Accepted Manuscript

A fast persistence-based segmentation of noisy 2D clouds with provable guarantees

Vitaliy Kurlin

Graphical Abstract (Optional)

A fast persistence-based segmentation of noisy 2D clouds with provable guarantees
Vitaliy Kurlin

Top: the only input is a cloud of unstructured points.

Bottom: the output is automatically closed regions.

Research Highlights (Required)

To create your highlights, please type the highlights against each \item command.

It should be a short collection of bullet points that convey the core findings of the article. It should include 3 to 5 bullet points (maximum 85 characters, including spaces, per bullet point.)

- A 2D cloud of points is automatically segmented without any extra input parameters.

- The output is a hierarchy of segmentations ordered by their 1D topological persistence.

- The running time is $O(n \log n)$ for a cloud $C \subset \mathbb{R}^2$ of n points with any real coordinates.

- For any ε-sample of a graph $G \subset \mathbb{R}^2$, the boundaries of all regions are 2ε-close to G.

- The publicly available C++ code can automatically segment any real-life images.

A fast persistence-based segmentation of noisy 2D clouds with provable guarantees

Vitaliy Kurlin[a,∗∗]

[a]*Microsoft Research, 21 Station Road, Cambridge CB1 2FB and Mathematical Sciences, Durham University, Durham DH1 3LE, UK.*

ABSTRACT

We design a new fast algorithm to automatically segment a 2D cloud of points into persistent regions. The only input is a dotted image without any extra parameters, say a scanned black-and-white map with almost closed curves or any image with detected edge points. The output is a hierarchy of segmentations into regions whose boundaries have a long enough life span (persistence) in a sequence of nested neighborhoods of the input points. We give conditions on a noisy sample of a graph, when the boundaries of resulting regions are geometrically close to original cycles in the unknown graph.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction: the problem, summary and applications

1.1. Problem: automatic segmentation of a noisy 2D cloud

Recognizing closed contours is a primary phenomenon according to Gestalt laws of perception. Humans can easily form closed loops from incomplete contours. Closed contours bound holes or semantic regions whose extraction is needed for higher level image understanding. So a robust location of holes and their boundaries is a key part of any low level vision system.

The problem is most similar to Saund (2003), who found perceptually closed contours using weights of various quality criteria. Completing closed contours is equivalent to finding regions bounded by these contours. The segmentation problem will be solved in the hardest form without using extra input parameters: *given only a noisy sample C of a graph $G \subset \mathbb{R}^2$, reconstruct closed contours approximating all cycles of G.*

The input is a *cloud C*, which is a finite set of points with any real coordinates in the plane $\mathbb{R}^2$, see Fig. 1. Such a cloud can be a noisy sample or a scan of a hand-drawn sketch or an artistic drawing. The input can also be a binary black-and-white image or a sparse cloud of points without any pixel connectivity.

The output is a hierarchy of segmentations into regions whose boundaries are non-self-intersecting cycles of straight edges connecting points in a given cloud. Fig. 1 shows a cloud C and the most persistent segmentation, which was found by using only C without parameters. All regions bounded by resulting contours are painted with random colors for simplicity.


Fig. 1. Input: a noisy cloud C randomly sampled near a graph $G \subset \mathbb{R}^2$. Output: randomly painted most persistent regions (best viewed in colour).

The algorithm finds the full hierarchy of segmentations in time $O(n \log n)$ for a cloud of n points. The boundary contours in a 1st segmentation of the hierarchy have approximation guarantees for noisy samples of a graph $G \subset \mathbb{R}^2$ in Theorem 11.

In addition to the motivations of Saund (2003) we give four more applications for locating holes and their boundaries.

Auto-closure of polygons. A typical difficulty for users of graphics software is to accurately match endpoints of a polygonal line for painting a resulting region. The suggested approach resolves this difficulty without asking users for extra input.

Hierarchical segmentation. When there is no single ideal segmentation as in most natural images, the algorithm can propose several most likely segmentations, which are rigorously quantified by their stable-under-noise topological persistence.

∗∗Corresponding author: Tel.: +44-1913343081; fax: +44-1913343051;
e-mail: vitaliy.kurlin@gmail.com (Vitaliy Kurlin)

Automatic colorization. Artistic drawings are line sketches that often contain gaps, but may give an impression of closed contours. We may enhance these black-and-white drawings by automatically painting regions that appear visually closed.

Map reading. Conventional paper maps contain many level set curves. These curves are often split into disjoint arcs by labels showing actual heights above the sea level. Hence auto-completion of contours can help faster digitize paper maps.

1.2. Related Work on Topology-based Image Segmentation

Pixel-based segmentation is a traditional approach to find boundary contours in a pixel-based image. Such a segmentation usually minimizes an energy function that contains a cost of assigning a single label and a cost of assigning different labels to neighboring pixels. The resulting minimization problem is often NP-hard. Including higher-order potentials between more than two pixels has even larger computational costs and still encodes only local properties. Chen et al. (2011) suggested the first binary segmentation with global topological constraints, e.g. when a foreground object is connected and has no holes.

Extra parameters are essentially needed for many algorithms including Chernov and Kurlin (2013) and the image segmentation by Letscher and Fritts (2007). They start with a Delaunay triangulation of a point cloud C as in our first step. Then small triangles merge to form persistent regions using two threshold parameters: α for the radius of disks centered at the points of C, and p for the desired level of persistence. We substantially improve this method to avoid any input parameters and prove guarantees by using the stability of 1D persistence.


Fig. 2. A cloud C and its (non-unique) Delaunay triangulation Del(C)

Fast 1D persistence for offsets of 2D point clouds can be computed in time $O(n \log n)$ by the standard algorithm of Attali et al. (2009). This approach was applied by Kurlin (2014b) for counting topologically persistent holes in noisy clouds. However, the 1D persistence diagram contains only unstructured data. A segmentation of a cloud requires more information about adjacency relations of persistent regions. We substantially extend the algorithm from Attali et al. (2009) by adding the new sophisticated data structure Map(α), see section 4.

Homology inference conditions were obtained in many cases to guarantee a correct topological reconstruction from a sample C. In addition to topological guarantees, for the first time Theorem 11 provides a lower bound for the unknown noise level ε and gives conditions on a graph G when reconstructed contours are geometrically 2ε-close to true cycles of G.

The metric graph reconstruction is a related problem solved in Aanjaneya et al. (2012). The input is a large abstract metric graph Y, the output is another smaller metric graph $\hat{X}$.

If Y is a good ε-approximation to an unknown graph X, then $X, \hat{X}$ are homeomorphic. So the input and output are in different spaces, not in the same $\mathbb{R}^2$ as in our segmentation problem.

1.3. Contributions: Parameterless Algorithm and Guarantees

The key differences between the new automatic solution and all the past segmentation methods above are the following.

• We solve the harder problem of completing closed contours or segmenting a cloud of points *without any extra parameters*.

• Input points can have any real coordinates in $\mathbb{R}^2$, the algorithm works for feature points in images at any *subpixel resolution*.

• The algorithm is unsupervised and outputs a hierarchy, where one can get a segmentation with a *required number of regions*.

• The quality of segmentations is measured by the *stable-under-noise persistence*, which leads to guarantees in Theorem 11.

No scale parameters are needed, because a given cloud is analyzed for all values of a radius α. The 1st output consists of those contours whose persistence is above a 1st widest gap in a persistence diagram, see Definition 2 and 5. The 2nd output has contours with a persistence above the lowest of the first 2 widest gaps etc. If exactly k regions are needed, one can select contours corresponding to k cycles with highest persistence.

Data-driven measurements are used for quantifying persistence of contours. If a point cloud C lives in a metric space, we have only a distance function for studying a shape of C. So a natural representation of such a shape is the α-offset C^α, which is the union of disks with a radius $\alpha > 0$ and centers at all points of C. This complicated α-offset C^α continuously deforms to the simpler α-complex $C(\alpha)$ for any α, see Fig. 3 and Definition 1.


Fig. 3. The α-offset C^α deforms to the α-complex $C(\alpha)$ for $\alpha = \sqrt{2}$.

The persistence of a contour is its life span in the sequence of nested offsets C^α. When α is increasing, a contour is *born* at some α = birth and dies in a larger offset at α = death. The *persistence* is death − birth, see Definition 2. For C in Fig. 2, the 1st contour is born at $\alpha = \sqrt{2}$ and persists until all triangles of Del(C) are covered by disks of a larger radius α, see Fig. 3.

Here is a high-level description of the key contributions.

• **Near linear time**: for any n points in $\mathbb{R}^2$, a hierarchy of segmentations is computed in time $O(n \log n)$, see Theorem 9.

• **Guarantees**: if a cloud C is densely sampled from a good enough graph $G \subset \mathbb{R}^2$, all contours of G can be geometrically approximated by using only the cloud C, see Theorem 11.

• **Stability**: the output is globally stable under noise in a cloud C, namely remains in a small offset of C, see Corollary 12.

• **Experiments** on synthetic and natural images in section 6 confirm that the results are robust even for really large noise.

2. Persistent homology of α-complexes $C(\alpha)$ in the plane

The α-offsets C^α have complicated shapes, but continuously deform to α-complexes, which are substructures of a Delaunay triangulation, see Edelsbrunner (1995).

Definition 1 (α-complexes $C(\alpha)$ and holes). *A* Delaunay triangulation Del(C) *on a cloud $C \subset \mathbb{R}^2$ consists of all triangles whose vertices are from C and whose circumcircles enclose no points of C. For any scale $\alpha > 0$, the α-complex $C(\alpha)$ is obtained from* Del(C) *by removing all open edges longer than 2α and all open triangles with circumradii more than α. A hole of $C(\alpha)$ is a bounded connected component of $\mathbb{R}^2 - C(\alpha)$. The boundaries of all holes are* boundary contours *of $C(\alpha)$.*

If $\alpha > 0$ is small, the α-complex $C(\alpha)$ consists of all isolated points of C. For any large enough α, the complex $C(\alpha)$ is Del(C). So the α-complex $C(\alpha)$ is built on isolated points of C by adding edges and triangles at the following critical values:

• an edge between points $p_i, p_j \in C$ is added at $\alpha = \frac{1}{2}d(p_i, p_j)$;

• an *acute* triangle (that has all angles less than $\frac{\pi}{2}$) is added at the critical value α equal to the circumradius of the triangle;

• a non-acute triangle T is added to $C(\alpha)$ at the scale α that is equal to the half-length of the largest side in the triangle T.

All α-complexes form the *filtration* $\{C(\alpha)\}$ that is a nested sequence $C = C(0) \subset \ldots \subset C(\alpha) \subset \ldots \subset C(+\infty) = $ Del(C).


Fig. 4. The α-offset C^α deforms to the α-complex $C(\alpha)$ for $\alpha = 2$.

The summary of topological changes in the filtration $\{C(\alpha)\}$ is quantified by life spans (from birth to death) of holes. For the cloud C in Fig. 2, the first hole is born at the scale $\alpha = \sqrt{2}$ in Fig. 3 and splits into 2 smaller holes at $\alpha = 2$, see Fig. 4.

The left-hand-side hole in Fig. 4 is a square, which shrinks to the triangle with sides $4, 2\sqrt{5}, 2\sqrt{5}$ in $C(\sqrt{5})$, see Fig. 5. The triangle has the circumradius $R_2 = 2.5$ and enters $C(\alpha)$ at the scale $\alpha = 2.5$, hence the hole persists over $2 \le \alpha < 2.5$.

The right-hand-side hole in Fig. 4 splits again at $\alpha = \sqrt{5}$ giving birth to the 3rd triangular hole with sides $2\sqrt{2}, 2\sqrt{5}, 2\sqrt{5}$. This acute triangle has the circumradius $R_3 = \frac{5}{3}\sqrt{2} \approx 2.357$ and enters $C(\alpha)$ at $\alpha = R_3$, so the 3rd hole persists over $\sqrt{5} \le \alpha < R_3$. The remaining hole in $C(2.5)$ contains the Delaunay triangle with the largest circumradius $R_1 = \frac{5}{7}\sqrt{26} \approx 3.642$, see 2. So the most persistent hole lives over $\sqrt{2} \le \alpha < R_1$.

The 1-dimensional homology $H_1(S)$ counts holes of a complex $S \subset \mathbb{R}^2$. Formally, $H_1(S)$ is a vector space generated by 1-dimensional cycles that enclose holes of S. If a connected complex $S \subset \mathbb{R}^2$ consists of V vertices, E edges, F triangles, then


Fig. 5. The α-complexes $C(\alpha)$ of the cloud C in Fig. 2 for $\alpha = \sqrt{5}, 2.5$.

$\chi(S) = V - E + F$ is the *Euler characteristic* and $\beta_1(S) = 1 - \chi(S)$ is the first *Betti* number equal to the number of holes of S.

The complex $C(\sqrt{2})$ in Fig. 3 has $V = E = 17$, $F = 0$, so $\beta_1 = 1$. The vector space $H_1(S)$ of linear combinations with coefficients in the field $\mathbb{Z}_2 = \{0, 1\}$ is determined by $\dim H_1(S) = \beta_1(S)$, so $H_1(S) = \mathbb{Z}_2^{\beta_1(S)}$. The α-complex $C(\sqrt{5}) \subset \mathbb{R}^2$ in Fig. 5 has 3 holes and $H_1(C(\sqrt{5})) = \mathbb{Z}_2^3$.

Since any α-offset C^α continuously deforms to the α-complex $C(\alpha)$, see Edelsbrunner (1995), they have the same homology H_1. In practice, $H_1(C^\alpha)$ is unstable under perturbations of a cloud C. This instability is resolved by the more advanced concept of persistent homology that tracks changes in the homology $H_1(C^\alpha)$ across all scales $\alpha > 0$.

Induced maps in homology. Any inclusion of complexes $\iota : X \to Y$ induces the linear map $\iota_* : H_1(X) \to H_1(Y)$. Namely, let a homology class $\gamma \in H_1(X)$ be represented by a cycle $\sum_i a_i e_i$, where $a_i \in \mathbb{Z}_2$ and e_i are edges of the complex X. Then the class $\iota_*(\gamma) \in H_1(Y)$ is represented by the linear combination $\sum_i a_i e_i$ considered in the different (smaller or larger) space $H_1(Y)$.

Inclusions $C(\alpha_i) \subset C(\alpha_j)$ for any $\alpha_i < \alpha_j$ induce a sequence of linear maps $H_1(C(\alpha_i)) \to H_1(C(\alpha_j))$ decomposable into a sum of elementary sequences over life spans of each hole. The α-complexes for the cloud C in Fig. 2 induce $H_1(C(0)) = 0 \to$

$$\to H_1(C(\sqrt{2})) = \mathbb{Z}_2 \to H_1(C(2)) = \mathbb{Z}_2^2 \to H_1(C(\sqrt{5})) = \mathbb{Z}_2^3 \to$$

$$\to H_1(C(R_3)) = \mathbb{Z}_2^2 \to H_1(C(R_2)) = \mathbb{Z}_2 \to H_1(C(R_1)) = 0,$$

which decomposes as a sum of 3 sequences $0 \to \mathbb{Z}_2 \xrightarrow{\text{id}} \mathbb{Z}_2 \to 0$ over 3 life spans $\sqrt{2} \le \alpha < R_1$, $2 \le \alpha < R_2$ and $\sqrt{5} \le \alpha < R_3$.

Definition 2 (persistence diagram PD$\{C(\alpha)\}$**).** *A homology class $\gamma \in H_1(C(\alpha_i))$ is born at a scale $\alpha_i = $ birth(γ) if γ is not in the image of the map $H_1(C(\alpha)) \to H_1(C(\alpha_i))$ for any $\alpha < \alpha_i$. The class γ dies at $\alpha_j = $ death(γ) $\ge \alpha_i$ when the image of γ under $H_1(C(\alpha_i)) \to H_1(C(\alpha_j))$ merges into the image of $H_1(C(\alpha)) \to H_1(C(\alpha_j))$ for some smaller $\alpha < \alpha_i$. For instance, $\gamma \in H_1(C(\alpha_i))$ dies at α_j if its image in $H_1(C(\alpha_j))$ vanishes. By the elder rule a younger hole born at α_i dies by merging another older hole, which was born earlier at $\alpha < \alpha_i$. All dots* (birth, death) *form the* persistence diagram PD$\{C(\alpha)\} \subset \mathbb{R}^2$.

The filtration $\{C(\alpha)\}$ of α-complexes for the cloud C in Fig. 2 has the persistence diagram in Fig. 6 containing 3 dots (birth, death) for the 3 life spans of holes or their boundaries.

Points near the diagonal have a low persistence death $-$ birth and are considered as noise. If $\mu > 1$ different holes are born and die at the same scales, the corresponding dot (birth, death)

Fig. 6. The persistence diagram of $\{C(\alpha)\}$ **for the cloud** C **in Fig. 2.**

has the *multiplicity* μ. It is safe to add to a persistence diagram all diagonal points (α, α) with infinite multiplicity, because they have persistence 0, so PD$\{C(\alpha)\}$ is a multi-set of dots.

Definition 3 (the bottleneck distance). *Let* $\|p - q\|_\infty = \max\{|x_1 - x_2|, |y_1 - y_2|\}$ *be the distance between* $p = (x_1, y_1)$, $q = (x_2, y_2)$ *in* $\mathbb{R}^2$. *The* bottleneck distance *between persistence diagrams* PD *and* PD′ *is* $d_B = \inf_\psi \sup_{q \in \text{PD}} \|q - \psi(q)\|_\infty$ *over all bijections* $\psi : \text{PD} \to \text{PD}'$ *of multi-sets* PD *and* PD′.

Stability Theorem 4 below is a key foundation of topological data analysis saying that the persistence diagram is stable under perturbations of original data. We quote only a simple version of the Stability Theorem for filtrations by the distance function.

Theorem 4 (stability for ε-samples). *Cohen-Steiner et al. (2007) Let a finite cloud* $C \subset \mathbb{R}^2$ *of points be an ε-sample of a plane graph* $G \subset \mathbb{R}^2$, *namely* $C \subset G^\varepsilon$ *and* $G \subset C^\varepsilon$. *Then the bottleneck distance* $d_B(\text{PD}\{G^\alpha\}, \text{PD}\{C^\alpha\})$ *is at most* ε.

For any cloud $C \subset \mathbb{R}^2$, all 1D homology classes in $H_1(C(\alpha))$ die when the α-complex $C(\alpha)$ becomes the full Delaunay triangulation. Hence the 1D persistence diagram PD$\{C(\alpha)\}$ contains finitely many off-diagonal dots (birth, death) and infinitely many diagonal dots (α, α). After the introduction above, we define the new concept of diagonal gaps in persistence diagrams.

Definition 5 (diagonal gaps and subdiagram PD_k). *For the filtration* $\{C(\alpha)\}$ *of a cloud* $C \subset \mathbb{R}^2$, *a diagonal gap is a largest strip* $\{0 \le a < y - x < b\}$ *that has dots on the boundary lines, but not inside the strip, see Fig. 6. For any integer* $k \ge 1$, *the k-th widest diagonal gap has the k-th largest vertical width* $b - a$. *If there are several widest gaps with the same width, we choose the lowest gap. The k-th subdiagram* $\text{PD}_k \subset \text{PD}\{C(\alpha)\}$ *consists of all dots above the lowest of the first k widest gaps.*

Each k-th widest gap separates dots (birth, death) with a higher persistence death − birth from dots with a lower persistence. The gap between the diagonal birth = death and a dot with a lowest persistence death − birth > 0 is also considered, e.g. this gap is 3rd widest in Fig. 6. Definition 6 introduces a new structure of persistent regions and their boundaries on C.

For the filtration $\{C(\alpha)\}$ of α-complexes, all dots (birth, death) in the full persistence diagram PD$\{C(\alpha)\}$ are in a 1-1 correspondence with all acute triangles in a Delaunay triangulation Del(C). Indeed, each hole dies inside one acute triangle T when α becomes equal to the circumradius of T.

Any non-acute triangle $T \subset \text{Del}(C)$ gives no birth to a hole, because T becomes completely covered by the 3 disks around the vertices of T when the radius α equals the half-length of a longest side, before the scale α reaches the circumradius of T.

Definition 6 (k-th segmentation of C into persistent regions). *To get the* initial segmentation *of a cloud* $C \subset \mathbb{R}^2$, *we recursively merge each acute triangle in* Del(C) *with all non-acute triangles along their longest sides until all non-acute triangles are merged. To get the k-th segmentation, we keep only those initial regions that correspond to* m_k *dots* (birth, death) *in the k-th subdiagram* $\text{PD}_k\{C(\alpha)\}$. *We recursively merge every remaining region along its longest boundary side with its adjacent region until we get exactly* m_k *final persistent regions.*

The persistence diagram in Fig. 6 has only one dot ($\sqrt{2}, R_1$) above the 1st widest gap. The boundary of the corresponding region is the closed contour $C(\sqrt{2})$ in Fig. 3. The remaining 3 right-angled triangles outside the complex $C(\sqrt{2})$ merge with the external region. By Definition 6 the 1st segmentation of C consists of the single region enclosed by $C(\sqrt{2})$, see Fig. 9.

The key information about merged regions is missing in the classical 1D persistence diagram PD$_k\{C(\alpha)\}$ consisting of only dots (birth, death). Hence the new algorithm in section 4 substantially extends the standard persistence computation to get a hierarchy of persistent structures directly on a cloud C.

3. Duality between α-complexes and α-graphs in the plane

We analyze the evolution of contours in the filtration of 2D α-complexes $C(\alpha)$ using the simpler filtration of 1D α-graphs $C^*(\alpha)$ that are dual to $C(\alpha)$. Let us associate a node v_i to every triangle in Del(C), call the external region of the triangulation Del(C) also a 'triangle' and represent it by an extra node v_0. So v_i are abstract nodes shown as small red circles in Fig. 7.

Definition 7 (α-graphs $C^*(\alpha)$ of a cloud C). *Let the metric graph* C^* *dual to* Del(C) *have the nodes* $v_0, v_1, \ldots, v_k$ *and edges of a length* $\frac{1}{2}d_{ij}$ *connecting nodes* v_i, v_j *such that the corresponding triangles in* Del(C) *share a longest side of the length* d_{ij}. *The α-graph* $C^*(\alpha)$ *is obtained from* C^* *by removing all edges not longer than* α. *Any isolated node* v *(except* v_0) *is removed from the graph* $C^*(\alpha)$ *if the corresponding triangle* T_v *in* Del(C) *is not acute or has a small* circumradius rad$(v) \le \alpha$.

The smallest graph $C^*(+\infty)$ is the isolated vertex v_0 corresponding to the external region of Del(C). When α drops from 2.5 to $\sqrt{5}$ in Fig. 7, two isolated nodes v_2 and v_3 enter $C^*(\alpha)$, because rad$(v_2) = R_2 = 2.5$, rad$(v_3) = R_3 = \frac{5}{3}\sqrt{2} > \sqrt{5}$. However, these nodes remain isolated in $C^*(\sqrt{5})$ because all sides of their triangles have half-lengths not longer than $\sqrt{5}$.

The ascending filtration of a Delaunay triangulation Del(C) by α-complexes gives rise to the descending filtration of α-graphs $C^* = C^*(0) \supset \ldots \supset C^*(\alpha) \supset \ldots \supset C^*(+\infty) = \{v_0\}$. Each connected component of $C^*(\alpha)$ has the corresponding region enclosed by a boundary contour of the complex $C(\alpha)$.

For instance, if $C(\alpha)$ contains a triangular cycle, but not the enclosed triangle T_v, then the circumradius rad$(v) > \alpha$ and the

Fig. 7. The complex $C(\alpha)$ and graph $C^*(\alpha)$. Left: $\alpha = 2.5$. Right: $\alpha = \sqrt{5}$.

corresponding node v belongs to $C^*(\alpha)$. So there is a 1-1 correspondence between all isolated nodes (except v_0) of the graph $C^*(\alpha)$ and all triangular boundaries of $C(\alpha)$. This duality extends to all components of $C^*(\alpha)$ and all boundaries of $C(\alpha)$. The following duality result is topological folklore and follows original ideas from Delfinado and Edelsbrunner (1995).

Lemma 8 (duality between C and C^*). *For any scale $\alpha > 0$, all components of the α-graph $C^*(\alpha)$ are in a 1-1 correspondence with all boundary contours of the α-complex $C(\alpha)$.*

Proof. When α is decreasing, the birth of a boundary contour in $C(\alpha)$ means that a small hole appears around the center (of the circumcircle) of an acute triangle $T_v \subset \text{Del}(C)$. By Definition 7 at the same time the isolated node v corresponding to the triangle T_v enters the graph $C^*(\alpha)$ as a new component.

The death of a boundary contour in $C(\alpha)$ means that the contour is torn at its longest edge e for $\alpha = $ half-length of e. If the edge e is shared by triangles T_u, T_v, the corresponding nodes u, v become linked, their components merge in $C^*(\alpha)$.

So there is a 1-1 correspondence between births of contours in $C(\alpha)$ and components in $C^*(\alpha)$, and similarly between their deaths. In general, for any fixed value of α, each component of $C^*(\alpha)$ containing nodes $v_1, \ldots, v_k$ is dual to the contour going along the boundary of the union $T_1 \cup \ldots \cup T_k \subset \text{Del}(C)$ of the triangles represented by the nodes $v_1, \ldots, v_k \in C^*(\alpha)$. $\square$

When α is decreasing from $+\infty$ to 0, the α-complex $C(\alpha) \subset \mathbb{R}^2$ is shrinking, while $C^*(\alpha)$ is growing by Duality Lemma 8. Initially, $C(+\infty) = \text{Del}(C)$ and we show all triangles of $\text{Del}(C)$ in blue. If a triangle disappears from $C(\alpha)$ at a critical value α, for all smaller α we put a small red circle in this triangle.

4. New data structure Map(α) for segmenting 2D clouds

The nice algorithm of Attali et al. (2009) for a fast 1D persistence requires an essential extension for segmenting 2D clouds, because the 1D persistence diagram $\text{PD}\{C^\alpha\}$ has no information about adjacency of required persistent regions. Informally, we will merge all noisy dots in $\text{PD}\{C^\alpha\}$ with a small number of dots above the widest gap, which leads to a final segmentation.

Nodes of the α-graph $C^*(\alpha)$ with attributes are stored in the array Forest(α). Components of $C^*(\alpha)$ describing persistent regions are in the array Map(α). Briefly, the algorithm maintains a union-find structure on Forest(α) and updates adjacency relations in Map(α) when one region mergers another one.

By Definition 7 all nodes of $C^*(\alpha)$ are in a 1-1 correspondence $v \leftrightarrow T_v$ with all triangles of a Delaunay triangulation

Del(C), where the external region of Del(C) is also called a 'triangle'. We call a node $u \in \text{Forest}(\alpha)$ blue if $T_u \in C(\alpha)$, otherwise u is red. Initially for a large α, all nodes are isolated and blue. When α is decreasing, the nodes start turning *red* and join each other to form *red* components of the α-graph $C^*(\alpha)$.


Fig. 8. The complex $C(\alpha)$ and graph $C^*(\alpha)$. Left: $\alpha = 2$. Right: $\alpha = \sqrt{2}$.

Merging two components. When two nodes become linked and their red components merge, a younger component dies. Since α is decreasing, a component is younger if its first node was born at a smaller scale α than the older component. The older component survives by the *elder rule* (Edelsbrunner and Harer, 2010, p. 150), which aims to maximize persistence.

The new larger red component contains the dead nodes from the younger component and the *live* nodes (or corresponding *live* triangles) from the older component. Since α is decreasing, any younger component dies at a smaller value of α than its birth. So the persistence of a red component in $C^*(\alpha)$ and the corresponding boundary contour in $C(\alpha)$ is birth $-$ death > 0.

The *height* of a directed tree with a root is the number of nodes in a longest directed path, so a single node has height 1.

In the structure Forest(α) any node v has these attributes:

- birth(v) = $\sup\{\alpha \mid v \in C^*(\alpha)\} = \sup\{\alpha \mid T_v \subset C(\alpha)\}$;
- uplink(v) is a unique *upward parent* node of $v \in \text{Forest}(\alpha)$;
- height(v) is the height of Tree(v) *going down* from a node v;
- Live(v) is the list of all triangles that are alive in Tree(v);
- bar(v) is the index of the region in Map(α) containing T_v.

Forest(α) with uplinks and heights is a union-find structure. This structure misses adjacency relations between geometric regions and will be extended to the new structure Map(α). Indeed, when a new node u is linked to another node v in Forest(α), it is usually not checked if the corresponding regions are adjacent.

For each acute triangle T_v in Del(C), the corresponding node v has birth(v) = circumradius rad(v) of T_v. For any non-acute triangle T_v, the node v is linked to an existing component when α is the half-length of the longest edge of T_v, so T_v merges with its neighbor. Starting from any node v, we come to root(v) by *going up* along uplinks until the node root(v) is a self-parent.

Then u, v belong to the same component of Forest(α) if and only if root(u) = root(v) as in a union-find structure. When we need to join nodes u, v from different components, we actually join their roots by adding a shorter tree to a taller tree.

Any root keeps the important information about its tree. For instance, the birth time for any node v is extracted as birth(root(v)). To justify the notation *bar*, any dot

$(\text{birth}, \text{death}) \in \mathbb{R}^2$ can be considered as the interval $[\text{birth}, \text{death}) \subset \mathbb{R}$. These intervals or bars form a *bar code* equivalent to PD. A value $\text{bar}(v) = k$ means that the triangle T_v is in the region whose boundary contour was the k-th to die.

Structure $\text{Map}(\alpha)$ **of persistent regions.** Any region R consists of triangles whose corresponding nodes $v_1, \ldots, v_k$ form a connected component of $C^*(\alpha)$, and R has these attributes:

• $\text{ind}(R)$ is the index of the region R in the array $\text{Map}(\alpha)$;

• $\text{birth}(R) = \text{scale } \alpha$ when a 1st triangle enters the region R;

• $\text{death}(R) = \alpha$ when R mergers with an older superior region;

• $\text{Core}(R)$ is a list of nodes whose triangles form the region R;

• $\text{sup}(R)$ is the index of the superior region merged with R.

• $\text{heir}(R)$ is the node $v \notin \text{Core}(R)$ adjacent to a node $u \in \text{Core}(R)$ such that linking u, v mergers R and its superior region (of v);

The input is a set of n points given by real coordinates $(x_1, y_1), \ldots, (x_n, y_n) \in \mathbb{R}^2$. We start by finding the Delaunay triangulation $\text{Del}(C)$ in time $O(n \log n)$ with $O(n)$ space.

Initialization. We set $\text{birth}(v_0) = +\infty$ for the external node $v_0 \in \text{Forest}(\alpha)$. After finding $\text{Del}(C)$, we go through each triangle T_v of $\text{Del}(C)$ and set the birth of the corresponding node $v \in \text{Forest}(\alpha)$ as the circumradius $\text{rad}(v)$ for acute T_v and $\text{birth}(T_v) = 0$ for non-acute T_v. All $\text{bar}(v)$ have the initial value 0 meaning that the bar indices are undefined. All arrays $\text{Live}(v)$ and $\text{Map}(\alpha)$ are empty. We sort all edges of the triangles of $\text{Del}(C)$ in the decreasing order starting from the longest edge of a length d. We start from the initial largest value $\alpha = \frac{1}{2}d$.

The 'while' loop goes through each edge e of $\text{Del}(C)$ in the decreasing order of length until $\text{Forest}(\alpha)$ becomes connected. Let the edge e be shared by adjacent triangles T_u, T_v. If α is going down through the critical value equal to the half-length of e, we link the corresponding nodes u, v in the α-graph $C^*(\alpha)$. This addition doesn't affect $\text{Forest}(\alpha)$ if u, v were already connected, namely they have a common root as in Case 1 below.

It is possible that one of the nodes, say u, was not included in $\text{Forest}(\alpha)$ at the initialization stage, because the corresponding triangle T_u is not acute, so u was blue. In this Case 2 discussed below we link the single node u to the red component of v.

If the nodes u, v are in different trees of $\text{Forest}(\alpha)$, we merge these trees by linking their roots in Case 3. Final Case 4 below studies the exception when both triangles T_u, T_v are not acute. Then T_u, T_v are right-angled and have the common hypotenuse e, otherwise T_v is enclosed by the circumcircle of T_u, which is forbidden in a Delaunay triangulation $\text{Del}(C)$ by Definition 1.

At the end $\text{Map}(\alpha)$ will contain adjacency relations of all regions in addition to the 1D persistence diagram $\text{PD}\{C(\alpha)\}$ as pairs (birth, death). Each region R will have the index $\text{sup}(R)$ of its adjacent more persistent region that merged with R.

Case 1: the edge e has the same region on both sides, namely the neighboring triangles T_u, T_v sharing the edge v have the same root $\text{root}(u) = \text{root}(v)$. This value of α is not critical, because both nodes u, v were already connected in $\text{Forest}(\alpha)$.

Case 2: e is the longest edge of a non-acute triangle T_u and another triangle T_v with $v \in \text{Forest}(\alpha)$. We find $\text{root}(v)$ and add u to $\text{Forest}(\alpha)$ setting $\text{uplink}(u) = \text{root}(v)$, $\text{birth}(u) =$

$\text{birth}(\text{root}(v))$. We increase $\text{height}(\text{root}(v))$ by 1 only if it was 1. If $\text{bar}(v)$ is defined, then the node v belongs to the already dead region $R \in \text{Map}(\alpha)$ with $\text{ind}(R) = \text{bar}(v)$. Hence u joins this region R and we set $\text{bar}(u) = \text{bar}(v)$. If $\text{bar}(v)$ is undefined, then T_v is a live triangle and we add u to $\text{Live}(\text{root}(v))$.

Case 3: both $\text{birth}(\text{root}(u)), \text{birth}(\text{root}(v)) > 0$. Then the two components of $\text{Forest}(\alpha)$ containing the nodes u, v merge. Assume that $\text{birth}(\text{root}(u)) \leq \text{birth}(\text{root}(v))$, so the component of u is younger, hence dies. We create a new region R in $\text{Map}(\alpha)$, say with an index i, by setting $\text{birth}(R) = \text{birth}(\text{root}(u))$, $\text{death}(R) = \alpha$ (the current value equals the half-length of e). All nodes $w \in \text{Live}(\text{root}(u))$ die and we set $\text{bar}(w) = i$ for them.

If $\text{bar}(v)$ is already defined, the component of v has died and we can set $\text{sup}(R) = \text{bar}(v)$. Otherwise the component of v will die later and we set $\text{heir}(R) = v$ remembering to update $\text{sup}(R)$ using $\text{heir}(R)$ after the 'while' loop is finished. Then we copy the list $\text{Live}(\text{root}(u))$ to $\text{Core}(R)$ so that the region R knows all its nodes (triangles) that were alive just before R died.

Subcase 3a: $\text{height}(\text{root}(u)) \leq \text{height}(\text{root}(v))$. Then we link $\text{root}(u)$ of the shorter tree to $\text{root}(v)$ of the taller tree to keep to maximum height of all trees in $\text{Forest}(\alpha)$ minimal, so $\text{root}(v)$ becomes $\text{uplink}(\text{root}(u))$. If the heights were equal, then $\text{height}(\text{root}(v))$ jumps up by 1. All live triangles from $\text{Live}(\text{root}(v))$ are kept at the root of the new larger tree.

Subcase 3b: $\text{height}(\text{root}(u)) > \text{height}(\text{root}(v))$. Then we link $\text{root}(v)$ of the shorter tree to $\text{root}(u)$ of the taller tree. Hence $\text{root}(u)$ becomes $\text{uplink}(\text{root}(v))$, but $\text{height}(\text{root}(u))$ remains the same. We should keep all triangles from $\text{Live}(\text{root}(v))$ at the root of the new tree replacing $\text{Live}(\text{root}(u))$ by $\text{Live}(\text{root}(v))$. So it is important to save $\text{Live}(\text{root}(u))$ in $\text{Core}(R)$ as we did.

Case 4: $\text{birth}(\text{root}(u)) = 0 = \text{birth}(\text{root}(v))$ means that both triangles T_u, T_v are not acute and share their longest edge. This is possible only if T_u, T_v are right-angled with a common hypotenuse. Hence u, v form a new red component of $C^*(\alpha)$.

The external region in $\text{Map}(\alpha)$. When $\text{Forest}(\alpha)$ becomes connected, it remains to add to $\text{Map}(\alpha)$ the last entry R corresponding to the external region of $\text{Del}(C)$. The last root v has the list $\text{Live}(v)$ containing the node v_0. Similarly to Case 3, $\text{Live}(v)$ is copied to $\text{Core}(R)$ and we set $\text{bar}(w)$ equal to the index $\text{ind}(R)$ for any node $w \in \text{Live}(v)$, but $\text{heir}(R)$ is not needed.

Initial segmentation. Each triangle from $\text{Del}(C)$ contributes to a single component of $\text{Map}(\alpha)$, namely all lists $\text{Core}(R)$ are disjoint. Hence $\text{Map}(\alpha)$ contains $m = O(n)$ entries and can be sorted in time $O(n \log n)$ in the decreasing order of $\text{pers} = \text{birth} - \text{death}$. We output the initial segmentation where all triangles from $\text{Core}(R)$ have a color associated with R. The 1st picture in Fig. 9 shows 3 initial regions corresponding to all 3 dots in the 1D persistence diagram $\text{PD}\{C^\alpha\}$ from Fig. 6.

The k-th widest gap in persistence. We find the k-th gap between decreasing persistences of regions in sorted $\text{Map}(\alpha)$ in time $O(n \log n)$. In the conditions of Theorem 11, the 1st gap separates m dots (birth, death) corresponding to true cycles of a graph G. Even if the conditions do not hold, the widest gap gives an approximation to the expected number m of regions, which can also be set by a user, for a final segmentation below.

Algorithm 1 Build Map(α) of persistent regions in a cloud C

1: Input: a cloud C of n points $(x_1, y_1), \ldots, (x_n, y_n)$
2: Compute Del(C) with k triangles on the n points of C
3: Sort edges of Del(C) in the decreasing order of length
4: Forest $\leftarrow$ isolated nodes $v_0, \ldots, v_k$ with all birth times 0 except birth(v_0) $\leftarrow +\infty$ and for each acute triangle $T_v \subset$ Del(C) we update birth(v) $\leftarrow$ circumradius of T_v
5: Set the number of links in Forest(α): $l \leftarrow 0$
6: **while** $l < k$ (stop when Forest(α) becomes a tree) **do**
7:     Take the next longest edge e, set $\alpha \leftarrow \frac{1}{2}$length($e$)
8:     Find u, v dual to the triangles T_u, T_v that share e
9:     Find root(u), root(v) going along uplinks from u, v
10:     **if** root(u) = root(v) **then** Case 1: no changes, hence
11:         continue the loop without increasing the number l.
12:     **end if**
13:     **if** birth(root(u)) = 0 and birth(root(v)) > 0 **then**
14:         (there is also a symmetric case with u, v swapped)
15:         Case 2: (u is blue, v is red) run Algorithm 2 below.
16:         The number of links $l \leftarrow l + 1$, continue the loop.
17:     **end if**
18:     **if** $0 <$ birth(root(u)) $\leq$ birth(root(v)) **then**
19:         (u younger than v, another case with u, v swapped)
20:         Case 3: components of u, v merge, run Algorithm 3
21:         The number of links $l \leftarrow l + 1$, continue the loop.
22:     **end if**
23:     **if** birth(root(u)) = 0 = birth(root(v)) **then**
24:         Case 4: (the triangles T_u, T_v are right-angled)
25:         Set birth(u) = birth(v) $\leftarrow \alpha$, uplink(v) $\leftarrow u$,
26:         height(u) $\leftarrow$ height(u)+1, add the node v to Live(u)
27:         The number of links $l \leftarrow l + 1$, continue the loop.
28:     **end if**
29: **end while**
30: Return array Map(α) of regions generated by Case 3

Algorithm 2 : link the node u to the component of the node v

1: Set uplink(u) $\leftarrow$ root(v), birth(u) $\leftarrow$ birth(root(v))
2: **if** height(root(v)) = 1 **then** height(root(v)) $\leftarrow$ 2
3: **end if**
4: **if** bar(v) is already defined **then** set bar(u) $\leftarrow$ bar(v)
5: **else** Add u to Live(root(v)) in the subtree at root(v)
6: **end if**

Algorithm 3 : the younger component becomes a region R

1: Create a new region $R \in$ Map(α) of the younger node u
2: Set birth(R) $\leftarrow$ birth(root(u)), death(R) = α, Core(R) $\leftarrow$ Live(root(u)), bar(w) $\leftarrow$ ind(R) for each $w \in$ Live(root(u))
3: **if** bar(v) is defined **then** sup(R) = bar(v)
4: **else** heir(R) = v (the region of v will enter Map later)
5: **end if**
6: **if** height(root(u)) $\leq$ height(root(v)) **then**
7:     uplink(root(u)) $\leftarrow$ root(v)
8:     height(root(v)) + + if height(root(u)) = height(root(v))
9: **else** uplink(root(v)) $\leftarrow$ root(u)
10:     Live(root(u)) $\leftarrow$ Live(root(v))
11:     birth(root(u)) $\leftarrow$ birth(root(v))
12: **end if**


Fig. 9. Initial segmentation: 3 regions. Final segmentation: 1 region.

Indices sup(R) **of superior regions.** We go through each region $R \in$ Map(α) and build the 1-1 correspondence old index $\mapsto$ new index in sorted Map(α). We go again through each R and access the node heir(R) whose bar index bar(heir(R)) is the original (non-sorted) index of the *superior* region that merged with R. Using the 1-1 correspondence of indices in Map(α) above, we know the new index sup(R) of this *superior* or more persistent region with a higher persistence that merged with R.

Final segmentation into m regions. Now we form m regions, where m can be the number of dots in the 1st (or k-th) subdiagram of PD$\{C^\alpha\}$ or m can be user-defined. We go through all regions R of sorted Map(α) starting from the least persistent region. If the current index ind(R) > m, we add the list Core(R) to Core(sup(R)), which enlarges the superior region that has the index sup(R) and a higher persistence. If ind(sup(R)) > m, the region sup(R) will also merge with its superior later. After merging all regions of ind(R) > m with their superiors one by one, we output lists Core(R) of triangles for ind(R) = $1, \ldots, m$.

5. Fast running time and reconstruction guarantees

The main results are Theorem 9, guarantees for boundaries in Theorem 11 and global stability of output in Corollary 12.

Theorem 9 (fast computation of persistent contours). *For any point cloud C of n points in the plane, the algorithm in section 4 computing the full hierarchy of segmentations has the time complexity $O(n \log n)$ and memory space $O(n)$.*

Proof of Theorem 9. A Delaunay triangulation Del(C) for a cloud $C \subset \mathbb{R}^2$ of n points has $l = O(n)$ triangles and is found in time $O(n \log n)$ (de Berg et al., 2008, section 9.1). The 'while' loop in Algorithm 1 goes once through not more than $O(n)$ edges in Del(C). We can associate to each edge e its two incident triangles $T_u, T_v \subset$ Del(C) in advance, so identifying the corresponding nodes u, v in line 8 is easy. We prove that finding root(u), root(v) in line 9 by going along uplinks in any tree of Forest(α) with l nodes requires $O(\log l) = O(\log n)$ steps.

The height of a tree can increase only in Case 2 (from 1 to 2) or in Case 3, where the height jumps by 1 after we link two trees of the same height. In Case 3 we always link a shorter tree to a taller one. So any two paths in a tree from a root to terminal nodes (leaves) can differ by at most 1, where we include all trivial paths consisting of a single node.

Hence almost any node is linked to at least nodes one level down, except all terminal nodes and some nodes only one level up. Then any tree of height $h \geq 1$ should contain at least 2^{j-1} nodes at level $1 \leq j \leq h - 1$ plus at least one node at level h,

so at least $1 + 2 + 2^2 + \cdots + 2^{h-2} + 1 = 2^{h-1}$ nodes in total. If $l \geq 2^{h-1}$, then the height is $h \leq 1 + \log_2 l = O(\log l)$.

All other steps in the 'while' loop from section 4 need $O(1)$ time. Then we spend $O(n \log n)$ time for sorting $O(n)$ entries in Map(α). The lists Live(v) of triangles are disjoint in Forest(α) as well as similar lists Core(R) in Map(α). Then $O(n)$ Delaunay edges, triangles or corresponding nodes with attributes need only $O(n)$ space in Forest(α) and similarly in Map(α). $\square$

Continuous maps between spaces $f_0, f_1 : X \to Y$ are called *homotopic* if they can be included into a continuous family of maps $f_t : X \to Y$, $t \in [0, 1]$. A set $S \subset \mathbb{R}^2$ is *contractible* to a point $q \in S$ if the identity id : $S \to S$ is homotopic to $S \to q$.

Let L be any closed non-self-intersecting loop in $\mathbb{R}^2$. We consider all offsets L^α when α is increasing. If $\alpha > 0$ is small, L^α is the thickening of L with the radius α. There is a first critical scale α when the internal boundary of L^α touches itself, so L^α is no longer a topological annulus. There is a last critical scale α when the internal boundary of L^α shrinks to a point, so L^α becomes contractible, which is true for all larger radii α.


Fig. 10. The 'heart' graph has $\theta = 0$. The 'figure-eight' graph has $\theta > 0$.

A round disk is contractible to its center, but we can not contract a circle to one point always staying within the circle. Below we consider closed loops that go along boundaries of regions in $\mathbb{R}^2 - G$. The internal loop of the graph $\ominus$ goes along the short vertical edge twice, first up and then down.

Definition 10 (radius ρ of a cycle, thickness θ of a graph G). *A cycle L of a graph $G \subset \mathbb{R}^2$ is basic if L is non-self-intersecting and encloses a bounded region of $\mathbb{R}^2 - G$. When α is increasing, the hole enclosed by the α-offset L^α is born at $\alpha = 0$ and dies at the scale $\alpha = \rho(L)$ that is called the radius of the cycle L. So the hole enclosed by L has the life span $0 \leq \alpha < \rho(L)$.*

In general, when α is increasing, new holes can be born in G^α, let they be enclosed by $L_1, \dots, L_k$ at their birth times. The thickness $\theta(G) = \max\limits_{j=1,\dots,k} \rho(L_j)$ is the maximum persistence of these smaller holes born during the evolution of offsets G^α. If no such holes appear, then $\theta = 0$, otherwise $\theta > 0$, see Fig. 10.

If a cycle $L \subset \mathbb{R}^2$ encloses a convex region, then the only hole of L^α completely dies when α is the radius $\rho(L)$, so $\theta(L) = 0$.

The heart-shaped cycle L in the first picture of Fig. 10 encloses a non-convex region, however no new holes are born in L^α, so $\theta(L) = 0$. The figure-eight-shaped graph G in the second picture of Fig. 10 has a positive thickness equal to the radius $\rho(L_1)$ of the largest cycle born in G^α when α is increasing.

A cloud C is an *ε-sample* of a graph $G \subset \mathbb{R}^2$ if $C \subset G^\varepsilon$ and $G \subset C^\varepsilon$. Theorem 11 gives conditions on C and G when the 1st segmentation in the output hierarchy for C has boundary contours close to all basic cycles of the unknown graph G.

Theorem 11 (guarantees for the 1st segmentation). *Let C be any ε-sample of a connected graph $G \subset \mathbb{R}^2$ with a thickness $\theta(G) \geq 0$ and $m \geq 1$ basic cycles having ordered radii $\rho_1 \leq \dots \leq \rho_m$. If $\rho_1 > 7\varepsilon + \theta(G) + \max\limits_{i=1,\dots,m-1} \{\rho_{i+1} - \rho_i\}$, then the 1st segmentation of C from Definition 6 has exactly m regions whose boundary contours are in the 2ε-offset $G^{2\varepsilon} \subset \mathbb{R}^2$. If $\alpha(C)$ is maximum birth over all dots in $\mathrm{PD}_1\{C(\alpha)\}$, then $\alpha(C) \leq \varepsilon$.*

The cloud C in Fig. 2 can be considered as an ε-sample of the graph $G = C(\sqrt{2})$ in Fig. 3 for $\varepsilon = \sqrt{2}$. Then $\theta(G) = R_2 - 2 = 0.5$ is the maximum persistence of a new hole born in G^α. Since G has only one radius $\rho_1 = R_1 \approx 3.642$, which is the minimum scale when G^α becomes contractible, the inequality of Theorem 11 fails. However PD$\{C^\alpha\}$ in Fig. 6 above the widest gap has 1 dot corresponding to one region enclosed by G in Fig. 9. So the algorithm may give a correct reconstruction beyond the guarantees of Theorem 11 when a noisy sample is 'uniform'.

Proof of Theorem 11. The homology $H_1(G)$ is generated by m basic cycles $L_1, \dots, L_m$ that enclose m holes (bounded regions in the complement $\mathbb{R}^2 - G$). These m cycles give dots $(0, \rho_i)$ in the vertical axis of the 1D persistence diagram $\mathrm{PD}\{G^\alpha\}$.


Fig. 11. Left: 1D persistence diagram $\mathrm{PD}\{G^\alpha\}$ for a graph G satisfying Theorem 9. Right: perturbed diagram $\mathrm{PD}\{C^\alpha\}$ for a noisy ε-sample C of G.

All other dots in the 1D persistence diagram $\mathrm{PD}\{G^\alpha\}$ come from smaller holes in the α-offsets G^α that were born later. The maximum persistence death $-$ birth of these holes is bounded above by the thickness $\theta(G)$, see Definition 8 and Fig. 11.

The given inequality $\rho_1 > 7\varepsilon + \theta(G) + \max\limits_{i=1,\dots,m-1} \{\rho_{i+1} - \rho_i\}$ guarantees that the widest diagonal gap $\{\theta(G) < y - x < \rho_1\}$ in $\mathrm{PD}\{G^\alpha\}$ is wider than any other gaps including the higher gaps $\rho_{i+1} - \rho_i$ between the dots $(0, \rho_i) \in \mathrm{PD}\{G^\alpha\}$, $i = 1, \dots, m - 1$.

By Stability Theorem 4 the perturbed diagram $\mathrm{PD}\{C(\alpha)\}$ is in the ε-offset of $\mathrm{PD}\{G^\alpha\} \subset \cup_{i=1}^m (0, \rho_i) \cup \{y - x < \theta(G)\}$ with respect to the L_∞ metric on $\mathbb{R}^2$. All noisy dots near the diagonal in $\mathrm{PD}\{C(\alpha)\}$ can not be higher than $\theta(G) + 2\varepsilon$ after projecting along the diagonal $\{x = y\}$ to the vertical axis $\{x = 0\}$.

The remaining dots can not be lower than $\rho_1 - 2\varepsilon$ after the same projection $(x, y) \mapsto y - x$. Hence the smaller diagonal strip $\{\theta(G) + 2\varepsilon < y - x < \rho_1 - 2\varepsilon\}$ of the vertical width $\rho_1 - 4\varepsilon - \theta(G)$ is still empty in the perturbed diagram $\mathrm{PD}\{C(\alpha)\}$.

By Stability Theorem 4 any dot $(0, \rho_i) \in \mathrm{PD}\{G^\alpha\}$, $i \geq 2$, can not jump lower than the line $y - x = \rho_i - 2\varepsilon$ or higher than $y - x = \rho_i + \varepsilon$. Then the widest diagonal gap between these perturbed dots has a vertical width at most $\max\limits_{i=1,\dots,m-1} \{\rho_{i+1} - \rho_i\} + 3\varepsilon$.

Since all dots near the diagonal have diagonal gaps not wider than $\theta(G) + 2\varepsilon$, the 2nd widest gap in the perturbed diagram $\text{PD}\{C(\alpha)\}$ always has a vertical width smaller than $\rho_1 - 4\varepsilon - \theta(G)$. Hence the 1st widest gap in $\text{PD}\{C(\alpha)\}$ covers the diagonal strip $\{\theta(G) + 2\varepsilon < y - x < \rho_1 - 2\varepsilon\}$, which is within the 1st widest gap $\{\theta(G) < y - x < \rho_1\}$ in the original diagram $\text{PD}\{G^\alpha\}$.

Then the subdiagram $\text{PD}_1\{C(\alpha)\}$ above the line $y - x = \rho_1 - 2\varepsilon$ contains exactly m perturbations (b_i, d_i) of the original dots $(0, \rho_i)$ in the vertical strip $\{0 \leq x \leq \varepsilon\}$. Hence the 1st segmentation of C from Definition 6 has exactly m final regions corresponding to the m dots in the subdiagram $\text{PD}_1\{C(\alpha)\}$.

It remains to prove that the boundary contours ∂C in the 1st segmentation of a cloud C are 2ε-close to G. All dots in $\text{PD}_1\{C(\alpha)\}$ are at most ε away from their corresponding dots $(0, \rho_i) \in \text{PD}\{G^\alpha\}$. Hence the critical scale $\alpha(C)$ is at most ε.

A longest boundary edge $e \subset \partial C$ of any region corresponding to a dot $(b_i, d_i) \in \text{PD}_1\{C(\alpha)\}$ has the half-length b_i, because adding e made the boundary contour closed. Moreover, merging other regions along their longest boundary sides can make a longest boundary edge of the final region only shorter.

Then all edges in the boundary contours ∂C have half-lengths at most $\alpha(C) \leq \varepsilon$. Hence ∂C is covered by the disks with the radius ε and centers at all points of C, so $\partial C \subset C^\varepsilon \subset G^{2\varepsilon}$. $\quad\square$

Since resulting contours pass through points of C, the output is locally sensitive to perturbations of C. However, the result below confirms a global stability of boundaries in a small offset.

Corollary 12. *In the conditions of Theorem 11 if another cloud $\tilde{C}$ is δ-close to C, then the boundaries in the 1st segmentation from Definition 6 for the perturbed cloud $\tilde{C}$ are within the $(2\delta + 4\varepsilon)$-offset of the boundaries in the 1st segmentation for C.*

Proof. Since $\tilde{C}$ is δ-close to C, which is ε-close to the graph $G \subset \mathbb{R}^2$, the perturbed cloud $\tilde{C}$ is $(\delta + \varepsilon)$-close to G.

Theorem 11 for the ε-sample C and $(\delta + \varepsilon)$-sample $\tilde{C}$ says that the boundaries ∂C of the 1st segmentation are 2ε-close to G and the boundaries $\partial \tilde{C}$ are $(2\delta + 2\varepsilon)$-close to the graph G.

Hence these boundaries are $(2\delta + 4\varepsilon)$-close to each other. $\quad\square$

6. Conclusions and experiments on synthetic and real data

Clouds C of $n = 1000$ points in Fig. 12–13 were randomly sampled around graphs $G \subset \mathbb{R}^2$ (the regular octagon with 4 big diagonals in Fig. 12, a lattice graph in Fig. 13) as follows:

(1) choose a random seed point p in the graph G so that the distribution of seeds over the total length of G is uniform;

(2) choose a final point q in $[p - \varepsilon, p + \varepsilon] \times [p - \varepsilon, p + \varepsilon]$, which is the square neighborhood of a size ε around the seed p.

The noise bound ε is needed only for generating a cloud C, the algorithm uses only C, not ε. The 2nd picture of Fig. 12 is the 1D persistence diagram $\text{PD}\{C^\alpha\}$ with exactly 8 dots above the yellow widest gap. The initial regions in the 3rd picture of Fig. 12 are in a 1-1 correspondence with all dots from $\text{PD}\{C^\alpha\}$. The final segmentation in the 4th picture has only 8 regions obtained by merging 'noisy' regions of lower persistence.


Fig. 12. Top: cloud C and 1D persistence diagram $\text{PD}\{C^\alpha\}$ **with yellow widest gap. Bottom: initial and final segmentations (best viewed in color).**


Fig. 13. Top: cloud C and 1D persistence diagram $\text{PD}\{C^\alpha\}$ **with yellow widest gap. Bottom: initial and final segmentations (best viewed in color).**

The initial segmentation in the 3rd picture of Fig. 13 contains more than 9 large regions. However only 9 regions have a large enough persistence, all others correctly merge them in the 4th picture. So the size doesn't matter, but the persistence does!

The clouds C in Fig. 14–15 are samples of binary images. C is rather sparse in Fig. 14, but the widest gap in $\text{PD}\{C^\alpha\}$ separates 2 dots corresponding to the final 2 non-convex regions.

The 1st pictures in Fig. 16–19 are from the Berkely segmen-

Fig. 14. Top: cloud C of $n = 159$ **points and diagram** $\mathrm{PD}\{C^\alpha\}$ **with the yellow widest gap. Bottom: initial and final segmentations (best viewed in color).**


Fig. 15. Top: cloud C of $n = 852$ **points and diagram** $\mathrm{PD}\{C^\alpha\}$ **with the yellow widest gap. Bottom: initial and final segmentations (best viewed in color).**

tation database, see Martin et al. (2001). The 2nd pictures are clouds C obtained as Canny edge points with the low threshold 75 and ratio 3. The clouds C in Fig. 16 and 19 have a lot of clutter, but the final segmentation in the 4th pictures correctly identify a shape of the object corresponding to the only dot above the 1st widest gap in the diagrams from the 3rd pictures.

The final 2 pictures in Fig. 17 and 19 show the 2nd segmentations from the hierarchy, which separate two hands from the face in Fig. 17 and the boat from its shadow in Fig. 19.

The summary of the key contributions is below.

• The $O(n \log n)$ time algorithm from section 4 accepts any cloud of points in $\mathbb{R}^2$ without extra parameters and outputs a hierarchy of segmentations selected by their persistence.

• Theorem 11 proves that the persistent contours approximate true contours of a graph G given only by a noisy sample C.

• Corollary 12 guarantees a global stability of boundaries in the 1st segmentation of C under perturbations of a given cloud C.

In comparison with the earlier conference version in Kurlin (2014a), Theorem 11 has been extended from the partial case $\theta(G) = 0$ to a wider class of graphs with any thickness $\theta(G) \geq 0$. The proposed algorithm has the C++ code at http://kurlin.org and outputs only boundaries of persistent holes. The extension to a graph with hanging vertices and branches is in Kurlin (2015a). Gaps in 1D persistence are studied for more general filtrations in Kurlin (2015b). Here are further open problems.

• Smooth boundary contours in a final segmentation of C.

• Extend Theorem 11 to noisy samples with unbounded noise.

• Use persistence to locate holes in high-dimensional clouds.

References

Aanjaneya, M., Chazal, F., Chen, D., Glisse, M., Guibas, L., Morozov, D., 2012. Metric graph reconstruction from noisy data. International J. Computational Geometry and Applications 22, 305–325.

Attali, D., Glisse, M., Hornus, S., Lazarus, F., Morozov, D., 2009. Persistence-sensitive simplification of functions on surfaces in linear time, in: TopoInVis: Topology-based Methods in Visualization, Springer.

de Berg, M., Cheong, O., van Kreveld, M., Overmars, M., 2008. Computational Geometry: Algorithms and Applications. Springer.

Chen, C., Freedman, D., Lampert, C., 2011. Enforcing topological constraints in random field image segmentation, in: Proceedings of CVPR: Computer Vision and Pattern Recognition, pp. 2089–2096.

Chernov, A., Kurlin, V., 2013. Reconstructing persistent graph structures from noisy images. Image-A 3, 19–22.

Cohen-Steiner, D., Edelsbrunner, H., Harer, J., 2007. Stability of persistence diagrams. Discrete and Computational Geometry 37, 103–120.

Delfinado, C., Edelsbrunner, H., 1995. An incremental algorithm for betti numbers of simplicial complexes on the 3-sphere. Computer Aided Geometric Design 12, 771–784.

Edelsbrunner, H., 1995. The union of balls and its dual shape. Discrete and Computational Geometry 13, 415–440.

Edelsbrunner, H., Harer, J., 2010. Computational topology. An introduction. AMS, Providence.

Kurlin, V., 2014a. Auto-completion of contours in sketches, maps and sparse 2d images based on topological persistence, in: Proceedings of CTIC: Computational Topology in Image Context, pp. 594–601.

Kurlin, V., 2014b. A fast and robust algorithm to count topologically persistent holes in noisy clouds, in: Proceedings of CVPR, pp. 1458–1463.

Kurlin, V., 2015a. A homologically persistent skeleton is a fast and robust descriptor of interest points in 2d images, in: LNCS, pp. 606 – 617.

Kurlin, V., 2015b. A one-dimensional homologically persistent skeleton of a point cloud in any metric space. Computer Graphics Forum 34, 253–262.

Letscher, D., Fritts, J., 2007. Image segmentation using topological persistence, in: Proceedings of CAIP: Comp. Anal. Images and Patterns, pp. 587–595.

Martin, D., Fowlkes, C., Tal, D., Malik, J., 2001. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics, in: Proc. ICCV, pp. 416–423.

Saund, E., 2003. Finding perceptually closed paths in sketches and drawings. Transactions Pattern Analysis and Machine Intelligence 25, 475–490.

**Fig. 16. Top: image 86016 from BSD500 and cloud C of 19627 Canny edge points. Bottom: PD$\{C^\alpha\}$ and final segmentation with a correct object.**


Fig. 18. Top: image 3096 from BSD500 and cloud C of 1614 Canny edge points. Bottom: PD$\{C^\alpha\}$ and final segmentations with a correct object.
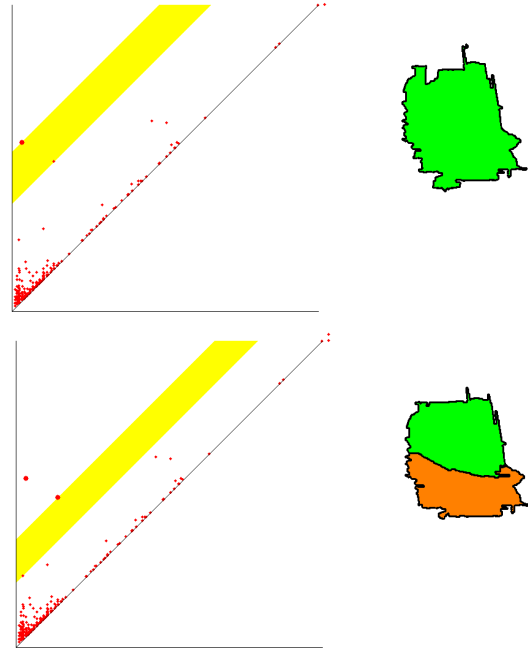

Fig. 17. Top: image 302003 from BSD500 and cloud C of 5388 Canny edge points. Middle: PD$\{C^\alpha\}$ with 1st widest gap and segmentation with 1 region Bottom: PD$\{C^\alpha\}$ with 2nd widest gap and segmentation with 4 regions.

Fig. 19. Top: image 15088 from BSD500 and cloud C of 6565 Canny edge points. Middle: PD$\{C^\alpha\}$ with 1st widest gap and segmentation with 1 region Bottom: PD$\{C^\alpha\}$ with 2nd widest gap and segmentation with 2 regions.