

Computing the Yolk in Spatial Voting Games without Computing Median Lines

Joachim Gudmundsson, Sampson Wong

University of Sydney
Sydney, Australia

joachim.gudmundsson@sydney.edu.au, swon7907@sydney.edu.au

Abstract

The yolk is an important concept in spatial voting games: the yolk center generalises the equilibrium and the yolk radius bounds the uncovered set. We present near-linear time algorithms for computing the yolk in the plane. To the best of our knowledge our algorithm is the first that does not precompute median lines, and hence is able to break the best known upper bound of $O(n^{4/3})$ on the number of limiting median lines. We avoid this requirement by carefully applying Megiddo’s parametric search technique, which is a powerful framework that could lead to faster algorithms for other spatial voting problems.

1 Introduction

Voting theory is concerned with preference aggregation and group decision making. A classic framework for aggregating voter’s preferences is the Downsian (1957), or spatial model of voting.

In this model, voters are positioned on a ‘left-right’ continuum along multiple ideological dimensions, such as economic, social or religious. These dimensions together form the policy space. Each voter is required to choose a single candidate from a set of candidates, and a common voter preference function is a metric/distance function within the policy space. An intuitive reason behind using metric preferences is that voters tend to prefer candidates ideologically similar to themselves.

The spatial model of voting with metric preferences have been studied extensively, both theoretically (McKelvey 1976; 1986; Enelow and Hinich 1984; Miller, Grofman, and Feld 1989; Tovey 1993) and empirically (Poole and Rosenthal 1984; 1991; 2001; Ordeshook 1993; Schofield, Miller, and Martin 2003; Schofield 2004; 2007). Recently, lower bounds were provided on the distortion of voting rules in the spatial model, and interestingly, metrics other than the Euclidean metric were considered (Anshelevich, Bhardwaj, and Postl 2015; Skowron and Elkind 2017; Goel, Krishnaswamy, and Munagala 2017).

We focus our attention on two-candidate spatial voting games, where the winner is the candidate preferred by a simple majority of voters. In a one dimension policy space,

Black’s Median Voter Theorem (1948) states that a voting equilibrium (alt. Condorcet winner, plurality point, pure Nash equilibrium) is guaranteed to exist and coincides with the median voter.

Naturally, social choice theorists searched for the equilibrium in the two dimensional policy space, but these attempts were shown to be fruitless (Plott 1967). The initial reaction was one of cynicism (McKelvey 1976), but in response a multitude of generalisations were developed, with the yolk being one such concept (McKelvey 1986; Miller, Grofman, and Feld 1989). The yolk in the Euclidean \mathcal{L}_2 metric is defined as the minimum radius disk that intersects all median lines of the voters.

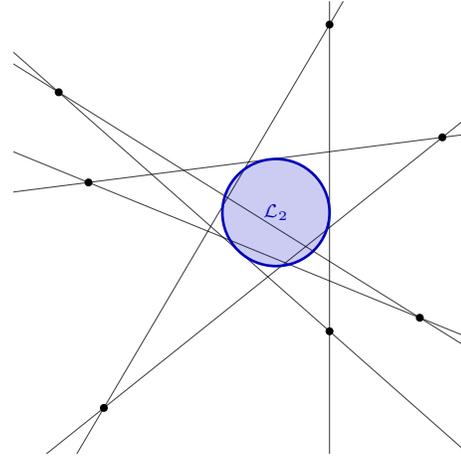


Figure 1: The \mathcal{L}_2 yolk intersects all median lines of voters.

The yolk is an important concept in spatial voting games due to its simplicity and its relationship to other concepts. The yolk radius provides approximate bounds on the uncovered set (Feld et al. 1987; Miller 1980; Miller, Grofman, and Feld 1989), limits on agenda control (Feld, Grofman, and Miller 1989), Shapley-Owen power scores (Feld and Grofman 1990), the Finagle point (Wuffle et al. 1989) and the ϵ -core (Tovey 2011). As such, studies on the size of the yolk (Feld, Grofman, and Miller 1988; Koehler 1990; Tovey 2010a) translate to these other concepts as well.

From the perspective of computational social choice, this

raises the following problem: Are there efficient algorithms for computing the yolk? Fast algorithms would, for instance, facilitate empirical studies on large data sets. Tovey (1992) provides the first polynomial time algorithm, which in two dimensions, computes the yolk in $O(n^4)$ time. De Berg et al. (2016) provides an improved $O(n^{4/3} \log^{1+\varepsilon} n)$ time algorithm for the same.

The shortcoming of existing algorithms is that they require the computation of all limiting median lines, which are median lines that pass through at least two voters (Stone and Tovey 1992). However, there are $\Omega(ne^{\sqrt{\log n}})$ (Tóth 2000) limiting median lines in the worst case. Moreover, the best known upper bound of $O(n^{4/3})$ seems difficult to improve on (Dey 1997). It is an open problem as to whether there is a faster algorithm that computes the yolk without precomputing all limiting median lines.

Problem Statement

Given a set V of n points in the plane, a median line of V is any line that divides the plane into two closed halfplanes, each with at most $n/2$ points. The yolk is a minimum radius disk in the \mathcal{L}_p metric that intersects all median lines of V .

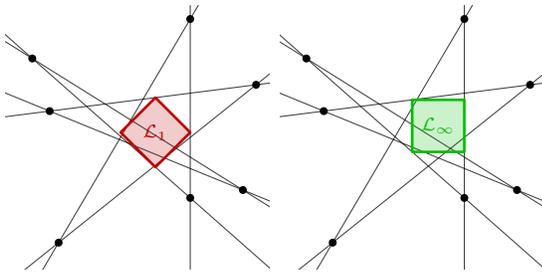


Figure 2: Example of yolks in the \mathcal{L}_1 and \mathcal{L}_∞ metrics.

We compute yolks in the \mathcal{L}_1 (Taxicab), the \mathcal{L}_2 (Euclidean), and the \mathcal{L}_∞ (Uniform) metrics. As shown in Figure 2, the yolk in \mathcal{L}_1 is the smallest 45° -rotated square and in \mathcal{L}_∞ the smallest axis-parallel square, that intersects all median lines of V .

Our Contributions and Results

Our contributions are, first, an algorithm that computes the yolk in the \mathcal{L}_1 and \mathcal{L}_∞ metrics in $O(n \log^7 n)$ time, and second, an algorithm that computes a $(1 + \varepsilon)$ -approximation of the yolk in the \mathcal{L}_2 metric in $O(n \log^7 n \cdot \log^4 \frac{1}{\varepsilon})$ time.

We achieve the improved upper bounds by carefully applying Megiddo's (1983) parametric search technique, which is a powerful yet complex technique and that could be useful for other spatial voting problems.

The parametric search technique is a framework for converting decision algorithms into optimisation algorithms. For the yolk problem, a decision algorithm would decide whether a given disk intersects all median lines. If this decision algorithm satisfies the three properties as specified by the framework, then Megiddo's result states that there is an efficient algorithm to compute the yolk.

For the purposes of designing a decision algorithm with the desired properties, we instead consider the more general problem of finding the smallest regular, k -sided polygon that intersects all median lines of V . The regular k -sided polygon $P_k(r, x, y)$ is shown in Figure 3 and is defined as:

Definition 1. Given an integer $k \geq 3$, construct the regular k -sided polygon $P_k(r, x, y)$ by:

- Constructing a circle with radius r and centered at (x, y) .
- Placing a vertex at the top-most point on the circle, i.e. at $(x, y + r)$.
- Placing the remaining $k - 1$ vertices around the circle so that the k vertices are evenly spaced.

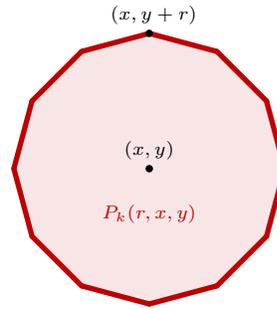


Figure 3: The regular, k -sided polygon $P_k(r, x, y)$.

In Section 2, we present the decision algorithm, which given a regular, k -sided polygon $P_k(r, x, y)$, decides whether the polygon intersects all median lines of V . Next, in Section 3, we apply Megiddo's technique to the decision algorithm and prove the convexity and parallelisability properties. This leaves one final property, the existence of critical hyperplanes, left to check. We prove this final property in Sections 4-6, thus completing the parametric search. Finally, in Section 7, we show that our general problem for the regular, k -sided polygon $P_k(r, x, y)$ implies the claimed running times by setting $k = 4$ for \mathcal{L}_1 and \mathcal{L}_∞ , and $k = \frac{1}{\varepsilon}$ for \mathcal{L}_2 .

2 Decision Algorithm

The aim of this section is to design an algorithm that solves the following decision problem:

Definition 2. Given an integer $k \geq 3$ and a set V of n points in the plane, the decision problem $D_{k,V}(r, x, y)$ is to decide whether the polygon $P_k(r, x, y)$ intersects all median lines of V .

We show that there is a comparison-based decision algorithm that solves $D_{k,V}(r, x, y)$ in $O(n \log n \cdot \log k)$ time, provided the following two comparison-based subroutines.

Subroutine 1. A comparison-based subroutine that, given a point p and a regular k -sided polygon $P_k(r, x, y)$, decides if p is outside $P_k(r, x, y)$ in $O(\log k)$ time.

Subroutine 2. A comparison-based subroutine that, given points p, q outside a regular k -sided polygon $P_k(r, x, y)$,

computes the relative clockwise order of the four tangent lines drawn from $\{p, q\}$ to $P_k(r, x, y)$ in $O(\log k)$ time.

Although the running time of these two subroutines are not too difficult to prove, we shall see in Section 3 that these subroutines must satisfy a stronger requirement for the parametric search technique to apply. We will formally define the stronger requirement in the next section. To avoid repetition, we simultaneously address the subroutine and the stronger requirement in Sections 5 and 6. But for now, we assume the subroutines exist and present the decision algorithm:

Theorem 1. *Given an integer $k \geq 3$ and a set V of n points in the plane, there is a comparison-based algorithm that solve the decision problem $D_{k,V}(r, x, y)$ in $O(n \log n \cdot \log k)$ time, provided that Subroutine 1 and Subroutine 2 exist.*

Proof. The proof comes in three parts. First, we transform the decision problem $D_{k,V}(r, x, y)$ into an equivalent form that does not have median lines in its statement. Then, we present a sweep line algorithm for the transformed version. Finally, we perform an analysis of the running time.

Consider for now a single median line m_g that has gradient g . Construct two parallel lines $t_U(g)$ and $t_D(g)$ that also have gradient g , but are tangent to $P_k(r, x, y)$ from above and below respectively. If the median line m_g intersects $P_k(r, x, y)$, as shown in Figure 4, then m_g must be in between $t_U(g)$ and $t_D(g)$.

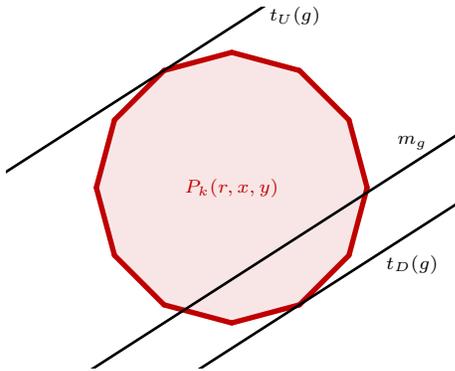


Figure 4: The relative positions of m_g , $t_U(g)$ and $t_D(g)$ if m_g intersects the k -sided regular polygon $P_k(r, x, y)$.

We will decide whether all median lines of gradient g are between $t_U(g)$ and $t_D(g)$, as this would immediately decide whether all median lines of gradient g intersects $P_k(r, x, y)$. We will solve this restricted decision problem by counting the number of points in V above $t_U(g)$ and the number of points in V below $t_D(g)$.

Let $t_U^+(g)$ be the number of points in V that are above $t_U(g)$, and similarly $t_D^-(g)$ for the points in V below $t_D(g)$. Suppose that $t_U^+(g) < n/2$ and $t_D^-(g) < n/2$. Then there cannot be a median line of gradient g above $t_U(g)$ or below $t_D(g)$, since one side of the median line, in particular the side that contains the polygon, will have more than $n/2$ points. Hence, if $t_U^+(g) < n/2$ and $t_D^-(g) < n/2$, then all median lines of gradient g must be between $t_U(g)$ and $t_D(g)$.

Conversely, suppose that all median lines of gradient g are between $t_U(g)$ and $t_D(g)$. Then if $t_U^+(g) \geq n/2$, we can move $t_U(g)$ continuously upwards until it becomes a median line, which is a contradiction. So in this case, we know $t_U^+(g) < n/2$ and $t_D^-(g) < n/2$.

In summary, we have transformed the decision problem into one that does not have median lines in its statement: All median lines intersect $P_k(r, x, y)$ if for all gradients g , the pair of inequalities $t_U^+(g) < n/2$ and $t_D^-(g) < n/2$ hold.

We present a sweep line algorithm that computes whether the pair of inequalities hold for all gradients g . Let t be an arbitrary line tangent to the polygon $P_k(r, x, y)$, and define t^+ to be the open halfplane that has t as its boundary and does not include the polygon $P_k(r, x, y)$. Then all median lines intersect $P_k(r, x, y)$ if and only if for all positions of t , the open halfplane t^+ contains less than $n/2$ points.

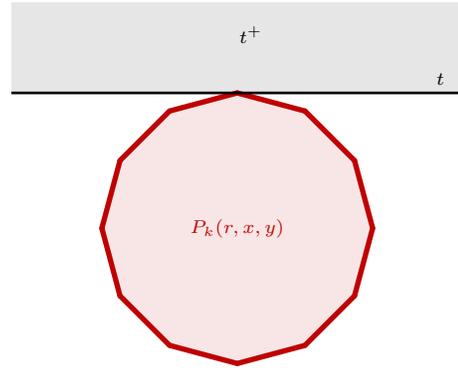


Figure 5: The rotating sweepline t and the open halfplane t^+ .

The tangent line t is a clockwise rotating sweep line and the invariant maintained by the sweep line algorithm is the number of points of V inside the region t^+ . Take any tangent line t_0 to be the starting line, and calculate the number of points in t_0^+ . From here, define an event to be when the line t passes through a point. There are two events for each point outside $P_k(r, x, y)$; there is one event for when the point enters the region t^+ , and one for when it exits. There are no events for points of V that lie inside $P_k(r, x, y)$. The unsorted set of event points can be computed by applying Subroutine 1 to each of point in V .

We sort the set of event points in a clockwise fashion. If we consider only two voters, their associated events can be sorted using Subroutine 2. We can extend this to sort the associated events of all voters with any standard comparison-based sorting algorithm, for example Merge sort.

Once the sorted set of events is computed, we process the events in order. At each new event we maintain our invariant, the number of points inside the region t^+ . This value increases by one at “entry” events and decreases by one at “exit” events. Finally, we return whether our invariant remained less than $n/2$ at all events.

The running time analysis for the algorithm is as follows. Computing the points outside $P_k(r, x, y)$ takes $O(\log k)$ time per point by Subroutine 1, so in total this

takes $O(n \log k)$ time. Computing the sorted order of the event points takes $O(\log k)$ time per comparison by Subroutine 2, which adds up to $O(n \log n \cdot \log k)$ time. Processing the sorted event points takes $O(n)$ time. Adding these gives the stated bound. \square

3 Parametric Search

Parametric search is a powerful yet complex technique for solving optimisation problems. The two steps involved in this technique are, first, to design a decision algorithm, and second, to convert the decision algorithm into an optimisation algorithm.

For example, our parameter space is $(r, x, y) \in \mathbb{R}^3$, our decision algorithm is stated in Theorem 1, and our optimisation objective is to minimise $r \in \mathbb{R}^+$.

Preliminaries

Megiddo's (1983) states the requirements for converting the decision algorithm into an optimisation algorithm. First, let us introduce some notation. Let \mathbb{R}^d be a parameter space, let $\lambda \in \mathbb{R}^d$ be a parameter and let $D(\lambda)$ be a decision problem that either evaluates to true or false. Then the first requirement is for the decision problem $D(\lambda)$.

Property 1. *The set of parameters $\{\lambda \in \mathbb{R}^d : D(\lambda)\}$ that satisfies the decision problem is convex.*

Convexity guarantees that the optimisation algorithm finds the global optimum.

The second property of the technique relates to the decision algorithm. Let $\mathcal{A}(\lambda)$ be a comparison-based decision algorithm that computes $D(\lambda)$. Let $C(\lambda)$ be any comparison in the comparison-based decision algorithm $\mathcal{A}(\lambda)$. The comparison $C(\lambda)$ is said to have an associated critical hyperplane in \mathbb{R}^d if the result of the comparison is linearly separable with respect to $\lambda \in \mathbb{R}^d$. Formally, suppose that the comparison $C(\lambda)$ evaluates to either $>$, $=$ or $<$. Then we say that the $(d - 1)$ -dimensional hyperplane $H \subset \mathbb{R}^d$ is the associated critical hyperplane of $C(\lambda)$ if C evaluates to $>$, $=$ or $<$ if and only if λ is above, on, or below H respectively. The comparisons of the decision algorithm must satisfy the following property.

Property 2. *Every comparison $C(\lambda)$ in the comparison-based decision algorithm $\mathcal{A}(\lambda)$ either (i) does not depend on λ , or (ii) has an associated critical hyperplane in \mathbb{R}^d .*

This requirement allows us to compute a large set of critical hyperplanes that determines the result of $\mathcal{A}(\lambda)$. Moreover, the optimum must lie on one of these critical hyperplanes, since the result of $\mathcal{A}(\lambda)$ locally changes sign at the optimum. The new search space now has dimension $d - 1$ instead of dimension d , and we can recursively apply this procedure to reduce the dimension further. For details see (Agarwal and Sharir 1998).

The final property speeds up the parametric search.

Property 3. *The decision algorithm has an efficient parallel algorithm.*

If the decision algorithm $\mathcal{A}(\lambda)$ runs in T_s time and runs on P processors in T_p parallel steps, then the paramet-

ric search over $\lambda \in \mathbb{R}^d$ runs in $O(T_p P + T_s(T_p \log P)^d)$ time (Agarwal and Sharir 1998).

Applying the technique

To apply the parametric search technique, we show that our decision problem $D_{k,V}(r, x, y)$ satisfies Properties 1-3.

Lemma 1. *Given an integer $k \geq 3$ and a set V of n points in the plane, the set of parameters $\{(r, x, y) : D_{k,V}(r, x, y)\}$ that satisfies the decision problem is convex.*

Proof. Suppose we are given a convex combination $\lambda_3 = \alpha\lambda_1 + (1 - \alpha)\lambda_2$ of the two parameters $\lambda_1, \lambda_2 \in \mathbb{R}^3$. Then the polygon $P_k(\lambda_3)$ is a convex combination of the polygons $P_k(\lambda_1)$ and $P_k(\lambda_2)$. It is easy to check that if a line m intersects both $P_k(\lambda_1)$ and $P_k(\lambda_2)$, then the line m must also intersect the convex combination $P_k(\lambda_3)$.

Now assume that both $D_{k,V}(\lambda_1)$ and $D_{k,V}(\lambda_2)$ are true. Then for any median line m both $P_k(\lambda_1)$ and $P_k(\lambda_2)$ intersect m . By the observation above, the convex combination $P_k(\lambda_3)$ must also intersect m . Repeating this fact for all median lines implies that $P_k(\lambda_3)$ intersects all median lines of V . So $D_{k,V}(\lambda_3)$ is true whenever $D_{k,V}(\lambda_1)$ and $D_{k,V}(\lambda_2)$ are true. Therefore, the set of parameters $\{(r, x, y) \subseteq \mathbb{R}^3 : D_{k,V}(r, x, y)\}$ is convex. \square

Lemma 2. *Every comparison in the decision algorithm in Theorem 1 either (i) does not depend on (r, x, y) , or (ii) has an associated critical hyperplane in \mathbb{R}^3 .*

Proof. Theorem 1 consists of three steps, computing the points outside the polygon, computing the event order, and processing the events. For the first two steps, the comparisons do depend on (r, x, y) and have associated critical hyperplanes. We defer the proof of this claim to Sections 5 and 6 respectively. For the third step, the comparisons do not depend on (r, x, y) but rather the event order, so there is no requirement that comparisons have critical hyperplanes. \square

Lemma 3. *The decision algorithm in Theorem 1 has an efficient parallel algorithm that runs on $O(n)$ processors and takes $O(\log n \cdot \log k)$ parallel steps per processor.*

Proof. Given $O(n)$ processors, we decide which points are outside the polygon in parallel by assigning a processor to each point. By Subroutine 1, this takes $O(\log k)$ parallel steps per processor. We compute the event order in parallel using Preparata's sorting scheme (Preparata 1978). Each processor requires $O(\log n)$ calls to Subroutine 2, so in total, each processor requires $O(\log n \cdot \log k)$ parallel steps. Finally, processing the events generates no critical hyperplanes, so this step does not require parallelisation. \square

Now we combine Properties 1-3 with Megiddo's result to obtain an optimisation algorithm for the smallest, regular, k -sided polygon $P_k(r, x, y)$ that intersects all median lines.

Theorem 2. *Given a set V of n points in the plane, there is an $O(n \log^7 n \cdot \log^4 k)$ time algorithm to compute the minimum r such that $D_{k,V}(r, x, y)$ is true for some regular, k -sided polygon $P_k(r, x, y)$.*

Proof. Megiddo’s multidimensional parametric search implies that there is an efficient optimisation algorithm. It only remains to show the running time of the technique.

The parallel algorithm runs on $P = O(n)$ processors in $T_p = O(\log n \cdot \log k)$ parallel steps, whereas the decision algorithm runs in $T_s = O(n \log n \cdot \log k)$ time. The dimension d of the parameter space is three. The running time of multidimensional parametric search is $O(T_p P + T_s (T_p \log P)^d)$ (Agarwal and Sharir 1998). Substituting our values into the above formula yields the required bound. \square

4 Computing Critical Hyperplanes

The only requirement left to check is Property 2 for the comparisons in the comparison-based subroutines, that is, Subroutine 1 and Subroutine 2. Before launching into the analysis of the two subroutines, we first prove a tool. We will use the tool repeatedly in the next two sections to simplify checking Property 2.

Lemma 4. *Let gradient $g \in \mathbb{R}$, point $p \in \mathbb{R}^2$ and vector $v \in \mathbb{R}^2$ be given, and let $(r, x, y) \in \mathbb{R}^3$ be a variable parameter. Let $L_{g,v}(r, x, y)$ be the line of gradient g through the point $(x, y) + r \cdot v$. Then p is above, on, or below $L_{g,v}(r, x, y)$ if and only if the point (r, x, y) is above, on, or below its associated critical hyperplane $H_{p,g,v}$.*

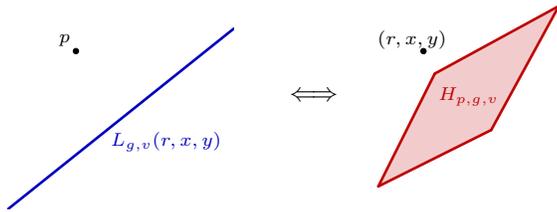


Figure 6: Point p is above $L_{g,v}(r, x, y)$ if and only if parameter (r, x, y) is above $H_{p,g,v}$.

Proof. Let point $p = (p_x, p_y)$ and vector $v = (v_x, v_y)$. Now, (p_x, p_y) is above the line through (q_x, q_y) of gradient g if $(p_x - q_x) - g \cdot (p_y - q_y) > 0$. Substituting the point $(q_x, q_y) = (x, y) + r \cdot (v_1, v_2)$, we get the inequality

$$(p_2 - y - rv_2) - g \cdot (p_2 - x - rv_2) > 0.$$

This inequality can be rearranged into the form $ax + by + cr + d > 0$, where

$$a = g, b = -1, c = (gv_1 - v_2), d = p_2 - gp_1.$$

In this form, we can see that the inequality is satisfied if and only if (r, x, y) lies above the hyperplane $H_{p,g,v} := (ax + by + cr + d = 0)$, where a, b, c, d are given above. Hence, the two conditions, p above a line and (r, x, y) above a hyperplane, can be decided with the same inequality, which completes the proof. \square

Now we are ready to address the subroutines.

5 Subroutine 1

Subroutine 1 decides whether a given point p is outside the k -sided, regular polygon $P_k(r, x, y)$. We present an $O(\log k)$ time comparison-based algorithm and show that Property 2 holds.

Lemma 5. *Subroutine 1 has an $O(\log k)$ time comparison-based algorithm, and comparisons in the algorithm that depend on the parameter (r, x, y) each have an associated critical hyperplane.*

Proof. We partition the polygon $P_k(r, x, y)$ into k triangles, and decide which partition the point p is in, if it indeed is in any of these partitions. For $1 \leq i \leq k$, the i^{th} partition of $P_k(r, x, y)$ is the triangle joining the i^{th} vertex, the $(i + 1)^{\text{th}}$ vertex and the center of $P_k(r, x, y)$. Figure 7 shows the i^{th} partition of $P_k(r, x, y)$.

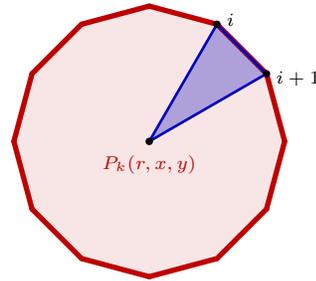


Figure 7: The i^{th} partition of $P_k(r, x, y)$.

Assume for now that the point p is indeed in the polygon $P_k(r, x, y)$ and hence in one of the k partitions. We decide whether p is in the i^{th} partition for some $i \leq j$, or for some $i > j$, and perform a binary search for the index i . This can be done by deciding if the point p is above, on, or below the line joining the center of $P_k(r, x, y)$ and its j^{th} vertex. The comparison depends on (r, x, y) , so we must compute its associated critical hyperplane using Lemma 4. Let $P_k(1, 0, 0)$ be the k -sided polygon of radius 1 and centered at the origin. Then set g to be the gradient of the line joining the center to the i^{th} vertex of $P_k(1, 0, 0)$, and vector $v = 0$ in Lemma 4 to obtain the associated critical hyperplane.

We have searched for the partition that p is in if it is indeed in $P_k(r, x, y)$. Hence, it only remains to decide whether p is indeed in that partition. This requires a constant number of comparisons, each of which depend on (r, x, y) . We have already computed associated critical hyperplanes for two of the sides. The last side joins two adjacent vertices of the polygon $P_k(r, x, y)$. Set g to be the gradient of the i^{th} side of polygon $P_k(1, 0, 0)$, and the vector v to be the i^{th} vertex of $P_k(1, 0, 0)$, to obtain the final associated critical hyperplane.

The running time is dominated by the binary search for the i^{th} partition, which takes $O(\log k)$ time. \square

6 Subroutine 2

Subroutine 2 computes the relative clockwise order of four tangent lines drawn from two points to polygon $P_k(r, x, y)$.

Lemma 6. *Subroutine 2 has an $O(\log k)$ -time comparison-based algorithm, and comparisons in the algorithm that depend on the parameter (r, x, y) each have an associated critical hyperplane.*

Proof. Draw two lines t_i, t_j tangent to $P_k(r, x, y)$ and parallel to pq , and let the points of tangency be vertex i and vertex j . If there are multiple points of tangency then choose any such point. Then without loss of generality, set ij to be horizontal, and assume further that p has a larger y coordinate than q . Then the t_i, t_j and ij partition the plane into the four regions, as shown in Figure 8. Region L is left of both tangents, R is right of both tangents, U is between the tangents and above ij , and D is between the tangents and below ij .

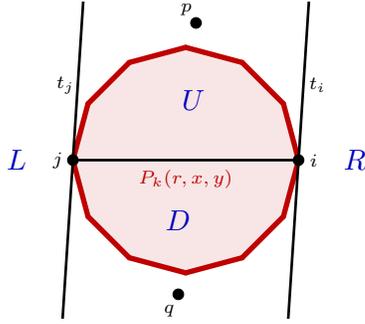


Figure 8: The lines t_i, t_j, ij partition the plane into regions L, R, U, D .

Then the relative clockwise order of the four lines drawn from p and q are determined by which of the four regions L, R, U or D the points p and q are located. See Figure 9.

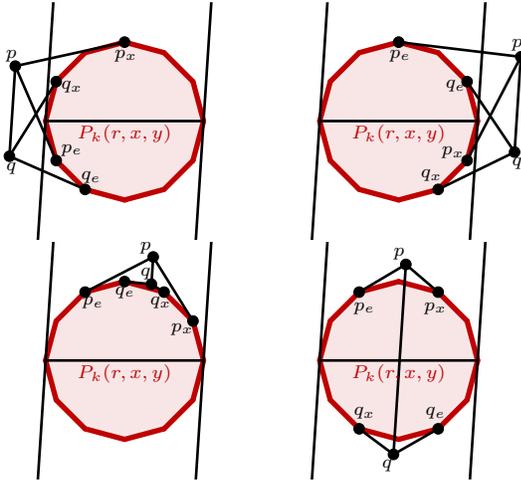


Figure 9: The relative orders shown for when (i) $p, q \in L$, (ii) $p, q \in R$, (iii) $p, q \in U$ and (iv) $p \in U, q \in D$.

Five cases follows. Let p_e and p_x points of tangency from p such that the points p_e, p, p_x are in clockwise order. If p, q are in the same region, then the containing re-

gion L, R, U , and D correspond to the relative clockwise orders $q_e p_e q_x p_x, p_e q_e p_x q_x, p_e q_e q_x p_x$, and $q_e p_e p_x q_x$ respectively. If p, q are in different regions, then they must be in U and D respectively, and the relative order is $p_e p_x q_e q_x$. The proof for case analysis for the five cases is omitted, but the diagrams in Figure 9 may be useful for the reader.

The running time of the algorithm is as follows. Given the gradient of pq , there is an $O(\log k)$ time algorithm to binary search the gradients of the sides of $P_k(r, x, y)$ to compute the vertices i and j . Then the remainder of the algorithm takes constant time: rotating the diagram so that ij is horizontal, deciding whether p or q has a larger y coordinate, and computing the region L, R, U, D that points p, q are in.

The proof of existence of critical hyperplanes is as follows. Since the gradients of pq and sides of P_k do not depend on (r, x, y) , computing i and j generates no critical hyperplanes. Similarly, rotating the diagram so that ij is horizontal and then deciding which of p or q have larger y coordinates also generates no critical hyperplanes. It only remains to decide which of the four regions L, R, U, D the point p , and respectively q , is in. Set g to the gradient of pq and vector v to be the i^{th} vertex of $P_k(1, 0, 0)$ in Lemma 4 to decide if p is to the left of the tangent through i . Do so similarly for j to decide if p is to the right of the tangent through j . Finally, set g to the gradient of ij and vector v to be either the i^{th} or j^{th} vertex of $P_k(1, 0, 0)$ to decide if p is above the chord ij . \square

Checking that Property 2 holds for the comparison-based subroutines, Subroutine 1 and Subroutine 2, completes the proof to Theorem 2. In the final section we will prove that Theorem 2 implies that we have an efficient algorithm for computing the yolk in the \mathcal{L}_1 and \mathcal{L}_∞ metrics, and an efficient approximation algorithm for the \mathcal{L}_2 metric.

7 Computing the Yolk in $\mathcal{L}_1, \mathcal{L}_2$, and \mathcal{L}_∞

It remains to show that our general problem for $P_k(r, x, y)$ implies the results as claimed in the introduction.

Theorem 3. *Given a set V of n points in the plane, there is an $O(n \log^7 n)$ time algorithm to compute the yolk of V in the \mathcal{L}_1 and \mathcal{L}_∞ metrics.*

Proof. Setting $k = 4$ in Theorem 2 gives an algorithm to compute the smallest $P_4(r, x, y)$ that intersects all median lines of V in $O(n \log^7 n)$ time. This rotated square coincides with yolk in the \mathcal{L}_1 metric, refer to Figure 2 and Definition 1.

Computing the yolk in the \mathcal{L}_∞ metric requires one extra step. Rotate the points of V by 45° clockwise, compute the smallest $P_4(r, x, y)$, and then rotate the square $P_4(r, x, y)$ back 45° anticlockwise to obtain the yolk in the \mathcal{L}_∞ metric. \square

Theorem 4. *Given a set V of n points in the plane and an $\varepsilon > 0$, there is an $O(n \log^7 n \cdot \log^4 \frac{1}{\varepsilon})$ time algorithm to compute a $(1 + \varepsilon)$ -approximation of the yolk in the \mathcal{L}_2 metric.*

Proof. Setting $k \approx \pi \cdot (1 + \frac{1}{\varepsilon})$ in Theorem 2 gives an algorithm to compute the smallest $P_k(r, x, y)$ that intersects all

median lines of V in the desired running time. It suffices to show that for this parameter set (r, x, y) , the disk centered at (x, y) with radius r is a $(1 + \varepsilon)$ -approximation for the yolk in the \mathcal{L}_2 metric.

First, note that $P_k(r, x, y)$ intersects all median lines, and $B(r, x, y)$ encloses $P_k(r, x, y)$, so the disk must also intersect all median lines of V . Hence, it suffices to show that the radius r of $B(r, x, y)$ satisfies $r \leq (1 + \varepsilon) \cdot r_2$, where r_2 is the radius of the true yolk in the \mathcal{L}_2 metric.

Let the yolk in the \mathcal{L}_2 metric be the disk $B(r_2, x_2, y_2)$. Consider the regular, k -sided polygon $P_k(r_2 \cdot \sec \frac{\pi}{k}, x_2, y_2)$, so that by construction, all sides of this polygon are tangent to $B(r_2, x_2, y_2)$.

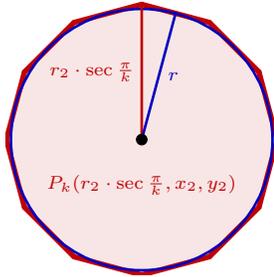


Figure 10: The polygon $P_k(r_2 \cdot \sec \frac{\pi}{k}, x_2, y_2)$ is externally tangent to the disk $B(r_2, x_2, y_2)$.

Now since $B(r_2, x_2, y_2)$ is the \mathcal{L}_2 yolk, it intersects all median lines and so does its enclosing polygon $P_k(r_2 \cdot \sec \frac{\pi}{k}, x_2, y_2)$. By the minimality of $P_k(r, x, y)$, we get $r \leq \sec \frac{\pi}{k} \cdot r_2$. But for $\theta \in [0, \frac{\pi}{3}]$, we have $\sec \theta \leq \frac{1}{1 - \theta}$. So,

$$\sec \frac{\pi}{k} \leq \frac{1}{1 - \frac{\pi}{k}} \leq 1 + \varepsilon,$$

which implies that $r \leq (1 + \varepsilon) \cdot r_2$, as required. \square

8 Concluding Remarks

Cole's (1987) extension to parametric search states that the running time of the parametric search may be reduced if certain comparisons are delayed. This is a direction for further research that could potentially improve the running time of our algorithms.

An open problem is whether one can compute the yolk in higher dimensions without precomputing all median hyperplanes. Avoiding the computation of median hyperplanes yields even greater benefits as less is known about bounds on the number of median hyperplanes in higher dimensions.

Similarly, our approximation algorithm for the \mathcal{L}_2 yolk in the plane is optimal up to polylogarithmic factors, however, it is an open problem as to whether there is a near-linear time exact algorithm. Our attempts to apply Megiddo's parametric search technique to the \mathcal{L}_2 yolk have been unsuccessful so far.

Finally, there are other solution concepts in computational spatial voting that currently lack efficient algorithms. The shortcomings of existing algorithms are: for the Shapley Owen power score there is only an approximate algorithm (Godfrey 2005), for the Finagle point only regular

polygons have been considered (Wuffle et al. 1989) and for the ε -core only a membership algorithm exists (Tovey 2010b). Since these problems have a close connection to either median lines or minimal radius, we suspect that Megiddo's parametric search technique could also be useful for these problems.

References

- Agarwal, P. K., and Sharir, M. 1998. Efficient algorithms for geometric optimization. *ACM Computing Surveys* 30(4):412–458.
- Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI 2015*, 777–783.
- Berg, M. D.; Gudmundsson, J.; and Mehr, M. 2016. Faster algorithms for computing plurality points. In *32nd International Symposium on Computational Geometry, SoCG 2016*, 32:1–32:15.
- Black, D. 1948. On the rationale of group decision-making. *Journal of Political Economy* 56(1):23–34.
- Cole, R. 1987. Slowing down sorting networks to obtain faster sorting algorithms. *Journal of the ACM* 34(1):200–208.
- Dey, T. K. 1997. Improved bounds on planar k -sets and k -levels. In *38th Annual Symposium on Foundations of Computer Science, FOCS 1997*, 156–161.
- Downs, A. 1957. An economic theory of political action in a democracy. *Journal of Political Economy* 65(2):135–150.
- Enelow, J. M., and Hinich, M. J. 1984. *The spatial theory of voting: An introduction*. CUP Archive.
- Feld, S. L., and Grofman, B. 1990. A theorem connecting Shapley-Owen power scores and the radius of the yolk in two dimensions. *Social Choice and Welfare* 7(1):71–74.
- Feld, S. L.; Grofman, B.; Hartly, R.; Kilgour, M.; Miller, N.; et al. 1987. The uncovered set in spatial voting games. *Theory and Decision* 23(2):129–155.
- Feld, S. L.; Grofman, B.; and Miller, N. 1988. Centripetal forces in spatial voting: on the size of the yolk. *Public Choice* 59(1):37–50.
- Feld, S. L.; Grofman, B.; and Miller, N. R. 1989. Limits on agenda control in spatial voting games. *Mathematical and Computer Modelling* 12(4-5):405–416.
- Godfrey, J. 2005. Computation of the Shapley-Owen power index in two dimensions. In *4th Annual workshop, University of Warwick*, 20–22.
- Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC 2017*, 287–304.
- Koehler, D. H. 1990. The size of the yolk: computations for odd and even-numbered committees. *Social Choice and Welfare* 7(3):231–245.

- McKelvey, R. D. 1976. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory* 12(3):472–482.
- McKelvey, R. D. 1986. Covering, dominance, and institution-free properties of social choice. *American Journal of Political Science* 283–314.
- Megiddo, N. 1983. Applying parallel computation algorithms in the design of serial algorithms. *Journal of the ACM* 30(4):852–865.
- Miller, N. R.; Grofman, B.; and Feld, S. L. 1989. The geometry of majority rule. *Journal of Theoretical Politics* 1(4):379–406.
- Miller, N. R. 1980. A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting. *American Journal of Political Science* 68–96.
- Ordeshook, P. C. 1993. The spatial analysis of elections and committees: Four decades of research. Technical report, California Institute of Technology, Division of the Humanities and Social Sciences.
- Plott, C. R. 1967. A notion of equilibrium and its possibility under majority rule. *The American Economic Review* 787–806.
- Poole, K. T., and Rosenthal, H. 1984. The polarization of American politics. *The Journal of Politics* 46(4):1061–1079.
- Poole, K. T., and Rosenthal, H. 1991. Patterns of congressional voting. *American Journal of Political Science* 228–278.
- Poole, K. T., and Rosenthal, H. 2001. *D*-nominate after 10 years: A comparative update to congress: A political-economic history of roll-call voting. *Legislative Studies Quarterly* 5–29.
- Preparata, F. P. 1978. New parallel-sorting schemes. *IEEE Transactions on Computers* (7):669–673.
- Schofield, N.; Miller, G.; and Martin, A. 2003. Critical elections and political realignments in the USA: 1860–2000. *Political Studies* 51(2):217–240.
- Schofield, N. 2004. Equilibrium in the spatial ‘valence’ model of politics. *Journal of Theoretical Politics* 16(4):447–481.
- Schofield, N. 2007. *The spatial model of politics*. Routledge.
- Skowron, P. K., and Elkind, E. 2017. Social choice under metric preferences: Scoring rules and STV. In *Thirty-First AAAI Conference on Artificial Intelligence, AAAI 2017*, 706–712.
- Stone, R. E., and Tovey, C. A. 1992. Limiting median lines do not suffice to determine the yolk. *Social Choice and Welfare* 9(1):33–35.
- Tóth, G. 2000. Point sets with many *k*-sets. In *16th Annual Symposium on Computational Geometry, SoCG 2000*, 37–42.
- Tovey, C. A. 1992. A polynomial-time algorithm for computing the yolk in fixed dimension. *Mathematical Programming* 57(1):259–277.
- Tovey, C. A. 1993. Some foundations for empirical study in the Euclidean spatial model of social choice. In *Political economy: institutions, competition, and representation: Proceedings of the Seventh International Symposium in Economic Theory and Econometrics*, 175. Cambridge Univ Pr.
- Tovey, C. A. 2010a. The almost surely shrinking yolk. *Mathematical Social Sciences* 59(1):74–87.
- Tovey, C. A. 2010b. A finite exact algorithm for epsilon-core membership in two dimensions. *Mathematical Social Sciences* 60(3):178–180.
- Tovey, C. A. 2011. The Finagle point and the epsilon-core: a comment on Bräuningner’s proof. *Journal of Theoretical Politics* 23(1):135–139.
- Wuffle, A.; Feld, S. L.; Owen, G.; and Grofman, B. 1989. Finagle’s law and the finagle point, a new solution concept for two-candidate competition in spatial voting games without a core. *American Journal of Political Science* 348–375.