

A fast and robust algorithm to count topologically persistent holes in noisy clouds

Vitaliy Kurlin

Durham University

Department of Mathematical Sciences, Durham, DH1 3LE, United Kingdom

vitaliy.kurlin@gmail.com, <http://kurlin.org>

Abstract

Preprocessing a 2D image often produces a noisy cloud of interest points. We study the problem of counting holes in noisy clouds in the plane. The holes in a given cloud are quantified by the topological persistence of their boundary contours when the cloud is analyzed at all possible scales.

We design the algorithm to count holes that are most persistent in the filtration of offsets (neighborhoods) around given points. The input is a cloud of n points in the plane without any user-defined parameters. The algorithm has $O(n \log n)$ time and $O(n)$ space. The output is the array (number of holes, relative persistence in the filtration).

We prove theoretical guarantees when the algorithm finds the correct number of holes (components in the complement) of an unknown shape approximated by a cloud.

1. Introduction: counting holes in noisy clouds

We apply methods from the new area of topological data analysis to counting persistent holes in a noisy cloud of points. Such a cloud can be obtained by selecting interest points in a gray scale or RGB image. Our region-based method uses global topological properties of contours.

By a *shape* we mean any subset $X \subset \mathbb{R}^2$ that can be split into finitely many (topological) triangles. Hence X is bounded, but may not be connected. Then a *hole* in a shape $X \subset \mathbb{R}^2$ is a bounded connected component of the complement $\mathbb{R}^2 - X$. Such a hole can be a disk, a ring or may have a more complicated topological form, see Fig. 1.

The α -offset X^α is the union $\cup_{p \in X} B(p; \alpha)$ of disks with the radius $\alpha \geq 0$ and centers at all $p \in X$. For instance, X^0 is the original shape $X \subset \mathbb{R}^2$. When α is increasing, the holes of $\mathbb{R}^2 - X^\alpha$ are shrinking, may split into smaller newborn holes and will eventually die, each at its own death time α , see Fig. 3. The *persistence* of a hole is its life span death – birth in the filtration $\{X^\alpha\}$ of all α -offsets. So we quantify holes by their persistence at different scales α .

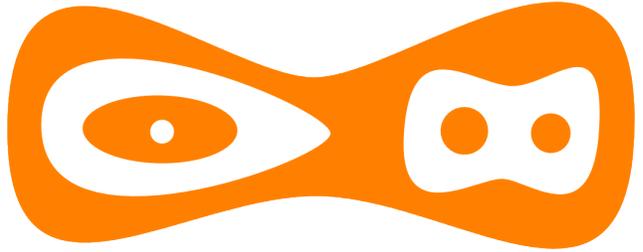


Figure 1. The orange shape $X \subset \mathbb{R}^2$ with 3 white holes of different forms: a small disk, a ring-like hole, a ‘figure-eight’ hole.

Hole counting problem. Let a shape X be represented by a finite sample C of points in \mathbb{R}^2 . Find conditions on X and its sample when one can quickly count persistent holes.

We solve the problem by the algorithm HOCTOP: Hole Counting based on Topological Persistence. The only input is a finite cloud C of n points approximating an unknown shape $X \subset \mathbb{R}^2$. The algorithm outputs the relative persistence of k holes in the filtration $\{C^\alpha\}$ for all $k \geq 0$. If the scale α is random and uniform, this output gives probabilities $P(k \text{ holes})$. The boundary edges of persistent holes can be quickly post-processed to extract all boundary contours.

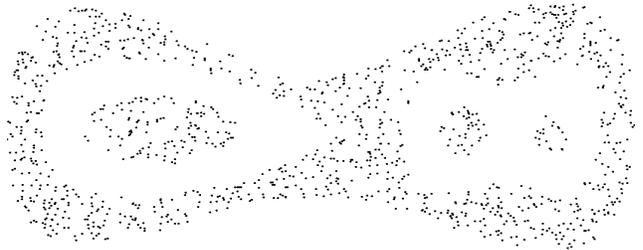


Figure 2. Input: cloud C of 1251 points uniformly sampled from the shape with 3 holes in Fig. 1. Output probabilities of HOCTOP: $P(3 \text{ holes}) \approx 24\%$, $P(2 \text{ holes}) \approx 13\%$, $P(8 \text{ holes}) \approx 11\%$.

Theorems 1, 4 say that the algorithm HOCTOP quickly and correctly finds all persistent holes using only a good enough sample C of an unknown shape X , see section 2.

2. Main results: the algorithm and guarantees

We start from a high-level description of our algorithm.

The topological persistence of contours in the filtration $\{C^\alpha\}$ is computed by using a Delaunay triangulation $\text{Del}(C)$ of a given cloud $C \subset \mathbb{R}^2$ of n points. By Nerve Lemma 8 the α -offsets C^α can be continuously deformed to the α -complexes $C(\alpha)$, which filter $\text{Del}(C)$ as follows: $C = C(0) \subset \dots \subset C(\alpha) \subset \dots \subset C(+\infty) = \text{Del}(C)$. Each $C(\alpha)$ has some edges and triangles from $\text{Del}(C)$.

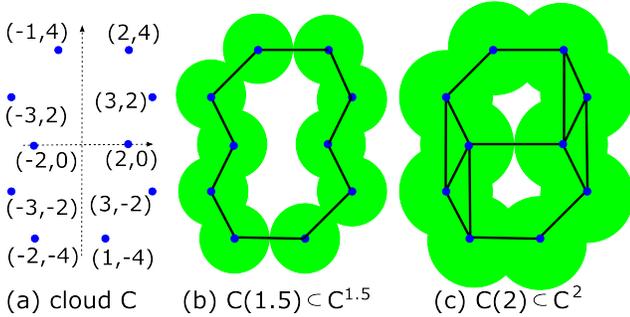


Figure 3. The big hole in the green offset C^α is born at $\alpha = 1.5$, splits into 2 smaller holes at $\alpha = 2$ and dies at $\alpha \approx 2.577$, so the topological persistence of this hole is death – birth ≈ 1.077 .

The graph dual to $\text{Del}(C)$ is filtered by the subgraphs $C^*(\alpha)$ whose connected components correspond to holes in $C(\alpha)$. When α is decreasing, $C(\alpha)$ is shrinking, so its holes are growing and corresponding components of $C^*(\alpha)$ merge at critical values of α , see Fig. 6. The persistence of cycles in the filtration $\{C^\alpha\}$ corresponds to the persistence of components in $\{C^*(\alpha)\}$, see Duality Lemma 14.

The pairs (birth, death) of connected components in $\{C^*(\alpha)\}$ are found via a union-find structure by adding edges and merging components. So computing the 1-dimensional persistence of cycles in $\{C^\alpha\}$ reduces to the 0-dimensional persistence of components in $\{C^*(\alpha)\}$.

Starting from a given cloud $C \subset \mathbb{R}^2$ of n points with real coordinates (x_i, y_i) , $i = 1, \dots, n$, we find a Delaunay triangulation $\text{Del}(C)$ in $O(n \log n)$ time with $O(n)$ space. Then we remove each edge of $\text{Del}(C)$ one by one in the decreasing order of their length. Removing an edge may break a contour when adjacent regions in $C(\alpha)$ and the corresponding components of $C^*(\alpha)$ merge. In the case of a merger, a younger component of $C^*(\alpha)$ and the corresponding hole in $C(\alpha)$ die. We note the birth and death of each dead hole. We get the probability of k holes as the relative length of all intervals of the scale α when $C^\alpha \subset \mathbb{R}^2$ has k holes.

Theorem 1. *The algorithm HOCTOP counts all holes in a given cloud $C \subset \mathbb{R}^2$ of n points in $O(n \log n)$ time with $O(n)$ space. All holes are ordered by their topological persistence in the ascending filtration $\{C^\alpha\}$ of the α -offsets.*

Definition 2 (ε -sample). *A cloud C is an ε -sample of a shape $X \subset \mathbb{R}^2$ if $X \subset C^\varepsilon$ and $C \subset X^\varepsilon$. So any point of C is within the distance ε from a point of X and any point of X is at most ε away from a point of C . Hence ε can be considered as the upper bound of some arbitrary noise.*

Definition 3 (min and max homological feature sizes). *For any shape $X \subset \mathbb{R}^2$, let $\alpha = \text{minhfs}(X)$ be the minimum homological feature size when a first hole is born or dies in X^α . Let $\alpha = \text{maxhfs}(X)$ be the maximum homological feature size after which no holes are born or die in X^α .*

Theorem 4 gives sufficient (not necessary) conditions when the algorithm finds the correct number of holes in an unknown shape $X \subset \mathbb{R}^2$ that is represented by its finite sample C . We extend the Homology Inference Theorem [4] to the case when the upper bound ε of noise is unknown.

Theorem 4. *Let a cloud C be an ε -sample of a shape $X \subset \mathbb{R}^2$ with an unknown parameter ε such that $\text{minhfs}(X) > \frac{1}{2}\text{maxhfs}(X) + 2\varepsilon$. Then the algorithm HOCTOP finds the correct number of holes in X by using only the cloud C .*

The condition $\text{minhfs}(X) > \frac{1}{2}\text{maxhfs}(X) + 2\varepsilon$ means that all holes of X , which are bounded components of $\mathbb{R}^2 - X$, have comparable sizes (neither tiny nor huge).

Even if the conditions of Theorem 4 are not satisfied, we can always find the number k of holes with the highest probability. The algorithm HOCTOP can also accept a signal-to-noise ratio τ and output all holes whose persistence is larger than τ . Alternatively, the user may prefer to get most likely outputs ordered by the probability $P(k \text{ holes})$.

3. Previous work on computing persistence

The offsets C^α of a finite cloud C are usually studied through the Čech or Rips complexes, which may contain up to $O(n^k)$ simplices in all dimensions $k \leq n - 1$ even if $C \subset \mathbb{R}^2$. A Delaunay triangulation has the advantage of a smaller size up to $m = O(n^2)$ in dimensions $n = 2, 3, 4$.

The fastest algorithm [8] for computing persistence of a filtration in all dimensions has the same running time $O(m^{2.376})$ in the number m of simplices as the best known time for the multiplication of two $m \times m$ matrices.

In dimension 0 the persistence can be computed in almost linear time [6, p. 6–8], which was used for simplifying functions on surfaces [1] and for approximating persistence of an unknown scalar field from its values on a sample [3].

Two extra parameters were used in a Delaunay-based image segmentation [7]: α for the radius of disks centered at points of a cloud C and p for a desired level of persistence.

4. Delaunay triangulation and α -complexes

Definition 5 (simplicial complex). *A simplicial 2-complex is a finite set of simplices (vertices, edges, triangles):*

- the sides of any triangle are included in the complex;
- the endpoints of any edge are included in the complex;
- two triangles can intersect only along a common edge;
- edges can meet only at a common endpoint (a vertex);
- an edge can not pierce through the interior of a triangle.

If a complex S is drawn in \mathbb{R}^n without self-intersections, we may call this image $|S|$ a *geometric realization* of S . We have defined a shape $X \subset \mathbb{R}^2$ as a geometric realization of a 2-complex. For instance, a round disk whose boundary is split into 3 edges by 3 vertices is a topological triangle.

A *cycle* in a complex is a sequence of edges e_1, \dots, e_m such that any consecutive edges e_i, e_{i+1} (in the cyclic order) have a common vertex. Any loop in a geometric realization $|S|$ continuously deforms to a cycle of edges in S .

Definition 6 (Delaunay triangulation Del). *For a point p_i in a cloud $C = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$, the Voronoi cell $V(p_i) = \{q \in \mathbb{R}^2 : d(p_i, q) \leq d(p_j, q) \forall j \neq i\}$ is the set of all points q that are (non-strictly) closer to p_i than to other points of C . The Delaunay triangulation $\text{Del}(C)$ is the nerve of the Voronoi diagram $\cup_{p \in C} V(p)$. Namely, $p, q, r \in C$ span a triangle if and only if $V(p) \cap V(q) \cap V(r) \neq \emptyset$.*

By another definition [2, section 9.1] the circumcircle of any *Delaunay triangle* in $\text{Del}(C)$ encloses no points of C .

For a cloud $C \subset \mathbb{R}^2$ of n points, let $\text{Del}(C)$ have k triangles and b boundary edges in the external region. Counting all E edges over triangles, we get $3k + b = 2E$. Euler's formula $n - E + (k + 1) = 2$ implies that $k = 2n - b - 2$, $E = 3n - b - 3$. So $\text{Del}(C)$ has $O(n)$ edges and triangles.

Definition 7 (α -complex $C(\alpha)$). *For a scale parameter $\alpha > 0$, the α -complex $C(\alpha)$ is the nerve of $\cup_{p \in C} (V(p) \cap B(p; \alpha))$, see [6, section III.4]. Points $p, q \in C$ are connected by an edge if $V(p) \cap B(p; \alpha)$ meets $V(q) \cap B(q; \alpha)$. Three points $p, q, r \in C$ span a triangle if the intersection $V(p) \cap B(p; \alpha) \cap V(q) \cap B(q; \alpha) \cap V(r) \cap B(r; \alpha) \neq \emptyset$.*

If $\alpha > 0$ is very small, all points of C are disjoint in $C(\alpha)$, while $C(\alpha) = \text{Del}(C)$ for any large enough α , see examples in Fig. 3. So all α -complexes form the filtration $C = C(0) \subset \dots \subset C(\alpha) \subset \dots \subset C(+\infty) = \text{Del}(C)$. Edges or triangles are added only at *critical values* of α .

Lemma 8 (Nerve of a ball covering [5]). *The union of balls $C^\alpha = \cup_{p \in C} B(p; \alpha)$ continuously deforms to (has the homotopy type of) a geometric realization of $C(\alpha)$.*

5. Persistent homology: definitions, examples

Definition 9 (1-dimensional homology H_1). *We consider the 1-dimensional homology group $H_1(S)$ only with coefficients in $\mathbb{Z}/2\mathbb{Z} = \{0, 1\}$. Cycles of a 2-dimensional complex S can be algebraically written as linear combinations*

of edges (with coefficients 0 or 1) and generate the vector space C_1 of cycles. The boundaries of all triangles in S (as cycles of 3 edges) generate the subspace $B_1 \subset C_1$. The quotient group C_1/B_1 is the homology group $H_1(S)$.

By a *filtration* $\{S(\alpha)\}$ we mean a sequence of nested complexes $S(0) \subset \dots \subset S(\alpha) \subset \dots$ that change only at finitely many *critical values* $\alpha_1, \dots, \alpha_m$. Then we get the induced linear maps $H_1(S(\alpha_1)) \rightarrow \dots \rightarrow H_1(S(\alpha_m))$.

Definition 10 (persistence diagram $\text{PD}\{S(\alpha)\}$). *In a filtration $\{S(\alpha)\}$ a homology class $\gamma \in H_1(S(\alpha_i))$ is born at $\alpha_i = \text{birth}(\gamma)$ if γ is not in the image of $H_1(S(\alpha)) \rightarrow H_1(S(\alpha_i))$ for any $\alpha < \alpha_i$. The class γ dies at the first time $\alpha_j = \text{death}(\gamma) \geq \alpha_i$ when the image of γ under $H_1(S(\alpha_i)) \rightarrow H_1(S(\alpha_j))$ merges into the image of $H_1(S(\alpha)) \rightarrow H_1(S(\alpha_j))$ for some $\alpha < \alpha_i$. The class γ has the persistence $\text{death}(\gamma) - \text{birth}(\gamma)$. The point (α_i, α_j) has the multiplicity μ_{ij} equal to the number of independent classes that are born at α_i and die at α_j . The persistence diagram $\text{PD}\{S(\alpha)\}$ in $\{(x, y) \in \mathbb{R}^2 : x \leq y\}$ is the multi-set consisting of all points (α_i, α_j) with the multiplicity μ_{ij} and all diagonal points (x, x) with the infinite multiplicity.*

Pairs with a low persistence $\text{death} - \text{birth}$ (close to the diagonal $\{x = y\}$ in PD) are treated as noise. Pairs with a high persistence represent persistent cycles in $\{S(\alpha)\}$.

We shall consider the filtrations of α -offsets $\{X^\alpha\}$ and $\{C^\alpha\}$ for a shape $X \subset \mathbb{R}^2$ and a finite cloud $C \subset \mathbb{R}^2$. Figures 4 and 5 show the persistence diagram PD for the filtration of the α -offsets C^α equivalent to $C(\alpha)$ by Lemma 8.

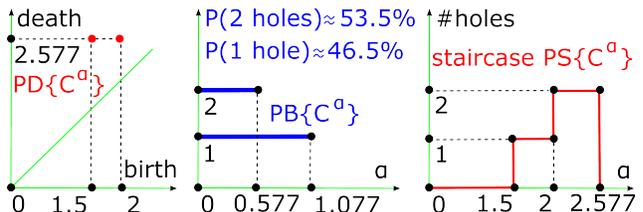


Figure 4. Extra outputs for the cloud C of 10 points in Fig. 3. Left: persistence diagram, middle: barcode, right: persistence staircase.

We can convert the persistence diagram into the persistence barcode $\text{PB}\{C^\alpha\}$. All pairs (birth, death) give horizontal bars ordered by their length $\text{death} - \text{birth}$. Usually the bars are drawn from the left endpoint 0 to the right endpoint $\text{death} - \text{birth}$, see the middle picture in Fig. 4.

We suggest one more way to visualize persistence. Each pair (birth, death) defines the function $f(\alpha) = 1$ for $\text{birth} \leq \alpha < \text{death}$ and $f(\alpha) = 0$ otherwise. The sum of these functions over all pairs gives the *persistence staircase* $\text{PS}\{C^\alpha\}$. The value of this piecewise constant function of α is the number of holes in the offset C^α . We have connected consecutive horizontal segments of $\text{PS}\{C^\alpha\}$ to get a ‘continuous’ staircase as in the right picture of Fig. 4.

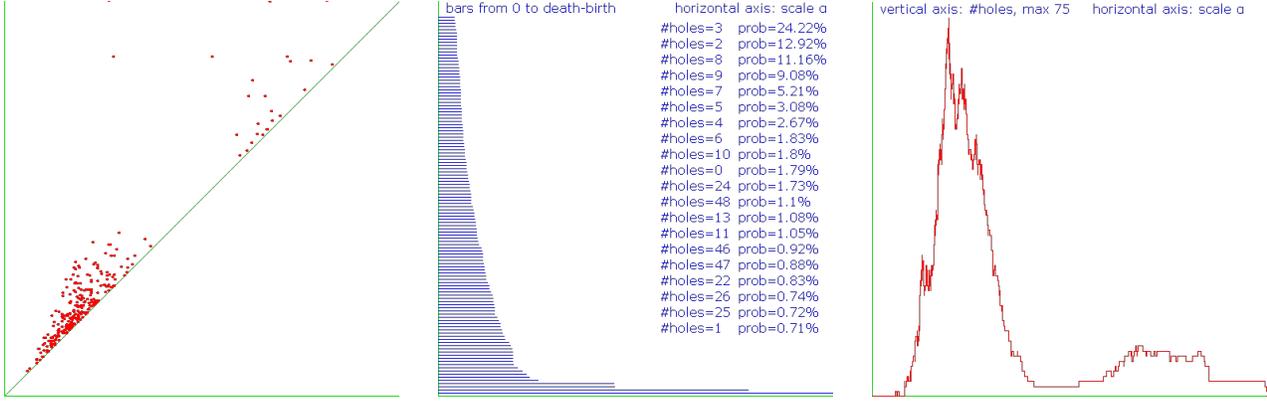


Figure 5. Extra outputs for the cloud C of 1251 points in Fig. 2. Left: persistence diagram PD, middle: barcode PB, right: staircase PS.

For the cloud C of 10 points in Fig. 3, the full range of the scale α is from the smallest critical value $\alpha = 1.5$ (when a first hole is born) to the largest critical value $\alpha = \frac{5}{8}\sqrt{17} \approx 2.577$ (when both final holes die). The output probability $P(1 \text{ hole}) \approx 46.5\%$ is the contribution of the interval $(1.5, 2)$ to the full range $1.5 \leq \alpha \leq \frac{5}{8}\sqrt{17}$. The largest probability $P(2 \text{ holes}) \approx 53.5\%$ is the contribution of the interval $(2, \frac{5}{8}\sqrt{17})$ when C^α has exactly 2 holes.

For the cloud C of 1251 points in Fig. 2, we scaled $\text{PB}\{C^\alpha\}$ and $\text{PS}\{C^\alpha\}$ along the horizontal α -axis and kept only the longest bars in the barcode $\text{PB}\{C^\alpha\}$ in Fig. 5.

6. Persistent homology: stability and duality

Definition 11 (bottleneck distance d_B). Let the distance between $p = (x_1, y_1)$, $q = (x_2, y_2)$ in \mathbb{R}^2 be $\|p - q\|_\infty = \max\{|x_1 - x_2|, |y_1 - y_2|\}$. The bottleneck distance is $d_B(D, D') = \inf_\varphi \sup_{p \in D} \|p - \varphi(p)\|_\infty$ over all bijections $\varphi : D \rightarrow D'$ between persistence diagrams D, D' .

We quote the simple version of the Persistence Stability Theorem [4] only for persistence diagrams of α -offsets.

Theorem 12. [4] If a finite cloud C of points is an ε -sample of a shape $X \subset \mathbb{R}^2$, then $d_B(\text{PD}\{X^\alpha\}, \text{PD}\{C^\alpha\}) \leq \varepsilon$.

Stability Theorem 12 implies for barcodes PB that the endpoints of all bars are perturbed by at most ε . So a long bar can become only a bit shorter after adding noise.

To every triangle in the Delaunay triangulation $\text{Del}(C)$, let us associate a single abstract vertex v_i , $i = 1, \dots, k$. It will be convenient to call the external region of $\text{Del}(C)$ also a ‘triangle’ and represent it by an extra vertex v_0 .

Definition 13 (graphs $C^*(\alpha)$). For any vertices v_i, v_j representing adjacent triangles in $\text{Del}(C)$, let d_{ij} be the length of the (longest) common edge of the triangles. The metric graph C^* dual to $\text{Del}(C)$ has the vertices v_0, v_1, \dots, v_k and

edges of the length d_{ij} connecting vertices v_i, v_j that represent adjacent triangles, see Fig. 6. The graph C^* is filtered by the subgraphs $C^*(\alpha)$ that have only the edges of a length $d_{ij} > 2\alpha$. We remove any isolated node v (except v_0) from $C^*(\alpha)$ if the corresponding triangle T_v is not acute or has a small circumradius $\text{rad}(v) \leq \alpha$. We get the filtration $C^* = C^*(0) \supset \dots \supset C^*(\alpha) \supset \dots \supset C^*(+\infty) = \{v_0\}$.

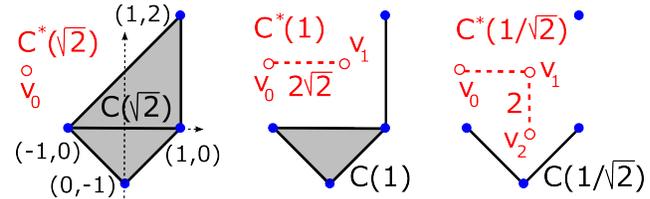


Figure 6. The complexes $C(\alpha)$ have solid edges and gray triangles. The graphs $C^*(\alpha)$ have circled vertices and red dashed edges.

Components of $C^*(\alpha)$ are called *white*, because they represent regions in $\mathbb{R}^2 - C(\alpha)$ (or holes in $\mathbb{R}^2 - C^\alpha$). A cycle $\gamma \subset C(\alpha)$ is called a *contour* if γ bounds a region in $\mathbb{R}^2 - C(\alpha)$, so γ ‘encloses’ the corresponding white component of $C^*(\alpha)$. Lemma 14 is an analogue of the Symmetry Theorem [6, p. 164] for a function on a closed manifold.

Lemma 14 (Duality). All contours of the complex $C(\alpha)$ are in a 1-1 correspondence with all connected components of the graph $C^*(\alpha)$ not containing the vertex v_0 . When α is decreasing, the contours of $C(\alpha)$ and the white components of $C^*(\alpha)$ have the corresponding critical moments:

- a birth of a contour \leftrightarrow a birth of a white component,
- a death of a contour \leftrightarrow a death of a white component. \square

7. The algorithm HOCTOP for counting holes

We build the union-find structure $\text{Forest}(\alpha)$ on the vertices of the graph $C^*(\alpha)$. All nodes and trees of $\text{Forest}(\alpha)$

will be in a 1-1 correspondence with all vertices and white components of $C^*(\alpha)$. Every node v in $\text{Forest}(\alpha)$ has

- a pointer to a unique parent of the node v in $\text{Forest}(\alpha)$;
- a pointer to the Delaunay triangle dual to the node v ;
- the weight (the number of nodes below v in its tree);
- the critical value (birth) $\alpha_v = \sup\{\alpha : v \in C^*(\alpha)\}$.

If a node v is a self-parent, we call v a *root*. We can find $\text{root}(v)$ of any node v by going up along parent links. If α is decreasing, α_v can be considered as the birth time when the vertex v joins $C^*(\alpha)$. The algorithm initializes $\text{Forest}(\alpha)$ as the set of isolated nodes v_0, \dots, v_k . If the triangle corresponding to v_k is acute, the birth time of v_k is the circumradius of the triangle, otherwise 0. We will go through all edges of $\text{Del}(C)$ in the decreasing order of their length and will update α_v when v enters the ascending filtration $\{v_0\} = C^*(+\infty) \subset \dots \subset C^*(\alpha) \subset \dots \subset C^*(0) = C^*$.

All triangles of $C(\alpha)$ and the corresponding nodes of $\text{Forest}(\alpha)$ are called *gray*. The remaining triangles and the external region of $\text{Del}(C)$ are called *white*. The external region has birth time $+\infty$ and is called a ‘triangle’ for simplicity. Initially all triangles with birth time 0 are gray.

The while loop. For each edge $e \subset \text{Del}(C)$ arriving in the decreasing order of length, we find two triangles T_u, T_v attached to e and check if they are gray or white. To determine if a triangle T_v represented by a node v is gray, we go up along parent links from v to $\text{root}(v)$. If the birth time of $\text{root}(v)$ is 0, the triangle T_v is still gray, otherwise white.

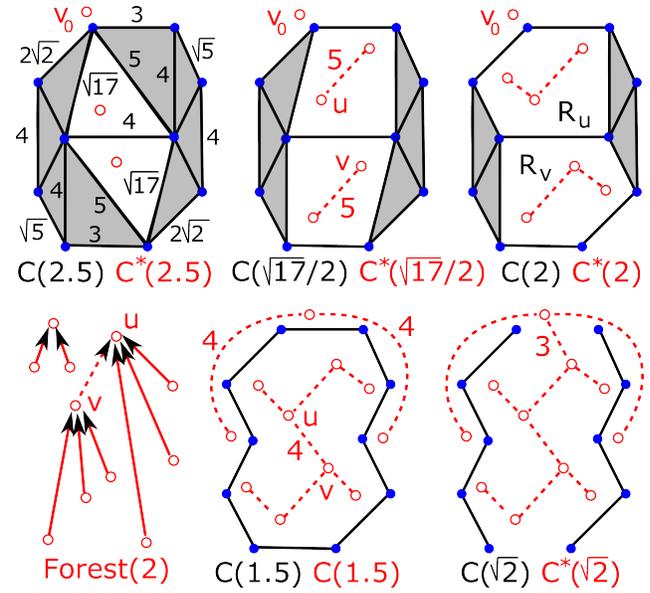
To distinguish Cases 1 and 4 below, we also check if the triangles T_u, T_v attached to the current edge e are in the same region of $\mathbb{R}^2 - C(\alpha)$. Case 1 means that the nodes $u, v \in \text{Forest}(\alpha)$ belong to the same tree, so $\text{root}(u) = \text{root}(v)$. In all 4 cases the scale α goes down through the half-length $\frac{1}{2}\text{length}(e)$ of the current edge e from $\text{Del}(C)$.

Case 1: e has the same white region on both sides of e . $C(\alpha)$ loses only the open edge e . The white components of $C^*(\alpha)$ are unchanged. Fig. 6 illustrates Case 1 for $\alpha = 1$ when $C(\alpha)$ loses the edge connecting $(1, 0)$ to $(1, 2)$.

Case 2: the edge e is in 1 gray triangle and 1 white triangle. Let $u, v \in C^*(\alpha)$ be the vertices dual to the gray triangle T_u and the white triangle T_v attached to the current edge e in $\text{Del}(C)$. Then the birth times are $\alpha_u = 0, \alpha_{\text{root}(v)} > 0$.

Since α is decreasing, the descending filtration $C(\alpha)$ loses the (open) edge e and the gray (open) triangle T_u . So the vertex u becomes connected by an edge with v and joins the white component of $C^*(\alpha)$ containing v . Then we link the isolated node u to the tree containing the older node v in $\text{Forest}(\alpha)$. So $\text{root}(v)$ becomes the parent of u and the weight of $\text{root}(v)$ jumps by 1. Fig. 7 illustrates Case 2 for $\alpha = \frac{\sqrt{17}}{2}$ when $C(\alpha)$ loses the 2 edges of length $\sqrt{17}$.

Case 3: the edge e is in the boundary of 2 gray triangles.



sistence is usually defined when the scale α is increasing. However, we need the ascending filtration $\{C^*(\alpha)\}$ to use a union-find structure, so α is decreasing in the algorithm.

Finally, to merge the trees with $\text{root}(u), \text{root}(v)$ in $\text{Forest}(\alpha)$, we compare the weights of the roots and set the root of the (non-strictly) larger tree as the parent for the root of another tree. So the size of any subtree grows by a factor of at least 2 each time when we pass to the parent. We get

Lemma 15. *By the above construction the longest path in any tree of size k from $\text{Forest}(\alpha)$ has length $O(\log k)$.* \square

8. Proofs of main results and our conclusion

Proof of Theorem 1. Constructing the Delaunay triangulation $\text{Del}(C)$ on a cloud of n points requires $O(n \log n)$ time [2, Chapter 9]. Sorting $O(n)$ edges of $\text{Del}(C)$ needs $O(n \log n)$ time. Then we go through the *while loop* analyzing each of the $O(n)$ edges of $\text{Del}(C)$. For the nodes $u, v \in \text{Forest}(\alpha)$ of triangles attached to each edge e , we find the roots of u, v by going up along $O(\log n)$ parent links by Lemma 15. All other steps in the *while loop* require only $O(1)$ time. Hence the total time is $O(n \log n)$. The sizes of all data structures are proportional to the numbers of edges or triangles in $\text{Del}(C)$, so we use $O(n)$ space. \square

The careful analysis of a union-find structure [9] says that $\text{Forest}(\alpha)$ can be built in time $O(nA^{-1}(n, n))$ time, where $A^{-1}(n, n)$ is the extremely slowly growing inverse Ackermann function. Our time $O(n \log n)$ is dominated by the construction of $\text{Del}(C)$ and sorting all $O(n)$ edges.

Proof of Theorem 4. The important critical values of α for the 1-dimensional homology of the filtration $\{X^\alpha\}$ are

- $\alpha = 0$ when the homology $H_1(X^0) = H_1(X)$ is correct;
- $\alpha = \text{minhfs}(X)$ is the 1st value when $H_1(X^\alpha)$ changes;
- $\alpha = \text{maxhfs}(X)$ is the last value when $H_1(X^\alpha)$ changes.

Then $H_1(X^\alpha) \cong H_1(X)$ for $0 \leq \alpha < \text{minhfs}(X)$. If a cloud C is an ε -sample of a shape $X \subset \mathbb{R}^2$, then due to Stability Theorem 12 [4] we have $H_1(C^\alpha) \cong H_1(X)$ over (possibly) shorter interval $\varepsilon \leq \alpha \leq \text{minhfs}(X) - \varepsilon$. Then all changes of $H_1(C^\alpha)$ will stop at $\alpha = \text{maxhfs}(X) + \varepsilon$.

The condition $\text{minhfs}(X) > \frac{1}{2}\text{maxhfs}(X) + 2\varepsilon$ guarantees that the persistence interval $[\varepsilon, \text{minhfs}(X) - \varepsilon]$ is the longest over the full range $[0, \text{maxhfs}(X) + \varepsilon]$, because $\varepsilon < \text{minhfs}(X) - 2\varepsilon > \text{maxhfs}(X) + \varepsilon - (\text{minhfs}(X) - \varepsilon)$. Without using ε , we can find this interval of α , where $C(\alpha)$ has the same number of holes as the unknown shape X . \square

Conclusion. Here are the key advantages of our approach:

- a cloud $C \subset \mathbb{R}^2$ of n points is simultaneously analyzed at all scales α without any extra user-defined parameters;
- the algorithm HOCTOP counts persistent holes of any topological form in $O(n \log n)$ time, see Theorem 1;

- theoretical guarantees for a correct number of holes are proved for ε -samples of unknown shapes, see Theorem 4;
- the output is stable under perturbations of a cloud C and the only parameter of noise is an unknown upper bound ε .

Fig. 8 shows extracted contours (with our uniform noise) of images at <http://www.lems.brown.edu/~dmc>.

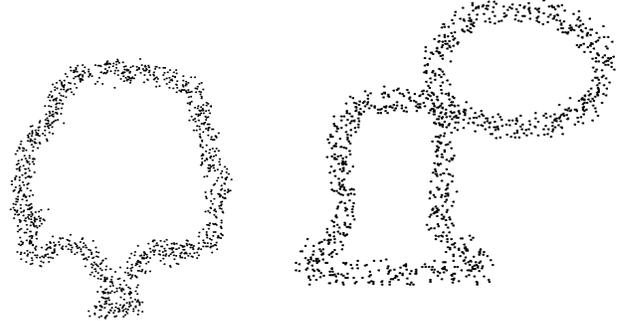


Figure 8. Output of HOCTOP for real noisy contours. Left: $P(1 \text{ hole}) \approx 90.5\%$, $P(2 \text{ holes}) \approx 3\%$, $P(4 \text{ holes}) \approx 0.6\%$. Right: $P(2 \text{ holes}) \approx 74.2\%$, $P(1 \text{ hole}) \approx 13\%$, $P(3) \approx 1.3\%$.

More details and experimental results are in the extended version on author's website <http://kurlin.org>. We thank the anonymous reviewers for their helpful suggestions. We are open to collaboration on any related projects.

References

- [1] D. Attali, M. Glisse, S. Hornus, F. Lazarus, and D. Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. *TopoInVis 2009: Topological Methods in Data Analysis and Visualization*. 2
- [2] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 2008. 3, 6, 7
- [3] F. Chazal, L. Guibas, S. Oudot, P. Skraba. Scalar Field Analysis over Point Cloud Data. *Discrete and Computational Geometry*, v. 46 (2011), p.743-775. 2
- [4] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, 37:103–130, 2007. 2, 4, 6
- [5] H. Edelsbrunner. The union of balls and its dual shape. *Discrete Computational Geometry*, 13:415–440, 1995. 3
- [6] H. Edelsbrunner and J. Harer. *Computational topology. An introduction*. AMS, Providence, 2010. 2, 3, 4, 5, 7
- [7] Letscher, D., Fritts, J. Image segmentation using topological persistence. *Proceedings of CAIP 2007: Computer Analysis of Images and Patterns*, pages 587–595. 2
- [8] N. Milosavljevic, D. Morozov, and P. Skraba. Zigzag persistent homology in matrix multiplication time. *Proceedings of SoCG 2011*, pages 216–225, ACM. 2
- [9] R. Tarjan. *Data structures and network algorithms*. SIAM, Philadelphia, 1983. 6

Appendix A: a pseudo-code of HOCTOP

Algorithm 1 below contains the pseudo-code of the our main algorithm HOCTOP. Cases 2–4 from the description in section 7 are covered in further Algorithms 2–4.

Algorithm 1 Find (birth, death) of all cycles in $C(\alpha)$

Require: a cloud C given as pairs $(x_1, y_1), \dots, (x_n, y_n)$

- 1: Build Delaunay triangulation $\text{Del}(C)$ with k triangles
- 2: Extract all edges with pointers to 2 adjacent triangles
- 3: Sort edges of $\text{Del}(C)$ in the decreasing order of length
- 4: Forest \leftarrow isolated nodes v_0, \dots, v_k with birth times 0
- 5: For the external node v_0 , update the birth $\alpha \leftarrow +\infty$
- 6: For each acute triangle T_v , $\alpha_v \leftarrow$ circumradius of T_v
- 7: Set the total number of links in Forest(α): $L \leftarrow 0$
- 8: **while** $L < k$ (we stop when Forest(α) is a tree) **do**
- 9: Take the next longest edge e from $\text{Del}(C)$
- 10: Set the current critical value: $\alpha \leftarrow \frac{1}{2}\text{length}(e)$
- 11: Find 2 nodes u, v dual to the triangles attached to e
- 12: Find the roots $\text{root}(u), \text{root}(v)$ of the nodes u, v
- 13: **if** $\text{root}(u) = \text{root}(v)$ (u, v in the same region) **then**
- 14: Case 1 (no changes): continue the *while loop*
- 15: **else if** $\alpha_{\text{root}(u)} = 0$ **and** $\alpha_{\text{root}(v)} > 0$ **then**
- 16: Case 2 (u gray, v white): run Algorithm 2
- 17: **else if** $\alpha_{\text{root}(u)} = 0$ **and** $\alpha_{\text{root}(v)} = 0$ **then**
- 18: Case 3 (both u, v are gray): run Algorithm 3
- 19: **else**
- 20: Case 4 ($\alpha_{\text{root}(u)}, \alpha_{\text{root}(v)} > 0$): run Algorithm 4
- 21: **end if**
- 22: $L \leftarrow L + 1$ (one link was added in Cases 2, 3, 4)
- 23: **end while**
- 24: **return** array of pairs (birth, death) from Case 4

Recall that a node u is gray if the birth time $\alpha_u = 0$. The case (u white, v gray) is symmetric to Case 2 below, so we simply denote the gray node by u when calling Algorithm 2.

Algorithm 2 Link 2 nodes u, v in Forest(α) in Case 2

Require: nodes u and $\text{root}(v)$ (so u is gray, v is white)

- 1: Set $\text{root}(v)$ as parent of u , set $\alpha_u \leftarrow \alpha_{\text{root}(v)}$
- 2: Add 1 (coming from u) to the weight of $\text{root}(v)$

In Algorithm 3 below any of the gray nodes u, v can be the parent of the other node, we have simply chosen u .

Algorithm 3 Link 2 nodes u, v in Forest(α) in Case 3

Require: α , nodes u, v dual to triangles (both u, v gray)

- 1: Set u as the parent of the node v in Forest(α)
- 2: Set: $\alpha_u, \alpha_v \leftarrow \alpha$, $\text{weight}(u) \leftarrow 1$, $\text{weight}(v) \leftarrow 0$

Algorithm 4 Update Forest and (birth, death) in Case 4

Require: α , roots $\text{root}(u), \text{root}(v)$ of white nodes u, v

- 1: **if** $\alpha_{\text{root}(u)} > \alpha_{\text{root}(v)}$ (so u is older than v) **then**
- 2: Add new pair $(\alpha, \alpha_{\text{root}(v)})$ to array (birth, death)
- 3: **else**
- 4: Add new pair $(\alpha, \alpha_{\text{root}(u)})$ to array (birth, death)
- 5: **end if**
- 6: **if** $\text{weight}(\text{root}(u)) > \text{weight}(\text{root}(v))$ **then**
- 7: $\text{root}(u)$ becomes the parent of $\text{root}(v)$ in Forest
- 8: Add $\text{weight}(\text{root}(v)) + 1$ to $\text{weight}(\text{root}(u))$
- 9: **else**
- 10: $\text{root}(v)$ becomes the parent of $\text{root}(u)$ in Forest
- 11: Add $\text{weight}(\text{root}(u)) + 1$ to $\text{weight}(\text{root}(v))$
- 12: **end if**

Appendix B: proofs of lemmas and theorems

Proof of Duality Lemma 14. The component of $C^*(\alpha)$ containing the node v_0 corresponds to the boundary contour of the external region of $\text{Del}(C)$. Any region of $\mathbb{R}^2 - C(\alpha)$ enclosed by a contour consists of several Delaunay triangles whose dual nodes form a white component of $C^*(\alpha)$.

A birth of a contour γ in the descending filtration $\{C(\alpha)\}$ means that γ now encloses a new region of $\mathbb{R}^2 - C(\alpha)$. Hence a new white component is born in the dual graph $C^*(\alpha)$, see the evolution of $C(\alpha), C^*(\alpha)$ in Fig. 7.

A death of a contour γ in $\{C(\alpha)\}$ means that γ is no longer encloses a region of $\mathbb{R}^2 - C(\alpha)$. Hence two white components merge into a big one. By the elder rule of persistence [6, p. 150], the youngest component dies, while the oldest component survives and inherits all nodes. \square

The elder rule is a preference for the case when one class has a high persistence and another has a lower persistence over the case when both classes have similar persistences.

Let us recall that Theorem 1 claims that the algorithm HOCTOP runs in $O(n \log n)$ times with $O(n)$ space.

Step-by-step proof of Theorem 1. Constructing the Delaunay triangulation $\text{Del}(C)$ on a cloud of n points with $O(n)$ edges and triangles requires $O(n \log n)$ time and $O(n)$ space [2, Chapter 9] in Steps 1–2 of Algorithm 1. Sorting all $O(n)$ edges in the decreasing order of the length needs $O(n \log n)$ time in Step 3. Going through each of $k = O(n)$ triangles to initialize Forest(α), we set each birth α_v in $O(1)$ time in Steps 4–6. Most expensive Step 12 in the *while loop* is finding $\text{root}(u), \text{root}(v)$. Each root is found recursively by going up along $O(\log n)$ parent links until we come to a self-parent pointing to itself. All other steps in Algorithms 1–4 require only $O(1)$ time. Hence the total time of the *while loop* and HOCTOP is $O(n \log n)$. \square

Appendix C: experiments on counting holes

The left hand side picture in Fig. 9 is horse2-068-180-contour.png from the database ETH80. The right hand side picture is a cloud around the contour with added noise. The captions contain output probabilities of HOCTOP for most likely numbers of holes when the scale α is uniform.

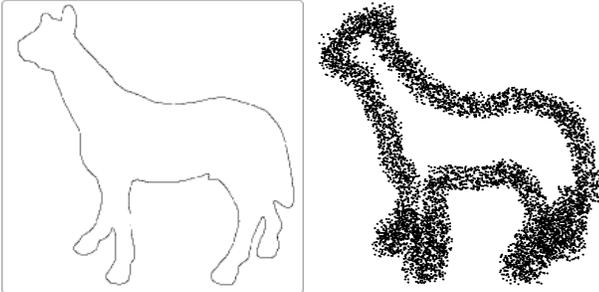


Figure 9. $P(1 \text{ hole}) \approx 52.7\%$, $P(2 \text{ holes}) \approx 25.8\%$, $P(3 \text{ holes}) \approx 9.4\%$, $P(4 \text{ holes}) \approx 2\%$, $P(5 \text{ holes}) \approx 0.5\%$.

The left hand side pictures in Fig. 10–14 are from <http://www.lems.brown.edu/~dmc>. The right hand side pictures are extracted contours with added noise.



Figure 10. $P(1 \text{ hole}) \approx 88.4\%$, $P(2 \text{ holes}) \approx 1.5\%$, $P(0 \text{ holes}) \approx 0.9\%$, $P(13 \text{ holes}) \approx 0.5\%$, $P(5 \text{ holes}) \approx 0.4\%$.



Figure 11. $P(1 \text{ hole}) \approx 66\%$, $P(2 \text{ holes}) \approx 11\%$, $P(3 \text{ holes}) \approx 3.8\%$, $P(4 \text{ holes}) \approx 3.3\%$, $P(6 \text{ holes}) \approx 1.1\%$.

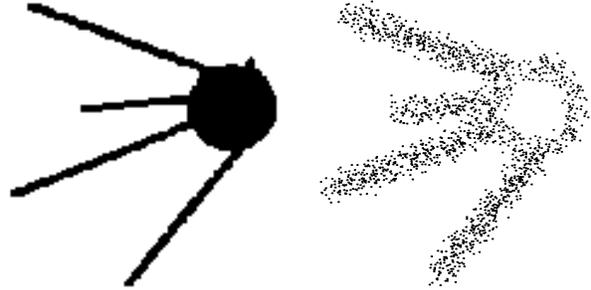


Figure 12. $P(1 \text{ hole}) \approx 58.3\%$, $P(2 \text{ holes}) \approx 19.3\%$, $P(3 \text{ holes}) \approx 4.2\%$, $P(4 \text{ holes}) \approx 1.6\%$, $P(8 \text{ holes}) \approx 0.8\%$.

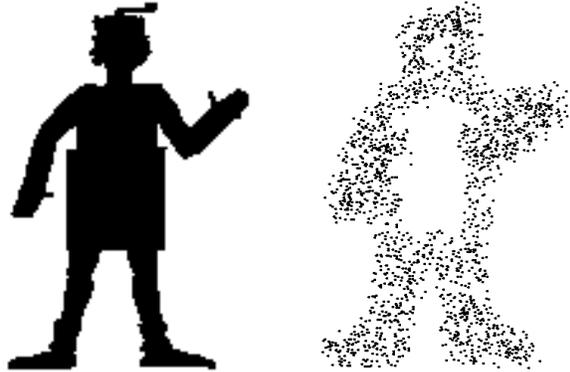


Figure 13. $P(1 \text{ hole}) \approx 49.6\%$, $P(2 \text{ holes}) \approx 21.1\%$, $P(3 \text{ holes}) \approx 4.7\%$, $P(4 \text{ holes}) \approx 3.3\%$, $P(5 \text{ holes}) \approx 1.8\%$.



Figure 14. $P(2 \text{ holes}) \approx 43.7\%$, $P(1 \text{ hole}) \approx 27.8\%$, $P(3 \text{ holes}) \approx 2.5\%$, $P(5 \text{ holes}) \approx 2.1\%$, $P(6 \text{ holes}) \approx 1.6\%$.

The left hand side pictures in Fig. 15–17 contain a cloud C uniformly generated around *wheels* (the boundaries of regular polygons with the radii to all vertices). The middle pictures show the persistence diagrams $\text{PD}\{C^\alpha\}$. The right hand side pictures are the staircases $\text{PS}\{C^\alpha\}$ giving the number of holes of C depending on the scale α .

The left hand side pictures in Fig. 18–20 are noisy clouds around square lattices containing 25, 36, 49 small squares. The algorithm HOCTOP finds the expected number 49 of holes in Fig. 20 when even humans may struggle.

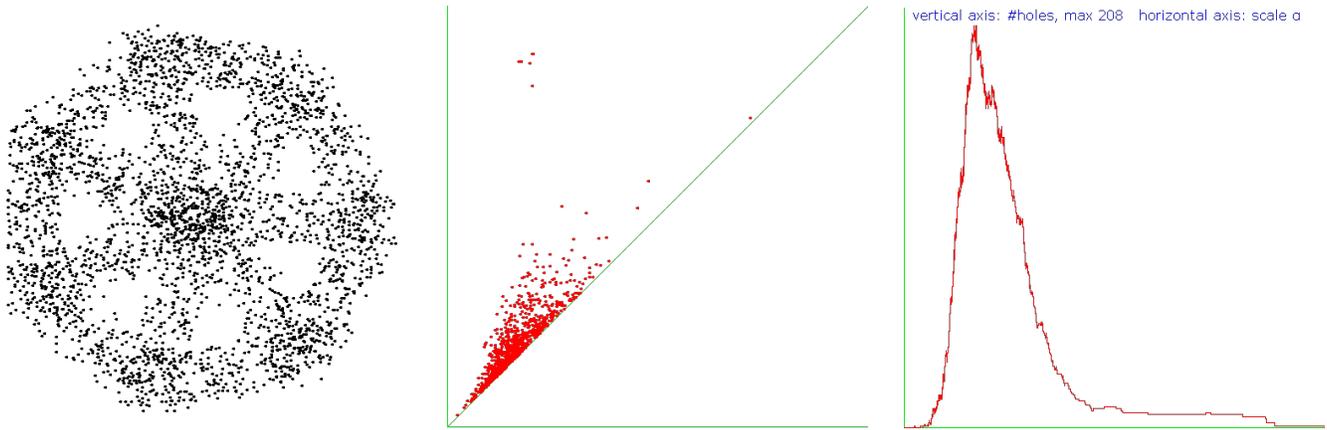


Figure 15. $P(7 \text{ holes}) \approx 22\%$, $P(1 \text{ hole}) \approx 14\%$, $P(8 \text{ holes}) \approx 7.5\%$, $P(6 \text{ holes}) \approx 5.8\%$, $P(10 \text{ holes}) \approx 4.4\%$.

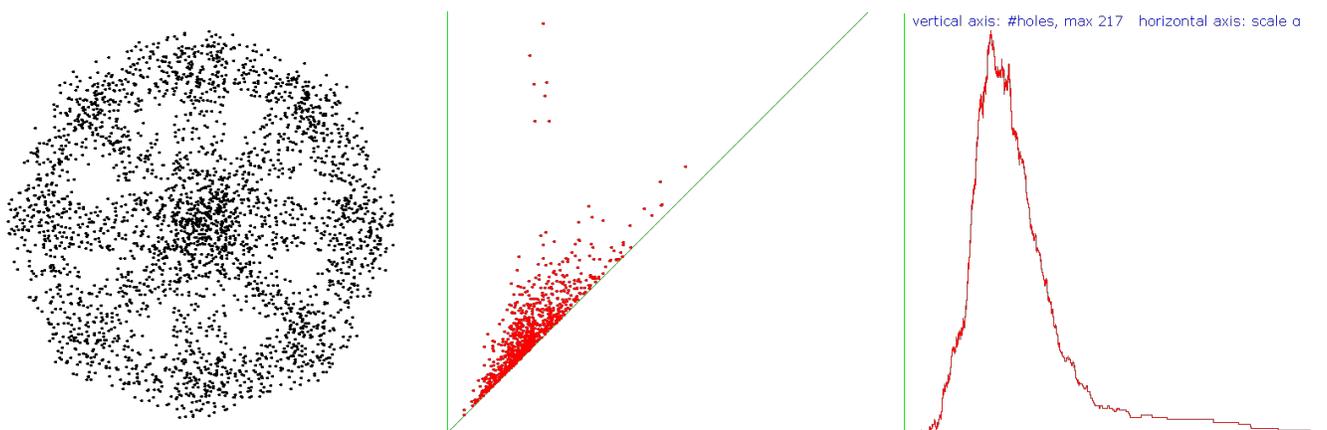


Figure 16. $P(8 \text{ holes}) \approx 11.5\%$, $P(2 \text{ holes}) \approx 8.5\%$, $P(3 \text{ holes}) \approx 7\%$, $P(9 \text{ holes}) \approx 6.8\%$, $P(6 \text{ holes}) \approx 6.5\%$.

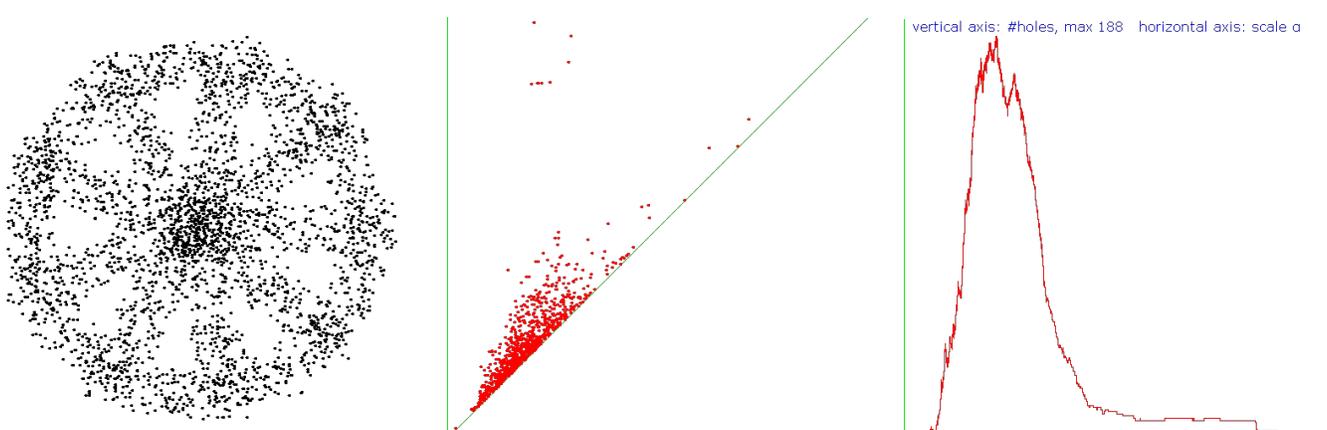


Figure 17. $P(9 \text{ holes}) \approx 18.5\%$, $P(10 \text{ holes}) \approx 11.3\%$, $P(3 \text{ holes}) \approx 6.8\%$, $P(3 \text{ holes}) \approx 6.8\%$, $P(4 \text{ holes}) \approx 5.3\%$.

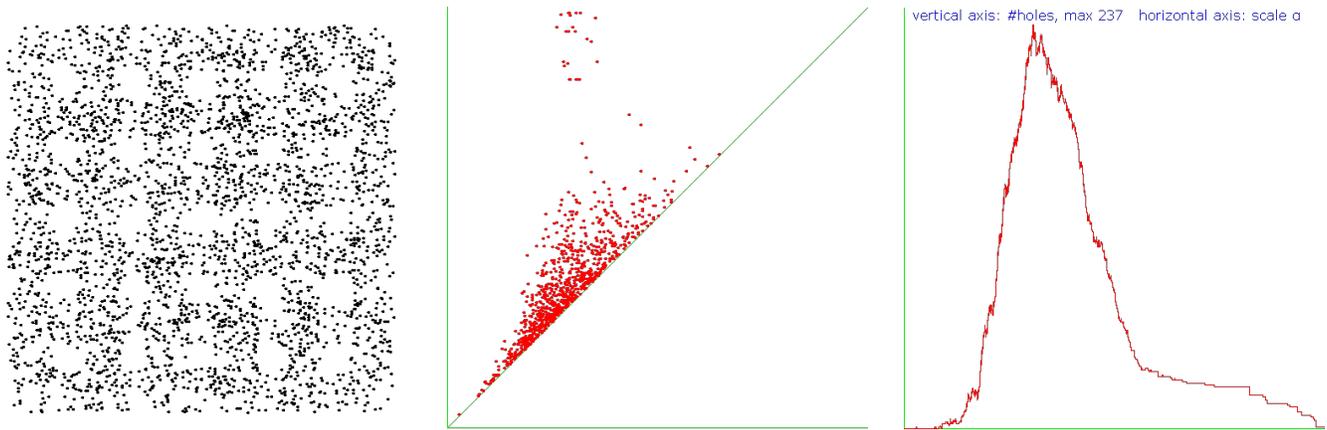


Figure 18. $P(25 \text{ holes}) \approx 8.8\%$, $P(0 \text{ holes}) \approx 5.4\%$, $P(15 \text{ holes}) \approx 5\%$, $P(27 \text{ holes}) \approx 4.6\%$, $P(20 \text{ holes}) \approx 3.5\%$.

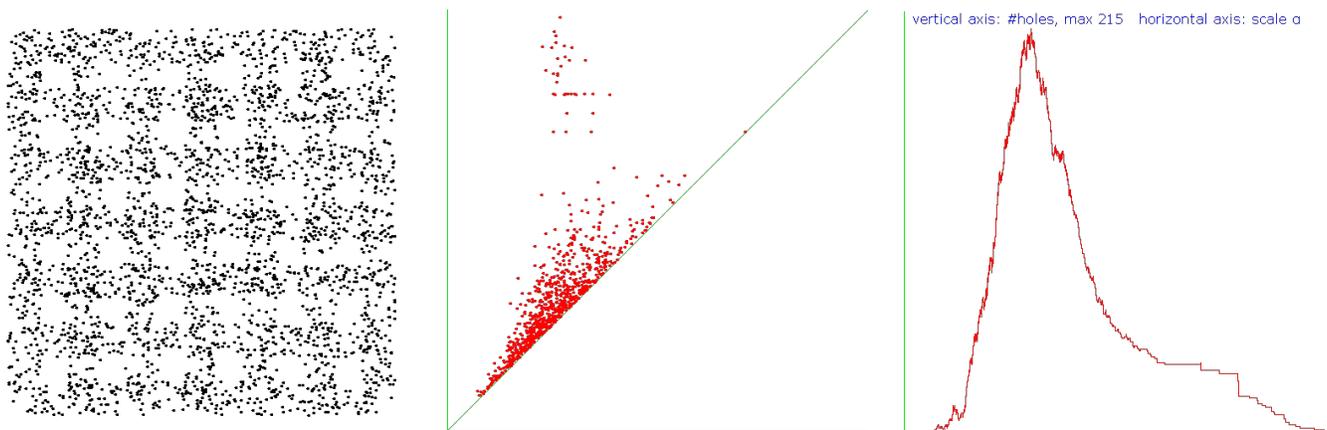


Figure 19. $P(36 \text{ holes}) \approx 9.4\%$, $P(31 \text{ holes}) \approx 4.8\%$, $P(33 \text{ holes}) \approx 4.8\%$, $P(2 \text{ holes}) \approx 4.6\%$, $P(1 \text{ hole}) \approx 3.2\%$.

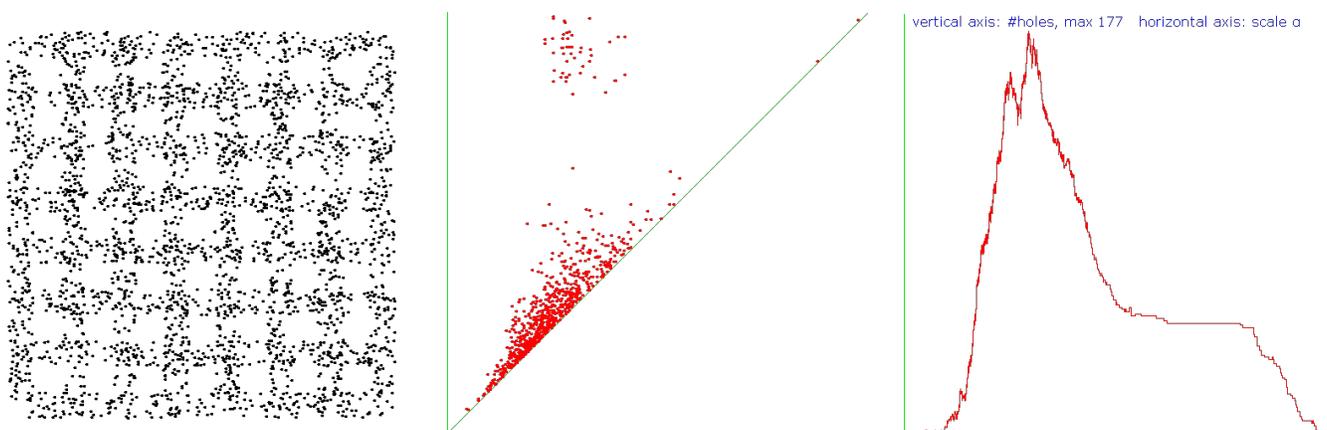


Figure 20. $P(49 \text{ holes}) \approx 18.4\%$, $P(52 \text{ holes}) \approx 4.9\%$, $P(1 \text{ hole}) \approx 3.3\%$, $P(6 \text{ holes}) \approx 2.8\%$, $P(0 \text{ holes}) \approx 2.7\%$.