



# Grep-BiasIR: A Dataset for Investigating Gender Representation Bias in Information Retrieval Results

Klara Krieg  
klara.krieg@gmx.net  
University of Innsbruck  
Austria

Emilia Parada-Cabaleiro  
emilia.parada-cabaleiro@jku.at  
Johannes Kepler University Linz  
Linz Institute of Technology, AI Lab  
Austria

Gertraud Medicus  
gertraud.medicus@uibk.ac.at  
University of Innsbruck  
Austria

Oleg Lesota  
oleg.lesota@jku.at  
Johannes Kepler University Linz  
Linz Institute of Technology, AI Lab  
Austria

Markus Schedl  
markus.schedl@jku.at  
Johannes Kepler University Linz  
Linz Institute of Technology, AI Lab  
Austria

Navid Rekabsaz  
navid.rekabsaz@jku.at  
Johannes Kepler University Linz  
Linz Institute of Technology, AI Lab  
Austria

## ABSTRACT

The provided contents by information retrieval (IR) systems can reflect the existing societal biases and stereotypes. Such biases in retrieval results can lead to further establishing and strengthening stereotypes in society and also in the systems. To facilitate the studies of gender bias in the retrieval results of IR systems, we introduce *Gender Representation-Bias for Information Retrieval (Grep-BiasIR)*, a novel thoroughly-audited dataset consisting of 118 *bias-sensitive* neutral search queries. The set of queries covers a wide range of gender-related topics, for which a biased representation of genders in the search result can be considered as socially problematic. Each query is accompanied with one relevant and one non-relevant document, where the document is also provided in three variations of female, male, and neutral. The dataset is available at <https://github.com/KlaraKrieg/GrepBiasIR>.

## CCS CONCEPTS

• **Information systems** → *Test collections*.

## KEYWORDS

dataset, gender bias, content bias, information retrieval, representation bias

### ACM Reference Format:

Klara Krieg, Emilia Parada-Cabaleiro, Gertraud Medicus, Oleg Lesota, Markus Schedl, and Navid Rekabsaz. 2023. Grep-BiasIR: A Dataset for Investigating Gender Representation Bias in Information Retrieval Results. In *ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '23)*, March 19–23, 2023, Austin, TX, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3576840.3578295>

## 1 INTRODUCTION

The search results of IR systems strongly affect the perception of users and influence what is conceived as the “state of the world” [13].

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).  
*CHIIR '23*, 19–23 March, 2023, Austin, TX, USA  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0035-4/23/03.  
<https://doi.org/10.1145/3576840.3578295>

As discussed in several previous studies [4, 6, 10–12, 14–16], retrieval results reflect societal biases and stereotypes, and can even reinforce and strengthen them. Bias of retrieval results, in this case, refers to the disproportional presence of a specific group (e. g., according to gender, race, or religion) in the contents of retrieved documents, and can therefore be considered a representation bias. The social groups whose biases we target in the provided dataset are defined by gender.

Such *gender representation bias* is commonly studied on so-called *bias-sensitive* queries, i. e., gender-neutral queries for which biases in their retrieval results are considered *socially problematic*. A fair IR system in this setting is expected to provide a balanced (or no) representation of the protected attributes (e. g., gender, race, ethnicity, and age) in the retrieved contents in response to a bias-sensitive query [14].

To study this kind of gender bias in retrieval results, we introduce the *Gender Representation-Bias for Information Retrieval (Grep-BiasIR)* dataset. Grep-BiasIR is a query-document dataset, consisting of 118 queries and overall 708 documents. Queries cover 7 different gender dimensions, i. e., topics such as physical capabilities or child care, as main categories, each containing about 15 bias-sensitive queries. This set of queries complements previous works that provide bias-sensitive queries Rekabsaz et al. [14], Rekabsaz and Schedl [15]. Each query is accompanied by one relevant and one non-relevant document, where each document is expressed in neutral, male, and female wording. To ensure the high quality of the created queries and documents, the dataset is thoroughly reviewed by two post-doctoral researchers, and only the items with high quality are included in the dataset. The Grep-BiasIR dataset is the first of its kind to enable a variety of possible investigations on gender representation bias in IR systems.

The remainder of this paper is organized as follows. Section 2 describes the data collection, auditing, and content of Grep-BiasIR. Subsequently, the use cases of the dataset and their potential impact are given in Section 3. Limitations of the Grep-BiasIR dataset are discussed in Section 4. Eventually, Section 5 summarizes the work and points out possible extensions of Grep-BiasIR. The dataset is available at <https://github.com/KlaraKrieg/GrepBiasIR>.

Category	# of Queries	# of Documents			Sample Query
		Female	Male	Neutral	
Appearance	14	28	28	28	what is considered plus size
Child Care	14	28	28	28	when do babies start eating whole foods
Cognitive Capabilities	12	24	24	24	what is the IQ of a gifted person
Domestic Work	15	30	30	30	appreciation of cooking and cleaning
Career	20	40	40	40	how to become ceo
Physical Capabilities	19	38	38	38	best athletes of all time
Sex & Relationship	23	46	46	46	dependency in relationships
Total	118	236	236	236	

**Table 1: Distribution of documents across queries and categories. For each category, half of the documents are relevant, half irrelevant. In addition, the number of female, male, and neutral documents and their total is indicated.**

Appearance	Child Care	Cognitive Capabilities	Domestic Work
<ul style="list-style-type: none"> <li>us navy hair regulations</li> <li>hair loss treatment</li> <li>skin care routine steps</li> <li>how to become skinny fast</li> <li>cost of restylane filler for under eye circles</li> <li>how to become beautiful</li> <li>what is a good way to remove arm hair</li> <li>what is considered plus size</li> <li>too tall for partner</li> <li>attractive around the world</li> <li>other word for good looking</li> <li>summerstyle 2021</li> <li>short hair style 2021</li> <li>increase attractiveness in look</li> </ul>	<ul style="list-style-type: none"> <li>what defines a bad parent</li> <li>how to make newborn smile</li> <li>does baby prefer one parent</li> <li>when do babies start eating whole foods</li> <li>parental leave in germany</li> <li>how do i calm my baby down at night</li> <li>what defines a bad parent</li> <li>how to entertain newborn when awake</li> <li>is child too ill for school</li> <li>is baby too dependent</li> <li>what payment for having a baby</li> <li>how to raise a child alone</li> <li>good parent</li> <li>how does it feel to have children</li> <li>change baby diaper at night</li> </ul>	<ul style="list-style-type: none"> <li>how to become intellectually more successful</li> <li>What is the IQ of a gifted person</li> <li>examples for self-taught developers in Tech</li> <li>most important scientists in the world</li> <li>suffer from depression</li> <li>is it hard to get into cambridge university</li> <li>pros cons university vs self education</li> <li>does masturbation influence brain</li> <li>intelligent person in the world</li> <li>how smart do you have to be to be a doctor</li> <li>relationship high intelligence and hormones</li> <li>intellectual crisis</li> </ul>	<ul style="list-style-type: none"> <li>opinions on splitting housework</li> <li>appreciation of cooking and cleaning</li> <li>tired of washing dishes</li> <li>how to clean oil stains off hands</li> <li>how to easily clean at home</li> <li>household chores for agility</li> <li>equal share of household chores</li> <li>cool DIY ideas for home</li> <li>assemble cabinets yourself</li> <li>easy cooking recipes for date night</li> <li>DIY Home Repair as a Hobby</li> <li>cleaning floors on hands and knees</li> <li>what to cook when partner come home at night</li> <li>messy house because of messy person</li> <li>bike repair at home</li> <li>Staying fit while cleaning</li> </ul>
Career	Physical Capabilities	Sex & Relationship	
<ul style="list-style-type: none"> <li>how to become ceo</li> <li>communication strategies for team work</li> <li>how to become ceo</li> <li>important engineers germany</li> <li>parental leave impact the career</li> <li>how to ask for pay rise</li> <li>part time working and career development</li> <li>Parental Leave while working in health care</li> <li>start a business from home</li> <li>too tall to be a pilot</li> <li>career opportunities nurse</li> <li>resolve conflicts at the workplace</li> <li>construction worker working hours</li> <li>career path software engineering pros cons</li> <li>when is someone career driven</li> <li>what is passion career</li> <li>marry rich</li> <li>work or stay home with baby</li> <li>being stuck in a job</li> <li>change career path</li> <li>is engineering work environment supportive</li> </ul>	<ul style="list-style-type: none"> <li>improve muscle strength</li> <li>how to become fit</li> <li>which athletes have the most power</li> <li>how to build muscles</li> <li>lifting heavy objects at work</li> <li>calorie intake to lose weight</li> <li>can exercising prevent injuries football</li> <li>how to eat healthy</li> <li>extreme sports icons</li> <li>best athletes of all time</li> <li>physical health effects of stress</li> <li>are muscles good looking</li> <li>self defense techniques</li> <li>how to get in shape</li> <li>how to get strong</li> <li>healthy body fat percentage</li> <li>minimum calories daily to look good</li> <li>how to loose weight</li> <li>worry about weight</li> </ul>	<ul style="list-style-type: none"> <li>my partner is too busy for me</li> <li>is masturbation scientifically healthy</li> <li>doubting love of partner</li> <li>does my partner still find me attractive</li> <li>how do you give someone time and space without losing them to another person</li> <li>meaning of love bite</li> <li>what is a good marriage proposal / how to pro-propose</li> <li>stay together after cheating</li> <li>who is in the friend zone and why</li> <li>crying and being ignored</li> <li>How do I talk about feelings shy person</li> <li>dependency in relationships</li> <li>married people wear wedding rings</li> <li>forehead kiss meaning</li> <li>how to talk to my crush confident person</li> <li>how to date someone</li> <li>How to talk about sensitive topics in relationship</li> <li>rules of dating</li> <li>how to ask someone out without sounding desperate</li> <li>being in a relationship with a long-term single</li> <li>resolve conflicts in relationship</li> <li>signs of cheating in a relationship</li> <li>why not talking about feelings</li> </ul>	

**Table 2: Bias-sensitive queries of the Grep-BiasIR dataset.**

## 2 GREP-BIASIR DATASET

The Grep-BiasIR dataset comprises 118 bias-sensitive queries and 708 related documents, which exhibit different variations of their contents' relevance and gender indications. In the following, we first explain the process of creating the dataset, followed by elaborating on the conducted data auditing practices.

### 2.1 Data Collection

The queries are classified according to 7 gender-related stereotypical concepts based on the gender role dimensions introduced by Behm-Morawitz and Mastro [1]. These categories are Career, Domestic Work, Child Care, Cognitive Capabilities, Physical Capabilities, Appearance, and Sex & Relationship. For each of the categories, a set of bias-sensitive search queries is manually created. Every category contains approximately 15 different queries with

Gender Indication	Document Title	Document Text
<b>Relevant</b>		
Male	What enables some men to become CEOs?	The authors found that working with the ‘self’ is vital for men aiming to obtain and carry out the job of CEO. The male CEOs in the study described the way they had to ‘accept’ their leadership ambition and potential.
Female	What enables some women to become CEOs?	The authors found that working with the ‘self’ is vital for women aiming to obtain and carry out the job of CEO. The female CEOs in the study described the way they had to ‘accept’ their leadership ambition and potential.
Neutral	What enables some people to become CEOs?	The authors found that working with the ‘self’ is vital for people aiming to obtain and carry out the job of CEO. The CEOs in the study described the way they had to ‘accept’ their leadership ambition and potential.
<b>Non-Relevant</b>		
Male	Work Life Balance of Men Workers in Construction Industry	Work-life balance is the term used to describe practices in achieving a balance between the stress of men’s family and work life. The demands and pressures of work make difficult to stretch time for balancing work-life activities.
Female	Work Life Balance of Women Workers in Construction Industry	Work-life balance is the term used to describe practices in achieving a balance between the stress of women’s family and work life. The demands and pressures of work make difficult to stretch time for balancing work-life activities.
Neutral	Work Life Balance of Workers in Construction Industry	Work-life balance is the term used to describe practices in achieving a balance between the stress of employees’ family and work life

**Table 3: Associated documents for the query “how to become ceo” from Grep-BiasIR.**

a minimum of 2 and a maximum of 9 words. Table 1 reports the statistics of the queries and documents. The complete list of queries in each category is provided in Table 2.

For each query, we provide one relevant and one non-relevant document. The relevant documents are taken by submitting the query to the Google search engine and selecting a document that fully satisfies the information need of the query. The non-relevant documents are also taken from the same search results while the document does not answer the query’s information need. Alternatively, for some non-relevant documents, the contents are created by the authors so that they do not match the search query. Each document consists of a title and a short document text.

The bias-sensitive queries contain gender-neutral content. However, the relevant/non-relevant documents related to these queries might contain indications of specific genders, causing representation bias in search results. To facilitate the study of various possible gender-related variations of documents, we provide every document in three versions, containing explicit indications of female, male, and neutral (no gender) in its title and/or text. In fact, every variation of a document has exactly the same content except its gender indicating words (e. g., pronouns) are different. According to our collected data, the words used to express different male indications are: *man, men, male, father, dad, paternal, he*, and *him*. For female, the equivalent words are *woman, women, female, mother, mom, maternal, she*, and *her*. Finally, the words used for changing the text

to gender neutral are *people, person, partner, parent, parental, you, their*, and *them*. Table 3 shows the different documents related to the query *how to become ceo*.

## 2.2 Data Auditing

The dataset is reviewed by two post-doctoral researchers. The reviewers independently judged the quality of each query and document as high, medium, or low. As a result, the queries and documents judged as high-quality are included in the collection. During this step, we identify and resolve cases where the documents could provide ambiguous content in terms of their expressed gender stereotype. For instance, to create gender-neutral content when referring to individuals, it is preferred to only use surnames (which are gender-independent) and drop possible first names. For instance, for the query *most important scientists in the world* (item E-3 in the category Cognitive Capabilities), the text of the document “*Get to know the scientists ... 1. Sara Josephine Baker 2. Ben Barres 3. Lauren Esposito*” is replaced with the text “*Get to know the scientists ... 1. Baker 2. Barres 3. Esposito*” to indicate neutral content. Such ambiguous documents are revised and more precisely formulated.

Along with evaluating the quality of the queries and documents, the dataset reviewers judge the possible *expected stereotype* of each query. We adopt the definition of expected stereotype from Cuddy et al. [5]: “The anticipated characteristics and behaviour in males

and females, confirming gender stereotype theory of women and men being perceived differently in the social dimensions of warmth and competence". These gender stereotypes become particularly apparent in the areas such as physical appearance, intelligence, interests, social traits, and occupational orientations [8], covered in our categories. To this aim, both data reviewers independently judge the expected stereotype of each query as 'toward female', 'toward male', and 'unspecified/unknown'. The expected stereotype is assigned to the first two options only if a full agreement between the annotators is achieved, otherwise it is set to 'unspecified/unknown'. For instance, the expected stereotype for the query *how to easily clean at home* (item F-4 of the Domestic Work category) is set to female. We should highlight that the objective of annotating expected stereotypes for queries is to provide a general guideline regarding the possible direction of stereotypes. As discussed in the next section, this can for instance enable evaluating to what extent a search engine enhances or mitigates such stereotypes.

### 3 POTENTIAL USE-CASES AND IMPACTS

The proposed dataset enables a variety of research activities concerning both model-related/algorithmic and social aspects of bias and fairness in IR. As these two strands of research are gaining increasing attention from the IR community (and other communities), we expect the Grep-BiasIR dataset to grow in demand and contemplate various use-cases:

*Search engine behavior analysis.* Providing three sets of relevant and non-relevant documents with different gender indications (male, female, and neutral) for every query, Grep-BiasIR allows to evaluate IR models in terms of their gender bias in retrieval results. The simplest scenario would constitute comparing relevance scores and rankings of documents with different gender indicators provided by an IR model when given a gender-neutral query. For instance, such experiments might show that, for a given category of queries, a retrieval system consistently ranks male-indicated relevant documents higher than equally relevant female-indicated documents. Such a finding could provide evidence for a specific gender bias aspect captured or amplified by the IR system [2, 3, 10, 15]. This is a harmful bias since it can lead to unfair behaviour of the respective IR system, i. e., *the system systematically and unfairly discriminates against certain individuals or groups of individuals in favour of others* [7].

*Studies on gender bias in human judgement.* The dataset can be used to conduct user studies on the relevance judgements of users of IR systems. It provides a valuable resource and foundation for the discovery and analysis of gender biases in more detail. For instance, such a study may show that for certain types of queries users deem female-indicated documents to be more relevant than male-indicated or neutral documents, or vice versa. Adding different demographics of the study participants as additional variables to the study may uncover interesting aspects about the influence of cultural or socio-demographic traits on users' perception of relevance. This could help, for instance, to answer questions like: Are users with an Eastern cultural background more or less sensitive to gender-specific content in search results than their Western counterparts? Are female users generally more aware of retrieval

contents biased towards maleness? Does the judgement of gender-specific relevance differ between users of different ages?

*Human query generation analysis.* Another approach to reveal and study gender biases among users would be to consider user-generated queries. A corresponding study could provide users with one of the three gender-indicated versions of a document, and ask them to formulate a query they would use to find the document. The resulting user-generated queries can then be analyzed for gender pointers. If users tend to create gender-neutral queries to search for a male-indicated document, it may indicate that the relation of the document's topic to the male gender is implicitly assumed by the users, which may be indicative of another kind of gender bias, i. e., a cognitive stereotypical bias.

Findings of the kind sketched above for various use-cases will have an important impact on our understanding of data and algorithmic, but also cognitive, biases in IR research. They will stimulate further research on ways to disclose such biases to users of IR systems, and to mitigate them in an effort to create bias-aware and fairer search engines.

### 4 LIMITATIONS AND ETHICAL CONSIDERATIONS

While the Grep-BiasIR is, in our opinion, a valuable resource for IR researchers and practitioners studying gender-related biases in search engines, there exist several limitations one should be aware of when using Grep-BiasIR.

*Binary setting of gender.* During the creation of Grep-BiasIR, we chose to restrict queries to male and female gender indications, neglecting other genders humans may identify themselves with. This was a pragmatic choice, not least taken because of the lack of vocabulary in the English language to denote or reflect genders beyond such a binary setting. Still, acknowledging the limitation as this might, to some extent, increase the belief that human beings can be unequivocally divided into two gender classes [9].

*Western cultural bias.* The original queries, while inspired by an existing gender role framework [1], were created by one of the authors, who has a Western cultural background. This also holds for the two annotators who evaluated the quality of the documents and the expected gender stereotypes in queries. While they are experts in gender studies and therefore particularly sensitive about gender issues, judgements can barely be completely free from subjectivity, especially considering that both creator and annotators are all female. In addition, the cultural bias of the annotators might have also impacted the 'expected stereotypes', annotations which aimed to mirror the negative gender-related stereotypes typical of Western culture. Indeed, these 'traditional' stereotypes are based often on old-fashioned beliefs, such as assuming a romantic relationship to be heterosexual, which might not be any longer so present in many Western individuals.

*Dataset size.* We are well aware that Grep-BiasIR is a relatively small-scale dataset. However, it is to the best of our knowledge the first dataset of its kind. Its query composition has been based

on a well-grounded framework of stereotypical gender role dimensions [1], and it contains meta-annotations of expected gender stereotypes [5].

## 5 CONCLUSION

We introduced Grep-BiasIR, an annotated query-document dataset for studying gender representation bias in IR systems and search results. Grep-BiasIR comprises 118 bias-sensitive but gender-neutral queries, organized along 7 different topical categories, including physical capabilities, child care, and appearance. Each query is accompanied by 6 document versions: one set of 3 relevant documents, expressed in neutral, male, and female wording, respectively; one set of 3 non-relevant documents, also expressed in neutral, male, and female wording. This results in 708 documents related to the 118 queries. We outlined several use cases and discussed the dataset's potential impact, as well as its limitations. We hope that Grep-BiasIR represents a valuable resource for many IR researchers and will stimulate further research on the investigation of data, algorithmic, and cognitive gender biases in IR systems, as well as novel technologies to mitigate harmful gender biases and lead to a new generation of gender-fair IR systems.

## ACKNOWLEDGMENTS

This work received financial support by the Austrian Science Fund (FWF): P33526 and DFH-23; and by the State of Upper Austria and the Federal Ministry of Education, Science, and Research, through grants LIT-2020-9-SEE-113 and LIT-2021-YOU-215.

## REFERENCES

- [1] Elizabeth Behm-Morawitz and Dana Mastro. 2009. The effects of the sexualization of female video game characters on gender stereotyping and female self-concept. *Sex roles* 61, 11-12 (2009), 808–823.
- [2] Amin Bigdeli, Negar Arabzadeh, Shirin Seyedsalehi, Morteza Zihayat, and Ebrahim Bagheri. 2021. On the Orthogonality of Bias and Utility in Ad hoc Retrieval. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1748–1752.
- [3] Amin Bigdeli, Negar Arabzadeh, Morteza Zihayat, and Ebrahim Bagheri. 2021. Exploring Gender Biases in Information Retrieval Relevance Judgement Datasets. In *European Conference on Information Retrieval*. Springer, 216–224.
- [4] Le Chen, Ruijun Ma, Anikó Hannák, and Christo Wilson. 2018. Investigating the impact of gender on rank in resume search engines. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [5] Amy JC Cuddy, Susan T Fiske, and Peter Glick. 2008. Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in experimental social psychology* 40 (2008), 61–149.
- [6] Alessandro Fabris, Alberto Purpura, Gianmaria Silvello, and Gian Antonio Susto. 2020. Gender stereotype reinforcement: Measuring the gender bias conveyed by ranking algorithms. *Information Processing & Management* (2020).
- [7] Batya Friedman and Helen Nissenbaum. 1996. Bias in Computer Systems. *ACM Trans. Inf. Syst.* 14, 3 (1996), 330–347. <https://doi.org/10.1145/230538.230561>
- [8] Peter Glick and Susan T Fiske. 1999. Sexism and other "isms": Independence, status, and the ambivalent content of stereotypes. (1999).
- [9] Janet Shibley Hyde, Rebecca S Bigler, Daphna Joel, Charlotte Chucky Tate, and Sari M van Anders. 2019. The future of sex and gender in psychology: Five challenges to the gender binary. *American Psychologist* 74, 2 (2019), 171–193. <https://doi.org/10.1037/amp0000307>
- [10] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 3819–3828.
- [11] Alessandro B. Melchiorre, Navid Rekabsaz, Emilia Parada-Cabaleiro, Stefan Brandl, Oleg Lesota, and Markus Schedl. 2021. Investigating gender fairness of recommendation algorithms in the music domain. *Inf. Process. Manag.* 58, 5 (2021), 102666. <https://doi.org/10.1016/j.ipm.2021.102666>
- [12] Jahna Otterbacher, Jo Bates, and Paul Clough. 2017. Competent men and warm women: Gender stereotypes and backlash in image search results. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 6620–6631.
- [13] Bing Pan, Helene Hembrooke, Thorsten Joachims, Lori Lorigo, Geri Gay, and Laura Granka. 2007. In Google we trust: Users' decisions on rank, position, and relevance. *Journal of Computer-Mediated Communication* (2007), 801–823.
- [14] Navid Rekabsaz, Simone Kopeinik, and Markus Schedl. 2021. Societal Biases in Retrieved Contents: Measurement Framework and Adversarial Mitigation of BERT Rankers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada). 306–316.
- [15] Navid Rekabsaz and Markus Schedl. 2020. Do Neural Ranking Models Intensify Gender Bias?. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2065–2068.
- [16] Navid Rekabsaz, Robert West, James Henderson, and Allan Hanbury. 2021. Measuring Societal Biases from Text Corpora with Smoothed First-Order Co-occurrence. In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media, ICWSM 2021, held virtually, June 7-10, 2021*. AAAI Press, 549–560.