

Faster Algorithms for Computing Plurality Points

Mark de Berg^{*1}, Joachim Gudmundsson^{†2}, and Mehran Mehr^{‡3}

1 Department of Computer Science, TU Eindhoven, The Netherlands
mdberg@win.tue.nl

2 School of IT, University of Sydney, Australia
joachim.gudmundsson@sydney.edu.au

3 Department of Computer Science, TU Eindhoven, The Netherlands
mmehr@tue.nl

Abstract

Let V be a set of n points in \mathbb{R}^d , which we call voters, where d is a fixed constant. 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 $O(n)$.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Computational geometry, computational social choice, voting theory, plurality points, Condorcet points

Digital Object Identifier 10.4230/LIPIcs.SoCG.2016.32

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 \mathcal{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 \mathcal{C} which one is preferred by v or whether both choices are equally preferable. A (weak) plurality point is now defined as a choice $p \in \mathcal{C}$ such that no alternative $p' \in \mathcal{C}$ is preferred by more voters.

When there are different issues on which the voters can decide, then the space \mathcal{C} becomes a multi-dimensional space. This has led to the study of plurality points in the setting where $\mathcal{C} = \mathbb{R}^d$ and each voter has an ideal choice which 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

* MdB is supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003 and 022.005025.

† JG supported under Australian Research Council's Discovery Projects funding scheme (project number DP150101134).

‡ MM is supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003 and 022.005025.



so we view each voter $v \in V$ as being a point in \mathbb{R}^d , the so-called *spatial model* in voting theory [15]. 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 [6, 11, 10, 15, 17]. McKelvey and Wendell [15] 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 \mathcal{C} such that she always wins at least as many clients (voters) as her competitor, no matter where the competitor places his 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 [13]. Competitive facility location has also been studied in a geometric, continuous setting under the name Voronoi games [1, 4]. 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 $V = \mathcal{C} = R$. The plurality-point problem in a geometric space lies in between the network setting and the fully continuous setting: the space \mathcal{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 [16]. 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 halfspace 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 $\lceil n/2 \rceil$. Wu *et al.* [19] showed that a point $p \in \mathbb{R}^d$ is a plurality point in the L_2 norm if and only if any open halfspace 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 [2], 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. This led Lin *et al.* [14] 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 [15, 17, 18]. 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 2-dimensional median is a plurality point. Unfortunately, this is no longer true when $d > 2$ [15, 17]. 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 $\lfloor n/2 \rfloor$. Wu *et al.* [19] 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.

We then turn our attention to the minimum-cost plurality problem. We solve the open problem of Lin *et al.* [14] by presenting an algorithm that solves the problem in $O(n^4)$ time, in any (fixed) dimension. 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 the plane: 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.

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 non-negative 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 that Lin *et al.* [14] 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.

2.1 Properties of plurality points in the L_2 norm

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

► **Fact 1** (Wu *et al.* [19]). 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 halfspace with p on its boundary contains at most $n/2$ voters.

Verifying the condition in Fact 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.* [14] 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, then 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 Lin *et al.* [14]. 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 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 . Note that if p has the alternating property then the number of unbalanced lines must be odd.

► **Theorem 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 $V \cap f$.*

For $d = 1$ the theorem is trivial, and for $d = 2$ – the condition in case 3 then simply states that p has the alternating property – the theorem was proved by Lin *et al.* Our contribution is the extension to higher dimensions. Before proving Theorem 2, we need the following lemma regarding the robustness of plurality points to dimension reduction.

► **Lemma 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 $V \cap f$.*

Proof. We prove the statement by induction on d . For $d = 1$ the lemma is trivially true, so now consider the case $d > 1$. We consider two cases.

The first case is that f is a hyperplane, that is, $\dim(f) = d - 1$. Let f^+ and f^- denote the open halfspaces 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. Then all voters in $f^+ \cap V$ end up in \hat{f}^+ . Moreover, we can choose the direction of the rotation such that the voters in $g^+ \cap V$ end up in \hat{f}^+ . But

then $|\hat{f}^+ \cap V| = |f^+ \cap V| + |g^+ \cap V| > |f^+ \cap V| + |f \cap V|/2 \geq n/2$, which contradicts the assumption that p is a plurality point.

The second case is that $\dim(f) < d - 1$. Let h be a hyperplane that contains f . From the first case we know that p must be a plurality point for $h \cap V$. Hence, we can apply our induction hypothesis to conclude that p must be a plurality point for $f \cap V$. ◀

► **Corollary 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 3, both p_1 and p_2 are plurality points for $f \cap V$. But this contradicts the result by Wu *et al.* [19] that any set of voters in the plane admits at most one plurality point. ◀

Now we are ready to prove Theorem 2.

Proof of Theorem 2. Since the case $d = 2$ was already proved by Lin *et al.* [14], and the case $d = 1$ is trivial, we assume $d \geq 3$. Below we prove part 3, the proof for parts 1 and 2 is given in the full version.

(3, ⇐). 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 $V \cap f$. Consider an arbitrary open halfspace h^+ whose bounding hyperplane h contains p , and let h^- be the opposite open halfspace. 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 we have the alternating property on f , we know that p is a plurality point on f . 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 halfspace h^+ we have $|h^+ \cap V| \leq n/2$, and so p is a plurality point.

(3, ⇒). 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 [9] this implies there is an ordinary line in g' , 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 3.

We just argued that all unbalanced lines must lie on a single 2-flat f . By Lemma 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. ◀

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* h for V , which is a hyperplane such that $|h^+ \cap V| < n/2$ and $|h^- \cap V| < n/2$, where h^+ and h^- denote the two open halfspaces bounded by h . The following lemma is easy to prove.

► **Lemma 5.** *Let $v \in V$ be a voter that lies on a hyperplane h_0 such that all voters either lie on h_0 or in h_0^+ . Then we can find a median hyperplane h containing v in $O(n)$ time.*

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, 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 follow.
 - a. Let $v_0 \in V$ be a voter with minimum x_d -coordinate. Find a median hyperplane h_0 containing v_0 using Lemma 5, and let $cand_0 := \text{FINDCANDIDATES}(h_0 \cap V)$.
 - b. If $cand_0$ is a single point or $cand_0 = \emptyset$ then return $cand_0$.
 - c. If $cand_0$ is a (non-degenerate) segment then let $v_1 \in V$ be a voter whose distance to h_0 is maximized. Find a median hyperplane h_1 containing v_1 using Lemma 5, and let $cand_1 := \text{FINDCANDIDATES}(h_1 \cap V)$. Return $cand_0 \cap cand_1$.

► **Lemma 6.** *Algorithm FINDCANDIDATES(V) returns in $O(n)$ 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 $|h_0^+ \cap V| < n/2$ and $|h_0^- \cap V| < n/2$, for any point $p \notin h$ there is an open halfspace containing p and bounded by a hyperplane parallel to h_0 that contains more than $n/2$ voters. Hence, by Fact 1 any plurality point for V must lie on h_0 . By Lemma 3, if a plurality point exists for V it must also be a plurality point for $h_0 \cap V$. By induction we can assume that $\text{FINDCANDIDATES}(h_0 \cap V)$ is correct. Hence, the result of the algorithm is correct when $cand_0$ is a single point or $cand_0 = \emptyset$. Note that when $cand_0$ is a (non-degenerate) segment – this only happens when all voters in $h_0 \cap V$ are collinear – we must have $V \neq h_0 \cap V$, otherwise V would be collinear and we would be done after Step 1. Hence, $v_1 \notin h_0$. By the same reasoning as above the median hyperplane h_1 must contain the plurality point of V (if it exists). But then the plurality point must lie in $cand_0 \cap cand_1$, and since $v_1 \notin h_0$ we know that $cand_0 \cap cand_1$ 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 non-degenerate candidate segment. This only happens when all voters in $h_0 \cap V$ are collinear, which implies the recursive call just needs to compute a median segment in $O(n)$ 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 $O(n)$ time, the bound follows. ◀

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 7.** *Given a set V of n voters in \mathbb{R}^d and a candidate point p , we can test in $O(n \log n)$ 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 $L(v)$, and for each line $\ell \in L(p)$ the number of voters on the rays $\rho(\ell)$ and $\bar{\rho}(\ell)$, in $O(n \log n)$ 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, 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 $O(n)$ 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 $O(n \log n)$ time. ◀

We obtain the following theorem. (See the full version for the $\Omega(n \log n)$ lower bound.)

► **Theorem 8.** *Let V be a set of n voters in \mathbb{R}^d , where $d \geq 2$ is a fixed constant. Then we can find in $O(n \log n)$ time the plurality point for V in the L_2 norm, if it exists, and this time bound is optimal.*

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 one of the voters, 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 22 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 9** (Lin et al. [14]). *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 .*

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 . This can be done in $O(n^4)$ time in total, by projecting all lines onto an arbitrary 2-flat, constructing the arrangement in this 2-flat in $O(n^4)$ time [5], and then checking which intersections on the 2-flat correspond to actual intersections in \mathbb{R}^d . 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 $O(n^3)$ time per voter – and take the cheapest of all n possibilities.

Consider a candidate p coinciding with some $v \in V$. By Theorem 2 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 9. This takes $O(n_f \log n_f)$ time, where $n_f := |V \cap f|$.
 - 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 and the discussion above. As for the running time, we note that Steps 1, 2, and 4 all run in $O(n^4)$ time. For Step 3, the time needed to compute the price of a single voter v is $\sum_f O(n_f \log n_f + n)$, which is bounded by $O(n^3 + \sum_f n_f \log n_f)$. 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 $O(n^4)$ time.

► **Theorem 10.** *Let V be a set of n voters in \mathbb{R}^d , each with a positive cost, where $d \geq 2$ is a fixed constant. Then we can compute in $O(n^4)$ 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 [8] and has an $\Omega(n^2)$ lower bound if only sidedness tests are used [7].

► **Theorem 11.** *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.*

3.2 An output-sensitive algorithm for unit-cost voters in the plane

The proof of Theorem 11 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, in the planar case and for unit-cost voters. More precisely, we consider the case where we are given a set V of unit-cost voters in the plane 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 both open halfplanes bounded by ℓ 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 12.** *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. Hence, $\ell(p, v) \in L(v)$. If $\ell(p, v) \notin L_k(v)$ then there is a halfplane bounded by $\ell(p, v)$ containing 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 $\ell(p, v) \in L_k(v)$. ◀

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 – this must then be the line through v_{2i-1} and v_{2i} – that is present in both $L_k(v_{2i-1})$ and $L_k(v_{2i})$. This case is handled using the following lemma.

32:10 Faster Algorithms for Computing Plurality Points

► **Lemma 13.** *Suppose we want to compute a cheapest plurality point on the line $\ell := \ell(v, v')$ through two voters $v, v' \in V$, where we are only interested in points of price at most k . Let ℓ^+ be any of the two open halfplanes bounded by ℓ , and let $V^+ := \ell^+ \cap V$. Then p is either an intersection point of ℓ and a line in $\bigcup_{v'' \in V^+} L_k(v'')$, or the point coinciding with any median 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 12. Otherwise all voters in V^+ are in W_p . But then any line through p and a voter in $\ell^- \cap V$ is unbalanced, which implies that all voters in $\ell^- \cap V$ except at most one must be in W_p . In this case the point coinciding with any of the at most two medians along ℓ is a cheapest plurality point on ℓ . ◀

We now present our algorithm. For technical reasons we assume $k \leq n/15$.

OUTPUTSENSITIVEMINCOSTPLURALITYPOINT(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. (We are guaranteed we can collect this many pairs since $k \leq n/15$.)
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 halfplane 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 9. Return the cheapest plurality point $p^* \in C$, provided $|W_{p^*}| \leq k$; if $|W_p| > k$ for all candidates, then report that it is not possible to obtain a plurality point by removing at most k voters.

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

► **Lemma 14.** *Let v be a voter of some pair in P . Then $|L_k(v)| = O(k)$.*

We can now prove the following theorem.

► **Theorem 15.** *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.*

Proof. To prove the time bound, we note that we can compute the convex layers in $O(n \log n)$ time [3]. Step 2 runs in $O(n)$ time. Computing the set $L_k(v)$ for a voter v can easily be done in $O(n \log n)$ time. By Lemma 14, 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 12 and 13, 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. 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. ◀

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 $b(p, r/2)$. Fact 1 can be generalized as follows.

► **Fact 16.** A ball $b(p, r)$ is a plurality ball if and only if every open halfspace 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 halfspace 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 , and consider the k -level in the arrangement $\mathcal{A}(V^*)$, that is, the set of points on the hyperplanes in V^* that have exactly k hyperplanes strictly below them. We associate each $(d - 1)$ -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 . 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 17.** 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 Fact 16 there must be an open halfspace 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 halfspace, which is a contradiction. On the other hand, if $b(p, r/2)$ does not intersect some $\lfloor n/2 \rfloor$ -cone $\text{cone}(f)$ then there is an open halfspace h^+ not

intersecting $b(p, r/2)$ containing all the points in $\text{cone}(f)$. Therefore $|h^+ \cap V| \geq \lfloor n/2 \rfloor + 1$, and so $b(p, r)$ is not a plurality ball. \blacktriangleleft

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 [12] we obtain the following result.

► **Theorem 18.** *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 $O(n \log n + m \log^{1+\varepsilon} n)$ time, where m is the complexity of the level. Our algorithm then runs in $O(n^{4/3} \log^{1+\varepsilon} n)$ time. We believe the algorithm described above can be generalized to higher dimensions, using a standard algorithm for computing k -levels, and generalized linear programming for the second part of the algorithm. We are currently verifying the details.

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 non-negative 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 [17]. Interestingly, in the personalized L_1 norm the statement already fails in the plane: in the full version we given 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 and Lemma 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 $V[p \succ q] := \{v \in V : \text{dist}_w(v, p) < \text{dist}_w(v, q)\}$. We also define $V[p \sim q] := \{v \in V : \text{dist}_w(v, p) = \text{dist}_w(v, q)\}$ and $V[p \succcurlyeq q] := V[p \succ q] \cup V[p \sim q]$. Let p be a candidate plurality point. We call a point q a *non-degenerate competitor for p* if $V[p \sim q] = \emptyset$, and we say that q is ε -close to p if $|pq| < \varepsilon$, 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 non-degenerate competitors that are ε -close to p .

► **Lemma 19.** *Let p be a candidate plurality point and let q be a competitor of p . For any $\varepsilon > 0$ there is a non-degenerate competitor q' that is ε -close to p such that $|V[q' \succ p]| \geq |V[q \succ p]| + \frac{1}{2} \cdot |V[q \sim p]|$.*

The following lemma helps us to narrow down our search for plurality points. Recall that a *multi-dimensional 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 20.** *Let p be a plurality point for V in the personalized L_1 norm. Then p is a multi-dimensional median for V .*

The set M_V of all multi-dimensional medians for V is an axis-aligned hyperrectangle in \mathbb{R}^d , that is, it can be written as $M_V = I_1 \times \cdots \times I_d$, 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 20 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 facets of M_V for $0 \leq k \leq d$, where each facet $f \in F$ is considered relatively open. Note that $|F| = 3^{d'}$, where d' is the number of non-degenerate intervals defining M_V .

► **Theorem 21.** *Let V be a set of voters in \mathbb{R}^d and let f be a relatively open facet of the median box M_V of V .*

- (1) *Either all points in f are plurality points in the personalized L_1 norm, or none of the points in f are.*
- (2) *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 (1) follows immediately from part (2). Next we prove part (2).

(2, \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 $\varepsilon > 0$ 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 Euclidean balls of radius $\min(\varepsilon, \varepsilon')$ and centered at p and p' , respectively. Since p is a plurality point we know that $|V[p \succ q]| \geq |V[q \succ p]|$ 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 the full version, this implies that $V[p \succ q] \subseteq V[p' \succ q']$ and $V[q' \succ p'] \subseteq V[q \succ p]$. Hence, $|V[p' \succ q']| \geq |V[p \succ q]| \geq |V[q \succ p]| \geq |V[q' \succ p']|$. Since q' is an arbitrary point in b' , the point p' must be a plurality point according to Lemma 19.

(2, \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 $\varepsilon > 0$) who 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 $M_V = I_1 \times \cdots \times I_d$, 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. ◀

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 $O(n)$ 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 21. The following theorem summarizes the result.

► **Theorem 22.** *Let V be a set of n voters in \mathbb{R}^d , where $d \geq 2$ is a fixed constant. Then we can compute in $O(n^{d-1})$ 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 $O(n)$.*

Proof. Below we show that we can test if a given vertex p of M_V is a plurality point in $O(n^{d-1})$ time (and in $O(n)$ 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 $|V[p \succ q]| \geq |V[q \succ p]|$. By Lemma 19 we only have to consider non-degenerate competitors. Such a competitor q beats p if $|V[q \succ p]| > n/2$. Because p is at the origin, a voter v is in $V[q \succ p]$ 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 of 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 $v \in V[q \succ p]$ 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, $H := \left\{ \sum_{i=1}^d w_i(v) \alpha_i x_i = 0 : v \in V \text{ and } (\alpha_1, \dots, \alpha_d) \in \{-1, +1\}^d \right\}$. The discussion above implies that if two competitors q, q' have the same position relative to every hyperplane in H , then $V[q \succ p] = V[q' \succ p]$. Hence, we can proceed as follows.

We first compute the set H in $O(n)$ 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 $O(n^{d-1})$ and it can be constructed in $O(n^{d-1})$ time [5]. Note that for any cell C of $\mathcal{A}(H)$, we have $V[q \succ p] = V[q' \succ p]$ for any two competitors q, q' in C . With a slight abuse of notation we denote this set by $V[C \succ p]$. The sets $V[C \succ p]$ and $V[C' \succ p]$ for neighboring cells C, C' differ by at most one voter (corresponding to the hyperplane that separates the cells).¹ Hence, we can compute for each cell C of $\mathcal{A}(H)$ the size of $V[C \succ p]$ in $O(n^{d-1})$ time, 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 $|V[C \succ p]| > n/2$ 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, the algorithm runs in $O(n)$ time. ◀

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

¹ Actually this is not quite true, as several voters could generate the same hyperplane. In this case the difference between $V[C \succ p]$ and $V[C' \succ p]$ 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.

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 multi-sets of voters, something which our current algorithms cannot do. Another direction for future research is to extend our output-sensitive algorithm for the minimum-cost problem and for plurality balls to higher dimensions, and to the personalized L_1 norm.

References

- 1 H. K. Ahn, S. W. Cheng, O. Cheong, M. Golin, and R. van Oostrum. Competitive facility location: the Voronoi game. *Theor. Comput. Sci.*, 310(1–3):457–467, 2004.
- 2 T. Chan. An optimal randomized algorithm for maximum tukey depth. In *Proc. 15th ACM-SIAM Symp. Discr. Alg. (SODA)*, pages 430–436, 2004.
- 3 B. Chazelle. On the convex layers of a planar set. *IEEE Trans. Inf. Theory*, 31(4):509–517, 1985.
- 4 O. Cheong, S. Har-Peled, N. Linial, and J. Matousek. The one-round voronoi game. *Discr. Comput. Geom.*, 31(1):125–138, 2004.
- 5 H. Edelsbrunner, J. O’Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. *SIAM J. Comput.*, 15(2):341–363, 1986.
- 6 H. A. Eiselt and G. Laporte. Sequential location problems. *Europ. J. Op. Res.*, 96(2):217–231, 1997.
- 7 J. Erickson. New lower bounds for convex hull problems in odd dimensions. *SIAM J. Comput.*, 28(4):1198–1214, 1999.
- 8 A. Gajentaan and M. H. Overmars. On a class of $o(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5(3):165–185, 1995.
- 9 T. Gallai. Solution to problem number 4065. *The American Mathematical Monthly*, 51(3):169–171, 1944.
- 10 P. Hansen and J. F. Thisse. Outcomes of voting and planning: Condorcet, weber and rawls locations. *Journal of Public Economics*, 16(1):1–15, 1981.
- 11 P. Hansen, J. F. Thisse, and R. E. Wendell. Equivalence of solutions to network location problems. *Mathematics of Operations Research*, 11(4):672–678, 1986.
- 12 S. Jadhav, A. Mukhopadhyay, and B. Bhattacharya. An optimal algorithm for the intersection radius of a set of convex polygons. *J. Alg.*, 20(2):244–267, 1996.
- 13 D. Kress and E. Pesch. Sequential competitive location on networks. *Europ. J. Op. Res.*, 217(3):483–499, 2012.
- 14 W. Y. Lin, Y. W. Wu, H. L. Wang, and K. M. Chao. Forming plurality at minimum cost. In *Proc. 9th Int. Workshop Alg. Comput. (WALCOM)*, LNCS 8973, pages 77–88, 2015.
- 15 R. D. McKelvey and R. E. Wendell. Voting equilibria in multidimensional choice spaces. *Math. Op. Res.*, 1(2):144–158, 1976.
- 16 J. W. Tukey. Mathematics and the picturing of data. In *In Proc. Int. Cong. Mathematicians*, volume 2, pages 523–531, 1975.
- 17 R. E. Wendell and R. D. McKelvey. New perspectives in competitive location theory. *Europ. J. Op. Res.*, 6(2):174–182, 1981.
- 18 R. E. Wendell and S. J. Thorson. Some generalizations of social decisions under majority rule. *Econometrica*, 42(5):893–912, 1974.
- 19 Y. W. Wu, W. Y. Lin, H. L. Wang, and K. M. Chao. Computing plurality points and condorcet points in euclidean space. In *Proc. 24th Int. Symp. Alg. Comput. (ISAAC)*, LNCS 8283, pages 688–698, 2013.