



Fairness in Preference Queries: Social Choice Theories Meet Data Management

Senjuti Basu Roy
NJIT
Newark, NJ, USA
senjutib@njit.edu

Baruch Schieber
NJIT
Newark, NJ, USA
sbar@njit.edu

Nimrod Talmon
Ben-Gurion University
Beer-Sheva, Israel
talmonn@bgu.ac.il

ABSTRACT

Given a large number (notationally m) of users' (members or voters) preferences as inputs over a large number of items or candidates (notationally n), preference queries leverage different preference aggregation methods to aggregate individual preferences in a systematic manner and come up with a single output (either a complete order or top- k , ordered or unordered) that is most representative of the users' preferences. The goal of this **1.5 hour lecture style tutorial** is to adapt different preference aggregation methods from social choice theories, summarize how existing research has handled fairness over these methods, identify their limitations, and outline new research directions.

PVLDB Reference Format:

Senjuti Basu Roy, Baruch Schieber, and Nimrod Talmon. Fairness in Preference Queries: Social Choice Theories Meet Data Management. PVLDB, 17(12): 4225 - 4228, 2024.

doi:10.14778/3685800.3685841

1 OVERVIEW & JUSTIFICATION

Preference queries [13, 18, 29] are prevalent in high fidelity data management tasks, including, search, ranking, and recommendations [2, 3, 11, 23, 25], for applications such as selecting a handful of candidates in domains where resource is scarce (such as hiring and admission) and electoral voting systems, to name a few. Preference queries leverage different preference aggregation methods to aggregate individual preferences in a systematic manner and come up with a single output (either a complete order or top- k , ordered or unordered) that is most representative of the users' preferences. The need to support preference queries has not gone unnoticed by the data management community, and a number of general frameworks emerged [6, 18]. Two related aspects are studied: (i) semantic clarity and adequacy, much of which are adapted from social choice theory, and (ii) computational efficiency. Algorithmic fairness has been receiving increasing attention [12, 26, 27, 33, 39, 48] in ranking and recommendation mostly focusing on how to change a single output (ranked or top- k) and make it fair. In contrast, fairness in preference queries (which involves multiple input preferences) [28, 30, 34, 35, 45] remains less explored in the literature.

The tutorial is likely to bring interdisciplinary perspectives from three different research communities - it will systematically identify

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 17, No. 12 ISSN 2150-8097.
doi:10.14778/3685800.3685841

fairness opportunities considering a wide variety of preference aggregation methods adapted from social choice theory, investigate their data management and computation implications.

The tutorial will be presented considering preference queries supporting four interspersed dimensions, as described below.

Preference Elicitation Models. Study difference preference elicitation processes that we broadly categorize as rank based and non rank based and their applicability to different applications. In rank based processes, the users can provide a fully ranked order over all items, a partial order, or a coarser preference (like item a ranked higher than item b , etc). In non rank based preferences, users can provide only likes, both likes and dislikes, or even an ordinal preference (likes item a as "excellent", b as "good", etc). Rank based ones are suitable in hiring/admission/electoral system, while non rank based ones are more relevant in obtaining user feedback from search results, user satisfaction survey, product reviews, etc.

Preference Aggregation Methods. The tutorial will identify appropriate preference aggregation methods that are most commensurate to the underlying preference elicitation process and underlying application. For example, when user preferences are given as ranked order, depending on the underlying application, we will aggregate them using existing single-round rank based methods (e.g., Kemeny, Spearman's footrule, or Borda), or multi-round based methods (STV, IRV). The former aggregation methods are suitable in hiring decision, whereas, the latter ones are gaining popularity in voting systems. On the other hand, when users provide non rank based preferences, we will study how Jaccard similarity or Hamming distances are suitable to aggregate them.

Produced Output Form. From the application point of view, the produced output may require an order over all n items (hiring/admission), or a small number k of n items as outputs. In case of top- k items requirement, the returned k -items may need to be ordered for certain applications (top- k web pages returned by the search engine), or in some cases it is fine to return them as a set (selecting a set of representatives or body to form certain committee).

Make original outcome fair. How to quantify the minimum effort needed to make the outcome fair. There are two ways to make the outcome fair: A. change inputs. B. change outputs. We will discuss both of these options and their fairness and computational implications.

2 TARGET AUDIENCE AND ASSUMED BACKGROUND

The tutorial will be of interest to both theoreticians and practitioners who are interested in the development of fair data-centric applications in the areas of databases, data mining, machine learning,

social science, and algorithms, ranging from large-scale analytics to emerging online applications. Tutorial attendees are expected to have basic knowledge in algorithms, and data management. Knowledge in constrained optimization is not necessary.

3 RELATED RECENT TUTORIALS & OVERLAP

The closest to our proposal is the tutorial [41] presented in SIGMOD 2023. It primarily focuses on classification framework for fairness-enhancing interventions, single score-based ranking, and supervised learning-to-rank. We, on the other hand, will focus on preference queries that take multiple input preferences of different forms. The overlap between our proposal and this tutorial is likely to be minimal.

4 SCOPE AND STRUCTURE

The tutorial will be divided in three parts.

4.1 Preference aggregation methods (30 minutes)

The basic model in computational social choice [15] is voting, in which – first, preferences are elicited from an agent community; and then, those preferences are aggregated. Formally,

Definition 4.1. An *election* is a tuple $E = (C, V)$ where C is a set of m *candidates* and V is a set of n voters; each voter $v \in V$ submits a *ballot*, denoted also by v . A *voting rule* (i.e., aggregation method) is a function that takes an election E as its input and outputs a *winner* of the election.

Input Formats - Preference Elicitation. The standard, most popular preference elicitation formats are:

- Approval ballots [14] – voters specify approve/disapprove for each alternative (so $v \in 2^C$);
- Ordinal ballots [5] – voters specify a linear order over the alternatives;
- Scoring ballots – voters specify a score for each alternatives (consider also the related cumulative ballot format [40]).

Output Formats - Type of Winner. As for the outputs, the most prominent output formats are:

- Single-winner elections [15, 20] – the winner is a single alternative (some $c \in C$; this corresponds to, e.g., selecting a president).
- Multiwinner elections [21] – the winner is a set of alternatives (some $c \subseteq C$ of some size; e.g., selecting a committee or a parliament).
- Participatory budgeting [7] – the winner is a set of alternatives but each alternative has a cost and the total cost of the selected alternatives cannot go beyond some given budget (this is usually done in municipal settings [42]).

Getting from Inputs to Outputs – Preference Aggregation.

Many voting rules (i.e., aggregation methods) have been devised and analyzed [15, 20]. As the most well-studied social choice setting is that of ordinal-based single-winner elections we mention some of the prominent ones for this setting (recall that, here, preferences are elicited as linear orders – ranking – over the alternatives – and the output is a single candidate. In the tutorial we will also cover cases where the output is a set of candidates, either ordered or unordered.):

- Plurality – the winner is the candidate ranked first the most times.
- Borda – each voter gives a score of $m-i$ to a candidate it ranks in the i th position and the candidate with the highest score wins.
- Copeland [36] – we create a directed graph with a vertex for each candidate and a directed edge from u to v if there are more voters ranking u before v ; and, then, the candidate with the highest out-degree wins.
- STV [43] – in each iteration, if there is a candidate ranked first by a majority, then it is selected as the winner; otherwise, a candidate appearing first the least number of times is eliminated, and the rankings are updated accordingly.
- Kemeny [1] – a Kemeny consensus ranking is the ranking that minimizes the sum of Kendall-Tau distances to all the input rankings and the winner is the candidate ranked in the first position of the Kemeny consensus ranking.

Analyzing Aggregation Methods. How to choose which aggregation method to use? The most popular analysis approaches are:

- An axiomatic approach – here, desired properties of aggregation methods are formally defined and methods are analyzed to whether they satisfy them. Classical results here includes May's theorem [32] (essentially showing that simple majority is the only reasonable voting rule when $|C| = 2$), Black's median theorem (essentially showing that taking the median is the only reasonable voting rule when C is an Euclidean line), Arrow's theorem [4] (essentially an impossibility result for the existence of an aggregation method that has no dictatorial features) and the theorems of Gibbard and Satterthwaite [22, 38] (essentially an impossibility result for the existence of an aggregation method that is resilient to strategic voting).
- A computational approach – here, computational features of aggregation methods are analyzed (showing, e.g., the fact that some aggregation methods are NP-hard to compute [16, 24, 46]).
- A simulation-based approach – here, aggregation methods are simulated and their results are being statically-analyzed and visualized [19].

4.2 Fairness in answering preference queries - existing research (30 minutes)

As demonstrated in the previous section one of the goals in devising and analyzing a preference aggregation method is to make sure that it is faithfully and fairly representing the opinion of the voters. In some applications the outcome of the preference queries needs also to ensure a *fair* representation of the *candidates*. For example, if the preference query is a ranked order of job applicants we may need to make sure that gender and race are appropriately represented in the ranked outcome.

Fairness constraints may be imposed in case the output of the query is a ranking (either full or partial ranking), but also in case of perpetual voting of a single candidate. Consider the case where viewers are polled to select the “movie of the day”, everyday. It makes sense to impose some fairness constraints on the selections to avoid the case that all the chosen movies are of the same genre whose number of followers is a majority.

We model fairness by *protected attributes*. Each item/candidate $c \in C$ is associated a set protected attributes A_C , where each protected attribute $A_C(i)$ can take any of ℓ_i different values. As an example, seniority level is a multi-valued protected attribute with three possible values Junior, Mid career, Senior, while gender is commonly a binary protected attribute with two values male and female.

Ensuring fairness. Recent work considered several ways to ensure fairness. Celis et al. [17] introduce a top- k fairness measure that ensures a given upper and lower bound of the representation of each of the protected attribute values in the top- k , for a fixed value of k . Zehlike et al. [47] extend group fairness using the standard notion of protected groups and ensure that the proportion of protected candidates in *every* top- k ranking remains statistically above a given minimum (while not ensuring any upper bound).

A more general way to ensure fairness of a ranking of size n is *proportionate fairness* that was introduced in the context of fair rankings by Wei et al. [45]. For any protected attribute value p , let $f(p)$ denote the fraction of items with this value. A ranking is proportionate fair or p -fair if for every $k \in [1..n]$, the number of items with protected attribute value p among the k top ranked items is either $\lfloor f(p) \cdot k \rfloor$ or $\lceil f(p) \cdot k \rceil$. (A relaxed p -fairness can be defined by introducing an integer tolerance to the constraint on the number of items with protected attribute value p among the k top ranked items.) P -fairness was introduced in the well known Chairman Assignment problem [44] that studies how to select a chairman for a union of states such that at any time the accumulated number of chairmen from each state is proportional to its weight. This notion has been studied in the context of resource allocation and scheduling [9, 10].

Multi protected attributes. In some applications fairness needs to be ensured for multiple protected attributes, such as gender, race, and income. In this case it makes a difference whether the protected attributes are independent or not. In case of independence multi protected attributes can be converted to a single attribute whose set of values is the Cartesian product of the original protected attributes, and the respective proportions are given by multiplying the proportions of the original values. An example to protected attributes that may be assumed to be independent are gender and race as in most applications the representation of race within each gender needs to be proportionate. In some cases it may not make sense to assume that the protected attributes are independent. For example, consider income and race. In this case ensuring fairness over multi attributes is more difficult computationally. As a matter of fact it is shown in [28] that for three or more dependent protected attribute even determining whether there exists a fair output is strong NP-Hard.

Producing a fair outcome. The preference aggregation methods introduced earlier may not necessarily produce a fair outcome. Thus, this outcome needs to be modified to obtain a fair outcome. Certainly, the goal is to minimize the modification in order to maintain the faithful and fair representation of the opinion of the voters. There are two approaches in making such modification; (1) modify the output of the preference aggregation method to produce a fair output, and (2) modify the input preferences so that the aggregated preference is guaranteed to be fair. In both cases the goal is to minimize the modification, and thus in the first approach it a metric

space needs to be defined over the set of possible preferences, and the goal is to find the fair preference that is closest to the aggregated preference. For example, in [45] the preferences are full rankings and the distance between two rankings (permutations) is Kendall Tau distance. In the second approach the distance may be measured by the number of votes modified/added/removed. This closely relates to the Margin of Victory [8, 31, 37] problem defined as the minimum number of vote changes needed to change the outcome of an election.

Both approaches introduce interesting problems and will be covered in the tutorial.

4.3 Future research directions (30 minutes)

We will focus on the three major aspects.

New preference aggregation methods. First, the design and analysis of new preference aggregation methods using both the axiomatic approach and the computational approach, specifically, for the cases of perpetual voting and participatory budget will be investigated in this discussion. We will also provide a roadmap that could be helpful to the practitioners in identifying the appropriate aggregation methods considering the application at hand.

Alternative models to enable fair outcome. Enabling fairness on preference queries could be studied as a bi-criteria optimization problem, that is, for a given pair $(\alpha > 1, \beta > 1)$ and a set of m input preferences make the outcome fair, such that the distance between the original output (ranked or top- k) and the produced output is at most α and its distance from a fair ranking is at most β , if such an output exists. These opportunities will be investigated in this section. Other than demographic parity [33], p -fairness [45], and top- k statistical parity [30], we will also explore what other fairness measures are applicable to preference queries.

Efficient solution design. Ensuring a fair outcome while maintaining adequate representation of the voters also poses several problems, mainly computational. These problems are closely related to the Margin of Victory problems that are known to be computationally hard. Finding an approximate solution to such problems is a challenge.

5 BIOGRAPHY OF THE PRESENTERS

Senjuti Basu Roy is the Panasonic Chair in Sustainability and an Associate Professor in the Department of Computer Science at the New Jersey Institute of Technology. Her research focus lies at the intersection of data management, data exploration, and AI, especially enabling human-machine analytics in scale. Senjuti has published more than 85 research papers in high impact data management and data mining conferences and journals. She has served as the tutorial co-chair of VLDB 2023.

Baruch Schieber is a Professor in the Department of Computer Science, NJ Institute of Technology. Before joining NJIT Baruch was a Distinguished Research Staff Member in IBM Research. His research interests are in theoretical computer science, including optimization under uncertainty, algorithms, mathematical programming, and high-performance computing. Baruch received his PhD in Computer Science from Tel Aviv University in 1987. He published more than 150 papers in scientific journals and conferences.

Nimrod Talmon is an Assistant Professor within the Industrial Engineering and Management Department, Ben-Gurion University, Israel. Before joining Ben-Gurion University, he was a Postdoctoral Fellow with the Computer Science and Applied Mathematics Department, Weizmann Institute of Science, Israel and received his Ph.D. degree in computer science from TU Berlin. His research interests include artificial intelligence, game theory, computational social choice, social networks, and combinatorial optimization. He has Erdos number 3, Sabbath number 7, and Bacon number 6.

ACKNOWLEDGMENTS

The work of Senjuti Basu Roy is supported by the following funding agencies: (1) National Science Foundation award number(s): 1942913, 2007935 (2) Office of Naval Research award number(s): N000142112966, N000142412466.

REFERENCES

- [1] Alnur Ali and Marina Meilă. 2012. Experiments with Kemeny ranking: What works when? *Mathematical Social Sciences* 64, 1 (2012), 28–40.
- [2] Sihem Amer-Yahia, Shady Elbassuoni, Behrooz Omidvar-Tehrani, Ria Mae Borromeo, and Mehrdad Farokhnejad. 2019. Grouptravel: Customizing travel packages for groups. In *22nd International Conference on Extending Database Technology (EDBT)*.
- [3] Sihem Amer-Yahia, Senjuti Basu Roy, Ashish Chawlat, Gautam Das, and Cong Yu. 2009. Group recommendation: Semantics and efficiency. *Proceedings of the VLDB Endowment* 2, 1 (2009), 754–765.
- [4] Kenneth J Arrow. 1950. A difficulty in the concept of social welfare. *Journal of political economy* 58, 4 (1950), 328–346.
- [5] Kenneth J Arrow, Amartya Sen, and Kotaro Suzumura. 2010. *Handbook of social choice and welfare*. Elsevier.
- [6] Anastasios Arvanitis and Georgia Koutrika. 2012. Towards preference-aware relational databases. In *2012 IEEE 28th International Conference on Data Engineering*. IEEE, 426–437.
- [7] Haris Aziz and Nisarg Shah. 2021. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science*. Springer, 215–236.
- [8] John Bartholdi and James Orlin. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8 (01 1991), 341–354. <https://doi.org/10.1007/BF00183045>
- [9] Sanjoy K Baruah, Neil K Cohen, C Greg Plaxton, and Donald A Varvel. 1996. Proportionate progress: A notion of fairness in resource allocation. *Algorithmica* 15, 6 (1996), 600–625.
- [10] Sanjoy K Baruah, Johannes E Gehrke, C Greg Plaxton, Ion Stoica, Hussein Abdell-Wahab, and Kevin Jeffay. 1997. Fair on-line scheduling of a dynamic set of tasks on a single resource. *Inform. Process. Lett.* 64, 1 (1997), 43–51.
- [11] Senjuti Basu Roy, Laks VS Lakshmanan, and Rui Liu. 2015. From group recommendations to group formation. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. 1603–1616.
- [12] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wu, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H Chi, et al. 2019. Fairness in recommendation ranking through pairwise comparisons. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2212–2220.
- [13] Ronen Brafman and Carmel Domshlak. 2004. *Database preference queries revisited*. Technical Report. Cornell University.
- [14] Steven Brams and Peter C. Fishburn. 2007. Approval voting.
- [15] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.
- [16] Ioannis Caragiannis, Jason A Covey, Michal Feldman, Christopher M Homan, Christos Kaklamanis, Nikos Karanikolas, Ariel D Procaccia, and Jeffrey S Rosenschein. 2012. On the approximability of Dodgson and Young elections. *Artificial Intelligence* 187 (2012), 31–51.
- [17] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. 2017. Ranking with fairness constraints. *arXiv preprint arXiv:1704.06840* (2017).
- [18] Jan Chomicki. 2003. Preference formulas in relational queries. *ACM Transactions on Database Systems (TODS)* 28, 4 (2003), 427–466.
- [19] Edith Elkind et al. 2017. What do multiwinner voting rules do? An experiment over the two-dimensional Euclidean domain. In *AAAI '17*. 494–501.
- [20] Ulle Endriss. 2017. *Trends in computational social choice*. AI Access.
- [21] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. In *Trends in computational social choice*, Ulle Endriss (Ed.). AI Access, 27–47.
- [22] Allan Gibbard. 1973. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society* (1973), 587–601.
- [23] Lei Guo, Hongzhi Yin, Qinyong Wang, Bin Cui, Zi Huang, and Lizhen Cui. 2020. Group recommendation with latent voting mechanism. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 121–132.
- [24] Edith Hemaspaandra, Holger Spakowski, and Jörg Vogel. 2005. The complexity of Kemeny elections. *Theoretical Computer Science* 349, 3 (2005), 382–391.
- [25] Zhenhua Huang et al. 2020. Social group recommendation with TrAdaBoost. *TCSS* (2020).
- [26] Md Mouinul Islam, Mahsa Asadi, and Senjuti Basu Roy. 2023. Equitable Top-k Results for Long Tail Data. *Proceedings of the ACM on Management of Data* 1, 4 (2023), 1–24.
- [27] Md Mouinul Islam, Soroush Vahidi, Baruch Schieber, and Senjuti Basu Roy. 2024. Promoting Fairness and Priority in k -Winners Selection Using IRV. In *KDD*. ACM. <https://doi.org/10.1145/3637528.3671735>
- [28] Md. Mouinul Islam, Dong Wei, Baruch Schieber, and Senjuti Basu Roy. 2022. Satisfying complex top- k fairness constraints by preference substitutions. *Proc. VLDB Endow.* 16, 2 (oct 2022), 317–329.
- [29] Georgia Koutrika and Yannis Ioannidis. 2005. Personalized queries under a generalized preference model. In *21st International Conference on Data Engineering (ICDE '05)*. IEEE, 841–852.
- [30] Caitlin Kuhlman and Elke Rundensteiner. 2020. Rank aggregation algorithms for fair consensus. *Proceedings of the VLDB Endowment* 13, 12 (2020).
- [31] Thomas R. Magrino, Ronald L. Rivest, Emily Shen, and David Wagner. 2011. Computing the margin of victory in IRV elections. In *Proceedings of the 2011 Conference on Electronic Voting Technology/Workshop on Trustworthy Elections* (San Francisco, CA) (EV/TWOTE '11). USENIX Association, USA, 4.
- [32] Kenneth O May. 1952. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society* (1952), 680–684.
- [33] Evangelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. 2022. Fairness in rankings and recommendations: an overview. *The VLDB Journal* (2022), 1–28.
- [34] Senjuti Basu Roy. 2022. Returning top- k : Preference aggregation or sortition, or is there a better middle ground. *SIGMOD Blog* (2022).
- [35] Senjuti Basu Roy. 2024. Fairness and Robustness in Answering Preference Queries. *Data Engineering* (2024), 36.
- [36] Donald G. Saari. 2006. Which is better: the Condorcet or Borda winner? *Social Choice and Welfare* 26, 1 (2006), 107.
- [37] Anand Sarwate, Stephen Checkoway, and Hovav Shacham. 2013. Risk-Limiting Audits and the Margin of Victory in Nonplurality Elections. *Statistics, Politics, and Policy* 4, 1 (Jan. 2013), 29–64.
- [38] Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* 10, 2 (1975), 187–217.
- [39] Ashudeep Singh, David Kempe, and Thorsten Joachims. 2021. Fairness in ranking under uncertainty. *Advances in Neural Information Processing Systems* 34 (2021), 11896–11908.
- [40] Piotr Skowron, Arkadii Slinko, Stanisław Szufa, and Nimrod Talmon. 2020. Participatory budgeting with cumulative votes. *arXiv preprint arXiv:2009.02690* (2020).
- [41] Julia Stoyanovich, Meike Zehlike, and Ke Yang. 2023. Fairness in Ranking: From Values to Technical Choices and Back. In *Companion of the 2023 International Conference on Management of Data*. 7–12.
- [42] Nimrod Talmon and Piotr Faliszewski. 2019. A framework for approval-based budgeting methods. In *AAAI '17*. 2181–2188.
- [43] Nicolaus Tideman. 1995. The single transferable vote. *Journal of Economic Perspectives* 9, 1 (1995), 27–38.
- [44] Robert Tijdeman. 1980. The chairman assignment problem. *Discrete Mathematics* 32, 3 (1980), 323–330.
- [45] Dong Wei, Md Mouinul Islam, Baruch Schieber, and Senjuti Basu Roy. 2022. Rank Aggregation with Proportionate Fairness. In *Proceedings of the 2022 International Conference on Management of Data* (Philadelphia, PA, USA) (SIGMOD '22). Association for Computing Machinery, New York, NY, USA, 262–275.
- [46] Peyton Young. 1995. Optimal voting rules. *Journal of Economic Perspectives* 9, 1 (1995), 51–64.
- [47] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. FA*IR: A fair top- k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1569–1578.
- [48] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2021. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000* (2021).