Shuffling a Few Stalls in a Crowded Bazaar

Potential Impact of Document-side Fairness on Unprivileged Info-seekers

Seán Healy <sean.healy@adaptcentre.ie>
ADAPT Centre (Dublin City University)
PhD Supervisors:
Prof. Gareth Jones (School of Computing, DCU)
Prof. Aphra Kerr (Department of Sociology, NUIM)
Prof. Annalina Caputo (School of Computing, DCU; **on leave**)

Abstract

This research aims to bridge gaps between fairness theory and web search system realities. There remain both practical and relatedly ethical barriers to the adoption of fairness practice within web search, especially within *open web search* (i.e. transparent, non-centralised search). Addressing practical barriers is often a task for industry. However, fairness lies outside the profit motive, and it is therefore the place of the academy and non-profits to identify and tackle these remaining problems. Otherwise, academic fairness efforts appear performative.

Background & Motivation

John Rawls, 1971:

"First maximize the welfare of the worst off."

Provider-side FairIR (pre-2024):

Purported Outcome:
Fairer allocation of society's resources.

Other Outcome:
Fewer public relations issues for powerful, non-transparent tech firms.

Recent EU funding for a fairer, more transparent web

NGI INTERNET OF HUMANS — Open WebSearch .eu — HORIZON EUROPE — The Connecting Europe Facility — findhr

"... aims to shape the development and evolution of an internet that responds to people's fundamental needs, including trust, security, and inclusion"

"openwebsearch.eu will create an open European infrastructure for internet search, based on European values and jurisdiction."

"... tackles climate change, helps to achieve the UN's Sustainable Development Goals and boosts the EU's competitiveness and growth."

"Supporting the creation of an ecosystem of interoperable digital services, that make the European Digital Single Market work in practice."

"... facilitate the prevention, detection, and management of discrimination in human recommendation."

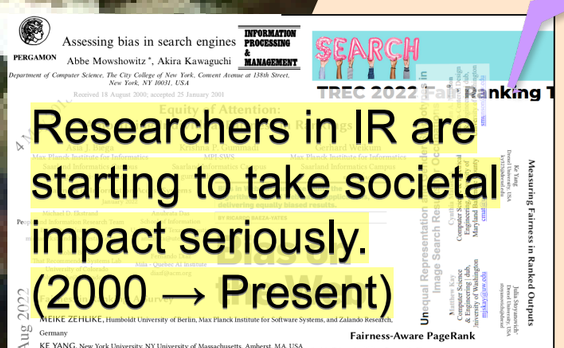

Researchers in IR are starting to take societal impact seriously. (2000 → Present)

However…

Most proposed methods work for recommendation, and not yet for ad-hoc web search.

Many proposed methods assume significant computing resources and centralized access to static identity information spanning searchers and search subjects.

Without decent fair ranking techniques that work on currently democratized hardware, small-to-medium sized businesses are locked out of the *fairness PR benefits* enjoyed by larger tech firms.

Small-to-medium sized firms employ 2/3 of the European workforce, and account for 99.8% of non-financial enterprises.

Small-to-medium sized business is thereby incentivized to imitate computing culture at larger tech firms: thus increasing carbon emissions and contributing to Big-tech cultural hegemony.

What would fair text retrieval research effort look like if the culture of computing tools and services looked like this?

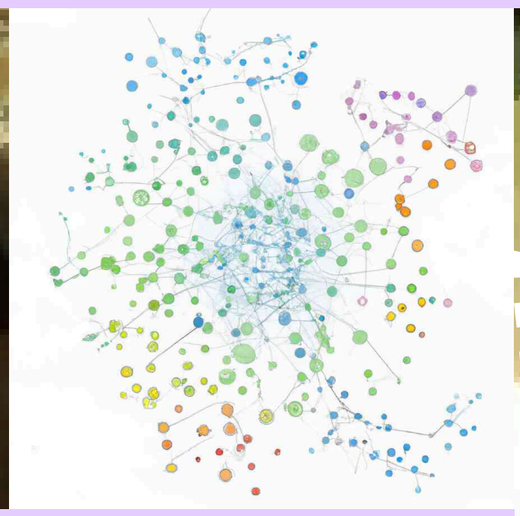

… and less like this?

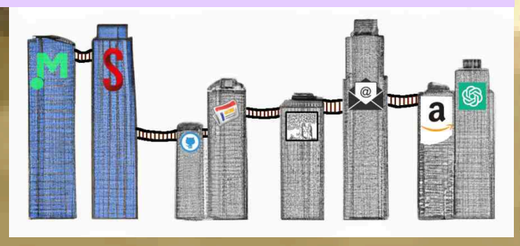

Current State-of-the-art

- Assume lots of duplicate queries (or none)
- Assume groups rely on the system evenly
- Assume salient biases are known
- More compute is better (available and necessary)

Research Questions

- How large is the user base?
- How much do different searcher groups rely on the system?
- What are the salient fairness considerations for the information need?
- What are the resources available/necessary?

Hypotheses

1. Query duplication levels are highly variable.
2. Different groups rely differently; quality decrease is a fairness issue.
3. Assumed unknown; estimation methods necessary.
4. There are efficient, negligibly sub-optimal interventions.

Method

Step #1

Simulation-based research justification
Hypotheses 2 and 3 (rephrased): Without novel methodologies, blindly applying fair ranking techniques to ad-hoc search can lead to disparate impact among unprivileged or understudied searchers.

variables and notation used in the simulation

Notation	Description
$\mathcal{U} = \{u_1 \dots u_{\lvert\mathcal{U}\rvert}\}$	seekers (also known as users in single stakeholder work)
$\mathcal{S} = \{s_1 \dots s_{\lvert\mathcal{S}\rvert}\}$	*systems* available to seekers
s_1	system where amortized individual fairness is applied
$\mathcal{G} = \{g_1, g_2\}$	protected *groups* for seekers
$\mathcal{M} : \mathcal{U} \to \mathcal{G}$	mapping seekers to their group *membership*
$C : \mathcal{U} \times \mathcal{S} \to \{0, 1\}$	availability of systems $\mathcal{S}$ to seekers $\mathcal{U}$, '*info capital*'
$\hat{r} : \mathcal{S} \times \mathbb{N}_1 \to [0, 1]$	estimated *relevance* function, according to systems in $\mathcal{S}$
$r : \mathbb{N}_1 \to [0, 1]$	ground-truth *relevance* function
τ	noticing *threshold*, i.e. minimum difference in NDCG across $\mathcal{S}$ before seekers consider results outside s_1
$\theta \in [0, 1]$	*threshold* for fairness intervention

Table 1. Table of notation

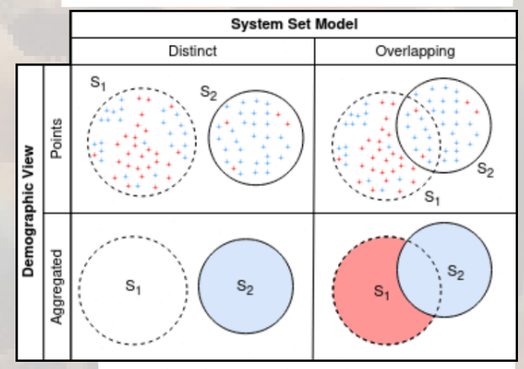

We simulate userbases in two competing systems, informed by polls and social science research into homophily and segregation across social networks and information systems [11, 12, 13].

modelling "empowered disillusionment" from a poorly performing system

$$v(u) \leftrightarrow \left(\max_{s' \in \mathcal{S}_u} \left[\text{NDCG}_{10}^{s'} \right] - \text{NDCG}_{10}^{s_1} > \tau \right)$$

modelling the information system choices made by users

$$\tilde{s}(u) = \begin{cases} s_1 & \text{when} \quad v(u) \\ \text{argmax}_{s \in \mathcal{S}_u} \left[\text{NDCG}_{10}^{s} \right] & \text{when} \quad \neg\, v(u) \end{cases}$$

modelling broader search engine utility to various user groups

$$R_g = \frac{\sum_{\{u_g \in \mathcal{U} \,:\, \mathcal{M}(u_g) = g\}} \text{NDCG}_n^{u_g}}{\sum_{\{u \in \mathcal{U}\}} \text{NDCG}_n^{u}}$$

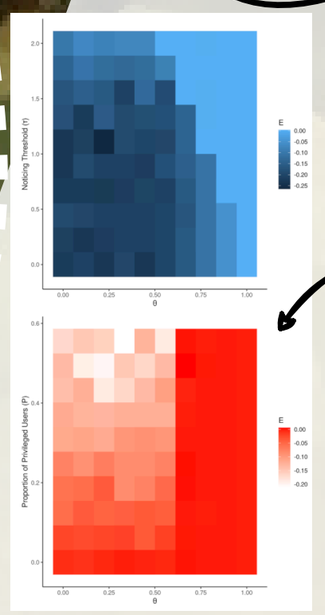

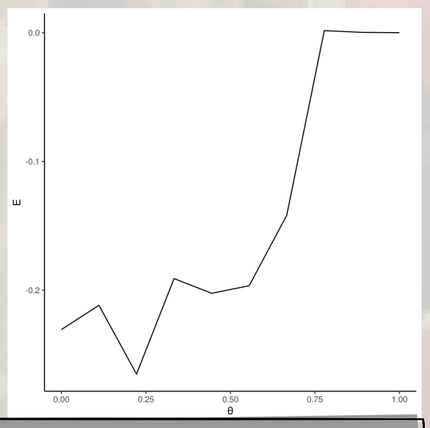

Under these assumptions, blind application of document-side fairness to ad-hoc search might disparately impact unprivileged info-seekers.

We repeated the simulation with varying values for the "noticing threshold" and privileged system segregation level. When simulated users are less prone to noticing quality decreases, there is less privilege-based disparate impact. When segregation levels are low, disparate impact is also low. However, quantitative social science research reveals that identity-based segregation, i.e. homophily, is both significant and widespread, except in ubiquitous networks [11].

***Conclusion*: Varied fairness interventions within web search should be applied at appropriate moments and with appropriate degree. We consider this an ethical issue as much as it is a system performance issue.**

i.e. not this

User: Show me images of German soldiers ca. 1944.

System:

Step #2, #3, #4

(future work)

- So, what are the salient biases for each web search query?
- Since this is ethical CS research, how can we calculate these bias saliencies without relying on Mechanical Turk economies?
- How can we perform these estimations (and the fairness intervention itself) in a way that is reproducible in businesses of all shapes and sizes, not just those with data centres?

preliminary progress: librecoir.com

References (1)

1. Biega, A.J., Gummadi, K.P., Weikum, G.: Equity of attention: Amortizing individual fairness in rankings. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. pp. 405–414 (2018)
2. Bonab, H., Aliannejadi, M., Vardasbi, A., Allan, E.K.J.: Cross-market product recommendation. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. pp. 110–119. CIKM '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3459637.3482493, https://doi.org/10.1145/3459637.3482493
3. Dinnissen, K., Bauer, C.: Fairness in music recommender systems: A stakeholder-centered mini review. Frontiers in Big Data **5** (2022). https://doi.org/10.3389/fdata.2022.913608, https://www.frontiersin.org/articles/10.3389/fdata.2022.913608
4. Fabbri, F., Bonchi, F., Boratto, L., Castillo, C.: The effect of homophily on disparate visibility of minorities in people recommender systems. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 14, pp. 165–175 (2020)
5. Fabbri, F., Croci, M.L., Bonchi, F., Castillo, C.: Exposure inequality in people recommender systems: The long-term effects. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 16, pp. 194–204 (2022)
6. de Jonge, T.: Exposure gerrymandering: Search engine manipulation flying under fairness' radar. In: FDIA '22: Future Directions in Information Retrieval. Lisbon, Portugal (2022)

References (2)

7. de Jonge, T., Hiemstra, D.: UNFair: Search engine manipulation, undetectable by amortized inequity. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. p. 830–839. FAccT '23, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3593013.3594046, https://doi.org/10.1145/3593013.3594046
8. Kearns, M., Neel, S., Roth, A., Wu, Z.S.: Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In: Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 2564–2572. PMLR (2018)
9. Lo, K., Wang, L.L., Neumann, M., Kinney, R., Weld, D.S.: S2orc: The semantic scholar open research corpus (2019)
10. Marsden, P.V.: Core discussion networks of americans. American Sociological Review **52**(1), 122–131 (1987), http://www.jstor.org/stable/2095397
11. McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a feather: Homophily in social networks. Annual review of sociology **27**(1), 415–444 (2001)
12. Overflow, S.: Stack overflow developer survey 2022. https://survey.stackoverflow.co/2022 (2022), accessed: 2023-11-14
13. Pseudomanifold: Poll: What are the gender demographics of hn? https://news.ycombinator.com/item?id=38031363 (2023), accessed: 2023-11-14
14. Rawls, J.: A theory of justice: Revised edition. Harvard university press (1999)
15. Roemer, J.E.: Equality of opportunity: A progress report. Social Choice and Welfare pp. 455–471 (2002)
16. Sakai, T., Kim, J.Y., Kang, I.: A versatile framework for evaluating ranked lists in terms of group fairness and relevance. arXiv preprint arXiv:2204.00280 (2022)
17. Selbst, A.D., Boyd, D., Friedler, S.A., Venkatasubramanian, S., Vertesi, J.: Fairness and abstraction in sociotechnical systems. In: Proceedings of the Conference on Fairness, Accountability, and Transparency. pp. 59–68. FAT* '19, Association for Computing Machinery, New York, NY, USA (2019)
18. Wu, H., Ma, C., Mitra, B., Diaz, F., Liu, X.: A multi-objective optimization framework for multi-stakeholder fairness-aware recommendation. ACM Trans. Inf. Syst. **41**(2) (dec 2022). https://doi.org/10.1145/3564285, https://doi.org/10.1145/3564285
19. Zehlike, M., Yang, K., Stoyanovich, J.: Fairness in ranking. ACM Comput. Surv. **55**(6) (2022). https://doi.org/10.1145/3533379