Finding Plurality Points in $\mathbb{R}^{d*}$

Mark de Berg[†] Joachim Gudmundsson[‡] Mehran Mehr[§]

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.

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. Thus 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 $p' \in C$ 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 so we view each voter $v \in V$ as being a point in $\mathbb{R}^d$, the so-called *spatial model* in voting theory [10]. 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 $|\{v \in V : \text{dist}(v,p) < \text{dist}(v,p')\}| \geqslant |\{v \in V : \text{dist}(v,p') < \text{dist}(v,p)\}|$.

Plurality points and related concepts were already studied in the 1970s in voting theory [4, 6, 7, 10, 12]. McKelvey and Wendell [10] 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 [8]. Competitive facility location has also been studied in a geometric, continuous setting under the name Voronoi games [1, 3]. 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 [11]. 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.* [13] showed that a point $p \in \mathbb{R}^d$ is a plurality point in the L_2 norm if

*MdB and MM are supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003 and 022.005025, respectively. JG supported under Australian Research Council's Discovery Projects funding scheme (project number DP150101134).

†Department of Computer Science, TU Eindhoven, the Netherlands, mdberg@win.tue.nl

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

§Department of Computer Science, TU Eindhoven, the Netherlands, mmehr@tue.nl

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 given set V of n voters in $\mathbb{R}^d$ admits a plurality point with respect to the L_2 norm and, if so, finds such a point. 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.

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.* [13] 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.

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.* [9] 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.* [13]) *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.* [9] 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 $\overline{\rho}(\ell)$. (The point p itself is not part of these rays.) We say that a line $\ell \in L(p)$ is *balanced* if $|\rho(\ell) \cap V| = |\overline{\rho}(\ell) \cap V|$. If n is odd, then p is 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 $\overline{\rho}(\ell)$. 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 $\overline{\rho}(\ell)$ 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.* [9]. 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 $d \geqslant 1$, 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 c then simply states that p has the alternating property—the theorem was proved by Lin *et al.* [9] 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 $d \geqslant 1$, 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 $|f^+ \cap V| \geqslant |f^- \cap V|$. 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 $|\hat{f}^+ \cap V| = |f^+ \cap V| + |g^+ \cap V| > |f^+ \cap V| + |f \cap V|/2 \geqslant n/2$, 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 $f \cap V$. ☐

Now we are ready to prove Theorem 2.

Proof. [Proof of Theorem 2] Since the case $d = 2$ was already proved by Lin *et al.* [9], and the case $d = 1$ is trivial, we assume $d \geqslant 3$. Below we prove part c. The proof of parts a and b is given in the full version.

(c,⇐). 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 $|h^+ \cap V| \leqslant n/2$. 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 $|h^+ \cap V| \leqslant n/2$, because all voters not in f lie on balanced lines. We conclude that for any open halfspace h^+ we have $|h^+ \cap V| \leqslant n/2$, and so p is a plurality point.

(c,⇒). 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 $L^*(g) \subset L^*(p)$ 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 $L^*(g)$. Each of the lines in $L^*(g)$ intersects f_1 in a single point, and these intersection points are not all collinear. According to the Sylvester-Gallai Theorem [5] 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 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 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 4 *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.*

For $d \geqslant 2$ the plurality point is unique, if it exists (the proof is given in the full version). 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 4, and let $cand_0 :=$ 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 4, and let $cand_1 :=$ FindCandidates($h_1 \cap V$). Return $cand_0 \cap cand_1$.

Lemma 5 *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) otherwise* cand *contains at most one point, and no other point can be a plurality point of V.*

Proof. 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 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 exist). 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. $\square$

Our algorithm to find a plurality point first calls FindCandidates(V). If all points in V are collinear we are done—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.

To this end we have to check the conditions of Theorem 2, which can easily be done in $O(n \log n)$ time.

Lemma 6 *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.*

We obtain the following theorem.

Theorem 7 *Let V be a set of n voters in $\mathbb{R}^d$, where $d \geqslant 2$ 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 Conclusion

Most point sets do not admit a plurality point in the L_2 norm. Hence, in the full version of the paper we also study several other problems concerning plurality points: we give fast algorithms to find the smallest subset $W \subset V$ such that $V \setminus W$ admits a plurality point, we show how to compute a so-called plurality ball in the plane, and we study plurality points in the 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] O. Cheong, S. Har-Peled, N. Linial, and J. Matousek. The one-round Voronoi game. *Discr. Comput. Geom.* 31(1):125–138, 2004.

[4] H. A. Eiselt and G. Laporte. Sequential location problems. *Europ. J. Op. Res.* 96(2):217–231, 1997.

[5] T. Gallai. Solution to problem number 4065. *The American Mathematical Monthly* 51:169–171, 1944.

[6] P. Hansen, J.-F. Thisse, and R. E. Wendell. Equivalence of solutions to network location problems. *Math. Op. Res.* 11(4):672–678, 1986.

[7] 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.

[8] D. Kress and E. Pesch. Sequential competitive location on networks. *Europ. J. Op. Res.* 217(3):483–499, 2012.

[9] 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.

[10] R.D. McKelvey and R.E .Wendell. Voting equilibria in multidimensional choice spaces. *Math. Op. Res.* 1(2):144–158, 1976.

[11] J. W. Tukey. Mathematics and the picturing of data. In *Proc. Int. Cong. Mathematicians*, volume 2, pages 523–531, 1975.

[12] R.E. Wendell and R.D. McKelvey. New perspectives in competitive location theory. *Europ. J. Op. Res.* 6(2):174–182, 1981.

[13] 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.