PhD Topic:
Computational Models of Fairness for Tomorrow's Web

Seán Healy B.A. (Mod.), M.Phil.
ADAPT Centre, Dublin City University

Engaging Content
Engaging People

Automatic Rankings Can Be Biased

- News aggregators, social media, streaming platforms, search engines, matchmaking apps, online marketplaces: all involve ranked lists of people, viewpoints, or items.
- If there is bias in these systems—and there usually is—it gets amplified across a large userbase.
- Counter-measures exist, but there are some remaining problems to address:

  (1) Efficiency issues: some of the current solutions require significant computing power (e.g. individual fairness; slide 3).

  (2) Multi-goal objectives: there is still limited ability to target multiple, competing fairness considerations in the present solutions.

  (3) Very little work looks at the task of identifying salient fairness issues in different scenarios. Automating this process would be necessary in web search settings.

  (4) Dealing with a scarcity of duplicate queries in web search is underexplored (slide 4).

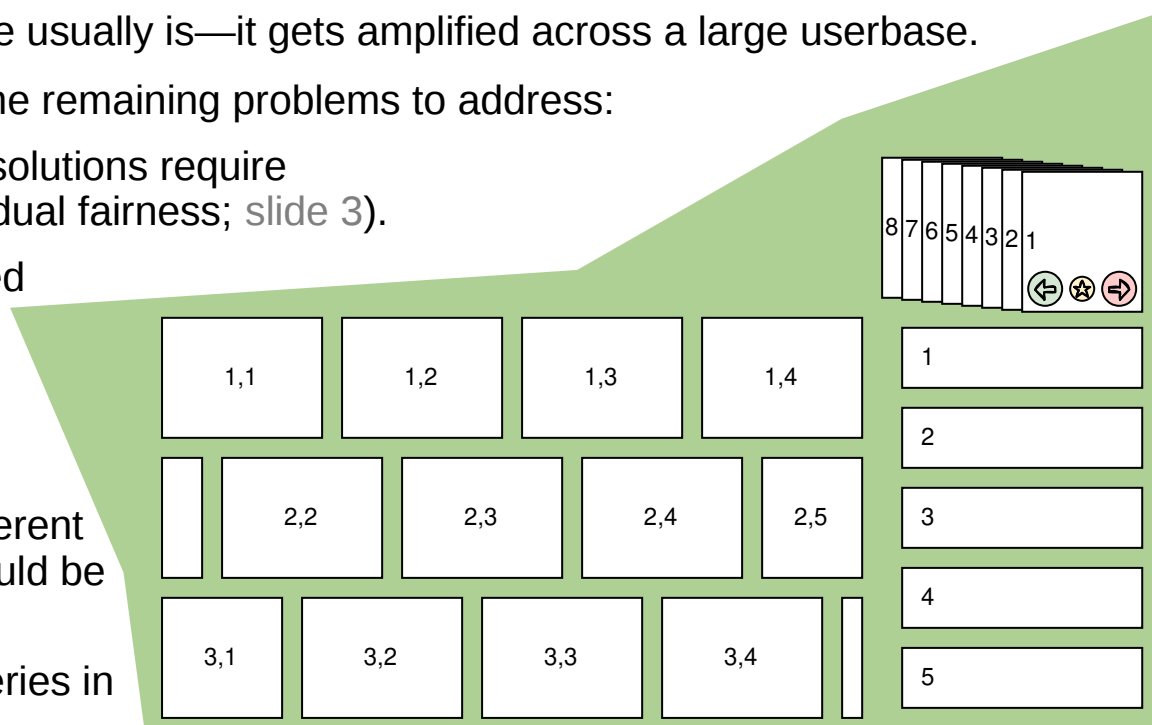

Individual vs. Group Fairness in Web Search

$$\forall\ \text{group(a)}, \text{group(b)}\ \ \frac{\text{relevance(a)}}{\text{exposure(a)}} = \frac{\text{relevance(b)}}{\text{exposure(b)}}$$

Examples:
– Ethnicity/gender (hiring)
– Political lean (news)
– Content genre (social feeds)

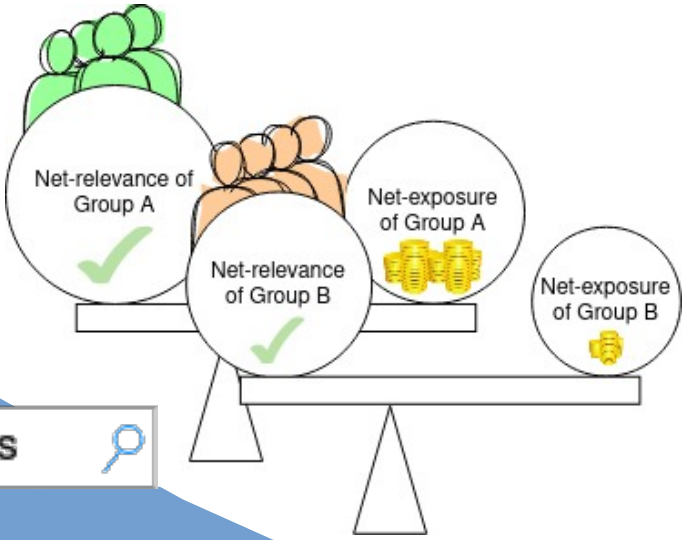

Engineering graduates

News

Music recommendations

… or *'user profile as query'*
(e.g. news aggregators, social feeds)

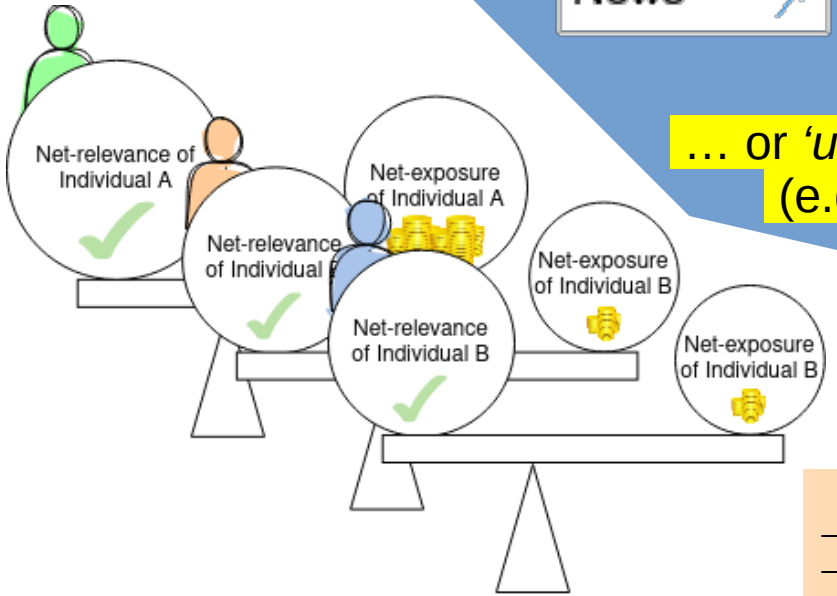

$$\forall\ \text{individual(a)}, \text{individual(b)}\ \ \frac{\text{relevance(a)}}{\text{exposure(a)}} = \frac{\text{relevance(b)}}{\text{exposure(b)}}$$

Examples:
– People, individual authors
– Entire web domains

Individual Fairness

Solutions to the individually fair ranking problem rely on a number of users having similar information needs: e.g. “Dublin + restaurant”. We call these “*query doppelgängers*”, and we call the process (top right) “*amortized fairness”*.

Results are returned in varying order to *query doppelgängers*, so that user attention (yellow bars) slowly mirrors relevance probability.

The currently proposed solution for individual fairness requires integer linear programming (bottom right), or other costly procedures.

(1) This only works well for up to 30 results.
(2) Requires strong computing hardware.
(3) Lacks integration with group fairness goals.
(4) Optimises around computationally simplified evaluation metrics (DCG).

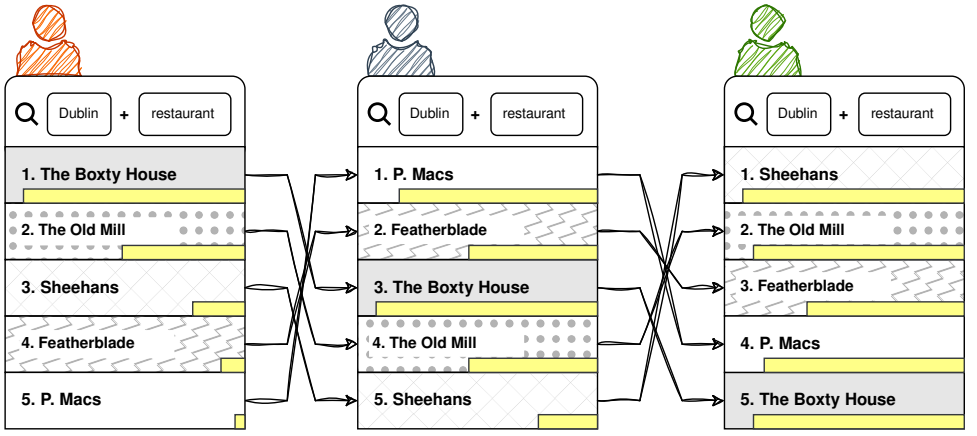

$$\text{minimize} \quad \sum_{i=1}^{n}\sum_{j=1}^{n} |A_i^{l-1} + w_j - (R_i^{l-1} + r_i^l)| \cdot X_{i,j}$$

$$\text{subject to} \quad \sum_{j=1}^{k}\sum_{i=1}^{n} \frac{2^{r_i^l} - 1}{log_2(j+1)} X_{i,j} \geq \theta \cdot IDCG@k$$

$$X_{i,j} \in \{0,1\}, \ \forall_{i,j}$$

$$\sum_i X_{i,j} = 1, \ \forall_j$$

$$\sum_j X_{i,j} = 1, \ \forall_i$$

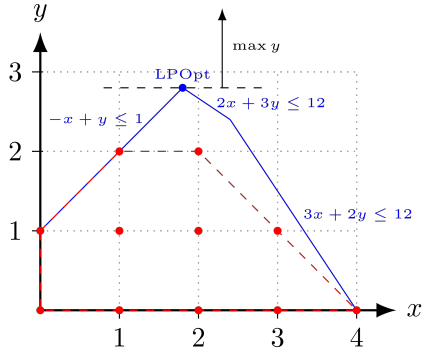

Individual Fairness: Broader Caveats

- Analysis of real search engine logs reveals a lack of *query doppelgängers*.
- Most queries to a search engine have only been repeated a handful of times.
- Fairness tasks in the IR community often hand-pick queries in a way that does not reflect real usage of search functionality within websites and web applications.
  - E.g. 15-30% of web queries are mispelt.
- Misspellings and re-wording can significantly degrade fairness interventions (right).

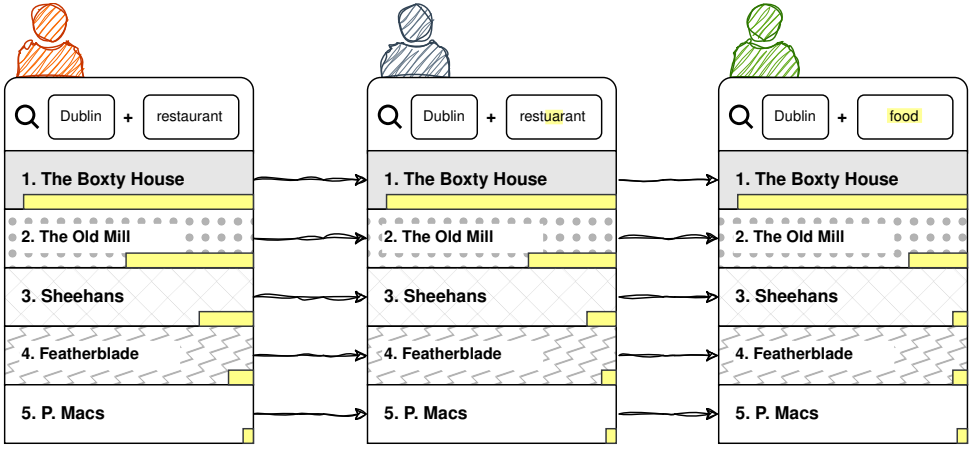

The *Matthew effect* in search—evading individual fairness intervention due to query reformulation.

- A research question emerges from this observation:
  - How far do we go to coalesce queries together?
    - Merging similar queries allows fair re-ranking to distribute exposure more evenly.
    - But doing this excessively will impact system functionality.

Computational Models of Fairness for Tomorrow's Web

Research Questions:

- Are the salient bias issues across varying search queries computable? [slide 1 (3), slide 2]
- Is there a theoretical model for individually fair re-ranking across longer lists? [slide 3 (2)]
- Can the decision of how much query merging to apply be partially automated? [slide 4]
  - Are there query merging techniques that will reliably never result in system degradation?

Conclusion

Renewed EU interest and funding for an open-web, based on digital sovereignty and informed user consent, poses new theoretical problems for fair re-ranking:

- We consider whether prior work pre-supposes a client-datacentre model.
- Fair ranking solutions can incur a large barrier of entry.
  - This centralises the state-of-the-art within the reach of large tech enterprises alone.
- Some examples of EU-funded open-web projects which often lack the structural means for current individual fairness interventions:
  - OpenWebSearch.eu – a federation of (high and low-resourced) web search providers.
  - DA VINCI – federated news aggregation system for journalists.
  - Spritely – decentralised social networking via modular components.

Bibliography (Progress)

[1] S. Healy, CO*IR: A Greedy and Individually Fair Re-ranker, AIMMES'24: Workshop on AI bias, Amsterdam, NL

[2] S. Healy, Shuffling a Few Stalls in a Crowded Bazaar: Potential Impact of Document-Side Fairness on Unprivileged Info-Seekers, ECIR 2024, Glasgow, Scotland

Grants received:

- €500 travel grant from FINDHR (AIMMES'24)
- €500 travel grant from Glasgow University (ECIR 2024)

This work is funded by the Science Foundation Ireland.

Project Website: librecoir.org

Engaging Content
Engaging People