

Fitting a Set of Points by a Circle*

J. García-López,¹ P. A. Ramos,² and J. Snoeyink³

¹Departamento de Matemática Aplicada, Escuela Universitaria de Informática,
Universidad Politécnica de Madrid,
28031 Madrid, Spain
jglopez@eui.upm.es

²Departamento de Matemática Aplicada, Facultad de Informática,
Universidad Politécnica de Madrid,
28040 Madrid, Spain
pramos@fi.upm.es

³Department of Computer Science,
University of British Columbia,
Vancouver, BC, Canada V6T 1Z4
snoeyink@cs.ubc.ca

Abstract. Given a set of points $S = \{p_1, \dots, p_n\}$ in Euclidean d -dimensional space, we address the problem of computing the d -dimensional annulus of smallest width containing the set. We give a complete characterization of the centers of annuli which are *locally minimal* in arbitrary dimension and we show that, for $d = 2$, a locally minimal annulus has two points on the inner circle and two points on the outer circle that interlace anglewise as seen from the center of the annulus. Using this characterization, we show that, given a circular order of the points, there is at most *one* locally minimal annulus consistent with that order and it can be computed in time $O(n \log n)$ using a *simple* algorithm. Furthermore, when points are in *convex position*, the problem can be solved in optimal $\Theta(n)$ time.

1. Introduction

Let $S = \{p_1, \dots, p_n\}$ be a set of n points in \mathbb{E}^d and let $d(x, y)$ denote the Euclidean distance between points x and y . If we denote by \mathcal{F} the family of hyperspheres in \mathbb{E}^d , the problem of finding the optimal Tchebyshev fit of S with a hypersphere can be formulated

* The research of the first and second authors was partially supported by DGES Grant PB96-0005-C02-01. The research of the third author was supported in part by an NSERC Research Grant.

as a min-max problem, that is, finding

$$\min_{\gamma \in \mathcal{F}} \max_{i=1, \dots, n} d(p_i, \gamma).$$

We call d -annulus with *center* c , *inner radius* r_1 , and *outer radius* r_2 , denoted by $\mathcal{A}(c, r_1, r_2)$, to the locus of points between two concentric spheres of radius r_1 and r_2 , that is,

$$\mathcal{A}(c, r_1, r_2) = \{x \in \mathbb{E}^d \mid r_1 \leq d(x, c) \leq r_2\}.$$

The *width* of the annulus is $w = r_2 - r_1$ and the problem that we address here can be formulated in a more geometric setting in the following way: compute the annulus of smallest width containing the set S . If we define the function $\mathcal{R}_S(x)$ as the difference of distances between the furthest and the closest neighbor of x in S , we say that an annulus is locally minimal if its center is a local minimum of $\mathcal{R}_S(x)$ and the problem of finding the annulus of smallest width containing the set of points is equivalent to finding the global minimum of $\mathcal{R}_S(x)$.

The main motivation for this problem comes from Computational Metrology. Given an object that has to be tested for circularity ($d = 2$), take a sample of points from the object and measure the circularity of this sample set; then accept the object if the circularity is good enough and reject otherwise. The measure for circularity recommended by international standards is the width of the annulus of smallest width containing S (see pp. 40–42 of [8] and p. 14 of [9]). Further information about the relation between Computational Geometry and Metrology can be found in [15] and [19].

The problem of computing the annulus of smallest width containing a set of points has been extensively studied in the Computational Geometry literature. The first nontrivial algorithm was independently proposed by Wainstein [18], by Roy and Zhang [14], and by Ebara et al. [6], observing that the center of the annulus is either an intersection of an edge of the closest-point Voronoi diagram with an edge of the furthest-point Voronoi diagram, or a vertex of one of these Voronoi diagrams. This leads to an algorithm whose complexity is $O(n \log n + I)$, where I is the number of intersections between closest- and furthest-point Voronoi diagrams and can be quadratic in the worst case (even for sets of points in convex position). In [16] Smid and Janardan point out that the optimal annulus can be a slab. We include this case into the general one by considering a slab as an annulus with center at ∞ .

This characterization of the center of the annulus is used by Agarwal et al. [2] in order to give a solution to the problem in $O(n^{8/5+\epsilon})$ time which has been improved to $O(n^{3/2+\epsilon})$ time by Agarwal and Sharir [1] by reducing the problem to a width-type problem in \mathbb{E}^3 . Both algorithms use parametric search and other sophisticated programming techniques.

Other recent papers deal with some related problems: Mehlhorn et al. [11] give an algorithm which takes into account the sampling process and not only the discrete optimization problem, while Duncan et al. [5] and de Berg et al. [4] solve the problem in $O(n \log n)$ time if the inner, the outer, or the medial radius is fixed.

The equivalent problem for simple polygons has been studied by Le and Lee [10] and has been solved in optimal $O(n)$ time for the special case of convex polygons and center of the annulus inside the polygon by Swanson et al. [17].

The problem has also been treated in a number of papers in different areas. Actually, a complete characterization of the local minima in the two-dimensional case was given by Rivlin [13], showing that a point is a local minimum if, and only if, there are at least

two points on each boundary circle (and, therefore, the center is always an intersection of one edge of the closest-point Voronoi diagram with one edge of the furthest-point Voronoi diagram) and, furthermore, closest and furthest points interlace anglewise as seen from the center of the annulus.

1.1. Our Results

We give necessary and sufficient conditions for a point to be the center of a locally minimal annulus in \mathbb{E}^d . This characterization has important implications on the distribution of local minima in some cases: we prove that, if points are in convex position, there is at most *one* local minimum inside the convex hull. More generally, we see that, given an angular order of a set of points, there is at most *one* local minimum compatible with the given order. Actually, this is the problem that has to be solved in practice, because sampled points are given in order around a *reference center* and the solution we are looking for should preserve this order.

In Section 5 we give a simple $O(n \log n)$ algorithm that computes the local minimum for a given angular order of the points and improve this to optimal linear time if points are in convex position.

Finally, in Section 6 we give an example of a set of points with a quadratic number of local minima for the roundness function.

Before going on, we state the formal definitions and introduce some notation in the next section.

2. Preliminaries

Let $S = \{p_1, \dots, p_n\}$ be a set of points in \mathbb{E}^d . The convex hull of S is denoted by $\text{conv } S$ and the *boundary* of $\text{conv } S$ is denoted by $CH(S)$. $\mathcal{V}_c(S)$ and $\mathcal{V}_f(S)$ denote the closest- and furthest-point Voronoi diagram of S , respectively.

Given a point $x \in \mathbb{E}^d$, $CN(x)$ and $FN(x)$ denote, respectively, the set of closest and furthest neighbors of x in S , that is,

$$CN(x) = \left\{ p_j \in S \mid d(x, p_j) = \min_{i=1, \dots, n} d(x, p_i) \right\},$$

$$FN(x) = \left\{ p_j \in S \mid d(x, p_j) = \max_{i=1, \dots, n} d(x, p_i) \right\}.$$

Then the *roundness function* can be defined in the following way:

$$\mathcal{R}_S(x) = d(x, FN(x)) - d(x, CN(x)).$$

The problem of computing the annulus of smallest width containing S is equivalent to the problem of finding the global minimum of $\mathcal{R}_S(x)$. The special case where the annulus of minimum width is a slab could be included in the general one by extending the definition of $\mathcal{R}_S(x)$ to the real projective d -space, $\mathbf{P}^d\mathbb{R}$. We recall that $\mathbf{P}^d\mathbb{R}$ can be obtained by adding to \mathbb{R}^d a point at ∞ for each family of parallel lines.

Given a point $x \in \mathbf{P}^d \mathbb{R} \setminus \mathbb{R}^d$, we define $\mathcal{R}_S(x)$ as the limit, when y goes to ∞ along a line of the corresponding family, of $\mathcal{R}_S(y)$. Clearly, this limit does not depend on the line and the value is the distance between the supporting hyperplanes of S orthogonal to the family of lines.

If $p, q \in \mathbb{E}^d$, \vec{pq} denotes the vector from p to q and, as usual, we identify point p and vector \mathbf{p} (for a fixed point o). The Euclidean scalar product is denoted by $\mathbf{u} \cdot \mathbf{v}$ and $\|\mathbf{u}\| = (\mathbf{u} \cdot \mathbf{u})^{1/2}$ is the length of the vector.

We use k -plane or k -sphere to denote a plane or a sphere of dimension k and the unit k -sphere is denoted by S^k . Given a set T , its complement is denoted by T^c . As usual, we say that a hyperplane separates two sets of points if each set has points in, at most, one open half-space and the same applies for infinite cones that divide the space into two open connected components.

Finally, we point out that if S is included in a $(d - 1)$ -sphere, the global minimum of $\mathcal{R}_S(x)$ is zero. This property can be checked in linear time so, from now on, we exclude this case from our study.

3. Characterizing Local Minima of $\mathcal{R}_S(x)$

The following lemma, whose proof follows easily from Taylor expansion, will be useful later:

Lemma 1. *Let $p \in \mathbb{E}^d$, $\mathbf{v} \in S^{d-1}$, and let α_p be the angle determined by \mathbf{p} and \mathbf{v} . Then*

$$\|\lambda \mathbf{v} - \mathbf{p}\| = \|\mathbf{p}\| - \lambda \cos \alpha_p + \frac{\lambda^2}{2} \frac{\sin^2 \alpha_p}{\|\mathbf{p}\|} + o(\lambda^2).$$

The next result gives a complete characterization of the local minima of $\mathcal{R}_S(x)$ in arbitrary dimension.

Theorem 1. *o is a local minimum of $\mathcal{R}_S(x)$ if, and only if, closest and furthest neighbors of o cannot be separated by a circular cone with apex at o .*

Proof. We focus on the finite case. If $o \in \mathbf{P}^d \mathbb{R} \setminus \mathbb{R}^d$, the circular cone is a cylinder and the proof follows from a complete analogous reasoning substituting the angles with respect to the axis of the cone by the distances to the axis of the cylinder and, instead of taking derivatives, taking points on the axis of the cylinder “close enough” to ∞ .

We first observe that, if \mathbf{v} defines the axis of the cone, closest and furthest neighbors are not separated if, and only if,

$$\min_{p_i \in FN(o)} \frac{\mathbf{v} \cdot \mathbf{p}_i}{\|\mathbf{p}_i\|} < \max_{p_i \in CN(o)} \frac{\mathbf{v} \cdot \mathbf{p}_i}{\|\mathbf{p}_i\|}. \tag{1}$$

For a given $\mathbf{v} \in S^{d-1}$, let $p_{c(\mathbf{v})}$ denote the point of $CN(o)$ that minimizes the angle between \mathbf{v} and \mathbf{p}_i for $p_i \in CN(o)$ and let $\alpha_{c(\mathbf{v})}$ be that angle, that is,

$$\cos \alpha_{c(\mathbf{v})} = \frac{\mathbf{v} \cdot \mathbf{p}_{c(\mathbf{v})}}{\|\mathbf{p}_{c(\mathbf{v})}\|} = \max_{p_i \in CN(o)} \frac{\mathbf{v} \cdot \mathbf{p}_i}{\|\mathbf{p}_i\|}.$$

In a similar way, $p_{f(\mathbf{v})}$ is the point of $FN(o)$ that maximizes the angle between \mathbf{v} and \mathbf{p}_i for $p_i \in FN(o)$ and let $\alpha_{f(\mathbf{v})}$ be that angle, that is,

$$\cos \alpha_{f(\mathbf{v})} = \frac{\mathbf{v} \cdot \mathbf{p}_{f(\mathbf{v})}}{\|\mathbf{p}_{f(\mathbf{v})}\|} = \min_{p_i \in FN(o)} \frac{\mathbf{v} \cdot \mathbf{p}_i}{\|\mathbf{p}_i\|}.$$

Then, condition (1) says that,

$$\forall \mathbf{v} \in \mathcal{S}^{d-1}, \quad \cos \alpha_{f(\mathbf{v})} < \cos \alpha_{c(\mathbf{v})}. \tag{2}$$

From Lemma 1, it follows that, for $\lambda > 0$ small enough, $p_{c(\mathbf{v})}$ and $p_{f(\mathbf{v})}$ are, respectively, the closest and furthest neighbors of $\lambda \mathbf{v}$ and, from Taylor expansion, we have

$$\mathcal{R}_S(\lambda \mathbf{v}) = \mathcal{R}_S(o) + \lambda (\cos \alpha_{c(\mathbf{v})} - \cos \alpha_{f(\mathbf{v})}) + \frac{\lambda^2}{2} \left(\frac{\sin^2 \alpha_{f(\mathbf{v})}}{\|\mathbf{p}_{f(\mathbf{v})}\|} - \frac{\sin^2 \alpha_{c(\mathbf{v})}}{\|\mathbf{p}_{c(\mathbf{v})}\|} \right) + o(\lambda^2).$$

Therefore, if condition (1) holds, we have $\cos \alpha_{c(\mathbf{v})} - \cos \alpha_{f(\mathbf{v})} > 0$ for all $\mathbf{v} \in \mathcal{S}^{d-1}$ and o is a local minimum. Conversely, if there exists $\mathbf{v} \in \mathcal{S}^{d-1}$ such that $\cos \alpha_{f(\mathbf{v})} \geq \cos \alpha_{c(\mathbf{v})}$, then $\mathcal{R}_S(x)$ decreases in that direction and o is not a local minimum. \square

Observing that the intersection between a circular cone and a hypersphere is contained in a hyperplane, the previous characterization can also be stated saying that if we project points of $CN(o)$ and $FN(o)$ onto a common sphere centered at o , then projections of points of $CN(o)$ cannot be separated from projections of points of $FN(o)$ by a hyperplane.

We also observe that Theorem 1 implies that $|CN(o)| + |FN(o)| \geq d + 2$ if o is a local minimum. To see this, we observe that two sets R and B in \mathbb{E}^d with $|R| + |B| \leq d + 1$ can be separated by a hyperplane (this is easy to prove, using induction on d).

Finally, we point out that $|CN(o)| > 1$ and $|FN(o)| > 1$ because, in other cases, points can be separated by a degenerate cone.

For $d = 2$, this is also equivalent to saying that there are two points on the outer circle and two points on the inner circle of the annulus that interlace anglewise as seen from the center of the annulus.

In the last section we shall see that this characterization can also be the basis for the construction of a point set in the plane with $\Omega(n^2)$ local minima for $\mathcal{R}_S(x)$.

4. Roundness in the Plane

From now on, $S = \{p_1, \dots, p_n\}$ will be a *labeled* set of points in the plane such that polygon \mathcal{P} with vertices p_1, \dots, p_n is *simple*. We also assume that points are ordered in the positive (counterclockwise) direction and indices are understood modulo n . We define the *kernel* of S , denoted by $\ker S$, as the locus of points from which the points of S are seen in the given angular order.

We denote by e_i the edge between p_i and p_{i+1} and by h_i the left half-plane defined by the directed line through p_i, p_{i+1} . For a given edge e_i in the convex hull of \mathcal{P} , we denote by

$$H_i = \bigcap_{j \neq i} h_j.$$




Fig. 1. (a) $\ker S$ is shaded with the ear corresponding to e_i darkened. (b) Illustration for the proof of Lemma 2.

The (possibly empty) part of H_i outside $\text{conv } S$ is called the *ear* of $\ker S$ corresponding to edge e_i (see Fig. 1(a)).

It is important to distinguish between the kernel of the set of points and the kernel of the simple polygon because in some relevant practical situations the solution to the problem is expected to be outside the polygon (for example, when the set is in the shape of an arc of a circle). The following result includes an alternative characterization of $\ker S$ and the relation between this concept and the well-known concept of a kernel of a simple polygon.

Lemma 2.

- (a) $\ker S = \bigcup_{e_i \in CH(\mathcal{P})} H_i$.
- (b) $\ker \mathcal{P} = \ker S \cap \text{conv } S$.
- (c) If \mathcal{P} is star-shaped, $\ker S$ is also star-shaped and $\ker \mathcal{P}$ is contained in the kernel of $\ker S$.
- (d) If \mathcal{P} is not star-shaped, $\ker S$ is convex (and contained in one of the ears).

Proof. Part (a) follows from the observation that, if $q \in \ker S$, points appear ordered and thus q is outside of, at most, one of the half-spaces which, in that case, corresponds to an edge of the convex hull (see Fig. 1(b)), while part (b) is obvious from (a) and the definition of $\ker \mathcal{P}$.

For (c), assume that $p \in \ker \mathcal{P}$ and $q \in \ker S$. Then $p, q \in H_i$ for some i and, because H_i is convex, $pq \in H_i$ and we have that $pq \in \ker S$.

Finally, in order to prove part (d), we observe that, if \mathcal{P} is not star-shaped,

$$\bigcap_{e_i \notin CH(\mathcal{P})} h_i \cap \mathcal{P} = \emptyset.$$

Let $q \in \ker S$ and let e_i be the edge of \mathcal{P} supporting the tangents from q to \mathcal{P} (see Fig. 1(b)). Because $\ker \mathcal{P}$ is empty, there exist j, k such that e_j and e_k are not on the convex hull of \mathcal{P} and $h_j \cap h_k \in h_i^c$. Therefore, $\ker S = H_i$ and is a convex set (the ear of e_i). □

The following theorem is the main result of this section:

Theorem 2. *Inside $\ker S$ there is, at most, one local minimum of $\mathcal{R}_S(x)$.*

Proof. Assume that $o \in \ker S$ is a local minimum of the roundness function and let $\mathbf{v} \in \mathcal{S}^1$. If we denote by $\ell_{\mathbf{v}}$ the half-line $\lambda \mathbf{v}$, $\lambda > 0$, it is enough to show that there is no other local minimum in $\ell_{\mathbf{v}} \cap \ker S$. Assume that \mathbf{v} defines the direction of the positive x axis, $CN(o) = \{p_i, p_j\}$ and $FN(o) = \{p_k, p_l\}$. If α_i denotes the angle between \mathbf{v} and \mathbf{p}_i we can further assume, without loss of generality, that $\alpha_i \leq \alpha_j$ and $\alpha_k \geq \alpha_l$. Then, from Theorem 1, it follows that $\alpha_i < \alpha_k$. Finally, for $q \in \mathbb{E}^2$,

$$\mathcal{C}(q) = \left\{ x \in \mathbb{E}^2 \mid \frac{\mathbf{v} \cdot \mathbf{x}}{\|\mathbf{x}\|} \geq \frac{\mathbf{v} \cdot \mathbf{q}}{\|\mathbf{q}\|} \right\}$$

denotes a cone with apex at o . We use $\mathcal{C}(q)^c$ to denote the complement.

We distinguish two cases:

- (i) If $\mathbf{v} \cdot \mathbf{p}_i \geq \mathbf{v} \cdot \mathbf{p}_k$, looking at circles centered at $\lambda \mathbf{v}$ and passing through p_k (see Fig. 2(a)) we have that $FN(\lambda \mathbf{v}) \subset \mathcal{C}(p_k)^c$, while looking at circles centered at $\lambda \mathbf{v}$ and passing through p_i it follows that the closest neighbors of $\lambda \mathbf{v}$ are contained in the half-plane

$$\mathcal{H} = \{q \mid \mathbf{q} \cdot \mathbf{v} > \mathbf{p}_i \cdot \mathbf{v}\}$$

and, because

$$\mathcal{H} \cap \mathcal{A}(o, \|\mathbf{p}_i\|, \|\mathbf{p}_k\|) \subset \mathcal{C}(p_k),$$

it follows that $CN(\lambda \mathbf{v}) \subset \mathcal{C}(p_k)$ and there cannot be any local minima on the half-line $\ell_{\mathbf{v}} \cap \ker S$.

- (ii) If $\mathbf{v} \cdot \mathbf{p}_i < \mathbf{v} \cdot \mathbf{p}_k$, it holds that $\mathbf{v} \cdot \mathbf{p}_k > 0$ and we claim that there exist two points $p_c \in CN(o)$ and $p_f \in FN(o)$, both in the upper or the lower half-plane and such that $\mathbf{v} \cdot \mathbf{p}_f > \mathbf{v} \cdot \mathbf{p}_c$ and $\alpha_f > \alpha_c$. In order to prove the claim assume that p_k is in the upper half-plane. If p_i is also in the upper half-plane, then $p_k = p_f$ and $p_i = p_c$




Fig. 2. Illustration for the proof of Theorem 2.

so assume that p_i is in the lower half-plane. In that case, if $\alpha_l > \alpha_i$ and p_l is in the lower half-plane, $p_l = p_f$ and $p_i = p_c$. Otherwise, p_j is between p_k and p_l and we have $p_c = p_j$ while p_f is p_k or p_l depending on the situation of p_j .

Without loss of generality, we assume that p_c and p_f are in the upper half-plane as in Fig. 2(b) and denote by q the intersection between the half-line ℓ_v and the line passing through p_c and p_f . Again, looking at circles centered at λv passing through p_c and p_f , it follows that, to the left of q , $CN(\lambda v) \subset \mathcal{C}(p_c)$ and $FN(\lambda v) \subset \mathcal{C}(p_f)^c$. Therefore, in order to finish the proof we only have to show that there are no points of $\ker S$ on the half-line ℓ_v to the right of q . However, this follows from Lemma 2: observing that in the chain between p_c and p_f there is at least one edge not belonging to the convex hull and intersecting the half-line ℓ_v to the left of q and that o is not inside an ear (in that case, the rest of the points should be inside the angular region defined by p_c , o , and p_f , contradicting the fact that the other furthest neighbor is outside that region). \square

If S is a set of points in convex position (i.e., they are the vertices of a convex polygon), because $\text{conv } S \subset \ker S$, we have:

Corollary 1. *If S is in convex position, $\mathcal{R}_S(x)$ has at most one local minimum inside $\text{conv } S$.*

It is worth noticing that Theorem 2 easily gives useful properties about the distribution of local minima in practical examples. For instance, if we assume that we have taken a sample of n points which are contained in an annulus of width ε centered at the origin (we normalize the problem in such a way that the inner radius is one) and denote by θ the smallest angle between two consecutive points, a straightforward computation shows that the disk centered at the origin with radius

$$\delta = \frac{(1 + \varepsilon) \sin \theta}{\sqrt{1 + (1 + \varepsilon)^2 - 2(1 + \varepsilon) \cos \theta}}$$

is contained in $\ker S$ and thus has at most one local minimum. If points are equally distributed and we assume a (huge) 5% error both in the angle and the width, for $n = 360$ we have $\delta \approx 0.33$. This result can explain why standard optimization methods perform well “in practice” to approximate the solution.

In the next section we see that an additional structure of the problem will allow us to design a simple and efficient *discrete* algorithm which can find the *exact* solution to the problem.

5. The Roundness Algorithm

Before we describe how to compute roundness when the circular order of the points is known, we need some additional notation and definitions. Consider the edges of the furthest-point Voronoi diagram that lie in the kernel of S , namely, $\mathcal{T} = \ker S \cap \mathcal{V}_f(S)$. Note that \mathcal{T} can be disconnected, but has linear size because the intersection between an




Fig. 3. (a) Illustration for the proof of Lemma 3. (b) Illustration for the proof of Lemma 4.

edge of $\mathcal{V}_f(S)$ and $\ker S$ is connected. We denote by e_{ij} the edge of \mathcal{T} contained in the bisector of segment $p_i p_j$. The endpoints of e_{ij} are denoted by c_{ij} (the closest end to p_i and p_j) and f_{ij} (the furthest end). This is well defined except for, at most, one edge of \mathcal{T} , which can be treated separately. Finally, we assume that points are labeled in such a way that triangle $c_{ij} p_i p_j$ has positive orientation and we denote by S_{ij} the chain between p_i and p_j (see Fig. 3(a)). Then we have:

Lemma 3. *The roundness function has a local minimum in e_{ij} if, and only if, $CN(c_{ij}) \subset S_{ij}$ and $CN(f_{ij}) \subset S_{ji}$.*

Proof. Using the same arguments as in the proof of Theorem 2, it follows that if the closest neighbor of c_{ij} is not in S_{ij} , then the closest neighbors of points of e_{ij} cannot be in S_{ij} and, therefore, there is no local minimum in this edge. Moreover, if the closest neighbors of c_{ij} and f_{ij} are in S_{ij} , then the closest neighbors of all the points in e_{ij} are in S_{ij} (if this was not true, we would get a contradiction with the preceding argument). Finally, if $CN(c_{ij}) \subset S_{ij}$ and $CN(f_{ij}) \subset S_{ji}$, then the local minimum is achieved when the closest neighbors switch between S_{ij} and S_{ji} . \square

Using this result, the algorithm is straightforward: as a first step, we compute Voronoi diagrams $\mathcal{V}_f(S)$ and $\mathcal{V}_c(S)$ and process $\mathcal{V}_c(S)$ for point location. This can be done in $O(n \log n)$ time [12], [7]. Because the number of intersections between $\mathcal{V}_f(S)$ and $\ker S$ is linear, \mathcal{T} can also be computed in time $O(n \log n)$ with a standard sweepline algorithm. We can process each edge of \mathcal{T} in time $O(\log n)$ by locating the closest neighbors of its endpoints. Finally, if we get an edge $e \in \mathcal{T}$ containing a local minimum, we can locate it traveling along e across $\mathcal{V}_c(S)$ in additional $O(n)$ time. Then we have shown:

Theorem 3. *If there is a local minimum of $\mathcal{R}_S(x)$ inside $\ker S$, it can be computed in $O(n \log n)$ time.*

It can be useful to point out that, if \mathcal{T} is connected (which is likely to be the case in most practical situations), the algorithm can be made even simpler within the same asymptotic complexity by avoiding the point location step. Furthermore, the complexity can be improved to optimal linear time if points are in convex position. Both results are based on the following lemma.

Lemma 4. *Let $q \in \ker S$ be a vertex of $\mathcal{V}_f(S)$ with $FN(q) = \{p_i, p_j, p_k\}$ where points are ordered according to their indices. Assume that \mathcal{T} is connected and let \mathcal{T}_{ij} denote the subtree corresponding to points in S_{ij} (see Fig. 3(b)). If the closest neighbor of q is not in S_{ij} , there is no local minimum of $\mathcal{R}_S(x)$ in \mathcal{T}_{ij} .*

Proof. If the angle $p_i q p_j$ is smaller than π , the result follows by repeatedly applying Lemma 3 to the edges of \mathcal{T}_{ij} starting at q . If this is not the case, as in Fig. 3(b) for \mathcal{T}_{ki} , we can use again the same arguments as in the proof of Theorem 2 to see that if the closest neighbor of q is not in S_{ki} , then the closest neighbor of points of e_{ki} cannot be in S_{ki} and iterate this argument to show that there is no local minimum in \mathcal{T}_{ki} . \square

Using this lemma, we compute \mathcal{T} as before and locate a vertex such that the subtrees that it connects have size smaller than cn for some constant c . Now, we compute the closest neighbor of the vertex by brute force and discard a linear fraction of the edges in linear time. Therefore, the edge containing the local minimum is located in $O(n \log n)$ time. We conclude the section observing that, if points of S are in convex position, both approaches can be combined in order to get:

Theorem 4. *If S is a set of points in convex position and $\mathcal{R}_S(x)$ has a local minimum inside $\text{conv } S$, it can be computed in time $\Theta(n)$.*

Proof. If points are in convex position, $\mathcal{V}_c(S)$ and $\mathcal{V}_f(S)$ can be computed and the former processed for point location in $O(n)$ time [3], [7]. Moreover, $\mathcal{T} = \mathcal{V}_f(S) \cap \text{conv } S$ is connected and can be computed in linear time. To see this, consider \mathcal{T} as a directed tree rooted at the center of the minimum enclosing circle of S and observe that an edge $e \in \mathcal{T}$ cannot enter $\text{conv } S$. Finally, applying Lemma 4 we can discard a linear fraction of the edges of \mathcal{T} in time $O(\max\{\log n, |\mathcal{T}_i|\})$ where \mathcal{T}_i is the tree after i prune steps and the result follows. \square

6. An Example with $\Omega(n^2)$ Local Minima

Previous works [2], [16] conjectured that $\mathcal{R}_S(x)$ may have $\Omega(n^2)$ local minima because $\mathcal{V}_c(S)$ and $\mathcal{V}_f(S)$ can have a quadratic number of intersections. We construct $n + 1$ closest and $n + 1$ furthest points whose respective Voronoi diagrams have n^2 intersections in the first quadrant, all of which will be local minima by Corollary 2.




Fig. 4. Closest- and furthest-point Voronoi edges in the first quadrant.

Corollary 2. *If edges of $\mathcal{V}_c(S)$ and $\mathcal{V}_f(S)$ intersect and the closest and furthest points interlace anglewise as seen from the intersection point, then $\mathcal{R}_S(x)$ has a local minimum at that point.*

First, we choose the set of $n + 1$ “closest points” along the y axis with y coordinates $y_i = (2c + 1)^{i-1}$, for $1 \leq i \leq n + 1$, where c is a positive constant to be chosen later. The bisectors of these points are the lines

$$y = \frac{y_i + y_{i+1}}{2} = (c + 1)y_i.$$

In the first quadrant, these bisectors will be edges of the closest-point Voronoi diagram. See Fig. 4.

We choose the $n + 1$ “furthest points” in the third quadrant of a circle centered at the origin with radius $r > y_n$ to be determined later. Thus, the furthest-point Voronoi edges, $\{\ell_j : 1 \leq j \leq n\}$, begin at the origin and extend to infinity in the first quadrant.

We choose “furthest points” so that the angle from the x axis to the first edge ℓ_1 is $\theta_1 = \pi/8$ and the angle from line ℓ_{j-1} to line ℓ_j is $\theta_j = \pi/2^{j+2}$. Then the angle from the x axis to ℓ_j is $\pi/4 - \theta_j$ for $1 \leq j \leq n$. Specifically, choose the j th “furthest point,” f_j , to lie on the line through the origin at angle $\pi/4 - 4\theta_j/3$, for $1 \leq j \leq n + 1$. Then the bisector of the points f_j and f_{j+1} has angle

$$\frac{1}{2} \left(\frac{\pi}{4} - \frac{4\theta_j}{3} + \frac{\pi}{4} - \frac{4\theta_{j+1}}{3} \right) = \frac{\pi}{4} - \frac{1}{2} \left(\frac{4\theta_j}{3} + \frac{2\theta_j}{3} \right) = \frac{\pi}{4} - \theta_j,$$

and the furthest-point Voronoi edge ℓ_j has the desired angle.

Theorem 5. *If the $n + 1$ “closest points” and $n + 1$ “furthest points” are defined as above, then for some constant c and radius r , the function $\mathcal{R}_S(x)$ has $\Theta(n^2)$ local minima.*

Proof. First, some additional notation. We denote the furthest-point Voronoi edge ℓ_j by its slope/intercept form $y = m_jx$, where the slope $m_j = \tan(\pi/4 - \theta_j)$. We can bound $\frac{1}{2} \leq m_j < 1$, for all j . We denote the slope of the line that defines the furthest point f_j by m'_j . We can express m'_j in terms of m_j :

$$m'_j = \tan\left(\frac{\pi}{4} - \frac{4\theta_j}{3}\right) = \frac{\tan(\pi/4 - \theta_j) - \tan(\theta_j/3)}{1 + \tan(\pi/4 - \theta_j)\tan(\theta_j/3)} = \frac{m_j - \tan(\theta_j/3)}{1 + m_j \tan(\theta_j/3)}.$$

Now, the i th closest-point edge intersects the j th furthest-point edge in the point $q_{ij} = ((c + 1)y_i/m_j, (c + 1)y_i)$. By Corollary 2, the point q_{ij} is a local minimum if the points $(0, y_{i+1}), f_j, (0, y_i)$, and f_{j+1} appear in counterclockwise order viewed from q_{ij} . It is clear that $(0, y_{i+1})$ is first; we need to check that f_j and f_{j+1} are on opposite sides of the line through q_{ij} and $(0, y_i)$. The equation of this line is

$$y = \left(\frac{(c + 1)y_i - y_i}{(c + 1)y_i}\right)m_jx + y_i = \left(\frac{c}{c + 1}\right)m_jx + y_i.$$

Notice that this line is above ℓ_j in the third quadrant, so f_{j+1} is below the line. Notice also that the slope of this line is determined only by the furthest-point Voronoi edge ℓ_j —all lines defined by intersection with ℓ_j are parallel. If we can choose f_j to lie above the highest line, $y = cm_jx/(c + 1) + y_n$, then the points q_{ij} will be local minima for all i .

For the furthest point f_j to be above the highest line, the slope m'_j of the line defining f_j must be less than the slope $cm_j/(c + 1)$ of the highest line. If we choose $c > 2/\tan(\theta_n/3)$, then

$$\frac{c}{c + 1}m_j > \frac{m_j}{1 + \tan(\theta_j/3)/2} \geq \frac{m_j - \tan(\theta_j/3)}{1 + m_j \tan(\theta_j/3)} = m'_j.$$

Thus the line defining f_j crosses from below to above the highest line at some point. As long as we choose r , the radius of the circle that defines f_j , larger than the distance to the intersection of these two lines, then f_j will be above the highest line. Thus, it is possible to construct the furthest points f_j so that the n^2 intersection points q_{ij} determined by a closest- and furthest-point Voronoi edge in the first quadrant are local minima. \square

To get an idea of the magnitude of the coordinates involved, choose $c = 2^{n+5}$. The circle must include all intersections of the lines $y = m'_jx$ with $y = cm_jx/(c + 1) + y_n$ for all $1 \leq j \leq n$. The largest y coordinate is at most the largest x coordinate, which the reader can check is

$$\frac{y_n}{m'_j - cm_j/(c + 1)} < \frac{(2c + 1)^{n-1}}{m_j/(1 + \tan(\theta_j/3)/2) - m_j/(1 + 1/c)} < 2^{(n+2)(n+5)}.$$

Therefore the radius can be chosen to be $r = 2^{\Theta(n^2)}$. Since the positions of points in this construction have been chosen to simplify the proof, we are certain that smaller

constructions of $\Omega(n^2)$ local minima can be given. We conjecture, however, that the magnitudes of points must grow exponentially in n .

We note that in this construction all of the points are on the boundary of the convex hull, and that a small perturbation could make all points be vertices of the convex hull. The $\Omega(n^2)$ minimal annuli all have centers outside the convex hull; we have seen that at most one center can be inside the hull.

7. Concluding Remarks

In this paper we have dealt with the problem of computing the annulus of minimum width containing a set of points. We have characterized the set of local minima of $\mathcal{R}_S(x)$ (feasible centers for the annulus) in the d -dimensional case and we have shown that, inside the *kernel* of the set, there is at most one local minimum which can be computed in $O(n \log n)$. We remark that the algorithm is more efficient than previous algorithms solving the general problem and that it is useful for solving the problem in practice, because in the metrology application the sample points are given in angular order. We also remark that the algorithm is simple and, thus, easy to implement. It is an interesting open problem to see whether or not a set of points in convex position in the space has only one local minimum inside its convex hull.

Acknowledgments

The authors thank Prof. G. Toussaint and C. K. Yap for bringing the problem to our attention during the Bellairs Workshop on Computational Geometry '95, colleagues M. Abellanas, G. Blanco, G. Hernández, F. Gómez, and O. Schwartzkopf for helpful discussions, P. Bose for letting us know about reference [13], and Prof. G. Rote for pointing out some minor gaps in a previous version of this work.

References

1. P. Agarwal and M. Sharir. Efficient randomized algorithms for some geometric optimization problems. *Discrete Comput. Geom.* **16** (1996), 317–337.
2. P. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *J. Algorithms* **17** (1994), 292–318.
3. A. Aggarwal, L. J. Guibas, J. Saxe, and P. W. Shor. A linear time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete Comput. Geom.* **4** (1989), 591–604.
4. M. de Berg, P. Bose, D. Bremner, S. Ramaswami, and G. Wilfong. Computing constrained minimum width annuli of point sets, *Proc. WADS 97, 1997*, pp. 3–16.
5. C. A. Duncan, M. T. Goodrich, and E. A. Ramos. Efficient approximation and optimization algorithms for computational metrology. *Proc. 8th ACM–SIAM Symp. on Discrete Algorithms*, 1997, pp. 121–130.
6. H. Ebara, N. Fukuyama, H. Nakano, and Y. Nakanishi. Roundness algorithms using Voronoi diagrams. *Abstracts 1st Canad. Conf. on Computational Geometry*, 1989, p. 41.
7. H. Edelsbrunner, L. J. Guibas, and J. Stolfi. Optimal point location in a monotone subdivision, *SIAM J. Comput.* **15** (1986), 317–340.
8. L. W. Foster. *GEO-METRICS II: The Application of Geometric Tolerancing Techniques*. Addison-Wesley, Reading, MA, 1982.

9. G. Henzold. *Handbook of Geometrical Tolerancing*, J Wiley, New York, 1995.
10. V. B. Le and D. T. Lee. Out-of-roundness problem revisited. *IEEE Trans. Pattern Anal. Mach. Intell.*, **13** (1991), 217–223.
11. K. Mehlhorn, T. Shermer, and C. Yap. A complete roundness classification procedure, *Proc. 13th ACM Ann. Symp. on Computational Geometry*, 1997, pp. 129–138.
12. F. P. Preparata and M. I. Shamos. *Computational Geometry*. Springer-Verlag, New York, 1985.
13. T. J. Rivlin. Approximation by circles, *Computing* **21** (1979), 93–104.
14. U. Roy and X. Zhang. Establishment of a pair of concentric circles with the minimum radial separation for assessing rounding error. *Computer-Aided Design* **24**(3) (1992), 161–168.
15. T. Shermer and C. K. Yap. Probing for near-centers and approximate relative roundness, June 21–23, 1995. (Invited talk) ASME Workshop on Tolerancing and Metrology, University of North Carolina at Charlotte.
16. M. Smid and R. Janardan. On the width and roundness of a set of points in the plane, *Proc. 7th. Canad. Conf. on Computational Geometry*, 1995, pp. 193–198.
17. K. Swanson, D. T. Lee, and V. L. Wu. An optimal algorithm for roundness determination of convex polygons. *Comput. Geom. Theory Appl.* **5** (1995), 225–235.
18. A. D. Wainstein. A non-monotonous placement problem in the plane. Software systems for solving optimal planning problems, *Abstracts, 9th All-Union Symp.*, Minsk, BSSR, 1986, pp. 70–71.
19. C. K. Yap. Exact computational geometry and tolerancing metrology. In *Snapshots on Computational Geometry*, Vol. 3, D. Avis and P. Bose (eds.). McGill School of Computer Science, 1994, pp. 34–48.

Received June 25, 1997, and in revised form March 5, 1998.