

Faster algorithms for computing plurality points

Citation for published version (APA):

de Berg, M. T., Gudmundsson, J., & Mehr, M. (2018). Faster algorithms for computing plurality points. *ACM Transactions on Algorithms*, 14(3), [36]. <https://doi.org/10.1145/3186990>

Document license:

TAVERNE

DOI:

[10.1145/3186990](https://doi.org/10.1145/3186990)

Document status and date:

Published: 01/07/2018

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Faster Algorithms for Computing Plurality Points

MARK DE BERG, TU Eindhoven, The Netherlands

JOACHIM GUDMUNDSSON, University of Sydney, Australia

MEHRAN MEHR, TU Eindhoven, The Netherlands

Let V be a set of n points in \mathbb{R}^d , which we call voters. A point $p \in \mathbb{R}^d$ is preferred over another point $p' \in \mathbb{R}^d$ by a voter $v \in V$ if $\text{dist}(v, p) < \text{dist}(v, p')$. A point p is called a plurality point if it is preferred by at least as many voters as any other point p' .

We present an algorithm that decides in $O(n \log n)$ time whether V admits a plurality point in the L_2 norm and, if so, finds the (unique) plurality point. We also give efficient algorithms to compute a minimum-cost subset $W \subset V$ such that $V \setminus W$ admits a plurality point, and to compute a so-called minimum-radius plurality ball.

Finally, we consider the problem in the personalized L_1 norm, where each point $v \in V$ has a preference vector $\langle w_1(v), \dots, w_d(v) \rangle$ and the distance from v to any point $p \in \mathbb{R}^d$ is given by $\sum_{i=1}^d w_i(v) \cdot |x_i(v) - x_i(p)|$. For this case we can compute in $O(n^{d-1})$ time the set of all plurality points of V . When all preference vectors are equal, the running time improves to $O(n)$.

CCS Concepts: • **Theory of computation** → **Computational geometry**;

Additional Key Words and Phrases: Computational geometry, computational social choice, voting theory, plurality points, Condorcet points, Voronoi games

ACM Reference format:

Mark de Berg, Joachim Gudmundsson, and Mehran Mehr. 2018. Faster Algorithms for Computing Plurality Points. *ACM Trans. Algorithms* 14, 3, Article 36 (June 2018), 23 pages.

<https://doi.org/10.1145/3186990>

1 INTRODUCTION

We study computational problems concerning plurality points, a concept arising in social choice and voting theory, defined as follows. Let V be a set of n voters and let C be a space of possible choices. Each voter $v \in V$ has a utility function indicating how much v likes a certain choice; i.e., the utility function of v determines for any two choices from C which one is preferred by v or whether both choices are equally preferable. A (weak) plurality point is now defined as a choice $p \in C$ such that no alternative $p' \in C$ is preferred by more voters.

MdB and MM are supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003 and 022.005.025, respectively. JG is supported under the Australian Research Council's Discovery Projects funding scheme (project number DP150101134).

A preliminary version of this article appeared in the proceedings of the *32nd International Symposium on Computational Geometry*, 2016.

Authors' addresses: M. de Berg, TU Eindhoven, Department of Computer Science, The Netherlands; email: mdberg@win.tue.nl; J. Gudmundsson, University of Sydney, School of IT, Australia; email: joachim.gudmundsson@sydney.edu.au; M. Mehr, TU Eindhoven, Department of Computer Science, The Netherlands; email: mmehr@tue.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 1549-6325/2018/06-ART36 \$15.00

<https://doi.org/10.1145/3186990>

When there are different issues on which the voters can decide, then the space C becomes a multidimensional space. This has led to the study of plurality points in the setting where $C = \mathbb{R}^d$ and each voter has an ideal choice that is a point in \mathbb{R}^d . To simplify the presentation, from now on we will not distinguish the voters from their ideal choice and so we view each voter $v \in V$ as being a point in \mathbb{R}^d , the so-called *spatial model* in voting theory [19]. Thus, the utility of a point $p \in \mathbb{R}^d$ for a voter v is inversely proportional to $\text{dist}(v, p)$, the distance from v to p under a given distance function, and v prefers a point p over a point p' if $\text{dist}(v, p) < \text{dist}(v, p')$. Now a point $p \in \mathbb{R}^d$ is a plurality point if for any point $p' \in \mathbb{R}^d$ we have $|\{v \in V : \text{dist}(v, p) < \text{dist}(v, p')\}| \geq |\{v \in V : \text{dist}(v, p') < \text{dist}(v, p)\}|$.

Plurality points and related concepts were already studied in the 1970s in voting theory [10, 14, 15, 19, 20, 22]. McKelvey and Wendell [19] define three different notions of plurality points—majority Condorcet, plurality Condorcet, and majority core—and for each notion they define a weak and a strong variant. Under certain assumptions on the utility functions, which are satisfied for the L_2 norm, the three notions are equivalent. Thus, for the L_2 norm we only have two variants: weak plurality points (which should be at least as popular as any alternative) and strong plurality points (which should be strictly more popular than any alternative). We focus on weak plurality points, since they are more challenging from an algorithmic point of view. From now on, whenever we speak of plurality points we refer to weak plurality points.

Plurality points represent a stable choice with respect to the opinions of the voters. One can also look at the concept from the viewpoint of competitive facility location. Here one player wants to place a facility in the space C such that he or she always wins at least as many clients (voters) as his or her competitor, no matter where the competitor places his or her facility. Competitive facility location problems have been studied widely in a discrete setting, where the clients and the possible locations for the facilities are nodes in a network; see the survey by Kress and Pesch [17]. Competitive facility location has also been studied in a geometric, continuous setting under the name Voronoi games [1, 6]. Here one is given a region R in \mathbb{R}^2 , say, the unit square, and the goal is to win the maximum area within R . In other words, the set V of voters is no longer finite, but we have $V = C = R$. The plurality-point problem in a geometric space lies in between the network setting and the fully continuous setting: the space C of choices is \mathbb{R}^d , but the set V of voters is finite.

When the L_2 norm defines the distance between voters and potential plurality points, then plurality points can be defined in terms of Tukey depth [21]. The *Tukey depth* of a point $p \in \mathbb{R}^d$ with respect to a given set V of n points is defined as the minimum number of points from V lying in any closed half-space containing p . A point of maximum Tukey depth is called a *Tukey median*. It is known that for any set V , the depth of the Tukey median is at least $\lceil n/(d+1) \rceil$ and at most $\lfloor n/2 \rfloor$. Wu et al. [24] showed that a point $p \in \mathbb{R}^d$ is a plurality point in the L_2 norm if and only if any open half-space with p on its boundary contains at most $n/2$ voters. This is equivalent to saying that the Tukey depth of p is $\lceil n/2 \rceil$. They used this observation to present an algorithm that decides in $O(n^{d-1} \log n)$ time if a plurality point exists for a given set V of n voters in \mathbb{R}^d . A slightly better result can be obtained using a randomized algorithm by Chan [4], which computes a Tukey median (together with its depth) in $O(n \log n + n^{d-1})$ time.

As is clear from the relation to Tukey depth, a plurality point in the L_2 norm does not always exist. In fact, the set V of voters must, in a certain sense, be highly symmetric to admit a plurality point. For example, when the number of voters is even, any line containing a plurality point p must contain the same number of voters on either side of p . This led Lin et al. [18] to study the *minimum-cost plurality problem*. Here each voter is assigned a cost, and the goal is to find a minimum-cost subset $W \subset V$ of voters such that if we ignore the voters in W —that is, if we consider $V \setminus W$ —then a plurality point exists. Lin et al. gave an $O(n^5 \log n)$ algorithm for the planar version of the problem; whether the problem in \mathbb{R}^3 can be solved in polynomial time was left as an open problem.

In the voting theory literature, plurality points in the L_1 norm have also been considered [19, 22, 23]. One advantage of the L_1 norm is that in \mathbb{R}^2 a plurality point always exists and can easily be found in $O(n)$ time: any point (p_x, p_y) such that p_x and p_y are medians of the multisets of x -coordinates and y -coordinates of the voters in V , respectively, is a plurality point. Unfortunately, this is no longer true when $d > 2$ [19, 22]. We are not aware of any existing algorithms for deciding whether a given set V in \mathbb{R}^d admits a plurality point in the L_1 norm.

Our Results. Currently the fastest algorithm for deciding whether a plurality point exists runs in $O(n \log n + n^{d-1})$ randomized time and actually computes a Tukey median. However, in the case of plurality points, we are only interested in the Tukey median if its depth is the maximum possible, namely, $\lceil n/2 \rceil$. Wu et al. [24] exploited this to obtain a deterministic algorithm, but their running time is $O(n^{d-1} \log n)$. This raises the question: can we decide whether a plurality point exists faster than by computing the depth of the Tukey median? We show that this is indeed possible: we present a deterministic algorithm that decides if a plurality point exists (and, if so, computes one) in $O(n \log n)$ time.¹

We then turn our attention to the minimum-cost plurality problem. We solve the open problem of Lin et al. [18] by presenting an algorithm that solves the problem in $O(n^4)$ time. Note that this even improves on the $O(n^5 \log n)$ running time for the planar case. We also consider the following problem for unit-cost voters in \mathbb{R}^d : given a parameter k , find a minimum-cost set W of size at most k such that $V \setminus W$ admits a plurality point, if such a set exists. Our algorithm for this case runs in $O(k^3 n \log n)$ time when $d = 2$ and in $O(k^5 \log k + k^3 n \log n)$ expected time when $d > 2$.

Ignoring some voters in order to have a plurality point is undesirable when almost all voters must be ignored. Instead of ignoring voters, we can work with *plurality balls*, as defined next. The idea is that if two points p and p' are very similar, then voters do not care much whether p or p' is chosen. Thus, we define a ball $b(p, r)$ centered at p and of radius r to be a plurality ball if the following holds: there is no point p' outside $b(p, r)$ that is preferred by more voters than p . Note that a plurality point is a plurality ball of zero radius. We show that in the plane, the minimum-radius plurality ball can be computed in $O(T(n))$ time, where $T(n)$ is the time needed to compute the $\lfloor n/2 \rfloor$ level in an arrangement of n lines.

Recall that the different dimensions represent different issues on which the voters can express their preferences. It is then natural to allow the voters to give different weights to these issues. This leads us to introduce what we call the *personalized L_1 norm*. Here each voter $v \in V$ has a *preference vector* $\langle w_1(v), \dots, w_d(v) \rangle$ of nonnegative weights that specifies the relative importance of the various issues. The distance of a point $p \in \mathbb{R}^d$ to a voter v is now defined as $\text{dist}_w(v, p) := \sum_{i=1}^d w_i(v) \cdot |x_i(v) - x_i(p)|$, where $x_i(\cdot)$ denotes the i th coordinate of a point. We present an algorithm that decides in $O(n^{d-1})$ time whether a set V of n voters admits a plurality point with respect to the personalized L_1 norm. For the special case when all preference vectors are identical—this case reduces to the normal case of using the L_1 norm—the running time improves to $O(n)$.

2 PLURALITY POINTS IN THE L_2 NORM

Let V be a set of n voters in \mathbb{R}^d . In this section, we show how to compute a plurality point for V with respect to the L_2 norm in $O(n \log n)$ time, if it exists. We start by proving several properties of the plurality point in higher dimensions, which generalize similar properties as Lin et al. [18] proved in \mathbb{R}^2 . These properties imply that if a plurality point exists, it is unique (unless all points are collinear). Our algorithm then consists of two steps: first it computes a single candidate point $p \in \mathbb{R}^d$, and then it decides if p is a plurality point.

¹Here and in the other bounds stated in the introduction we assume the dimension d is a fixed constant, to make it easier to compare our results to earlier work. In the theorems stated later in the article, we also analyze the dependency on d .




Fig. 1. A plurality point p with alternating property.

2.1 Properties of Plurality Points in the L_2 Norm

As remarked in the introduction, plurality points can be characterized as follows.

PROPOSITION 2.1 (WU ET AL. [24]). *A point p is a plurality point for a set V of n voters in \mathbb{R}^d with respect to the L_2 norm if and only if every open half-space with p on its boundary contains at most $n/2$ voters.*

Verifying the condition in Proposition 2.1 directly is not efficient. Hence, we will prove alternative conditions for a point p to be a plurality point in \mathbb{R}^d , which generalize the conditions Lin et al. [18] stated for the planar case. First, we define some concepts introduced by Lin et al.

Let V be a set of n voters in \mathbb{R}^d , and consider a point $p \in \mathbb{R}^d$. Let $L(p)$ be the set of all lines passing through p and at least one voter $v \neq p$. The point p partitions each line $\ell \in L(p)$ into two opposite rays, which we denote by $\rho(\ell)$ and $\bar{\rho}(\ell)$. (The point p itself is not part of these rays.) We say that a line $\ell \in L(p)$ is *balanced* if $|\rho(\ell) \cap V| = |\bar{\rho}(\ell) \cap V|$. When n is odd, p turns out to be a plurality point if and only if every line $\ell \in L(p)$ is balanced (which implies that we must have $p \in V$). When n is even, the situation is more complicated. Let $R(p)$ be the set of all rays $\rho(\ell)$ and $\bar{\rho}(\ell)$. Label each ray in $R(p)$ with an integer, which is the number of voters on the ray minus the number of voters from V on the opposite ray. Thus, a line ℓ is balanced if and only if its rays $\rho(\ell)$ and $\bar{\rho}(\ell)$ have label zero. Let $L^*(p)$ be the set of all unbalanced lines in $L(p)$ and let $R^*(p)$ be the corresponding set of rays. We now define the so-called *alternating property*, as introduced by Plott [20] and later by Lin et al. [18]. This property is restricted to the 2-dimensional setting, where we can order the rays in $R^*(p)$ around p . In this setting, the point p is said to have the *alternating property* if the following holds: the circular sequence of nonzero labels of the rays in $R^*(p)$, which we obtain when we visit the rays in $R^*(p)$ in clockwise order around p , alternates between labels $+1$ and -1 (see Figure 1). Note that if p has the alternating property, then the number of unbalanced lines must be odd.

Recall that for $0 \leq k \leq d$ a k -dimensional flat, or a k -flat for short, is the affine hull of $k + 1$ affinely independent points. For example, a line is a 1-flat and a hyperplane is a $(d - 1)$ -flat. The number k is called the dimension of the flat and is denoted by $\text{dim}(f)$. The following theorem gives necessary and sufficient conditions for a point p to be a plurality point for a set V of voters.

THEOREM 2.2. *Let V be a set of n voters in \mathbb{R}^d , with $d \geq 1$, and let p be an arbitrary point.*

- (a) *If n is odd, p is a plurality point if and only if $p \in V$ and every line in $L(p)$ is balanced.*
- (b) *If n is even and $p \notin V$, then p is a plurality point if and only if every line in $L(p)$ is balanced.*
- (c) *If n is even and $p \in V$, then p is a plurality point if and only if all unbalanced lines in $L(p)$ are contained in a single 2-dimensional flat f and p has the alternating property for the set $f \cap V$.*

For $d = 1$ the theorem is trivial, and for $d = 2$ —the condition in case (c) then simply states that p has the alternating property—the theorem was proved by Lin et al. [18] and by Plott [20]. Our contribution is the extension to higher dimensions. Before proving Theorem 2.2, we need the following lemma regarding the robustness of plurality points to dimension reduction.

LEMMA 2.3. *Let p be a plurality point for a set V in \mathbb{R}^d , with $d \geq 1$, and let f be any lower-dimensional flat containing p . Then p is a plurality point for $f \cap V$.*

PROOF. We prove the statement by induction on $d - \dim(f)$.

Base case: $d - \dim(f) = 1$; that is, f is a hyperplane.

Let f^+ and f^- denote the open half-spaces bounded by f , and assume without loss of generality that $|f^+ \cap V| \geq |f^- \cap V|$. Suppose for a contradiction that p is not a plurality point for $f \cap V$. Then there must be a $(d - 2)$ -flat $g \subset f$ containing p such that, within the $(d - 1)$ -dimensional space f , the number of voters lying strictly to one side of g is greater than $|f \cap V|/2$. Let $g^+ \subset f$ denote the part of f lying to this side of g . Now imagine rotating f around g by an infinitesimal amount. Let \hat{f} denote the rotated hyperplane, and \hat{f}^+ be the open half-space that contains all voters in f^+ . Observe that we can choose the direction of the rotation such that the voters in $g^+ \cap V$ end up in \hat{f}^+ . But then $|\hat{f}^+ \cap V| \geq |f^+ \cap V| + |g^+ \cap V| > |f^+ \cap V| + |f \cap V|/2 \geq n/2$, which contradicts the assumption that p is a plurality point.

Induction step: $d - \dim(f) > 1$.

Let h be a hyperplane that contains f . From the base case we know that p must be a plurality point for $h \cap V$. Hence, we can apply our induction hypothesis to the flat f in the $(d - 1)$ -dimensional space h to conclude that p must be a plurality point for $f \cap V$. \square

COROLLARY 2.4. *Let V be a set of voters in \mathbb{R}^d , for $d \geq 2$, that are not collinear. Then V has at most one plurality point.*

PROOF. Suppose for a contradiction that V has two distinct plurality points p_1 and p_2 . Let f be a 2-flat containing p_1 and p_2 , and a voter v not collinear with p_1 and p_2 . By Lemma 2.3, both p_1 and p_2 are plurality points for $f \cap V$. But this contradicts the result by Wu et al. [24] that any set of voters in the plane admits at most one plurality point. \square

Now we are ready to prove Theorem 2.2.

PROOF OF THEOREM 2.2. Since the case $d = 2$ was already proved by Lin et al. [18], and the case $d = 1$ is trivial, we assume $d \geq 3$.

- (a, \Leftarrow) and (b, \Leftarrow).** Let p be a point such that all lines in $L(p)$ are balanced. Consider an arbitrary open half-space h^+ whose bounding hyperplane h contains p , and let h^- be the open half-space opposite to h^+ . Since every line in $L(p)$ is balanced, there is a bijection from $h^+ \cap V$ to $h^- \cap V$. Hence, $|h^+ \cap V| \leq n/2$. Since this holds for any open half-space h^+ , the point p is a plurality point.
- (b, \Rightarrow).** Let $p \notin V$ be a plurality point, and consider a line $\ell \in L(p)$. By Lemma 2.3, the point p is also a plurality point on ℓ . Since $p \notin V$, this implies that ℓ is balanced.
- (a, \Rightarrow).** Assume n is odd, and let p be a plurality point. Suppose for a contradiction there is an unbalanced line $\ell \in L(p)$. First observe that we must have $p \in V$; otherwise, p is not a plurality point on ℓ , which contradicts Lemma 2.3. Let h be a hyperplane containing ℓ that contains no voters except those on ℓ , and assume without loss of generality that $|h^+ \cap V| \geq |h^- \cap V|$. Let $\rho(\ell)$ and $\bar{\rho}(\ell)$ denote the opposite rays along ℓ emanating

from p . Since ℓ is unbalanced, these rays do not contain the same number of voters, so we can assume $|\rho(\ell) \cap V| > |\bar{\rho}(\ell) \cap V|$. We consider two cases.

The first case is that $|h^+ \cap V| = |h^- \cap V|$. Since n is odd, this means that $|h \cap V|$ is odd, and so $|(h \cap V) \setminus \{p\}|$ is even. But then $|\rho(\ell) \cap V| \geq |\bar{\rho}(\ell) \cap V| + 2$. Now imagine rotating h infinitesimally around the $(d-2)$ -flat through p orthogonal to ℓ . For the resulting hyperplane \hat{h} we then have $|\hat{h}^+ \cap V| \geq |\hat{h}^- \cap V| + 2$. Since p is the only voter on \hat{h} , this implies $|\hat{h}^+ \cap V| > n/2$, which contradicts that p is a plurality point.

The second case is that $|h^+ \cap V| > |h^- \cap V|$. Now we can use a similar argument: this time we only have $|\rho(\ell) \cap V| \geq |\bar{\rho}(\ell) \cap V| + 1$, but we also have $|h^+ \cap V| > |h^- \cap V|$ and so we still have $|\hat{h}^+ \cap V| \geq |\hat{h}^- \cap V| + 2$, leading to the desired contradiction.

(c, \Leftarrow). Assume n is even and let p be a point such that all unbalanced lines in $L(p)$ are contained in a single 2-dimensional flat f and p has the alternating property for the set $f \cap V$. Since all the lines $L(p)$ not contained in f are balanced, this means $|f^+ \cap V| = |f^- \cap V|$, and so $|f \cap V|$ is even. Consider an arbitrary open half-space h^+ whose bounding hyperplane h contains p , and let h^- be the opposite open half-space. If h contains f , then all unbalanced lines lie in h and so $|h^+ \cap V| = |h^- \cap V|$, which implies $|h^+ \cap V| \leq n/2$. If h does not contain f , we can argue as follows. Let $\ell := h \cap f$. Since the theorem is true for $d = 2$ and as we have the alternating property on f and $|f \cap V|$ is even, we know that p is a plurality point for $f \cap V$. Hence, the number of voters on f on either side of ℓ is at most $|f \cap V|/2$. But then we have $|h^+ \cap V| \leq n/2$, because all voters not in f lie on balanced lines. We conclude that for any open half-space h^+ we have $|h^+ \cap V| \leq n/2$, and so p is a plurality point.

(c, \Rightarrow). Assume n is even and let p be a plurality point. We first argue that all unbalanced lines must lie on a single 2-flat. Assume for a contradiction that there are three unbalanced lines that do not lie on a common 2-flat. Let g be the 3-flat spanned by these lines, and let $L_g^*(p) \subset L^*(p)$ be the set of all unbalanced lines contained in g . Let $f_1 \subset g$ be a 2-flat not containing p and not parallel to any of the lines in $L_g^*(p)$. Each of the lines in $L_g^*(p)$ intersects f_1 in a single point, and these intersection points are not all collinear. According to the Sylvester-Gallai Theorem [13], this implies there is an ordinary line in f_1 , that is, a line containing exactly two of the intersection points. Thus, we have an ordinary 2-flat in g , that is, a flat f_2 containing exactly two lines from $L^*(p)$. This implies that $f_2 \cap V$ does not have the alternating property, and since we know by the result of Lin et al. that the theorem holds when $d = 2$, this implies that p is not a plurality point in f_2 . However, this contradicts Lemma 2.3.

We just argued that all unbalanced lines must lie on a single 2-flat f . By Lemma 2.3, the point p is a plurality point on f . Since the theorem holds for $d = 2$, we can conclude that $f \cap V$ has the alternating property. \square

2.2 Finding Plurality Points in the L_2 Norm

We now turn our attention to finding a plurality point. Our algorithm needs a subroutine for finding a *median hyperplane* for V , which is a hyperplane h such that $|h^+ \cap V| < n/2$ and $|h^- \cap V| < n/2$, where h^+ and h^- denote the two open half-spaces bounded by h . The following lemma is easy to prove.

LEMMA 2.5. *Let $v \in V$ be a voter that lies on a hyperplane h_0 such that no two voters strictly lie on opposite sides of h_0 (i.e., either h_0^+ does not contain voters or h_0^- does not contain voters). Then we can find a median hyperplane h containing v in $O(dn)$ time.*

PROOF. Assume without loss of generality that v is the origin and h_0 is the hyperplane $x_d = 0$. Define $V_0 := V \cap h_0$ and $V_1 := V \setminus V_0$. Note that all voters in V_1 lie to the same side of h_0 , say, in h_0^+ .

If $|V_1| < n/2$, then h_0 is a median hyperplane. Otherwise, let f_0 be the $(d-2)$ -flat $x_d = x_{d-1} = 0$, and note that f_0 contains v . We can now rotate a hyperplane h around f_0 until it satisfies the requirements. Next we make this precise.

Define n_0^- to be the number of voters in V_0 that lie (within the $(d-1)$ -dimensional space h_0) strictly on the negative side of f_0 . Project all voters in V_1 onto the hyperplane $h_1 : x_d = 1$, using v as the center of projection. Let X be the multiset of all x_{d-1} -coordinates of the projected voters. Compute an element $x^* \in X$ such that

$$|\{x \in X : x < x^*\}| < n/2 - n_0^- \leq |\{x \in X : x \leq x^*\}|.$$

This can be done in $O(dn)$ time using a standard rank selection algorithm. Note that there is at least one voter $v_1 \in V_1$ whose projection has an x_{d-1} -coordinate equal to x^* . We claim that the hyperplane h containing f_0 and v_1 is a median hyperplane. To see this, we first observe that

$$|h^- \cap V| = n_0^- + |\{x \in X : x < x^*\}| < n/2.$$

Second, since h contains v as well as the voters in V_1 whose projection has an x_{d-1} -coordinate x^* , we have

$$|h^- \cap V| + |h \cap V| \geq n_0^- + (n/2 - n_0^-) + 1 = n/2 + 1.$$

Thus, we have $|h^+ \cap V| < n/2$, which proves the claim and, hence, the lemma. \square

Recall that for $d \geq 2$ the plurality point is unique, if it exists. The algorithm below either reports a single candidate point p —we show later how to test if the candidate is actually a plurality point or not—or it returns \emptyset to indicate that it already discovered that a plurality point does not exist. When called with a set V of n collinear voters, the algorithm will return the set of all plurality points; if n is even, the set is a segment connecting the two median voters, and if n is odd, the set is a degenerate segment consisting of the (in this case unique) median voter. We call this segment the *median segment*.

FINDCANDIDATES(V)

- (1) If all voters in V are collinear, then return the median segment of V .
- (2) Otherwise, proceed as follows.
 - (a) Let $v_0 \in V$ be a voter with a minimum x_d -coordinate. All voters in V lie either on the plane $h_0 : x_d = x_d(v_0)$ or in the half-space above h_0 , which we denote by h_0^+ . Find a median hyperplane m_0 containing v_0 using Lemma 2.5, and let $cand_0 := \text{FINDCANDIDATES}(m_0 \cap V)$.
 - (b) If $cand_0$ consists of a single point or $cand_0 = \emptyset$, then return $cand_0$.
 - (c) If $cand_0$ is a (nondegenerate) segment, then let $v_1 \in V$ be a voter whose distance to m_0 is maximized and h_1 be the hyperplane parallel to m_0 containing v_1 . All voters in V lie either on h_1 or in the half-space defined by h_1 and containing m_0 , which we denote by h_1^+ . Find a median hyperplane m_1 containing v_1 using Lemma 2.5, and let $cand_1 := \text{FINDCANDIDATES}(m_1 \cap V)$. Return $cand_0 \cap cand_1$.

LEMMA 2.6. *Algorithm FINDCANDIDATES(V) returns in $O(dn)$ time a set $cand$ of candidate plurality points such that*

- (i) *if all voters in V are collinear, then $cand$ is the set of all plurality points of V ;*
- (ii) *if not all voters in V are collinear, then $cand$ contains at most one point, and no other point can be a plurality point of V .*

PROOF. We first prove the correctness of the algorithm and then consider the time bound.

If all voters in V are collinear, then the algorithm returns the correct result in Step 1, so assume not all voters are collinear. Consider the median hyperplane computed in Step 2a. Since $|m_0^+ \cap V| < n/2$ and $|m_0^- \cap V| < n/2$, for any point $p \notin m_0$ there is an open half-space containing p and bounded by a hyperplane parallel to m_0 that contains more than $n/2$ voters. Hence, by Proposition 2.1, any plurality point for V must lie on m_0 . By Lemma 2.3, if a plurality point exists for V , it must also be a plurality point for $m_0 \cap V$. By induction, we can assume that $\text{FINDCANDIDATES}(m_0 \cap V)$ is correct. Hence, the result of the algorithm is correct when cand_0 consists of a single point or $\text{cand}_0 = \emptyset$. Note that when cand_0 is a (nondegenerate) segment—this only happens when all voters in $m_0 \cap V$ are collinear—we must have $V \neq m_0 \cap V$, as otherwise V would be collinear and we would be done after Step 1. Hence, $v_1 \notin m_0$. By the same reasoning as above, the median hyperplane m_1 must contain the plurality point of V (if it exists). But then the plurality point must lie in $\text{cand}_0 \cap \text{cand}_1$, and since $v_1 \notin m_0$, we know that $\text{cand}_0 \cap \text{cand}_1$ is either a single point or it is empty. This proves the correctness.

To prove the time bound, we note that we only have two recursive calls when the first recursive call reports a nondegenerate candidate segment. This only happens when all voters in $m_0 \cap V$ are collinear, which implies the recursive call just needs to compute a median segment in $O(dn)$ time—it does not make further recursive calls. Thus, we can imagine adding this time to the original call, so that we never make more than one recursive call. Since the recursion depth is at most d , and each call needs $O(dn)$ time, the bound follows. \square

Our algorithm to find a plurality point first calls $\text{FINDCANDIDATES}(V)$. If all points in V are collinear, we are done— $\text{FINDCANDIDATES}(V)$ then reports the correct answer. Otherwise, we either get a single candidate point p or we already know that a plurality point does not exist. It remains to test if a candidate point p is a plurality point or not.

LEMMA 2.7. *Given a set V of n voters in \mathbb{R}^d and a candidate point p , we can test in $O(dn \log n)$ time if p is a plurality point in the L_2 norm.*

PROOF. First compute the set $L(p)$ of lines containing p and at least one voter. We can compute $L(p)$, and for each line $\ell \in L(p)$ the number of voters on the rays $\rho(\ell)$ and $\bar{\rho}(\ell)$, in $O(dn \log n)$ time. (To this end, we take the line ℓ_v through p and v for each voter $v \neq p$ and group these into subsets of identical lines.) According to Theorem 2.2, we can now immediately decide if p is a plurality point when n is odd, or when n is even and $p \notin V$. When n is even and $p \in V$, we first check in $O(dn)$ time if all unbalanced lines lie in a 2-flat f . If not, then p is not a plurality point. Otherwise, we check the alternating property in $O(dn \log n)$ time. \square

We obtain the following theorem.

THEOREM 2.8. *Let V be a set of n voters in \mathbb{R}^d , where $d \geq 2$. Then we can find in $O(dn \log n)$ time the plurality point for V in the L_2 norm, if it exists, and this time bound is optimal for $d = 2$.*

PROOF. The time bound and correctness of our algorithm follow from Lemmas 2.6 and 2.7. The optimality of the algorithm follows from an easy reduction from SET EQUALITY, which takes $\Omega(n \log n)$ time in the algebraic decision tree model [3]. To see this, consider an instance A, B of SET EQUALITY, where $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ are sets of n numbers. Define $V := \{(a_i, 1) : a_i \in A\} \cup \{(0, 0)\} \cup \{(-b_i, -1) : b_i \in B\}$. Then $A = B$ if and only if V admits a plurality point (which must then be the point $(0, 0)$), as follows immediately from Theorem 2.2a. \square

3 DEALING WITH POINT SETS THAT DO NOT ADMIT A PLURALITY POINT

Most point sets do not admit a plurality point in the L_2 norm. In this section, we consider two ways of dealing with this: we present algorithms to compute a minimum-cost subset $W \subset V$ such that

$V \setminus W$ admits a plurality point, and we present an algorithm for computing a minimum-radius plurality ball in \mathbb{R}^2 .

3.1 The Minimum-Cost Plurality Problem

Let V be a set of n voters in \mathbb{R}^d , where each voter v has a cost $\text{cost}(v) > 0$ associated to it. For a candidate plurality point p —here we consider all points in \mathbb{R}^d as candidates—we define W_p to be a minimum-cost subset of V such that p is a plurality point for $V \setminus W_p$. We define the *price* of p to be the cost of W_p . Our algorithm will report a pair (p, W_p) , where p is a cheapest candidate plurality point. The algorithm has two main parts: one finds the cheapest candidate that does not coincide with any voters, and the other finds the cheapest candidate that coincides with a voter.

Let $L(V)$ be the set of lines passing through at least two voters in V , and let $P(V)$ be the set of all intersection points of the lines in $L(V)$, excluding the intersection points coinciding with a voter. To find the cheapest candidate p that does not coincide with a voter, we only have to consider points in $P(V)$. Indeed, if all points in $V \setminus W_p$ are collinear, then we can pick p to coincide with a voter; otherwise, we know by Theorem 2.2b that p must be an intersection point of two lines in $L(V \setminus W_p)$, and so $p \in P(V)$. We will need the following lemma for the planar case.

LEMMA 3.1 (LIN ET AL. [18]). *Let p be a candidate plurality point for a set V in \mathbb{R}^2 . Then we can compute in $O(n \log n)$ time the price of p , together with the subset W_p .*

We also need the following lemma to compute all the intersections of each line with other lines in sorted order.

LEMMA 3.2. *Given a set \mathcal{L} of m lines in \mathbb{R}^d , we can list all intersection points of each line ℓ with the other lines in sorted order along ℓ in total time $O(d^2 m^2)$.*

PROOF. Consider the set \mathcal{F} of all 2-flats defined by any two axes x_i and x_j ($1 \leq i, j \leq d$). We have $O(d^2)$ of them in total. For any such 2-flat $f \in \mathcal{F}$, we project all lines in \mathcal{L} onto f and construct the arrangement in f in $O(m^2)$ time [9] and then check which intersections on f correspond to actual intersections in \mathbb{R}^d . Finally, for each line ℓ we merge the $O(d^2)$ lists of intersection points to obtain a complete, sorted list of intersection points along ℓ . In total, this takes $O(d^2 m^2)$ time.

Any two intersecting lines $\ell_1, \ell_2 \in \mathcal{L}$, define a 2-flat; as this 2-flat cannot be orthogonal to all the 2-flats in \mathcal{F} , it will be projected as a plane onto at least one of them. Therefore, we don't miss any intersection and this proves the correctness of the algorithm. \square

In the algorithm below, we use $L(p, V')$ to denote the set of lines through a point p and at least one voter in a set $V' \subseteq V$.

MINCOSTPLURALITYPOINT(V)

- (1) Compute the set $L(V)$. If $|L(V)| = 1$ —that is, all voters lie on a common line ℓ —then compute a median p along ℓ as a plurality point and report (p, \emptyset) .
- (2) Compute a cheapest candidate p that does not coincide with a voter, as follows. Compute the set $P(V)$. For each line $\ell \in L(V)$, sort the intersection points along ℓ . Using Lemma 3.2, this can be done in $O(d^2 n^4)$ time in total. Let $C := \sum_{v \in V} \text{cost}(v)$ be the total cost of all voters. For each intersection point $p \in P(V)$, let $\gamma(p)$ be the total cost of all voters v for which there is no line in $L(v, V \setminus \{v\})$ that contains p ; we can compute $\gamma(p)$ in $O(n^4)$ time in total by determining the total cost of all voters that *do* have a line in $L(v, V \setminus \{v\})$ that contains p and then subtracting this cost from C .
 - (a) Traverse each line $\ell \in L(V)$ to visit the intersection points along ℓ in order. During the traversal, maintain the number of voters on ℓ on either side of the current intersection point p . Thus, we know how many voters we have to remove to make ℓ balanced,

and also from which side we should remove them. If we have to remove k voters, we have to remove the k cheapest voters on the relevant side. The subset $W_p(\ell)$ that we have to remove to make ℓ balanced only changes when p passes over a voter v on ℓ . When this happens, we can compute the new $W_p(\ell)$ in linear time. In this way, the traversal of ℓ takes $O(n^2)$ time in total, so over all $\ell \in L(V)$ we spend $O(n^4)$ time.

- (b) For each intersection point p compute the price of p . This price has two components: the price to make every line $\ell \in L(V)$ that contains p balanced, and the price to remove any voter v for which the line $\ell(v, p)$ through v and p is not a line in $L(V)$. The first component equals $\sum_{\ell \ni p} \text{cost}(W_p(\ell))$. The second component equals $\gamma(p)$, which we precomputed.
- (3) Compute a cheapest candidate p that coincides with a voter $v \in V$. To this end, compute for each voter v the price of setting $p := v$ —below we describe how to do this in $O(dn^3)$ time per voter—and take the cheapest of all n possibilities. Consider a candidate p coinciding with some $v \in V$. By Theorem 2.2, all unbalanced lines in $L(p, V \setminus W_p)$, if any, lie on a single 2-flat f . We will compute $\text{price}_p(f)$, the price to make p into a plurality point under the condition that all unbalanced lines lie in f , over all 2-flats f spanned by two lines from $L(p, V \setminus \{v\})$. Then we take the best of the results.
 - Fix a 2-flat f spanned by two lines $\ell_1, \ell_2 \in L(p, V \setminus \{v\})$. Let L_f be the subset of lines from $L(p, V \setminus \{v\})$ contained in f , and let L'_f be the subset of lines not contained in f .
 - (a) Compute $\text{price}_p(L_f)$, the price of making p a plurality point on f , using Lemma 3.1. This takes $O(n_f \log n_f)$ time, where $n_f := |f \cap V|$.
 - (b) For each line $\ell \in L'_f$, compute $\text{price}_p(\ell)$, the price of making ℓ balanced. (This can easily be done in $O(n)$ time: we compute the smallest number k of voters we have to remove on ℓ to make ℓ balanced, and then find the k cheapest voters on the heavier of the two rays along ℓ and emanating from p .) Let $\text{price}_p(L'_f) := \sum_{\ell \in L'_f} \text{price}_p(\ell)$.
 - (c) Set $\text{price}_p(f) := \text{price}_p(L_f) + \text{price}_p(L'_f)$.
- (4) Let p be the cheaper of the two candidates found in Steps 2 and 3, respectively. Compute the set W_p for this candidate—this takes $O(n \log n)$ time—and report (p, W_p) .

The correctness of the algorithm follows from Theorem 2.2 and the discussion above. As for the running time, we note that Steps 1, 2, and 4 all run in $O(d^2n^4)$ time. For Step 3, the time needed to compute the price of a single voter v is $\sum_f O(n_f \log n_f + dn)$, which is bounded by $O(dn^3 + \sum_f n_f \log n_f)$. Because every voter $v' \in V$ lies on at most n of the flats f (generated by the lines $\ell(v, v')$ and $\ell(v, v'')$ through v, v' and v, v'' , respectively), we have $\sum_f n_f = O(n^2)$ and so $\sum_f n_f \log n_f = O(n^2 \log n)$. Hence, the whole algorithm runs in $O(d^2n^4)$ time.

THEOREM 3.3. *Let V be a set of n voters in \mathbb{R}^d , each with a positive cost, where $d \geq 2$. Then we can compute in $O(d^2n^4)$ time a minimum-cost subset $W \subset V$ such that $V \setminus W$ admits a plurality point in the L_2 norm.*

Our algorithm for finding a minimum-cost plurality point checks $O(n^4)$ candidate points. The algorithm from the previous section for deciding if a plurality point exists avoids this, resulting in a near-linear running time. An obvious question is if a faster algorithm is also possible for the minimum-cost plurality-point problem. While we do not have the answer to this question, we can show that, even in the plane and when all voters have unit cost, it is unlikely that the problem can be solved in truly subquadratic time. We do this by a reduction from the problem **THREE CONCURRENT LINES**, which is to decide if a set of n lines has three or more lines meeting

in a single point. THREE CONCURRENT LINES is 3SUM-hard [12] and has an $\Omega(n^2)$ lower bound if only sidedness tests are used [11].

THEOREM 3.4. *Suppose we have an algorithm solving the minimum-cost plurality-point problem for any set of n unit-cost voters in the plane in time $T(n)$. Then there is a probabilistic algorithm solving THREE CONCURRENT LINES with probability 1 in $O(n \log n + T(n))$ time.*

PROOF. Let L be a set of n lines for which we want to solve THREE CONCURRENT LINES. We will generate a set V of $2n$ voters such that L contains three or more concurrent lines if and only if there is a subset $W \subset V$ of size at most $2n - 6$ such that $V \setminus W$ admits a plurality point. To this end, we first compute in $O(n \log n)$ time a large circle C that contains all vertices of the arrangement $\mathcal{A}(L)$ in its interior. The circle C can be any circle containing the axis-parallel bounding box of the vertices of $\mathcal{A}(L)$. (It is well known that the bounding box of the vertices of a planar arrangement can be found in $O(n \log n)$ time; see Exercise 8.4 from [7].)

For each line $\ell \in L$ construct two small segments contained in ℓ , a segment $s_1(\ell) \subset \ell$ around the first intersection point of ℓ and C and another segment $s_2(\ell) \subset \ell$ around the second intersection point. Finally, we pick a voter uniformly at random in each segment $s_i(\ell)$ with $i \in \{1, 2\}$ and $\ell \in L$, thus creating a set V of $2n$ voters. We pick the segments $s_i(\ell)$ sufficiently small so that the set V is guaranteed to be in convex position. (Computing suitable lengths can be done in linear time, after we have sorted the intersection points of the lines in L with the circle C in clockwise order along C .)

If L has three or more lines meeting at a point p , then we can make p into a plurality point by removing at most $2n - 6$ points on the lines not passing through p . To prove the other direction, note that to make any of the voters in V into a plurality point, we have to delete all but one of the remaining $2n - 1$ voters (since the voters are in convex position). Let $L(V) := \{\ell(v, v') : v, v' \in V\}$ be the set of all lines through a pair v, v' of voters. To finish the proof, we observe that by picking the voters randomly on the segments $s_i(\ell)$, the following holds with probability 1: the set $L(V)$ has no triple intersections except those already present in L . \square

3.2 A Fixed-Parameter Algorithms for Unit-Cost Voters

The proof of Theorem 3.4 uses a problem instance where many voters must be removed to obtain a plurality point. Below we show that a plurality point for which only a few voters have to be removed can be found in near-linear time for unit-cost voters. More precisely, we consider the case where we are given a set V of unit-cost voters in \mathbb{R}^d and a parameter k , and we want to compute a smallest subset $W \subset V$ of size at most k (if it exists) such that $V \setminus W$ admits a plurality point.

Define a k -line to be a line ℓ such that for any hyperplane H containing ℓ , both open half-spaces bounded by H contain at most $n/2 + k$ voters. For a voter v , let $L(v)$ be the set of lines containing v and at least one other voter, and let $L_k(v)$ be the set of all k -lines in $L(v)$.

LEMMA 3.5. *Let p be a plurality point of $V \setminus W$, for some subset W of size at most k . If $v \notin W$, then p lies on one of the lines in $L_k(v)$.*

PROOF. Assume $v \notin W$. If $\ell(p, v)$, the line through p and v , is not a line in $L(v)$, then p does not coincide with a voter and $\ell(p, v)$ is unbalanced. But then p cannot be a plurality point, by Theorem 2.2. Hence, $\ell(p, v) \in L(v)$. If $\ell(p, v) \notin L_k(v)$, then there is an open half-space bounded by a hyperplane containing p with more than $n/2 + k$ voters. But since $|W| \leq k$, such a point p cannot be a plurality point in $V \setminus W$, which implies we must have $\ell(p, v) \in L_k(v)$. \square

The idea of our algorithm is now as follows. Consider a set $P := \{(v_{2i-1}, v_{2i}) : 1 \leq i \leq k+1\}$ of disjoint pairs of voters. Then there must be a pair $(v_{2i-1}, v_{2i}) \in P$ such that neither v_{2i-1} nor v_{2i} is in W . Hence, the point p we are looking for must lie on one of the lines in $L_k(v_{2i-1})$ and one of

the lines in $L_k(v_{2i})$. So we check all intersection points between these lines, for every pair in P . The key to obtain an efficient algorithm is to generate P such that all sets $L_k(v_i)$ are small. There is one case that needs special attention, namely, when there is a line—that is present in both $L_k(v_{2i-1})$ and $L_k(v_{2i})$. This case is handled using the following lemma.

LEMMA 3.6. *Suppose there exists a plurality point on $\ell := \ell(v, v')$ through two voters $v, v' \in V$ of price at most k . Let H be an arbitrary hyperplane containing ℓ but no other voter in $V \setminus \ell$. Let H^+ be any of the two open half-spaces bounded by H , and let $V^+ := H^+ \cap V$. Then either an intersection point of ℓ and a line in $\bigcup_{v'' \in V^+} L_k(v'')$, or one of the medians of the voters located on ℓ is a cheapest plurality point.*

PROOF. Let p be a cheapest plurality point and W_p the corresponding subset of voters to be removed, where $|W_p| \leq k$. If there is a voter $v'' \in V^+ \setminus W_p$, then $p \in L_k(v'')$ by Lemma 3.5. Otherwise, all voters in V^+ are in W_p . But then any line through p and a voter in $H^- \cap V$ is unbalanced, which implies that all voters in $H^- \cap V$ except at most one must be in W_p . If all voters in $H^- \cap V$ are in W_p , one of the medians of the voters located on ℓ is a cheapest plurality point. Otherwise, the number of voters located on ℓ must be odd and the median is the cheapest plurality point as any other point violates the alternating property. \square

We now present our algorithm. We first explain the algorithm for the planar case and then show how to extend it to higher dimensions. For technical reasons we assume $k \leq n/15$; see the proof of Theorem 3.8 for more details.

FIXEDPARAMETERMINCOSTPLURALITYPOINT(V, k)

- (1) Compute the set of convex layers of V . Let V_1, V_2, \dots be the sets of voters in these layers, where V_1 is the outermost layer, V_2 the next layer, and so on. Set $P := \emptyset$ and $i := 1$.
- (2) Visit the voters in V_i in clockwise order, starting at the lexicographically smallest voter in V_i . Put the first and second visited voters, the third and fourth visited voters, and so on as pairs into P until either P contains $k + 1$ pairs or we run out of voters in V_i . In the former case, we are done; in the latter case, we start collecting pairs from the next layer, by setting $i := i + 1$ and repeating the process. This continues until we have collected $k + 1$ pairs. Note that in the latter case, exactly one voter remains unpaired if $|V_i|$ is odd.
- (3) Set $C := \emptyset$; the set C will contain candidates for the cheapest plurality points. For each pair $(v_{2i-1}, v_{2i}) \in P$, proceed as follows.
 - (a) Compute the sets $L_k(v_{2i-1})$ and $L_k(v_{2i})$, and put all intersection points between two lines in $L_k(v_{2i-1}) \cup L_k(v_{2i})$ as candidates into C .
 - (b) Let $\ell := \ell(v_{2i-1}, v_{2i})$. If ℓ is present in both $L_k(v_{2i-1})$ and $L_k(v_{2i})$, and ℓ contains at least $n/2 - 7k - 1$ voters, then proceed as follows. (We assume that the same line ℓ has not already been handled in this manner for a pair (v_{2j-1}, v_{2j}) with $j < i$; otherwise, we can skip it now.) Assume without loss of generality that ℓ^+ , the open half-plane above ℓ , contains at most as many voters as ℓ^- . For each voter $v \in \ell^+ \cap V$, compute $L_k(v)$ and add the intersection points of the lines in $L_k(v)$ with ℓ to the candidate set C . In addition, put a median voter along ℓ into C .
- (4) For each candidate point $p \in C$, compute a minimum-size subset W_p that makes p into a plurality point, using Lemma 3.1. Return the cheapest plurality point $p^* \in C$, provided $|W_{p^*}| \leq k$; if $|W_p| > k$ for all candidates, and then report that it is not possible to obtain a plurality point by removing at most k voters.




Fig. 2. Sketch of the proof for Lemma 3.7.

The efficiency of our algorithm is based on the following lemma, which uses the fact that we constructed P using the convex layers of V .

LEMMA 3.7. *Let v be a voter of some pair in P . Then $|L_k(v)| = O(k)$.*

PROOF. Let V_i be the layer containing v , and let ℓ be a line tangent to $\text{CH}(V_i)$ at v , the convex hull of V_i . Assume without loss of generality that $\text{CH}(V_i)$ lies above ℓ . Then $\ell^- \cap V \subset V_1 \cup \dots \cup V_{i-1}$, and so $|\ell^- \cap V| \leq 3k$. Moreover, the number of voters from V_i that lie on ℓ but “before” v is at most $2k$. (This is true because when we start visiting the voters in V_i in Step 2, we never start in the middle of a sequence of collinear voters.) Thus, we can rotate ℓ slightly around v , so that v is the only voter on ℓ and $|\ell^- \cap V| \leq 5k$. Turn ℓ into a directed line in such a way that ℓ^- lies to the left of ℓ , and imagine rotating ℓ around v by π radians in the counterclockwise direction. Let $\ell(\phi)$ be the line resulting from rotating ℓ over an angle ϕ , and define $F^-(\phi)$ and $F^+(\phi)$ to be the number of voters to the left and right of $\ell(\phi)$, respectively (see Figure 2). The line $\ell(\phi)$ is a k -line if it passes through a voter, $F^-(\phi) \leq n/2 + k$ and $F^+(\phi) \leq n/2 + k$. Note that $F^-(0) \leq 5k$ and $F^+(0) \geq n - 5k$, and $F^-(\pi) \geq n - 5k$ and $F^+(\pi) \leq 5k$. Since $F^-(\phi)$ can only decrease (and, similarly, $F^+(\phi)$ can only increase) when $\ell(\phi)$ passes through a voter that was to the left of $\ell(0)$, this implies that the number of times at which $\ell(\phi)$ is a k -line for v is $O(k)$. \square

We can now prove the following theorem.

THEOREM 3.8. *Let V be a set of n voters in the plane, and let k be a parameter with $k \leq n/15$. Then we can compute in $O(k^3 n \log n)$ time a minimum-size subset $W \subset V$ such that $V \setminus W$ admits a plurality point in the L_2 norm or report that there is no such subset W of size at most k .*

PROOF. To prove the time bound, we note that we can compute the convex layers in $O(n \log n)$ time [5]. Step 2 runs in $O(k)$ time. Computing the set $L_k(v)$ for a voter v can easily be done in $O(n \log n)$ time by using a sweeping algorithm that rotates a line around v and keeps track of the number of points on each of the half-spaces defined by that line. By Lemma 3.7, Step 3a takes $O(n \log n + k^2)$ time. To bound the running time of Step 3b, we observe that there can be at most $O(1)$ pairs (v_{2i-1}, v_{2i}) to which this case applies. Indeed, $\ell(v_{2i-1}, v_{2i})$ should contain at least $n/2 - 7k - 1$ voters, and since $k \leq n/15$, there can only be $O(1)$ such lines. Thus, Step 3b needs

$O(kn)$ time, and so Step 3 needs $O(k^2n + kn \log n)$ time in total over all pairs in P , to generate $O(k^3)$ candidate points. Checking each of the candidates takes $O(n \log n)$, which proves the time bound.

The correctness of the algorithm follows from Lemmas 3.5 and 3.6, except for one thing: in Step 3b we only handle a line $\ell := \ell(v_{2i-1}, v_{2i})$ when it contains at least $n/2 - 7k - 1$ voters. This is allowed for the following reason. Note that ℓ is tangent to a convex hull $\text{CH}(V_i)$ and on the side of ℓ that does not contain $\text{CH}(V_i)$, say, ℓ^+ , there are at most $3k$ voters. Now consider a plurality point $p \in \ell$. Then there can be at most $3k + 1$ voters in $(V \setminus W_p) \cap \ell^-$, by Theorem 2.2. Since $|W_p| \leq k$, we thus have $|V \cap \ell^-| \leq 4k + 1$, which means that ℓ must contain at least $n - 7k - 1$ voters. \square

Next we show how to generalize the algorithm to higher dimensions.

Let V be a set of voters in \mathbb{R}^d , with $d \geq 3$. Let f be any lower-dimensional flat in \mathbb{R}^d . Let v^f be the projection of voter v onto flat f and V^f be the set of voters obtained by projecting all the voters in V onto flat f . Assume no two voters in V are projected to a single voter in V^f . Let $L_k(v^f)$ be the set of k -lines of V^f on f containing v^f , and let $L_k^f(v)$ be all lines in $L(v)$ whose projection is in $L_k(v^f)$. We need the following lemma about the conservation of the plurality point property under projection.

LEMMA 3.9. *Let $V \in \mathbb{R}^d$ be a set of voters, where $d > 2$, and let f be a lower-dimensional flat in \mathbb{R}^d . Then*

- (i) $L_k(v) \subseteq L_k^f(v)$; and
- (ii) if p is a plurality point for V , then p^f is a plurality point for V^f .

PROOF. Consider a line $\ell \in L(v)$ and its projection $\ell^f \in L(v^f)$ on f . If $\ell^f \notin L_k(v^f)$, then there is a hyperplane H^f in f such that the number of voters from V^f on one of the relatively open half-spaces bounded by H^f is greater than $n/2 + k$. Let H be a hyperplane such that the projection of H on f is H^f . H contains ℓ and the number of voters on an open half-space bounded by H is more than $n/2 + k$, which means $\ell \notin L_k(v)$. This proves the first part of the lemma.

A point p is a plurality point for a set of voters V if and only if for any $v \in V$, all lines $\ell \in L(v)$ containing p are in $L_0(v)$. For any line $\ell \in L(v)$, by a similar argument to the previous part, $\ell^f \notin L_0(v^f)$ implies $\ell \notin L_0(v)$. This proves the second part of the lemma. \square

Using this lemma, we can extend Theorem 3.8 to higher dimensions.

THEOREM 3.10. *Let V be a set of n voters in \mathbb{R}^d where $d > 2$, and let k be a parameter with $k \leq n/15$. Then we can compute in $O(dk^5 \log k + dk^3 n \log n)$ expected time a minimum-size subset $W \subset V$ such that $V \setminus W$ admits a plurality point in the L_2 norm.*

PROOF. We extend the algorithm for \mathbb{R}^2 to \mathbb{R}^d . The algorithm has four steps. In steps 1 and 2, we find the set of pairs P and a set of lines passing through the voters in these pairs. In step 3, we find the set of candidate points using these two sets. Finally, in step 4, we verify all the candidate points in order to find a plurality point (if one exists).

We modify steps 1 and 2 for \mathbb{R}^d as follows. We consider a 2-flat $f \subset \mathbb{R}^d$ and project all the voters in V onto flat f in order to get V^f . We will need certain nondegeneracy assumptions on the projections of the voters from V and on the projections of lines through pairs of voters onto f . If we take f to be a random 2-flat, then these assumptions are satisfied with probability 1. During the remainder of the algorithm, we will test for nondegeneracy at the appropriate places; see below for details. When we find that the flat f is not suitable after all, we start from scratch with a new random 2-flat f . Since the success probability is 1, we expect to succeed in one trial.

The first nondegeneracy assumption is that no two voters in V are mapped to a single voter in V^f . We can test this condition in $O(dn \log n)$ time. Now (assuming f passed the test) we apply the

(2-dimensional version of) the algorithm on V^f , giving us a set P^f of $k + 1$ pairs (v_{2i-1}^f, v_{2i}^f) and also the sets $L_k(v^f)$ for all v^f that appear in some pair in P^f .

Since the first nondegeneracy assumption is met, any voter $v^f \in V^f$ is the projection of a unique voter $v \in V$. Therefore, we can uniquely identify the set P of $k + 1$ pairs (v_{2i-1}, v_{2i}) in V that are projected to P^f . However, a line $\ell \in L_k(v^f)$ or $\ell(v_{2i-1}, v_{2i})$ might still be a projection of many lines in $L_k^f(v)$. We need another nondegeneracy assumption to get around this problem, as explained next.

Having the pairs P , the lines $\ell(v_{2i-1}, v_{2i})$ for the pairs $(v_{2i-1}, v_{2i}) \in P$, and the lines in $L_k^f(v)$ for all v appearing in these pairs, we can perform step 3 of FIXEDPARAMETERMINCOSTPLURALITY-POINT and find the set of all candidate points C . The only difference is that instead of $L_k(v)$, we work with $L_k^f(v)$. Therefore, the second nondegeneracy assumption is that the number of voters on any line ℓ we work with in step 3 of the algorithm must be equal to the number of voters on the projected line ℓ^f . As the total number of lines to be tested is $O(k^3)$, we can test this condition in $O(dk^3n)$ in total. (Again, we restart the algorithm if the test fails.)

Now each line in $L_k^f(v)$ uniquely corresponds to a line in $L_k(v^f)$, and $|L_k^f(v)| = |L_k(v^f)|$; moreover, each line $\ell(v_{2i-1}, v_{2i})$ has a unique projection $\ell(v_{2i-1}^f, v_{2i}^f)$. By Lemma 3.9, we have $L_k(v) \subseteq L_k^f(v)$. As each candidate point in C corresponds to a candidate point on f , the number of candidates in C is $O(k^3)$.

In step 4, to verify a candidate point $p \in C$, we find the set $U = \{u_1, \dots, u_{|U|}\}$ of all unbalanced lines containing p . According to Theorem 2.2, in order to have a plurality point, all unbalanced lines must be on a single 2-flat. We devise our verification algorithm based on the following two observations:

- (1) If a 2-flat f contains more than $k + 1$ lines in U , then f is the only possible 2-flat for unbalanced lines. Indeed, any other 2-flat B' has at most one common unbalanced line with f , and to make at least $k + 1$ remaining unbalanced lines in f balanced, we would have to remove more than k voters (which is not allowed).
- (2) For a 2-flat f , the number of lines in U that are not contained in f cannot be more than k —otherwise, we would have to remove more than k voters (which is not allowed).

We define a set \mathcal{F} of candidate 2-flats that together contain all unbalanced lines. We consider two cases.

Case A: $|U| \leq 2k + 2$. In this case, we simply put all 2-flats defined by any two lines in U in \mathcal{F} .

Case B: $|U| > 2k + 2$. In this case, we check all 2-flats defined by any two lines in $\{u_1, \dots, u_{2k+3}\}$. If there is a 2-flat that contains more than $k + 1$ lines in U , we put one such flat f into \mathcal{F} . By the first observation above, this is sufficient. Otherwise, the second observation implies that we cannot have a plurality point by removing only k voters.

Having the candidate 2-flats, we can compute the cost of the minimum-cost plurality point on each candidate 2-flat f . To this end, we first compute the price of p on f using Lemma 3.1. To get the real price of p , we add the cost of balancing the unbalanced lines that are not on f .

To prove the time bound, we note that the main difference between this algorithm and the algorithm for the 2D case is in step 4. We can compute the set U in $O(dn \log n)$ time. If \mathcal{F} contains a single candidate, the complexity of computing the price of p will be $O(dn \log n)$. Otherwise, there are $O(k^2)$ candidate 2-flats. However, since each line appears in at most k 2-flats, the time complexity of the algorithm for each candidate plurality point will be $O(dk^2 \log k)$. For $O(k^3)$ candidate plurality points, the total complexity of the algorithm is therefore $O(dk^5 \log k + dk^3 n \log n)$.




Fig. 3. A 3-set, $\{v_6, v_7, v_8\}$, of $V = \{v_1, \dots, v_8\}$ (left) and the 3-level in $\mathcal{A}(V^*)$ (right).

The correctness of steps 1, 2, and 3 follows from Lemmas 3.5 and 3.6 and the correctness of the algorithm for the 2D case. The correctness of step 4 follows from Lemma 3.1. \square

3.3 The Minimum-Radius Plurality-Ball Problem

Let V be a set of n voters in \mathbb{R}^d . A closed ball $b(p, r)$ of radius r and centered at a point p is a *plurality ball* if for any point $q \notin b(p, r)$ the number of voters who prefer p over q is at least the number of voters who prefer q over p . Note that for any point p the ball $b(p, r)$ is a plurality ball if r is sufficiently large, and that a plurality ball with $r = 0$ is a plurality point. Below we describe an algorithm to compute a minimum-radius plurality ball for V . If all voters are collinear, then any point on the median segment of V is a plurality ball of radius 0, so in the remainder we assume not all voters are collinear.

We define the *core* of a ball $b(p, r)$ as $b(p, r/2)$. Note that for any point $q \notin b(p, r)$ the locus of points that are strictly closer to q than p is an open half-space that does not intersect the core $b(p, r/2)$, because obviously the bisector of p and q (i.e., the hyperplane of all points equidistant to p and q) does not intersect $b(p, r/2)$. Hence, Proposition 2.1 can be generalized as follows.

PROPOSITION 3.11. *A ball $b(p, r)$ is a plurality ball if and only if every open half-space that does not intersect the core $b(p, r/2)$ contains at most $n/2$ voters.*

To check this condition, we use the concept of k -set and k -level and their duality. A k -set of V , for some $0 \leq k \leq n - d$, is defined as a subset $V' \subset V$ of size k such that there is an open half-space h^+ with $h^+ \cap V = V'$ and with at least d points from V on its boundary. Let V^* be the set of hyperplanes dual to the voters in V . The k -level in the arrangement $\mathcal{A}(V^*)$ is the set of points on the hyperplanes in V^* that have exactly k hyperplanes strictly below them (see Figure 3).

We call a k -dimensional feature of a d -dimensional arrangement (or some other d -dimensional structure) a k -face. A $(d - 1)$ -face is also called a *facet*. We associate each facet f of the k -level of $\mathcal{A}(V^*)$ to a cone in the primal space, as follows. Let $V^*(f)$ be the set of hyperplanes strictly below f , and consider the hyperplanes (in primal space) dual to the vertices of f . Then $\text{cone}(f)$ is the closed cone defined by these hyperplanes that contains the k voters whose dual hyperplanes are in $V^*(f)$ plus, at its apex, the voter whose dual hyperplane contains f ; see Figure 3. We call $\text{cone}(f)$ a k -cone. A k -cone contains exactly $k + 1$ voters including the voter at its apex, and the other $n - k - 1$ voters all lie in the opposite cone.

LEMMA 3.12. *A ball $b(p, r)$ is a plurality ball if and only if its core $b(p, r/2)$ intersects all $\lfloor n/2 \rfloor$ -cones of V .*

PROOF. Assume all $\lfloor n/2 \rfloor$ -cones are intersected by $b(p, r/2)$ and suppose for a contradiction that p is not a plurality point. By Proposition 3.11, there must be an open half-space h^+ not intersecting $b(p, r/2)$ and containing more than $n/2$ voters. But then there must be a $\lfloor n/2 \rfloor$ -cone contained inside the half-space, which is a contradiction. On the other hand, if $b(p, r/2)$ does not intersect some $\lfloor n/2 \rfloor$ -cone $cone(f)$, then, since any two disjoint convex sets can be separated by a hyperplane, there is an open half-space h^+ not intersecting $b(p, r/2)$ containing all the points in $cone(f)$. Therefore, $|h^+ \cap V| \geq \lfloor n/2 \rfloor + 1$, and so $b(p, r)$ is not a plurality ball. \square

Our algorithm is now easy: we compute all $\lfloor n/2 \rfloor$ -cones of V by computing the $\lfloor n/2 \rfloor$ -level in the dual arrangement $\mathcal{A}(V^*)$. Then we compute the minimum-radius ball $b(p, r/2)$ intersecting all these cones, and report $b(p, r)$ as the minimum-radius plurality ball. Since in \mathbb{R}^2 a minimum-radius disk intersecting all the cones is computable in linear time [16], we obtain the following result.

THEOREM 3.13. *Let V be a set of n voters in the plane. Then we can compute the minimum-radius plurality ball for V in $O(T(n))$ time, where $T(n)$ is the time needed to compute the $\lfloor n/2 \rfloor$ -level in an arrangement of n lines in the plane.*

Computing the k -level in an arrangement of lines can be done in $O(n \log n + m \log^{1+\epsilon} n)$ time, where m is the complexity of the level. Our algorithm then runs in $O(n^{4/3} \log^{1+\epsilon} n)$ time. In higher dimensions, this problem fits within the generalized linear programming (GLP) framework [2]. However, as the complexity of $\lfloor n/2 \rfloor$ -cones in higher dimensions is high, *violation test* and *basis computation* will be costly; therefore, a GLP-based algorithm would not be very efficient.

4 PLURALITY POINTS IN THE PERSONALIZED L_1 NORM

Let V be a set of n voters in \mathbb{R}^d , where each voter $v \in V$ has a preference vector $\langle w_1(v), \dots, w_d(v) \rangle$ of nonnegative weights. Define $\text{dist}_w(v, p) := \sum_{i=1}^d w_i(v) \cdot |x_i(v) - x_i(p)|$. In this section, we study plurality points for this personalized L_1 distance. As mentioned in the introduction, a plurality point in the L_1 norm always exists in \mathbb{R}^2 , but not in higher dimensions [22]. Interestingly, in the personalized L_1 norm, the statement already fails in the plane: in Section 4.3, we give an example of a weighted point set V in \mathbb{R}^2 that does not admit a plurality point in the personalized L_1 norm.

4.1 Properties of Plurality Points in the Personalized L_1 Norm

Our goal is to formulate conditions that help us to find candidate plurality points and to decide if a given candidate is actually a plurality point. For the L_2 norm we used Theorem 2.2 and Lemma 2.3 for this. Here we need a different approach. Recall that a candidate point $p \in \mathbb{R}^d$ is a plurality point if, for any point $q \in \mathbb{R}^d$, the number of voters who prefer p is at least the number of voters who prefer q . From now on we refer to the point q as a *competitor*.

For two points p and q , define $V[p > q] := \{v \in V : \text{dist}_w(v, p) < \text{dist}_w(v, q)\}$. We also define $V[p \sim q] := \{v \in V : \text{dist}_w(v, p) = \text{dist}_w(v, q)\}$ and $V[p \succcurlyeq q] := V[p > q] \cup V[p \sim q]$. Let p be a candidate plurality point. We call a point q a *nondegenerate competitor* for p if $V[p \sim q] = \emptyset$, and we say that q is ϵ -close to p if $|pq| < \epsilon$, where $|pq|$ denotes the Euclidean distance between p and q . The following lemma implies that to test if a point p is a plurality point, we only have to consider nondegenerate competitors that are ϵ -close to p .

LEMMA 4.1. *Let p be a candidate plurality point and let q be a competitor for p . For any $\epsilon > 0$, there is a nondegenerate competitor q' that is ϵ -close to p such that $|V[q' > p]| \geq |V[q > p]| + \frac{1}{2} \cdot |V[q \sim p]|$.*




Fig. 4. Sketch of the proof for Lemma 4.1.

PROOF. We prove the lemma in two steps: first we show there is an $(\varepsilon/2)$ -close competitor q'' with $V[q > p] \subseteq V[q'' > p]$ and $V[q \succcurlyeq p] \subseteq V[q'' \succcurlyeq p]$, and then we show that this implies the lemma. We assume without loss of generality that p lies at the origin.

For the first step, let $\varepsilon' > 0$ be small enough so that $q'' := (\varepsilon' \cdot x_1(q), \dots, \varepsilon' \cdot x_d(q))$ is $(\varepsilon/2)$ -close to p . Let v be a voter who prefers q over p , and define $b_v := b(v, \text{dist}_w(v, p))$ to be the ball (in the personalized L_1 metric of v) centered at v and with radius $\text{dist}_w(v, p)$. Then $p \in \partial b_v$ and $q \in \text{int}(b_v)$, where ∂b_v and $\text{int}(b_v)$ denote the boundary and interior of b_v , respectively. Since b_v is convex—this directly follows from the fact that balls in the standard L_1 norm are convex—we can conclude that $q'' \in \text{int}(b_v)$. Hence, $V[q > p] \subseteq V[q'' > p]$. The proof that $V[q \succcurlyeq p] \subseteq V[q'' \succcurlyeq p]$ is similar.

For the second step, define $B := \{b_v : v \in V[p \sim q'']\}$. Note that each ball $b_v \in B$ has both p and q'' on its boundary. Without loss of generality, we can assume $x_1(p) \neq x_1(q)$. Take two points q_1, q_2 such that $x_1(q_1) = x_1(q'') - \varepsilon'$, and $x_1(q_2) = x_1(q'') + \varepsilon'$, and $x_i(q_1) = x_i(q_2) = x_i(q'')$ for all $2 \leq i \leq d$ (see Figure 4). By taking $\varepsilon' > 0$ sufficiently small, we can ensure that (1) q_1 and q_2 are ε -close to p , and (2) if $q'' \in \text{int}(b_v)$ for some ball $b_v \in B$, then $q_1, q_2 \in \text{int}(b_v)$, and if $q'' \in \text{ext}(b_v)$, then $q_1, q_2 \in \text{ext}(b_v)$, where $\text{ext}(b_v)$ is the set of points not on the boundary or interior of b_v . Now consider a voter v with $q'' \in \partial b_v$. Since q'', q_1, q_2 are ε -close to p , we can assume $x_1(v) \neq x_1(q'')$ and that either $x_1(v) < x_1(q_1) < x_1(q'') < x_1(q_2)$ or $x_1(q_1) < x_1(q'') < x_1(q_2) < x_1(v)$. This implies that either $\text{dist}_w(v, q_1) < \text{dist}_w(v, q'')$ and $\text{dist}_w(v, q_2) > \text{dist}_w(v, q'')$, or $\text{dist}_w(v, q_2) < \text{dist}_w(v, q'')$ and $\text{dist}_w(v, q_1) > \text{dist}_w(v, q'')$. But then q_1, q_2 are both nondegenerate, and each voter $v \in V[q'' \sim p]$ is either in $V[q_1 > p]$ or in $V[q_2 > p]$. Assume without loss of generality that at least half of the voters from $V[q'' \sim p]$ are in $V[q_1 > p]$. Then

$$|V[q_1 > p]| \geq |V[q'' > p]| + \frac{1}{2} \cdot |V[q'' \sim p]| \geq |V[q > p]| + \frac{1}{2} \cdot |V[q \sim p]|,$$

where the last inequality holds because $V[q > p] \subseteq V[q'' > p]$ and $V[q \succcurlyeq p] \subseteq V[q'' \succcurlyeq p]$. Hence, we can take $q' := q_1$. \square

The following lemma helps us to narrow down our search for plurality points. Recall that a *multidimensional median* for V is a point $p \in \mathbb{R}^d$ such that, for all $1 \leq i \leq d$, we have $|\{v \in V : x_i(v) < x_i(p)\}| \leq n/2$ and $|\{v \in V : x_i(v) > x_i(p)\}| \leq n/2$.

LEMMA 4.2. *Let p be a plurality point for V in the personalized L_1 norm. Then p is a multidimensional median for V .*

PROOF. Suppose for a contradiction that p is not a multidimensional median. Assume without loss of generality that p lies at the origin and that $|\{v \in V : x_1(v) > 0\}| > n/2$. Consider a competitor $q := (\varepsilon, 0, \dots, 0)$ for p , where ε is small enough so that no voter v has an x_1 -coordinate in the range $(0, \varepsilon)$. For every voter v with $x_1(v) > 0$, we then have $|x_1(v) - x_1(q)| < |x_1(v) - x_1(p)|$ and $|x_i(v) - x_i(q)| = |x_i(v) - x_i(p)|$ for all $2 \leq i \leq d$. Hence, every such voter prefers q over p , contradicting the condition that p is a plurality point. \square

The set M_V of all multidimensional medians for V is an axis-aligned hyperrectangle in \mathbb{R}^d ; that is, it can be written as $M_V = I_1 \times \dots \times I_d$, where each I_i is a closed interval that may degenerate into a single value. We call M_V the *median box* of V . Lemma 4.2 states that we only have to look at points in M_V when searching for plurality points. The next theorem implies that we only have to check which vertices of M_V are plurality points to fully classify the set of all plurality points. Let F be the set of all k -dimensional faces of M_V for $0 \leq k \leq d$, where each face $f \in F$ is considered relatively open. Note that $|F| = 3^{d'}$, where d' is the number of nondegenerate intervals defining M_V .

THEOREM 4.3. *Let V be a set of voters in \mathbb{R}^d and let f be a relatively open face of the median box M_V of V .*

- (a) *Either all points in f are plurality points in the personalized L_1 norm or none of the points in f are.*
- (b) *The points in f are plurality points in the personalized L_1 norm if and only if all vertices of f are plurality points.*

PROOF. Part (a) follows immediately from part (b). Next we prove part (b).

(b, \Rightarrow). Suppose $p \in f$ is a plurality point, and consider a vertex p' of f . We need to prove that p' is also a plurality point. Let $\varepsilon > 0$ be small enough so that for each $1 \leq i \leq d$ and every voter $v \in V$ with $x_i(v) \neq x_i(p)$ we have $|x_i(v) - x_i(p)| > \varepsilon$; define ε' similarly for p' . Let b and b' be the Euclidean balls of radius $\min(\varepsilon, \varepsilon')$ and centered at p and p' , respectively. Since p is a plurality point, we know that $|V[p > q]| \geq |V[q > p]|$ for all $q \in b$. Now consider an arbitrary point $q' \in b'$, and define $q := q' + (p - p')$. Note that $q \in b$ and that the relative position of q and p is the same as the relative position of q' and p' . In particular, $x_i(p) - x_i(q) = x_i(p') - x_i(q')$ for all $1 \leq i \leq d$. As shown in Lemma 4.4, this implies that $V[p > q] \subseteq V[p' > q']$ and $V[q' > p'] \subseteq V[q > p]$. Hence, $|V[p' > q']| \geq |V[p > q]| \geq |V[q > p]| \geq |V[q' > p']|$. Since q' is an arbitrary point in b' , the point p' must be a plurality point according to Lemma 4.1.

(b, \Leftarrow). To prove this, it suffices to show the following: if p is a point in the relative interior of f that is not a plurality point, then there is a vertex p' of f that is not a plurality point. To this end let q be an ε -close competitor (for a sufficiently small $\varepsilon > 0$) that beats p . We will argue that the vertex p' on the opposite side of p as compared to q' —this is defined more precisely below—is not a plurality point.

Let $M_V = I_1 \times \dots \times I_d$, where $I_i = [\min_i, \max_i]$ (possibly with $\min_i = \max_i$). For each $1 \leq i \leq d$, we pick $x_i(p')$ as follows. If $x_i(p) = \min_i$ or $x_i(p) = \max_i$, then we set $x_i(p') := x_i(p)$. Otherwise, we have $\min_i < x_i(p) < \max_i$; then, if $x_i(p) \leq x_i(q)$, we set $x_i(p') := \min_i$, and if $x_i(p) > x_i(q)$, we set $x_i(p') := \max_i$. Now consider the competitor q' of p' defined by $q' := q + (p' - p)$, so that q' has the same positive relative to p as q has relative to p . We claim the following: for any voter v and any $1 \leq i \leq d$, we have $|x_i(v) - x_i(p')| - |x_i(v) - x_i(q')| = |x_i(v) - x_i(p)| - |x_i(v) - x_i(q)|$. Since q beats p , this claim implies that

q' beats p' , thus finishing the proof. Indeed, if $x_i(p') = x_i(p)$, we have $x_i(q') = x_i(q)$ and the claim holds. Otherwise, assume without loss of generality that $x_i(p') < x_i(p)$, and recall that q is ε -close to p . By taking ε sufficiently small, we can thus ensure that $\min_i = x_i(p') < x_i(q') < x_i(p) < x_i(q) < \max_i$. Since there are no voters v with $\min_i < x_i(v) < \max_i$, this implies the claim. \square

The following lemma is used in the proof of Theorem 4.3.

LEMMA 4.4. $V[p > q] \subseteq V[p' > q']$ and $V[q' > p'] \subseteq V[q > p]$.

PROOF. The claim follows if we can show that for any voter v ,

$$\sum_{i=1}^d w_i(v) \cdot (|x_i(v) - x_i(p)| - |x_i(v) - x_i(q)|) \geq \sum_{i=1}^d w_i(v) \cdot (|x_i(v) - x_i(p')| - |x_i(v) - x_i(q')|).$$

(Indeed, the sign of the left-hand side is negative if and only if $v \in V[p, q]$ and positive if and only if $v \in V[q, p]$, and similarly for the right-hand side and $V[p', q']$ and $V[q', p']$.) Since all weights are nonnegative, this holds if for all $1 \leq i \leq d$ we have

$$|x_i(v) - x_i(p)| - |x_i(v) - x_i(q)| \geq |x_i(v) - x_i(p')| - |x_i(v) - x_i(q')|. \quad (1)$$

To prove this we consider three cases.

Case 1: $x_i(p) = x_i(p')$. Now $x_i(q) = x_i(q')$ and the inequality in Equation (1) holds with equality.

Case 2: $x_i(v) = x_i(p')$. Now we have

$$\begin{aligned} |x_i(v) - x_i(p)| - |x_i(v) - x_i(q)| &\geq -|x_i(p) - x_i(q)| \text{ (Triangle Inequality)} \\ &= -|x_i(p') - x_i(q')| \\ &= |x_i(v) - x_i(p')| - |x_i(p') - x_i(q')|, \end{aligned}$$

and so the inequality in Equation (1) also holds in this case.

Case 3: $x_i(p) \neq x_i(p')$ and $x_i(v) \neq x_i(p')$. Now $x_i(p)$ lies in the interior of I_i , the interval defining the i th coordinate of M_V . Our choice of ε guarantees that either $x_i(v) < \min(x_i(p), x_i(q))$ or $x_i(v) > \max(x_i(p), x_i(q))$. Similarly, $x_i(v) < \min(x_i(p'), x_i(q'))$ or $x_i(v) > \max(x_i(p'), x_i(q'))$. (Note that $x_i(p)$ is an endpoint of I_i , and so we can also have $x_i(v) = x_i(p')$, but this was covered in Case 2.) Moreover, $x_i(v) < \min(x_i(p), x_i(p'))$ or $x_i(v) > \max(x_i(p), x_i(p'))$. Hence,

$$x_i(v) < \min(x_i(p), x_i(p'), x_i(q), x_i(q')) \text{ or } x_i(v) > \max(x_i(p), x_i(p'), x_i(q), x_i(q')),$$

which proves the inequality in Equation (1) (as in Case 1 with equality) and, hence, the lemma. \square

4.2 Finding all the Plurality Points in the Personalized L_1 Norm

Our algorithm for finding the set of all plurality points is quite simple: we compute the median box M_V in $O(dn)$ time, then we check for each vertex of M_V if it is a plurality point (as described in the proof of the theorem below), and finally we report the set of all plurality points using Theorem 4.3. The following theorem summarizes the result.

THEOREM 4.5. *Let V be a set of n voters in \mathbb{R}^d , where $d \geq 2$. Then we can compute in $O(C_d n^{d-1})$ time the set of all plurality points for V in the personalized L_1 norm. When all voters have the same preferences, the time bound reduces to $O(C_d + d^2 n)$. Here C_d is a constant depending on the dimension d .*

PROOF. Let C'_d be a constant (depending on the dimension d) such that the arrangement defined by a set of n hyperplanes in \mathbb{R}^d can be computed in $O(C'_d n^d)$ time. (The fact that an arrangement

in \mathbb{R}^d can be constructed in this time bound is well known [8, 9], but the constant C'_d has not been explicitly specified.) We will prove the theorem with $C_d := C'_{d-1} 2^{d^2}$.

Below we show that we can test if a given vertex p of M_V is a plurality point in $O(C'_{d-1}(2^d n)^{d-1})$ time (and in $O(C'_{d-1} 2^{d(d-1)} + dn)$ time if all voters have the same preferences), from which the theorem readily follows.

Assume without loss of generality that p lies at the origin. We need to check if for any competitor q we have $|V[p > q]| \geq |V[q > p]|$. By Lemma 4.1, we only have to consider non-degenerate competitors. Such a competitor q beats p if $|V[q > p]| > n/2$. Because p is at the origin, a voter v is in $V[q > p]$ if $\sum_{i=1}^d w_i(v) \cdot (|x_i(v) - x_i(q)| - |x_i(v)|) < 0$. When q is an ε -close competitor for p , we have $|x_i(q)| < \varepsilon$, and then $|x_i(v) - x_i(q)| - |x_i(v)| \in \{-x_i(q), x_i(q)\}$. Hence, whether or not $v \in V[q > p]$ depends on the position of q relative to the hyperplane $h := \sum_{i=1}^d w_i(v) \alpha_i x_i = 0$, where $\alpha_i = +1$ if $x_i(q) > x_i(v)$ and $\alpha_i = -1$ if $x_i(q) < x_i(v)$. Each voter $v \in V$ thus generates a set of 2^d hyperplanes. Let H be the total set of hyperplanes generated, that is, $H := \{\sum_{i=1}^d w_i(v) \alpha_i x_i = 0 : v \in V \text{ and } (\alpha_1, \dots, \alpha_d) \in \{-1, +1\}^d\}$. The discussion above implies that if two competitors q, q' have the same position relative to every hyperplane in H , then $V[q > p] = V[q' > p]$. Hence, we can proceed as follows.

We first compute the set H in $O(2^d n)$ time. Next we compute the arrangement $\mathcal{A}(H)$ defined by the hyperplanes in H . Since all hyperplanes pass through the origin, $\mathcal{A}(H)$ is effectively a $(d-1)$ -dimensional arrangement, so it has complexity $O(2^d (2^d n)^{d-1})$ and it can be constructed in $O(C'_{d-1} (2^d n)^{d-1})$ time [8, 9]. Note that for any cell C of $\mathcal{A}(H)$, we have $V[q > p] = V[q' > p]$ for any two competitors q, q' in C . With a slight abuse of notation we denote this set by $V[C > p]$. The sets $V[C > p]$ and $V[C' > p]$ for neighboring cells C, C' differ by at most one voter (corresponding to the hyperplane that separates the cells).² Hence, we can compute for each cell C of $\mathcal{A}(H)$ the size of $V[C > p]$ in $O(2^d (2^d n)^{d-1})$ time in total, by performing a depth-first search on the dual graph of $\mathcal{A}(H)$ and updating the size as we step from one cell to the next. When we find a cell C with $|V[C > p]| > n/2$, we report that p is not a plurality point; otherwise, we report that p is a plurality point.

When all preferences are equal—after appropriate scaling this reduces to the case where we simply use the standard L_1 norm—then all voters $v \in V$ define the same set of 2^d hyperplanes, and so $|H| = 2^d$. Hence, when all preferences are equal, the algorithm for testing runs in $O(C'_{d-1} 2^{d(d-1)} + dn)$ time. \square

4.3 A Plurality Point Does Not Always Exist in the Personalized L_1 Norm

In this section, we give an example showing that a plurality point doesn't always exist for the personalized L_1 norm in \mathbb{R}^2 .

We define the set V of voters as follows (see Figure 5). First, we add $(1, -1)$ and $(-1, 1)$ to V , with preference vector $\langle 1, 1 \rangle$; these points will define opposite corners of the median box M_V . Then we add 6 points, as follows:

- 3 points in the $(+, -)$ quadrant of voter $(1, -1)$, with preference vector $\langle 2, 1 \rangle$;
- 3 points in the $(-, +)$ quadrant of voter $(-1, 1)$, with preference vector $\langle 1, 2 \rangle$.

²Actually this is not quite true, as several voters could generate the same hyperplane. In this case, the difference between $V[C > p]$ and $V[C' > p]$ can be more than one voter. Thus, the time needed to step from C to C' is linear in the number of voters who generate the separating hyperplane of C and C' . It is easy to see that this does not influence the final time bound.

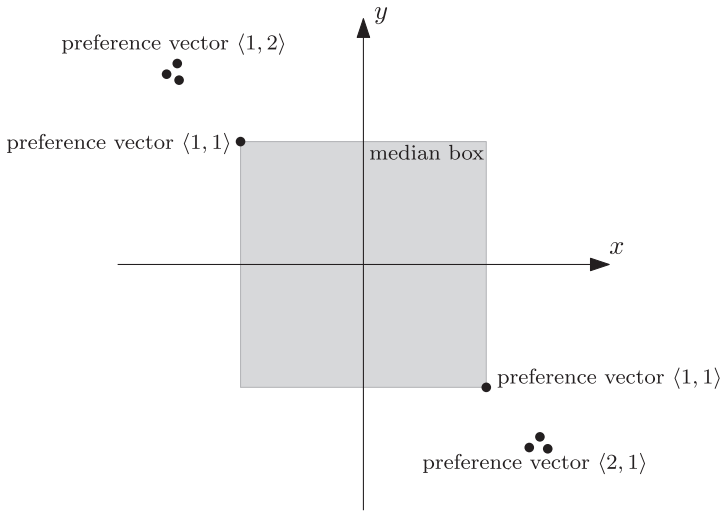


Fig. 5. An example where there are no plurality points in the general case.

Now consider any point $p = (x_1(p), x_2(p))$ inside (or on the border of) the median box. It is evident that the competitor $q = (x_1(p) + \varepsilon, x_2(p) + \varepsilon)$ (for a sufficiently small $\varepsilon > 0$) falsifies that p is a plurality point as all the 6 lately added points belong to $V[q > p]$. Therefore, there will be no plurality point in this setting.

5 CONCLUDING REMARKS

We presented efficient algorithms for a number of problems concerning plurality points. It would be interesting to generalize these to the setting where the voters have weights—not to be confused with the weights defining the personal preferences—and a point is a plurality point if there is no other point that is preferred by a set of voters of higher total weight. This would also allow us to deal with multisets of voters, something our current algorithms cannot do. Another direction for future research is to extend our fixed-parameter algorithm for the minimum-cost problem and the algorithm for plurality balls to the personalized L_1 norm.

REFERENCES

- [1] H. K. Ahn, S. W. Cheng, O. Cheong, M. Golin, and R. van Oostrum. 2004. Competitive facility location: The Voronoi game. *Theor. Comput. Sci.* 310, 1–3 (2004), 457–467.
- [2] N. Amenta. 1994. Helly-type theorems and generalized linear programming. *Discrete Comput. Geom.* 12, 3 (1994), 241–261.
- [3] M. Ben-Or. 1983. Lower bounds for algebraic computation trees. In *Proc. 15th ACM Symp. Theory Comp. (STOC'83)*. 80–86.
- [4] T. Chan. 2004. An optimal randomized algorithm for maximum Tukey depth. In *Proc. 15th ACM-SIAM Symp. Discr. Alg. (SODA'04)*. 430–436.
- [5] B. Chazelle. 1985. On the convex layers of a planar set. *IEEE Trans. Inf. Theory* 31, 4 (1985), 509–517.
- [6] O. Cheong, S. Har-Peled, N. Linial, and J. Matoušek. 2004. The one-round voronoi game. *Discr. Comput. Geom.* 31, 1 (2004), 125–138.
- [7] M. De Berg, M. Van Kreveld, M. Overmars, and O. C. Schwarzkopf. 2000. Computational geometry. In *Computational Geometry*. Springer, 1–17.
- [8] H. Edelsbrunner. 2012. *Algorithms in Combinatorial Geometry*. Vol. 10. Springer Science & Business Media.
- [9] H. Edelsbrunner, J. O'Rourke, and R. Seidel. 1986. Constructing arrangements of lines and hyperplanes with applications. *SIAM J. Comput.* 15, 2 (1986), 341–363.
- [10] H. A. Eiselt and G. Laporte. 1997. Sequential location problems. *Europ. J. Op. Res.* 96, 2 (1997), 217–231.

- [11] J. Erickson. 1999. New lower bounds for convex hull problems in odd dimensions. *SIAM J. Comput.* 28, 4 (1999), 1198–1214.
- [12] A. Gajentaan and M. H. Overmars. 1995. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.* 5, 3 (1995), 165–185.
- [13] T. Gallai. 1944. Solution to problem number 4065. *Am. Math. Monthly* 51, 3 (1944), 169–171.
- [14] P. Hansen and J. F. Thisse. 1981. Outcomes of voting and planning: Condorcet, weber and rawls locations. *J. Public Econ.* 16, 1 (1981), 1–15.
- [15] P. Hansen, J. F. Thisse, and R. E. Wendell. 1986. Equivalence of solutions to network location problems. *Math. Oper. Res.* 11, 4 (1986), 672–678.
- [16] S. Jadhav, A. Mukhopadhyay, and B. Bhattacharya. 1996. An optimal algorithm for the intersection radius of a set of convex polygons. *J. Alg.* 20, 2 (1996), 244–267.
- [17] D. Kress and E. Pesch. 2012. Sequential competitive location on networks. *Europ. J. Op. Res.* 217, 3 (2012), 483–499.
- [18] W. Y. Lin, Y. W. Wu, H. L. Wang, and K. M. Chao. 2015. Forming plurality at minimum cost. In *Proc. 9th Int. Workshop Alg. Comput. (WALCOM'15)*, LNCS 8973. 77–88.
- [19] R. D. McKelvey and R. E. Wendell. 1976. Voting equilibria in multidimensional choice spaces. *Math. Op. Res.* 1, 2 (1976), 144–158.
- [20] C. R. Plott. 1967. A notion of equilibrium and its possibility under majority rule. *Am. Econ. Rev.* 57, 4 (1967), 787–806.
- [21] J. W. Tukey. 1975. Mathematics and the picturing of data. In *Proc. Int. Cong. Mathematicians*, Vol. 2. 523–531.
- [22] R. E. Wendell and R. D. McKelvey. 1981. New perspectives in competitive location theory. *Eur. J. Op. Res.* 6, 2 (1981), 174–182.
- [23] R. E. Wendell and S. J. Thorson. 1974. Some generalizations of social decisions under majority rule. *Econometrica* 42, 5 (1974), 893–912.
- [24] Y. W. Wu, W. Y. Lin, H. L. Wang, and K. M. Chao. 2013. Computing plurality points and Condorcet points in euclidean space. In *Proc. 24th Int. Symp. Alg. Comput. (ISAAC'13)*, LNCS 8283. 688–698.

Received January 2017; revised January 2018; accepted February 2018