

Metric Distortion with Elicited Pairwise Comparisons

Soroush Ebadian¹, Daniel Halpern², Evi Micha²

¹University of Toronto

²Harvard University

soroush@cs.toronto.edu, dhalpern@g.harvard.edu, emicha@cs.toronto.edu

Abstract

In many social choice applications, information about individuals’ preferences can only be elicited using a limited number of pairwise comparisons. In these cases, the task is twofold: we must first choose the queries, and then second, we must aggregate the responses to choose an outcome. We study the problem of designing algorithms for this setting. To compare the effectiveness of different outcomes, we use the metric distortion framework. In addition, we consider various constraints on the query algorithms, namely, placing restrictions on how the choice of the next query may depend on previous answers. Our main contributions are nearly optimal algorithms for all settings considered.

1 Introduction

When aggregating a population’s preferences to make a decision, we must contend with the fact that our access to preferences may be limited. Take as a motivating example *Reinforcement Learning with Human Feedback (RLHF)*, a process used to fine-tune large language models. At a high level, the goal is to choose model parameters that yield desirable outcomes according to the participants. We typically learn what is “desirable” by asking comparison queries, i.e., on prompt x , do you prefer output y or output z ? However, notice that these responses only give us partial information about the individuals’ true underlying preferences. It would be completely infeasible to learn, say, a ranking for each person over all possible prompt-output combinations. As another example, consider the opinion aggregation platform *Polis* [Small *et al.*, 2021]. Here, we would like to summarize a community’s opinions on an issue (e.g., how should we regulate Uber?). Participants submit free-form comments (“I think Passenger Liability Insurance should be mandatory”) and are shown comments of other users to vote on, either agree or disagree. However, similar to the fine-tuning example, there are simply too many submissions from other users to feasibly ask every individual about every comment. In both these examples, human input is a scarce resource that must be carefully allocated.

A unifying theme of the applications discussed is that there are two stages to the aggregation pipeline: first, we must decide which queries to ask each participant, and second, given the responses, we must choose an outcome. While the field of computational social choice [Brandt *et al.*, 2016] gives many tools to understand and design systems for the latter stage, the main focus of this work is the former. In particular, we consider the problem of *choosing a limited number of pairwise-comparison queries to effectively aggregate heterogeneous preferences*.

When selecting the queries, it is crucial to consider various constraints that a designer might need to satisfy. For example, on platforms like *Polis*, it is not guaranteed that we can interact with a user, then acquire more information from other users, and later return to the original user for a second round of queries. For this, we may want algorithms that are *single-pass*: the queries asked to earlier users do not depend on queries to ones arriving later. Another desirable property a practitioner might want to satisfy is a notion of fairness, say, the queries asked for each agent to be independent of the responses of other agents; this can help ensure that interaction with the system remains unaffected by the arrival time. The same need may also arise from temporal concerns: if agents all answer simultaneously, we simply do not have data from the other agents. Conversely, we may want to have queries be independent of an agent’s *own* answers, ensuring an unbiased interaction with the system. Therefore, our main research questions are:

How can we design a limited number of pairwise comparison queries to achieve efficient outcomes? Furthermore, how well can we do so under various restrictions such as limited rounds of interactions and which answers a query can depend on?

To understand the effectiveness of various query choices, we need a way to measure the quality of the outcome. For this, we turn to the celebrated *distortion framework* [Procaccia and Rosenschein, 2006]. Here, it is assumed that the comparison feedback submitted by the agents is derived from more expressive underlying preferences, typically scalar values. The more expressive values give rise to concrete objective values for each of the outcomes. The goal is to design a procedure that minimizes *distortion*, the worst-case ratio between the best possible objective value and the objective value of the chosen outcome. Without additional as-

sumptions, achieving bounded distortion becomes impossible, even with complete knowledge of each agent’s ranking. In this work, we make the frequently-made natural assumption that agents have costs induced by distances in an underlying metric space [Anshelevich *et al.*, 2018]. Beyond just a theoretical necessity, this restriction seems especially aligned with the motivating examples discussed.¹ To summarize, our goal is to optimize metric distortion using a limited number of pairwise comparison queries.

1.1 Our Approach and Results

We assume there are m alternatives, and we can ask each agent up to t pairwise comparison queries. A pairwise comparison query (or simply a query) involves two alternatives that an agent is asked to rank. Note that if t is quite large relative to m (e.g., at least $m \log_2(m)$), we can learn the full ranking of an agent.² However, in applications like RLHF and Polis, t is much smaller.

We first show that a query algorithm making at most t queries to each agent, even if it can ask these in an arbitrary order with arbitrary dependencies, will incur distortion $\Omega(m/t)$ (Section 3.1).

Next, we consider algorithms that match this lower bound. We are especially interested in algorithms that require few rounds of interactions, as it may not be reasonable to assume that agents return later in the process after answering queries. Of particular importance is the extreme case: *single-pass* algorithms needing only one round of interaction. Strikingly, we show that there exists a single-pass algorithm that achieves optimal distortion of $O(m/t)$ (Section 3.2).

Finally, we consider other kinds of restrictions on the adaptivity of the algorithm, i.e., how the queries depend on previous answers. We focus on two types of adaptivity. The first, called *inter-agent* adaptivity, indicates whether the queries made to an agent can depend on responses from other agents. The second, called *intra-agent* adaptivity, indicates whether the queries made to an agent can depend on previous responses from that agent. Here, we devise nearly tight (up to logarithmic factors) upper and lower bounds for the optimal distortion achievable under all combinations of these restrictions (Section 4). The results are summarized in Table 1. Note that all upper bounds are induced by single-pass algorithms, while the lower bounds hold for all algorithms, regardless of the number of passes.

¹In Polis, as part of the aggregation process, users and comments are mapped to a common low-dimensional feature space representing “opinion” space summarizing the various views. Within opinion space, voters are more likely to approve comments near them. Hence, Polis is implicitly assuming it is possible to map users/comments to a metric space. In RLHF, the possible models generating input/output prompts live in a feature space. An agent is assumed to give feedback according to a reward function of these features. If the reward function is concave, it is natural to assume that an agent has an ideal point maximizing their reward function and prefers inputs closer to this point.

²This is equivalent to using comparison-based queries to sort a list. Hence, this can be done via any sorting algorithm, e.g., Merge-Sort which takes at most $m \log_2(m)$ queries [Cormen *et al.*, 2001].

| Inter-agent | Intra-agent | |
|-------------|--|--|
| | Yes | No |
| Yes | $\Theta\left(\frac{m}{t}\right)$ | $O\left(\frac{m}{t} + \log m\right), \Omega\left(\frac{m}{t}\right)$ |
| No | $O\left(\frac{m^2 \log^2 t}{t^2}\right), \Omega\left(\frac{m^2}{t^2}\right)$ | $\Theta\left(\frac{m^2}{t}\right)$ |

Table 1: Summary of results for different inter- and intra-agent adaptivity combinations. All upper bounds are achieved with single-pass rules. All lower bounds hold regardless of the number of passes.

1.2 Related Work

The most closely related work is that of Anagnostides *et al.* [2022], who also consider the metric distortion problem while eliciting preferences. However, the preference elicitation method differs from ours; they assume that an algorithm can query two candidates $\{a, b\}$ and, in response, learns which of the two candidates a majority of voters prefer. While this could be implemented in our model, note that this would require multiple rounds of interaction and also makes additional restrictions, e.g., all voters will be asked the same queries, and only the majority winner is given, rather than the magnitude of this majority. Hence, we are interested in similar questions under a different query model, which we view as more suited for our motivating applications.

Slightly further afield is considering metric distortion, where only partial information about each agent’s ranking is given. For example, the same work [Anagnostides *et al.*, 2022] builds on work by Kempe [2020b] where voters reveal their top- k candidates for a fixed value $k \leq m$. Kempe [2020b] also considers the more general problem of b bits of information being communicated about a ranking using a communication scheme fixed upfront.

Metric distortion itself (using complete rankings) was introduced by Anshelevich *et al.* [2015]. They show that no social choice rule has distortion better than 3, while the well-known Copeland rule achieves distortion 5, among bounds for other well-known rules. Followup work proved bounds for even more rules [Skowron and Elkind, 2017; Goel *et al.*, 2017], but none beat this barrier of 5. Significant progress was made by Munagala and Wang [2019], who introduced a rule achieving distortion 4.236. The gap was finally closed by Gkatzelis *et al.* [2020], who gave a rule achieving distortion 3. Anshelevich *et al.* [2021] provide a survey of these results, as well as progress in non-metric distortion more broadly.

In addition to improvements in lowering distortion, other progress has been made. Kempe [2020a] gives an LP-duality framework for analyzing the distortion of existing rules. Kizilkaya and Kempe [2022] and Kizilkaya and Kempe [2023] both provide new rules achieving optimal metric distortion requiring much simpler analyses. We build on techniques from Kempe [2020a] as well as the rule of Kizilkaya and Kempe [2022]; a more in-depth description can be found when we use them.

Preference elicitation, more broadly, has a rich history in computational social choice with several variations. A more general summary can be found in Brandt *et al.* [2016];

however, we focus on the most closely related works here. Lu and Boutilier [2011] has the most similar style pairwise queries to our model; however, their objective is quite different from metric distortion. Amanatidis *et al.* [2021] study the tradeoff of distortion and query complexity; however, their queries are quite dissimilar, directly eliciting cardinal values and relative magnitudes rather than information about the ranking. Preference elicitation questions are also often studied through the lens of communication complexity, i.e., the number of bits that need to be revealed rather than the number of queries of a specific type. For example, Conitzer and Sandholm [2005] give the communication complexity of determining the winner of several common voting rules (where each agent privately knows their ranking). This has also been applied to non-metric distortion, where voters can reveal a certain number of bits about their utilities rather than just the rankings [Mandal *et al.*, 2019; Mandal *et al.*, 2020].

2 Model

For $s \in \mathbb{N}$, define $[s] = \{1, 2, \dots, s\}$.

Metric voting. We consider the setting of single-winner elections where we have a set C of m candidates and a set $V = [n]$ of n voters. We assume there is a (pseudo)metric³ over $C \sqcup V$ characterized by a distance function $d : (C \sqcup V) \times (C \sqcup V) \rightarrow \mathbb{R}_{\geq 0}$. That is, for all $i, j, k \in C \sqcup V$, we have that $d(i, i) = 0$, $d(i, j) = d(j, i)$, and $d(i, j) \leq d(i, k) + d(k, j)$ (the triangle inequality). Furthermore, each voter $i \in V$ has a *preference ranking* $\sigma_i = \sigma_i(1) \succ_i \dots \succ_i \sigma_i(m)$ such that $a \succ_i b$ implies that $d(i, a) \leq d(i, b)$ — we allow ties to be broken arbitrarily. We write $a \succeq_i b$ for $a \succ_i b$ or $a = b$. The collection of all voters' preference rankings form a *preference profile* $\sigma = (\sigma_1, \dots, \sigma_n)$. An *instance* is a tuple (d, σ) .

Utilitarianism. Given an instance (d, σ) , the *social cost* (or simply *cost*) of a candidate $c \in C$ is defined as $\text{cost}_d(c) := \sum_{i \in V} d(i, c)$. When d is clear from context, we may drop it from the notation by simply writing $\text{cost}(c)$. Under this measure, the optimal candidate is the one minimizing social cost. We will use the notation $\text{opt}(d, \sigma) := \min_{c \in C} \text{cost}_d(c)$ for the minimum cost.

Algorithms with pairwise comparison queries. We consider algorithms that do not have direct access to the underlying metric nor the preference profile; rather, they are only aware of C and V and can learn of voters' preferences via a sequence of *pairwise comparison queries*. That is, the algorithm can choose a voter i and two candidates $\{a, b\}$ and learn which candidate voter i prefers, $a \succ_i b$ or $b \succ_i a$. After receiving the responses, the algorithm (i.e., a voting rule) returns a candidate c as the winner. More formally, by a slight abuse of notation, we let $\mathcal{A}(d, \sigma)$ denote the candidate returned by an algorithm \mathcal{A} on instance (d, σ) .

We are interested in algorithms that do not make too many queries to each voter. If an algorithm makes at most t queries to each voter i , we call it a *t-query* algorithm.

³A pseudometric is a generalization of a metric that allows two distinct points to be distance 0 from each other. In other words, we allow voters and candidates to be at the same location.

Distortion. The quantitative measure of efficiency we investigate in this work is *distortion*. For a given instance (d, σ) , the distortion of candidate c is defined as

$$\text{dist}(c \mid (d, \sigma)) = \frac{\text{cost}_d(c)}{\text{opt}(d, \sigma)},$$

i.e., the ratio of the social cost of c to the optimal social cost. This value is always at least one, and the lower the ratio, the more efficient the candidate. The distortion of an algorithm \mathcal{A} is defined as

$$\text{dist}_m(\mathcal{A}) = \sup_{(d, \sigma), |C|=m} \text{dist}(\mathcal{A}(d, \sigma) \mid (d, \sigma)),$$

where the supremum is over all instances with m candidates and any number of voters. In other words, the distortion of an algorithm \mathcal{A} is its worst-case guarantee in terms of the multiplicative approximation to the optimal social cost.

2.1 Algorithm Restrictions

In the setup described so far, an algorithm can make queries to voters in an arbitrary order, with the option to revisit voters it has already queried or choose each subsequent query based on arbitrary prior responses. This flexibility may be costly or even impractical in certain scenarios. Hence, we are interested in the optimal distortion of algorithms subject to various constraints. We describe the restrictions of interest below.

Number of passes. We will say an algorithm is *p-pass* if the sequence of queries can be partitioned into at most p contiguous sequences where agent indices are non-decreasing. We will call it *single-pass* if it is 1-pass.

Inter-agent adaptivity. An algorithm is *inter-agent non-adaptive* if the queries it makes to voter i do not depend on responses from other voters $i' \neq i$. In other words, the choice of queries can be decomposed into n subprocedures A_1, \dots, A_n such that A_i only makes queries and learns the responses of voter i (the aggregation given these responses can be arbitrary). Note that an inter-agent nonadaptive algorithm can always be implemented as a single-pass algorithm because we can run these procedures sequentially (first A_1 , then A_2 , and so on).

Intra-agent adaptivity. An algorithm is *intra-agent non-adaptive* if the queries asked to voter i do not depend on responses to previous queries to that agent.

Full nonadaptivity. If an algorithm is inter- and intra-agent nonadaptive, we will call it *fully nonadaptive*. Such an algorithm can be implemented such that all queries are fixed upfront without depending on any of the responses.

3 t-Query Algorithms

We begin by considering optimal distortion bounds for *t*-query algorithms. First, we establish a lower bound for any algorithm. Then, we present an algorithm that not only matches this lower bound but is also single-pass.

3.1 Lower Bound

We start by establishing a lower bound on distortion when each agent can answer up to t queries without any other restrictions on the algorithm.

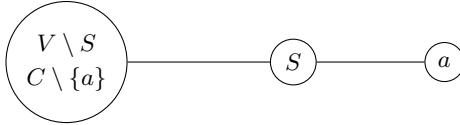


Figure 1: Metric for the case where the algorithm returns a .

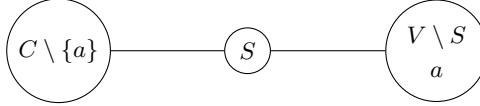


Figure 2: Metric for the case where the algorithm does not return a .

Theorem 1. All t -query algorithms incur a distortion of at least $\frac{m}{t} - 1$.

In fact, we show an even stronger result. This holds for any algorithm that makes at most tn queries in total *regardless* of how they are distributed among the agents. Any algorithm making t queries per voter trivially makes at most tn queries.

Proof. Let $R = \{a_1 \succ_{i_1} b_1, a_2 \succ_{i_2} b_2, \dots\}$ with $|R| \leq tn$ be an arbitrary sequence of responses the algorithm has queried and learned, and let c be the candidate returned given these responses. Since there are at most $2|R| \leq 2tn$ candidate appearances in $|R|$, there exists a candidate $a \in C$ that appears in less than $2tn/m$ many queries. Let $S \subseteq V$ be the set of voters that were ever queried on a . We have that $|S| \leq 2tn/m$.

We do a case analysis based on $c = a$ or $c \neq a$. The metrics are shown in Figures 1 and 2, where the distance function is defined by the path length. First, we show that there are profiles for both metrics that are consistent with the responses in R . In both metrics, voters in S are equidistant from all candidates, and voters in $V \setminus S$ are equidistant from all candidates except a . We construct the profile as follows. We choose rankings for voters in S in an arbitrary way consistent with R . For the voters in $V \setminus S$, depending on the case, we either place a at the top or bottom of their ranking and complete the rest of the ranking arbitrarily in a way consistent with R . Since no voter in $V \setminus S$ was queried on a , this is always possible, and by the equal distances, these profiles are consistent with the metrics.

Case $c = a$. Consider the metric in Figure 1. Note that $\text{cost}(a) = 2(|V| - |S|) + |S| = 2n - |S|$. For any candidate $c' \in C \setminus \{a\}$, $\text{cost}(c') = |S|$. Therefore,

$$\text{dist}(c) \geq \frac{\text{cost}(c)}{\text{cost}(c')} = \frac{2n - |S|}{|S|} \geq \frac{m}{t} - 1.$$

Case $c \neq a$. Consider the metric in Figure 2. Note that $\text{cost}(a) = |S|$ and, since $c \neq a$, $\text{cost}(c) = |S| + 2(|V| - |S|)$. As in the previous case, we get a lower bound of $m/t - 1$. \square

3.2 A Matching Upper Bound

Next, we consider designing algorithms to match this lower bound. Recall that using $m \log_2 m$ queries, it is possible to learn an agent's entire ranking. Hence, if t is sufficiently large ($\geq m \log_2 m$), we can simply learn all of the agents' complete rankings and then run any constant-distortion rule

Algorithm 1 Single-Pass Plurality Veto

```

1: Initialize  $\text{plu}(c) = 0$ 
2: for each voter  $i$  among the first  $\lceil |V|/2 \rceil$  voters do
3:    $c_i \leftarrow \text{top}(\sigma_i, C)$ 
4:    $\text{plu}(c_i) \leftarrow \text{plu}(c_i) + 1$ 
5: for each voter  $i$  among the next  $\lceil |V|/2 \rceil - 1$  voters do
6:    $C \leftarrow \{c \in C \mid \text{plu}(c) \geq 1\}$ 
7:    $c'_i \leftarrow \text{bottom}(\sigma_i, C)$ 
8:    $\text{plu}(c'_i) \leftarrow \text{plu}(c'_i) - 1$ 
return the unique candidate  $a$  with  $\text{plu}(a) \geq 1$ .

```

Algorithm 2 t -Query Single-Pass Adaptive

```

1:  $R \leftarrow \lceil (m-1)/t \rceil$ ;  $C_{\text{active}} \leftarrow C$ 
2: for each round  $r = 1$  to  $R$  do
3:    $C_r \leftarrow \min\{t+1, |C_{\text{active}}|\}$  candidates from  $C_{\text{active}}$ 
4:    $V_r \leftarrow$  the next  $\lceil n/R \rceil$  voters
5:    $c_r^* \leftarrow \text{Algorithm 1}(C \leftarrow C_r, V \leftarrow V_r)$ 
6:    $C_{\text{active}} \leftarrow (C_{\text{active}} \setminus C_r) \cup \{c_r^*\}$ 
return the unique candidate  $a \in C_{\text{active}}$ 

```

such as *Copeland* [Anshelevich *et al.*, 2015] or *Plurality Veto* [Kizilkaya and Kempe, 2022].

In fact, Plurality Veto can be implemented using $O(m)$ queries. As this will be useful to our later analysis, we describe the rule and query implementation in more detail. Plurality Veto runs in two phases. In the first, it determines the *plurality score* of each candidate, i.e., the number of voters that ranks that candidate first. The second begins with each candidate having a number of points equal to their plurality score. It sequentially goes through the voters (in an arbitrary order), and for each, asks the voter their *least favorite* candidate who still has a nonzero number of points; it then decrements that candidate by one point. The winning candidate is the last one remaining just before the final voter is asked to decrement.⁴

Note that both determining a favorite or least favorite candidate among a set S requires $|S| - 1$ pairwise queries. Hence, this can be implemented by asking each voter at most $2m - 2$ pairwise queries. Therefore, when $t \geq 2m - 2$, Plurality Veto has optimal distortion.

However, there are still two drawbacks to this result. First, even though Plurality Veto requires only $O(m)$ queries, it needs to do this in two phases. Hence, its naïve implementation requires two passes. We may wonder if the same can be done with a single-pass algorithm. Second, and more important, this gives us no insight for how to handle $t < 2m - 2$, a crucial case for many applications. Our next result handles both of these matters, devising a single-pass algorithm that yields asymptotically optimal distortion for all t .

Theorem 2. For $n \geq \lceil \frac{m-1}{t} \rceil$ and $t \leq m - 1$, Algorithm 2 is a single-pass, t -query algorithm achieving distortion $6 \lceil \frac{m-1}{t} \rceil + 1$.

⁴Note that the total number of points given out is n , and each voter decreases by one. Hence, just before the final voter is asked, exactly one candidate will remain with a single point.

The formal proof of the theorem appears later in this section. Algorithm 2 uses Algorithm 1 as a subroutine on subsets of candidates over a sequence of $\lceil \frac{m-1}{t} \rceil$ rounds.⁵ This subroutine is essentially a variant of Plurality Veto which runs in a single pass. In the special case of $t = m - 1$, there is only one round, and the subroutine is used “as is” on the entire candidate set. To better illustrate the ideas, we first prove lemmas that are sufficient to show Algorithm 1 achieves the desired distortion bound in this special case. We then build on these ideas to handle the general case of $t \leq m - 1$.

Warm-up Case: $t = m - 1$

In this case, Algorithm 2 simply runs Algorithm 1 and returns the chosen candidate. Hence, we directly analyze Algorithm 1.

Algorithm Description. Algorithm 1 uses two sequences of adaptive queries:

- $\text{top}(\sigma_i, C')$: A sequence of adaptive queries that finds the most preferred candidate of voter i among a subset of candidates $C' \subseteq C$. This can be done using $|C'| - 1$ queries (compare the first two candidates, then compare the preferred one with the third candidate, and so on).
- $\text{bottom}(\sigma_i, C')$: Similar to above but for the least preferred alternative of voter i among $C' \subseteq C$.

Algorithm 1 accumulates plurality scores of the first $\lceil |V|/2 \rceil$ voters in $\text{plu}(c)$ (lines 2–4). Then, for each of the second $\lceil |V|/2 \rceil - 1$ voters, among the remaining candidates with positive plurality score C_{active} , the voter indicates her least preferred candidate c'_i by decreasing the candidate’s plurality score by one (lines 5–8). Since the number of increments is exactly one more than the number of decrements, only one candidate remains with a positive score, which subsequently is returned.

Distortion Analysis. To prove the distortion bound, we need the following technical definition and lemmas.

Definition 1 $((a, b)\text{-path})$. Fix two candidates $a, b \in C$. A sequence of distinct voters v_1, \dots, v_k is called a voter path from a to b (or an $(a, b)\text{-path}$ for short) if there exists a sequence of $k + 1$ not necessarily distinct candidates beginning with a and ending with b , $a = c_0, c_1, \dots, c_{k-1}, c_k = b$ such that, for each $j \in [k]$, $c_{j-1} \succeq_{i_j} c_j$. Slightly abusing notation, we will refer to a set of voters $V' \subseteq V$ as an $(a, b)\text{-path}$ if some order of V' is an $(a, b)\text{-path}$. We will say two $(a, b)\text{-paths}$ are disjoint if they do not share any voters.

The following is essential to the distortion analysis of Algorithms 1 and 2. The proof appears in the full version.⁶

Lemma 1. Fix $a, b \in C$. Suppose there exist ℓ disjoint $(a, b)\text{-paths}$. Then,

$$\frac{\text{cost}(a)}{\text{cost}(b)} \leq \frac{2n}{\ell} + 1.$$

The definition of $(a, b)\text{-paths}$ and the result of Lemma 1 are similar to the definition of chains of preferences and Corol-

⁵Note that the $n \geq \lceil \frac{m-1}{t} \rceil$ requirement simply ensures that we can use at least one distinct voter per round.

⁶Full version is available at <https://daniel-halpern.com/files/metric-distortion-with-queries.pdf>.

lary 5.3 of Kempe [2020a]. The key difference is that we require the voters along each path to be distinct, which allows us to prove stronger properties.

Next, using Lemma 1, we show that there are $\lceil n/2 \rceil$ disjoint voter paths from the candidate returned by Algorithm 1 to each other candidate.

Lemma 2. There exist $\lceil |V|/2 \rceil$ disjoint voter paths from the candidate selected by Algorithm 1 to any other candidate.

Proof. Let a be the candidate returned by Algorithm 1 and $b \in C \setminus \{a\}$ be an arbitrary candidate. Let V_1 be the first $\lceil |V|/2 \rceil$ voters initializing the plurality scores and V_2 be the second set of voters decreasing the scores. Note that $|V_1| - 1 = |V_2|$. Create a one-to-one function $g: V_2 \rightarrow V_1$ as follows. As per the algorithm, for $i \in V_1$, c_i is the most favorite candidate for i ; and, for $i \in V_2$, c'_i is the candidate indicated by i as her least favorite among the still active candidates (line 7). Map i to a corresponding voter $g(i)$ increasing the plurality score of c'_i . This way, g is injective and $c'_i = c_{g(i)}$.

Next, we show there exist $\lceil n/2 \rceil$ disjoint voter paths from a to b . First, we claim that for each $i \in V_2$, $\{i, g(i)\}$ is an $(a, b)\text{-path}$ via the candidate sequence a, c'_i, b . Indeed, for the first comparison, $a \succ_i c'_i$ because i indicates c'_i as her least preferred alternative among the active set which necessarily contains a as it remained active until the end. For the second, $c'_i \succ_{g(i)} b$ because $c'_i = c_{g(i)}$, and $c_{g(i)}$ was $g(i)$ ’s most preferred alternative. In addition, consider the single voter $i^* \in V_1 \setminus \{g(i') \mid i' \in V_2\}$ who does not appear in the range of g . Her top vote must have been a as a is the only candidate with positive score in the end. Thus, $a \succ_{i^*} b$ and the set $\{i^*\}$ alone is a voter path from a to b . Note that each of the above voter paths are disjoint, so we have constructed $|V_2| + |\{i^*\}| = \lceil |V|/2 \rceil$ disjoint voter paths from a to b . \square

As an aside, note that if we directly apply Lemma 1 to Lemma 2, this shows that Algorithm 1 achieves distortion 5.

General Case: $t \leq m - 1$

Now, we analyze Algorithm 2 in full generality.

Algorithm Description. Algorithm 2 runs in $R = \lceil (m-1)/t \rceil$ rounds. In each round $r \in [R]$, it takes up to $t + 1$ candidates from the active candidate set C_{active} , denoted C_r . It then runs Algorithm 1 with a new set of $\lceil n/R \rceil$ voters, denoted V_r , along with the candidates in C_r , which return a candidate $c_r^* \in C_r$. Then, Algorithm 2 removes all candidates in C_r from the active candidate set C_{active} except for c_r^* . Because of this removal procedure, in each of the first $R - 1$ rounds, it eliminates t candidates from the active set. In the final round, there are at most $m - (R-1)t \leq 1 + Rt - (R-1)t \leq t + 1$ candidates remaining, from which a single winning candidate is kept in the active set and subsequently returned.

Distortion Analysis. To utilize Lemma 1 similar to that in Lemma 2, we find a sufficient number of disjoint voter paths from the candidate returned by Algorithm 2 to each of the other candidates.

Proof of Theorem 2. Let a be the candidate returned by Algorithm 2 and $b \in C \setminus \{a\}$ be an arbitrary candidate.

Consider a directed graph G over candidates described as follows. At each round r , draw a directed edge from the winning candidate c_r^* to all candidates $C_r \setminus \{c_r^*\}$. Note that each candidate $c \neq a$ has exactly one incoming edge, i.e., if c is eliminated at round r , the directed edge $c_r^* \rightarrow c$ is in G . Thus, G forms a directed rooted tree at a . Let $a = c_0 \rightarrow \dots \rightarrow c_k = b$ be the unique directed path in G from a to b . For all $j \in [k]$, the edge $c_{j-1} \rightarrow c_j$ is created in a different round r_j . By Lemma 2, there are $\ell := \lceil \lfloor n/R \rfloor / 2 \rceil$ voter paths from c_{j-1} to c_j on a disjoint set of voters V_{r_j} . We have that $\ell \geq n/(3R)$ because $n \geq R$ and $\lceil \lfloor x \rfloor / 2 \rceil \geq x/3$ for $x \geq 1$. Note that V_{r_j} 's are disjoint as the algorithm selects a new set of voters at each round. Thus, we can extend the ℓ paths from c_0 to c_1 with the ℓ paths from c_1 to c_2 by arbitrarily matching those two sets of paths and get ℓ disjoint paths from c_0 to c_2 . Then, we extend the paths from $c_0 = a$ to c_3 and so on to $c_k = b$. Therefore, by Lemma 1 and that this holds for all $b \in C \setminus \{c\}$ we have

$$\text{dist}(a) \leq \frac{2n}{\ell} + 1 \leq 6R + 1 = 6 \left\lceil \frac{m-1}{t} \right\rceil + 1. \quad \square$$

where in the last inequality we used $n \geq R = \lceil \frac{m-1}{t} \rceil$ and that $\lfloor x \rfloor \geq \frac{x}{2}$ for $x \geq 1$.

4 Algorithms with Adaptivity Constraints

We now turn to varying restrictions on algorithm adaptivity. We begin with a tight analysis of the fully nonadaptive setting (Section 4.1). Then, we relax the constraints by requiring only one of inter- and intra-agent nonadaptivity at a time (Sections 4.2 and 4.3).

Note that when algorithms are intra-agent nonadaptive (Sections 4.1 and 4.3), we can no longer learn each agent's full ranking with $O(m \log m)$ queries. To do so, we instead may have to ask about all $\binom{m}{2} = \Theta(m^2)$ possible pairwise comparisons. Hence, even for $t = \omega(m \log m)$, we may not be able to achieve constant distortion. Furthermore, recall that any inter-agent nonadaptive algorithms can trivially be implemented to be single-pass.

4.1 Fully Nonadaptive Algorithms

We show that the optimal distortion value for t -query fully nonadaptive algorithms is $\Theta(m^2/t)$. This is a factor of m worse than the unrestricted setting.

Theorem 3. *All fully nonadaptive t -query algorithms incur a distortion of at least $\frac{m^2}{t} - 1$. Furthermore, for $t \leq \binom{m}{2}$ and $n \geq \binom{m}{2}/t$, there exists a fully nonadaptive t -query algorithm bound with distortion at most $3m^2/t + 1$.*

We defer the lower bound to the full version. Below, we outline the algorithm and then analyze its distortion.

Algorithm Description. There are $\binom{m}{2}$ pairs of candidates. The algorithm distributes queries among voters as equally as possible in a way that each pair of candidates $(c, c') \in \binom{C}{2}$ is given to at least $\lfloor nt/\binom{m}{2} \rfloor$ voters. This can be done in various ways. For instance, concatenate $\lceil nt/\binom{m}{2} \rceil$ copies of the list of all $\binom{m}{2}$ pairs and assign the i th voter the pairs in range $[(i-1)t+1, (i+1)t]$. Such a distribution can be determined beforehand and, thus, is fully nonadaptive.

The algorithm aggregates the results as follows. For each pair $c, c' \in C$, draw a directed edge from c to c' if a majority of voters queried on $\{c, c'\}$ preferred c to c' (in the case of a tie, pick an arbitrary direction). Note that this is now a *tournament* graph, where between every two nodes, there is one directed edge. Finally, from this graph, select a *king* vertex a . A vertex is a king if it reaches all other vertices by at most two edges. In other words, a is a king if for all other candidates $b \in C \setminus \{a\}$, either a has an edge to b or there exists another candidate c such that a has an edge to c and c has an edge to b . It is well-known that king vertices exist in tournament graphs (e.g., West [2001]).

Distortion Analysis. To prove Theorem 3, we utilize the following lemma derived in Kempe [2020a] (a special case of Corollary 5.3).

Lemma 3. *If at least ℓ voters prefer a to b , then $\frac{\text{cost}(a)}{\text{cost}(b)} \leq \frac{2n}{\ell} - 1$. Furthermore, if there is a candidate $c \in C \setminus \{a, b\}$ such that f voters prefer a to c , and also ℓ voters prefer c to b , then $\frac{\text{cost}(a)}{\text{cost}(b)} \leq \frac{2n}{\ell} + 1$.*

Proof of Theorem 3. Let a be the king vertex in the pairwise-majority graph described above. Fix a candidate $b \in C \setminus \{a\}$. Since a is a king vertex, there is a path of length at most 2 from a to b . Each edge indicates the existence of $\lceil \frac{1}{2} \lfloor nt/\binom{m}{2} \rfloor \rceil$ voters preferring the candidate on the tail of the edge to the candidate on the head. Since $n \geq \binom{m}{2}/t$ and $\lceil \frac{1}{2} \lfloor x \rfloor \rceil \geq x$ for $x \geq 1$, $\lceil \frac{1}{2} \lfloor nt/\binom{m}{2} \rfloor \rceil \geq nt/(3\binom{m}{2}) \geq (2nt)/(3m^2)$. By Lemma 3, $\text{dist}(a) \leq 3m^2/t + 1$. \square

4.2 Inter-Agent Nonadaptive Algorithms

Moving away from the most restricted case of full nonadaptivity, we find that by allowing intra-agent adaptivity, we can design algorithms that achieve better distortion by a factor of roughly $\tilde{O}(t)$ (up to logarithmic factors). Only proof sketches are given. The complete proofs of Theorems 4 and 5 can be found in the full version. Our upper bound builds on the ideas from the fully nonadaptive setting.

Theorem 4. *For $t \leq m \log_2 m$ and $n \geq 2m^2 \log_2^2(t+1)/t^2$, there exists an inter-agent nonadaptive t -query algorithm achieving a distortion of $\frac{10m^2 \log_2^2(t+1)}{t} + 1$.*

Proof Sketch. The key idea of the algorithm from the fully nonadaptive case was to have a sufficient number of pairwise comparisons between all pairs. Exploiting intra-agent adaptivity, for each voter, we can sort ℓ candidates based on her preferences using $O(\ell \log \ell)$ queries. From a sorted list of ℓ candidates, we can infer $\binom{\ell}{2}$ pairwise comparisons — candidates appearing earlier are preferred to the ones appearing later. This is significantly stronger than the one learned comparison per query in the fully nonadaptive setting. Using this observation, we query voters to sort different sets of ℓ candidates. Similar to Theorem 3, we select a king vertex from the majority pairwise comparison graph. \square

We complement the above theorem by showing its optimality up to logarithmic factors with a surprisingly-intricate lower bound.

Theorem 5. All inter-agent nonadaptive algorithms have distortion at least $m^2/(4t^2)$.

Proof Sketch. Recall query choices of an inter-agent nonadaptive algorithm \mathcal{A} can be decomposed into n subroutines A_1, \dots, A_n , one for each agent, that chooses the next query based on previous responses of that agent. The proof incorporates ideas from showing lower bounds on comparison-based sorting algorithms, namely, representing these algorithms as binary decision trees. Each node corresponds to a query, and the two branches correspond to the two possible responses. With this representation, we are able to show that, for many pairs of candidates $\{a, b\}$, each algorithm A_i will find two rankings of the form $\sigma^1 = a \succ b \succ \dots$ and $\sigma^2 = b \succ a \succ \dots$ indistinguishable, i.e., they will both be asked the same queries and receive the same responses.

From this, we can find a single pair $\{a, b\}$ such that this property holds for a large fraction of the agents. The algorithm will be unable to distinguish the case of a large fraction of agents having a as their first choice and b as their second vs b as their first choice and a as their second. Since it must make the same choice in these two instances, one of these must be suboptimal. The final piece is to construct a simple metric in which choosing the “wrong” candidate on one of these profiles leads to large distortion. \square

4.3 Intra-Agent Nonadaptive Algorithms

We now turn to designing t -query algorithms that are *intra-agent* nonadaptive, i.e., each voter is asked up to t pairwise comparison queries all at once and returns her answers to all the t queries together. Its distortion guarantee is as follows.

Theorem 6. For $n \geq \lfloor m/t \rfloor + \lceil \log_2 t \rceil$ and $t \leq m - 1$, Algorithm 3 is a single-pass, intra-agent nonadaptive, t -query algorithm that achieves a distortion of at most $6(\lfloor m/t \rfloor + \lceil \log_2 t \rceil) + 1$.

Algorithm Description. Algorithm 3 runs in R rounds (for R that will be determined later) querying an equal number of $\lfloor n/R \rfloor$ new voters in each. For round $r \in [R]$, the corresponding group of voters are queried on ℓ_r disjoint pairs of candidates, arbitrarily selected from the active candidate set C_{active} . For each of the ℓ_r pairs, the candidate losing the majority of the $\lfloor n/R \rfloor$ votes is eliminated from C_{active} . This process goes on until one candidate remains, which is subsequently returned.

Since we want at most t queries made to each voter, in each round, we can ask about at most t pairs (i.e., $\ell_r \leq t$). While $|C_{\text{active}}| \geq 2t$, there are enough candidates such that the algorithm can, in fact, make all t queries (i.e., $\ell_r = t$). However, after $\lfloor m/t \rfloor - 1$ rounds, fewer than $2t$ candidates will remain. Then, the algorithm pairs up as many candidates as possible in C_{active} (one candidate will be unpaired when $|C_{\text{active}}|$ is odd), and $\ell_r = \lfloor |C_{\text{active}}|/2 \rfloor$ candidates are eliminated. This can go on for at most $\lceil \log_2 2t \rceil \leq \lceil \log_2 t \rceil + 1$ many rounds until a single candidate remains. Hence, there will be at most $\lfloor m/t \rfloor + \lceil \log_2 t \rceil$ rounds.

The complete proof can be found in the full version.

While the lower bound of Theorem 1 shows the optimality of the above result for $t = O(m/\log m)$, it is unclear what the best achievable distortion is for $t = \omega(m/\log m)$.

Algorithm 3 t -Query Single-Pass Intra-agent Nonadaptive

```

1:  $R \leftarrow \lfloor m/t \rfloor + \lceil \log_2 t \rceil$ ;  $C_{\text{active}} \leftarrow C$ 
2: for each round  $r = 1$  to  $R$  do
3:    $\ell_r \leftarrow \min\{t, \lfloor |C_{\text{active}}|/2 \rfloor\}$ 
4:   Let  $C_r \leftarrow \{(a_{r,1}, b_{r,1}), (a_{r,2}, b_{r,2}), \dots, (a_{r,\ell_r}, b_{r,\ell_r})\}$ 
      be  $\ell_r$  pairs of distinct candidates from  $C_{\text{active}}$ 
5:   for each voter  $i$  among the next  $\lfloor n/R \rfloor$  voters do
6:     Present the  $\ell_r$  pairs  $(a_{r,1}, b_{r,1}), \dots, (a_{r,\ell_r}, b_{r,\ell_r})$ 
7:   for each pair in  $\{(a_j, b_j) \mid j \in [\ell_r]\}$  do
8:     Let  $c_{r,j}$  be the candidate losing the majority of
       comparisons (ties broken arbitrarily)
9:    $C_{\text{active}} \leftarrow C_{\text{active}} \setminus \{c_{r,j}\}$ 
return the unique candidate  $a \in C_{\text{active}}$ 

```

Theorem 6 continues to give a $O(\log m)$ upper bound, but the lower bound becomes $O(m/t)$, and only constant once $t = \Omega(m)$. For $t = \Omega(m^2)$, we know it is possible to achieve $O(1)$ distortion because, at this point, it is possible to learn all agents’ complete rankings and run any of the well-known constant-distortion rules. However, we leave it as an interesting open question to close this gap for intermediate values of t , or find the minimum t with constant distortion.

5 Discussion

To summarize, we have studied the design of a limited number of pairwise queries to optimize metric distortion. We considered a variety of constraints on our algorithms regarding both the rounds of interactions and the dependency of queries on previous answers. In all cases, we provide either tight or nearly tight upper and lower bounds. In the case of inter- and intra-agent nonadaptivity, there remain logarithmic gaps.

As for follow-up work, the most direct would be to close these gaps. However, more broadly, while requiring few rounds of interaction and both inter- and intra-agent nonadaptivity capture an interesting set of constraints on the algorithms, there are a plethora of reasonable notions one could consider. For example, we may wish for the algorithm to be *anonymous*, demanding that each agent’s queries depend only on their ranking, σ_i , rather than their identity. An even stricter property is to require that all agents receive the *exact* same queries, ensuring a strong notion of equality in how different agents interact with the system. We leave optimizing distortion in these settings and the development of new desirable properties as intriguing directions for future work.

Lastly, in certain applications, there can be such an enormous number of alternatives that our bounds become completely impractical. For instance, in variations of Polis, the domain of alternatives can be extended beyond just submitted comments to *any* statement from a natural language, a seemingly limitless domain. In such cases, when it is impossible to ensure that each candidate is queried by at least one voter, achieving bounded distortion is hopeless. We need to either (i) make additional assumptions about the domain of alternatives (say, they arise from a low-dimensional feature space), (ii) consider queries that are more powerful than pairwise comparisons, or (iii) use other measures of efficiency. We again leave these as exciting directions for further work.

References

[Amanatidis *et al.*, 2021] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, and A. A. Voudouris. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. *Artificial Intelligence*, 296:103488, 2021.

[Anagnostides *et al.*, 2022] I. Anagnostides, D. Fotakis, and P. Patsilinakos. Metric-distortion bounds under limited information. *Journal of Artificial Intelligence Research*, 74:1449–1483, 2022.

[Anshelevich *et al.*, 2015] E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 777–783, 2015.

[Anshelevich *et al.*, 2018] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.

[Anshelevich *et al.*, 2021] E. Anshelevich, A. Filos-Ratsikas, N. Shah, and A. A. Voudouris. Distortion in social choice problems: The first 15 years and beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4294–4301, 2021. Survey Track.

[Brandt *et al.*, 2016] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[Conitzer and Sandholm, 2005] V. Conitzer and T. Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 78–87, 2005.

[Cormen *et al.*, 2001] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2nd edition, 2001.

[Gkatzelis *et al.*, 2020] V. Gkatzelis, D. Halpern, and N. Shah. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, pages 1427–1438, 2020.

[Goel *et al.*, 2017] A. Goel, A. K. Krishnaswamy, and K. Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 287–304, 2017.

[Kempe, 2020a] D. Kempe. An analysis framework for metric voting based on LP duality. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2079–2086, 2020.

[Kempe, 2020b] D. Kempe. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2087–2094, 2020.

[Kizilkaya and Kempe, 2022] F. E. Kizilkaya and D. Kempe. Plurality veto: A simple voting rule achieving optimal metric distortion. In *Proceedings of the 31th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 349–355, 2022.

[Kizilkaya and Kempe, 2023] F. E. Kizilkaya and D. Kempe. Generalized veto core and a practical voting rule with optimal metric distortion. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, page 913–936, 2023.

[Lu and Boutilier, 2011] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 287–293, 2011.

[Mandal *et al.*, 2019] D. Mandal, A. D. Procaccia, N. Shah, and D. P. Woodruff. Efficient and thrifty voting by any means necessary. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 7178–7189, 2019.

[Mandal *et al.*, 2020] D. Mandal, N. Shah, and D. P. Woodruff. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 795–813, 2020.

[Munagala and Wang, 2019] K. Munagala and K. Wang. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 245–262, 2019.

[Procaccia and Rosenschein, 2006] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006.

[Skowron and Elkind, 2017] P. Skowron and E. Elkind. Social choice under metric preferences: scoring rules and STV. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 706–712, 2017.

[Small *et al.*, 2021] C. Small, M. Björkegren, T. Erkkilä, L. Shaw, and C. Megill. Polis: Scaling deliberation by mapping high dimensional opinion spaces. *Revista De Pensament I Anàlisi*, 26(2), 2021.

[West, 2001] D. B. West. *Introduction to Graph Theory*, volume 2. Prentice Hall Upper Saddle River, 2001.