CO*IR: A Greedy and Individually Fair Re-ranker

Seán Healy, ADAPT Centre, Dublin City University

<sean.healy@adaptcentre.ie>

Background

What I talk about when I talk about "Individual Fairness"

"Individual fairness in ranking" is all about ensuring that search result exposure is allocated proportionally to merit.

This approach requires application across a number of duplicated queries.

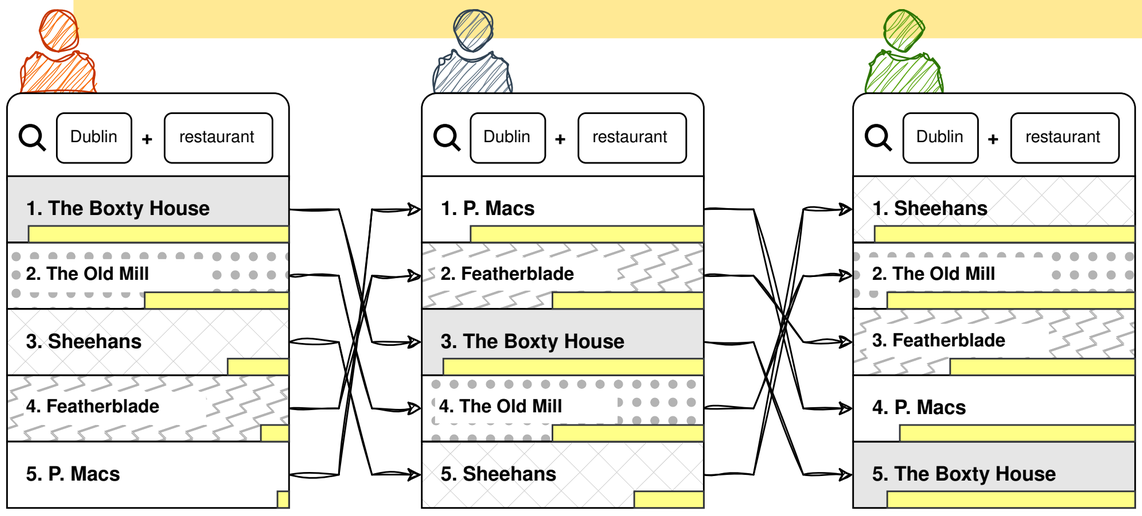

Fig 1. Amortized individual fairness (Biega, Gummandi and Weikum, 2018) in action.

Motivation

How can we increase fair ranking adoption in small-to-medium sized businesses?

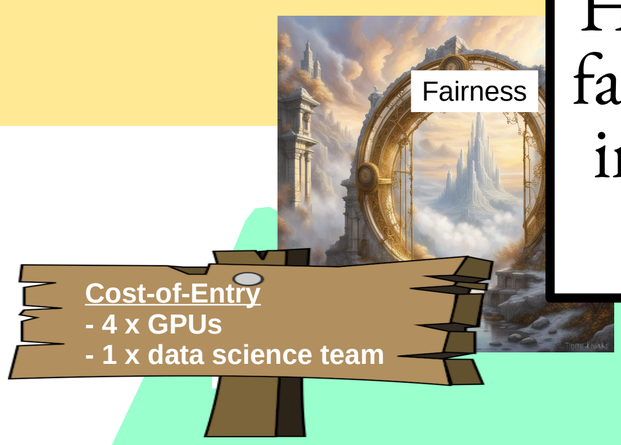

#1

- Fairness solutions can come without built-in quality constraints.
- Quality is a central concern of information system users and operators (Al-Maskari et al., 2010)

#2

- Academic lock-in: fairness solutions can be difficult to deploy.
- Contemporary research often places compatibility issues 'out-of-scope'.

#3

- Published fairness methodologies can be computationally intensive.
- This restricts fairness interventions to larger organisations.

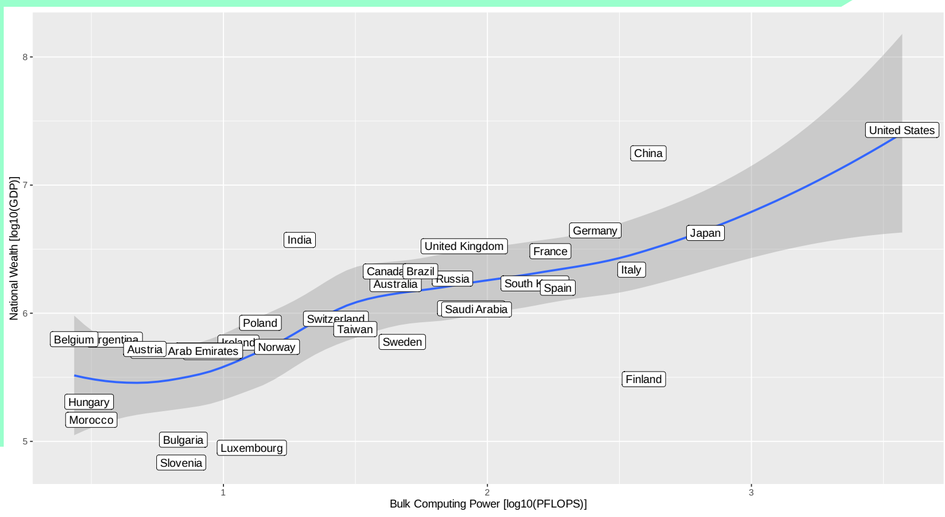

Fig 2. DL-based fairness interventions as bourgeois solutionism: Computational power is correlated with national wealth.

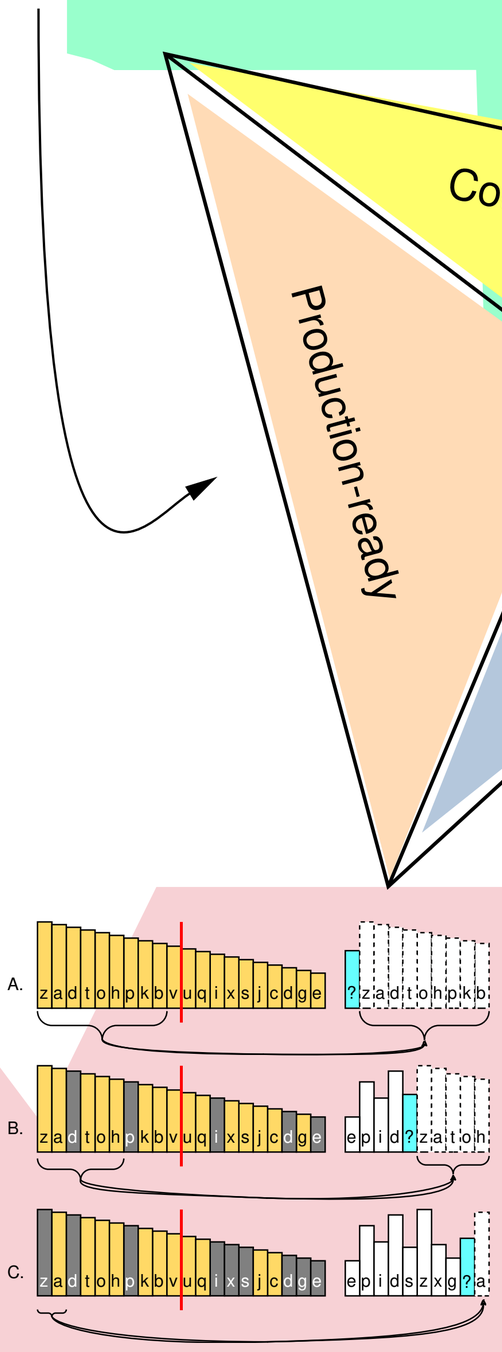

Fig 3. Visualising greedily fair re-ranking.

Proposed Methods

#1

- Ensure proposed methods include a quality threshold.

#2

- Package proposed solutions for immediate deployment; otherwise nobody else will ever do it.

#3

- Use an algorithmically 'greedy' approach; This reduce computational cost-of-entry.
- Focus on individually fair re-ranking first; This doesn't require identity information.

Explaining Greedy Ranking to a Multi-disciplinary Audience

- Produce a ranking item-by-item (highest ranks done first).
- Choose the most under-represented item at the current moment (except if it's really bad).
  - I.e. stay optimistic (but not unrealistic) that we will meet the quality threshold in the end.
- Expect some increase in how many duplicate queries are needed to achieve fairness. (sub-optimality)

Software package homepage:
https://librecoir.com/

References

Biega, A. J., Gummadi, K. P., & Weikum, G. (2018). Equity of attention: Amortizing individual fairness in rankings. In The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (pp. 405–414).
Zehlike, M., Yang, K., & Stoyanovich, J. (2022). Fairness in ranking. ACM Computing Surveys, 55(6), Article 118. Association for Computing Machinery.
Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M., & Baeza-Yates, R. (2017). FA*IR: A fair top-k ranking algorithm. In Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (pp. 1569–1578). Association for Computing Machinery.
Al-Maskari, A., & Sanderson, M. (2010). A review of factors influencing user satisfaction in information retrieval. Journal of the American Society for Information Science and Technology, 61(5), 859–868.