

Language Fairness in Multilingual Information Retrieval

Eugene Yang
eugene.yang@jhu.edu
Johns Hopkins University
Baltimore, Maryland, USA

Thomas Jänich
t.jaenich.1@research.gla.ac.uk
University of Glasgow
Glasgow, United Kingdom

James Mayfield
mayfield@jhu.edu
Johns Hopkins University
Baltimore, Maryland, USA

Dawn J. Lawrie
lawrie@jhu.edu
Johns Hopkins University
Baltimore, Maryland, USA

ABSTRACT

Multilingual information retrieval (MLIR) considers the problem of ranking documents in several languages for a query expressed in a language that may differ from any of those languages. Recent work has observed that approaches such as combining ranked lists representing a single document language each or using multilingual pretrained language models demonstrate a preference for one language over others. This results in systematic unfair treatment of documents in different languages. This work proposes a language fairness metric to evaluate whether documents across different languages are fairly ranked through statistical equivalence testing using the Kruskal-Wallis test. In contrast to most prior work in group fairness, we do not consider any language to be an unprotected group. Thus our proposed measure, PEER (Probability of Equal Expected Rank), is the first fairness metric specifically designed to capture the language fairness of MLIR systems. We demonstrate the behavior of PEER on artificial ranked lists. We also evaluate real MLIR systems on two publicly available benchmarks and show that the PEER scores align with prior analytical findings on MLIR fairness. Our implementation is compatible with `ir-measures` and is available at http://github.com/hltcoe/peer_measure.

CCS CONCEPTS

• **Information systems** → **Multilingual and cross-lingual retrieval**; **Evaluation of retrieval results**.

KEYWORDS

language fairness, multilingual retrieval, statistical testing

ACM Reference Format:

Eugene Yang, Thomas Jänich, James Mayfield, and Dawn J. Lawrie. 2024. Language Fairness in Multilingual Information Retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3626772.3657943>

1 INTRODUCTION

Multilingual information retrieval searches a multilingual document collection and creates a unified ranked list for a given query [4–6, 16, 20, 24]. In tasks like navigational search [7], known item retrieval [1, 21], and retrieval for question-answering [10, 19], the user only needs a handful or just one relevant document to satisfy

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGIR '24, July 14–18, 2024, Washington, DC, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0431-4/24/07.
<https://doi.org/10.1145/3626772.3657943>

the information need, and the language of that document does not matter. In contrast, users interested in gaining a broad problem understanding prefer seeing how coverage varies across languages.

Analysis of retrieval results has shown that MLIR systems often show a preference for certain languages [16, 20]; we call this the *MLIR Fairness problem*. Such preference can bias a user's understanding of the topic [28, 30]. This problem is particularly apparent in models built on top of multilingual pretrained language models (mPLM) [20], which inherit bias from the text used to build them [11, 17]. This paper presents a new metric to allow quantitative study of MLIR Fairness.

Prior work in fairness evaluation focuses on either individual or group fairness [33]. Individual fairness ensures that similar documents receive similar treatment; this often corresponds to a Lipschitz condition [3, 15]. Group fairness ensures that a protected group receives treatment at least as favorable as unprotected groups [27, 31, 32]. Group fairness metrics designed for protecting specific groups are not directly applicable to the MLIR fairness problem because the latter has no protected language; we want all languages to be treated equally in a ranked list.

To operationalize our notion of MLIR fairness, we propose the Probability of Equal Expected Rank (PEER) metric. By adopting the Kruskal-Wallis H test, which is a rank-based, non-parametric variance analysis for multiple groups, we measure the probability that documents of a given relevance level for a query are expected to rank at the same position irrespective of language. We compare PEER to previous fairness metrics, and show its effectiveness on synthetic patterned data, on synthetic assignment of language to real retrieval ranked lists, and on system output for the CLEF 2003 and NeuCLIR 2022 MLIR benchmarks.

2 RELATED WORK

There is no universally accepted definition of fairness. This paper views languages as groups within a ranking, and characterizes MLIR Fairness as a group fairness problem.

Existing group fairness metrics fall into two categories: those that assess fairness independent of relevance, and those that take relevance into account. *Ranked group fairness*, based on statistical parity proposed by Zehlike et al. [31, 32], demands equitable representation of protected groups in ranking without explicitly considering relevance through statistical testing.

Attention Weighted Ranked Fairness (AWRF), introduced by Sapiezynski et al. [27], compares group exposure at certain rank cutoffs against a pre-defined target distribution. It uses the same distribution for both relevant and nonrelevant documents. This means for example that if utility is defined as finding the most relevant documents, a system can gain utility by including more documents from the language with the most relevant documents early in the rankings. In doing so, more nonrelevant documents

from that language are placed above relevant documents from the other languages. From a fairness perspective, this should be penalized as unfair. Our proposed metric does not rely on a target distribution, so it does not suffer from this utility/fairness tradeoff.

Among metrics that incorporate relevance, Singh and Joachims [29] introduced the Disparate Treatment Ratio, which measures the equality of exposure of two groups. This metric is not well suited to MLIR though, since it handles only two groups. Adjacent to fairness, Clarke et al. [13] extended Normalized Discounted Cumulative Gain (nDCG) to incorporate diversity. Their metric, α -nDCG, assigns document weights based on both relevance and diversity. Diversity though applies to *user* utility where fairness applies to *documents* (in our case, the languages of the returned documents) [8]. We nonetheless report α -nDCG to contextualize our results.

Related work on fairness over sequences of rankings [14, 23] requires both more evidence and distributional assumptions compared to fairness of a specific ranking. While similar, our method assumes the position of each document is a random variable.

3 PROBABILITY OF EQUAL EXPECTED RANK

In this section, we describe the proposed measure – PEER: Probability of Equal Expected Rank. We first introduce our notation and the fairness principle, followed by forming the statistical hypothesis of the system’s fairness across document languages. Finally, we define PEER as the p -value of the statistical test.

Let $d_i \in \mathcal{D}$ be the i -th document in the collection \mathcal{D} of size N . We define the language that d_i is written in as $l_{d_i} \in \{\mathcal{L}_1, \dots, \mathcal{L}_M\}$. For convenience, we define the set $L_j = \{d_i \mid l_{d_i} = \mathcal{L}_j\}$ to be all documents in language \mathcal{L}_j .

For a given query $q \in \mathcal{Q}$, we define the degree of document d_i being relevant (or the relevance grade) to the query q as $y_i^q \in \{\mathcal{R}^{(0)}, \mathcal{R}^{(1)}, \dots, \mathcal{R}^{(K)}\}$, where $\mathcal{R}^{(0)}$ indicates not relevant and $\mathcal{R}^{(K)}$ is the most relevant level, i.e., graded-relevance with K levels. Similarly, we define the set $R^{(q,k)} = \{d_i \mid y_i^q = \mathcal{R}^{(k)}\}$ to be all documents at the $\mathcal{R}^{(k)}$ relevance level. Furthermore, we define the documents in \mathcal{L}_j with relevance level $\mathcal{R}^{(k)}$ for a query q as $D_j^{(q,k)} = L_j \cap R^{(q,k)}$.

In this work, we consider a ranking function $\pi : \mathcal{D} \times \mathcal{Q} \rightarrow [1 \dots N]$ that produces the rank $r_i^q \in [1 \dots N]$.

3.1 Fairness through Hypothesis Testing

We define MLIR fairness using the following principle: *Documents in different languages with the same relevance level, in expectation, should be presented at the same rank.* We measure the satisfaction level of this principle by treating it as a testable hypothesis.

For relevance level $\mathcal{R}^{(k)}$, assuming r_i^q is a random variable over $[1 \dots N]$, we implement the principle using the null hypothesis:

$$H_0 : \mathbb{E}_{d_i \in D_a^{(q,k)}} [r_i^q] = \mathbb{E}_{d_j \in D_b^{(q,k)}} [r_j^q] \quad \forall \mathcal{L}_a \neq \mathcal{L}_b, \quad (1)$$

which is the equivalence of the expected rank among documents in each language with the given relevance level and given query q .

Such null hypotheses can be tested with the Kruskal-Wallis H test (K-W test) [18]. The null hypothesis of this test is that all groups have the same mean (i.e., equivalent mean ranks). The K-W test is like a non-parametric version of the ANOVA F-test, which tests

whether each group (languages in our case) comes from the same distribution. Since the K-W test does not assume any particular underlying distribution, it uses the ranking of the data points to make this determination. Unlike prior work such as Zehlike et al. [32] that assumes a binomial distribution for each document over the groups, not assuming the distribution of the query-document scores used for ranking and instead operating directly on ranks yields a robust statistical test.

Conceptually, the test statistics H for the K-W test is the ratio between the sum of group rank variance and the total rank variance. The variance ratio obeys a chi-squared distribution; we use its survival function to derive the p -value. Specifically, we can express the test statistic H as

$$H = \left(|R^{(q,k)}| - 1 \right) \frac{\sum_{j=1}^M |D_j^{(q,k)}| \left(\bar{r}_j^{q,k} - \bar{r} \right)^2}{\sum_{j=1}^M \sum_{d_i \in D_j^{(q,k)}} \left(r_i^q - \bar{r} \right)^2} \quad (2)$$

$$\text{where } \bar{r}_j^{q,k} = \frac{1}{|D_j^{(q,k)}|} \sum_{d_i \in D_j^{(q,k)}} r_i^q \quad (3)$$

$$\text{and } \bar{r} = \frac{1}{|R^{(q,k)}|} \sum_{j=1}^M \sum_{d_i \in D_j^{(q,k)}} r_i^q \quad (4)$$

for a given query q and relevance level $\mathcal{R}^{(k)}$. Recall that $R^{(q,k)}$ and $D_j^{(q,k)}$ are sets. For each query q and given relevance level, we report the p -value of the K-W test, which is the Probability of Equal Expected Rank, by comparing the H statistic against a chi-squared distribution with $M - 1$ degrees of freedom. The p -value provides us with the probability that documents in different languages are ranked fairly within a given relevance level. We denote the p -value for a given query q and a relevance level $\mathcal{R}^{(k)}$ as $p^{(q,k)}$.

Our fairness notion is similar to the one proposed by Diaz et al. [14]. However, we operationalize the principle by treating each document as a sample from a distribution given the language and relevance level, instead of assuming the entire ranked list is a sample from all possible document permutations.

3.2 Fairness at Each Relevance Level

The impact of unfairly ranking documents in different languages may differ at each relevance level. Such differences can be linked to a specific user model or application. For example, for an analyst actively seeking information for which each language provides different aspects, ranking nonrelevant documents of a particular language at the top does not degrade fairness; finding disproportionately fewer relevant documents in a certain language, on the other hand, may yield biased analytical conclusions.

In contrast, for a user seeking answers to a specific question who views the language as just the content carrier, reading more nonrelevant documents from a language may degrade that language’s credibility, leading the user eventually to ignore all content in that language. In this case, we do consider the language fairness of the ranking of nonrelevant documents; in the former case we do not.

To accommodate different user models, we define the PEER score as a linear combination of the p -value of each relevance level $\mathcal{R}^{(k)}$.




Figure 1: Ranked lists with different fairness patterns between two languages and binary relevance.

Let $w^{(k)} \in [0, 1]$ be the weights and $\sum_{k=1}^K w^{(k)} = 1$, the overall weighted PEER for query q is

$$PEER^{(q)} = \sum_{k=1}^K w^{(k)} p^{(q,k)} \quad (5)$$

3.3 Rank Cutoff and Aggregation

While a ranking function π ranks each document in collection \mathcal{D} , in practice, a user only examines results up to a certain cutoff. Some IR effectiveness measurements consider elaborate browsing models, such as exponential-decreasing attention in Ranked-biased precision (RBP) [22] or patience-based attention in expected reciprocal rank (ERR) [9]; user behavior though is perpendicular to language fairness, so we consider only a simple cutoff model.

With a rank cutoff X , we treat only the top- X documents as the sampling universe for the K-W test. However, since disproportionately omitting documents of a certain relevance level is still considered unfair, before conducting the hypothesis test, we concatenate unretrieved documents (or those ranked below the cutoff) at that relevance level to the ranked list, assigning them a tied rank of $X + 1$. This is optimistic, since these documents might rank lower in the actual ranked list. However, this provides a robust penalty for any ranking model that provides only a truncated ranked list. We define the p -value as 1.0 when no document is retrieved at a given relevance level in spite of their presence in the collection; from the user perspective, no document at that level is presented, so it is *fair* (albeit ineffective) across languages. We denote the weighted p -value calculated on the top- X documents as $PEER^{(q)}@X$.

Overall, we report the average weighted PEER over all queries at rank X , i.e., $PEER@X = |\mathcal{Q}|^{-1} \sum_{q \in \mathcal{Q}} PEER^{(q)}@X$. Since we treat each document as a random variable of position in a ranked list, what we are measuring is *how likely* a system is fair between languages instead of how fair each ranked list is. A higher PEER score indicates that the measured MLIR system is *more likely* to place documents written in different languages but with the same relevance level at similar ranks.

4 EXPERIMENTS AND RESULTS

4.1 Synthetic Data

To demonstrate PEER behavior, we create ranked lists of two languages and binary relevance with four methods, each creating lists from very unfair to very fair. Results are illustrated in Figure 1. *Shifting* starts with all documents in one language ranking higher than those in the other, and slowly interleaves them until the two are alternating. Figure 1(a) shows for fifty documents that when no documents are interleaved (left), fairness is low, with PEER close to 0. As alternation increases, the PEER score increases. In *Moving Single*, the ranked list consists entirely of one language except for one document. That single document moves from the top (unfair) to the middle of the ranking (fair). In Figure 1(b) with 99 majority language documents, the PEER scores increase as the singleton moves from the top to the middle. Figure 1(c) shows *Interleaving*, in which the languages alternate and the number of retrieved documents slowly increases. With odd lengths the highest and the lowest ranked documents are in the same language, giving them the same average and 1.0 PEER scores. With even lengths, one language has a slightly higher rank than the other. The difference shrinks with longer ranked lists, resulting in increased PEER scores. In *Increasing Length*, 100 retrieved documents comprise first an alternating section followed by all documents in a single language, and the size of the alternating section is gradually increased. This is similar to *shifting*, but with overlapping languages at the top instead of in the middle of the rank list. At the left of Figure 1(d) only the document at rank 1 is minority language, followed by minority language at ranks 1 and 3, and so on. The right of the graph is identical to the right of Figure 1(a). These four patterns demonstrate that PEER scores match our intuition of fairness between languages. The next section evaluates real MLIR retrieval systems on two MLIR evaluation collections.

4.2 Assigning Languages to a Real Ranked List

We used NeuCLIR'22 runs to create new synthetic runs with relevant documents in the same positions, but with languages assigned to the relevant documents either fairly or unfairly. We randomly

Table 1: Effectiveness and fairness results. Both AWRF and PEER exclude nonrelevant documents. They are removed in AWRF calculation. For PEER, the importance weights of nonrelevant documents are set to 0.

Collection	Rank Cutoff	20				1000			
	Measure	nDCG	α -nDCG	AWRF	PEER	Recall	α -nDCG	AWRF	PEER
CLEF 2003	QT » BM25	0.473	0.444	0.513	0.239	0.743	0.579	0.788	0.202
	DT » BM25	0.636	0.640	0.623	0.243	0.857	0.747	0.895	0.299
	DT » ColBERT	0.669	0.674	0.658	0.293	0.889	0.768	0.904	0.328
	ColBERT-X ET	0.591	0.592	0.610	0.215	0.802	0.695	0.845	0.327
	ColBERT-X MTT	0.643	0.658	0.649	0.318	0.827	0.748	0.860	0.362
NeuCLIR 2022	QT » BM25	0.305	0.447	0.537	0.453	0.557	0.569	0.752	0.383
	DT » BM25	0.338	0.448	0.542	0.497	0.633	0.580	0.809	0.421
	DT » ColBERT	0.403	0.539	0.635	0.449	0.708	0.652	0.842	0.426
	ColBERT-X ET	0.299	0.447	0.578	0.458	0.487	0.561	0.745	0.421
	ColBERT-X MTT	0.375	0.545	0.621	0.425	0.612	0.644	0.786	0.386

selected how many relevant documents would be assigned to each language, and created that many language labels. For each label we drew from a normal distribution with either the same mean for the two languages (fair), or different means (unfair). We assigned a drawn number to each label, sorted the labels by that number, and assigned the labels to the relevant documents in the resulting order. We did the same for the nonrelevant documents, ensuring that each language was assigned at least 45% of those documents.

Figures 1(e) and (f) vary the sampling mean of the second language’s relevant and nonrelevant documents, respectively, while keeping a first language sampling mean of 1.0. The figures show that PEER captures fairness independently for each relevance level. Since there are far fewer relevant documents, the evidence for fairness is also weaker, resulting in slower decay when changing the sampling mean for relevant documents.

4.3 Real MLIR Systems

We evaluate five MLIR systems, including query translation (QT) and document translation (DT) with BM25, DT with English ColBERT [26], and ColBERT-X models trained with English triples (ET) and multilingual translate-train (MTT) [20], on CLEF 2003 (German, Spanish, French, and English documents with English queries) and NeuCLIR 2022 (Chinese, Persian, and Russian documents with English queries). For QT, English queries are translated into each document language and monolingual search results from each language are fused by score. We report α -nDCG, AWRF (with number of relevant documents as target distribution), and the proposed PEER with rank cutoffs at 20 and 1000. Along with nDCG@20 and Recall@1000, we summarize the results in Table 1.

Logically, merging ranked lists from each monolingual BM25 search with translated queries purely by scores is inviting unfair treatment, as scores from each language are incompatible with different query lengths and collection statistics [25]. We observed this trend in both PEER and AWRF, while α -nDCG strongly correlates with the effectiveness scores and does not distinguish fairness.

Neural MLIR models trained with only English text and transferred zero-shot to MLIR with European languages exhibit a strong language bias compared to those trained with document languages

[16, 20]. PEER exhibits a similar trend in CLEF 2003, showing ColBERT-X ET is less fair than the MTT counterpart, while AWRF is less sensitive. Lawrie et al. [20] show that preference for English documents in the ET model causes this unfair treatment; this suggests that MLIR tasks without the training language (English) in the document collection would not suffer from such discrepancy.

AWRF and PEER disagree on the comparison between English ColBERT on translated documents (DT) and ColBERT-X models. While AWRF suggests DT » ColBERT¹ is fairer than ColBERT-X MTT in CLEF03, DT creates a larger difference among languages [20]. PEER, in contrast, aligns with prior analysis, giving a lower score to DT » ColBERT. According to Huang et al. [16], QT » BM25 has a similar language bias compared to mDPR [12], which was trained with English MS MARCO [2]. PEER suggests a similar conclusion between QT » BM25 and ColBERT-X ET, which AWRF assigns a larger difference between the two with a rank cutoff of 20.

With a rank cutoff of 1000, AWRF strongly correlates with recall (Pearson $r = 0.93$ over both collections), while PEER does not (Pearson $r = -0.55$). The 0.904 AWRF value (range 0-1) of DT » ColBERT on CLEF03 suggests a fair system, while the ranked list does not. This strong relationship shows that AWRF, with target distribution being the ratio of relevant documents, is indeed measuring recall instead of fairness. While it is an artifact of the choice of target distribution, the need to define a target distribution reduces the robustness of AWRF in measuring MLIR Fairness.

5 SUMMARY

We propose measuring the Probability of Equal Expected Rank (PEER) for MLIR fairness. As PEER measures the weighted p -value of a non-parametric group hypothesis test, it neither requires a target distribution nor makes distributional assumptions; this makes the metric robust. Through comparison to prior analytical work in MLIR Fairness, we conclude that PEER captures the differences and nuances between systems better than other fairness metrics.

¹The (ET)+ITD setting in Lawrie et al. [20].

REFERENCES

- [1] Jaime Arguello, Adam Ferguson, Emery Fine, Bhaskar Mitra, Hamed Zamani, and Fernando Diaz. 2021. Tip of the tongue known-item retrieval: A case study in movie identification. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. 5–14.
- [2] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016).
- [3] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *Proc. SIGIR*. 405–414.
- [4] Martin Braschler. 2001. CLEF 2001—Overview of Results. In *Workshop of the Cross-Language Evaluation Forum for European Languages*. Springer, 9–26.
- [5] Martin Braschler. 2002. CLEF 2002—Overview of results. In *Workshop of the Cross-Language Evaluation Forum for European Languages*. Springer, 9–27.
- [6] Martin Braschler. 2003. CLEF 2003—Overview of results. In *Workshop of the Cross-Language Evaluation Forum for European Languages*. Springer, 44–63.
- [7] Andrei Broder. 2002. A taxonomy of web search. In *ACM Sigir forum*, Vol. 36. ACM New York, NY, USA, 3–10.
- [8] Carlos Castillo. 2019. Fairness and transparency in ranking. In *ACM SIGIR Forum*, Vol. 52. ACM New York, NY, USA, 64–71.
- [9] Olivier Chapelle, Donald Metzler, Ya Zhang, and Pierre Grinspan. 2009. Expected reciprocal rank for graded relevance. In *Proceedings of the 18th ACM conference on Information and knowledge management*. 621–630.
- [10] Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*. Association for Computational Linguistics (ACL), 1870–1879.
- [11] Monojit Choudhury and Amit Deshpande. 2021. How Linguistically Fair Are Multilingual Pre-Trained Language Models?. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 12710–12718.
- [12] Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages. *Transactions of the Association for Computational Linguistics* (2020).
- [13] Charles L. A. Clarke, Maheedhar Kolla, Gordon V. Cormack, Olga Vechtomova, Azin Ashkan, Stefan B'uttcher, and Ian MacKinnon. 2008. Novelty and diversity in information retrieval evaluation. In *SIGIR*.
- [14] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 275–284.
- [15] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 214–226.
- [16] Zhiqi Huang, Hansi Zeng, Hamed Zamani, and James Allan. 2023. Soft Prompt Decoding for Multilingual Dense Retrieval. *arXiv preprint arXiv:2305.09025* (2023).
- [17] Nora Kassner, Philipp Dufter, and Hinrich Schütze. 2021. Multilingual LAMA: Investigating knowledge in multilingual pretrained language models. *arXiv preprint arXiv:2102.00894* (2021).
- [18] William H Kruskal and W Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* 47, 260 (1952), 583–621.
- [19] Heeyoung Kwon, Harsh Trivedi, Peter Jansen, Mihai Surdeanu, and Niranjan Balasubramanian. 2018. Controlling information aggregation for complex question answering. In *Advances in Information Retrieval: 40th European Conference on IR Research, ECIR 2018, Grenoble, France, March 26-29, 2018, Proceedings 40*. Springer, 750–757.
- [20] Dawn Lawrie, Eugene Yang, Douglas W Oard, and James Mayfield. 2023. Neural Approaches to Multilingual Information Retrieval. In *European Conference on Information Retrieval*. Springer, 521–536.
- [21] Jin Ha Lee, Allen Renear, and Linda C Smith. 2006. Known-item search: Variations on a concept. *Proceedings of the american society for information science and technology* 43, 1 (2006), 1–17.
- [22] Alistair Moffat and Justin Zobel. 2008. Rank-biased precision for measurement of retrieval effectiveness. *ACM Transactions on Information Systems (TOIS)* 27, 1 (2008), 1–27.
- [23] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling fairness and bias in dynamic learning-to-rank. In *Proc. SIGIR*. 429–438.
- [24] Razieh Rahimi, Azadeh Shakery, and Irwin King. 2015. Multilingual information retrieval in the language modeling framework. *Information Retrieval Journal* 18, 3 (2015), 246–281.
- [25] Stephen E Robertson and Ian Soboroff. 2002. The TREC 2002 Filtering Track Report. In *Proceedings of the Eleventh Text Retrieval Conference*.
- [26] Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, 3715–3734.
- [27] Piotr Sapiezynski, Wesley Zeng, Ronald E Robertson, Alan Mislove, and Christo Wilson. 2019. Quantifying the impact of user attention on fair group representation in ranked lists. In *Companion Proceedings of WWW*. 553–562.
- [28] Stefan Schulz-Hardt, Dieter Frey, Carsten Lüthgens, and Serge Moscovici. 2000. Biased information search in group decision making. *Journal of personality and social psychology* 78, 4 (2000), 655.
- [29] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proc. KDD*.
- [30] Ryan White. 2013. Beliefs and biases in web search. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. 3–12.
- [31] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. Fa* ir: A fair top-k ranking algorithm. In *Proc. of CIKM*.
- [32] Meike Zehlike, Tom Sühr, Ricardo Baeza-Yates, Francesco Bonchi, Carlos Castillo, and Sara Hajian. 2022. Fair Top-k Ranking with multiple protected groups. *Information processing & management* 59, 1 (2022), 102707.
- [33] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2021. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000* (2021).