

Chapter 8

Conclusion



Abstract Our book brings together the multidisciplinary insights, methods, and empirical findings related to bounded rationality and human biases in decision-making and presents a behavioral economics research agenda under which a series of specific research questions, new directions, and methodological challenges can be further investigated by students and researchers in future IR studies. In this final chapter, we summarize the contents of previous chapters and discuss the contributions, practical implications, and related new directions under our behavioral economics research approach to IR problems. We hope that this book can serve as a useful starting point for studying bias-aware IR and motivate students and researchers from diverse backgrounds to further explore and advance the science and technology on supporting boundedly rational people interacting with information.

Understanding how people behave and why they behave in such ways is a central topic to information seeking and retrieval research. The knowledge learned about users' search behavioral patterns, strategies of search result judgments, and evaluation of system performances is essential for not only predicting users' in situ search actions and feedback but also developing built-in formal user models for retrieval and ranking algorithms, search recommendation techniques, as well as scalable evaluation metrics. In contrast to the rational assumptions of existing formal models, people tend to be boundedly rational and are affected by a series of human biases and heuristics when making decisions under uncertainty (Kahneman, 2003; Simon, 1955). Behavioral economics researchers have explored and empirically tested a broad range of human biases and factors that contribute to bounded rationality in a variety of real-life and simulated simple decision-making scenarios (Kahneman, 2003; Thaler, 2016; Weber & Camerer, 2006). Although the operation of System 1 and the adoption of mental shortcuts enable individuals to simplify the decision-making process and make quick judgments without processing a large amount of new information (Gigerenzer & Brighton, 2009), the decisions and associated outcomes tend to be affected by higher error rates and deviate from the optimal results predicted by rational models. Many of these systematic deviations have been ignored or abstracted out from formal user models and simulation-based experiments, which

largely restricts the actual contributions from the advances in IR algorithms and systems to understanding and supporting real-world users engaging in information search interactions.

To address the above research gap, our book brings together the multidisciplinary insights, methods, and empirical findings related to bounded rationality and human biases in decision-making and presents a behavioral economics research agenda under which a series of specific research questions, new directions, and methodological challenges can be further investigated by students and researchers in future IR studies. Specifically, Chap. 1 offers an overview of the theoretical basis and involved disciplines related to the problem of bias-aware IR and clarifies the structure of this book. Chapter 2 thoroughly reviews the basic structures and recent advances in a series of mainstream formal user models applied in various sub-areas of IR (e.g., click modeling, simulation of search sessions, offline evaluation experiments) and highlights their contributions in modeling search behaviors and limitations with respect to accommodating biased human decisions. Based on this review on formal models, Chap. 3 briefly introduces the gaps between simulated rational agents or assumptions and empirically confirmed human biases that frequently appear in real-world decision-making activities.

To further enrich our discussions on the identified gaps and bring in the relevant insights from behavioral economics, Chap. 4 goes beyond rational agents and presents a comprehensive overview of behavioral experiments and findings on the human biases and heuristics emphasized in Chap. 3. Built upon the identified gaps and knowledge regarding both formal modeling and human bounded rationality, Chap. 5 revisits the rational assumptions underpinning user models and evaluation metrics and proposes reasonable approaches to revising and extending the rational oversimplified assumptions and offering the assumptions a more solid behavioral and psychological basis.

Built upon the knowledge synthesized in previous chapters, Chap. 6 moves forward by introducing the progress we as a research community have made on understanding human biases and bounded rationality in IR and related fields (e.g., Azzopardi, 2021; Liu & Han, 2020), including information seeking and recommendations (e.g., Agosto, 2002; Ge et al., 2020; Xu et al., 2020). Based on the identified research gaps and existing findings, Chap. 7 presents a full behavioral economics research agenda that addresses three different aspects of bias-aware IR, including characterizing bounded rationality, building search systems, and developing bias-aware search evaluation. In particular, we highlight the importance of going beyond traditional evaluation metrics focusing on relevance-based search effectiveness and discuss a new vision named BITS system, which can proactively address the negative impacts from both algorithmic biases and human biases and offer unbiased support for users engaging in complex tasks. With respect to developing reliable, ethical, and trustworthy IR and AI in general (cf. Schwartz et al., 2022), we also discuss how the studies on human bounded rationality could further extend current conceptualization and research on FATE in IR and redefine the assessment and regulation of AI-assisted interactive search systems and retrieval algorithms.

Compared to mature standardized IR evaluation experiments (e.g., TREC¹) and recent fast-growing research on algorithmic bias and fairness,² the research on human bounded rationality and its applications in IR problems is still at a very early stage. However, with the increasing interests on human perceptions and cognitive biases in multiple fields of computing research (e.g., Barbosa & Chen, 2019; Dingler et al., 2020; Draws et al., 2021; Lee & Rich, 2021; Taniguchi et al., 2018; Saab et al., 2019), it is an appropriate timing to draw attention to and further investigate the intersection between bounded rationality research and IR experiments. Our book contributes to this line of research mainly by clarifying the related theoretical roots and technical basis, synthesizing the insights and empirical findings from multiple disciplines that may be useful for IR modeling and evaluation, and developing a bias-aware research approach with specific open problems and new directions. To address the identified research problems (see Chap. 7), future studies will need to make further progresses on four aspects:

1. Further studying the search behavioral patterns, cognitive activities, and decision-making models of boundedly rational users through user studies conducted in naturalistic settings
2. Designing, testing, and fine-tuning different forms of bias-aware formal user models and assessing their performances in predicting user behaviors (e.g., examination, clicking, query reformulation, and search stopping) and facilitating in situ adaptive ranking and recommendations
3. Developing and implementing experimental search systems of varying modalities (e.g., desktop search, mobile search, conversational search) that can detect potential human and algorithmic biases in real-time search sessions
4. Designing and meta-evaluating bias-aware search evaluation metrics that measure the actual contributions of search systems and retrieval algorithms to improving search effectiveness and addressing both human and algorithmic biases in decision-making activities.

In addition, to achieve these four goals, researchers also need to go beyond existing user study design and tools (e.g., Kelly, 2009; Liu & Shah, 2019) and overcome a series of new methodological obstacles, such as collecting users' in situ feedback on the role of biases in searching at different moments of sessions, designing realistic search tasks that could trigger the adoption of mental shortcuts in decision-making, and disambiguating the divergent effects that come from different cognitive biases.

Current methods for predicting and resolving the potential harmful impacts of algorithmic biases in IR mainly focus on computational components from ML pipelines (Mehrabi et al., 2021). However, human biases and societal factors are significant sources of AI biases in intelligent information systems of varying types

¹Text Retrieval Conference (TREC): <https://trec.nist.gov>

²A new research community emerged and is growing rapidly around the ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT): <https://faccconference.org/>

and are usually overlooked (Schwartz et al., 2022). Therefore, to address the challenges in developing unbiased, reliable search systems, IR researchers need to take all forms of biases into consideration. We hope that this book can serve as a useful starting point for the above research journey and motivate students and researchers from diverse backgrounds to further explore and advance the science and technology on supporting boundedly rational people interacting with information.

References

- Agosto, D. E. (2002). Bounded rationality and satisficing in young people's web-based decision making. *Journal of the American Society for Information Science and Technology*, 53(1), 16–27. <https://doi.org/10.1002/asi.10024>
- Azzopardi, L. (2021). Cognitive biases in search: A review and reflection of cognitive biases in information retrieval. In *Proceedings of the 2021 ACM SIGIR conference on human information interaction and retrieval* (pp. 27–37). ACM. <https://doi.org/10.1145/3406522.3446023>
- Barbosa, N. M., & Chen, M. (2019). Rehumanized crowdsourcing: A labeling framework addressing bias and ethics in machine learning. In *Proceedings of the 2019 ACM SIGCHI conference on human factors in computing systems* (pp. 1–12). ACM. <https://doi.org/10.1145/3290605.3300773>
- Dingler, T., Tag, B., Karapanos, E., Kise, K., & Dengel, A. (2020). Workshop on detection and design for cognitive biases in people and computing systems. In *Extended abstracts of the 2020 ACM SIGCHI conference on human factors in computing systems* (pp. 1–6). ACM. <https://doi.org/10.1145/3334480.3375159>
- Draws, T., Rieger, A., Inel, O., Gadiraju, U., & Tintarev, N. (2021). A checklist to combat cognitive biases in crowdsourcing. In *Proceedings of the AAAI conference on human computation and crowdsourcing* (Vol. 9, pp. 48–59) <https://ojs.aaai.org/index.php/HCOMP/article/view/18939>
- Ge, Y., Zhao, S., Zhou, H., Pei, C., Sun, F., Ou, W., & Zhang, Y. (2020). Understanding echo chambers in e-commerce recommender systems. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval* (pp. 2261–2270). <https://doi.org/10.1145/3397271.3401431>
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1(1), 107–143. <https://doi.org/10.1111/j.1756-8765.2008.01006.x>
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5), 1449–1475. <https://doi.org/10.1257/000282803322655392>
- Kelly, D. (2009). Methods for evaluating interactive information retrieval systems with users. *Foundations and Trends in Information Retrieval*, 3(1–2), 1–224. <https://doi.org/10.1561/15000000012>
- Lee, M. K., & Rich, K. (2021). Who is included in human perceptions of AI? Trust and perceived fairness around healthcare AI and cultural mistrust. In *Proceedings of the 2021 ACM SIGCHI conference on human factors in computing systems* (pp. 1–14). ACM. <https://doi.org/10.1145/3411764.3445570>
- Liu, J., & Han, F. (2020). Investigating reference dependence effects on user search interaction and satisfaction: A behavioral economics perspective. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval* (pp. 1141–1150). ACM. <https://doi.org/10.1145/3397271.3401085>

- Liu, J., & Shah, C. (2019). Interactive IR user study design, evaluation, and reporting. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 11(2), i–93. <https://doi.org/10.2200/S00923ED1V01Y2019051CR067>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Saab, F., Elhadj, I. H., Kayssi, A., & Chehab, A. (2019). Modelling cognitive bias in crowdsourcing systems. *Cognitive Systems Research*, 58, 1–18. <https://doi.org/10.1016/j.cogsys.2019.04.004>
- Schwartz, R., Vassilev, A., Greene, K., Perine, L., Burt, A., & Hall, P. (2022). *Towards a standard for identifying and managing bias in artificial intelligence*. NIST special publication, 1270.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- Taniguchi, H., Sato, H., & Shirakawa, T. (2018). A machine learning model with human cognitive biases capable of learning from small and biased datasets. *Scientific Reports*, 8(1), 1–13. <https://doi.org/10.1038/s41598-018-25679-z>
- Thaler, R. H. (2016). Behavioral economics: Past, present, and future. *American Economic Review*, 106(7), 1577–1600. <https://doi.org/10.1257/aer.106.7.1577>
- Weber, R. A., & Camerer, C. F. (2006). “Behavioral experiments” in economics. *Experimental Economics*, 9(3), 187–192. <https://doi.org/10.1007/s10683-006-9121-5>
- Xu, Z., Han, Y., Zhang, Y., & Ai, Q. (2020). E-commerce recommendation with weighted expected utility. In *Proceedings of the 29th ACM international conference on information & knowledge management* (pp. 1695–1704). ACM. <https://doi.org/10.1145/3340531.3411993>