_{k,l}(C) = a$.

In Definition 12 if there are different vertical gaps with the same horizontal width, we say that the leftmost vertical gap has a larger width. So we prefer the leftmost vertical gap, while in Definition 10 we prefer the lowest diagonal gap.




Figure 5: The graph $\theta \subset \mathbb{R}^2$ with the α -offset θ^1 and the 1-dimensional persistence diagram $\text{PD}\{\theta^\alpha\}$.

We allow the case $b = +\infty$, so the widest vertical gap $\text{vgap}_{k,1}(C)$ always has the form $\{x > a\}$, $\text{VS}_{1,1}(C) = \text{DS}_1(C)$ and we set $|\text{vgap}_{k,1}(C)| = +\infty$. If $\text{DS}_k(C)$ has dots in $m \geq 1$ different lines $\{x = b_l \geq 0\}$, $l = 1, \dots, m$, then $\text{DS}_k(C)$ has exactly m vertical gaps $\text{vgap}_{k,l}(C)$.

For the cloud C in Fig. 1, the diagonal subdiagram $\text{DS}_1(C)$ has $\text{vgap}_{1,1}(C) = \{x > 2\}$, $\text{vgap}_{1,2}(C) = \{1.5 < x < 2\}$, $\text{VS}_{1,1}(C) = \{(1.5, 2.577), (2, 2.577)\}$ and $\text{VS}_{1,2}(C) = \{(1.5, 2.577)\}$, so $\text{vs}_{1,1}(C) = 2$, $\text{vs}_{1,2}(C) = 1.5$, see Fig. 4. Any finite cloud C has no cycles at $\alpha = 0$, hence $\text{PD}\{C^\alpha\}$ has no dots in the vertical death axis, so any $\text{vgap}_{k,l}(C)$ has a left boundary line $\{x = a > 0\}$ and all $\text{vs}_{k,l}(C) > 0$.

Definition 12 makes sense for any persistence diagram,

say for α -offsets G^α of a graph G in a metric space. The diagonal subdiagram $\text{DS}_1(\theta)$ for the graph θ in Fig. 5 consists of the doubled dot $(0, 2.577)$. The only vertical subdiagram $\text{VS}_{1,1}(\theta)$ is within the vertical death axis and consists of the same doubled dot $(0, 2.577)$, so $\text{vs}_{1,1}(\theta) = 0$.

Proposition 13 Let $\psi : \text{DS}_k(G) \rightarrow \text{DS}_k(C)$ be a bijection such that $\|q - \psi(q)\|_\infty \leq \varepsilon$ for all dots $q \in \text{DS}_k(G)$ as in Proposition 11. If $|\text{vgap}_{k,l}(G)| - |\text{vgap}_{k,l+1}(G)| > 4\varepsilon$ for some $l \geq 1$, then ψ restricts to a bijection $\text{VS}_{k,l}(G) \rightarrow \text{VS}_{k,l}(C)$ between smaller vertical subdiagrams.

7. Reconstructions by derived skeletons $\text{HoPeS}_{k,l}(C)$

Definition 14 (derived skeletons $\text{HoPeS}_{k,l}$) Let C be a finite cloud in a metric space. For integers $k, l \geq 1$, the derived skeleton $\text{HoPeS}_{k,l}(C)$ is obtained from a full Homologically Persistent Skeleton $\text{HoPeS}(C)$ by removing all edges that are longer than $2\text{vs}_{k,l}(C)$ and keeping only critical edges with (birth, death) $\in \text{VS}_{k,l}(C)$ and with death $> \text{vs}_{k,l}(C)$.

For the cloud C in Fig. 1, we have $\text{vs}_{1,1}(C) = 2$ and $\text{vs}_{1,2}(C) = 1.5$. Hence the derived skeleton $\text{HoPeS}_{1,1}(C)$ coincides with full $\text{HoPeS}(C)$, while $\text{HoPeS}_{1,2}(C)$ misses only the middle edge of length 4. The critical scale $\text{vs}_{k,l}(C)$ is defined in such a way that the derived skeleton $\text{HoPeS}_{k,l}(C)$ is within the reduced skeleton $\text{HoPeS}(C; \text{vs}_{k,l}(C))$.

Our second main Theorem 15 guarantees a reconstruction of a good enough graph G from any noisy ε -sample C up to homotopy equivalence. The homotopy type of a connected graph G is completely determined by its homology $H_1(G)$. Namely, G is homotopy equivalent to a wedge of $\beta(G)$ loops, where the first Betti number is $\beta(G) = \dim H_1(G)$.

Theorem 15 (reconstruction by derived skeletons) Let C be any ε -sample of an unknown graph G in a metric space, so $C \subset G^\varepsilon$ and $G \subset C^\varepsilon$. Let G satisfy the following conditions.

- (1) All cycles $L \subset G$ are ‘persistent’, namely $\text{death}(L) \geq \text{ds}_k(G)$ for some index $k \geq 1$.
- (2) The width $|\text{dgap}_k(G)|$ ‘jumps’, namely $|\text{dgap}_k(G)| - |\text{dgap}_{k+1}(G)| > 8\varepsilon$ for the same k as in (1).
- (3) No cycles are born in α -offsets G^α for ‘small’ $\alpha > 0$, namely $\text{vs}_{k,l}(G) = 0$ for some $l \geq 1$.
- (4) The width $|\text{vgap}_{k,l}(G)|$ ‘jumps’, so $|\text{vgap}_{k,l}(G)| - |\text{vgap}_{k,l+1}(G)| > 4\varepsilon$ for the same k, l as above.

Then we get the lower bound for noise $\text{vs}_{k,l}(C) \leq \varepsilon$ and the derived skeleton $\text{HoPeS}_{k,l}(C) \subset G^{2\varepsilon}$ has the same H_1 as G .

Conditions (1)–(4) of Theorem 15 are stated in terms of the persistence diagram $\text{PD}\{G^\alpha\}$ for simplicity. However, all characteristics like $|\text{dgap}_k(G)|$ and $|\text{vgap}_{k,l}(G)|$ can be defined purely in terms of α -offsets $G^\alpha \subset M$. Indeed, any dot (birth, death) $\in \text{PD}\{G^\alpha\}$ corresponds to the life span of a homology class in $H_1(G^\alpha)$ over birth $\leq \alpha < \text{death}$.

For a fixed graph G , Theorem 15 can be satisfied by many

different indices $k, l \geq 1$. Hence, starting from only a cloud C , we get many approximations by $\text{HoPeS}_{k,l}(C)$ to an unknown graph G within the 2ϵ -offset $G^{2\epsilon} \subset M$. Moreover, we can estimate the closeness of our approximation due to the lower bound $\text{vs}_{k,l}(C) \leq \epsilon$ of an unknown noise level ϵ .

In the proof of Theorem 15 condition (1) implies that the homology class of each cycle $L \subset G$ is captured in the subdiagrams $\text{DS}_k(G), \text{VS}_{k,l}(G)$. The diagonal scales $\text{ds}_k(G)$ are expected to be low, because the diagonal subdiagram $\text{DS}_k(G)$ is above the lowest of the first k widest $\text{dgap}_i(G)$, $i = 1, \dots, k$. If we consider C in Fig. 1 as an ϵ -sample of the graph θ in Fig. 5, then condition (1) holds, because all cycles of θ have death = 2.577 and both $\text{ds}_1(\theta) = \text{ds}_2(\theta) = 2.577$.

Let $G \subset \mathbb{R}^2$ be a distant union of 2 circles with radii $R_1 < R_2$. Then $\text{PD}\{G^\alpha\}$ has only 2 off-diagonal dots $(0, R_1), (0, R_2)$. If $R_2 > 2R_1$, then $\text{dgap}_1(G) = \{R_1 < y - x < R_2\}$, so $\text{ds}_1(G) = R_2$ is larger than the death radius of the 1st circle. Here condition (1) fails for $k = 1$, because the 1st circle is ‘too small’ in comparison with the 2nd circle.

However, the 2nd widest diagonal gap is $\text{dgap}_2(G) = \{0 < y - x < R_1\}$, so $\text{ds}_2(G) = R_1$ and condition (1) holds. Namely both dots $(0, R_1)$ and $(0, R_2)$ are captured in the diagonal subdiagram $\text{DS}_2(G)$. For radii $R_1 < R_2 \leq 2R_1$, condition (1) holds for both possible indices $k = 1, k = 2$.

Conditions (2) and (4) are the same as in Propositions 11, 13, which will imply that the vertical subdiagrams $\text{VS}_{k,l}(G), \text{VS}_{k,l}(C)$ are ϵ -close in the bottleneck distance d_B . All diagonal gaps are ordered by their vertical widths: $|\text{dgap}_1(G)| \geq |\text{dgap}_2(G)| \geq \dots$, where we can set the last width as 0, say for the empty diagonal gap $\{0 < y - x < 0\}$.

So if $\text{PD}\{G^\alpha\}$ isn’t empty, we can always find different successive widths $|\text{dgap}_k(G)| > |\text{dgap}_{k+1}(G)|$. Hence Condition (2) holds for any $\epsilon < \frac{1}{8} (|\text{dgap}_k(G)| - |\text{dgap}_{k+1}(G)|)$. The same arguments apply to condition (4), which certainly holds for $l = 1$. Indeed, the 1st widest vertical gap has the infinite width $|\text{vgap}_{k,1}(G)| = +\infty$ by Definition 12.

If a graph G is a tree without cycles, so $H_1(G) = 0$ and no 1-dimensional homology classes are born at $\alpha = 0$, then $\text{PD}\{G^\alpha\}$ has no dots in the vertical axis and any $\text{vs}_{k,l}(G) > 0$, hence condition (3) fails. Actually, our attempt to approximate a tree G from only a cloud C fails naturally, because the dot closest to the vertical death axis in $\text{PD}\{G^\alpha\}$ is considered as a ‘noisy’ perturbation of the dot $(0, \text{death}(L))$ representing a non-existing cycle L of the tree G .

For a graph G with cycles, we explain why condition (3) isn’t restrictive. Let $\{0 < x < w\}$ be the leftmost vertical gap in the diagonal subdiagram $\text{DS}_k(G)$. Denote by l_{\min} the minimum index with $|\text{vgap}_{k,l}(G)| \leq w$. Then all further vertical gaps can be only thinner, hence the leftmost of any first $l \geq l_{\min}$ vertical gaps is $\{0 < x < w\}$, so $\text{vs}_{k,l}(G) = 0$.

Corollary 16 In the conditions of Theorem 15 if another cloud \tilde{C} is δ -close to C , then the perturbed derived skeleton $\text{HoPeS}_{k,l}(\tilde{C})$ is $(2\delta + 4\epsilon)$ -close to $\text{HoPeS}_{k,l}(C)$.

The skeletons $\text{HoPeS}_{k,l}(C)$ have vertices at all points of C and are locally sensitive to perturbations of C . However, Corollary 16 justifies that $\text{HoPeS}_{k,l}(C)$ are globally stable in the most practical case for noisy samples C of graphs G .

8. Algorithms, experiments, discussion and problems

Fig. 6–8 show the derived skeletons $\text{HoPeS}_{1,1}(C)$, which were computed for clouds $C \subset \mathbb{R}^2$ of n points in time $O(n \log n)$, see details in [Kur15b, Appendix D]. This algorithm uses a duality between 0-dimensional and 1-dimensional persistence for α -complexes obtained from a Delaunay triangulation $\text{Del}(C) \subset \mathbb{R}^2$ [EH10]. Starting from $\text{Del}(C)$, we consider all edges in the decreasing order of length and find all critical edges when two components of $\mathbb{R}^2 - C^\alpha$ merge. At the end $\mathbb{R}^2 - C^\alpha$ becomes a single component, we get $\text{MST}(C)$ and add all critical edges. The simplified graphs are obtained by collapsing edges of half-lengths shorter than $\text{vs}_{k,l}(C) \leq \epsilon$ justified by Theorem 15.




Figure 6: Top: cloud C of 999 points and $\text{PD}\{C^\alpha\}$. Bottom: derived skeleton $\text{HoPeS}_{1,1}(C)$ and its rough simplification.

For a cloud C of n points in any metric space, the 1-dimensional persistence diagram is found [EH10] in time $O(m^3)$, where $m = O(n^3)$ is the size of the largest 2-dimensional complex in a given filtration on C . This algorithm allows us in parallel to record in all critical edges, usually called *positive*, because they create 1D cycles, while edges of $\text{MST}(C)$ are *negative*, because they lead to a merger of components. Hence the skeletons $\text{HoPeS}(C), \text{HoPeS}(C; \alpha), \text{HoPeS}_{k,l}(C)$ are found in the same time.

In Fig. 9–12 clouds C are extracted from images using the Canny edge detector with the low threshold 150 and ratio 3.




Figure 7: Top: cloud C of 999 points and $\text{PD}\{C^\alpha\}$. Bottom: derived skeleton $\text{HoPeS}_{1,1}(C)$ and its rough simplification.

This work was supported by the EPSRC-funded secondment at Microsoft Research Cambridge, UK. A C++ code is at <http://kurlin.org/persistent-skeletons.php>. The author thanks all reviewers for their helpful suggestions and is open for collaboration on problems below and any related projects.

- Simplify $\text{HoPeS}(C)$ to get a smoother and locally stable skeleton that has fewer vertices than C , but still guarantees a close geometric approximation as in Theorem 15.
- Generalize Theorem 15 for unbounded noise similarly to the recent advance in dealing with outliers [BCD*15].

References

- [ABE09] ATTALI D., BOISSONNAT J.-D., EDELSBRUNNER H.: Stability and computation of medial axes - a state-of-the-art. In *Math. Foundations of Scientific Visualization, Comp. Graphics, Massive Data Exploration* (2009), Springer, pp. 109–125. 4
- [ACC*12] AANJANEYA M., CHAZAL F., CHEN D., GLISSE M., GUIBAS L., MOROZOV D.: Metric graph reconstruction from noisy data. *Int. J. Comp. Geometry Appl.* 22 (2012), 305–325. 2
- [BCD*15] BUCHET M., CHAZAL F., DEY T. K., FAN F., OUDOT S., WANG Y.: Topological analysis of scalar fields with outliers. In *Proceedings of SoCG* (2015). 8
- [CHS15] CHAZAL F., HUANG R., SUN J.: Gromov-hausdorff approximation of filament structure using reeb-type graph. *Discrete Computational Geometry* 53 (2015), 621–649. 2
- [CK13] CHERNOV A., KURLIN V.: Reconstructing persistent graph structures from noisy images. *Image-A* 3 (2013), 19–22. URL: <http://munkres.us.es/image-a>. 2




Figure 8: Top: cloud C of 999 points and $\text{PD}\{C^\alpha\}$. Bottom: derived skeleton $\text{HoPeS}_{1,1}(C)$ and its rough simplification.

- [DFW13] DEY T. K., FAN F., WANG Y.: Graph induced complex on data points. In *Proceedings of SoCG* (2013), pp. 107–116. 2
- [EH10] EDELSBRUNNER H., HARER J.: *Computational topology. An introduction*. AMS, Providence, 2010. 4, 5, 7
- [EHZ03] EDELSBRUNNER H., HARER J., ZOMORODIAN A.: Hierarchical morse-smale complexes for piecewise linear 2-manifolds. *Discrete Comput. Geometry* (2003), 87–107. 2
- [GSBW11] GE X., SAFA I., BELKIN M., WANG Y.: Data skeletonization via reeb graphs. In *Proceedings of NIPS* (2011). 2
- [KK02] KÉGL B., KRZYŻAK A.: Piecewise linear skeletonization using principal curves. *Tran. PAMI* 24 (2002), 59–74. 2
- [Kur14a] KURLIN V.: Auto-completion of contours in sketches, maps and sparse 2d images based on topological persistence. In *Proceedings of SYNASC workshop CTIC: Computational Topology in Image Context* (2014), pp. 594–601. 2
- [Kur14b] KURLIN V.: A fast and robust algorithm to count topologically persistent holes in noisy clouds. In *Proceedings of CVPR* (2014), pp. 1458–1463. URL: <http://kurlin.org/projects/holes-in-noisy-clouds.php>. 2
- [Kur15a] KURLIN V.: A homologically persistent skeleton is a fast and robust descriptor for a sparse cloud of interest points in noisy 2d images. In *Proceedings of CAIP* (2015). URL: <http://kurlin.org/projects/hopes.pdf>. 2
- [Kur15b] KURLIN V.: A one-dimensional homologically persistent skeleton of an unstructured point cloud in any metric space (full version of this paper). URL: <http://kurlin.org/projects/persistent-skeleton.pdf>. 3, 4, 5, 7
- [LLT04] LEWINER T., LOPES H., TAVARES G.: Applications of forman's discrete morse theory to topology visualization and mesh compression. *IEEE Transactions on Visualization and Computer Graphics* 10 (2004), 499–508. 2
- [SCP00] SINGH R., CHERKASSKY V., PAPANIKOLOPOULOS N.: Self-organizing maps for the skeletonization of sparse shapes. *IEEE Tran. Neural Networks* 11 (2000), 241–248. 2

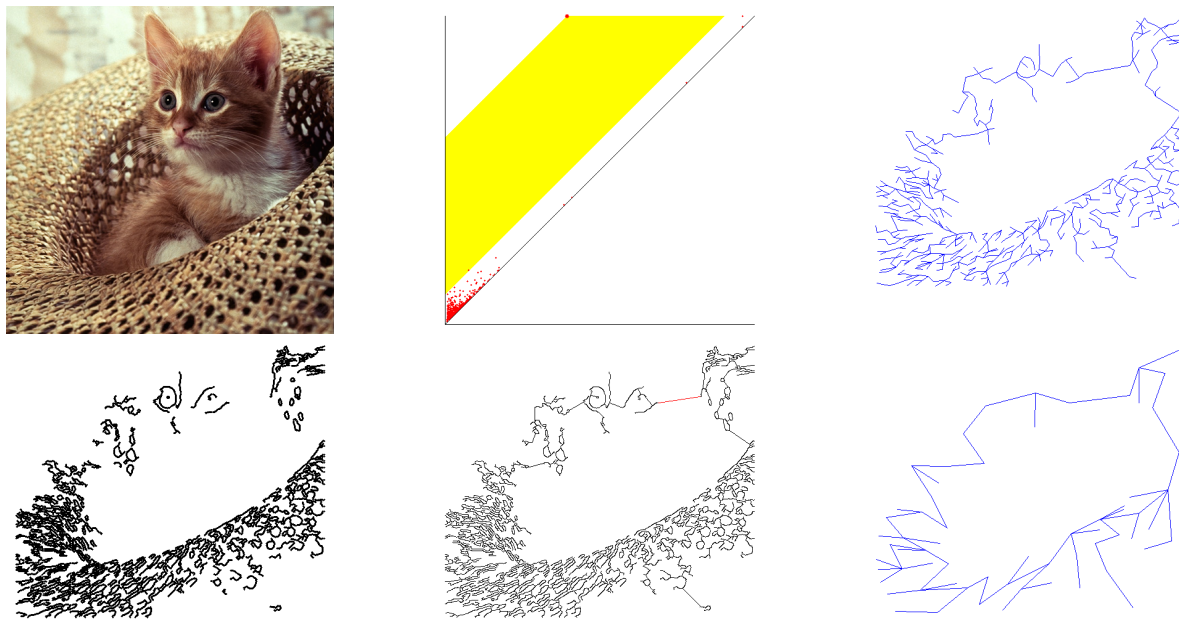


Figure 9: *Left column.* Top: *cat.png*. Bottom: *cloud C* of $n = 14272$ Canny edge points. *Middle column.* Top: persistence diagram $PD\{C^\alpha\}$ with the diagonal gap for $k = 1$ and vertical gap for $l = 1$. Bottom: $HoPeS_{1,1}(C)$ with 14272 edges including 1 red critical edge. Despite many noisy cycles the round hat in the image is successfully detected as the most persistent cycle. *Right column.* Top: simplified skeleton $HoPeS_{1,1}(C)$ with 645 edges. Bottom: simplified $HoPeS_{1,1}(C)$ with 65 edges.

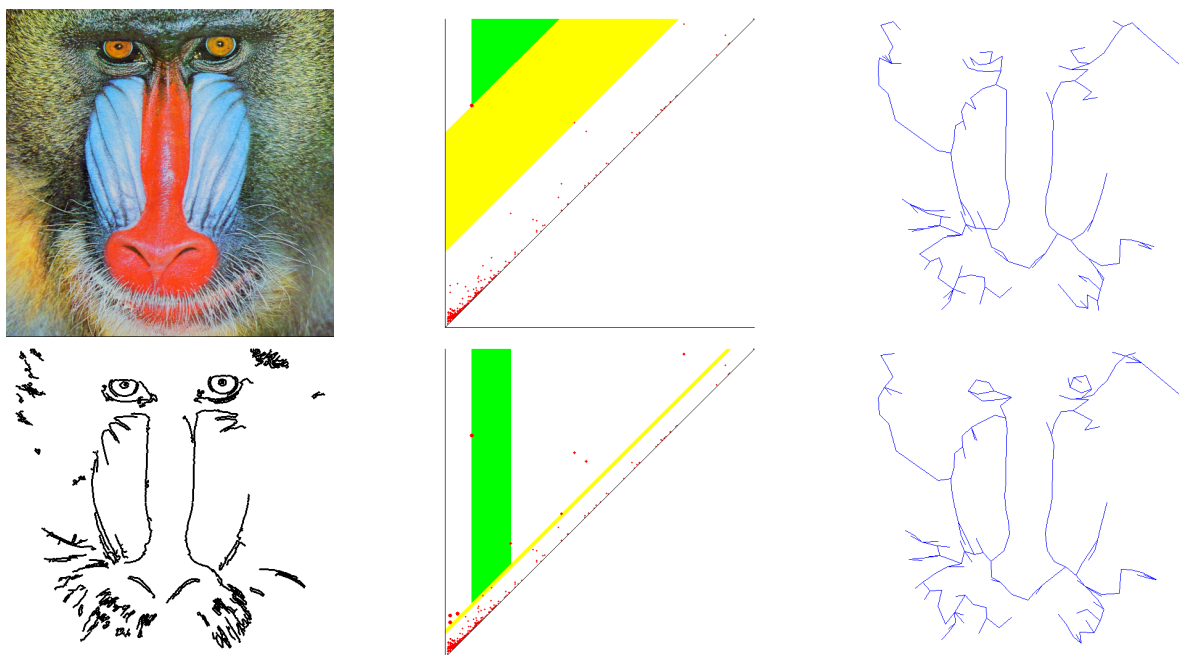


Figure 10: *Left column.* Top: *mandrill.png*. Bottom: *cloud C* of $n = 7490$ Canny edge points. *Middle column.* Top: persistence diagram $PD\{C^\alpha\}$ with colored gaps for $k = 1$ and $l = 1$. Bottom: persistence diagram with colored gaps for $k = 7$ and $l = 4$. *Right column.* Top: simplified skeleton $HoPeS_{1,1}(C)$ with 177 edges. Bottom: simplified $HoPeS_{7,4}(C)$ with 186 edges.




Figure 11: *Left column.* Top: *peppers.png*. Bottom: cloud C of $n = 7128$ Canny edge points. *Middle column.* Top: persistence diagram $\text{PD}\{C^\alpha\}$ with colored gaps for $k = 2$ and $l = 1$ (two largest peppers are detected as two most persistent cycles). Bottom: persistence diagram with colored gaps for $k = 3$ and $l = 1$ (four largest peppers are detected as four most persistent cycles). *Right column.* Top: simplified skeleton $\text{HoPeS}_{2,1}(C)$ with 141 edges. Bottom: simplified $\text{HoPeS}_{3,1}(C)$ with 106 edges.

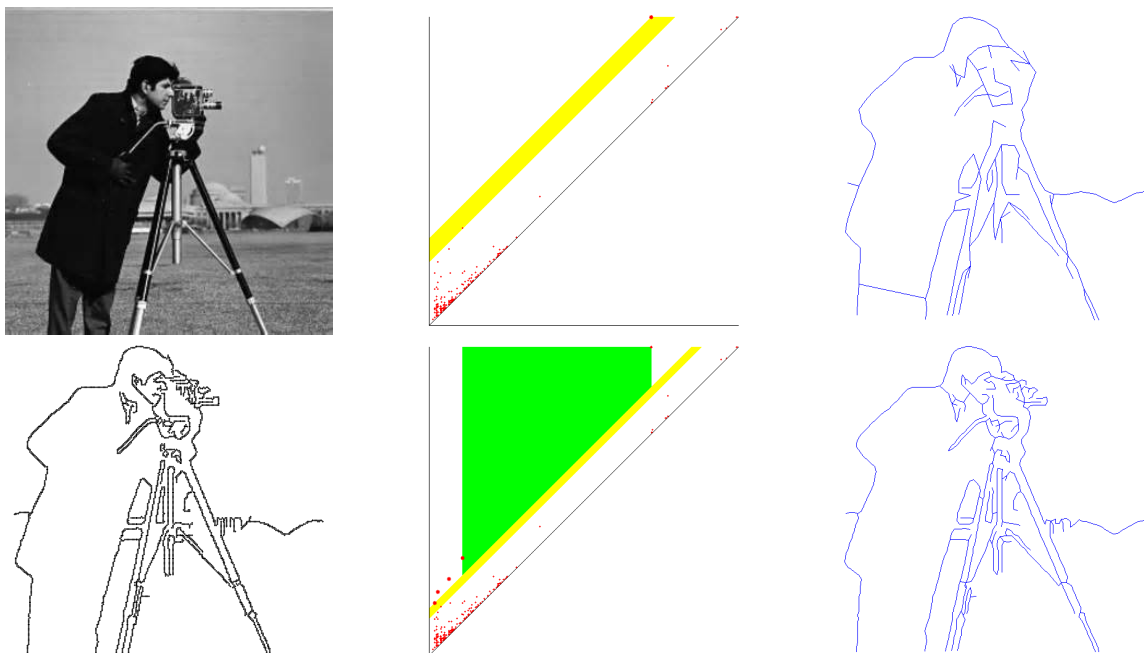


Figure 12: *Left column.* Top: *cameraman.jpg*. Bottom: cloud C of $n = 3151$ Canny edge points. *Middle column.* Top: persistence diagram $\text{PD}\{C^\alpha\}$ with colored gaps for $k = 1$ and $l = 1$. Bottom: persistence diagram with colored gaps for $k = 4$ and $l = 2$. *Right column.* Top: simplified skeleton $\text{HoPeS}_{1,1}(C)$ with 201 edges. Bottom: simplified $\text{HoPeS}_{4,2}(C)$ with 678 edges.