Metric Distortion of Line-up Elections: The Right Person for the Right Job

Christopher Jerrett, Yue Han, Elliot Anshelevich

Rensselaer Polytechnic Institute
110 Eighth Street
Troy, New York 12180
jerrec@rpi.edu, hany4@rpi.edu, eanshel@cs.rpi.edu

Abstract

We provide mechanisms and new metric distortion bounds for line-up elections. In such elections, a set of n voters, m candidates, and ℓ positions are all located in a metric space. The goal is to choose a set of candidates and assign them to different positions in order to minimize the total cost of the voters. The cost of each voter consists of the distances from itself to the chosen candidates (measuring how much the voter likes the chosen candidates or how similar it is to them), as well as the distances from the candidates to the positions they are assigned to (measuring the fitness of the candidates for their positions). Our mechanisms, however, do not know the exact distances and instead produce good outcomes while only using a smaller amount of information, resulting in small distortion.

We consider several different types of information: ordinal voter preferences, ordinal position preferences, and knowing the exact locations of candidates and positions, but not those of voters. In each of these cases, we provide constant distortion bounds, thus showing that only a small amount of information is enough to form outcomes close to optimum in line-up elections.

1 Introduction

Consider the well-known spacial model of voter preferences (Arrow 1990; Carroll et al. 2013; Enelow and Hinich 1984; Schofield 2007). In this model, both voters and candidates are located in an arbitrary metric space, with the distance from a voter to a candidate representing how much the voter prefers them: the closer the candidate to the voter, the better. These distances could correspond to ideological differences or something more concrete, such as when candidates correspond to placing facilities (e.g., new post offices) and voters want a facility close to them (Anshelevich et al. 2018). However, in many such settings, it is too difficult, expensive, or impossible to calculate the exact distances from each voter to each candidate, or the voters may be reluctant to provide such detailed information. On the other hand, obtaining *ordinal* knowledge about voter preferences (i.e., voter i prefers A to B, and B to C) is often much easier. This fact gave rise to a large line of work on *metric distortion* (see Related Work), which is a measure of how well a mechanism knowing only ordinal preferences can perform compared to an omniscient algorithm which knows the true distances between voters and candidates. When attempting to minimize distortion, we are concerned with finding a mechanism that does well in minimizing the social cost (the total distances from the voters to the chosen candidate(s), which come from the underlying metric) while only knowing ordinal information, as compared with the true optimum solution. For more motivation and examples see the full version.

Previously, most work on metric distortion has considered selecting a single candidate, or selecting a committee of k candidates. In this work we instead consider *line-up elections.* In such elections, there are multiple positions that need to be filled (e.g., different spots on a sports team, different duties in a club or a committee, etc.), and a shared pool of available candidates to fill them. The goal is to choose which candidate will be assigned to each position, with the constraint that a candidate can only hold a single position. Not all candidates are equally fit for every position, as some may be more qualified for some positions than others. Thus, unlike in the work mentioned above, where we only cared about the total distance from voters to the chosen candidates, we are now concerned both about choosing candidates who are qualified for the position and *also* about what candidate best represents the views of the voters.

More formally, we have n voters $\mathcal{V}$, m candidates $\mathcal{C}$, and ℓ positions $\mathcal{P}$ located in some metric space with distance function d: distances between voters and candidates represent how much the voter likes them, and distances between candidates and positions represent how qualified the candidate is for this position. The goal is to form a matching so that every position has a candidate assigned to it, and every candidate is assigned to at most one position. How much a voter $v \in \mathcal{V}$ likes when a candidate $c \in \mathcal{C}$ is assigned to position $p \in \mathcal{P}$ depends on both the similarity between v and c, as well as c's fitness for the position p, and so is measured by $d(v,c) + d(c,p)$: the smaller this quantity the happier voter v becomes. Thus the total cost for a voter is the sum of this quantity over all positions and candidates assigned to those positions, and the Social Cost for this model is defined as the sum of all voter costs, as in all previous work mentioned above.

The optimal solution to this problem is the matching minimizing the social cost, and is simple to compute given the

	$(m,1)$	(m,m)	(m,ℓ)	(2,1)
VP	3 (3)	3	7 (3)	3 (3)
PP	3 (3)	$3-\frac{1}{2^{m-2}}$	5 (3)	3 (3)
LOC	3 (3)	1	3 (3)	3 (3)
VP+PP	-	-	-	2 (2)
VP+LOC	$\frac{25}{9}$	-	-	$\frac{5}{3}(\frac{5}{3})$

Table 1: Summary of results showing upper/lower bounds on distortion for lineup elections with different types of available information. VP stands for mechanisms that operate knowing voter preferences, PP for position preferences, and LOC for candidate and position locations. Numbers in parentheses are lower bounds.

distances or the voter costs. However, we are interested in forming mechanisms with small distortion and thus assume that we *do not know* the distances between voters and candidates. Instead, we only use a smaller amount of information (see Section 1.1), and form simple and deterministic voting rules that achieve outcomes only a small factor away from the socially optimal outcome.

In a sense, line-up elections are a combination of multi-winner elections (as a set of candidates of size $|\mathcal{P}|$ is chosen), and bipartite matching (as the chosen candidates are matched with the positions). Because of this, we use techniques optimizing distortion for both multi-winner elections and matching in this paper. As we establish below, however, the fact that both are being optimized simultaneously allows us to form better distortion results than for either in isolation. Many of our proof techniques first show the loss from choosing the incorrect matching, and then combine it with the loss of choosing the incorrect set of candidates. By carefully combining these, however, and showing that the loss of one can often be bounded by the loss of the other, we are able to form new proof techniques and distortion bounds.

Importantly, a limitation of our model is that we require the sets $\mathcal{V},\mathcal{P}$, and $\mathcal{C}$ to all lie in a common metric space. This is a crucial requirement: if we allow for multiple arbitrary metric spaces then it is easy to show that in the worst case our model cannot give bounds better than the standard matching model, which has exponential distortion (Caragiannis et al. 2022). That said, there are numerous settings where this assumption holds, such as in distribution networks where all locations $\mathcal{P}$ creating the goods, locations $\mathcal{V}$ consuming goods, and intermediate distribution centers $\mathcal{C}$ all lie in a common space. For more discussion and further examples see the full version of this paper.

1.1 Our Contributions

See Table 1 for a summary of our results.

In Section 3 we consider the classic notion of distortion, where only the ordinal voter preferences are known, i.e., we know the order of which candidates each voter prefers for each position. It is not difficult to show that for (m,1) line-up elections, i.e., when there is only a single position, our model is equivalent to usual elections where a single candidate is chosen. Thus, existing results for distortion in social choice apply for this case. For the special case of (m,m) line-up elections, we can also show a distortion of at most 3: this case is simpler than the general one since all candidates must be selected, and we only need to form an appropriate matching without knowing the true distances between the candidates and positions. One of our main results is from combining techniques for social choice and matching to give a distortion bound of 7 for general (m,ℓ) line-up elections. The mechanism to achieve this bound is simple: it is an iterative election which chooses a candidate to assign to each position separately by running a single-winner election using the voter preferences. Proving the distortion bound, however, requires a careful charging argument between the costs of various components of the matching between candidates and positions, and the distances from voters to candidates. We first show a simpler distortion bound of 9 to build intuition, and then improve it to the bound of 7 by more carefully bounding the combination of loss from choosing the incorrect matching, and choosing the incorrect set of candidates to match to positions.

We then consider different types of information. In Section 4, we study mechanisms which know ordinal information about the fitness candidates for each position (which we refer to as "position preferences"). In social choice, it is often reasonable to know the relative fitness of each candidate for each position (e.g., how much experience they have or their qualifications), even if we don't know the exact fitness levels or the voter preferences. Perhaps surprisingly, this small amount of ordinal information (as compared to knowing the preferences of all the voters) is enough to provide *better* distortion than if we knew all voter preferences. In fact, for (m,m) line-up elections, we can form distortion better than the best possible distortion for standard (non-line-up) elections, and for general (m,ℓ) line-up elections we can improve our bound from 7 to 5. Once again, the mechanism achieving these bounds is simple: it is serial dictatorship by using the position preferences. We are able to achieve these better results by taking existing bounds on the quality of matching and of social choice, and showing tradeoffs between them which combine to form better results than they can separately.

Finally, in Section 5, we consider the case where instead we are given the exact locations of all candidates and positions (and thus we know the distances between them), but we know absolutely nothing about the voter preferences. This information is enough to form a distortion of 3 for (m,ℓ) line-up elections, even without knowing anything about voters or their preferences. Interestingly, if we *also* know the ordinal preferences of the voters, then for a single position we can improve distortion to $\frac{5}{3}$ for (2,1) line-up elections and to $\frac{25}{9}$ for (m,1) line-up elections. Contrast this with distortion for standard single-winner elections, where knowing the exact locations of all the candidates in addition to ordinal voter preferences does *not* improve the possible distortion beyond 3 as compared to only knowing voter preferences (Anshelevich and Zhu 2021).

1.2 Related Work

The line-up election model was first introduced by Boehmer et al. (2020), where they formalized the concept of a multi-winner election with a shared candidate pool. They analyzed several types of voting rules for such elections from an axiomatic approach and showed several desirable properties of voting rules. Our work builds on theirs by considering the distortion of such elections, and forming mechanisms which perform well despite not knowing the exact scores that voters assign to each candidate and position.

Mechanisms for minimizing distortion, i.e., forming mechanisms with access to only ordinal information which perform almost as well as mechanisms with full information, have received a lot of attention. See Anshelevich, Filos-Ratsikas, Shah, and Voudouris (2021) for a survey on this topic. Much of this work specifically focuses on *metric distortion*, i.e., on forming such mechanisms with the assumption that voters and candidates are located in some unknown metric space (see e.g. Anagnostides, Fotakis, and Patsilinakos (2022); Anshelevich and Postl (2017); Anshelevich et al. (2018); Anshelevich and Zhu (2021); Caragiannis, Shah, and Voudouris (2022); Charikar and Ramakrishnan (2022); Cheng, Dughmi, and Kempe (2017); Feldman, Fiat, and Golomb (2016); Ghodsi, Latifian, and Seddighin (2019); Goel, Krishnaswamy, and Munagala (2017); Pierczynski and Skowron (2019); Skowron and Elkind (2017); Anshelevich et al. (2024); Charikar et al. (2024)). Most of these works concerns single-winner elections, with Gkatzelis, Halpern, and Shah (2020); Kizilkaya and Kempe (2022) showing how to achieve a metric distortion of 3 via a deterministic mechanism in single-winner elections. Few papers have looked at distortion of multi-winner elections, with the following main exceptions. Chen, Li, and Wang (2020) look at selecting a committee equal in size to the total number of candidates minus one. They show a tight bound of 3 on distortion for the case when the cost of a voter equals the distance to the closest member of the committee. Caragiannis, Shah, and Voudouris (2022) also consider the problem of picking a committee of k candidates, but such that the cost for each voter is given by the distance to the q'th closest candidate. They then try to minimize the sum of the costs of all agents, and show that the distortion may become arbitrarily large when $q \leq \frac{k}{3}$, asymptotically linear in the number of agents when $\frac{k}{3} < q \leq \frac{k}{2}$ and constant when $q > \frac{k}{2}$. Our model differs from that of Caragiannis, Shah, and Voudouris (2022) in a few ways. First, we must both choose k candidates to select, as well as which of these is assigned to which position. Second, the cost of a voter is the sum of its costs for *all* positions, so it is the sum of distances to the selected candidates together with their distances to the assigned positions. Thus, our model is both more difficult (because we must form an assignment of candidates to positions as well), and simpler (since we are looking at the sum of distances instead of the q'th closest distance). Note that if there were no positions and we simply wished to choose k candidates $K \subseteq \mathcal{C}$ to minimize $\sum_{v \in \mathcal{V}} \sum_{c \in K} d(v, c)$, then it is easy to achieve a distortion of 3 by using plurality veto (Kizilkaya and Kempe 2022) to select k candidates one at a time, as in Goel, Hulett, and Krishnaswamy (2018).

As discussed above, our model is essentially a combination of classic social choice committee selection, together with forming a matching between selected candidates and positions. The distortion of matching in metric spaces has been previously studied in Anshelevich and Sekar (2016a,b); Bhalgat, Chakrabarty, and Khanna (2011); Caragiannis et al. (2022); Christodoulou et al. (2016). Most of this work considers computing the maximum weight matching, however, with the main exception of Caragiannis et al. (2022). They study the problem of minimizing the total social cost of a resource assignment problem, where users are matched with capacitated resources, using access only to ordinal preferences. As one of their results, they show that the distortion of serial dictatorship is at most $2^k - 1$ for metric bipartite matching. Recently, Anari, Charikar, and Ramakrishnan (2023) gave a better algorithm achieving distortion of $O(k^2)$ for this matching problem, as well as a lower bound of $\Omega(\log k)$. In contrast to the above results for matching, by matching candidates to positions at the same time as selecting candidates preferred by voters, and thus considering the total cost of both, in our model we are able to achieve much better distortion.

2 Model and Preliminaries

We now formally introduce the (m,ℓ) line-up election and define some terms that will be useful later on in this paper. Let $\mathcal{P}$ be a set of ℓ positions, $\mathcal{C}$ be a set of m candidates, and $\mathcal{V}$ be the set of n voters in a metric space with distance function $d : \mathcal{P} \cup \mathcal{C} \cup \mathcal{V} \to \mathbb{R}^+$. Throughout this paper we will assume $\ell \leq m$. The goal in a line-up election is to form an assignment from the set of positions to the set of candidates, so that each candidate is assigned to at most one position, and each position has a candidate assigned to it. In other words, the goal is to form an injective matching $M : \mathcal{P} \to \mathcal{C}$. If we assign position p to candidate c, then the cost for voter v is given by $d(v, c) + d(c, p)$; this means that the voter cares both about her proximity to the chosen candidate as well as the candidate's fitness for the position. Thus the total cost of a voter v for a matching M is defined as the following.

$$\text{cost}(v, d) = \sum_{p \in \mathcal{P}} [d(v, M(p)) + d(M(p), p)]$$

Thus the total cost of a voter depends both on which candidates are selected, and which positions these candidates are assigned to. The social cost of a matching M is simply the sum of all voter costs, as in most previous work:

$$\begin{aligned}\text{cost}(M, d) &= \sum_{v \in \mathcal{V}} \text{cost}(v) \\ &= \sum_{v \in \mathcal{V}} \sum_{p \in \mathcal{P}} \left[d\left(v, M(p)\right) + d\left(M(p), p\right) \right].\end{aligned}$$

We will sometimes write simply $\text{cost}(M)$ when d is clear from context. We will denote the set of candidates selected by matching M as $\text{cand}(M) = \{c \in \mathcal{C} \colon \exists p \in \mathcal{P} \text{ with } M(p) = c\}$. The goal of a line-up election is to choose an outcome M which minimizes the social cost.

Note that the total cost of M depends on two distinct parts, which we will sometimes consider separately in our proofs. Because of this, we define $\text{cost}_V(M)$ and $\text{cost}_P(M)$ so that $\text{cost}(M) = \text{cost}_V(M) + \text{cost}_P(M)$ and

$$\text{cost}_V(M) = \sum_{v \in \mathcal{V}} \sum_{p \in \mathcal{P}} d(v, M(p))$$
$$\text{cost}_P(M) = \sum_{v \in \mathcal{V}} \sum_{p \in \mathcal{P}} d(M(p), p).$$

In this paper we are interested in forming mechanisms which only use a limited amount of information, instead of knowing the exact distances d. Knowing the ordinal preferences of the voters for the candidates, instead of their true distances, is the most common type of information used in previous work on distortion. For line-up elections, this means knowing the voter preferences over candidates for each position. In other words, for each position p and voter v, and candidates c_1 and c_2, we know which of these candidates v prefers for position p, i.e., which of $d(v, c_1) + d(c_1, p)$ or $d(v, c_2) + d(c_2, p)$ is smaller. We only know the ordering of these when choosing an outcome of the election, however, and not their actual values.

In later sections, we also consider mechanisms with ordinal knowledge of how qualified different candidates are for each position. In other words, we know, for each position p and candidates c_1, c_2, whether $d(c_1, p)$ or $d(c_2, p)$ is smaller, but not the actual distance values. It is often less difficult to sort candidates in order of their qualifications for a position, instead of figuring out the exact numerical level of fitness. We refer to this knowledge as "position preferences" for the candidates.

Let I be the set of information we have about a line-up election instance. The set I could consist of some combination of ordinal voter preferences, ordinal position preferences, or exact candidate and position locations. We do not necessarily restrict what form the information takes, and consider only the previous three types of information in this work. We now introduce the distortion of a matching M given information set I.

Definition 2.1. Let $\mathcal{D}$ be the set of all metric spaces that are consistent with information set I and $M^*_{d'}$ be the optimal matching for metric d'. The *distortion* of a matching M is

$$\sup_{d' \in \mathcal{D}} \frac{\text{cost}(M, d')}{\text{cost}(M^*_{d'}, d')}.$$

In other words, distortion of M is the worst case ratio between the cost of M and the optimal solution over any distances which are consistent with the information we possess. Similarly, the distortion of a mechanism is the worst-case distortion of any matching M produced by this mechanism. In this paper, we only focus on the distortion of deterministic mechanisms, and leave the analysis of randomized mechanisms for future work.

3 Knowing Ordinal Voter Preferences

We first consider the case where for each position, we only know the ordinal preferences of each voter towards each candidate for this position. In other words, for each position p and voter v, we know which candidate would be v's first choice for p based on v's cost, which would be the second choice, etc. First we will show some easy reductions to well known problems, as well as some lower bounds, and then show more difficult and general theorems afterward.

We will call the following a *standard election*, to differentiate it from line-up elections which are the focus of this paper. Let $\mathcal{C}$ be a set of candidates and $\mathcal{V}$ be a set of voters in some metric space d. For any $c \in \mathcal{C}$, let the social cost for each voter v be $d(v, c)$. Then the total social cost of the election is the sum of the individual costs for each voter, $\sum_{v \in \mathcal{V}} d(v, c)$. The goal is to find a candidate $w \in \mathcal{C}$ such that w minimizes the total social cost. Having defined the standard election we can now present our first simple result for the case of a single position.

Lemma 3.1. *For each (m,1) line-up election there exists a standard election with the same candidate set $\mathcal{C}$, such that choosing $c \in \mathcal{C}$ in the line-up election is equal to the cost of choosing c in the standard election.*

All our detailed proofs can be found in the full version. Because of the above lemma, all the known distortion results for standard elections also hold for (m,1) line-up elections. In particular, since there exists a mechanism that gives a distortion of three for standard elections (Gkatzelis, Halpern, and Shah 2020; Kizilkaya and Kempe 2022), we immediately have the following theorem.

Theorem 3.2. *There exists a mechanism for the the (m,1) line-up election that achieves a distortion of 3. Moreover, no other deterministic mechanism can achieve a worst-case distortion bound less than 3, even on a line with 2 candidates.*

Therefore, by Theorem 3.2 we know that we can solve (m,1) line-up elections and achieve the optimal distortion of 3 using standard techniques from metric distortion. Before we look at the general case where we have arbitrary ℓ positions and m candidates, we first consider the special case where $\ell = m$.

Lemma 3.3. *If $\ell = m$ then for any two matchings M_1 and M_2, we have* $\text{cost}(M_1) \leq 3\,\text{cost}(M_2)$.

By including the voters in our model we can bound the distances between candidates and positions using the distance from the voters to the candidates, allowing us to achieve a constant factor approximation. In contrast, no such mechanism is currently known for the the analogous bipartite perfect matching problem.

Corollary 3.4. *For any (m,m) line-up election all matchings are at most a factor of 3 away from optimal.*

After having looked at special cases of $(m, 1)$ and (m,m) line-up elections, we will now look at a mechanism for the more general (m,ℓ) line-up election problem. We introduce *iterative election*, a natural mechanism for the (m,ℓ) line-up election problem. For a position p reduce the problem to a (m,1) line-up election problem using the transformation given in Lemma 3.1, and then use a mechanism for the standard election problem. If c is selected using the standard election mechanism we assign $M_1(p) = c$. We then remove the selected candidate c, and repeat the procedure for each remaining position, removing the selected candidate each time. Importantly, this mechanism is quite general, as we do not assume anything about the mechanism we use at each step or the specific information set I. One could achieve optimal metric distortion for each position by using plurality veto (Kizilkaya and Kempe 2022) for instance, or use a different mechanism that has attractive properties like strategyproofness. Lastly, iterative election is also an online algorithm, meaning that positions can arrive overtime, and as long as we have the appropriate information we can immediately assign each position at the time of arrival. We will now analyze the case where we utilize a standard election mechanism with distortion δ_{itr} and show an upper bound for the iterative election procedure on the (m,ℓ) line-up election problem.

Theorem 3.5. *Let M_1 be the matching returned by the iterative election procedure and M^* be the optimal matching minimizing the social cost. Then* $\mathrm{cost}(M_1) \leq 3\delta_{itr}\,\mathrm{cost}(M^*)$.

Proof Sketch. Let $S_1 = \mathrm{cand}(M_1) \setminus \mathrm{cand}(M^*)$, $S_2 = \mathrm{cand}(M_1) \cap \mathrm{cand}(M^*)$, and $S_3 = \mathrm{cand}(M^*) \setminus \mathrm{cand}(M_1)$. Now consider some $c \in S_1$ and $c' \in S_3$. c' was not selected when c was selected to be assigned to position p by a standard election mechanism with distortion δ_{itr}. This means that $\sum_{v\in\mathcal{V}} [d(v,c) + d(c,p)] \leq \delta_{itr} \sum_{v\in\mathcal{V}} [d(v,c') + d(c',p)]$, i.e., the cost of c for position p is within a factor δ_{itr} of the cost of c'. Now consider the matching M_2 obtained by swapping the assignment of $c \in S_1$ with some $c' \in S_3$. Then $\mathrm{cost}(M_1) \leq \delta_{itr}\,\mathrm{cost}(M_2)$. Now, by Lemma 3.3, we have $\mathrm{cost}(M_2) \leq 3\,\mathrm{cost}(M^*)$. Thus $\mathrm{cost}(M_1) \leq 3\delta_{itr}\,\mathrm{cost}(M^*)$. □

We know that plurality veto gives the best possible upper bound of 3 on the distortion of standard elections. Unfortunately, using the above with $\delta_{itr} = 3$ only gives us a factor of 9. To improve the distortion factor to 7, we require a far more careful analysis, which results in the following theorem.

Theorem 3.6. *Let M_1 be the matching returned by the iterative election procedure using a mechanism with distortion δ_{itr}, and M^* be the optimum matching. Then* $\mathrm{cost}(M_1) \leq (2\delta_{itr} + 1)\,\mathrm{cost}(M^*)$.

We will show the proof sketch for Theorem 3.6 below, but due to this theorem and the fact that Gkatzelis, Halpern, and Shah (2020); Kizilkaya and Kempe (2022) provide algorithms with $\delta_{itr} = 3$ for standard elections, we immediately get the following corollary.

Corollary 3.7. *Iterative election with plurality veto gives a distortion of at most 7.*

The proof of Theorem 3.6 actually follows from a more general statement, that will also be crucial to proving Theorem 4.3 later. We state this somewhat technical theorem below, and give a proof sketch to build intuition (for the full proof, see the full version).

Theorem 3.8. *Suppose some mechanism gives a matching M_1, and assume M^* is the optimal matching. Let $S_1 = \mathrm{cand}(M_1) \setminus \mathrm{cand}(M^*)$, $S_2 = \mathrm{cand}(M_1) \cap \mathrm{cand}(M^*)$, and $S_3 = \mathrm{cand}(M^*) \setminus \mathrm{cand}(M_1)$. If*

$$\sum_{v\in\mathcal{V}} (d(v,c) + d(c, {M_1}^{-1}(c))) \leq \alpha \sum_{v\in\mathcal{V}} d(v,c') + \beta \sum_{v\in\mathcal{V}} d(c', {M_1}^{-1}(c))$$

for every $c \in S_1$ and $c' \in S_3$, and with $\alpha + \beta \geq 2$, then $\mathrm{cost}(M_1) \leq (\alpha + \beta + 1)\,\mathrm{cost}(M^*)$.

Proof Sketch. At a high level, the proof structure parallels that of Theorem 3.5: we first transform M_1 into a matching M_2 which has the same chosen candidates as M^*, and bound the difference in cost due to this transformation, followed by the difference in cost between M_2 and M^*. One difference is that we will have to be significantly more careful about how we transform M_1 into an intermediate matching M_2.

Let us define the matching M_2 as follows:

1. For all positions p with $M_1(p) \in S_2$ assign $M_2(p) = M_1(p)$.
2. For all positions p with $M_1(p) \in S_1$ and $M^*(p) \in S_3$ assign $M_2(p) = M^*(p)$.
3. For any positions p that is not assigned by step 1 or step 2 assign $M_2(p) = c$ for some arbitrary unassigned candidate $c \in S_3$.

With this matching M_2, we can show the following.

Lemma 3.9. *Consider the matchings M_1 and M_2 given above and let $f(c) = M^*({M_2}^{-1}(c))$. We then have the following properties:*

1. $\mathrm{cand}(M_2) = S_2 \cup S_3 = \mathrm{cand}(M^*)$.
2. *The function $f\colon S_2 \cup S_3 \to S_2 \cup S_3$ is bijective.*
3. *If $c \in S_3$ then $f(c) \in S_2$ or $f(c) = c$.*

In other words, there exists a matching M_2 such that everything in S_1 has its assignment changed to the corresponding position in M^* if possible, ensuring that two different candidates in S_3 with never be matched to the same position in matchings M_2 and M^*.

Just as in the proof of Theorem 3.5, we have now transformed matching M_1 into matching M_2 with the same candidate set as M^*. We then have that $\mathrm{cost}(M_1)$ is at most

$$\mathrm{cost}(M_2|S_2) + \alpha\,\mathrm{cost}_V(M_2|S_3) + \beta\,\mathrm{cost}_P(M_2|S_3)$$

where $\mathrm{cost}(M|S) = \sum_{v\in\mathcal{V}} \sum_{c\in S} \left(d(v,c) + d(c, M^{-1}(c))\right)$ is the cost of M restricted to the candidates in S (see Figure 1a).

We then further partition S_2 and S_3 as shown in Figure 1a. We partition S_2 into sets S_2^2 (the set of candidates in S_2 with $f^{-1}(c) \in S_2$) and S_2^3 (the set of candidates in S_2 with $f^{-1}(c) \in S_3$). We also partition S_3 into S_3' (the set of candidates in S_3 with $f^{-1}(c) \neq M^*(c)$) and S_3^* (the set of candidates in S_3 with $f^{-1}(c) = M^*(c)$). Through a series of transformations utilizing the properties of the function f defined in Lemma 3.9, we can obtain that the cost of M_1 is

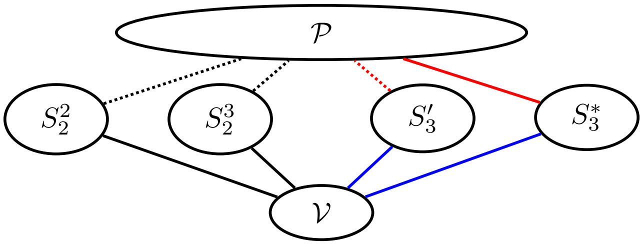

(a) Our argument shows that $\text{cost}(M_1)$ is at most the cost of M_2, with coefficient multipliers shown in this figure. The black lines indicate a factor of 1, red a factor of β, and blue a factor of α. Dotted lines indicate candidates in those sets may not be matched to the same position as in M^*.

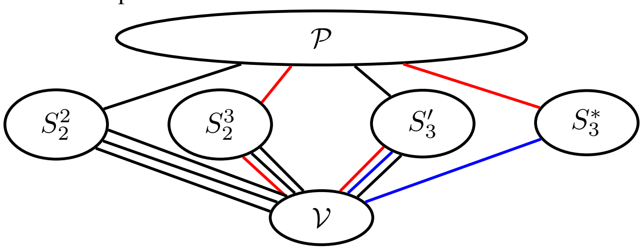

(b) We have now shown that the cost of the original matching is at most the cost of various M^* components. The component which appears the most is the distance from voters to S_3': this appears at most $\alpha + \beta + 1$ times.

Figure 1: Costs in Proof Sketch of Theorem 3.8.

at most the following (see Figure 1b for a graphical illustration):

$$\begin{aligned}
&3\,\text{cost}_V(M^*|S_2^2) + \text{cost}_P(M^*|S_2^2)\\
&+ (\beta+2)\,\text{cost}_V(M^*|S_2^3) + \beta\,\text{cost}_P(M^*|S_2^3)\\
&+ (\alpha+\beta+1)\,\text{cost}_V(M^*|S_3') + \text{cost}_P(M^*|S_3')\\
&+ \alpha\,\text{cost}_V(M^*|S_3^*) + \beta\,\text{cost}_P(M^*|S_3^*).
\end{aligned}$$

Looking at the largest of these factors, we can bound the distortion by at most $\alpha + \beta + 1$, as desired. □

Now, we can prove Theorem 3.6 using Theorem 3.8.

Proof Sketch of Theorem 3.6. Using the definitions of S_1, S_2 and S_3 in Theorem 3.8, it is easy to see that for $c \in S_1$ then for any $c' \in S_3$, the conditions for Theorem 3.8 hold with $\alpha = \beta = \delta_{itr}$ since c was selected instead of c'. Therefore, the distortion is bounded above by $2\delta_{itr} + 1$ for any iterative election by Theorem 3.8. □

4 Knowing Ordinal Position Preferences

Now that we have discussed the scenario where we only know the preferences of voters for each candidate, we look into the case where we know the ordinal preferences of the positions for the candidates. In other words, for each position p and candidates c_1 and c_2, we know which of $d(c_1, p)$ or $d(c_2, p)$ is smaller. So we only know the relative fitness of candidates for each position, but nothing about the voter preferences.

Theorem 4.1. *There exists a mechanism that achieves a distortion of* 3 *for* $(m,1)$ *line-up elections while knowing only the position preferences. Moreover, no deterministic mechanism can have smaller worst-case distortion while knowing only the position preferences, even on a line and with only 2 candidates.*

Note that the above mechanism is simply just picking the closest candidate to the position. For the rest of this section, we will be using the more general *serial dictatorship* mechanism (see e.g., Biró (2017); Caragiannis et al. (2022)): for each position p in some arbitrary order, we assign $M(p)$ to be the top choice of p among remaining unassigned candidates.

After considering the case where we have m candidates and one position, we now extend the setting to m candidates and m positions. In this case, each candidate will be matched to one position and vice versa. We will first show that serial dictatorship gives a trade off between the cost of the position side of the matching and the cost of the voter side of the matching. Then, we will utilize this trade off to obtain an upper bound of the distortion smaller than 3.

Theorem 4.2. *Serial dictatorship gives a distortion of at most* $3 - \frac{1}{2^{m-2}}$ *for* (m,m) *line-up elections.*

Proof Sketch. We begin with a more careful analysis of Lemma 3.3 to improve its result, and derive that $\text{cost}(M_1) \leq 3\,\text{cost}_V(M_2) + \text{cost}_P(M_2)$ for any two matchings M_1 and M_2. Now, we combine this result with the result from Caragiannis et al. (2022) showing that serial dictatorship on bipartite graphs has distortion at most $2^m - 1$. Then, we can see that the distortion is bounded by the minimum of

$$\frac{3\,\text{cost}_V(M^*) + \text{cost}_P(M^*)}{\text{cost}_V(M^*) + \text{cost}_P(M^*)}$$

and

$$\frac{\text{cost}_V(M^*) + (2^k - 1)\,\text{cost}_P(M^*)}{\text{cost}_V(M^*) + \text{cost}_P(M^*)}.$$

We can then bound the minimum of these two terms by $3 - \frac{1}{2^{m-2}}$. □

Theorem 4.2 shows that simply by using serial dictatorship we can get a distortion better than 3. For $m = 2$ we get a distortion of 2, and for $m = 3$ we get $\frac{5}{2}$, so for small m the distortion remains relatively low despite approaching 3 as m becomes larger.

We now consider the general case where we have m candidates and ℓ positions, with $m \geq \ell$. We will now show that simply running serial dictatorship on the set of positions results in an upper bound on the distortion of 5. Having the extra information on the ordering of candidates by fitness for each position not only allows us to improve the bounds from Theorem 3.6 when we only know voter preferences, but also does so without needing voter preferences at all. Note that serial dictatorship then is also strategy proof, since we do not consider the preferences of the candidates or the voters, and even the positions are not incentivized to lie (if such a thing were possible) to achieve a closer candidate.

Theorem 4.3. *Serial dictatorship achieves a distortion of at most* 5 *for* (m,ℓ) *line-up elections.*

Proof Sketch. Let M_1 be the matching returned by serial dictatorship, and use the same definitions of S_1, S_2 and S_3 as in Theorem 3.8 and. Suppose that $c \in S_1$, $c' \in S_3$ and $p = {M_1}^{-1}(c)$. Then by the mechanism $d(c,p) \leq d(c',p)$, thus $d(v,c) + d(c,p) \leq d(v,c') + d(c',p) + 2d(c,p) \leq d(v,c') + 3d(c',p)$. Thus $\sum_{v \in \mathcal{V}} (d(v,c) + d(c, {M_1}^{-1}(c))) \leq \alpha \sum_{v \in \mathcal{V}} d(v,c') + \beta \sum_{v \in \mathcal{V}} d(c', {M_1}^{-1}(c))$ for every $c \in S_1$ and $c' \in S_3$ for $\alpha = 1$ and $\beta = 3$. Thus by Theorem 3.8 the distortion is bounded by $\alpha + \beta + 1$. □

We will now briefly consider the case for when we also know the voter ordinal preferences for the candidates in addition to the ordinal position preferences. We show that we can improve the upper and lower bounds from 3 to 2 for (2,1) line-up elections.

Theorem 4.4. *Consider a (*2*,*1*) line-up election knowing position preference information* and *voter preference information, with $\mathcal{P} = \{1\}$ and $\mathcal{C} = \{A, B\}$. Without loss of generality assume we have $d(A,1) \leq d(B,1)$ and there are n_a voters who prefer A to B and n_b who prefer B to A. Consider the mechanism which selects candidate A if $n_b \leq 2n_a$, otherwise select candidate B. Then, this mechanism achieves a distortion of at most* 2*. Moreover, no deterministic mechanism can do better, even on a line.*

Proof Sketch. First we show that for any (2,1) line-up election instance there exists another (2,1) line-up election instance that lies on a line with the same information set such that it has a distortion no less than the original instance. If A wins, we use a charging argument from those who prefer B to those who prefer A, with charging at most 2 voters to each voter who prefers A. Likewise if B wins, we use a charging argument from those who prefer A to those who prefer B, so that each is charged to two unique voters who prefer B. In total this results in a distortion bound of 2. □

5 Knowing Candidate and Position Locations

In this section, we look at what is possible when candidate and position locations are known exactly (and thus all the distances between them are known), but absolutely nothing is known about voter locations or preferences. This is often possible since information about candidates and positions can be public knowledge, but voter locations and preferences are either private, or surveying a large number of people about their preferences is impractical.

Theorem 5.1. *Let M be the matching that minimizes $\sum_{v \in \mathcal{V}} \sum_{p \in \mathcal{P}} d(v, M(p))$. Then M gives distortion at most* 3 *for the (k,ℓ) line-up election while knowing only candidate-position distance pairs. Moreover, no deterministic mechanism can achieve a distortion bound less than* 3 *while knowing only position and candidate locations.*

Proof Sketch. The result follows almost directly from the properties of M as well as the triangle inequality applied to $d(v, M(p))$. □

Much like the mechanisms in Theorem 4.1 and 4.2 this mechanism is strategyproof and does not rely on the preferences of the voters. This means that we can achieve low distortion without surveying voters, and instead focusing on increasing our knowledge of the pairwise distances between candidates and positions. Note that for (*k*,*k*) line-up elections, the above mechanism trivially computes the optimum solution, since all candidates must always be chosen in such elections, and thus $\text{cost}_V(M)$ of any matching M is always the same, and only $\text{cost}_P(M)$ needs to be minimized.

We will now consider the case when we know the exact locations of candidates and positions *in addition* to voter ordinal preferences. We will consider the case where we have a single position, and first look at the case where there are 2 candidates.

Theorem 5.2. *Consider the (*2*,*1*) line-up election. Without loss of generality assume $\mathcal{P} = \{1\}$ and $\mathcal{C} = \{A, B\}$ and let n_a be the number of people that prefer A and n_b be the number that prefer B. The mechanism that assigns $M(1) = A$ if $n_b d(A,1) \leq n_a d(B,1)$, and otherwise assigns $M(1) = B$, gives distortion at most $\frac{5}{3}$. Moreover, there does not exist a deterministic mechanism that achieves a lower distortion bound.*

The method in Theorem 5.2 can be generalized to more than 2 candidates as follows.

Theorem 5.3. *There exists a mechanism that achieves a distortion of $\frac{25}{9} \approx 2.78$ for (m,1) line-up elections.*

Proof Sketch. We start by constructing a graph $G = (\mathcal{C}, E)$ similar to a weighted majority graph. Instead we orient arcs such that instead of arcs going from winners according to a plurality election, they are oriented based on the winner given by Theorem 5.2. So if $(x,y) \in E$ then x was selected by the mechanism used in Theorem 5.2. We then select the client with the largest out degree. Let O be the optimal candidate and A be the candidate with largest out-degree. It is a well-known observation that there exists a path of at most length 2 from A to O in G. Then, the worst case scenario is when there is a path of length two from A to O, which gives us a $\frac{5}{3} \cdot \frac{5}{3} = \frac{25}{9}$ distortion by Theorem 5.2. □

Recall from Anshelevich and Zhu (2021) that in the standard election problem if we are given exact information about candidate locations, we cannot improve the distortion bounds beyond 3, even for 2 candidates, which is the same distortion achievable when given only voter preferences. In other words, for standard elections having candidate locations does not give any advantage in terms of worst-case distortion. As we have seen, however, for line-up elections this is no longer true: having candidate and position locations allows us to improve the distortion below 3, at least for the case when $\ell = 1$.

Acknowledgments

This work was partially supported by NSF award CCF-2006286.

References

Anagnostides, I.; Fotakis, D.; and Patsilinakos, P. 2022. Metric-distortion bounds under limited information. *Journal of Artificial Intelligence Research (JAIR)*, 74: 1449–1483.

Anari, N.; Charikar, M.; and Ramakrishnan, P. 2023. Distortion in metric matching with ordinal preferences. In *Proceedings of ACM Conference on Economics and Computation (EC)*.

Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence (AIJ)*, 264: 27–51.

Anshelevich, E.; Filos-Ratsikas, A.; Jerrett, C.; and Voudouris, A. A. 2024. Improved Metric Distortion via Threshold Approvals. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 9460–9468.

Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research (JAIR)*, 58: 797–827.

Anshelevich, E.; and Sekar, S. 2016a. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Anshelevich, E.; and Sekar, S. 2016b. Truthful mechanisms for matching and clustering in an ordinal world. In *The 12th Conference on Web and Internet Economics (WINE)*.

Anshelevich, E.; and Zhu, W. 2021. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. *ACM Transactions on Economics and Computation, (TEAC)*, 9(2): 1–24.

Arrow, K. 1990. *Advances in the spatial theory of voting*. Cambridge University Press.

Bhalgat, A.; Chakrabarty, D.; and Khanna, S. 2011. Social welfare in one-sided matching markets without money. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*.

Biró, P. 2017. Applications of Matching Models under Preferences. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 18, 345–373. AI Access.

Boehmer, N.; Bredereck, R.; Faliszewski, P.; Kaczmarczyk, A.; and Niedermeier, R. 2020. Line-Up Elections: Parallel Voting with Shared Candidate Pool. In *International Symposium on Algorithmic Game Theory (SAGT)*.

Caragiannis, I.; Filos-Ratsikas, A.; Frederiksen, S. K. S.; Hansen, K. A.; and Tan, Z. 2022. Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship. *Mathematical Programming*, 1–30.

Caragiannis, I.; Shah, N.; and Voudouris, A. A. 2022. The metric distortion of multiwinner voting. *Artificial Intelligence (AIJ)*, 313.

Carroll, R.; Lewis, J. B.; Lo, J.; Poole, K. T.; and Rosenthal, H. 2013. The structure of utility in spatial models of voting. *American Journal of Political Science*, 57(4): 1008–1028.

Charikar, M.; and Ramakrishnan, P. 2022. Metric distortion bounds for randomized social choice. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Charikar, M.; Ramakrishnan, P.; Wang, K.; and Wu, H. 2024. Breaking the metric voting distortion barrier. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1621–1640. SIAM.

Chen, X.; Li, M.; and Wang, C. 2020. Favorite-candidate voting for eliminating the least popular candidate in a metric space. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 34, 1894–1901.

Cheng, Y.; Dughmi, S.; and Kempe, D. 2017. Of the people: voting is more effective with representative candidates. In *ACM Conference on Economics and Computation (EC)*.

Christodoulou, G.; Filos-Ratsikas, A.; Frederiksen, S. K. S.; Goldberg, P. W.; Zhang, J.; and Zhang, J. 2016. Social welfare in one-sided matching mechanisms. In *Autonomous Agents and Multiagent Systems Conference (AAMAS)*.

Enelow, J.; and Hinich, M. 1984. *The Spatial Theory of Voting: An Introduction*. Cambridge University Press.

Feldman, M.; Fiat, A.; and Golomb, I. 2016. On voting and facility location. In *ACM Conference on Economics and Computation (EC)*.

Ghodsi, M.; Latifian, M.; and Seddighin, M. 2019. On the distortion value of the elections with abstention. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Gkatzelis, V.; Halpern, D.; and Shah, N. 2020. Resolving the optimal metric distortion conjecture. In *IEEE Annual Symposium on Foundations of Computer Science (FOCS)*.

Goel, A.; Hulett, R.; and Krishnaswamy, A. K. 2018. Relating metric distortion and fairness of social choice rules. In *Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation (NetEcon)*.

Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *ACM Conference on Economics and Computation (EC)*.

Kizilkaya, F. E.; and Kempe, D. 2022. Plurality Veto: A Simple Voting Rule Achieving Optimal Metric Distortion. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Pierczynski, G.; and Skowron, P. 2019. Approval-based elections and distortion of voting rules. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Schofield, N. 2007. *The spatial model of politics*. Routledge.

Skowron, P.; and Elkind, E. 2017. Social choice under metric preferences: Scoring rules and STV. In *AAAI Conference on Artificial Intelligence (AAAI)*.