Aggregating Quantitative Relative Judgments: From Social Choice to Ranking Prediction

Yixuan (Even) Xu
Carnegie Mellon University

Hanrui Zhang
Chinese University of Hong Kong

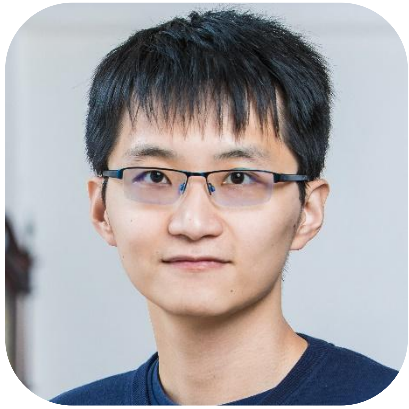

Yu Cheng
Brown University

Vincent Conitzer
Carnegie Mellon University

Ranking Aggregation

- Aggregating multiple input rankings into an integrated one

- The problem is of interest in multiple research communities
  - **Voting theory:** each **voter** ranks the candidates, and a **voting rule** decides a winning candidate or a ranking of all candidates
  - **Learning-to-rank:** ranking **web pages** in response to a search query, or ranking **recommendations** to a user

- **Common ground:** there is a latent “true” ranking of the elements, of which all inputs are just noisy observations

Focuses of Different Communities

- **Voting theory** (social choice)
  - Inputs of aggregation are usually **subjective**
  - Desiderata: **transparency**, **simple** voting rule, **strategy-proofness**
- **Learning-to-rank**
  - Inputs of aggregation are usually **objective**
  - Desiderata: **relevance** to the search, recommendation **quality**
- This work: an attempt to bridge the two communities

Judgment Aggregation

- Quantitative judgment aggregation
  - A way to think of ranking aggregation **in social choice**
  - Inputs: quantitative relative judgments $\{(a, b, y)\}$
    - "Candidate a is better than candidate b by y units quantitatively"

- We observe that the relative "judgments" can be produced by an **objective process** other than a subjective agent reporting
  - Applying formulations from social choice to learning-to-rank inputs
  - This conceptually bridges the two communities

Example Application: Races

- Races are one example of objective judgments

- Simple methods like mean / median are not good enough
  - Bob seems to be faster than Charlie judging from the Chicago race
  - But mean / median draws the opposite conclusion

	Boston	New York	Chicago
Alice	4:00:00	4:10:00	3:50:00
Bob	4:11:00	4:18:00	4:01:00
Charlie	N/A	N/A	4:09:00

QRJA Problem Formulation

- Given a set of m quantitative relative judgments $\{(a_i, b_i, y_i)\}$ and their weights $\{w_i\}$, find a vector $\boldsymbol{x} \in R^n$ that minimizes

$$\sum_{i=1}^{m} w_i \cdot f\left(\left|(x_{a_i} - x_{b_i}) - y_i\right|\right)$$

- $f: R \to R$ maps the inconsistency with inputs to loss
  - $\boldsymbol{f(x) = x}$**:** prior work (Conitzer et al., 2016, Zhang et al., 2019)
  - **If** $\boldsymbol{f(x)}$ **is convex:** solvable in polynomial time
  - $\boldsymbol{f(x) = x^p}$**:** **The focus of this work**

Computational Complexity

- We provide a tight characterization of ℓ_p QRJA's complexity
  - When $p \geq 1$, ℓ_p QRJA can be solved in **almost-linear** time $O(m^{1+o(1)})$
  - When $p < 1$, ℓ_p QRJA is **NP-Hard**, and there is no FPTAS

- Additionally, we show that when $p \in [1, 2]$ and $m \gg n$, we can **reduce $\boldsymbol{m}$ to $\boldsymbol{\tilde{O}(n)}$** while incurring a small error

Experiments

- We conduct experiments on real-world race data
  - **Datasets:** F1 races, marathon, programming contests, chess, etc.
  - **Our algorithms:** ℓ_1 and ℓ_2 QRJA

- Benchmarks
  - **Simple** benchmarks: Mean, Median
  - From **social choice:** Borda, Kemeny-Young
  - From **learning-to-rank:** Matrix Factorization

- We look at ordinal accuracy and quantitative loss

Experiments

- Both MF and QRJA are never significantly worse than the best-performing algorithm on any of the tested datasets, and QRJA additionally offers an interpretable model

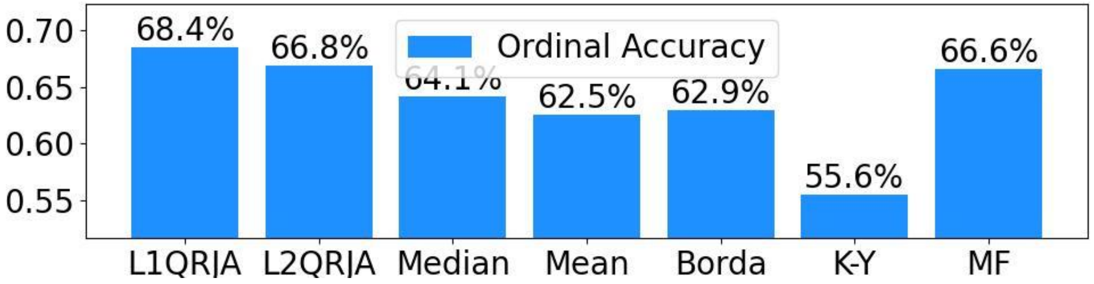

Chess

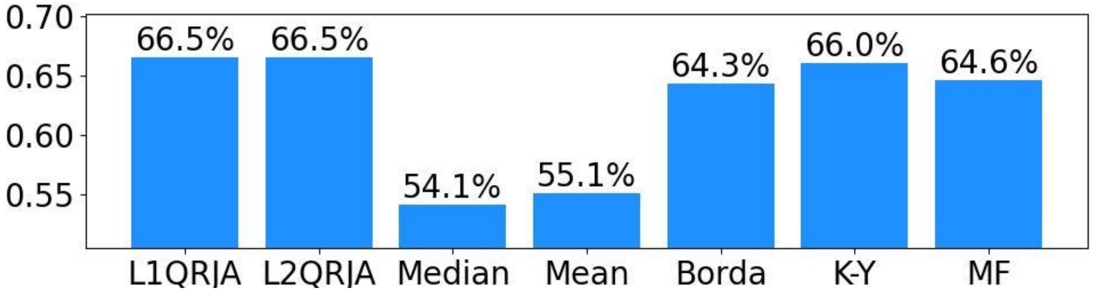

F1

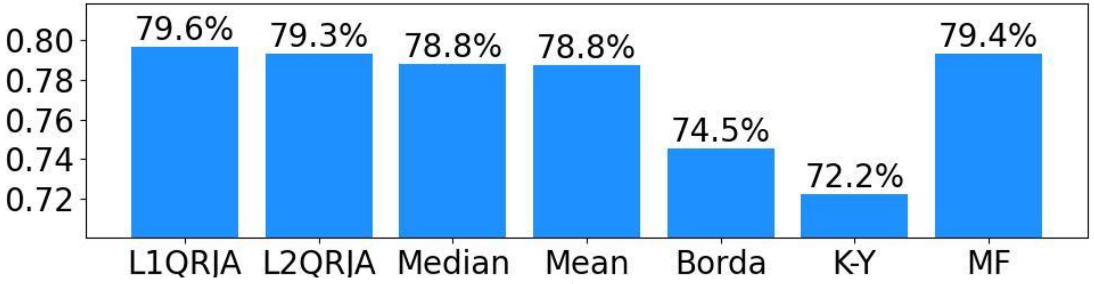

Marathon

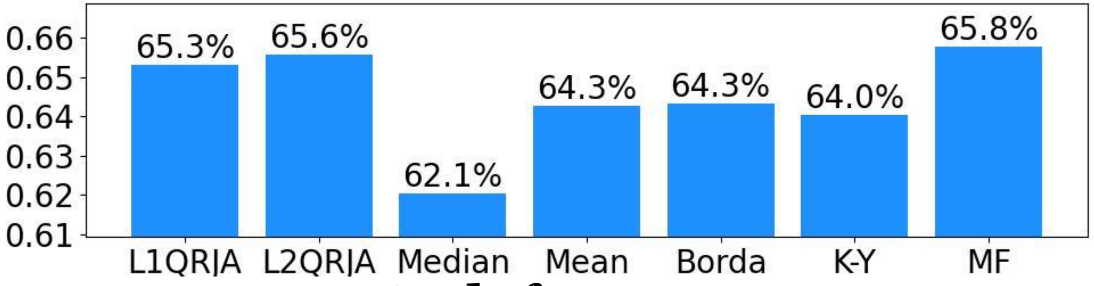

Codeforces

Our Contributions

- **We propose and study the QRJA problem**
  - Conceptually, this bridges social choice and learning-to-rank

- **We thoroughly study a subclass, ℓ_p QRJA**
  - Theoretically, we provide a tight characterization of its complexity
  - Empirically, we conduct experiments to demonstrate its effectiveness

Acknowledgements: *This work was supported by NSF IIS-1814056, the Center for Emerging Risk Research, the Cooperative AI Foundation and NSF Award CCF-2307106. Funding to attend this conference was provided by the CMU GSA Provost Conference Funding.*