

Order- k α -Hulls and α -Shapes

Dmitry Krasnoshchekov*

Valentin Polishchuk†

Abstract We introduce *order- k α -hulls and α -shapes* — generalizations of α -hulls and α -shapes. Being also a generalization of k -hull (known in statistics as “ k -depth contour”), order- k α -hull provides a link between shape reconstruction and statistical depth. As a generalization of α -hull, order- k α -hull gives a robust shape estimation by ignoring locally up to k outliers in a point set. Order- k α -shape produces an “inner” shape of the set, with the amount of “digging” into the points controlled by k . As a generalization of k -hull, order- k α -hull is capable of determining “deep” points amidst samples from a *multimodal* distribution: it correctly identifies points which lie outside clusters of samples.

The order- k α -hulls and α -shapes are related to order- k Voronoi diagrams in the same way in which α -hulls and α -shapes are related to Voronoi diagrams. This implies that order- k α -hull and α -shape can be readily built from the order- k Voronoi diagram, and that the number of different order- k α -shapes for all possible values of α is proportional to the complexity of the order- k Voronoi diagram.

1 Introduction

Convex hull of a point set is a truly classical, multidisciplinary concept. One definition of the convex hull is that it is the complement of the union of all halfspaces that contain no points from the set. Over years, several generalizations of the convex hull have been proposed; among them are the α -hull and α -shape and the k -hull (or k -depth contour).

α -Hull and α -Shape

The α -hull of a point set is the complement of the union of all balls of radius α that contain no points from the set; α -shape is a “straight-line version” of the α -hull (see Section 2 for formal definitions). The α -hull and α -shape formally capture an intuitive notion of “shape” of the point set. Since the original paper by Edelsbrunner, Kirkpatrick and

Seidel [16], a variety of extensions of α -hull and α -shape have been introduced [9, 27, 35].

Often the set of points representing a geometric object is obtained by sampling the object in the presence of noise, which can introduce outliers in the data. Handling noise and/or outliers is studied in [10, 21, 29]; similarly to us, [10, 21] rely on constructions involving k nearest neighbors. Except for this reliance, however, our idea is completely independent from the techniques developed in [10, 21]. Robust fitting of shapes to point sets was studied in [4, 22]; unlike works on α -shapes, the latter papers considered fitting a *given* shape to the point cloud.

In this paper we define yet another extension of α -shapes and α -hulls. Our extension is capable of ignoring a certain amount of outliers, which results in a more robust shape reconstruction.

Thinning

There exists several concepts capturing the idea of shape of a polygon, e.g., the Voronoi diagram of its vertices and edges, the medial axis, the skeleton [5, 11, 18]. In graphics, skeletonizing is commonly done via “iterative thinning” which produces a series of objects each being a “shaved” version of the original one [23, 25]. An example of thinned polygon is the offset curve; efficient algorithms for computing the curve were obtained by Eppstein and Erickson [18] (see also [6]).

Our order- k α -shape produces a thinned version of a point set, with the amount of thinning determined by k . The thinning method employed by order- k α -shape may be classified as parallel [30] because it collects all points to be removed from consideration before producing the final shape.

k -Hull and Statistical Depth

The k -hull of a point set, a generalization of the convex hull introduced by Cole, Sharir and Yap in [13], is the locus of points such that any (closed) halfspace through a point in the k -hull contains at least k points from the set. k -hulls are widely used in statistics where they are known under the name of *k -depth contours* because they are the level sets of the “location depth” (also known as halfspace

*Institute for Geospheres Dynamics, Russian Academy of Sciences. krasnd@idg.chph.ras.ru

†Helsinki Institute for Information Technology, CS Dept, University of Helsinki. polishch@cs.helsinki.fi

depth) [28]. Besides the location depth, the rich field of statistical data depth has produced a number of other depth functions — Tukey depth, statistical depth, travel depth, Mahalanobis depth, majority depth, simplicial depth, convex hull peeling depth, proximity depth, cone (wedge) depth; see [2] for references on depths. Any depth is a measure of “interiority” of a point with respect to the given point set. The purpose of depths is to identify “non-relevant” data elements. Many of the depths generalize convex hull in the sense that the points on the convex hull are those of depth 1.

Zuo and Serfling [37] describe several desirable properties of a depth. One of the 5 properties, Maximality at center, requires that the depth has a uniquely defined center; the depth must attain its maximum at the center. Another property, Monotonicity relative to deepest point, requires that the depth monotonically decreases with distance from the center. In general, depths are assumed (often implicitly) to be unimodal functions of distance to a “center”.

Our extension of the depth function allows it to have local maxima. This means that the point set in question may consist of several “clusters” each having its own depth contours (isolines).

Other Related Work

The Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [34] is similar to our order- k α -hulls in that it has the same two parameters: the ball radius and the number of data points inside it. The objective of the DBSCAN is to discover clusters of arbitrary shape. The techniques and the intuition behind the DBSCAN are completely different from those in the present paper.

Our video about order- k α -hulls and α -shapes appeared as [24].

2 Definitions

Let $P \subset \mathbb{R}^2$ be a set of n points in the plane. We will assume that no three points of P lie on a line and that no four points are cocircular. Let $\alpha > 0$ be a positive real number; let $k \geq 1$ be a positive integer. The next definition recapitulates several concepts related to the subject of this paper (see Fig. 1 for an illustration):

Definition 2.1. [13, 16, 26] *The α -hull of P is the complement of the union of all disks of radius α that contain no points of P in their interior. The α -shape of P is a straight-line graph whose vertices are points of P and whose edges connect two points $p, q \in P$ whenever there exists a disk of radius α*

that has p, q on the boundary and contains no points of P in its interior. The k -hull of P is the complement of the union of all halfplanes that contain less than k points of P . The order- k Voronoi diagram of P is the decomposition of the plane into cells such that within one cell the k closest points of P are the same.

We now introduce our new notions.

Definition 2.2. *An α -ball is an (open) disc of radius α . A k -empty α -ball is an α -ball containing (in its interior) less than k points of P . A k -maximal α -ball is an α -ball containing (in its interior) exactly $k - 1$ points of P .*

Definition 2.3. [order- k α -hull.] *The order- k α -hull of P is the complement of the union of all k -empty α -balls.*

Let $p, q \in P$ be a pair of distinct points from P .

Definition 2.4. *An α -ball is leaning on p, q if it has p, q on its boundary. Points p, q are k - α -neighbors if there exists a k -maximal α -ball leaning on p, q .*

The next definition captures the notion of the points “outside” and “inside” P .

Definition 2.5. *The points of P contained in a k -maximal α -ball leaning on k - α -neighbors p, q , are called (k, α, p, q) -outside. The points that are (k, α, p, q) -outside for at least one pair of k - α -neighbors p, q , are called (k, α) -outside. The points of P that are not (k, α) -outside, are called (k, α) -inside.*

Definition 2.6. [order- k α -shape.] *The order- k α -shape of P is the α -shape of the (k, α) -inside points.*

With the above definitions, the α -hull and α -shape are order-1 α -hull and order-1 α -shape respectively, and the k -hull is order- k ∞ -hull.

In a view introduced by Fischer [19] inspired by Edelsbrunner and Mücke [17], the plane \mathbb{R}^2 is an infinite piece of icecream, and P is a set of chocolate chips in it. With this interpretation the convex hull, the k -hull and the α -hull of P can be obtained, by a person allergic to chocolate, using the following kitchenware (Fig 2):

knife \rightarrow **convex hull** A knife is a line. The person is allowed to cut out an arbitrary halfplane empty of chocolate chips and eat it. The convex hull is what remains after all such halfplanes are eaten up.

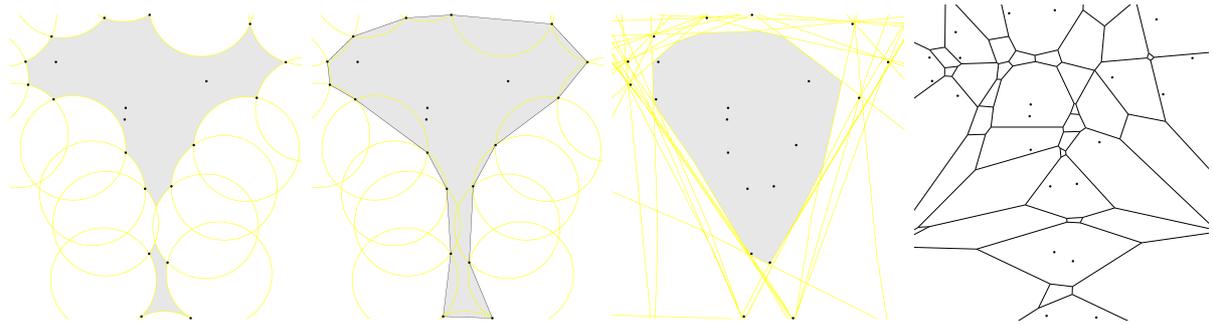


Figure 1: From left to right: α -hull, α -shape, 2-hull and order-2 Voronoi diagram of the same set.

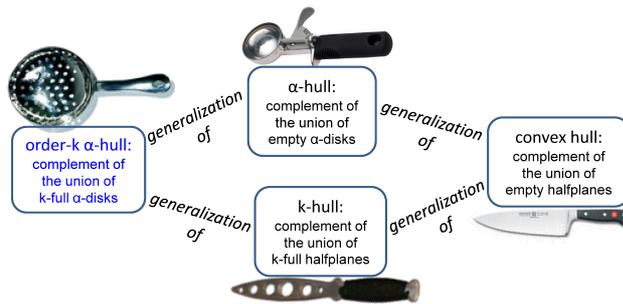


Figure 2: The family picture.

cheeseknife \rightarrow k -**hull** A cheeseknife has k holes in the blade. Thus a halfplane cut out by the cheeseknife may contain up to k chips; the chips are "sieved" through the holes and remain at their places as the halfplane is cut out. The k -hull is the icecream that remains after all such halfplanes are eaten up.

scoop \rightarrow α -**hull** A scoop is a disk of radius α . The person is allowed to scoop out an arbitrary disk empty of chocolate and eat it. The α -hull is what remains after all such disks are eaten up.

The order- k α -hull of P is what remains after eating up icecream with a **colander** of radius α having k holes; i.e, every scoop is a disk of radius α containing at most k chips (sieved through the colander's holes as the disk is removed).

3 Applications

In this section we describe situations in which using the above defined generalizations leads to different (subjectively — more appealing) output than when using the basic concepts from which our generalizations are derived.

Ignoring Outliers

In the literature there is an abundance of theorems and pictures showing how α -hull correctly captures the shape of a point set (see, e.g., [16], [14, Fig. 13.6]). It is possible though to construct examples where α -hull does less than perfect job. Figure 3(a) shows α -hull of a point set nicely reconstructing the set shape. In Figure 3(b) two outliers are added. Now the α -hull appears to miss an important feature of the set — the inner hole. The idea of order- k α -hull was prompted by trying to cope with outliers by allowing few points to reside inside the α -balls defining the hull. In Fig. 3(c,d), order-2 and order-3 α -hulls of the set with the outliers are shown. It can be seen how order- k α -hull with $k > 1$, being less sensitive to the outliers, provides a more robust result.¹ Another example is in Figure 4.

Reconstructing Inner Shape

The original α -shape, not allowing to have points of P inside the α -balls, follows the outer boundary of the point set. With $k > 1$, order- k α -shape "shaves off" k points from outside of P — specifically, the points that are (k, α) -outside. As a result the inner shape of the point set is obtained. See Figure 5 for an example.

Depth Function for Clustered Data

Order- k α -hull, being a generalization of k -hull, allows one to define a generalization of the halfspace depth function. Specifically, for a point $q \in \mathbb{R}^2$ we define the α -depth of q with respect to P to be the

¹In [31], Radke observes that an example in [32] "suggests that structures such as convex hull and α -shapes are often not good boundary generators in real world situations as points on the periphery of a point set lie on the boundary itself. The argument suggests that in biological and ecological studies, a buffer zone is needed between the extreme point locations and the boundary of the point set." Using our order- k α -hull with $k > 1$ serves the purpose of providing the "buffer zone".

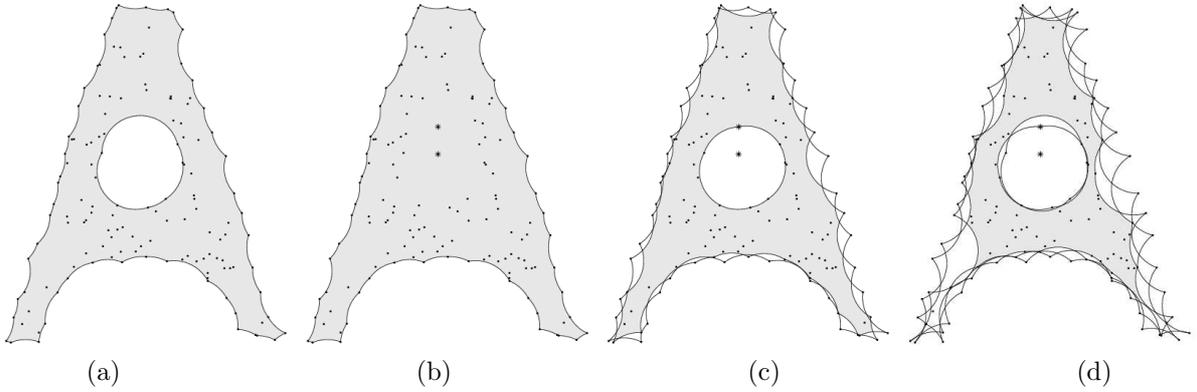


Figure 3: Order- k α -hull ignores the outliers. (a): An α -hull. (b): Two outliers change the picture. (c) and (d): Order- k α -hull with $k = 2, 3$ brings it back.

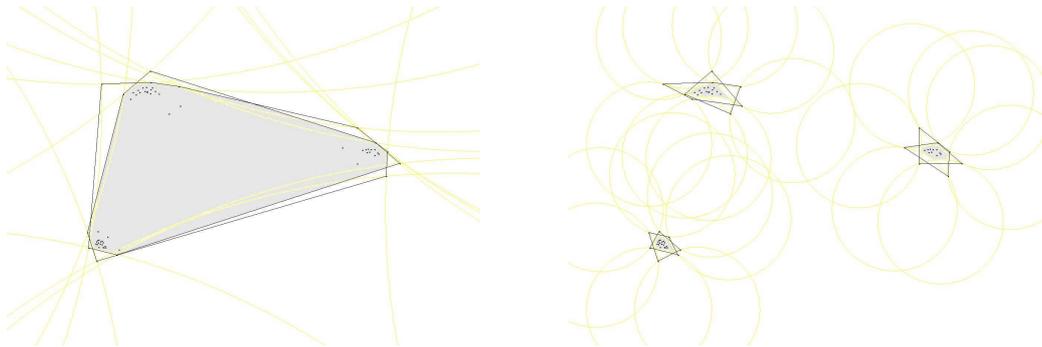


Figure 6: Points come in 3 clusters. Left: The 2-hull (2-depth contour) oversmooths the data. Right: Order-2 α -hull (for a suitable $\alpha < \infty$) better fits the data by identifying 2-depth contours independently within each cluster.

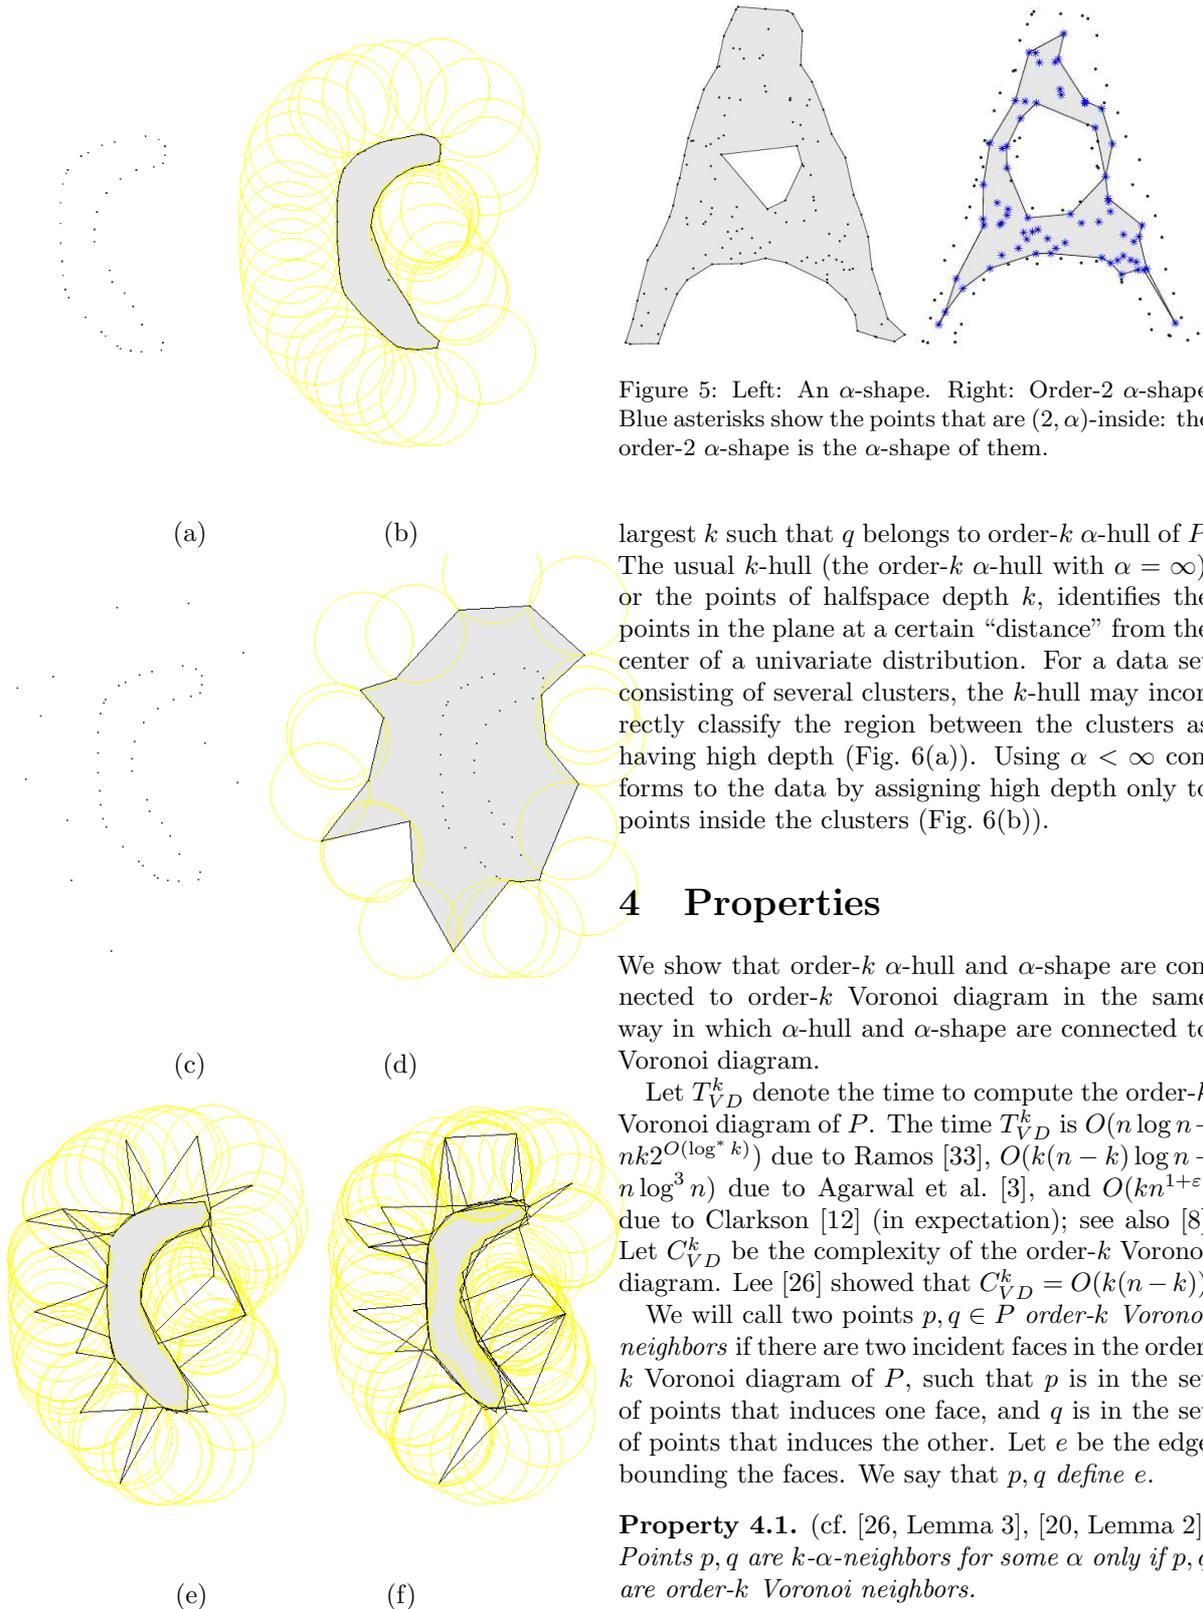


Figure 4: (a) Samples from a “C”-shape. (b) α -shape of the samples. (c) Outliers are added to the picture. (d) The outliers distort the α -shape. (e,f) With $k = 2$ or 3, influence of the outliers is diminished.

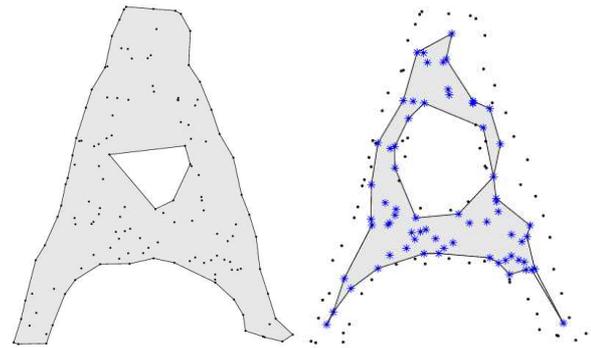


Figure 5: Left: An α -shape. Right: Order-2 α -shape. Blue asterisks show the points that are $(2, \alpha)$ -inside: the order-2 α -shape is the α -shape of them.

largest k such that q belongs to order- k α -hull of P . The usual k -hull (the order- k α -hull with $\alpha = \infty$), or the points of halfspace depth k , identifies the points in the plane at a certain “distance” from the center of a univariate distribution. For a data set consisting of several clusters, the k -hull may incorrectly classify the region between the clusters as having high depth (Fig. 6(a)). Using $\alpha < \infty$ conforms to the data by assigning high depth only to points inside the clusters (Fig. 6(b)).

4 Properties

We show that order- k α -hull and α -shape are connected to order- k Voronoi diagram in the same way in which α -hull and α -shape are connected to Voronoi diagram.

Let T_{VD}^k denote the time to compute the order- k Voronoi diagram of P . The time T_{VD}^k is $O(n \log n + nk2^{O(\log^* k)})$ due to Ramos [33], $O(k(n-k) \log n + n \log^3 n)$ due to Agarwal et al. [3], and $O(kn^{1+\varepsilon})$ due to Clarkson [12] (in expectation); see also [8]. Let C_{VD}^k be the complexity of the order- k Voronoi diagram. Lee [26] showed that $C_{VD}^k = O(k(n-k))$.

We will call two points $p, q \in P$ *order- k Voronoi neighbors* if there are two incident faces in the order- k Voronoi diagram of P , such that p is in the set of points that induces one face, and q is in the set of points that induces the other. Let e be the edge bounding the faces. We say that p, q *define* e .

Property 4.1. (cf. [26, Lemma 3], [20, Lemma 2]) *Points p, q are k - α -neighbors for some α only if p, q are order- k Voronoi neighbors.*

Proof. Let B be the k -maximal α -ball leaning on p, q ; let c be its center, let P' be the $k-1$ points of P in the interior of B . Since no four points of P are cocircular, c can be moved slightly so that B includes P' and p — then the k points of P ,

closest to c , will be $P' \cup p$. Similarly, c can be moved slightly so that the k points, closest to c are $P' \cup q$. Thus, c lies on the boundary of the cells of order- k Voronoi diagram, corresponding to $P' \cup p$ and $P' \cup q$. \square

As a corollary of the above property we have:

Property 4.2. *The total (for all possible α) number of pairs of points of P that can be k - α -neighbors is $O(C_{VD}^k)$.*

The next property follows from the proof of Property 4.1.

Property 4.3. *Points p, q are k - α -neighbors if and only if the center of the k -maximal α -ball, leaning on p, q , belongs to an order- k Voronoi edge defined by p, q .*

As a corollary of the last property we have:

Property 4.4. *Let $I = \{|xp| : x \in e\}$ be the set of distances from points on an order- k Voronoi edge e , defined by p, q , to p (or to q); then for any $\alpha \in I$, p, q are k - α -neighbors.*

An important property of the (order-1) α -shape is that for any pair of points from P there exists an interval of α s (possibly, empty or unbounded) such that the points are α -neighbors if and only if α is within the interval [16, Lemma 3], [14, Lemma 13.24]. The reason for that is that the bisector between two Voronoi neighbors appears in the Voronoi diagram as only one contiguous interval (i.e., a segment). For $k > 1$ the bisector between order- k Voronoi neighbors may appear in the order- k Voronoi diagram more than once (see, e.g., Fig. 4 in [26]), i.e., order- k Voronoi diagram may have more than one edge defined by the same pair. Thus the generalization of the property to order- k α -shape is as follows:

Property 4.5. *For any order- k Voronoi neighbors p, q there are $O(k)$ intervals of the real line such that p, q are k - α -neighbors if and only if α is within one of the intervals.*

Proof. Suppose that p, q are k - α -neighbors for some α ; and let B be the k -maximal α -ball leaning on p, q . Without loss of generality assume that pq is horizontal and that the center of B is above pq . Let P_a, P_b be the points in B that are below and above pq resp. (Fig. 7). By Property 4.3, the center of the k -maximal α -ball leaning on p, q , belongs to the edge of the order- k Voronoi diagram of P . As α changes, the center of the α -ball moves continuously along the edge. As α increases, the points in P_a always stay inside B . Hence the number of points inside B

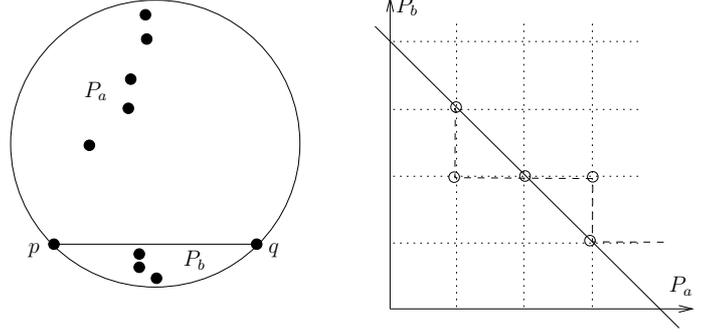


Figure 7: As α changes, the number of points in B below and above pq changes, which defines a path in the (P_a, P_b) -plane. The path has at most $k - 1$ bends, and hence intersects the line $P_a + P_b = k - 1$ at most k times.

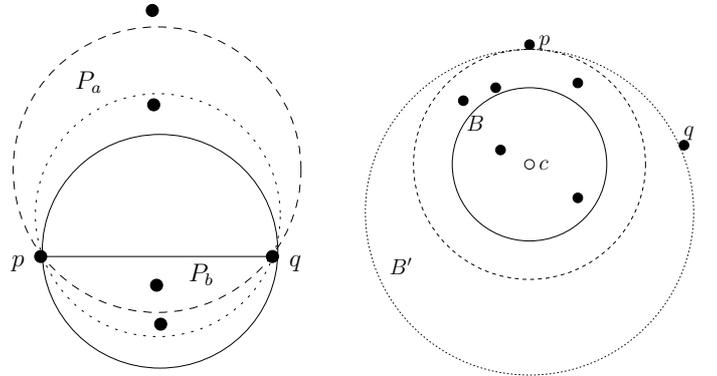


Figure 8: As α increases, gaining points to P_a alternates with losing points from P_b . Each gain makes p, q cease to be k - α -neighbors; each loss makes p, q k - α -neighbors. In this example $k = 3$.

Figure 9: Either B' leans on p, q while having $k - 1$ points inside, or there is a half-space through p that contains B and contains k points from P (p and the $k - 1$ points inside B .)

can go down only if a point of P_b leaves B , which can happen at most $|P_b|$ times. Similarly, as α decreases, the number of points inside B can go down at most P_a times, as a point in P_a leaves B . Thus, the “ α -neighborhood” status of p, q can change at most $|P_a| + |P_b| = k - 1$ times. Similar reasoning applies to the α -ball with the center below pq . \square

We call the intervals from the above property the *neighborhood intervals* of p, q . The neighborhood intervals for all pairs of P overlaid on the real line, decompose the line into finer intervals which we call the *shape intervals* (cf. the “shape spectrum” [16, Definition 7]). Within a shape interval every pair of points of P either forms k - α -neighbors or does not, for any α in the interval. A trivial bound on

the total number of the shape intervals, following from Properties 4.2 and 4.5, is $O(kC_{VD}^k)$ (it is easy to come up with examples in which for a pair of points there exist $\Omega(k)$ neighborhood intervals, see, e.g., Fig. 8). The next property asserts that it is actually just $O(C_{VD}^k)$.

Property 4.6. *The number of shape intervals is $O(C_{VD}^k)$. All the intervals can be found in $O(T_{VD}^k + C_{VD}^k \log C_{VD}^k)$ time.*

Proof. By Properties 4.3 and 4.4, each appearance of the bisector between p and q as an edge of the order- k Voronoi diagram gives rise to one neighborhood interval. Thus the total number of neighborhood intervals over all possible pairs of k - α -neighbors, is $O(C_{VD}^k)$. To find the intervals consider edges of the order- k Voronoi diagram one by one. Let p, q be the points defining an edge e . During construction of the diagram we can store the maximum and minimum distances from e to p, q . By Property 4.4, these distances define the endpoints of a neighborhood interval. To get the shape intervals we overlay the neighborhood intervals in additional $O(N \log N)$ time where $N = O(C_{VD}^k)$ is the number of the neighborhood intervals. \square

It follows from the proof of Property 4.5 that the set of (k, α, p, q) -outside points is the same for any α within a neighborhood interval of (p, q) . Thus, the set of (k, α) -outside points stays the same for any α within a shape interval; in (other) words, the same points of P are treated as outliers. Since the order- k α -shape stays the same as long as the (k, α) -outside points are the same, there is one order- k α -shape per shape interval. Hence,

Property 4.7. *The number of order- k α -shapes for all values of α is $O(C_{VD}^k)$.*

5 Construction

We use properties from the previous section to give algorithms for constructing order- k α -shape and α -hull from order- k Voronoi diagram of P .

Theorem 5.1. *Order- k α -shape can be built in $O(T_{VD}^k + C_{VD}^k \log C_{VD}^k)$ time.*

Proof. We show how to identify all (k, α) -outside points within the claimed time bounds. Then the order- k α -shape of P , which is the α -shape of the remaining points, can be built in $O(n \log n)$ time using the original α -shape algorithm [16].

Go through all edges of the order- k Voronoi diagram of P . For every pair of points defining an edge check whether the center of each α -ball leaning on

the points belongs to the edge. Let C be the set of the centers that do. By Property 4.3, the union of α -balls centered on C contains all (k, α) -outside points.

Build the Voronoi diagram of C and preprocess it for point location; by Property 4.7, $|C| = O(C_{VD}^k)$, so this can be done in $O(C_{VD}^k \log C_{VD}^k)$ time. Then for any point of P test whether it is closer than α to some point in C . If yes, mark the point as (k, α) -outside; this can be done in $O(n \log C_{VD}^k) = O(C_{VD}^k \log C_{VD}^k)$ time. \square

Next we give an algorithm to construct the order- k α -hull. Following [16], we first state an auxiliary lemma:

Lemma 5.2. (cf. [16, Lemma 5]) *Let B be a k -empty α -ball. Either B lies outside the k -hull of P , or there exists a k -empty α -ball B' , with the center on an edge of the order- k Voronoi diagram of P and such that $B \subseteq B'$.*

Proof. Let c be the center of B . Increase the radius of B until it contains $k-1$ points of P . Continue increasing the radius until B hits a point $p \in P$. Now start growing B , moving its center along pc (away from p), so that it is always touching p (Fig. 9). Note that during the growing the original B always stays inside the grown ball.

If such growing can continue indefinitely, with B never hitting another point of P , then there is a halfspace (the one bounded by the perpendicular to pc through p) that contains k points of P : the $k-1$ points that are inside B plus p itself. Thus, the halfspace is outside the k -hull of P . Since B belongs to the halfspace, B is outside the k -hull.

Otherwise let B' be the ball obtained by the growing procedure at the time when B hits another point $q \in P$. Now the center of B' is on the edge of order- k Voronoi diagram of P defined by p and q . \square

Theorem 5.3. *Order- k α -hull can be built in $O(T_{VD}^k + C_{VD}^k \log C_{VD}^k)$ time.*

Proof. Following [16], we redefine order- k α -hull as the complement of the union of all k -empty balls of radius at least α (this is an equivalent definition since any ball can be represented as the union of balls of smaller radii). By Lemma 5.2, order- k α -hull is the complement of the union of the k -empty balls of radius at least α whose centers lie on the edges of order- k Voronoi diagram. For each edge the union of such balls is the union of at most 4 balls (cf. Facts 8 and 9, and Lemma 9 in [16]). Thus, the (complement of) order- k α -hull can be expressed as the union of $O(C_{VD}^k)$ balls. \square

In fact, the union of the k -empty balls with centers on the edges of the order- k Voronoi diagram is the union of at most *two* balls (not four, as in the above theorem and in [16]). Indeed, for each vertex v of the order- k Voronoi diagram let the *Voronoi ball* of v be the largest k -empty ball centered at v ; the ball is the witness that v is a vertex of the diagram. It is easy to see that the order- k α -hull is the union of the Voronoi balls whose radius is at least α . The Voronoi balls can be stored with each Voronoi vertex during the construction of the diagram. Then, given a value of α , one can build the α -hull by computing the union of the relevant α -balls — those with radius at least α (which includes the possibility of radius being infinite, in which case the ball is just a halfplane).

6 Discussion

Any generalization of α -shape (e.g., anisotropic α -shape [35]) can be enhanced by considering its order- k counterpart. It is also possible to vary k depending, e.g., on the local sampling density or on accuracy (for outlier detection), or on feature size.

Other generalizations along the same lines are possible. For instance, one can define the α -rank of an (ordered) pair of (distinct) points (p, q) from P as the number of points in P that belong to the α -ball leaning on p and q and having its center to the left of the directed line \overrightarrow{pq} . Then the rank in the sense of Edelsbruner [14, Section 12.5] is the ∞ -rank. As with the order- k α -hull, using $\alpha < \infty$ more consistently determines ranks of points in data sets consisting of several clusters.

Finally, one can define similar order- k versions of other graphs on point sets, e.g., β -skeletons; we coded a Java applet [1] to demonstrate this.

Implementation The algorithms from Section 5 are based only on standard computational geometry tools—order- k Voronoi diagrams, point location data structure, α -shapes, etc.—for which a variety of implementations are available; we thus expect our constructions to be practical. (Experimenting with real-world datasets is beyond the scope of this note.) A Java applet with a brute-force implementation [36] easily handles sets with few hundred points almost in real time, as the points are clicked in by the user and k and α are changed.

Open problems It would be interesting to see how order- k α -shapes inherit various properties of the standard α -shapes, such as:

- homotopy equivalence to the Minkowski sum of P and an α -ball [15, p. 124],
- existence of intervals for α classifying the simplices of the “ α -complex” [7],
- properties of higher-dimensional α -shapes [17] (the definitions of order- k α -hull and α -shape extend verbatim to higher dimensions, but it is not obvious whether the constructions can be used in practice even in dimension 3).

Acknowledgments

We thank the anonymous reviewers for their valuable comments and David Kirkpatrick, Sheung-Hung Poon and Don Sheehy for discussions. The work was supported by the Academy of Finland grants 1138520 and 261019.

References

- [1] <http://www.cs.helsinki.fi/group/compgeom/applets/shapes/>, <http://tinyurl.com/k-order-graphs>.
- [2] Statistical Data Depth. <http://www.cs.tufts.edu/research/geometry/depth/>.
- [3] P. K. Agarwal, M. de Berg, J. Matoušek, and O. Schwarzkopf. Constructing levels in arrangements and higher order Voronoi diagrams. *SIJCOMP*, 27:654–667, 1998.
- [4] P. K. Agarwal, S. Har-Peled, and H. Yu. Robust shape fitting via peeling and grating coresets. *Discrete Comput. Geom.*, 39(1):38–58, 2008.
- [5] O. Aichholzer and F. Aurenhammer. Straight skeletons for general polygonal figures in the plane. In *COCOON '96*, pages 117–126.
- [6] G. Barequet, D. Eppstein, M. T. Goodrich, and A. Vaxman. Straight skeletons of three-dimensional polyhedra. In *ESA 2008*, pages 148–160.
- [7] F. Bernardini and C. L. Bajaj. Sampling and reconstructing manifolds using alpha-shapes. In *CCCG'97*.
- [8] J.-D. Boissonnat, O. Devillers, and M. Teilaud. A semidynamic construction of higher-order voronoi diagrams and its randomized analysis. *Algorithmica*, 9:329–356, 1993.
- [9] F. Cazals, J. Giesen, M. Pauly, and A. Zomorodian. The conformal alpha shape filtration. *The Visual Computer*, 22(8):531–540, 2006.

- [10] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.
- [11] F. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. *Discrete Comput. Geom.*, 21(3):405–420, 1999.
- [12] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.
- [13] R. Cole, M. Sharir, and C. K. Yap. On k -hulls and related problems. *SIJCOMP*, 16(1):61–77, 1987.
- [14] H. Edelsbrunner. *Algorithms in combinatorial geometry*. Springer-Verlag, 1987.
- [15] H. Edelsbrunner. Shape reconstruction with delaunay complex. In *LATIN*, 1998.
- [16] H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory*, IT-29:551–559, 1983.
- [17] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13(1):43–72, 1994.
- [18] D. Eppstein and J. Erickson. Raising roofs, crashing cycles, and playing pool: applications of a data structure for finding pairwise interactions. In *SoCG '98*.
- [19] K. Fischer. Introduction to alpha shapes, 2000.
- [20] J. Gudmundsson, M. Hammar, and M. van Kreveld. Higher order Delaunay triangulations. *Computational Geometry*, 23:85–98, 2002.
- [21] L. J. Guibas, Q. Mérigot, and D. Morozov. Witnessed k -distance. In *SoCG*, pages 57–64, 2011.
- [22] S. Har-Peled and Y. Wang. Shape fitting with outliers. *SIJCOMP*, 33(2):269–285, 2004.
- [23] T. Ju, M. L. Baker, and W. Chiu. Computing a family of skeletons of volumetric models for shape description. *Comput. Aided Des.*, 39(5):352–360, 2007.
- [24] D. N. Krasnoshchekov, V. Polishchuk, and A. Vihavainen. Shape approximation using k -order alpha-hulls. In *SoCG'10*.
- [25] L. Lam, S.-W. Lee, and C. Y. Suen. Thinning methodologies—a comprehensive survey. In *Document image analysis*, pages 61–77. 1995.
- [26] D. T. Lee. On k -nearest neighbor Voronoi diagrams in the plane. *IEEE Trans. Comput.*, C-31:478–487, 1982.
- [27] N. Lomenie, G. Stamon, L. Gallo, and N. Cambou. Morphological operations on Delaunay triangulations. In *ICPR '00*, page 3556.
- [28] K. Miller, S. Ramaswami, P. Rousseeuw, J. A. Sellarès, D. Souvaine, I. Streinu, and A. Struyf. Efficient computation of location depth contours by methods of computational geometry. *Statistics and Computing*, 13(2):153–162, 2003.
- [29] P. Niyogi, S. Smale, and S. Weinberger. A topological view of unsupervised learning from noisy data. *SIAM J. Comput.*, 40(3):646–663, 2011.
- [30] K. Palágyi and A. Kuba. A parallel 3d 12-subiteration thinning algorithm. *Graph. Models Image Process.*, 61(4):199–221, 1999.
- [31] J. Radke. On the shape of a set of points. In G. Toussaint, editor, *Computational Morphology*.
- [32] J. Radke. On measuring the influence of boundaries. In *Annual Canadian Association of Geographer*, May 1984.
- [33] E. A. Ramos. On range reporting, ray shooting and k -level construction. In *SoCG '99*, pages 390–399, 1999.
- [34] J. Sander, M. Ester, H.-P. Kriegel, and X. Xu. Density-based clustering in spatial databases: The algorithm GDBSCAN and its applications. *Data Mining & Knowledge Discovery*, 2(2):169–194, 1998.
- [35] M. Teichmann and M. V. Capps. Surface reconstruction with anisotropic density-scaled alpha shapes. In *IEEE Visualization*, pages 67–72, 1998.
- [36] A. Vihavainen. <http://www.cs.helsinki.fi/group/compgeom/kapplet/>.
- [37] Y. Zuo and R. Serfling. General notions of statistical depth function. *Ann. Statist.*, 28:461–482, 2000.