

Overview

Under-representation of particular characteristics in a ranked list of Wikipedia articles can result in systematic biases that can have human, social and economic impact. We treat the Task 1 of Fair Ranking Track as an optimisation problem and use a popular diversification technique for re-ranking.

While there has been work investigating the relationship diversity and fairness metrics [3,4], this work investigates the impact of diversification in a fair ranking problem.

This work hypothesizes that: ***“When the protected group definitions are unknown, diversification techniques using the implicit features of the Wikipedia articles can achieve fairer results”.***

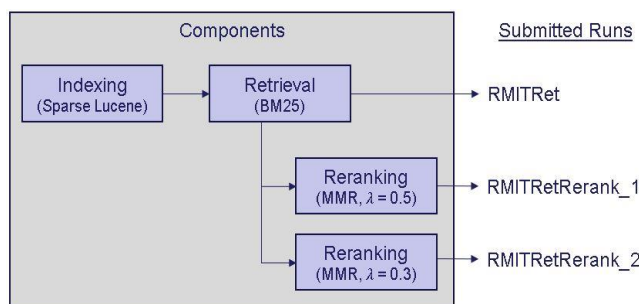
Challenges

- To optimise rankings for both fairness and relevance.
- Handle a broad range of group definitions (i.e., beyond associated geographic locations).

Methodology

In an effort, to handle a broad range of group definitions while optimising ranking for fairness and relevance, we take advantage of the implicit features of the Wikipedia articles using a diversity-based re-ranking technique - Maximal Marginal Relevance (MMR) - which tries to maximise novelty, diversity and relevance by picking documents most dissimilar to the previous documents in the ranked list [1]. The initial ranked list was obtained using BM25 [2].

Runs



In MMR, λ -parameter determines the degree of diversification. $\lambda=0$ gives maximal diversity, whereas $\lambda=1$ gives a standard relevance based ranked list. Our study varies the λ -values to investigate whether increased diversification leads to fairer results.

Results

| Run | System | Mean | | |
|-----------------|--------------------------------|--------|--------|--------|
| | | nDCG | AWRF | Score |
| RMITRet | BM25 | 0.2075 | 0.6413 | 0.1317 |
| RMITRetRerank_1 | BM25 + MMR ($\lambda = 0.5$) | 0.1760 | 0.6577 | 0.1144 |
| RMITRetRerank_2 | BM25 + MMR ($\lambda = 0.3$) | 0.1768 | 0.6582 | 0.1146 |

- Diversification using implicit features provided fairer results than the baseline RMITRet.
- Increasing the diversification aspect ($\lambda=0.3$) showed fairer results.
- Although implicit diversification in general showed fairer results, there was no statistically significant improvement (based on the independent t-tests) along any metrics, especially fairness.

Future Work

- In implicit diversification, the content of the Wikimedia articles prevails over the fields in the article which are potential candidates to be a protected characteristic. Hence we propose investigating how different fields in the articles can provide fairer results.

[1] Jaime Carbonell and Jade Goldstein. 1998. The use of MMR, diversity-based reranking for reordering documents and producing summaries. (SIGIR '98). DOI:<https://doi.org/10.1145/290941.291025>

[2] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at TREC-3. Nist Special Publication Sp109 (1995), 109.

[3] Ruoyuan Gao and Chirag Shah. 2020. Toward creating a fairer ranking in search engine results. Information Processing & Management 57, 1 (2020), 102138. <https://doi.org/10.1016/j.ipm.2019.102138>.

[4] Sachin Pathiyan Cherumanal, Damiano Spina, Falk Scholer, and W. Bruce Croft. 2021. Evaluating Fairness in Argument Retrieval. (CIKM'21). <https://doi.org/10.1145/3459637.3482099>.