4. What does it mean if we cannot achieve optimal fairness on the given dataset? What are the implications of the solution space?

If the solution space shows that with high probability, the solution set with optimal fairness does not exist, then one may need to consider the bias in the data collection and retrieval algorithms. From the perspective of data, if the data is highly unbalanced for different groups, or is too biased according to other fairness definitions, then no matter what algorithms we use, we are unlikely to surface enough items from the minority for a fair exposure. From the perspective of system design and algorithms, the core retrieval framework depends on many basic system components such as data collection, pre-processing, indexing, searching, ranking, and personalization. As a result, any steps that fail to consider the minority groups (e.g., dialects compared to the standard language) and subsequently fail to capture enough representations from the minority, will lead to the biased search space, which may contribute to the low possibility of achieving optimal fairness.

6 CONCLUSION

In this paper, we proposed a novel perspective of analyzing the fairness problems in IR. We presented a framework that first depicts the solution space on any given dataset by estimating theoretical boundaries and optimal solution values, and then we utilized the solution space to facilitate a variety of analysis and decision making which would otherwise be considerably time and resource consuming. This framework has the advantage of being simple to deploy, explain, is reliable with theoretical guarantees, and it is easy to generalize to account for various applications. Researchers and system designers can plug into this framework customized utility functions and fairness constraints, and apply to data of different distributions. We demonstrated the application of our framework with synthetic and real world data. We hope that through our exploration of examples of research questions that can be answered with this framework, we can inspire a broad range of analysis that could potentially benefit from this framework.

The theoretical results presented in this paper do call for some diligence while putting them in practice. The problem setting requires some assumptions that may not always hold. For example, the independence assumption between multi-parties or between relevance and fairness may not be valid for some real world situations. To address this, one can first perform the correlation analysis between different components and then depict the solution boundary by analyzing the characteristics of the joint distribution. We emphasize the contribution of this framework concept and leave more inclusive theoretical analysis such as multinomial distributions and multi-dimension optimizations for future work.

REFERENCES

[1] Mahmoudreza Babaei, Abhijnan Chakraborty, Juhi Kulshrestha, Elissa M Redmiles, Meeyoung Cha, and Krishna P Gummadi. 2018. Analysing Biases in Perception of Truth in News Stories and their Implications for Fact Checking. In *Proceedings of FAT'18*.

[2] Ricardo Baeza-Yates. 2018. Bias on the Web. *Commun. ACM* 61, 6 (2018), 54–61.

[3] Judit Bar-Ilan, Kevin Keenoy, Mark Levene, and Eti Yaari. 2009. Presentation bias is significant in determining user preference for search results – A user study. *JASIST* 60, 1 (2009), 135–149.

[4] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *California Law Review* 104 (2016), 671.

[5] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *Proceedings of SIGIR'18*.

[6] Ivar Bråten, Helge I Strømsø, and Ladislao Salmerón. 2011. Trust and mistrust when students read multiple information sources about climate change. *Learning and Instruction* 21, 2 (2011), 180–192.

[7] L Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth K Vishnoi. 2018. An Algorithmic Framework to Control Bias in Bandit-based Personalization. *arXiv preprint arXiv:1802.08674* (2018).

[8] L. Elisa Celis, Damian Straszak, and Nisheeth K. Vishnoi. 2018. Ranking with Fairness Constraints. In *ICALP'17*.

[9] Le Chen, Ruijun Ma, Anikó Hannák, and Christo Wilson. 2018. Investigating the Impact of Gender on Rank in Resume Search Engines. In *Proceedings of CHI'18*.

[10] J. Shane Culpepper, Fernando Diaz, and Mark D. Smucker. 2018. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne. *SIGIR Forum* 52, 1 (2018), 34–90.

[11] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science advances* 4, 1 (2018).

[12] Robert Epstein and Ronald E Robertson. 2015. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *PNAS* 112, 33 (2015), E4512–E4521.

[13] James Grimmelmann. 2011. Some skepticism about search neutrality. *The Next Digital Decade: Essays on the Future of the Internet* (Jan 2011).

[14] Alexander Haas and Julian Unkel. 2017. Ranking versus reputation: perception and effects of search result credibility. *Behaviour & Information Technology* 36, 12 (2017), 1285–1298.

[15] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. 2016. Fairness in learning: Classic and contextual bandits. In *Proceedings of NeurIPS'16*.

[16] Yvonne Kammerer and Peter Gerjets. 2012. Chapter 10 How Search Engine Users Evaluate and Select Web Search Results: The Impact of the Search Engine Interface on Credibility Assessments. In *Web search engine research*. Emerald Group Publishing Limited, 251–279.

[17] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of CHI'15*.

[18] Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. 2017. Meritocratic fairness for cross-population selection. In *Proceedings of ICML'17*.

[19] Juhi Kulshrestha, Motahhare Eslami, Johnnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In *Proceedings of CSCW'17*.

[20] Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. 2017. Auditing Search Engines for Differential Satisfaction Across Demographics. In *Proceedings of WWW'17*.

[21] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *Proceedings of CIKM'18*.

[22] Arvind Narayanan. 2018. 21 fairness definitions and their politics. Retrieved April 26, 2019, from https://www.youtube.com/watch?v=jIXIuYdnyyk.

[23] Alamir Novin and Eric Meyers. 2017. Making sense of conflicting science information: Exploring bias in the search engine result page. In *Proceedings of CHIIR'17*.

[24] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. 2016. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *SSRN Electron. J.* (2016), 1–47.

[25] Jahna Otterbacher, Jo Bates, and Paul Clough. 2017. Competent Men and Warm Women: Gender Stereotypes and Backlash in Image Search Results. In *Proceedings of CHI'17*.

[26] Jahna Otterbacher, Alessandro Checco, Gianluca Demartini, and Paul Clough. 2018. Investigating user perception of gender bias in image search: the role of sexism. In *Proceedings of SIGIR'18*.

[27] Milad Shokouhi, Ryen White, and Emine Yilmaz. 2015. Anchoring and adjustment in relevance estimation. In *Proceedings of SIGIR'15*.

[28] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of SIGKDD'18*.

[29] Jackie Snow. 2018. Bias already exists in search engine results, and it's only going to get worse. https://www.technologyreview.com/s/610275/meet-the-woman-who-searches-out-search-engines-bias-against-women-and-minorities/.

[30] Herman Tavani. 2016. Search Engines and Ethics. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

[31] 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 CIKM'17*.

[32] Yi Zhang. 2005. Bayesian Graphical Models for Adaptive Information Filtering. https://users.soe.ucsc.edu/~yiz/papers/data/YOWStudy/