

CS634 Course Project Proposal: Metric Distortion

Qilin Ye

October 19, 2025

The Voting Problem & The Metric Distortion Framework

In this project, I will consider the **metric distortion** problem. Collective decision-making lies at the core of both democratic governance and algorithmic social choice, and has been studied for decades across multiple disciplines. In many decision-making settings, given a collection of candidates (alternatives) and a population of voters, one wants to design an algorithm that selects the best candidate with respect to the electorates overall welfare or satisfaction.

In computer science, especially in recent years, this problem is widely examined under the metric distortion framework. A social choice rule observes only *ordinal/qualitative* rankings of candidates and must remain oblivious to any *cardinal/quantitative* information about feasible latent metrics; the natural objective is therefore to design a decision rule that preserves social welfare as well as possible given only ordinal data. The unavoidable information loss incurred by downgrading from cardinal to ordinal inputs is captured by the *distortion* of the rule. This viewpoint emphasizes geometric reasoning—triangle inequalities, medians, and proximity structure—as a substitute for the missing cardinal information and provides a principled standard for comparing rules.

Formally, let V be the set of voters, C the set of candidates, and let d be a metric on $V \cup C$. Each voter $v \in V$ induces an ordinal ranking of C by sorting alternatives in increasing order of distance $d(v, \cdot)$. A rule f takes only these ordinal rankings as input and outputs a winner $X \in C$. Fix an instance I consisting of voters, candidates, and ordinal preferences, and consider a social-cost objective such as the utilitarian cost $SC(X) = \sum_{v \in V} d(v, X)$. Put in simple words, the distortion of f on this instance is the ratio between $SC(X)$ and the cost of the optimal alternative $X^* \in \arg \min_{Y \in C} SC(Y)$, and the distortion of f as a rule itself is the supremum of this ratio over all metrics consistent with the ordinal rankings of the voters. The lower the distortion, the better.

It is well known that any deterministic rule that uses only ordinal information must have distortion at least 3, which sets a benchmark for what can be guaranteed in the absence of cardinal data.

Tournament rules. A prominent and well-studied subclass of ordinal rules are *tournament rules*, which base decisions on the outcomes of pairwise contests between alternatives. Tournament rules are attractive for their analytical simplicity and low cognitive burden, since voters only compare two candidates at a time. However, when restricted to pairwise information alone, deterministic tournament rules face a lower bound of about 3.11, slightly worse than the deterministic optimum achievable by general ordinal rules.

Deliberation. While the rule itself should not access cardinal distances, preferences in practice are not static: individuals exchange arguments and often refine their judgments after discussion. A line of work studies *deliberation*,

where small groups of voters refine their pairwise comparisons and report improved *ordinal* information to the rule. Intuitively, these group-based comparisons are indirectly informed by cardinal proximity, and even groups of size $k = 3$ breaks the lower bound 3 of any deterministic ordinal-information-only rule. That said, purely two-person deliberation may not suffice for improvements in all models, and some prior approaches rely on numerically solving non-convex programs, which obscures analytical insight.

Within the metric distortion framework, we ask:

Can pairwise ($k = 2$) deliberation, where voters refine their pairwise comparisons through discussion, improve the distortion of tournament rules while preserving their simplicity?

Expected Outcomes

A disclaimer first: I chose metric distortion because it is closely related my ongoing project with Professor Kamesh Munagala, in which I have shown some powerful and elegant results using mostly geometric analysis¹. I will make sure to include additional results specific to CS 634, sufficiently substantial to count toward a course project.

The project's deliverable will be a self-contained, polished exposition on metric distortion². In the write-up, I will:

- (Status: TODO) Provide a concise overview of the evolution of algorithms and impossibility results within the metric distortion framework.
- (Status: in-progress, half finished) More importantly, introduce a simple and transparent tournament-rule-based *deliberation via matching* protocol. So far, I have proven (just hours ago, assuming the proof is valid...) that it attains distortion 3, thereby improving upon the 3.11 lower bound for deliberation-free tournament rules and matching the deterministic optimum achievable without deliberation. In the meantime, I will continue to look for further improvements to my algorithm, and any additional findings can go solely into this project.

This part of deliverable will emphasize a novel geometric analysis of deliberation that, to my knowledge, is the first analytic framework reducing a highly nonlinear, non-convex formulation to a convex — even linear — one that is tractable and intuitive.

- (Status: TODO) If time and scope permit, I will also briefly analyze and present a randomized variant of *deliberation via matching*. Randomized algorithms typically enjoy superior expected performance. This step is purely speculative, but I expect similar results to hold in the *deliberation via matching* protocol.

¹If this is advised against, I'm more than willing to choose a less related topic.

²No in-class presentation, right?