On the Distortion of Single Winner Elections with Aligned Candidates

JAAMAS Track

Dimitris Fotakis
National Technical University of Athens
Athens, Greece
fotakis@cs.ntua.gr

Laurent Gourvès
Université Paris Dauphine-PSL, CNRS, LAMSADE
Paris, France
laurent.gourves@dauphine.fr

ABSTRACT

We study the problem of selecting a single element from a set of candidates on which a group of agents has some spatial preferences. The exact distances between agent and candidate locations are unknown but we know how agents rank the candidates from the closest to the farthest. Whether it is desirable or undesirable, the winning candidate should either minimize or maximize its aggregate distance to the agents. The goal is to understand the optimal distortion, which evaluates how good an algorithm that determines the winner based only on the agent rankings performs against the optimal solution. We give a characterization of the distortion in the case of latent Euclidean distances such that the candidates are aligned, but the agent locations are not constrained. This setting generalizes the well-studied setting where both agents and candidates are located on the real line. Our bounds on the distortion are expressed with a parameter which relates, for every agent, the distance to her best candidate to the distance to any other alternative.

KEYWORDS

Distortion; Single Winner Election; Obnoxious Facility

ACM Reference Format:

Dimitris Fotakis and Laurent Gourvès. 2023. On the Distortion of Single Winner Elections with Aligned Candidates: JAAMAS Track. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 CONTEXT

The problem of electing a set of representatives is central in social choice theory. Some voters (a.k.a. agents) express their preferences over a set of candidates and one has to aggregate the voters' preferences to identify the winners (see e.g., [28]). In typical voting scenarios, the voters can only express *ordinal* preferences over the candidates, which are consistent and summarize their *cardinal* preferences.

In a recent stream of articles (see, for example, [3] for a recent survey), researchers study problems where some agents have latent distances over a set of candidates but these distances are unknown. Nevertheless, each agent has reported a ranking of the candidates. Though these rankings are consistent with the latent distance function, we are not guaranteed to find the candidates whose aggregate distance to the agents is minimum, even if we aim to choose a single candidate [4].

Similar to the approximation ratio [27], the *distortion* measures the worst-case performance of an algorithm due to lack of cardinal information [6, 25]. The intriguing question of determining the best distortion for selecting a single winner (called the *metric distortion problem*) has attracted a lot of attention [1, 2, 24]. Gkatzelis et al. [16] proposed a deterministic algorithm with distortion 3, which is optimal because no deterministic algorithm has distortion less than 3 [1, 2]. Determining the best possible distortion for randomized algorithms is considered as a major open problem [3]. Nevertheless, the special cases of 2 and 3 candidates are resolved [4, 8, 19, 20].

More insight into the problem can be gained when more information on the instance is available. In this respect, *α-decisiveness*, where α is a real in $[0, 1]$, plays a key role [4]. This parameter captures how much more the agents prefer their best candidate to any other alternative. In an α-decisive instance, every agent's distance to her closest candidate is at most the distance to her second closest candidate multiplied by α. Then, every agent is co-located with her top choice when $\alpha = 0$. For the other extreme ($\alpha = 1$), α-decisiveness does not constrain the agents' locations at all.

The algorithm of Gkatzelis et al. has distortion $2+\alpha$ for α-decisive instances with at least 3 candidates [16]. The deterministic lower bound of 3, which relies on a two-candidate instance, can be extended to show that when the number of candidates m is at least 2, no deterministic algorithm has α-distortion less than $1 + 2\alpha$. The upper and lower bounds do not match anymore under the α-decisiveness framework, but Gkatzelis et al. proposed a lower bound which approaches $2 + \alpha$ when the number of candidates m tends to infinity [16]. When $m = 2$, the deterministic algorithm which outputs the top choice of a majority of agents has distortion $1 + 2\alpha$ [4, 16]. Regarding randomized algorithms parameterized by α, the best lower and upper bounds, for any number of candidates m, are $2 + \alpha - 2(1 - \alpha)/m$ and $2 + \alpha - 2/m$, respectively [16].

Besides α-decisiveness, the metric distortion problem has been studied in the well-known case where agents and candidates are located on a real number line. The locations are unknown but the agents rank the candidates from the closest to the farthest. The preferences induced by this setting (a.k.a. 1-Euclidean because the distances are Euclidean and there is only one dimension) possess nice properties (namely, *single-peakedness* [5] and *single-crossingness* [18, 23]) which can be favorably exploited by an algorithm. Anshelevich and Postl proposed a randomized algorithm with an optimal distortion of $1+\alpha$ for α-decisive instances on a line [4]. They exploit the possibility to efficiently identify a set of (at most) two candidates which are consecutive on the line and to which the optimum must

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

belong. Regarding deterministic algorithms, the aforementioned lower bound of $1+2\alpha$ deriving from the lower bound of 3, applies to the case where agents and candidates are on a line. On the contrary, the candidates are not aligned in the lower bound approaching $2+\alpha$ presented in [16].

Elections share similarities with k-median and facility location problems [7, 12, 14]. The goal is to choose a subset of candidate locations where desirable facilities (e.g., schools) can be built. The total distance to some given agent set has to be minimized, assuming that each agent is connected to the nearest facility. Sometimes, the candidate to be selected is undesirable (e.g., a garbage depot or a candidate to leave a group of people). In this case, one wants to select a candidate of *maximum* total distance to the agents (see [11] for a recent survey on *obnoxious* facility location). Obnoxious facility location problems have previously received attention from several viewpoints. In a "pure" optimization framework, one wants to choose the location of the facilities and the true distances are accessible (see e.g., [17, 26] and the references therein). In the field of algorithmic mechanism design, the agents may misreport their preferences over the set of candidates so as maximize their individual distance to the winner(s). The authors of [10, 21, 22] pursue the goal of designing (group) strategyproof mechanisms[1] with the best possible approximation ratio. Recently, Chen et al. [9] studied the distortion of algorithms in a setting where the location of the candidates is known but the location of every agent is private.[2] They resolved the deterministic case for which the best distortion is 3. For randomized algorithms, a general lower bound of 1.5 is given, together with upper bounds for well studied special cases. In particular, they proposed two randomized mechanisms for building a single facility on the real line. The first mechanism is strategyproof and its distortion is 2. The second mechanism is not strategyproof but its distortion is lower: 13/7.

2 CONTRIBUTION

Interested readers can refer to [15] for the full version of this article.

We consider the metric distortion problem in α-decisive instances. The distances between agent and candidate locations are unknown but every agent has reported a strict preference over the candidate set. The influence of α-decisiveness on the agents' locations is clear when $\alpha = 0$ or $\alpha = 1$, but no previous work precisely explains (to our best knowledge) how α-decisiveness rules the agents' locations when $\alpha \in (0, 1)$. Our first contribution is to fill this gap by showing that agents lie inside some spheres under Euclidean distances. This characterization is interesting on its own and we exploit it in the full version of this article [15].

Our second contribution is the definition of a domain which generalizes the well-studied case where both agents and candidates are located on the real line (1-Euclidean). In this generalization called AC for "Aligned Candidates", the candidates are aligned but the agent locations are not constrained. As for the 1-Euclidean case, the distances in the AC setting are Euclidean. As an application of AC, one can think of a straight road that crosses a region. The agents can be located anywhere in the region but the candidates must be along the road. One can also interpret the AC domain from an electoral perspective: every candidate lies on a left right political axis while the voters' ideological positions are more complex and require more dimensions.

We demonstrate that, as for the 1-Euclidean domain and under the assumption that no agent is equidistant from two distinct candidates, preferences remain single-peaked and single-crossing under the AC domain. Hence, when one wants to select a desirable candidate to which the agents want to be as close as possible, the set of potential optima can be reduced to two contiguous candidates, as for the 1-Euclidean case [4, 13]. Since the metric distortion problem is resolved when $m = 2$ by selecting the candidate supported by a majority of the agents [4, 16], we get a deterministic algorithm with distortion at most $1+2\alpha$ for any number of aligned candidates. This is the best possible ratio because the aforementioned lower bound of $1 + 2\alpha$ applies to the setting of aligned candidates.

Afterwards, we investigate the distortion of choosing a single *undesirable* candidate. The aim is to determine the candidate that *maximizes* the total distance to the agents. As opposed to the setting studied in [9], we do not assume that the location of the candidates are public. We generalize the notion of α-decisiveness to the case of selecting an undesirable candidate. Namely, an instance is $\bar{\alpha}$-decisive, for some $\bar{\alpha} \in [0, 1]$, if every agent prefers her best candidate (now, this is the farthest one) at least $1/\bar{\alpha}$ times more than her second best (the second farthest). Though this definition reads similar to that of decisiveness for a desirable facility, $\bar{\alpha}$-decisiveness constrains the instances in a very different way. We obtain tight bounds on the distortion of undesirable single winner election by deterministic algorithms, as a function of $\bar{\alpha}$, in two cases. When there are only two candidates and the latent distance function d is a metric (d is not necessarily Euclidean), we show that the simple algorithm which outputs the candidate ranked last by a majority of agents has distortion $1 + 2\bar{\alpha}$ and this is the best possible ratio.

We finally consider the AC domain with any number of candidates. As for the case of selecting a desirable candidate, the set of possible optima of the undesirable case with aligned candidates reduces to (at most) two elements which can be efficiently identified from the preference profile. However, these possible optima are typically different. All our bounds on the distortion for selecting an undesirable alternative are tight and summarized in the following table when $m > 2$ (the distortion is $1 + 2\bar{\alpha}$ when $m = 2$).

$\bar{\alpha} < \frac{1}{3}$	$\frac{1}{3} \leq \bar{\alpha} \leq \sqrt{2} - 1$	$\sqrt{2} - 1 < \bar{\alpha}$
1	$\frac{3\bar{\alpha}-\bar{\alpha}^2}{2-3\bar{\alpha}-\bar{\alpha}^2}$	$1 + 2\bar{\alpha}$

Regarding these bounds, note that $\frac{3\bar{\alpha}-\bar{\alpha}^2}{2-3\bar{\alpha}-\bar{\alpha}^2} = 1$ when $\bar{\alpha} = 1/3$, $\frac{3\bar{\alpha}-\bar{\alpha}^2}{2-3\bar{\alpha}-\bar{\alpha}^2} = 1 + 2\bar{\alpha}$ when $\bar{\alpha} = \sqrt{2} - 1$, and $\frac{3\bar{\alpha}-\bar{\alpha}^2}{2-3\bar{\alpha}-\bar{\alpha}^2} < 1 + 2\bar{\alpha}$ for all $\bar{\alpha} \in [1/3, \sqrt{2} - 1)$. Since $\bar{\alpha} \in [0, 1]$, the distortion is always below 3, which is consistent with the results of [9].

In conclusion, all our bounds on the distortion, for both selecting a desirable or undesirable candidate, are best possible and derive from the same simple algorithm: identify a set of two candidates containing the optimum and return the one that is supported by a majority of agents.

[1] There is no incentive for a single agent or a group of agents to misreport their true rankings.

[2] In the present work, the location of the agents and the candidates are private.

ACKNOWLEDGMENTS

Laurent Gourvès is supported by Agence Nationale de la Recherche (ANR), project THEMIS ANR-20-CE23-0018.

REFERENCES

[1] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.

[2] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. 2015. Approximating Optimal Social Choice under Metric Preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence.* AAAI Press, Palo Alto, California, 777–783.

[3] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021.* ijcai.org, California, 4294–4301.

[4] Elliot Anshelevich and John Postl. 2017. Randomized Social Choice Functions Under Metric Preferences. *Journal of Artificial Intelligence Research* 58 (2017), 797–827.

[5] Duncan Black. 1948. On the rationale of group decision-making. *Journal of Political Economy* 56, 1 (1948), 23–34.

[6] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227 (2015), 190–213.

[7] Jaroslaw Byrka, Thomas W. Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. 2017. An Improved Approximation for k-Median and Positive Correlation in Budgeted Optimization. *ACM Transactions on Algorithms* 13, 2 (2017), 23:1–23:31.

[8] Moses Charikar and Prasanna Ramakrishnan. 2022. Metric Distortion Bounds for Randomized Social Choice. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022.* SIAM, Philadelphia, 2986–3004.

[9] Xujin Chen, Minming Li, and Chenhao Wang. 2020. Favorite-Candidate Voting for Eliminating the Least Popular Candidate in a Metric Space. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020.* AAAI Press, Palo Alto, California, 1894–1901.

[10] Yukun Cheng, Wei Yu, and Guochuan Zhang. 2013. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science* 497 (2013), 154–163.

[11] Richard L. Church and Zvi Drezner. 2022. Review of obnoxious facilities location problems. *Computers & Operations Research* 138 (2022), 105468.

[12] Zvi Drezner and Horst W. Hamacher. 2004. *Facility Location: Applications and Theory.* Springer, Switzerland.

[13] Edith Elkind and Piotr Faliszewski. 2014. Recognizing 1-Euclidean Preferences: An Alternative Approach. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT 2014).* Springer, Switzerland, 146–157.

[14] Michal Feldman, Amos Fiat, and Iddan Golomb. 2016. On Voting and Facility Location. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016.* ACM, New York City, 269–286.

[15] Dimitris Fotakis and Laurent Gourvès. 2022. On the distortion of single winner elections with aligned candidates. *Auton. Agents Multi Agent Syst.* 36, 2 (2022), 37. https://doi.org/10.1007/s10458-022-09567-5

[16] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the Optimal Metric Distortion Conjecture. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020.* IEEE, New York City, 1427–1438.

[17] Sara Hosseini and Ameneh Moharerhaye Esfahani. 2009. *Obnoxious Facility Location.* Physica-Verlag HD, Heidelberg, 315–345.

[18] Samuel Karlin. 1968. *Total Positivity.* Stanford University Press, California.

[19] David Kempe. 2020. An Analysis Framework for Metric Voting based on LP Duality. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 2020).* AAAI Press, Palo Alto, California, 2079–2086.

[20] David Kempe. 2020. Communication, Distortion, and Randomness in Metric Voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 2020).* AAAI Press, Palo Alto, California, 2087–2094.

[21] Lili Mei, Deshi Ye, and Guochuan Zhang. 2018. Mechanism design for one-facility location game with obnoxious effects on a line. *Theoretical Computer Science* 734 (2018), 46–57.

[22] Lili Mei, Deshi Ye, and Yong Zhang. 2018. Approximation strategy-proof mechanisms for obnoxious facility location on a line. *Journal of Combinatorial Optimization* 36, 2 (2018), 549–571.

[23] James A. Mirrlees. 1971. An exploration in the theory of optimal income taxation. *Review of Economic Studies* 38 (1971), 175––208.

[24] Kamesh Munagala and Kangning Wang. 2019. Improved Metric Distortion for Deterministic Social Choice Rules. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019.* ACM, New York City, 245–262.

[25] Ariel D. Procaccia and Jeffrey S. Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *Cooperative Information Agents X, 10th International Workshop, CIA 2006, Edinburgh, UK, September 11-13, 2006, Proceedings.* Springer, Switzerland, 317–331.

[26] Arie Tamir. 1991. Obnoxious Facility Location on Graphs. *SIAM Journal on Discrete Mathematics* 4, 4 (1991), 550–567.

[27] Vijay V. Vazirani. 2001. *Approximation algorithms.* Springer, Switzerland.

[28] William S. Zwicker. 2016. Introduction to the Theory of Voting. In *Handbook of Computational Social Choice.* Cambridge University Press, Cambridge, England, 23–56.