

Evaluating Fairness in Argument Retrieval

Sachin Pathiyan Cherumanal

RMIT University
Melbourne, Australia
sachin.pathiyan.cherumanal@student.rmit.edu.au

Falk Scholer

RMIT University
Melbourne, Australia
falk.scholer@rmit.edu.au

Damiano Spina

RMIT University
Melbourne, Australia
damiano.spina@rmit.edu.au

W. Bruce Croft

University of Massachusetts
Amherst, MA, USA
croft@cs.umass.edu

ABSTRACT

Existing commercial search engines often struggle to represent different perspectives of a search query. Argument retrieval systems address this limitation of search engines and provide both positive (PRO) and negative (CON) perspectives about a user’s information need on a controversial topic (e.g., climate change). The effectiveness of such argument retrieval systems is typically evaluated based on topical relevance and argument quality, without taking into account the often differing number of documents shown for the argument stances (PRO or CON). Therefore, systems may retrieve relevant passages, but with a biased exposure of arguments. In this work, we analyze a range of non-stochastic fairness-aware ranking and diversity metrics to evaluate the extent to which argument stances are fairly exposed in argument retrieval systems.

Using the official runs of the argument retrieval task Touché at CLEF 2020, as well as synthetic data to control the amount and order of argument stances in the rankings, we show that systems with the best effectiveness in terms of topical relevance are not necessarily the most fair or the most diverse in terms of argument stance. The relationships we found between (un)fairness and diversity metrics shed light on how to evaluate group fairness – in addition to topical relevance – in argument retrieval settings.

CCS CONCEPTS

• Information systems → Evaluation of retrieval results.

KEYWORDS

evaluation; fairness; argument retrieval

ACM Reference Format:

Sachin Pathiyan Cherumanal, Damiano Spina, Falk Scholer, and W. Bruce Croft. 2021. Evaluating Fairness in Argument Retrieval. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21)*, November 1–5, 2021, Virtual Event, QLD, Australia. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3459637.3482099>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '21, November 1–5, 2021, Virtual Event, QLD, Australia

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8446-9/21/11...\$15.00

<https://doi.org/10.1145/3459637.3482099>

1 INTRODUCTION

The top search results in web search engines can be riddled with populism, conspiracy theories, and one-sidedness, not reflecting effectively the argumentative landscape for controversial topics (e.g., climate change) [22]. Recently, argument retrieval systems such as *args.me* [26] have been proposed. Argument retrieval systems aim to provide documents that cover both positive (PRO) and negative (CON) perspectives in search query results.

While argument retrieval systems aim to tackle the problem of providing multiple perspectives on user queries, the arguments that are retrieved for a given topic may only cover one point of view. For example, a query such as ‘is universal basic income good’ may return a list of relevant arguments supporting the statement but may expose fewer (or no) results with arguments against a universal basic income. Even when all stakeholders are identified and their stances on a given argumentative topic are summarized [22], the question arises whether rankings produced by such systems provide *fair* exposure of the different relevant stances for a user’s search on controversial topics [18]. Moreover, there is no consensus on how *fairness-aware argument retrieval* should be defined, and which metrics should be used to measure fairness in argument retrieval.

In this paper, we study the argument retrieval task – where systems aim to provide assistance to users searching for relevant PRO and CON arguments on various societal topics – proposed for the Touché Lab at CLEF 2020 [5] as a fairness-aware ranking problem. We explore fairness metrics for evaluating argument retrieval systems by defining the protected groups for this particular problem. Our results on the Touché submissions show that the retrieved results from an argument retrieval system are typically biased, and demonstrate that systems with the best effectiveness are not necessarily the most fair. Our analysis shows that fairness-aware rankings do not guarantee diverse results and vice-versa.

2 RELATED WORK

Fairness in Information Retrieval (IR). Algorithm fairness has been studied across multiple disciplines and has attracted attention from researchers in IR [2, 7, 27]. Most of these recent studies, regardless of the discipline [12, 20, 21], have considered *individual fairness* and *group fairness*. In this work we consider the group fairness scenario which requires that the demographics of those receiving a certain treatment is proportional to the demographics in the overall population – also known as statistical parity [13]. Most work on fairness in IR either measures the level of bias based on protected

attributes or aims to mitigate such biases using fairness-aware algorithms or optimization techniques [16, 24]. Our work is closely related to the former, as we measure group fairness in ranked outputs, but uses a slightly modified form of fairness measures that evaluate single-list rankings [28]. In a group fairness setting, a certain subset of the population is defined as a protected group. Most recent work uses characteristics such as race, gender, and disability status to define protected attributes, and therefore, protected groups [30]. For example, Geyik et al. [16] propose a fairness-aware ranking algorithm and fairness measures for talent search. Diaz et al. [9] and Biega et al. [4] use protected attributes such as age or gender to analyze fairness across distributions of rankings, rather than in a single fixed ranking. Our focus is on evaluating group fairness for argument retrieval in a Cranfield paradigm offline evaluation setting, and we apply the notion of protected groups to analyze unbiased exposure of different argument stances.

Argument Retrieval. An argument consists of a claim (e.g., “We should abandon fossil fuels”) and a premise that justifies either attacking or supporting the claim [11]. From an IR perspective, argument retrieval is the task of retrieving relevant supporting (PRO) and attacking (CON) justifications (premises) for a given query (claim). The motivation behind research in argument retrieval comes from the fact that search engines lack the capability to provide results that inform users about the premises and their stance (PRO or CON) towards the claim, as they currently ignore the argumentative nature of discussions in sources such as community question answering websites and debate portals [5]. An example of an argument search engine is *args.me* [25] which retrieves relevant arguments on a given (controversial) query from a focused collection of arguments crawled from a selection of debate portals.

3 FAIRNESS AND DIVERSITY METRICS

To study fairness in argument retrieval systems, we consider existing evaluation measures proposed to evaluate (un)fairness in rankings. As the task of balancing argument stances in a ranking can be seen as a particular case of search result diversification [10, 15], we also consider a commonly used diversity metric, α -nDCG [8].

3.1 (Un)fairness Metrics

We analyze three single-list metrics proposed by Yang and Stoyanovich [28]: normalized discounted difference (rND), normalized discounted KL-divergence (rKL), and normalized discounted ratio (rRD). We define these metrics using the notation described in Table 1, and explicitly indicate a cutoff k . Given that argument retrieval is a top-heavy ad-hoc retrieval task (e.g., the cutoff used at Touché challenge is $k = 5$), we adapt set-based fairness at discrete points ranging from top-1 to top-5.¹ These (un)fairness metrics measure the statistical parity in ranked outputs, i.e., they quantify the relative representation of the protected group S^+ in a ranking r . **Normalized Discounted Difference (rND@k).** rND@k computes the discounted difference between the proportion of arguments

¹The set-based fairness used by Yang and Stoyanovich [28] are defined at increments of 10 rank positions, ranging from 10 to N . The measures provide fairness values in the range $[0, 1]$, where 0 means most fair and 1 means most unfair. Therefore we refer to this family of metrics as (un)fairness metrics.

Table 1: Notation used in (un)fairness metrics.

Notation	Description
N	Number of relevant arguments in ground truth
r	A ranking
S^+	Protected group
S^-	Unprotected group

belonging to the protected group in the top- i subsets (represented by $\frac{|S^+_{1..i}|}{i}$) and the overall population ($\frac{|S^+|}{N}$):

$$\text{rND@k}(r) = \sum_{i=1}^k \frac{1}{\log_2(i+1)} \left| \frac{|S^+_{1..i}|}{i} - \frac{|S^+|}{N} \right| \quad (1)$$

Normalized Discounted KL-divergence (rKL@k). rKL@k calculates the Kullback-Leibler (KL) divergence to measure the distance between two probability distributions P and Q :

$$P = \left(\frac{|S^+_{1..i}|}{i}, \frac{|S^-_{1..i}|}{i} \right), Q = \left(\frac{|S^+|}{N}, \frac{|S^-|}{N} \right) \quad (2)$$

where P represents the proportion of the protected group until rank i in top- i and Q represents the proportion of the protected group in the overall population. Intuitively, rKL@5 measures the difference between the proportion of the protected group in the top- i ranks, and in the overall population:

$$\text{rKL@k}(r) = \sum_{i=1}^k \frac{D_{KL}(P||Q)}{\log_2(i+1)} \quad (3)$$

where $D_{KL}(P||Q)$ is the KL-divergence between P and Q .

Normalized Discounted Ratio (rRD@k). rRD@k computes the difference between the ratio of arguments belonging to the protected group to those that are not protected until rank i , and the ratio of the same in the overall population:

$$\text{rRD@k}(r) = \sum_{i=1}^k \frac{1}{\log_2(i+1)} \left| \frac{|S^+_{1..i}|}{|S^-_{1..i}|} - \frac{|S^+|}{|S^-|} \right| \quad (4)$$

As mentioned by Yang and Stoyanovich [28], when either the numerator or the denominator of a fraction is 0, the entire fraction is set to 0. After computation we apply min-max normalization to the three (un)fairness metrics.

3.2 Diversity Metric: α -nDCG

Exposing the different stances of an argument in a ranking can be seen as a search result diversification problem, i.e., rankings with diversified stances. Different diversity metrics have been proposed and analyzed in the literature [3, 23]. In our study, we consider α -nDCG proposed by Clarke et al. [8] to analyze the relationship between (un)fairness and diversity.

The parameter $\alpha \in [0, 1]$ controls the reward of diversity and novelty in the ranking. When $\alpha = 0$, α -nDCG is equivalent to nDCG, and the higher the α value, the more reward is obtained by diversifying the ranking—at the expense of reward obtained by relevance.




Figure 1: Permutations of argument stances across the top-5 positions in the generated synthetic data.

4 EXPERIMENTAL SETUP

4.1 Real Scenario

Test Collection. In our analysis, we use the test collection used at the Touché Lab at CLEF 2020, comprising 387,740 documents (arguments) from the args.me corpus [1], and a ground truth of 2,964 relevance judgments across 49 topics with relevant arguments graded in the range of 1 (least relevant) to 5 (most relevant).

Systems. We use the 21 official runs submitted to the Touché argument retrieval task. The systems cover a variety of retrieval, query expansion, and re-ranking techniques [5]. Note that neither fairness nor diversity were considered in the original setup of the challenge, and only topical relevance and argument quality needed to be optimized.

Protected Group. In this paper, we make the assumption that, for each topic, the stance with the lowest number of arguments judged as relevant is considered as the protected group. Protected group in the fairness literature is identified as a group that is disadvantaged/under-represented or a minority. Inspired from this definition, in this paper, we make the assumption that, for each topic, the stance with the lowest number of relevant arguments (i.e., minority) is considered as the protected group.

4.2 Controlled Scenario

To systematically study the behaviour of (un)fairness and diversity measures, we additionally created a synthetic test collection which consists of the 32 possible permutations of argument stances and rankings of length 5. Figure 1 shows the rankings generated for the controlled scenario. For the sake of simplicity, we only consider binary relevance in the synthetic data. Each permutation represents a topic and a ranking of 5 relevant arguments. The number of relevant judgments for each stance across topics vary according to the three definition of protected group described below.

Protected Groups. We consider three different definitions of protected group.

- *Minority Group.* For all topics, the proportion of PRO/CON arguments judged as relevant has a 4/1 ratio. The minority group (CON) is considered as protected S^+ .²
- *Proportion-Agnostic.* All topics have a 1/1 ratio of PRO/CON arguments judged as relevant in the ground truth. The protected group is then defined independently of the the proportion (e.g., a possible scenario would be that this decision is informed by a policy). We set PRO as the protected group, although choosing one stance or the other would lead to the same outcome in this case.

²Note that, as the distribution of PRO and CON across topics is symmetric, a ratio of 1/4 and choosing PRO as protected group would lead to the same results.




Figure 2: nDCG@5 (higher is better) and rKL@5 (un)fairness (lower is better) for official runs of the Touché argument retrieval task at CLEF 2020. Systems (x-axis) are shown in decreasing order of nDCG@5.

- *Majority Group.* For all topics, the proportion of PRO/CON arguments judged as relevant has a 4/1 ratio. The majority group (PRO) is considered as protected S^+ .

4.3 Evaluation Metrics

In our analysis we consider the official effectiveness evaluation of the Touché evaluation campaign, which only considers relevance, nDCG@5. Scores were computed using `trec_eval`. For (un)fairness metrics, we computed rND@5, rKL@5, and rRD@5 using our own implementation.³ To compute the diversity metric α -nDCG@5, we used `ndeal` from the TREC 2014 Web Track.⁴ Unless otherwise specified, we use the default parameter value $\alpha = 0.5$.

4.4 Meta-Evaluation

In order to analyze changes on ranking of systems, we compare evaluation measures using Kendall’s τ correlation [17].⁵ Kendall’s τ provides an intuitive way to compare the rankings of systems, i.e., $\frac{1-\tau}{2}$ % of system pairs are those with a different relative ordering in the two rankings. τ values range between $[-1, 1]$ and does not consider the magnitude of the item values but only their relative ordering. For the real scenario, scores are averaged across topics and system rankings are compared. For the controlled scenario, correlations are computed using the scores obtained for synthetic rankings.

5 RESULTS AND ANALYSIS

5.1 Real Scenario: Touché Lab at CLEF 2020

The effectiveness of runs submitted to Touché are shown in Figure 2, sorted by the official evaluation measure of the challenge: nDCG@5. Comparing nDCG@5 scores with rKL@5 we can see that rKL@5 scores do not increase monotonically. The same trend was observed for rND@5 and rRD@5. This suggests that the system rankings would be altered if both relevance and (un)fairness metrics are considered. However, it is worth noting that some of the ranked systems at Touché would remain unchanged, e.g., the top-3 systems are the

³The source code is available at <https://github.com/sachinpc1993/fair-arguments>.

⁴<https://github.com/trec-web/trec-web-2014/blob/master/src/eval/ndeal.c>

⁵We are aware of alternative ways of comparing rankings [6, 14, 19, 29]. As we are not defining an absolute threshold in correlations scores – and for the sake of interpretability of the results – we chose Kendall’s τ to compare evaluation measures.

Table 2: Kendall’s τ correlation between (un)fairness metrics, nDCG@5, and α -nDCG@5 for the Touché runs (real scenario). * indicates statistical significance ($p < 0.05$).

(un)fairness	nDCG@5	α -nDCG@5	rND@5	rKL@5
rND@5	-0.0762	-0.0667		
rKL@5	-0.2667	-0.3333*	0.3524*	
rRD@5	-0.2571	-0.3238*	-0.1143	0.2857

same if the harmonic mean of α -nDCG@5 and $1-rKL@5$ is used to rank the systems.

The Kendall’s τ correlations between the relevance, (un)fairness, and diversity metrics considered in our study are shown in Table 2. As expected, the (un)fairness metrics are negatively correlated with the relevance-based metric nDCG@5 (optimal scores are 0 and 1, respectively). The correlations are low and not statistically significant. α -nDCG@5 is more similar to rKL@5 and rRD@5, although the correlations are still weak. This suggests that while novelty and diversity are partially captured by these (un)fairness metrics, they still measure different dimensions. We also observed low correlations between rND@5 and the other metrics.

Meta-evaluation of the Touché challenge shows that: (i) diversity is related but not equivalent to (un)fairness; (ii) (un)fairness metrics lead to different system rankings; and (iii) (un)fairness metrics therefore provide complementary information to relevance-based metrics, and considering both would lead to changes in the leaderboard. From the real scenario, where different topics have different proportions of PRO/CON stances, it is not clear whether a particular (un)fairness metric is more appropriate than another. To understand how the positioning of the argument stance can affect the behavior of diversity and (un)fairness metrics, we next analyze these metrics in a controlled scenario with synthetic data.

5.2 Controlled Scenario: Synthetic Data

Kendall’s τ correlations between (un)fairness metrics and α -nDCG@5 are shown in Table 3 for three different definitions of protected groups. Here, we compare metrics over all possible permutations of depth 5 argument rankings – with all arguments in the rankings being equally relevant. The following observations can be made. rND@5 has a higher negative correlation with α -nDCG@5 for the *Minority* setting, but not for the other settings. The symmetric treatment of the groups S^+ and S^- in the rKL@5 measure is reflected in the correlations scores for the *Minority* and *Majority* settings.

It is also worth noting that α -nDCG@5 scores for the three settings are the same: this suggests that, if the definition of protected group is based on proportion of the population (e.g., statistical parity), α -nDCG@5 cannot be used. When α -nDCG@5 is calculated using $\alpha = 0.9$ (Table 4), higher negative correlations are obtained for rKL@5, but not for rND@5 and rRD@5. rRD@5 has the same correlation with α -nDCG@5 for different values of α (in both *Minority* and *Majority* settings), while the correlation is substantially lower for the *Proportion-Agnostic* setting.

A further examination of the individual topic scores showed that, for extreme cases where all the top-5 arguments belong to either S^+ or S^- (i.e., topics 1 and 32 in Figure 1), rKL@5 does not respond

Table 3: Kendall’s τ correlation between (un)fairness metrics and α -nDCG@5 diversity metric for different definitions of protected groups in the synthetic data (controlled scenario). * indicates statistical significance ($p < 0.05$).

(un)fairness	Protected Group S^+		
	Minority	Proportion-Agnostic	Majority
rND@5	-0.4812*	-0.1434	0.0000
rKL@5	-0.1747	-0.8378*	-0.1747
rRD@5	0.1032	-0.6924*	-0.2572

Table 4: Kendall’s τ correlation between (un)fairness metrics and α -nDCG@5 diversity metric with $\alpha = 0.9$. * indicates statistical significance ($p < 0.05$).

(un)fairness	Protected Group S^+		
	Minority	Proportion-Agnostic	Majority
rND@5	-0.4554*	-0.1520	0.0000
rKL@5	-0.2003	-1.0000*	-0.2003
rRD@5	0.1032	-0.5799*	-0.2572

as one would expect. For instance, in the *Minority* setting, rKL@5 shows low (un)fairness (0.1463) for the topics that has all arguments belonging to S^- , and high (un)fairness (1.0000) for the topic that has all S^+ . rND@5 shows similar but less skewed behavior, and the difference in rND@5 scores for symmetric topics is substantially lower than is observed for rKL@5. On the other hand, rRD@5 gives high (un)fairness scores for the extreme cases.

Our analysis in a controlled scenario shows that, although diversification in terms of argument stances may lead to a fair exposure of perspectives, diversity and (un)fairness metrics are not necessarily measuring the same dimensions.

6 CONCLUSION AND FUTURE WORK

To our knowledge, the notion of fair exposure of different stances in argument retrieval systems have not been explored before. We analyze the relation between relevance, group fairness, and diversity measures in the context of argument retrieval. Our analysis of the systems submitted to the Touché argument retrieval evaluation campaign shows that the outputs from argument retrieval systems can be biased towards a particular stance, and system rankings change when (un)fairness metrics are considered. Our correlation analysis using synthetic data corroborates that fairness-aware rankings do not guarantee diverse results and vice-versa.

In this work, we considered only a binary protected attribute (PRO and CON). Given that Normalized Discounted KL-divergence is a fairness measure capable of handling multi-level protected attributes, it is worth exploring how the complexity of such protected attributes affects the relationship between (un)fairness and diversity in future work. Stochastic fairness metrics [9] or other diversity measures [3, 15, 23] are yet to be investigated for argument retrieval. However, this requires re-purposing the existing test collections.

Acknowledgments. The authors thank the organizers of Touché@CLEF 2020 for making the runs available. This work is partially supported by the Australian Research Council (DE200100064, DP190101113).

REFERENCES

- [1] Yamen Ajjour, Henning Wachsmuth, Johannes Kiesel, Martin Potthast, Matthias Hagen, and Benno Stein. 2019. Data Acquisition for Argument Search: The args.me Corpus. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*. Springer, 48–59. https://doi.org/10.1007/978-3-030-30179-8_4
- [2] James Allan, Jaime Arguello, Leif Azzopardi, Peter Bailey, Tim Baldwin, Krisztian Balog, Hannah Bast, Nick Belkin, Klaus Berberich, Bodo von Billerbeck, Jamie Callan, Rob Capra, Mark Carman, Ben Carterette, Charles L. A. Clarke, Kevyn Collins-Thompson, Nick Craswell, W. Bruce Croft, J. Shane Culpepper, Jeff Dalton, Gianluca Demartini, Fernando Diaz, Laura Dietz, Susan Dumais, Carsten Eickhoff, Nicola Ferro, Norbert Fuhr, Shlomo Geva, Claudia Hauff, David Hawking, Hideo Joho, Gareth Jones, Jaap Kamps, Noriko Kando, Diane Kelly, Jaewon Kim, Julia Kiseleva, Yiqun Liu, Xiaolu Lu, Stefano Mizzaro, Alistair Moffat, Jian-Yun Nie, Alexandra Olteanu, Iadh Ounis, Filip Radlinski, Maarten de Rijke, Mark Sanderson, Falk Scholer, Laurianne Sitbon, Mark Smucker, Ian Soboroff, Damiano Spina, Torsten Suel, James Thom, Paul Thomas, Andrew Trotman, Ellen Voorhees, Arjen P. de Vries, Emine Yilmaz, and Guido Zuccon. 2018. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *SIGIR Forum* 52, 1 (Aug. 2018), 34–90. <https://doi.org/10.1145/3274784.3274788>
- [3] Enrique Amigó, Damiano Spina, and Jorge Carrillo-de Albornoz. 2018. An Axiomatic Analysis of Diversity Evaluation Metrics: Introducing the Rank-Biased Utility Metric. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) (SIGIR '18). Association for Computing Machinery, 625–634. <https://doi.org/10.1145/3209978.3210024>
- [4] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) (SIGIR '18). Association for Computing Machinery, 405–414. <https://doi.org/10.1145/3209978.3210063>
- [5] Alexander Bondarenko, Maik Fröbe, Meriem Beloucif, Lukas Gienapp, Yamen Ajjour, Alexander Panchenko, Chris Biemann, Benno Stein, Henning Wachsmuth, Martin Potthast, et al. 2020. Overview of Touché 2020: Argument Retrieval. In *International Conference of the Cross-Language Evaluation Forum for European Languages (CLEF '20)*. Springer, 384–395. https://doi.org/10.1007/978-3-030-58219-7_26
- [6] Ben Carterette. 2009. On Rank Correlation and the Distance between Rankings. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Boston, MA, USA) (SIGIR '09). Association for Computing Machinery, 436–443. <https://doi.org/10.1145/1571941.1572017>
- [7] Carlos Castillo. 2019. Fairness and Transparency in Ranking. *SIGIR Forum* 52, 2 (Jan. 2019), 64–71. <https://doi.org/10.1145/3308774.3308783>
- [8] Charles L.A. Clarke, Maheedhar Kolla, Gordon V. Cormack, Olga Vechtomova, Azin Ashkan, Stefan Büttcher, and Ian MacKinnon. 2008. Novelty and Diversity in Information Retrieval Evaluation. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Singapore, Singapore) (SIGIR '08). Association for Computing Machinery, 659–666. <https://doi.org/10.1145/1390334.1390446>
- [9] 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* (Virtual Event, Ireland) (CIKM '20). Association for Computing Machinery, 275–284. <https://doi.org/10.1145/3340531.3411962>
- [10] Tim Draws, Nava Tintarev, and Ujwal Gadiraju. 2021. Assessing Viewpoint Diversity in Search Results Using Ranking Fairness Metrics. *SIGKDD Explor. NewsL.* 23, 1 (May 2021), 50–58. <https://doi.org/10.1145/3468507.3468515>
- [11] Lorik Dumani, Patrick J Neumann, and Ralf Schenkel. 2020. A Framework for Argument Retrieval. In *Proceedings of the 42nd European Conference on Information Retrieval (ECIR '20)*. Springer, 431–445. https://doi.org/10.1007/978-3-030-45439-5_29
- [12] 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* (Cambridge, Massachusetts) (ITCS '12). Association for Computing Machinery, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [13] 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* (Cambridge, Massachusetts) (ITCS '12). Association for Computing Machinery, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [14] Nicola Ferro. 2017. What Does Affect the Correlation Among Evaluation Measures? *ACM Trans. Inf. Syst.* 36, 2, Article 19 (Aug. 2017), 40 pages. <https://doi.org/10.1145/3106371>
- [15] 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>
- [16] Sahin Cem Geyik, Stuart Ambler, and Krishnamurthy Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Anchorage, AK, USA) (KDD '19). Association for Computing Machinery, 2221–2231. <https://doi.org/10.1145/3292500.3330691>
- [17] Maurice G Kendall. 1938. A New Measure of Rank Correlation. *Biometrika* 30, 1/2 (1938), 81–93. <https://www.jstor.org/stable/2332226>
- [18] Johannes Kiesel, Damiano Spina, Henning Wachsmuth, and Benno Stein. 2021. The Meant, the Said, and the Understood: Conversational Argument Search and Cognitive Biases. In *CUI 2021 - 3rd Conference on Conversational User Interfaces* (Bilbao (online), Spain) (CUI '21). Association for Computing Machinery, Article 20, 5 pages. <https://doi.org/10.1145/3469595.3469615>
- [19] Mucahid Kutlu, Tamer Elsayed, Maram Hasanain, and Matthew Lease. 2018. When Rank Order Isn't Enough: New Statistical-Significance-Aware Correlation Measures. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) (CIKM '18). Association for Computing Machinery, 397–406. <https://doi.org/10.1145/3269206.3271751>
- [20] Dino Pedreschi, Salvatore Ruggieri, and Franco Turini. 2009. Measuring Discrimination in Socially-Sensitive Decision Records. In *Proceedings of the 2009 SIAM International Conference on Data Mining (SDM '09)*. SIAM, 581–592. <https://doi.org/10.1137/1.9781611972795.50>
- [21] Dino Pedreschi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-Aware Data Mining. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Las Vegas, Nevada, USA) (KDD '08). Association for Computing Machinery, 560–568. <https://doi.org/10.1145/1401890.1401959>
- [22] Martin Potthast, Lukas Gienapp, Florian Euchner, Nick Heilenkötter, Nico Weidmann, Henning Wachsmuth, Benno Stein, and Matthias Hagen. 2019. Argument Search: Assessing Argument Relevance. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR '19). Association for Computing Machinery, 1117–1120. <https://doi.org/10.1145/3331184.3331327>
- [23] Tetsuya Sakai and Zhaohao Zeng. 2019. Which Diversity Evaluation Measures Are "Good?". In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR '19). Association for Computing Machinery, 595–604. <https://doi.org/10.1145/3331184.3331215>
- [24] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (London, United Kingdom) (KDD '18). Association for Computing Machinery, 2219–2228. <https://doi.org/10.1145/3219819.3220088>
- [25] Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. 2017. Computational Argumentation Quality Assessment in Natural Language. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers (ACL '17)*. Association for Computational Linguistics, 176–187. <https://www.aclweb.org/anthology/E17-1017>
- [26] Henning Wachsmuth, Martin Potthast, Khalid Al-Khatib, Yamen Ajjour, Jana Puschmann, Jiani Qu, Jonas Dorsch, Viorel Morari, Janek Bevendorf, and Benno Stein. 2017. Building an Argument Search Engine for the Web. In *4th Workshop on Argument Mining (ArgMining 2017) at EMNLP*, Kevin Ashley, Claire Cardie, Nancy Green, Iryna Gurevych, Ivan Habernal, Diane Litman, Georgios Petasis, Chris Reed, Noam Slonim, and Vern Walker (Eds.). Association for Computational Linguistics, 49–59. <https://www.aclweb.org/anthology/W17-5106>
- [27] Colin Wilkie and Leif Azzopardi. 2014. Best and Fairest: An Empirical Analysis of Retrieval System Bias. In *Proceedings of the 36th European Conference on IR Research on Advances in Information Retrieval - Volume 8416* (Amsterdam, The Netherlands) (ECIR '14). Springer-Verlag, 13–25.
- [28] Ke Yang and Julia Stoyanovich. 2017. Measuring Fairness in Ranked Outputs. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management* (Chicago, IL, USA) (SSDBM '17). Association for Computing Machinery, Article 22, 6 pages. <https://doi.org/10.1145/3085504.3085526>
- [29] Emine Yilmaz, Javed A. Aslam, and Stephen Robertson. 2008. A New Rank Correlation Coefficient for Information Retrieval. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Singapore, Singapore) (SIGIR '08). Association for Computing Machinery, 587–594. <https://doi.org/10.1145/1390334.1390435>
- [30] Meike Zehlke, Ke Yang, and Julia Stoyanovich. 2021. Fairness in Ranking: A Survey. arXiv:2103.14000 [cs.IR]