



On Geodesic Disks Enclosing Many Points*

Prosenjit Bose 


School of Computer Science, Carleton University, Ottawa, Canada

Guillermo Esteban 


Departamento de Física y Matemáticas, Universidad de Alcalá, Alcalá de Henares, Spain

David Orden 

Departamento de Física y Matemáticas, Universidad de Alcalá, Alcalá de Henares, Spain

Rodrigo I. Silveira 

Departament de Matemàtiques, Universitat Politècnica de Catalunya, Barcelona, Spain

Tyler Tuttle 

School of Computer Science, Carleton University, Ottawa, Canada

Abstract

Let $\Pi(n)$ be the largest number such that for every set S of n points in a polygon P , there always exist two points $x, y \in S$, where every geodesic disk containing x and y contains $\Pi(n)$ points of S . We establish upper and lower bounds for $\Pi(n)$, and show that $\lceil \frac{n}{5} \rceil + 1 \leq \Pi(n) \leq \lceil \frac{n}{4} \rceil + 1$. We also show that there always exist two points $x, y \in S$ such that every geodesic disk with x and y on its boundary contains at least $\frac{16}{665}(n-2) \approx \lceil \frac{n-2}{41.6} \rceil$ points both inside and outside the disk. For the special case where the points of S are restricted to be the vertices of a geodesically convex polygon we give a tight bound of $\lceil \frac{n}{3} \rceil + 1$. We provide the same tight bound when we only consider geodesic disks having x and y as diametral endpoints. Finally, we give a lower bound of $\lceil \frac{n-2}{36} \rceil + 2$ for the two-colored version of the problem.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Enclosing disks, Geodesic disks, Bichromatic

Digital Object Identifier 10.4230/LIPIcs.WADS.2025.10

Related Version *Full Version:* <https://doi.org/10.48550/arXiv.2506.06477> [11]

Funding G. E., D. O. and R. S. are partially supported by grants PID2019-104129GB-I00 and PID2023-150725NB-I00 funded by MICIU/AEI/10.13039/501100011033. P. B. and T. T. are supported in part by NSERC.

1 Introduction

Given a set S of n points in the plane in general position – no three of them are collinear and no four of them are cocircular – there always exist two points $u, v \in S$ such that any disk that contains u and v also contains a constant fraction of S . Neumann-Lara and Urrutia [19] proved that this constant fraction is at least $\lceil \frac{n-2}{60} \rceil$. This bound was improved in a series of papers [14, 15] culminating in the best known lower bound of about $\frac{n}{4.7}$ by Edelsbrunner et al. [13]. Hayward et al. [15] gave an upper bound by constructing a set of n points such that for every pair of points u, v , there exists a disk with u, v on its boundary containing less than $\lceil \frac{n}{4} \rceil$ points, thereby showing that $\lceil \frac{n}{4} \rceil + 1$ is an upper bound for the problem. In addition, they studied the problem for point sets in convex position, giving a tight bound of $\lceil \frac{n}{3} \rceil + 1$.

* Since the acceptance of this paper by the WADS Program Committee, the authors have improved some bounds that can be found in the full version [11].



10:2 On Geodesic Disks Enclosing Many Points

■ **Table 1** Summary of our results, comparing the Euclidean and geodesic settings. Note that non-integer values in the table are approximations, and refer to the corresponding reference for the exact value.

Variant (see Def. 1)	Euclidean setting		Geodesic setting	
	Lower bound	Upper bound	Lower bound	Upper bound
$\Pi(n)$	$\lfloor \frac{n}{4.7} \rfloor$ [13]	$\lceil \frac{n}{4} \rceil + 1$ [15]	$\lceil \frac{n}{5} \rceil + 1$ (Thm. 10)	$\lceil \frac{n}{4} \rceil + 1$
$\bar{\Pi}(n)$	$\lceil \frac{n}{3} \rceil + 1$ [15]	$\lceil \frac{n}{3} \rceil + 1$ [15]	$\lceil \frac{n}{3} \rceil + 1$ (Thm. 14)	$\lceil \frac{n}{3} \rceil + 1$
$\Pi^{in-out}(n)$	$\lfloor \frac{n}{4.7} \rfloor$ [23]	$\lceil \frac{n}{4} \rceil + 1$ [15]	$\lceil \frac{n-2}{41.6} \rceil + 1$ (Thm. 12)	$\lceil \frac{n}{4} \rceil + 1$
$\Pi^{diam}(n)$	$\lceil \frac{n}{3} \rceil + 1$ [2]	$\lceil \frac{n}{3} \rceil + 1$ [2]	$\lceil \frac{n}{3} \rceil + 1$ (Thm. 16)	$\lceil \frac{n}{3} \rceil + 1$
$\Pi^{bichrom}(n)$	$\lfloor \frac{n}{6.8} \rfloor$ [12]	$\lceil \frac{n}{4} \rceil + 1$ [12]	$\lceil \frac{n-2}{36} \rceil + 2$ (Thm. 18)	$\lceil \frac{n}{4} \rceil + 1$

Another version of the problem is to consider disks in the plane having the points u and v as diametral endpoints [2]. In this case, a tight bound of $\lceil \frac{n}{3} \rceil + 1$ for both the convex and non-convex cases was shown.

The best known lower bound of $\left(\frac{1}{2} - \frac{1}{\sqrt{12}}\right)n + O(1) \approx \frac{n}{4.7}$ was obtained in multiple ways, see [12, 13, 23]. Edelsbrunner et al. [13] were the first to show this using techniques related to the order- k Voronoi diagram. Ramos and Viaña [23] used known results about j -facets of point sets in \mathbb{R}^3 . In fact, these results allowed Ramos and Viaña to prove that there is always a pair of points such that any disk through them has at least $\frac{n}{4.7}$ points both inside and outside the disk.

Generalizations of the original problem to higher dimensions have been studied by Bárány et al. [6] who proved that any set $S \subset \mathbb{R}^d$ of n points in general position contains a subset A of size $m = \lfloor \frac{1}{2}(d+3) \rfloor$ such that, any ball with A on its boundary, contains at least $\frac{m!(d-m-1)!}{d!}n$ other points of S . Moreover, Bárány and Larman extended the results in [19] from Euclidean balls to ellipses in \mathbb{R}^2 and, more generally, to quadrics in \mathbb{R}^d [5]. In particular, they showed that any set $S \subset \mathbb{R}^d$ of n points contains a subset A , of size $\lfloor \frac{1}{4}d(d+3) \rfloor + 1$, with the property that any quadric through A , contains at least $\frac{n}{s2^{s+1}}$ points of S , where $s = \frac{(d+1)(d+2)}{2}$.

A colored version of the problem has also been studied in the literature: Given a set S of $\frac{n}{2}$ red and $\frac{n}{2}$ blue points in the plane, there always exists a bichromatic pair of points $u, v \in S$ such that any disk containing u and v contains at least a constant fraction of the n points. Prodromou [22] proved¹ that any set $S \subset \mathbb{R}^d$ colored with $k = \lfloor \frac{d+3}{2} \rfloor$ colors contains a subset $A \subset S$ of k points, one of each color, such that any ball through all points in A contains at least $\frac{|S|}{2k3^k}$ points of S , which gives a fraction $\frac{n}{36}$ for points in the plane. This fraction for the two-dimensional case was later improved in [12] to $\left(\frac{1}{2} - \frac{1}{\sqrt{8}}\right)n - o(n) \approx \frac{n}{6.8}$.

We note that all the upper bounds in the Euclidean setting translate directly to the geodesic setting by enclosing the constructions in a large enough triangle. Thus, in this paper, we focus mainly on the lower bounds and address all variants of the problem in the case where S is a set of n points inside a simple polygon P . In addition, throughout the paper, we consider that $n \geq 21$, so that all our results hold.

We summarize all our results and compare them with the Euclidean setting in Table 1. In Section 3, we present an overview of key properties of shortest paths in P (which we refer to as the *geodesic setting*), emphasizing both their similarities to and differences from

¹ Using techniques from Edelsbrunner et al. [13], a PhD thesis written in Spanish and not published elsewhere had claimed before a lower bound of $\frac{23n}{250} \approx \frac{n}{11}$ [25].

the Euclidean case. It is these differences that pose challenges we address to extend the results to the geodesic setting. Using these properties, in Section 4, we extend the lower bounds from the Euclidean case [14, 19] to the geodesic setting. While the upper bounds carry over directly, the lower bounds require new arguments due to the differences outlined in Section 3. For instance, unlike the Euclidean metric, which has bounded doubling dimension, the doubling dimension of the geodesic metric is not necessarily bounded since it depends on the number of reflex vertices of P . Additionally, we establish a tight bound for the case where the point set is in geodesically convex position. In Section 5, we show that the techniques in [2], where only diametral disks are considered, can be generalized to the geodesic setting. Finally, in Section 6 we extend the results in [22, 25] for two-colored points.


The problem of studying the largest number of points contained in any disk through two points inside a polygon is distinct from the planar case. As illustrated in Figure 1, a geodesic disk of radius r in a polygon P can be strictly contained in a Euclidean disk of radius r . As such, in this setting, it is unclear whether there always exists a pair of points such that any geodesic disk through them can contain the same number of points as in the Euclidean setting.

2 Preliminaries and notation

A sequence of points in the plane, p_1, \dots, p_n , forms a *polygonal chain* where the edges of the chain are the segments $p_i p_{i+1}$, for $i \in \{1, \dots, n-1\}$. A polygonal chain is *simple* if consecutive edges intersect at their common point and non-consecutive segments do not intersect. In a graph theoretic sense, this is a path. A simple polygon is a simple closed polygonal chain, i.e. a cycle. Informally, a polygon P is a *weakly-simple* polygon provided that a slight perturbation of the vertices of the polygon results in a simple polygon. See Akitaya et al. [1] for a formal definition of weakly-simple as well as an algorithm to quickly recognize weakly-simple polygons. Given a pair of points u, v in a polygon P , the *geodesic* (or *shortest*) *path* from u to v in P is denoted by $g(u, v)$. It is well-known that $g(u, v)$ is a polygonal chain. The length of this path is the sum of the lengths of its edges and is denoted by $|g(u, v)|$.

A subset P' of P is called *geodesically convex* if, for all pairs of points $u, v \in P'$, the geodesic path $g(u, v)$ is contained in P' . The *geodesic convex hull* of a set S of points in P is the geodesically convex set of minimum area that contains S . We say that S is in *geodesically convex position* if every point of S is on the boundary of the geodesic convex hull of S . A *geodesic triangle* on three points $a, b, c \in P$, denoted $\Delta(a, b, c)$, is a weakly-simple polygon whose boundary consists of $g(a, b)$, $g(b, c)$, and $g(c, a)$. The geodesic paths $g(a, b)$ and $g(a, c)$ follow a common route from a until they diverge at a point a' (which may be a itself). Similarly, let b' be the point where $g(b, a)$ and $g(b, c)$ diverge, and c' be the point where $g(c, a)$ and $g(c, b)$ diverge. The geodesic triangle $\Delta(a', b', c')$ is a simple polygon that has a', b' and c' as its convex vertices. The geodesic triangle $\Delta(a', b', c')$ is referred to as the *geodesic core* of $\Delta(a, b, c)$ with the vertices a', b', c' referred to as the *core vertices* corresponding to the vertices a, b, c , respectively. The geodesic core of $\Delta(a, b, c)$ is denoted as $\nabla(a, b, c)$ [9, 21]. In addition, $\angle abc$ denotes the convex angle formed by the two edges of the geodesic core adjacent to the core vertex b' in $\nabla(a, b, c)$. See Figure 2.

A *geodesic disk* centered at $c \in P$ with radius $r \geq 0$ is the set $D(c, r) = \{x \in P : |g(c, x)| \leq r\}$. In the plane, there always exists a disk through three given points, assuming they are not collinear. When we refer to a disk *through* a set of points, we mean that the points are on the boundary of the disk. In a simple polygon, it is no longer true that there always exists a geodesic disk through three given points. This essentially means that given



■ **Figure 1** Geodesic disk centered at c , with radius $|g(c, z)|$, shown solid, and the equivalent Euclidean disk, shown dashed.

three points in a simple polygon, there does not always exist a point in the polygon that is geodesically equidistant to all three. However, if there is a geodesic disk through three points in a polygon, then it is unique [4]. Given two points u and v , there always exists a geodesic disk through u and v centered at c_{uv} , the midpoint of $g(u, v)$, called the *diametral geodesic disk*.

A set S' of at least three points in P is *geodesically collinear* if $\exists x, y \in S'$ such that $S' \subset g(x, y)$, see set $\{x, c, y\}$ in Figure 1. Similarly, a set S'' of at least four points is *geodesically cocircular* if there exists a geodesic disk with S'' on its boundary, see set $\{w, x, y, z\}$ in Figure 1. A set S of n points in a polygon P is said to be in *general position* if no three points in S are geodesically collinear and no four points in S or the boundary of P are geodesically cocircular. In this paper we only consider sets S in general position in polygons P with the extra condition that there are no two points $x, y \in S$ and a reflex vertex p on the boundary of P such that $y \in g(x, p)$.

Given two points u and v , the *extension path* $\ell(u, v)$ is defined as follows: extend the first and last segment of $g(u, v)$ until they properly intersect with P . Denote these extension points on the boundary of P as \hat{u} and \hat{v} , respectively. See the red path in Figure 2. Note that if u (or v) is on the boundary of P then $\hat{u} = u$ (or $\hat{v} = v$). Also note that, because of the extra condition added to the general position, the extensions cannot touch a reflex vertex p on the boundary of the polygon P . The extension path $\ell(u, v)$ naturally partitions the polygon into two weakly-simple polygons. We say that a point w is to the left (resp., right) of $g(u, v)$ if w belongs to the weakly-simple polygon to the left (resp., right) of $\ell(u, v)$, when considering $\ell(u, v)$ oriented from u to v . In addition, we define the left (resp., right) half-polygon of the path $g(u, v)$ as the set of points that are to the left (resp., right) of its extension path $\ell(u, v)$ and denote it as $P^+(u, v)$ (resp., $P^-(u, v)$).

Given two points $u, v \in S$, the bisector $b(u, v)$ is the set of all points that are equidistant to u and v . In this paper, as in [3], we assume that no bisector $b(u, v)$ passes through a vertex of P . Hence, $b(u, v)$ is a continuous curve connecting two points on the boundary of P . Moreover, $b(u, v)$ can be decomposed into $O(|P|)$ pieces [3], each of which is a subarc of a hyperbola (that could degenerate to a line segment).

Proof. Let a be an arbitrary point in $D \cap P^+(p, q)$. By definition of $P^+(p, q)$, we have that a is to the left of $\ell(p, q)$. We consider two cases: $a \in \Delta(c', p, q)$ and $a \notin \Delta(c', p, q)$.

If a is in $\Delta(c', p, q)$, see Figure 2, then we use the fact that, for every point $x \in g(p, q)$, the length of $g(c', x)$ as x moves from p to q is a convex function with the maximum occurring at p or q [21]. In particular, we have that $|g(c', x)| \leq |g(c', q)| = |g(c', p)|$, for $x \in g(p, q)$. If we move x from p to q along $g(p, q)$, at some point $g(c', x)$ will go through the point a , so $|g(c', a)| \leq |g(c', x)|$. Hence, $|g(c', a)| \leq |g(c', q)| = |g(c', p)|$, which means that $a \in D'$.

Now consider the case where the point a is outside $\Delta(c', p, q)$, like a' in Figure 2. Without loss of generality, assume a is to the right of $g(p, c')$. (Otherwise, we can consider a and q , and the proof is analogous). In this case, $g(c', p)$ must intersect $g(c, a)$. Let x be a point in this intersection. By definition, we have $|g(c', x)| + |g(x, p)| = |g(c', p)|$. Since $a \in D$, we have that $|g(c, a)| = |g(c, x)| + |g(x, a)| \leq |g(c, p)|$ and by the triangle inequality, we have $|g(c, p)| \leq |g(c, x)| + |g(x, p)|$. This means that $|g(x, a)| \leq |g(x, p)|$. Therefore, $|g(c', a)| \leq |g(c', x)| + |g(x, a)| \leq |g(c', x)| + |g(x, p)| = |g(c', p)|$, implying $a \in D'$ in this case. \blacktriangleleft

With this property of geodesic disks in hand, we are now ready to prove the corresponding property of intersecting segments in the geodesic setting.

► **Lemma 3.** *Let a, b, p, q be four distinct points in general position in P such that $g(p, q)$ and $g(a, b)$ intersect. Then every geodesic disk with p and q on its boundary contains at least one endpoint of $g(a, b)$, or every geodesic disk with a, b on its boundary contains at least one endpoint of $g(p, q)$.*



Proof. First, suppose that the bisectors $b(a, b)$ and $b(p, q)$ intersect, see Figure 3a, and let x be a point in their intersection. Then we have that $|g(x, a)| = |g(x, b)|$ and $|g(x, p)| = |g(x, q)|$. We can assume, without loss of generality, that $|g(x, p)| < |g(x, a)|$ since the points in S are in general position, and a, b, p, q are not co-circular. The case where $|g(x, p)| > |g(x, a)|$ is symmetric.

When $|g(x, p)| < |g(x, a)|$, the geodesic disk through a and b with center x contains p and q . Since $g(a, b)$ and $g(p, q)$ intersect, p and q are on different sides of the extension path $\ell(a, b)$. Hence, by Lemma 2, if we move the center of this disk along the bisector $b(a, b)$ in one direction while keeping a, b on its boundary, the geodesic disk will always contain p in its interior, and if we move the center in the other direction, the geodesic disk will always contain q .

Now, suppose that the bisectors $b(a, b)$ and $b(p, q)$ do not intersect, see Figure 3b. Hence, all points of the bisector $b(a, b)$ are closer to one endpoint of $g(p, q)$ than to the other one. Assume, without loss of generality, that this point is p , i.e., $|g(x', p)| \leq |g(x', q)|$, $\forall x' \in b(a, b)$. Similarly, all points of the bisector $b(p, q)$ are closer to one endpoint of $g(a, b)$ than to the other one. Assume, without loss of generality, that this point is b .

The bisector $b(a, b)$ intersects the boundary of the polygon. Also, the extension path $\ell(a, b)$ divides the polygon P into two half-polygons, one containing p and the other containing q . Let x be the intersection of $b(a, b)$ with the polygon boundary that is contained in the same half-polygon as q . Similarly, the extension path $\ell(p, q)$ divides the polygon into two half-polygons, and let y be the intersection of $b(p, q)$ with the polygon contained in the same half-polygon as a .

By definition, we have that $|g(x, a)| = |g(x, b)|$, and $|g(x, p)| < |g(x, q)|$. If $|g(x, p)| \leq |g(x, b)|$, the geodesic disk with center at x and radius $|g(x, b)|$ contains a, b and p . By Lemma 2, this means that every geodesic disk through a, b contains p .

(a) Bisectors $b(a, b)$ and $b(p, q)$ intersect.(b) Bisectors $b(a, b)$ and $b(p, q)$ do not intersect.

■ **Figure 3** Any geodesic disk through a and b contains at least one endpoint of $g(p, q)$.

The only remaining case to consider is if $|g(x, p)| > |g(x, b)|$. Since the bisectors do not intersect, we have that $g(x, b)$ intersects $g(y, p)$. Let z be a point in this intersection. By definition, we have $|g(x, b)| = |g(x, z)| + |g(z, b)|$ and $|g(y, p)| = |g(y, z)| + |g(z, p)|$. By the triangle inequality, we have $|g(x, p)| \leq |g(x, z)| + |g(z, p)|$, which with our assumption that $|g(x, p)| > |g(x, b)|$ implies that $|g(z, b)| < |g(z, p)|$. However, this means that $|g(y, b)| \leq |g(y, z)| + |g(z, b)| < |g(y, z)| + |g(z, p)| = |g(y, p)|$. Therefore, when $|g(x, p)| > |g(x, b)|$, we have $|g(y, b)| < |g(y, p)|$. By Lemma 2, this means that every geodesic disk through p, q contains b . ◀

We now prove a complementary lemma about points that always remain outside disks through endpoints of geodesic paths that intersect. The proof is similar to that of Lemma 3, and we include it here for completeness.


► **Lemma 4.** *Let a, b, p, q be four distinct points in general position in P such that $g(p, q)$ and $g(a, b)$ intersect. Then every geodesic disk with p and q on its boundary does not contain at least one endpoint of $g(a, b)$, or every geodesic disk with a, b on its boundary does not contain at least one endpoint of $g(p, q)$.*

Proof. Suppose that the bisectors $b(a, b)$ and $b(p, q)$ intersect, see Figure 3a, and let x be a point in their intersection. Then we have that $|g(x, a)| = |g(x, b)|$ and $|g(x, p)| = |g(x, q)|$. We can assume, without loss of generality, that $|g(x, a)| < |g(x, p)|$ since the points in S are in general position, and a, b, p, q are not co-circular. The case where $|g(x, a)| > |g(x, p)|$ is symmetric.

When $|g(x, a)| < |g(x, p)|$, the geodesic disk through a and b with center x contains neither p nor q . Since $g(a, b)$ and $g(p, q)$ intersect, p and q are on different sides of the extension path $\ell(a, b)$. Note that if we exchange the roles of p by q , and D by D' in Lemma 2, we have that $D' \cap P^-(p, q) \subset D$. Hence, by Lemma 2, if we move the center of this disk along the bisector $b(a, b)$ in one direction while keeping a, b on its boundary, the geodesic disk will never contain p , and if we move the center in the other direction, the geodesic disk will never contain q .

Now, suppose that the bisectors $b(a, b)$ and $b(p, q)$ do not intersect, see Figure 3b. Hence, all points of the bisector $b(a, b)$ are closer to one endpoint of $g(p, q)$ than to the other one. Assume, without loss of generality, that this point is p , i.e., $|g(x', p)| \leq |g(x', q)|$, $\forall x' \in b(a, b)$. Similarly, all points of the bisector $b(p, q)$ are closer to one endpoint of $g(a, b)$ than to the other one. Assume, without loss of generality, that this point is b .

The bisector $b(a, b)$ intersects the boundary of the polygon. Also, the extension path $\ell(a, b)$ divides the polygon P into two half-polygons, one containing p and the other containing q . Let x be the intersection of $b(a, b)$ with the polygon boundary that is contained in the



■ **Figure 4** The angle $\angle uvw$ is greater than $\frac{\pi}{3}$, and $|vq| < |vp|$.

same half-polygon as q . Similarly, the extension path $\ell(p, q)$ divides the polygon into two half-polygons, and let y be the intersection of $b(p, q)$ with the polygon contained in the same half-polygon as a .

By definition, we have that $|g(x, a)| = |g(x, b)|$, and $|g(x, p)| < |g(x, q)|$. If $|g(x, b)| < |g(x, q)|$, the geodesic disk with center at x through a, b does not contain q . By Lemma 2, this means that every geodesic disk through a, b does not contain q .

The only remaining case to consider is when $|g(x, b)| \geq |g(x, q)|$. By the fact that a, b, p, q are not co-circular, we have that in fact $|g(x, b)| > |g(x, q)|$. Since the bisectors do not intersect, we have that $g(x, q)$ intersects $g(y, b)$. Let z be a point in this intersection. By definition, we have $|g(x, q)| = |g(x, z)| + |g(z, q)|$ and $|g(y, b)| = |g(y, z)| + |g(z, b)|$. By the triangle inequality, we have $|g(x, b)| \leq |g(x, z)| + |g(z, b)|$, which with our assumption that $|g(x, b)| > |g(x, q)|$ implies that $|g(z, b)| > |g(z, p)|$. However, this means that $|g(y, q)| < |g(y, z)| + |g(z, q)| \leq |g(y, z)| + |g(z, b)| = |g(y, b)|$. Therefore, when $|g(x, b)| \geq |g(x, q)|$, we have $|g(y, p)| < |g(y, b)|$. By Lemma 2, this means that every geodesic disk through p, q does not contain b . ◀

We now show two results about the length of the sides of a triangle. These results are trivial in the plane, but in the geodesic setting have not been proven before. In particular, they will be useful in Section 5 to prove the lower bounds on the diametral case.

► **Lemma 5.** *Let $\triangle(u, v, w)$ be a geodesic triangle, such that v is a convex vertex of $\nabla(u, v, w)$ and $\angle uvw \geq \frac{\pi}{3}$. Then $|g(u, w)| \geq \min\{|g(u, v)|, |g(v, w)|\}$.*

Proof. Consider the geodesic core $\nabla(u, v, w) = \triangle(u', v, w')$ of $\triangle(u, v, w)$. Recall that $\angle uvw$ is the angle at vertex v in $\nabla(u, v, w)$. Let q be the intersection of the geodesic $g(u', w')$ with the line containing the segment of $g(v, u)$ adjacent to v . Symmetrically, let p be the intersection of the geodesic $g(u', w')$ with the line containing the segment of $g(v, w)$ adjacent to v , see Figure 4.

Since $\triangle(v, p, q)$ is a Euclidean triangle, we know that $|pq| \geq \min\{|vp|, |vq|\}$. We can assume, without loss of generality, that $|pq| \geq |vq|$. Let x be the vertex that is the endpoint of the first edge of $g(v, u)$. By construction $|vq| = |vx| + |xq|$. Note that $|g(p, q)| \geq |pq| \geq |vq| = |vx| + |xq|$ (1). By the triangle inequality, $|g(x, u')| \leq |xq| + |g(q, u')|$. This implies that $|g(x, u')| - |xq| \leq |g(q, u')|$ (2). Putting it all together, we have

$$\begin{aligned} |g(v, u)| &= |g(u, u')| + |g(v, u')| \text{ by definition} \\ &= |g(u, u')| + |vx| + |g(x, u')| \text{ by definition} \\ &= |g(u, u')| + |vx| + |g(x, u')| - |xq| + |xq| \end{aligned}$$

$$\begin{aligned}
&\leq |g(u, u')| + |vx| + |g(q, u')| + |xq| \text{ by (2)} \\
&\leq |g(q, u)| + |g(p, q)| \text{ by (1)} \\
&\leq |g(p, u)| \\
&\leq |g(w, u)|
\end{aligned}$$

We conclude that $|g(u, w)| \geq \min\{|g(u, v)|, |g(v, w)|\}$. \blacktriangleleft

► **Lemma 6.** *Let $\triangle(u, v, w)$ be a geodesic triangle, such that $\angle uvw \leq \frac{\pi}{3}$. Then we have that $|uw| \leq \max\{|g(u, v)|, |g(v, w)|\}$.*

Proof. Since the three points u, v, w form a Euclidean triangle, we know that $|uw| \leq \max\{|uv|, |vw|\}$. We can assume, without loss of generality, that $|uw| \leq |uv|$. Since $|uv| \leq |g(u, v)|$, we thus have that $|uw| \leq |g(u, v)| \leq \max\{|g(u, v)|, |g(v, w)|\}$. \blacktriangleleft

3.2 Some differences

In contrast to the similarities in the previous section, there are properties that cannot be generalized from the Euclidean to the geodesic setting. We highlight a few differences. One difference which we alluded to in Section 1 is that the geodesic metric, unlike the Euclidean metric, does not necessarily have bounded doubling dimension. Another difference is that the geodesic bisector of a pair of points inside a simple polygon can intersect a line segment more than once [10].

The Voronoi diagrams of order k were one of the tools used to prove the lower bounds in the Euclidean setting, see, e.g., [12, 13, 23]. However, some important properties of order- k Voronoi diagrams do not extend to the geodesic setting. For instance, in the Euclidean setting, the bisector of a pair of points is a line, whereas in the geodesic setting, the bisector is a finite curve. Moreover, this means that in the Euclidean case, three points in general position uniquely define a disk with those points on the boundary. On the other hand, given three points inside a polygon P , there may not exist a point in P equidistant to all three. Thus, there are cases where three points do not define a geodesic disk with the three points on the boundary [4].

Another difference between the two settings is that the well-known lifting transformation that maps a point $(x, y) \in \mathbb{R}^2$ to the point $(x, y, x^2 + y^2) \in \mathbb{R}^3$ used in [23] to obtain the lower bound of $\frac{n}{4.7}$ is no longer applicable in the geodesic setting. The exact number of edges of the order- k Voronoi diagram is known in the Euclidean case [18]. The formula for this exact number of edges was fundamental to Edelsbrunner et al.'s lower bound proof [13]. However, only upper bounds are known in the geodesic case [8]. In fact, this is one of the reasons why the lower bounds we obtain in the geodesic setting are not as strong as in the Euclidean setting.

4 The general case

In this section we provide a generalization to the geodesic setting of lower bounds that are known for the original problem in the plane. In particular, we generalize results from [14, 15, 19]. We note that upper bounds for the Euclidean setting are in general also valid for the geodesic setting, since one can take the constructions provided in [2, 12, 14], and enclose the relevant set of points in a bounding polygon, big enough so that the geodesic disks through every three points coincide with the corresponding Euclidean disks.

10:10 On Geodesic Disks Enclosing Many Points

We begin by adapting some definitions and results from [19] to our setting. The *intersection number* $I(S)$ of a set of points S in a polygon is defined as the number of different geodesic paths $g(u, v)$, $g(x, y)$ for $u, v, x, y \in S$ such that $g(u, v) \cap g(x, y) \neq \emptyset$. Then, the *intersection graph* $G(S) = (V(G(S)), E(G(S)))$ is defined with $V(G(S)) = \{g(u, v) | u, v \in S, u \neq v\}$, and two vertices $g(u, v)$, $g(x, y)$ being joined by an edge in $G(S)$ if $g(u, v) \cap g(x, y) \neq \emptyset$. Finally, the oriented graph $\vec{G}(S)$ of $G(S)$ is defined orienting an edge $g(u, v) - g(x, y)$ as $g(u, v) \rightarrow g(x, y)$ if any geodesic disk through u and v contains x or y , or otherwise orienting $g(x, y) \rightarrow g(u, v)$. This orientation is consistent because of Lemma 3.

► **Theorem 7.** *For any set S of $n \geq 5$ points in a polygon, $\Pi(n) \geq \lceil \frac{n-2}{60} \rceil \approx 0.0166n$.*

Proof. By Kuratowski's Theorem [17], each subset $S' \subset S$, with $|S'| = 5$, contains four points w, x, y, z such that $g(w, x) \cap g(y, z) \neq \emptyset$. Moreover, the subset $\{w, x, y, z\}$ appears in exactly $n - 4$ subsets of S with five elements. Thus, $I(S) \geq \frac{\binom{n}{5}}{n-4}$. By the definition of the intersection graph $G(S)$, $|E(G(S))| = I(S)$, thus $|E(G(S))| \geq \frac{\binom{n}{5}}{n-4}$.

Let $d^+(g(x, y))$ be the out-degree of a vertex $g(x, y)$ in $\vec{G}(S)$. Then,

$$\sum_{g(x, y) \in V(\vec{G}(S))} d^+(g(x, y)) = |E(G(S))| \geq \frac{\binom{n}{5}}{n-4}.$$

Since we have exactly $\binom{n}{2}$ vertices in $\vec{G}(S)$, there is a vertex $g(u_0, v_0)$ with

$$d^+(g(u_0, v_0)) \geq \frac{\frac{\binom{n}{5}}{n-4}}{\binom{n}{2}} = \frac{(n-2)(n-3)}{60}.$$

This means that any geodesic disk through u_0 and v_0 contains at least one endpoint of $\frac{(n-2)(n-3)}{60}$ geodesic paths $g(x, y)$, $x, y \in S$. Since each point could appear in at most $n - 3$ of these geodesic paths, any geodesic disk through u_0 and v_0 contains at least $\left\lceil \frac{\frac{(n-2)(n-3)}{60}}{n-3} \right\rceil = \left\lceil \frac{n-2}{60} \right\rceil$ points of S different from u_0 and v_0 . ◀

We now establish an improved lower bound on the number of points contained in any disk with x and y on its boundary, by generalizing the results by Hayward in [14] to the geodesic setting. We begin with a preliminary combinatorial result.

► **Lemma 8.** *Among any set S' of k points in a polygon P , there are at least $\binom{k-3}{2}$ pairs of points $x, y \in S'$ such that every geodesic disk through x and y contains at least one of the other $k - 2$ points of S' .*

Proof. Let $G = (V, E)$ be the graph where $V = S'$, and two vertices are joined by an edge if at least one of the geodesic disks through them does not contain any of the other $k - 2$ points of S' . The edge between u and v is drawn as the geodesic path $g(u, v)$. We claim that G is planar, and the edges can be drawn without crossings inside P . Suppose there are two edges $g(u, v), g(x, y)$ in G , for $u, v, x, y \in S'$, that intersect. By Lemma 3, one of the edges has the property that every geodesic disk through both of its endpoints contains at least one of the endpoints of the other diagonal, which contradicts the definition of the edges of G . Thus, G is a planar graph with k vertices. By Euler's formula, $|E| \leq 3k - 6$, so S' contains at least $\binom{k}{2} - (3k - 6) = \binom{k-3}{2}$ pairs that satisfy the property. ◀

Now, using Lemma 8, we obtain a tighter lower bound for the number of points contained in any geodesic disk with two points u and v on its boundary.

► **Theorem 9.** *For any set S of $n \geq 8$ points in a polygon, $\Pi(n) \geq \lceil \frac{5}{84}(n-2) \rceil = \lceil \frac{n-2}{16.8} \rceil \approx 0.0595n$.*

Proof. Two points in a subset $S' \subset S$ of k points *dominate* S' if every geodesic disk through them contains at least one of the other $k-2$ points of S' . There are $\binom{n}{k}$ subsets of size k . By Lemma 8, each of these sets has $\binom{k-3}{2}$ dominating pairs. The average number of subsets dominated by a pair is $\frac{\binom{k-3}{2}\binom{n}{k}}{\binom{n}{2}}$. Thus, there is a pair of points $\{u, v\} \in S'$ that dominates at least $\frac{\binom{k-3}{2}\binom{n}{k}}{\binom{n}{2}}$ subsets of S .

For each of these subsets, every geodesic disk through u and v contains at least one of the other $k-2$ points of S' . Since a point of $S \setminus \{u, v\}$ can be in at most $\binom{n-3}{k-3}$ subsets that include u and v , every geodesic disk through u and v contains at least

$$\frac{\binom{k-3}{2}\binom{n}{k}}{\binom{n}{2}\binom{n-3}{k-3}} = (n-2) \frac{(k-3)(k-4)}{k(k-1)(k-2)}$$

other points of S . The largest coefficient of $n-2$ is $\frac{5}{84}$, obtained for $k=8$ and for $k=9$. ◀

We now present our strongest lower bound for $\Pi(n)$ by generalizing the approach by Edelsbrunner et al. [13] to the geodesic setting.

► **Theorem 10.** *For any set S of n points in a polygon, $\Pi(n) \geq \lceil \frac{n}{5} \rceil + 1$.*

Proof. Unless specified otherwise, in this proof, order- k diagrams refer to order- k geodesic Voronoi diagrams. Recall that a geodesic disk through two points containing $k-1$ points in its interior has its center on an edge of the order- k diagram. Consider the set of all disks that pass through two points $p, q \in S$ in the polygon. These disks are all centered on the bisector $b(p, q)$, and there are at most $n-2$ points on the bisector where a third point of S is on the boundary of the disk. This partitions the bisector into *at most* $n-1$ pieces such that any circle through p and q centered on any point belonging to a fixed piece s contains in its interior the same number of points $\rho(s)$. The pieces s such that $\rho(s) = k-1$ are precisely the edges of the order- k diagram.

Now consider some fixed bisector $b(p, q)$. The edges of the order-1 through order- k diagrams on $b(p, q)$ form a number of connected components (which is strictly less than the total number of edges of all the diagrams). Let λ_k be the number of connected components formed by *all* the bisectors of the order-1 through order- k diagrams. The goal is to find the largest value of k such that $\lambda_k < \binom{n}{2}$. Every bisector is completely covered by a single connected component when $k = n-2$. When $\lambda_k < \binom{n}{2}$, by the pigeonhole principle there exists a bisector $b(p, q)$ that does not have any edges of the order-1 through order- k diagrams. Therefore, every geodesic disk with p and q on its boundary must contain in its interior at least k points.

We derive a formula for λ_k in terms of the number of edges in the order-1 to the order- k diagram by considering how the pieces must “pair up”, in the sense that a piece s with $\rho(s) = i$ can only be adjacent to pieces s' with $\rho(s') = i \pm 1$. If a piece s with $\rho(s) = k$ is adjacent to two pieces s_1 and s_2 with $\rho(s_1) = \rho(s_2) = k-1$, the number of connected components is decreased by one. If s is adjacent to only one piece s_1 with $\rho(s_1) = k-1$, it extends an existing component. Otherwise, if it is not adjacent to any piece s_1 with $\rho(s_1) = k-1$, it creates a new connected component. Let ϵ_a, ϵ_b and ϵ_c be, respectively, the number of pieces s with $\rho(s) = k$ of the first, second, and third type. Then we have

$$\lambda_k = \lambda_{k-1} + \epsilon_c - \epsilon_a \tag{1}$$

10:12 On Geodesic Disks Enclosing Many Points

Let e_k be the number of edges in an order- k diagram and let S_k be the number of edges of the order- k diagram that intersect the boundary. The total number of endpoints of all order- k edges is given by $2e_k - S_k$ since the edges that intersect the boundary of the polygon only contribute one endpoint. The edges intersecting the boundary are analogous to the unbounded edges in the Euclidean setting. Moreover, we note that $e_k = \epsilon_a + \epsilon_b + \epsilon_c$.

Let ϕ_i be the number of unpaired endpoints in λ_i . Consider all the unpaired endpoints of the connected components of λ_{k-1} . These endpoints can only come from edges in the order- $k-1$ diagram since all edges of the order- i diagrams for $i \in \{1, \dots, k-2\}$ are all paired. Therefore, these unpaired endpoints can only pair up with edges from the order- k diagram, which implies that $\phi_{k-1} = 2\epsilon_a + \epsilon_b$. Thus, from Equation (1) and the fact that $e_k = \epsilon_a + \epsilon_b + \epsilon_c$, we have that $\lambda_k = \lambda_{k-1} + e_k - \phi_{k-1}$. If we iterate over k , we get that $\lambda_k = (e_k - \phi_{k-1}) + (e_{k-1} - \phi_{k-2}) + \dots + (e_2 - \phi_1) + e_1 = \sum_{i=1}^k e_i - \sum_{i=1}^{k-1} \phi_i$. Note that ϕ_{i-1} is the number of unpaired endpoints in λ_{i-1} , which are the paired endpoints of λ_i . Therefore, $\phi_i + \phi_{i-1} = 2e_i - S_i$ —the number of endpoints in order- i diagrams. We use this to rewrite the equation for λ_k in terms of e_i and S_i .

$$\begin{aligned} \lambda_k &= \sum_{i=1}^k e_i - [(\phi_{k-1} + \phi_{k-2}) + (\phi_{k-3} + \phi_{k-4}) + \dots] \text{ (end of sum depends on parity of } k) \\ &= \sum_{i=1}^k e_i - [(2e_{k-1} - S_{k-1}) + (2e_{k-3} - S_{k-3}) + \dots + (2e_{2-or-1} - S_{2-or-1})] \text{ (depending on parity of } k) \\ &= \sum_{i=1}^k (-1)^{k-i} e_i + (S_{k-1} + S_{k-3} + \dots + S_{2-or-1}) \text{ (} S_2 \text{ or } S_1 \text{ depending on parity of } k) \end{aligned}$$

To complete the bound on λ_k , we need to bound the size of the order- j diagrams. Exact bounds on the size of order- j Euclidean Voronoi diagrams are known but do not apply since unlike the Euclidean case, the geodesic order- j Voronoi diagram has no unbounded edges. We can, however, establish upper and lower bounds using abstract Voronoi diagrams [16]: Geodesic bisectors fulfill all the axioms for an abstract Voronoi diagram except for the property of being unbounded, which in the geodesic setting corresponds to hitting the boundary of the polygon. We can crop this diagram with a closed Jordan curve that encloses all the intersections between bisectors (as is the case for the polygon and the geodesic bisectors) to have an abstract Voronoi diagram [20]. It is this cropping that no longer allows us to have an exact formula. Since an order- j geodesic Voronoi diagram is an order- j abstract Voronoi diagram, we have that the number of edges $e_j = 3(F_j - 1) - S_j$, where $F_j = 2jn - j^2 - n + 1 - \sum_{i=1}^{j-1} S_i$ is the number of faces of an order- j geodesic Voronoi diagram, see [8, Lemmas 15 and 16]. Hence, $e_j = 6jn - 3j^2 - 3n - 3\sum_{i=1}^{j-1} S_i - S_j$. Then we get

$$e_j - e_{j-1} = 6n - 6j + 3 - 2S_{j-1} - S_j. \quad (2)$$

In the final calculation, as in the proof of the lower bound in [13], we distinguish the case that k is even from k odd. However, in both cases we have that

$$\lambda_k = 3kn - \frac{3}{2}k^2 - \frac{3}{2}k - \sum_{i=1}^k S_i.$$

We use the fact that $\sum_{i=1}^j S_i \geq j(j+1)$, see [8, Lemma 11], which implies that

$$\lambda_k \leq 3kn - \frac{3}{2}k^2 - \frac{3}{2}k - k(k+1) = 3kn - \frac{5}{2}k^2 - \frac{5}{2}k,$$

which is a quadratic polynomial in k . We note that $\lambda_k < \binom{n}{2}$ whenever $k < \frac{n}{5}$, i.e., $k \leq \lceil \frac{n}{5} \rceil - 1$. Then, as mentioned above, the pigeonhole principle implies that every geodesic disk with p and q on its boundary must contain in its interior at least $k \leq \lceil \frac{n}{5} \rceil - 1$ points. Thus, we get that $\Pi(n) \geq \lceil \frac{n}{5} \rceil + 1$, with the plus one accounting for p and q . ◀

4.1 A Constant Fraction of Points Inside and Outside Geodesic Disks

In this section we show that for a point set S , there is a pair of points $x, y \in S$ such that any disk through them contains a constant fraction of points of S , both inside and outside the disks. We obtain this bound by modifying the previous proof of Theorem 9 by using Lemma 11. For completeness, we add here the proof of the bound.

► **Lemma 11.** *Among any set S' of k points in a polygon, there are at least $\binom{k-3}{2}$ pairs of points $x, y \in S'$ such that every geodesic disk through x and y leaves outside at least one of the other $k - 2$ points of S' .*

Proof. Let $G = (V, E)$ be the graph where $V = S'$, and two vertices are joined by an edge if at least one of the geodesic disks through them contains one of the other $k - 2$ points of S' . The edge between u and v is drawn as the geodesic path $g(u, v)$. We claim that G is planar, and the edges can be drawn without crossings inside P . Suppose there are two edges $g(u, v), g(x, y)$ in G , for $u, v, x, y \in S'$, that intersect. By Lemma 3, one of the edges has the property that every geodesic disk through both of its endpoints leaves outside at least one of the endpoints of the other diagonal, which contradicts the definition of the edges of G . Thus, G is a planar graph with k vertices. By Euler's formula, $|E| \leq 3k - 6$, so S' contains at least $\binom{k}{2} - (3k - 6) = \binom{k-3}{2}$ pairs that satisfy the property. ◀

► **Theorem 12.** *For any set S of $n \geq 21$ points in a polygon, $\Pi^{in-out}(n) \geq \lceil \frac{16}{665}(n - 2) \rceil \approx \lceil \frac{n-2}{41.6} \rceil \approx 0.024n$.*

Proof. In the same spirit as in the proof Theorem 9, we define two points in a subset $S' \subset S$ of k points to be *dominated by S'* if every geodesic disk through them leaves outside at least one of the other $k - 2$ points of S' .

Among any set of k points in P there are at most $\binom{k}{2} - \binom{k-3}{2} = 3k - 6$ pairs of points that do not satisfy the property of Lemma 8 and at most $3k - 6$ pairs of points (possibly different than the previous pairs) that do not satisfy the property of Lemma 11. For $k \geq 11$, it holds that $2(3k - 6) < \binom{k}{2}$, so that there is at least one pair of points $\{u, v\} \in S'$ that satisfies both properties, hence dominates and is dominated by at least $\frac{[2\binom{k-3}{2} - \binom{k}{2}]\binom{n}{k}}{\binom{n}{2}}$ subsets of S . For each of these subsets, every geodesic disk through u and v has at least one of the other $k - 2$ points of S' , both inside and outside. Since a point of $S \setminus \{u, v\}$ could be in at most $\binom{n-3}{k-3}$ subsets containing u and v , every geodesic disk through u and v has at least

$$\frac{\left[2\binom{k-3}{2} - \binom{k}{2}\right]\binom{n}{k}}{\binom{n}{2}\binom{n-3}{k-3}} = (n - 2) \frac{2(k - 3)(k - 4) - k(k - 1)}{k(k - 1)(k - 2)}$$

other points of S , both inside and outside. The largest coefficient of $n - 2$ is obtained for $k = 21$, thus $k = 21$ leads to the largest possible coefficient, the one in the statement. ◀

4.2 Sets in geodesically convex position

In the Euclidean setting, Hayward et al. [15] proved that when points are in convex position, for every disk containing two points on the boundary, there always exists a pair of points such that any disk containing that pair contains at least $\lceil \frac{n}{3} \rceil + 1$. We generalize that result to the geodesic setting by showing that $\bar{\Pi}(n) \geq \lceil \frac{n}{3} \rceil + 1$. The geodesic case poses its unique challenges, for example, it is not always possible to find an enclosing geodesic disk with three points on the boundary. However, when no such enclosing geodesic disk exists, we show that there exists a pair of points such that every geodesic disk through them is an enclosing geodesic disk.

► **Lemma 13.** *Let S be a set of points in a polygon P . If there is no enclosing geodesic disk through three points of S , then there exist two points of S such that any geodesic disk through them is an enclosing geodesic disk.*

Proof. Define the notation $D(c)$ to be the disk $D(c, \max_{p \in S} |g(c, p)|)$ which is the geodesic disk centered at an arbitrary point $c \in P$, and having as radius the maximum geodesic distance from c to any other point $p \in S$. By construction, this is an enclosing geodesic disk. Consider an arbitrary disk $D(x)$ centered on a point $x \in S$ with a point $y \in S$ on the boundary of $D(x)$. Either $D(x)$ has a second point z on its boundary or only has y on its boundary. If the latter is true, then we show how to obtain an enclosing disk with a second point on its boundary. Consider the disk $D(x)$ as the center x moves along $g(x, y)$ towards y . Since the radius is shrinking, eventually, a second point $z \in S$ will reach the boundary of $D(x)$. Note that $D(x)$ always remains enclosing.

Thus, $D(x)$ is an enclosing disk with y and z on its boundary. We now show that every geodesic disk through y and z is enclosing. Let w be any point on $b(y, z)$. $D(w)$ has y and z on its boundary. Moreover, $D(w)$ must be enclosing. If $D(w)$ is not enclosing, then we contradict the fact that there is no enclosing geodesic disk through three points. The lemma follows. ◀


We now generalize the result from Hayward et al. [15] for point sets in convex position to the geodesic setting. In particular, when S is a set of points in P that is in geodesically convex position, then there is a pair of points such that any geodesic disk through them contains at least $\lceil \frac{n}{3} \rceil + 1$ points of S .

► **Theorem 14.** *For any set S of n points in geodesically convex position in a polygon, $\bar{\Pi}(n) \geq \lceil \frac{n}{3} \rceil + 1$.*

Proof. By Lemma 13, if there is no enclosing geodesic disk through three points, there exist two points of S such that any geodesic disk through them is an enclosing geodesic disk. Then all such disks contain all points of S , and the inequality holds.

Otherwise, there is an enclosing disk D that has three points on its boundary, namely x, y, z in clockwise order along the boundary of D . Let $f(\Delta(x, y, z))$ be the maximum among the number of points from S contained in each of the three regions $D \cap P^+(x, y)$, $D \cap P^+(y, z)$ and $D \cap P^+(z, x)$. Among all enclosing disks with 3 points on the boundary, consider disk D_1 with points $u, v, w \in S$ on its boundary which minimizes $f(\Delta(u, v, w))$. There must be at least one region among $D_1 \cap P^+(u, v)$, $D_1 \cap P^+(v, w)$ and $D_1 \cap P^+(w, u)$ whose interior has at least $\lceil \frac{n-3}{3} \rceil$ points. Assume, without loss of generality, that $D_1 \cap P^+(u, v)$ does.

If the interior of $D_1 \cap P^-(u, v)$ contains fewer than $\lceil \frac{n-3}{3} \rceil$ points, then we continuously change D_1 by the disks through u, v obtained when moving the center c of D_1 along the part of the bisector $b(u, v)$ to the right of $\ell(u, c)$, as in Figure 5. While doing this, all those



■ **Figure 5** The geodesic disk D_2 is obtained by moving the center c of D_1 along the part of the bisector $b(u, v)$ (shown dashed) to the right of $\ell(u, c)$ until hitting point w_1 .

disks will keep being enclosing disks, by Lemma 2. As we move c , we may arrive at one of two situations: Either the center c hits the boundary of the polygon, or the boundary of the current disk, D_2 , hits a point $w_1 \in D_1 \cap P^+(u, v)$. We show, by contradiction, that the latter case cannot occur. If it occurred, the geodesic triangle $\triangle(u, w_1, v)$ divides D_2 into three regions $D_2 \cap P^+(u, w_1)$, $D_2 \cap P^+(w_1, v)$ and $D_2 \cap P^+(v, u)$. Note that $D_2 \cap P^+(v, u) = D_2 \cap P^-(u, v)$ contains fewer than $\lceil \frac{n-3}{3} \rceil$ points of S , since it is contained in $D_1 \cap P^-(u, v)$. Also, $D_2 \cap P^+(u, w_1)$ and $D_2 \cap P^+(w_1, v)$ both contain fewer points than $D_1 \cap P^+(u, v)$. Thus, $f(\triangle(u, w_1, v)) < f(\triangle(u, v, w))$, and contradicting the minimality of $f(\triangle(u, v, w))$.

Therefore, we will eventually hit the boundary. Then, every geodesic disk through u and v with center in the part of $b(u, v)$ to the right of $\ell(u, c)$ contains all the points in S , since the current D_1 is an enclosing disk. Also, by Lemma 2, every geodesic disk through u and v with center in the part of $b(u, v)$ to the left of $\ell(u, c)$ contains $D_1 \cap P^+(u, v)$, thus at least $\lceil \frac{n-3}{3} \rceil$ points of S .

It follows that there exist two points u, v such that any geodesic disk through u and v contains at least $\lceil \frac{n-3}{3} \rceil$ points of S in its interior. By accounting for u and v , on the disk boundary, we conclude that $\overline{\Pi}(n) \geq \lceil \frac{n}{3} \rceil + 1$. ◀

5 Diametral geodesic disk

In the Euclidean setting, Akiyama et al. [2] showed that for points in the plane (both in convex position and not) there always exists a diametral disk that contains $\lceil \frac{n}{3} \rceil + 1$ points of S and that this is a tight bound. In this section, we consider the problem of proving lower bounds on the values of $\Pi^{diam}(n)$ and $\overline{\Pi}^{diam}(n)$, generalizing the results in [2]. We prove that, among a set S of n points, there is a pair of points such that the diametral geodesic disk through them contains at least $\lceil \frac{n}{3} \rceil + 1$ points of S .

In order to obtain these lower bounds, we are going to use an enclosing geodesic disk with its center inside the geodesic convex hull of the points on its boundary.

10:16 On Geodesic Disks Enclosing Many Points

► **Lemma 15.** *Let S be a set of points in a polygon. There exists an enclosing geodesic disk D with (i) three points x, y, z on the boundary, and the center of D inside $\triangle(x, y, z)$, or (ii) two points x, y on the boundary, and the center of D at the midpoint of $g(x, y)$.*

Proof. Let D be an enclosing geodesic disk with three points on its boundary. If such a disk does not exist, by Lemma 13, there exist two points $x, y \in S$ such that every geodesic disk through them is a geodesic enclosing disk. In particular, the disk $D\left(c_{xy}, \frac{|g(x, y)|}{2}\right)$ is an enclosing diametral disk with x and y as its diametral endpoints.

Otherwise, let $x, y, z \in S$ be the three points on the boundary of D , in clockwise order. If the center of the disk is inside the geodesic triangle $\triangle(x, y, z)$, we are done. Otherwise, the center is inside one of the regions $D \cap P^+(x, y)$, $D \cap P^+(y, z)$ or $D \cap P^+(z, x)$. Assume, without loss of generality, it is in $D \cap P^+(x, y)$. Then we move the center c of D along $b(x, y)$ towards c_{xy} . If the center c reaches c_{xy} , then $D\left(c_{xy}, \frac{|g(x, y)|}{2}\right)$ is a diametral enclosing geodesic disk with x and y as its diametral endpoints. If the disk hits a third point $w \in S$, and the center of the disk is inside $\triangle(x, y, w)$, we are done. Otherwise, we repeat the process with the points x, y, w as the new points x, y, z until the center of the disk through the three points on the boundary of the disk lies inside the geodesic triangle formed by the three points on the boundary. This process ends because the distance to the center from the closest shortest path is decreasing at every iteration. ◀


► **Theorem 16.** *For any set S of n points in a polygon, $\Pi^{diam}(n) = \overline{\Pi}^{diam}(n) \geq \lceil \frac{n}{3} \rceil + 1$.*

Proof. Consider an enclosing disk D as guaranteed by Lemma 15. Either (i) D has exactly three points of S on the boundary, and the center of the geodesic disk is inside the geodesic triangle formed by the three points, or (ii) D has two points of S as diametral endpoints. For case (ii), D having two points on the boundary, the claimed result holds, since D contains all the points.

For case (i), let a, b and c be the three points on the boundary of the geodesic disk. We will prove that any point $x \in S$ is contained in at least one of the geodesic disks $D\left(c_{ab}, \frac{|g(a, b)|}{2}\right)$, $D\left(c_{bc}, \frac{|g(b, c)|}{2}\right)$ and $D\left(c_{ca}, \frac{|g(c, a)|}{2}\right)$.

First, consider the case where x is outside the geodesic triangle $\triangle(a, b, c)$. Then, x belongs to one of the regions $D \cap P^+(a, b)$, $D \cap P^+(b, c)$, and $D \cap P^+(a, c)$. Without loss of generality, assume $x \in D \cap P^+(a, b)$. By Lemma 15, the center of the geodesic disk is inside $\triangle(a, b, c)$, so it does not belong to $D \cap P^+(a, b)$. Then, by Lemma 2, x is contained in the diametral geodesic disk $D\left(c_{ab}, \frac{|g(a, b)|}{2}\right)$.

Now, we prove the case where $x \in \triangle(a, b, c)$. Consider the geodesic paths $g(a, x)$, $g(b, x)$ and $g(c, x)$. One of the angles $\angle axb$, $\angle bxc$ or $\angle axc$ is greater than or equal to $\frac{2\pi}{3}$, and the other two angles are smaller than or equal to $\frac{\pi}{3}$. Assume, without loss of generality, that $\angle bxc \geq \frac{2\pi}{3}$, see blue angle in Figure 6. Let o be the midpoint of the shortest path $g(b, c)$. We also know that one of the angles $\angle bxo$ or $\angle oxc$ is greater than or equal to $\frac{\pi}{3}$. Assume, without loss of generality, that it is the angle $\angle oxc$. When o is not visible from x , the angle $\angle oxc$ is equal to the angle $\angle bxc \geq \frac{2\pi}{3} > \frac{\pi}{2}$. Hence, by [21, Corollary 2], $|g(x, o)| < |g(o, c)| = |g(o, b)|$, and x is contained in the geodesic disk $D\left(c_{bc}, \frac{|g(b, c)|}{2}\right)$. In the case of o being visible from x , by Lemma 5, $|g(o, c)| \geq \min\{|g(c, x)|, |g(x, o)|\}$ and, by Lemma 6, $|g(x, o)| \leq \max\{|g(c, x)|, |g(o, c)|\}$, hence $|g(x, o)| \leq |g(o, c)| = |g(o, b)|$, and x is contained in the geodesic disk $D\left(c_{bc}, \frac{|g(b, c)|}{2}\right)$.



■ **Figure 6** The point o is visible from x .

Then, the union of the geodesic disks $D\left(c_{ab}, \frac{|g(a,b)|}{2}\right)$, $D\left(c_{bc}, \frac{|g(b,c)|}{2}\right)$ and $D\left(c_{ca}, \frac{|g(c,a)|}{2}\right)$ covers all the points in P . Since the points a, b and c are counted twice,

$$\left|D\left(c_{ab}, \frac{|g(a,b)|}{2}\right) \cap P\right| + \left|D\left(c_{bc}, \frac{|g(b,c)|}{2}\right) \cap P\right| + \left|D\left(c_{ca}, \frac{|g(c,a)|}{2}\right) \cap P\right| > n + 3.$$

By the pigeonhole principle, the result follows. ◀

6 Bichromatic geodesic disks

In this section, we consider the bichromatic version of the problem, which has also received attention in the Euclidean setting. Let B be a set of $\frac{n}{2}$ blue points and R be a set of $\frac{n}{2}$ red points inside a polygon P , and let $S = B \cup R$.

Following the procedure in [25], we obtain a similar result to Theorem 7, and we prove that it is possible to find a bichromatic pair such that any geodesic disk containing them contains at least $\left\lceil \frac{|S|-2}{72} \right\rceil$ points of S .

► **Theorem 17.** *For any set S of $\frac{n}{2} \geq 3$ red points and $\frac{n}{2} \geq 3$ blue points in a polygon $\Pi^{bichrom}(n) \geq \left\lceil \frac{n-2}{72} \right\rceil \approx 0.0139n$.*

Proof. Let $G(S) = (V(G(S)), E(G(S)))$ be the intersection graph defined as follows: $V(S) = \{g(u, v) | u, v \in S, u \text{ is red, and } v \text{ is blue}\}$, and two vertices $g(u, v), g(x, y)$ are joined by an edge in $G(S)$ if $g(u, v) \cap g(x, y) \neq \emptyset$. The oriented graph $\vec{G}(S)$ of $G(S)$ is defined orienting an edge $g(u, v) - g(x, y)$ as $g(u, v) \rightarrow g(x, y)$ if any geodesic disk through u and v contains x or y , or otherwise orienting $g(x, y) \rightarrow g(u, v)$. By Lemma 3, if two of these paths intersect, then one of the two bichromatic pairs fulfills the property that any geodesic disk through it will contain one of the other two endpoints, so this orientation is consistent.

We now obtain the minimum number of intersections between all pairs of geodesic paths. Let, as in Theorem 7, $I(S)$ be the number of pairs of geodesic paths that intersect. For each $S' \subset S$ with exactly 3 blue and 3 red points, by Kuratowski's Theorem [17], there are two blue points $w, x \in S'$ and two red points $y, z \in S'$ such that $g(w, y) \cap g(x, z) \neq \emptyset$. There are $\binom{\frac{n}{2}}{3}^2$ subsets of this type. Furthermore, every intersection belongs to exactly $\left(\frac{n}{2} - 2\right)^2$ subsets of this type. Dividing by this number, we remove repetitions, and we obtain that $I(S) \geq \left(\frac{\binom{\frac{n}{2}}{3}}{\frac{n}{2} - 2}\right)^2$.

By definition of the intersection graph $G(S)$, $|E(G(S))| = I(S)$ and $|V(G(S))| = \left(\frac{n}{2}\right)^2$. Hence, there exists one vertex $g(u_0, v_0) \in \vec{G}$ with out-degree at least $\frac{I(S)}{\left(\frac{n}{2}\right)^2}$.

10:18 On Geodesic Disks Enclosing Many Points

Then, we have that $\frac{I(S)}{\binom{n}{2}} \geq \left(\frac{\frac{n}{2}-1}{6}\right)^2$. This means that there exist two points u_0 and v_0 , one of each color, such that the geodesic path between them intersects at least $\left(\frac{\frac{n}{2}-1}{6}\right)^2$ other geodesic paths. Since every point is the endpoint of at most $\frac{n}{2}-1$ geodesic paths, every geodesic disk through this pair of points contains at least $\left(\frac{\frac{n}{2}-1}{6}\right)^2 \frac{1}{\frac{n}{2}-1} = \frac{\frac{n}{2}-1}{36} = \frac{n-2}{72}$ points of S . ◀

It was proved in [22] that any set of bichromatic points S , contains a pair $\{p, q\} \subset S$, one point of each color, such that any disk through p and q contains a positive fraction of the points of S , in particular $\frac{1}{36}$. We present the analogous colored result in the geodesic setting.

► **Theorem 18.** *For any set S of $\frac{n}{2} \geq 3$ blue points and $\frac{n}{2} \geq 3$ red points in a polygon, $\Pi^{bichrom}(n) \geq \lceil \frac{n-2}{36} \rceil + 2 \approx 0.0277n$.*

Proof. Let B be the set of blue points, and let R be the set of red points. By Kuratowski's Theorem and Lemma 3, for every set $Z \subset S$ of six points, three being blue and three being red, there is a pair $\{u, v\}$ of bichromatic points such that any geodesic disk with u and v on its boundary contains at least one point of $Z \setminus \{u, v\}$. Hence, there is a family \mathcal{Z} and a pair $\{u, v\} \subset S$ of bichromatic points such that $\{u, v\}$ belongs to every $Z \in \mathcal{Z}$ and

$$|\mathcal{Z}| \geq \frac{\binom{\frac{n}{2}}{3}^2}{\binom{\frac{n}{2}}{1}} = \frac{\left(\frac{n}{2}-2\right)^2 \left(\frac{n}{2}-1\right)^2}{6^2}. \quad (3)$$

Now we will upper-bound the value $|\mathcal{Z}|$. Let D be a geodesic disk with u and v on its boundary, and $m-2$ points in its interior. Each $Z \in \mathcal{Z}$ contains a point of $D \setminus \{u, v\}$. The set $\{u, v\}$ can be extended to Z , given that $D \cap (Z \setminus \{u, v\}) \neq \emptyset$, by (i) choosing a point of Z from the remaining $m-2$ points in D , (ii) choosing a point of the same color class from the remaining $\frac{n}{2}-2$ points in S , and (iii) choosing two points in S from the remaining color class. Thus, we get that

$$|\mathcal{Z}| \leq \binom{m-2}{1} \binom{\frac{n}{2}-2}{1} \binom{\frac{n}{2}-1}{2} = (m-2) \frac{\left(\frac{n}{2}-2\right)^2 \left(\frac{n}{2}-1\right)}{2}. \quad (4)$$

Finally, from equations (3) and (4) we get $m-2 \geq \frac{\frac{n}{2}-1}{18} \implies m \geq \frac{\frac{n}{2}+35}{18} = \frac{n+70}{36}$. ◀

7 Conclusion

We study different variants of the following question: Given a set S of n points in a polygon P , does there always exist two points $x, y \in S$, such that every geodesic disk containing x and y contains a constant fraction of the points of S ? This question has been studied intensely in the Euclidean setting [2, 12–15, 19, 22, 23, 25]. We focus on the geodesic versions of this problem. We note that all the solutions to these problems lend themselves to polynomial time algorithms since the characterizations are based on standard structures in Computational Geometry such as order- k Voronoi diagrams. An obvious open problem is the following: Can our lower bounds be improved? Basically, some of our geodesic lower bounds are not as strong as in the Euclidean case. It would be interesting to determine whether there is a separation between the bounds in the Euclidean and geodesic case. Since the acceptance of

this paper by the WADS Program Committee, we have improved the bounds in Theorems 12 and 18 by modifying the technique using order- k geodesic Voronoi diagrams as was done in the proof of Theorem 10. The proof of these improved results can be found in [11]. Finally, the main open problem which has been open since 1988 is whether we can find a tight bound for the original question?

References

- 1 Hugo A. Akitaya, Greg Aloupis, Jeff Erickson, and Csaba D. Tóth. Recognizing weakly simple polygons. *Discrete & Computational Geometry*, 58(4):785–821, 2017. doi:10.1007/S00454-017-9918-3.
- 2 Jin Akiyama, Yoshiyasu Ishigami, Masatsugu Urabe, and Jorge Urrutia. On circles containing the maximum number of points. *Discrete Mathematics*, 151(1-3):15–18, 1996. doi:10.1016/0012-365X(94)00076-U.
- 3 Boris Aronov. On the geodesic Voronoi diagram of point sites in a simple polygon. *Algorithmica*, 4(1-4):109–140, 1989. doi:10.1007/BF01553882.
- 4 Boris Aronov, Steven Fortune, and Gordon Wilfong. The furthest-site geodesic Voronoi diagram. *Discrete & Computational Geometry*, 9:217–255, 1993. doi:10.1007/BF02189321.
- 5 Imre Bárány and David G. Larman. A combinatorial property of points and ellipsoids. *Discrete & Computational Geometry*, 5:375–382, 1990. doi:10.1007/BF02187798.
- 6 Imre Bárány, James H. Schmerl, Stuart J. Sidney, and Jorge Urrutia. A combinatorial result about points and balls in Euclidean space. *Discrete & Computational Geometry*, 4:259–262, 1989. doi:10.1007/BF02187727.
- 7 Ahmad Biniaz, Prosenjit Bose, Anil Maheshwari, and Michiel Smid. Plane geodesic spanning trees, Hamiltonian cycles, and perfect matchings in a simple polygon. *Computational Geometry*, 57:27–39, 2016. doi:10.1016/J.COMGEO.2016.05.004.
- 8 Cecilia Bohler, Panagiotis Cheilaris, Rolf Klein, Chih-Hung Liu, Evanthia Papadopoulou, and Maksym Zavershynskiy. On the complexity of higher order abstract Voronoi diagrams. *Computational Geometry*, 48(8):539–551, 2015. doi:10.1016/J.COMGEO.2015.04.008.
- 9 Prosenjit Bose, Paz Carmi, and Thomas C. Shermer. Piercing pairwise intersecting geodesic disks. *Computational Geometry*, 98:101774, 2021. doi:10.1016/J.COMGEO.2021.101774.
- 10 Prosenjit Bose, Anthony D’Angelo, and Stephane Durocher. Approximating the smallest k -enclosing geodesic disc in a simple polygon. In *WADS*, volume 14079 of *Lecture Notes in Computer Science*, pages 179–192. Springer, 2023. doi:10.1007/978-3-031-38906-1_13.
- 11 Prosenjit Bose, Guillermo Esteban, David Orden, Rodrigo Silveira, and Tyler Tuttle. On geodesic disks enclosing many points, 2025. doi:10.48550/arXiv.2506.06477.
- 12 Mercè Claverol, Clemens Huemer, and Alejandra Martínez-Moraian. On circles enclosing many points. *Discrete Mathematics*, 344(10):112541, 2021. doi:10.1016/J.DISC.2021.112541.
- 13 Herbert Edelsbrunner, Nany Hasan, Raimund Seidel, and Xiao Jun Shen. Circles through two points that always enclose many points. *Geometriae Dedicata*, 32(1):1–12, 1989.
- 14 Ryan Hayward. A note on the circle containment problem. *Discrete & Computational Geometry*, 4(3):263–264, 1989. doi:10.1007/BF02187728.
- 15 Ryan Hayward, David Rappaport, and Rephael Wenger. Some extremal results on circles containing points. *Discrete & Computational Geometry*, 4(3):253–258, 1989.
- 16 Rolf Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes in Computer Science*. Springer, 1989. doi:10.1007/3-540-52055-4.
- 17 Casimir Kuratowski. Sur le probleme des courbes gauches en topologie. *Fundamenta mathematicae*, 15(1):271–283, 1930.
- 18 Der-Tsai Lee. On k -nearest neighbor Voronoi diagrams in the plane. *IEEE transactions on computers*, 100(6):478–487, 1982. doi:10.1109/TC.1982.1676031.
- 19 Victor Neumann-Lara and Jorge Urrutia. A combinatorial result on points and circles on the plane. *Discrete mathematics*, 69(2):173–178, 1988. doi:10.1016/0012-365X(88)90015-5.

- 20 Evanthia Papadopoulou. Abstract Voronoi-like graphs: Extending Delaunay’s theorem and applications. In *SoCG*, volume 258 of *LIPICs*, pages 52:1–52:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.SOCG.2023.52.
- 21 Richard Pollack, Micha Sharir, and Günter Rote. Computing the geodesic center of a simple polygon. *Discrete & Computational Geometry*, 4(6):611–626, 1989. doi:10.1007/BF02187751.
- 22 Maria N. Prodromou. A combinatorial property of points and balls, a colored version. *Discrete & Computational Geometry*, 38(4):641–650, 2007. doi:10.1007/S00454-007-9003-4.
- 23 Pedro A. Ramos and Raquel Viaña. Depth of segments and circles through points enclosing many points: a note. *Computational Geometry*, 42(4):338–341, 2009. doi:10.1016/J.COMGEO.2008.07.001.
- 24 Godfried Toussaint. Computing geodesic properties inside a simple polygon. *Revue d’intelligence artificielle*, 3(2):9–42, 1989.
- 25 Pilar Valencia Saravia. *Problemas de cobertura circular*. PhD thesis, Universidad Nacional Autónoma de México, 2005.