

Distortion in Social Choice Problems: An Annotated Reading List

ELLIOT ANSHELEVICH

Rensselaer Polytechnic Institute

and

ARIS FILOS-RATSIKAS

University of Liverpool

and

NISARG SHAH

University of Toronto

and

ALEXANDROS A. VOUDOURIS

University of Essex

The notion of *distortion* in social choice problems has been defined to measure the loss in efficiency—typically measured by the utilitarian social welfare, the sum of utilities of the participating agents—due to having access only to limited information about the preferences of the agents. Here, we provide a comprehensive reading list on the related literature.

General Terms: Algorithms

Additional Key Words and Phrases: Social choice; Distortion

1. INTRODUCTION

Social choice theory is concerned with aggregating the preferences of individuals into a desirable collective decision, and has many applications such as choosing an electoral candidate, a public policy, the recipient of an award, or something as simple as the most appropriate time for a meeting. These preferences are typically assumed to be captured by *utility functions*, which assign numerical values to the different options, indicating the intensity by which an individual prefers one possible outcome to another. While the existence of such a *cardinal* utility structure is rarely disputed, the predominant approach in social choice theory is to elicit more limited preference information from the participants; in particular, they are usually required to provide *ordinal* preference rankings over the different outcomes. This is primarily due to cognitive reasons, as it is much more conceivable to come up with a ranking based on comparisons rather than a numerical utility structure.

The inevitable loss of information due to the restricted expressiveness of the elicited preferences makes it rather challenging to optimize objectives of a cardinal nature. A natural such objective is the maximization of the (*utilitarian*) *social welfare*, defined as the sum of the individual utilities for the chosen outcome. *How can*

Authors' addresses: eanshel@cs.rpi.edu, aris.filos-ratsikas@liverpool.ac.uk,
nisarg@cs.toronto.edu, alexandros.voudouris@essex.ac.uk

we make the right choice when we do not have access to the utilities themselves? Driven by the principles of worst-case analysis and approximation algorithms, the notion of *distortion* has been defined to measure the worst-case deterioration of an aggregate cardinal objective, such as the utilitarian social welfare, due to having access to limited preference information, particularly ordinal rankings. The outcome minimizing the distortion then represents the best possible choice given the available information.

Over the past 15 years, distortion has been studied extensively in the context of many fundamental social choice problems, such as *single-winner and multi-winner elections*, *participatory budgeting*, and *matching*, thus giving rise to a rich and vibrant literature at the intersection of economics, computation, and artificial intelligence. The following articles provide a helpful introduction to this literature.

2. READING LIST

- * Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. Distortion in social choice problems: The first 15 years and beyond. arXiv:2103.00911, 2021.

This is a short survey on this topic.

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

This is one of the first papers on this topic. It introduces and motivates the basic concept of distortion. This paper studies the distortion of social choice problems when agent utilities are normalized, and bounds the quality of solutions which can be computed given only the ordinal preferences of the agents.

- * Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.

This paper continues the line of inquiry set by the previous paper by considering problems where a committee of candidates of a given size must be chosen. It provides bounds on the optimal distortion as a function of the committee size, and studies additive regret as an alternative measure of efficiency loss.

- * Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, Piotr Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.

This paper considers distortion in social choice under the metric setting, where voters and candidates are points in an arbitrary metric space, and a voter’s cost for a candidate being selected corresponds to the distance between the voter and the candidate. The paper motivates this setting, and provides a variety of lower and upper bounds on the distortion of many classic voting rules.

- * Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. Resolving the optimal met-

ric distortion conjecture. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1427–1438, 2020.

This paper resolves a central conjecture in social choice under the metric setting. It introduces a new, computationally efficient voting rule for aggregating voters’ ranked preferences to select a candidate, and proves that it is optimal among deterministic rules in terms of distortion.

- * Elliot Anshelevich and Wennan Zhu. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. *ACM Transactions on Economics and Computation*, 9.2:9, 2021.

This paper considers social choice in the metric setting when the candidate locations in the metric space are known, but for the voters, only their ordinal preferences over the candidates are known. It provides stronger bounds on distortion using this information, and generalizes the results to many other graph problems, including matching and facility location.

- * Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. *Artificial Intelligence*, 296:103488, 2021.

This paper considers the tradeoff between distortion and elicitation, where besides the ordinal preferences, the voters can also provide limited cardinal information by means of answering value or comparison queries. It provides upper and lower bounds on the possible tradeoffs between the number of queries per agent and distortion, and shows that a relatively small number of queries suffices to achieve desirable distortion bounds.

- * Debmalya Mandal, Nisarg Shah, and David P. Woodruff. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 795–813, 2020.

This paper and its predecessor also consider the aforementioned tradeoff between distortion and elicitation, but allow arbitrary queries to the voters and use the number of bits of information elicited as a measure of cognitive burden. By uncovering novel connections to the literature on streaming algorithms and communication complexity, they identify the optimal tradeoff between distortion and elicitation for single-winner and multi-winner elections.