

# **Fairness in Information Access Systems**

**Other titles in Foundations and Trends® in Information Retrieval**

*Deep Learning for Dialogue Systems: Chit-Chat and Beyond*

Rui Yan, Juntao Li and Zhou Yu

ISBN: 978-1-63828-022-4

*Search Interface Design and Evaluation*

Chang Liu, Ying-Hsang Liu, Jingjing Liu and Ralf Bierig

ISBN: 978-1-68083-922-7

*Psychology-informed Recommender Systems*

Elisabeth Lex, Dominik Kowald, Paul Seitlinger, Thi Ngoc Trang Tran,  
Alexander Felfernig and Markus Schedl

ISBN: 978-1-68083-844-2

*Search and Discovery in Personal Email Collections*

Michael Bendersky, Xuanhui Wang, Marc Najork and Donald Metzler

ISBN: 978-1-68083-838-1

*Extracting, Mining and Predicting Users' Interests from Social Media*

Fattane Zarrinkalam, Stefano Faralli, Guangyuan Piao and Ebrahim Bagheri

ISBN: 978-1-68083-738-4

# Fairness in Information Access Systems

---

**Michael D. Ekstrand**

People and Information Research Team (PIReT)  
Boise State University  
ekstrand@acm.org

**Anubrata Das**

School of Information  
University of Texas at Austin  
anubrata@utexas.edu

**Robin Burke**

Department of Information Science  
University of Colorado  
robin.burke@colorado.edu

**Fernando Diaz**

Mila - Quebec AI Institute  
diazf@acm.org

**now**

the essence of knowledge

Boston — Delft

## Foundations and Trends<sup>®</sup> in Information Retrieval

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
United States  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto:sales@nowpublishers.com)

*Outside North America:*

now Publishers Inc.  
PO Box 179  
2600 AD Delft  
The Netherlands  
Tel. +31-6-51115274

The preferred citation for this publication is

M. D. Ekstrand *et al.*. *Fairness in Information Access Systems*. Foundations and Trends<sup>®</sup> in Information Retrieval, vol. 16, no. 1-2, pp. 1–177, 2022.

ISBN: 978-1-63828-041-5

© 2022 M. D. Ekstrand *et al.*

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: [www.copyright.com](http://www.copyright.com)

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; [www.nowpublishers.com](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, [www.nowpublishers.com](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

# Foundations and Trends<sup>®</sup> in Information Retrieval

## Volume 16, Issue 1-2, 2022

### Editorial Board

#### Editors-in-Chief

**Yiqun Liu**

Tsinghua University  
China

**Diane Kelly**

University of Tennessee  
USA

#### Editors

Barbara Poblete

*University of Chile*

Claudia Hauff

*Delft University of Technology*

Ellen M. Voorhees

*National Institute of Standards and  
Technology*

Hang Li

*Bytedance Technology*

Isabelle Moulinier

*Capital One*

Jaap Kamps

*University of Amsterdam*

Jimmy Lin

*University of Waterloo*

Leif Azzopardi

*University of Glasgow*

Lorraine Goeriot

*Université Grenoble Alpes*

Lynda Tamine

*University of Toulouse*

Maarten de Rijke

*University of Amsterdam and Ahold  
Delhaize*

Rodrygo Luis Teodoro Santos

*Universidade Federal de Minas Gerais*

Ruihua Song

*Renmin University of China*

Shane Culpepper

*RMIT*

Soumen Chakrabarti

*Indian Institute of Technology*

Xiangnan He

*University of Science and Technology of  
China*

Xuanjing Huang

*Fudan University*

Yubin Kim

*Etsy*

Zi Helen Huang

*University of Queensland*

## Editorial Scope

### Topics

Foundations and Trends® in Information Retrieval publishes survey and tutorial articles in the following topics:

- Applications of IR
- Architectures for IR
- Collaborative filtering and recommender systems
- Cross-lingual and multilingual IR
- Distributed IR and federated search
- Evaluation issues and test collections for IR
- Formal models and language models for IR
- IR on mobile platforms
- Indexing and retrieval of structured documents
- Information categorization and clustering
- Information extraction
- Information filtering and routing
- Metasearch, rank aggregation and data fusion
- Natural language processing for IR
- Performance issues for IR systems, including algorithms, data structures, optimization techniques, and scalability
- Question answering
- Summarization of single documents, multiple documents, and corpora
- Text mining
- Topic detection and tracking
- Usability, interactivity, and visualization issues in IR
- User modelling and user studies for IR
- Web search

### Information for Librarians

Foundations and Trends® in Information Retrieval, 2022, Volume 16, 5 issues. ISSN paper version 1554-0669. ISSN online version 1554-0677. Also available as a combined paper and online subscription.

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	Abstracting Information Access . . . . .	6
1.2	A Brief History of Fairness . . . . .	7
1.3	Fairness and Bias . . . . .	9
1.4	Fairness and Other Responsibility Concerns . . . . .	10
1.5	Running Examples . . . . .	11
1.6	How to Use This Monograph . . . . .	14
1.7	Our Perspective . . . . .	15
1.8	Some Cautions . . . . .	16
<b>2</b>	<b>Information Access Fundamentals</b>	<b>21</b>
2.1	System Overview . . . . .	22
2.2	Repository of Items . . . . .	23
2.3	Users and Information Needs . . . . .	25
2.4	Presentation . . . . .	26
2.5	Evaluation . . . . .	28
2.6	Algorithmic Foundations . . . . .	32
<b>3</b>	<b>Fairness Fundamentals</b>	<b>38</b>
3.1	Sources of Unfairness . . . . .	41
3.2	Problems and Concepts . . . . .	47
3.3	Mitigation Methods . . . . .	56
3.4	Wrapping Up . . . . .	59

<b>4</b>	<b>The Problem Space</b>	<b>61</b>
4.1	What Breaks in Information Access? . . . . .	64
4.2	Kinds of Harms in Information Access . . . . .	68
4.3	Fair for Who? . . . . .	70
4.4	Fair How? . . . . .	78
4.5	Fair from What Vantage Point? . . . . .	84
4.6	Fair on What Time Scale? . . . . .	86
4.7	Fairness and the System Pipeline . . . . .	87
4.8	Fairness and Other Concerns . . . . .	89
4.9	Contributing Back to ML Fairness . . . . .	91
4.10	Navigating the Problem Space . . . . .	92
<b>5</b>	<b>Consumer Fairness</b>	<b>93</b>
5.1	Individual Fairness . . . . .	94
5.2	Group Fairness through Disaggregated Utility . . . . .	95
5.3	Disparate Effectiveness . . . . .	97
5.4	Providing Fair Utility . . . . .	99
5.5	Fairness Beyond Accuracy . . . . .	100
5.6	More Complex Scenarios . . . . .	102
<b>6</b>	<b>Provider Fairness</b>	<b>103</b>
6.1	Provider Representation . . . . .	105
6.2	Provider Exposure and Utility . . . . .	112
6.3	Fair Accuracy and Pairwise Fairness . . . . .	119
6.4	Related Problem: Subject Fairness . . . . .	121
<b>7</b>	<b>Dynamic Fairness</b>	<b>123</b>
7.1	Feedback Loops . . . . .	124
7.2	Dynamic Evaluation . . . . .	126
7.3	Opportunities in Feedback Loops . . . . .	126
<b>8</b>	<b>Next Steps for Fair Information Access</b>	<b>127</b>
8.1	Directions for Future Research . . . . .	128
8.2	Recommendations for Studying Fairness . . . . .	131
8.3	Concluding Remarks . . . . .	133
	<b>Acknowledgements</b>	<b>134</b>



<b>Appendix</b>	<b>135</b>
<b>A Resources for Fair Information Access</b>	<b>136</b>
A.1 Data Sets . . . . .	136
A.2 Software . . . . .	137
<b>References</b>	<b>138</b>
<b>Index</b>	<b>175</b>

# Fairness in Information Access Systems

Michael D. Ekstrand<sup>1</sup>, Anubrata Das<sup>2</sup>, Robin Burke<sup>3</sup> and Fernando Diaz<sup>4</sup>

<sup>1</sup>*People and Information Research Team (PIReT), Boise State University, USA; ekstrand@acm.org*

<sup>2</sup>*School of Information, University of Texas at Austin, USA; anubrata@utexas.edu*

<sup>3</sup>*That Recommender Systems Lab, Department of Information Science, University of Colorado, USA; robin.burke@colorado.edu*

<sup>4</sup>*Mila - Quebec AI Institute, Canada; diaz@acm.org*

---

## ABSTRACT

Recommendation, information retrieval, and other information access systems pose unique challenges for investigating and applying the fairness and non-discrimination concepts that have been developed for studying other machine learning systems. While fair information access shares many commonalities with fair classification, there are important differences: the multistakeholder nature of information access applications, the rank-based problem setting, the centrality of personalization in many cases, and the role of user response all complicate the problem of identifying precisely what types and operationalizations of fairness may be relevant.

In this monograph, we present a taxonomy of the various dimensions of fair information access and survey the literature to date on this new and rapidly-growing topic. We

---

Michael D. Ekstrand, Anubrata Das, Robin Burke and Fernando Diaz (2022), “Fairness in Information Access Systems”, Foundations and Trends® in Information Retrieval: Vol. 16, No. 1-2, pp 1–177. DOI: 10.1561/1500000079.

©2022 M. D. Ekstrand *et al.*

preface this with brief introductions to information access and algorithmic fairness to facilitate the use of this work by scholars with experience in one (or neither) of these fields who wish to study their intersection. We conclude with several open problems in fair information access, along with some suggestions for how to approach research in this space.

---

## List of Key Terms

---

<i>Term</i>	<i>Defined in</i>	<i>Page</i>
bias	1.3	9
disparate impact	3.2.3	51
disparate mistreatment	3.2.3	52
disparate treatment	3.2.3	50
fairness	1.3	9
group fairness	3.2.3	50
individual fairness	3.2.2	48
information access	1.1	7
information need	2.3	25

The index provides a more comprehensive cross-reference of terms used in this monograph.

# 1

---

## Introduction

---

As long as humans have recorded information in durable form, they have needed tools to access it: to locate the information they seek, review it, and consume it. Digitally, tools to facilitate information access take a variety of forms, including information retrieval and recommendation systems; these tools have been powered by technologies built on various paradigms, from heuristic metrics and expert systems to deep neural networks with sophisticated rank-based objective functions. Fundamentally, these technologies take a user's *information need* (an explicit and/or implicit need for information for some purpose (Kuhlthau, 1993), such as filling in knowledge or selecting a product) and locate documents or items that are *relevant* (that is, will meet the user's need).

Throughout the history of these technologies — which we treat under the integrated banner of **information access systems** — both research and development have been concerned with a range of effects beyond a system's ability to locate individual items that are relevant to a user's information need. Research has examined the *diversity* and *novelty* of results (Santos *et al.*, 2015; Hurley and Zhang, 2011) and the *coverage* of the system, among other concerns. In recent years, this concern has extended to the *fairness* of an information access system: are

the benefits and resources it provides fairly allocated between different people or organizations it affects? Does it introduce or reproduce harms, particularly harms distributed in an unfair or unjust way? This challenge is connected to the broader set of research on fairness in sociotechnical systems generally and AI systems more particularly (Mitchell *et al.*, 2020; Barocas *et al.*, 2019), but information access systems have their own set of particular challenges and possibilities.

Fairness is not an entirely new concern for information access; various fairness problems can be connected to topics with long precedent in the information retrieval and recommender systems literature. In the context of information retrieval, Friedman and Nissenbaum (1996) and Introna and Nissenbaum (2000) recognized the potential for search engines to embed social, political, and moral values in their ranking functions. In order to assess the impact of such values, Mowshowitz and Kawaguchi (2002) developed a metric to measure a search engine's deviation from an ideal exposure of content. Although conversations often focus on bias in algorithmic ranking, Vaughan and Zhang (2007) and Vaughan and Thelwall (2004) note that bias can be introduced because of biased crawling and indexing; in particular, they describe, writing in the 2000s, how Chinese webpages were under-indexed by search engines. These observations led to discussion amongst legal scholars about the regulation of search engines (Goldman, 2005; Pasquale, 2006). Azzopardi and Vinay (2008) proposed the notion of document *retrievability* and investigated the skew in this distribution for different retrieval systems. Work on *popularity bias* (Celma and Cano, 2008; Zhao *et al.*, 2013; Cañamares and Castells, 2018) and rich-get-richer effects (Cho *et al.*, 2005), along with attempts to ensure quality and equity in *long-tail recommendations* (Ferraro, 2019), can be viewed as a type of fairness problem: the system should not inordinately favor popular, well-known, and possibly well-funded content creators. In a group recommendation, one common objective is to ensure that the various members of a group are treated fairly (Kaya *et al.*, 2020).

The work on fair information access that we present here goes beyond these problems to examine how various forms of unfairness — particularly those that arise from *social biases* (Olteanu *et al.*, 2019) — can make their way in to the data, algorithms, and outputs of informa-

tion access systems. These biases can affect many different stakeholders of an information access system; Burke (2017) distinguishes between *provider-* and *consumer-*side fairness, and other individuals or organizations affected by an information access system may have further fairness concerns.

In this monograph, we provide an introduction to fairness in information access, aiming to give students, researchers, and practitioners a starting point for understanding the problem space, the research to date, and a foundation for their further study. Fairness in information access draws heavily from the fair machine learning literature, which we summarize in Section 3; researchers and practitioners looking to study or improve the fairness of information access will do well to pay attention to a broad set of research results. For reasons of scope, we are primarily concerned here with the fairness of the information access transaction itself: providing results in response to a request encoding an information need. Fairness concerns can also arise in other aspects of the system, such as the representation and presentation of documents themselves, or in support facilities such as query suggestions (Noble, 2018). We provide brief pointers on these topics, but a detailed treatment is left for future synthesis, noting that they have not yet received as much attention in the research literature. We are also specifically concerned with fairness-related harms, and not the broader set of harms that may arise in information access such as the amplification of disinformation.

Throughout this work, we use the term **system** to describe an algorithmic system that performs some task: retrieving information, recommending items, classifying or scoring people based on their data. These systems are embedded in social contexts, operating on human-provided inputs and producing results acted upon by humans. The technical system forms one part of a broader socio-technical system.

## 1.1 Abstracting Information Access

Our choice to title this monograph “Fairness in *Information Access*” is quite deliberate. While there is significant technical and social overlap between information retrieval, recommender systems, and related fields, they are distinct communities with differences in terminology, problem

definitions, and evaluation practices. However, there are fundamental commonalities, and they present many of the same problems that complicate notions of fairness, including ranked outputs, personalized relevance, repeated decision-making, and multistakeholder structure. We therefore refer to them together as **information access systems** — algorithmic systems that mediate the interaction between a repository of documents or items and a user’s information need.

This information access umbrella includes information retrieval, recommender systems, information filtering, and some applications of natural language processing. In Section 2, we present a fuller treatment of this integration and reviews the fundamentals of information access, both to introduce the concepts to readers who come to this paper from a general fairness background and to lay out consistent terminology for our readers from information retrieval or recommender systems backgrounds.

## 1.2 A Brief History of Fairness

In the pursuit of fairness in algorithmic systems and the society more generally, the authority of Aristotle’s citation of Plato “treat like cases alike” is a key touchstone: a normative requirement that those who are equal before the law should receive equal treatment (Gosepath, 2011). In more recent scholarship, the study of distributive welfare extends these concepts considerably, recognizing four distinct concepts of fairness: “exogenous rights, compensation, reward, and fitness.” (Moulin, 2004). *Exogenous rights*, as the term suggests, relate to external claims that a system must satisfy: equal shares in property as defined by contract, for example, or equality of political rights in democratic societies. *Compensation* recognizes that fairness may require extra consideration for parties where costs are unequal — affirmative action in hiring and college admissions are well-known examples. *Reward* justifies inequality on the basis of differing contributions: for example, increased bonuses to employees with greater contribution to the bottom line. Finally, we have *fitness*, the most nebulous category, and the one that many information access systems inhabit. The fitness principle holds that goods be distributed to those most fit to use, appreciate, or



derive benefit from them. It is an efficiency principle, where the fairest use is the one that allocates goods where the distribution achieves the maximum utility. Fitness has a natural application to information access, as we seek to locate documents and make them visible based on their utility to the user's information need.

U.S. legal theory has developed a rich tradition of anti-discrimination law, aimed at ensuring that people are not denied certain benefits (housing, work, education, financial services, etc.) on the basis of **protected characteristics** (race, color, religion, gender, disability, age, and in many jurisdictions, sexual orientation). It has given rise to several important concepts, such as the **disparate impact** standard (the idea that an allegedly discriminatory practice can be legally challenged on the grounds that it has disproportionate adverse impact on a protected group, without needing to show intent to discriminate<sup>1</sup>). Crenshaw (1989) points out some of the limitations of this legal framework; in particular, it has often focused on discrimination on the basis of *individual* protected characteristics, and people who have suffered harm as a result of combinations of protected characteristics (e.g. Black women being denied promotions given to both Black men and White women) have difficulty proving their case and obtaining relief. This theory of particular harms deriving from combinations of characteristics is called **intersectionality**.

Questions of fairness and discrimination have been the subject of significant discussion in many other communities as well. Educational testing, for example, has several decades of research on the fairness of various testing and assessment instruments; this history is summarized for computer scientists by Hutchinson and Mitchell (2019). Friedman and Nissenbaum (1996) provide one of the earlier examples of addressing questions of bias in computer science, pointing out how even seemingly-innocuous technical decisions may result in biased effects when a computing system is used in its social context. The last ten years have seen significant new activity on fairness in machine learning that

---

<sup>1</sup>Disparate impact is not sufficient basis to *win* a discrimination lawsuit; rather, it is the first step in a multi-stage burden-shifting framework used to decide discrimination cases under the standard. Barocas and Selbst (2016) provide an overview of the process.

forms the primary stream of algorithmic fairness research; in Section 3 we provide an introduction to this literature.

### 1.3 Fairness and Bias

There are many overlapping terms used to discuss issues of fairness, bias, and discrimination. While we give a fuller treatment of the vocabulary in Section 3, we will here introduce how we use these terms in this monograph. Work we cite may use them differently.

When we refer to **fairness**, we are talking about the ways a system treats people, or groups of people, in a way that is considered “unfair” by some moral, legal, or ethical standard. This is typically through effects or impacts that are not experienced in an equitable way, but can sometimes arise through the system’s internal operation or representations. This definition is similar to how Friedman and Nissenbaum (1996) use the term “bias”. There is not one particular definition of what constitutes fairness, as Selbst *et al.* (2019) and many others have noted; for the purpose of terminology, the important point is that we use the term to refer to normative ideas of what it means to treat people “fairly”, no matter their source.

When we talk about **bias**, we are using the term in something closer to its statistical sense: we mean properties of estimators, models, measurements, and data that systematically deviate from their intended ideal target. As detailed in Section 3.1, we share an expansive view of bias with Mitchell *et al.* (2020, Section 2.2.1), noting that these biases can be the kinds of statistical biases familiar to science (systematic discrepancies between data or outputs and the underlying observable world), but they can also be societal biases in the form of systematic discrepancies between the observable world and the arguable ideal world that would arise if society eliminated all forms of illegitimate discrimination.

The key distinction in our work is that we use the term “bias” to refer to a fact of the system without making any inherently normative judgment, and “fairness” to discuss the normative aspects of the system and its effects. Some biases are themselves fairness problems; some biases cause fairness problems; some have no effect with regards to

the concerns of fairness; and some may be intentionally introduced to address a fairness problem, often by correcting for another bias. Most fairness problems arise from biases somewhere in the system, its data, or its evaluation, but we find it useful to distinguish between the technical fact and the moral, ethical, or legal concern.

#### 1.4 Fairness and Other Responsibility Concerns

Fairness is commonly grouped together with other concerns under the banner of *responsibility* in computing systems. These concerns include:

**Accountability** Research on accountability examines the legal, social, and technical mechanisms by which computing systems and their operators, developers, and providers may be held accountable, usually for the human effects of their systems. This can connect directly to fairness when considering how to hold organizations accountable for ensuring their systems uphold societal goals to be fair. Such accountability can be through formal structures, such as applying anti-discrimination law to computing systems, or through informal structures such as applying pressure through publicizing the results of third-party audits.

**Transparency** Transparency (and its close cousin explainability) seeks to make the operation and results of algorithmic systems scrutable to users, developers, auditors, and other stakeholders so that it can be understood, reviewed, and contested. This relates to long-standing concern in information access on explanation (Tintarev and Masthoff, 2007), as well as ideas such as scrutable user models (Kay *et al.*, 2002).

**Safety** Information access systems can be harmful. They can distribute false information, promote fake or dangerous products, and provide support for illegal or malicious activities. These problems have received attention in the research literature, often under the general heading of *adversarial information retrieval*. See related workshops AIRWeb (Fetterly and Gyöngyi, 2009) and WebQuality (Nielek *et al.*, 2016).

**Privacy** Aspects of users' profiles including queries, interaction history, and usage patterns may be highly revealing of sensitive personal information: consider queries about medical symptoms or clicks on web pages for addiction counseling. It follows that information access systems have a duty to protect such information from harmful disclosure. Research on privacy-preserving recommendation seeks technical solutions to this challenge. Friedman *et al.* (2015) provide a survey of this area.

**Ethics** Computing ethics is concerned broadly with ensuring that the practice and products of computing adhere to appropriate ethical principles. The ACM Code of Ethics (ACM Council, 2018) specifically calls out non-discrimination, along with attention to potential harms, as an ethical obligation for computing professionals.

The report on the FACTS-IR Workshop on Fairness, Accountability, Confidentiality, Transparency, and Safety in Information Retrieval (Roegiest *et al.*, 2019) discusses how many of these concepts play out in information retrieval. In this work we are concerned with fairness, but bring in other concerns as well when they relate to fairness.

## 1.5 Running Examples

Throughout this monograph, we will use several examples to motivate and explain the various concepts we discuss.

**Job and Candidate Search** Many online platforms attempt to connect job-seekers and employment opportunities in some way. Some of these are dedicated employment-seeking platforms, while others, such as LinkedIn and Xing, are more general-purpose professional networking platforms for which job-seeking is one important component.

Job-seeking is a multisided problem — people need good employment and employers need good candidates — and also has significant fairness requirements that are often subject to regulation in various jurisdictions. Some of the specific fairness concerns for this application include:

- Do users receive a fair set of job opportunities in the recommendations or ads in their feed?

- If the system assesses a match or fit score for a candidate and a job, is this score fair, or does it under- or over-estimate scores for particular candidates or groups of candidates?
- Do users have a fair opportunity to appear in search lists when recruiters are looking for candidates for a job opening (Geyik and Kenthapadi, 2018)?
- Do employers in protected groups (minority-owned businesses, for example) have their jobs fairly promoted to qualified candidates?
- What fairness concerns come from regulatory requirements?

**Music Discovery** The search and recommendation systems in music platforms, such as Spotify, Pandora, and BandCamp, connect listeners with artists. These discovery tools have a significant impact not only on a user's listening experience and musical enjoyment, but also on artists' financial and career prospects, due both to direct revenue from listening and the commercial and reputational effects of visibility. Some specific fairness concerns include:

- Do artists receive fair exposure in the system's search results, recommendation lists, or streamed programming?
- Does the system systematically over- or under-promote particular groups of artists or songwriters through recommendations, search results, and other discovery surfaces (Epps-Darling *et al.*, 2020)?
- Do users receive fair quality of service, or does the system systematically do a better job of modeling some users' tastes and preferences than others?
- Do recommendations reflect well a user's preferences and if not, are there systematic errors due to stereotypes of gender, ethnicity, location, or other attributes?

**News** News search and recommendation influences user exposure to news articles on social media, news aggregation applications, and

search engines. Such influence extends to social and political choices users might make (Kulshrestha *et al.*, 2017; Epstein and Robertson, 2015). Additionally, the filter bubble effect (Pariser, 2011; Alstynne and Brynjolfsson, 2005) may cause users to be exposed primarily to news items that reinforce their beliefs and increase polarization. Depending on the journalistic policy of the provider, news platforms may want to facilitate balanced exposure to news from across the social, political, and cultural spectrum, but this may need to be balanced with the need to de-rank malicious and low-credibility sources.

Specific fairness concerns in news discovery include:

- Does the system provide fair exposure to news on different topics or affected groups?
- Do journalists from different perspectives receive fair visibility or exposure for their content?
- Does the system reward original investigators or primarily direct readers to tertiary sources?
- Do users receive a balanced set of news content?
- Are users in different demographics or locations equally well-served by their news recommendations?

**Philanthropic Giving** Online platforms are increasingly a site for philanthropic giving (Goecks *et al.*, 2008), and therefore recommendation is expected to be an increasing driver of donations. Sites may take an explicitly “peer-to-peer” approach to such giving, as in the educational charity site DonorsChoose.org; this results in many possible donation opportunities for donors to select from, requiring recommendation or sophisticated search to help match donors and opportunities. As many philanthropic organizations have a social justice focus, fairness concerns are essential in developing and evaluating their information access solutions, in particular to avoid potential positive feedback loops in which a subset of causes comes to dominate results and rankings.

In philanthropic settings, we would expect fairness issues to include:

- Does the system provide fair opportunities for the various recipients / causes to have their needs supported?
- Are specific groups of recipients under- or over-represented in the recommendation results?

## 1.6 How to Use This Monograph

We have written this monograph with two audiences in mind:

- Researchers, engineers, and students in information retrieval, recommender systems, and related fields who are looking to understand the literature on fairness, bias, and discrimination, and how it applies to their work.
- Researchers in algorithmic fairness who are looking to understand information access systems, how existing fairness concepts do or do not apply to this set of applications, and the things that information access brings to the research space that may differ from the application settings in which fairness is usually studied.

Due to our interest in serving both of these audiences, we do not expect our readers to have significant familiarity with either information retrieval or algorithmic fairness, although some background in machine learning will be helpful. We have organized the material as follows:

- **Section 2** rehearses the fundamentals of information access systems. This will be a review for most information retrieval and recommender systems researchers; such readers should read it for the terminology we use to integrate the fields, but may wish to focus their study energy elsewhere.
- **Section 3** provides an overview of research on fairness in machine learning generally, particularly in classification. Algorithmic fairness researchers will likely find this section to be a review.
- **Section 4** lays out the problem space of fair information access, providing a multi-faceted taxonomy of the problems in evaluating and removing discrimination and related harms in such systems.

- **Sections 5 and 6** survey key literature to date (as of 2021) on fairness in information access, with pointers to research working on many of the problems identified in Section 4, focused on the two most commonly-studied stakeholders: consumers and providers (with discussion of [subjects](#) in Section 6.4).
- **Section 7** discusses the need to go beyond point-in-time views of fairness to understand fairness over time how the temporal dynamics of an information access system affect fairness.
- **Section 8** looks to future work and provides tips for research and engineering on fair information access.

Section 4 is the keystone of this work that ties the rest together; subsequent sections work out details in the form of a literature survey of several of the problems discussed in Section 4, and the preceding sections set up the background needed to understand it. For readers looking to budget their time, we recommend they ensure they have the necessary background from Sections 2 and 3, read Sections 4 and 8, and read the later sections that are relevant to their work.

## 1.7 Our Perspective

While we have written this monograph to be useful for researchers approaching the topic of fairness from a variety of perspectives, we think it is helpful to explicitly describe our own perspectives and motivations, as well as the position from which we approach this work and some limitations it may bring.

Information access systems need to meet a variety of objectives from multiple stakeholders. They need to deliver relevant results to their users, business value for their operators, and visibility to the creators of the documents they present; they often also need to meet a variety of other goals and constraints, such as diversity across subtopics, regulatory compliance, and reducing avoidable harm to users or society. [Fairness](#), as we conceive of it and present it in this monograph, is not a be-all end goal, but rather another family of objectives to be considered in the design and evaluation of information access systems, and a



collection of techniques for enabling those objectives. It also does not encompass the totality of social or ethical objectives guiding a system's design. Researchers and developers need to work with experts in ethics, policy, sociology, and other relevant fields to identify relevant harms and appropriate objectives for any particular application; the concepts we discuss will be relevant to some of those harms and objectives.

We also emphasize the importance of starting with a robust *problem framing*: Section 4 is intended to help readers think about the fairness problem they are trying to solve, and position it in a landscape of information access; we have then organized our survey in Sections 5-7 around aspects of problem definition, instead of underlying techniques. Metrics and mitigations are best developed and assessed in the context of a specific, well-defined problem.

Finally, all four authors work in North America and approach the topic primarily in that legal and moral context. A Western focus, and particularly concepts of bias and discrimination rooted in United States legal theory, currently dominates thinking and research on algorithmic fairness in general. This is a limitation of the field that others have noted and critiqued (Sambasivan *et al.*, 2020); our present work acknowledges but does not correct this imbalance. While we attempt to engage with definitions and fairness objectives beyond the U.S., this work admittedly has a Western and especially U.S. focus in its treatment of the material. We look forward to seeing other scholars survey this topic from other perspectives.

## 1.8 Some Cautions

We hope that this monograph will help scholars from a variety of backgrounds to understand the emerging literature on fairness in information access and to advance the field in useful directions. In addition to the general concerns of careful, thoughtful science, work on fairness often engages with data and constructs that touch on fundamental aspects of human identity and experience. This work must also be done with great care and compassion to ensure that users, creators, and other stakeholders are treated with respect and dignity and to avoid various traps that result in overbroad or ungeneralizable claims.

We argue that there is nothing particularly new about this, but that thinking about the fairness of information access brings to the surface issues that should be considered in all research and development on information systems.

### 1.8.1 Beware Abstraction Traps

Our first caution is to beware of the allure of abstraction. Selbst *et al.* (2019) describe several specific problems that arise from excessive or inappropriate abstraction in fairness research in general. Their core argument is that the tendency in computer science to seek general, abstract forms of problems, while useful for developing tools and results that can be applied to a wide range of tasks, can cause important social aspects of technology and its impacts to be obscured.

One reason for this is that social problems that appear to be structurally similar arise from distinct (though possibly intertwined) causes and mechanisms, and may require different solutions. Sexism and anti-Black racism, for example, are both types of discrimination and fall into the “group fairness” category of algorithmic fairness, but they are not the same problem and have not been reinforced by the same sets of legal and social processes. Discrimination also varies by culture and jurisdiction, and oppression of what appears to be same group may arise from different causes and through different mechanisms in the different places in which it appears. Kohler-Hausmann (2019) argues that social constructivist frameworks for understanding group identities and experiences imply that even understanding what constitutes a group, let alone the discrimination it experiences, is inextricably linked with understanding how that group is constructed and treated in a particular society — an understanding that is inherently bound to the society in question, although there may be similarities in group construction in different contexts.

The result is that unfairness needs to be measured and addressed in each specific way in which it may appear. While general solutions for detecting and mitigating fairness-related harms may arise and be very useful, their effectiveness needs to be re-validated in context for the harms they are meant to address, a point reiterated by Dwork and Ilvento (2018).

Hoffmann (2019) similarly provides several warnings against overly simple ideas of the harms that can arise from discrimination and bias. Computational fairness inherits some of these limitations from its reference material, such as limitations of anti-discrimination law; others arise from what Hoffmann, Selbst *et al.* (2019), and others argue are reductionistic operationalizations of rich concepts. Hoffmann (2019) notes in particular—and we agree—that treating categories of personal identity as objective features in a multi-dimensional space (a natural move for computer scientists) obfuscates the role of technical and social systems in enacting and producing such categories. This move also has the effect of reducing [intersectionality](#) concerns to what can be captured by a subspace projection or similar formal operation, whether or not that corresponds to individual’s lived experience.

We believe computing systems in general, and information access systems in particular, have the opportunity to *advance* the discussion of emancipation and justice, not just bring existing constructs into a new domain. Information professionals have long been concerned about issues of ethics and justice. Just as two examples, we note that Edmund Berkeley, one of the founders of the Association for Computing Machinery, was an outspoken advocate for the ethical responsibilities of computer scientists as far back as the 1960s (Longo, 2015), and the creation of Computer Professionals for Social Responsibility in the mid-1980s (Finn and DuPont, 2020). The call here is to realize that vision fully and for all people affected by information access systems.

### 1.8.2 Beware Limits

It is crucial to be clear about the limitations of particular fairness studies and methods. Any work will be limited, if for no other reason than the impossibility of completely solving the problem of discrimination. Those limitations should not paralyze the research community or keep researchers from doing the most they can to advance equity and justice with the resources available to them; rather, work in this space needs to be forthright and thorough about the limitations of its approach, data, and findings. Some limitations common to this space include:

- Single-dimensional attributes for which fairness is considered, when in reality people experience discrimination and oppression along multiple simultaneous dimensions.
- Binary definitions of attributes, when in reality many social dimensions have more than two categories or exist on a continuum.
- Taking attributes as fixed and exogenous, when social categories are complex and socially constructed (Hanna *et al.*, 2020).
- Incomplete, erroneous, and/or biased data (Olteanu *et al.*, 2019; Ekstrand and Kluver, 2021).

This is not to say that work on single binary attributes is not useful; research must start somewhere. But it should not *stop* there, and authors need to be clear about the relationship their work in its broader context and provide a careful accounting of its known limitations.

Some methods are so limited that we advise against their use. For example, some work on fair information access has used statistical gender recognition based on names or computer vision techniques for gender recognition based on profile pictures.<sup>2</sup> This source of data is error-prone, subject to systemic biases (Buolamwini and Gebru, 2018), reductionistic (Hamidi *et al.*, 2018), and fundamentally denies subjects control over their identities, so we do not consider it good practice.

### 1.8.3 Beware Convenience

Researchers working in this problem space also need to be careful to do the *best* research possible with available resources, and work to expand those resources to increase the quality and social fidelity of their work, and not take the path of least resistance.

One particular application pertains to this monograph itself and to its proper use and citation. It is convenient and common practice to cite survey papers to quickly summarize a topic or substantiate its

---

<sup>2</sup>We do not provide citations to support the claim that this is in use because our purpose in this paragraph is to critique a general trend, not to focus on any specific paper. Elsewhere in this monograph, we cite work making use of these techniques where it makes a relevant contribution.

relevance. While we naturally welcome citations of our work, we would prefer to be cited specifically for our contributions to the organization and synthesis of fair information access research. The purpose of much of this monograph is to point our readers to the work that others have done, and we specifically ask that you **cite those papers**, instead of — or in addition to — this one when that work is relevant to your writing and research.

## Acknowledgements

---

Michael Ekstrand's contributions are based upon work supported by the National Science Foundation under Grant No. IIS 17-51278. Anubrata Das is supported in part by the Micron Foundation, Wipro, and Good Systems<sup>1</sup>, a University of Texas at Austin Grand Challenge to develop responsible AI technologies. Robin Burke's work was supported by the National Science Foundation under Grant No. IIS 19-11025. We also thank James Atwood, Asia Biega, Yoni Halpern, Matt Lease, Hansa Srinivasan, and the anonymous reviewers for providing additional feedback and suggestions.

---

<sup>1</sup><https://goodsystems.utexas.edu/>

## **Appendix**

# A

---

## Resources for Fair Information Access

---

In this appendix, we collect pointers to several resources for studying and working on fair information access. We have made every effort to ensure these links are current as of the time of publication, but they may degrade more quickly than the references in the rest of the publication.

### A.1 Data Sets

- The TREC Fair Ranking track (launched in 2019) provides data sets for provider fairness in search rankings, both in academic search (2019–2020) and Wikipedia article search (2021). The data is available in TREC (<https://trec.nist.gov/results.html>), with the track web site at <https://fair-trec.github.io>.
- The PIRET Book Data Tools at <https://bookdata.piret.info> provide tools to integrate book recommendation data sets (including from BookCrossing, Amazon, and GoodReads) with publicly-available book and author metadata to study provider fairness in book recommendation, as used by Ekstrand and Kluver (2021).
- Ghosh *et al.* (2021) develop a number of data sets for fair ranking, using various methods and studying the errors of demographic inference for data augmentation.



## A.2 Software

There are not yet widely-distributed open-source software for fair recommendation and retrieval; the available code is mostly embedded in published experiment scripts, or general-purpose systems repurposed for fair information access.

- Terrier (<http://terrierteam.dcs.gla.ac.uk/research.html>) provides xQuAD, a diversification technique that has been successfully applied for fair search ranking (McDonald and Ounis, 2020).
- Experimental scripts are available for the fair recommendation studies of Ekstrand and Kluver (2021) (<https://md.ekstrandom.net/pubs/bag-extended>) and Ekstrand *et al.* (2018b) (<https://md.ekstrandom.net/pubs/cool-kids>).
- librec-auto (<https://librec-auto.readthedocs.io/en/latest/>) provides automated support for running recommender systems experiments, including fairness metrics.

## References

---

- Abdollahpouri, H. (2020). “Popularity Bias in Recommendation: A Multi-stakeholder Perspective”. *PhD thesis*. University of Colorado Boulder. URL: <https://arxiv.org/pdf/2008.08551.pdf>.
- Abdollahpouri, H., G. Adomavicius, R. Burke, I. Guy, D. Jannach, T. Kamishima, J. Krasnodebski, and L. Pizzato. (2020). “Multi-stakeholder Recommendation: Survey and Research Directions”. en. *User Modeling and User-Adapted Interaction*. 30(1): 127–158. DOI: [10.1007/s11257-019-09256-1](https://doi.org/10.1007/s11257-019-09256-1).
- ACM Council. (2018). “ACM Code of Ethics and Professional Conduct”. *Tech. rep.* Association for Computing Machinery. URL: <https://www.acm.org/about-acm/acm-code-of-ethics-and-professional-conduct>.
- Albright, A. (2019). “If You Give a Judge a Risk Score: Evidence from Kentucky Bail Decisions”. *Harvard John M. Olin Fellow’s Discussion Paper*. 85. URL: [https://thelittledataset.com/about\\_files/albright\\_judge\\_score.pdf](https://thelittledataset.com/about_files/albright_judge_score.pdf).
- Ali, M., P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove, and A. Rieke. (2019). “Discrimination through Optimization: How Facebook’s Ad Delivery Can Lead to Biased Outcomes”. *Proceedings of the ACM on Human-Computer Interaction*. 3(CSCW): 1–30. DOI: [10.1145/3359301](https://doi.org/10.1145/3359301).

- Alstyne, M. van and E. Brynjolfsson. (2005). “Global Village or Cyber-Balkans? Modeling and Measuring the Integration of Electronic Communities”. *Management Science*. 51(6): 851–868. DOI: [10.1287/mnsc.1050.0363](https://doi.org/10.1287/mnsc.1050.0363).
- Angwin, J., J. Larson, L. Kirchner, and S. Mattu. (2016). “Machine Bias”. URL: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Aziz, H. (2019). “Developments in Multi-Agent Fair Allocation”. co.RR Nov. arXiv: [1911.09852](https://arxiv.org/abs/1911.09852) [cs.GT].
- Azzopardi, L. and V. Vinay. (2008). “Retrievability: An Evaluation Measure for Higher Order Information Access Tasks”. In: *Proceedings of the 17th ACM Conference on Information and Knowledge Management. CIKM '08*. New York, NY, USA: Association for Computing Machinery. 561–570.
- Barocas, S. and A. D. Selbst. (2016). “Big Data’s Disparate Impact”. *California Law Review*. 104(3): 671. DOI: [10.15779/Z38BG31](https://doi.org/10.15779/Z38BG31).
- Baeza-Yates, R. (2018). “Bias on the web”. en. *Communications of the ACM*. May. DOI: [10.1145/3209581](https://doi.org/10.1145/3209581).
- Bagdasaryan, E., O. Poursaeed, and V. Shmatikov. (2019). “Differential Privacy Has Disparate Impact on Model Accuracy”. In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc. 15479–15488. URL: <http://papers.nips.cc/paper/9681-differential-privacy-has-disparate-impact-on-model-accuracy.pdf>.
- Barocas, S., A. Guo, E. Kamar, J. Krones, M. R. Morris, J. W. Vaughan, W. D. Wadsworth, and H. Wallach. (2021). “Designing Disaggregated Evaluations of AI Systems: Choices, Considerations, and Tradeoffs”. In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society. AIES '21*. Virtual Event, USA: Association for Computing Machinery. 368–378. DOI: [10.1145/3461702.3462610](https://doi.org/10.1145/3461702.3462610).
- Barocas, S., M. Hardt, and A. Narayanan. (2019). *Fairness and Machine Learning: Limitations and Opportunities*. URL: <https://fairmlbook.org>.

- Basilico, J. and Y. Raimond. (2017). “Déjà Vu: The Importance of Time and Causality in Recommender Systems”. In: *Proceedings of the Eleventh ACM Conference on Recommender Systems. RecSys '17*. Como, Italy: Association for Computing Machinery. 342. DOI: [10.1145/3109859.3109922](https://doi.org/10.1145/3109859.3109922).
- Becker, C. D. and E. Ostrom. (1995). “HUMAN ECOLOGY AND RESOURCE SUSTAINABILITY: The Importance of Institutional Diversity”. *Annual Review of Ecology and Systematics*. 26(1): 113–133. DOI: [10.1146/annurev.es.26.110195.000553](https://doi.org/10.1146/annurev.es.26.110195.000553).
- Becker, H. (1982). *Art Worlds*. University of California Press.
- Beel, J. and V. Brunel. (2019). “Data Pruning in Recommender Systems Research: Best-Practice or Malpractice?” In: *ACM RecSys 2019 Late-Breaking Results*. URL: <http://ceur-ws.org/Vol-2431/paper6.pdf>.
- Belkin, N. J. and S. E. Robertson. (1976). “Some ethical and political implications of theoretical research in information science”. In: *Proceedings of the ASIS Annual Meeting*.
- Bender, E. M., T. Gebru, A. McMillan-Major, and S. Shmitchell. (2021). “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜”. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT '21*. Virtual Event, Canada: Association for Computing Machinery. 610–623. DOI: [10.1145/3442188.3445922](https://doi.org/10.1145/3442188.3445922).
- Beutel, A., J. Chen, Z. Zhao, and E. H. Chi. (2017). “Data Decisions and Theoretical Implications when Adversarially Learning Fair Representations”. July. arXiv: [1707.00075](https://arxiv.org/abs/1707.00075) [cs.LG].
- Beutel, A., E. H. Chi, C. Goodrow, J. Chen, T. Doshi, H. Qian, L. Wei, Y. Wu, L. Heldt, Z. Zhao, and L. Hong. (2019). “Fairness in Recommendation Ranking through Pairwise Comparisons”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM Press. DOI: [10.1145/3292500.3330745](https://doi.org/10.1145/3292500.3330745).
- Biega, A. J., F. Diaz, M. D. Ekstrand, and S. Kohlmeier. (2020). “Overview of the TREC 2019 Fair Ranking Track”. In: *The Twenty-Eighth Text REtrieval Conference (TREC 2019) Proceedings*. URL: <https://trec.nist.gov/pubs/trec28/papers/OVERVIEW.FR.pdf>.

- Biega, A. J., K. P. Gummadi, and G. Weikum. (2018). “Equity of Attention: Amortizing Individual Fairness in Rankings”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM. 405–414. DOI: [10.1145/3209978.3210063](https://doi.org/10.1145/3209978.3210063).
- Bigdeli, A., N. Arabzadeh, S. Seyedsalehi, M. Zihayat, and E. 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. SIGIR '21*. Virtual Event, Canada: Association for Computing Machinery. 1748–1752. DOI: [10.1145/3404835.3463110](https://doi.org/10.1145/3404835.3463110).
- Binns, R. (2020). “On the Apparent Conflict between Individual and Group Fairness”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 514–524. DOI: [10.1145/3351095.3372864](https://doi.org/10.1145/3351095.3372864).
- Bodenhausen, G. V. and M. Lichtenstein. (1987). “Social Stereotypes and Information-Processing Strategies: The Impact of Task Complexity”. *Journal of Personality and Social Psychology*. 52(5): 871–880. DOI: [10.1037/0022-3514.52.5.871](https://doi.org/10.1037/0022-3514.52.5.871).
- Bolukbasi, T., K.-W. Chang, J. Zou, V. Saligrama, and A. Kalai. (2016). “Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings”. In: *Advances in Neural Information Processing Systems 29 (NIPS 2016)*. Ed. by D. D. Lee and M. Sugiyama and U. V. Luxburg and I. Guyon and R. Garnett. Curran Associates, Inc. URL: <http://papers.nips.cc/paper/6227-man-is-to-computer-programmer-as-woman-is-to-homemaker-debiasing-word-embeddings>.
- Browne, S. (2015). *Dark Matters: On the Surveillance of Blackness*. en. Duke University Press. URL: <https://play.google.com/store/books/details?id=snmJCgAAQBAJ>.
- Budish, E. and E. Cantillon. (2012). “The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard”. *The American Economic Review*. 102(5): 2237–2271.

- Buolamwini, J. and T. Gebru. (2018). “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification”. In: *Proceedings of the 1st Conference on Fairness, Accountability, and Transparency*. Vol. 81. *Proceedings of Machine Learning Research*. PMLR. 77–91. URL: <http://proceedings.mlr.press/v81/buolamwini18a.html>.
- Burke, R. (2017). “Multisided Fairness for Recommendation”. July. arXiv: [1707.00093](https://arxiv.org/abs/1707.00093) [cs.CY].
- Burke, R., N. Sonboli, and A. Ordonez-Gauger. (2018). “Balanced Neighborhoods for Multi-sided Fairness in Recommendation”. In: *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Ed. by S. A. Friedler and C. Wilson. Vol. 81. *Proceedings of Machine Learning Research*. New York, NY, USA: PMLR. 202–214. URL: <http://proceedings.mlr.press/v81/burke18a.html>.
- Burke, V. I. and R. D. Burke. (2019). “Powerlessness and Personalization”. *The International Journal of Applied Philosophy*. 33(2): 319–343. DOI: [10.5840/ijap202034131](https://doi.org/10.5840/ijap202034131).
- Campos, P. G., F. Díez, and I. Cantador. (2014). “Time-Aware Recommender Systems: A Comprehensive Survey and Analysis of Existing Evaluation Protocols”. *User Modeling and User-Adapted Interaction*. 24(1): 67–119. DOI: [10.1007/s11257-012-9136-x](https://doi.org/10.1007/s11257-012-9136-x).
- Cañamares, R. and P. Castells. (2018). “Should I Follow the Crowd?: A Probabilistic Analysis of the Effectiveness of Popularity in Recommender Systems”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: ACM. 415–424. DOI: [10.1145/3209978.3210014](https://doi.org/10.1145/3209978.3210014).
- Carbonell, J. and J. Goldstein. (1998). “The Use of MMR, Diversity-based Reranking for Reordering Documents and Producing Summaries”. In: *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '98*. New York, NY, USA: ACM. 335–336. DOI: [10.1145/290941.291025](https://doi.org/10.1145/290941.291025).

- Carterette, B. (2011). “System Effectiveness, User Models, and User Utility: A Conceptual Framework for Investigation”. In: *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '11*. Beijing, China: Association for Computing Machinery. 903–912. DOI: [10.1145/2009916.2010037](https://doi.org/10.1145/2009916.2010037).
- Caton, S. and C. Haas. (2020). “Fairness in Machine Learning: A Survey”. Oct. arXiv: [2010.04053](https://arxiv.org/abs/2010.04053) [cs.LG].
- Celis, L. E. and V. Keswani. (2019). “Improved Adversarial Learning for Fair Classification”. Jan. arXiv: [1901.10443](https://arxiv.org/abs/1901.10443) [cs.LG].
- Celis, L. E. and V. Keswani. (2020). “Implicit Diversity in Image Summarization”. *Proceedings of the ACM on Human-Computer Interaction*. 4(CSCW2): 1–28. DOI: [10.1145/3415210](https://doi.org/10.1145/3415210).
- Celis, L. E., D. Straszak, and N. K. Vishnoi. (2018). “Ranking with Fairness Constraints”. In: *45th International Colloquium on Automata, Languages, and Programming*. Ed. by I. Chatzigiannakis, C. Kaklamanis, D. Marx, and D. Sannella. Vol. 107. *Leibniz International Proceedings in Informatics (LIPIcs)*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany. DOI: [10.4230/LIPIcs.ICALP.2018.28](https://doi.org/10.4230/LIPIcs.ICALP.2018.28).
- Celma, Ò. and P. Cano. (2008). “From hits to niches? or how popular artists can bias music recommendation and discovery”. In: *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition. NETFLIX '08*. No. Article 5. Las Vegas, Nevada: Association for Computing Machinery. 1–8. DOI: [10.1145/1722149.1722154](https://doi.org/10.1145/1722149.1722154).
- Chakraborty, A., A. Hannak, A. Biega, and K. Gummadi. (2017). “Fair Sharing for Sharing Economy Platforms”. *Fairness, Accountability and Transparency in Recommender Systems*. Aug. URL: <http://scholarworks.boisestate.edu/fatrec/2017/1/6>.
- Chandar, P., F. Diaz, and B. St. Thomas. (2020). “Beyond Accuracy: Grounding Evaluation Metrics for Human-Machine Learning Systems”. URL: <https://github.com/pchandar/beyond-accuracy-tutorial>.

- Chaney, A. J. B., B. M. Stewart, and B. E. Engelhardt. (2018). “How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility”. In: *Proceedings of the 12th ACM Conference on Recommender Systems. RecSys '18*. Vancouver, British Columbia, Canada: Association for Computing Machinery. 224–232. DOI: [10.1145/3240323.3240370](https://doi.org/10.1145/3240323.3240370).
- Chen, I., F. D. Johansson, and D. Sontag. (2018). “Why Is My Classifier Discriminatory?” In: *Advances in Neural Information Processing Systems 31*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Curran Associates, Inc. 3539–3550. URL: <http://papers.nips.cc/paper/7613-why-is-my-classifier-discriminatory.pdf>.
- Cho, J., S. Roy, and R. E. Adams. (2005). “Page Quality: In Search of an Unbiased Web Ranking”. In: *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data. SIGMOD '05*. Baltimore, Maryland: Association for Computing Machinery. 551–562. DOI: [10.1145/1066157.1066220](https://doi.org/10.1145/1066157.1066220).
- Cho, S., K. W. Crenshaw, and L. McCall. (2013). “Toward a Field of Intersectionality Studies: Theory, Applications, and Praxis”. *Signs: Journal of Women in Culture and Society*. 38(4): 785–810. DOI: [10.1086/669608](https://doi.org/10.1086/669608).
- Chouldechova, A. (2017). “Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments”. en. *Big Data*. 5(2): 153–163. arXiv: [1610.07524](https://arxiv.org/abs/1610.07524) [stat.AP]. URL: <http://dx.doi.org/10.1089/big.2016.0047>.
- Crawford, K. (2017). “The Trouble with Bias”. *Neural Information Processing Systems 2017*. URL: [https://youtu.be/fMym\\_BKWQzk](https://youtu.be/fMym_BKWQzk).
- Crenshaw, K. (1989). “Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics”. *The University of Chicago legal forum*. 1989: 139–168. URL: <https://heinonline.org/HOL/P?h=hein.journals/uchclf1989&i=143>.
- Croft, W. B., D. Metzler, and T. Strohman. (2010). *Search Engines: Information Retrieval in Practice*. Pearson Education, Inc.



- Cui, S., W. Pan, C. Zhang, and F. Wang. (2021). “Towards Model-Agnostic Post-Hoc Adjustment for Balancing Ranking Fairness and Algorithm Utility”. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. KDD '21*. Virtual Event, Singapore: Association for Computing Machinery. 207–217. DOI: [10.1145/3447548.3467251](https://doi.org/10.1145/3447548.3467251).
- D’Amour, A., H. Srinivasan, J. Atwood, P. Baljekar, D. Sculley, and Y. Halpern. (2020). “Fairness Is Not Static: Deeper Understanding of Long Term Fairness via Simulation Studies”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 525–534. DOI: [10.1145/3351095.3372878](https://doi.org/10.1145/3351095.3372878).
- D’Ignazio, C. and L. F. Klein. (2020). *Data Feminism*. MIT Press. URL: <https://data-feminism.mitpress.mit.edu/>.
- Das, A. and M. Lease. (2019). “A Conceptual Framework for Evaluating Fairness in Search”. July. arXiv: [1907.09328](https://arxiv.org/abs/1907.09328) [cs.IR].
- Deerwester, S., S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. (1990). “Indexing by Latent Semantic Analysis”. *Journal of the American Society for Information Science*. 41(6): 391–407. URL: <http://www3.interscience.wiley.com/journal/10049585/abstract?CRETRY=1&SRETRY=0>.
- Deldjoo, Y., V. W. Anelli, H. Zamani, A. Bellogin, and T. Di Noia. (2019). “Recommender Systems Fairness Evaluation via Generalized Cross Entropy”. In: *Proceedings of the Workshop on Recommendation in Multi-stakeholder Environments at RecSys '19*. Vol. 2440. CEUR-WS. URL: <http://arxiv.org/abs/1908.06708>.
- Deshpande, M. and G. Karypis. (2004). “Item-based Top-N Recommendation Algorithms”. *ACM Transactions on Information Systems*. 22(1): 143–177. DOI: [10.1145/963770.963776](https://doi.org/10.1145/963770.963776).
- Diakopoulos, N. (2015). “Algorithmic Accountability: Journalistic Investigation of Computational Power Structures”. en. *Digital Journalism*. 3(3): 398–415. DOI: [10.1080/21670811.2014.976411](https://doi.org/10.1080/21670811.2014.976411).

- Diaz, F., B. Mitra, M. D. Ekstrand, A. J. Biega, and B. Carterette. (2020). “Evaluating Stochastic Rankings with Expected Exposure”. In: *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*. ACM. DOI: [10.1145/3340531.3411962](https://doi.org/10.1145/3340531.3411962).
- Dragovic, N., I. Madrazo Azpiazu, and M. S. Pera. (2016). ““Is Sven Seven?”: A Search Intent Module for Children”. In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '16*. Pisa, Italy: ACM. 885–888. DOI: [10.1145/2911451.2914738](https://doi.org/10.1145/2911451.2914738).
- Dutta, S., D. Wei, H. Yueksel, P.-Y. Chen, S. Liu, and K. Varshney. (2020). “Is There a Trade-Off Between Fairness and Accuracy? A Perspective Using Mismatched Hypothesis Testing”. In: *Proceedings of the 37th International Conference on Machine Learning*. Ed. by H. D. Iii and A. Singh. Vol. 119. *Proceedings of Machine Learning Research*. PMLR. 2803–2813. URL: <http://proceedings.mlr.press/v119/dutta20a.html>.
- Dwork, C., M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. (2012). “Fairness Through Awareness”. In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. ITCS '12*. Cambridge, Massachusetts: Association for Computing Machinery. 214–226. DOI: [10.1145/2090236.2090255](https://doi.org/10.1145/2090236.2090255).
- Dwork, C. and C. Ilvento. (2018). “Fairness Under Composition”. In: *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*. Ed. by A. Blum. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. DOI: [10.4230/LIPICS.ITCS.2019.33](https://doi.org/10.4230/LIPICS.ITCS.2019.33).
- Edwards, H. and A. Storkey. (2016). “Censoring Representations with an Adversary”. In: *Proceedings of the International Conference on Learning Representations (ICLR 2016)*. URL: <http://arxiv.org/abs/1511.05897>.

- Ekstrand, M. D., R. Joshaghani, and H. Mehrpouyan. (2018a). “Privacy for All: Ensuring Fair and Equitable Privacy Protections”. In: *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Ed. by S. A. Friedler and C. Wilson. Vol. 81. *Proceedings of Machine Learning Research*. New York, NY: PMLR. 35–47. URL: <https://proceedings.mlr.press/v81/ekstrand18a.html>.
- Ekstrand, M. D. and D. Kluver. (2021). “Exploring Author Gender in Book Rating and Recommendation”. *User Modeling and User-Adapted Interaction*. 31(3). DOI: [10.1007/s11257-020-09284-2](https://doi.org/10.1007/s11257-020-09284-2).
- Ekstrand, M. D., M. Tian, I. M. Azpiazu, J. D. Ekstrand, O. Anuyah, D. McNeill, and M. S. Pera. (2018b). “All The Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness”. In: *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Ed. by S. A. Friedler and C. Wilson. Vol. 81. *Proceedings of Machine Learning Research*. New York, New York: PMLR. 172–186. URL: <https://proceedings.mlr.press/v81/ekstrand18b.html>.
- Ekstrand, M. D., M. Tian, M. R. I. Kazi, H. Mehrpouyan, and D. Kluver. (2018c). “Exploring Author Gender in Book Rating and Recommendation”. In: *Proceedings of the 12th ACM Conference on Recommender Systems*. Vancouver British Columbia Canada: ACM. DOI: [10.1145/3240323.3240373](https://doi.org/10.1145/3240323.3240373).
- Ensign, D., S. A. Friedler, S. Neville, C. Scheidegger, and S. Venkatasubramanian. (2018). “Runaway Feedback Loops in Predictive Policing”. In: *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Ed. by S. A. Friedler and C. Wilson. Vol. 81. *Proceedings of Machine Learning Research*. New York, NY, USA: PMLR. 160–171. URL: <http://proceedings.mlr.press/v81/ensign18a.html>.
- Epps-Darling, A., R. T. Bouyer, and H. Cramer. (2020). “Artist Gender Representation in Music Streaming”. In: *Proceedings of the 21st International Society for Music Information Retrieval Conference*. ISMIR. 248–254. URL: [https://program.ismir2020.net/poster\\_2-11.html](https://program.ismir2020.net/poster_2-11.html).

- Epstein, R. and R. E. Robertson. (2015). “The Search Engine Manipulation Effect (SEME) and its Possible Impact on the Outcomes of Elections”. en. *Proceedings of the National Academy of Sciences of the United States of America*. 112(33): E4512–21. DOI: [10.1073/pnas.1419828112](https://doi.org/10.1073/pnas.1419828112).
- Evans, D., R. Schmalensee, M. Noel, H. Chang, and D. Garcia-Swartz. (2011). *Platform Economics: Essays on Multi-Sided Businesses*. Competition Policy International. URL: <http://ssrn.com/abstract=1974020>.
- Evans, D. S. and R. Schmalensee. (2016). *Matchmakers: The New Economics of Multisided Platforms*. Harvard Business Review Press.
- Feldman, M., S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. (2015). “Certifying and Removing Disparate Impact”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM. 259–268. DOI: [10.1145/2783258.2783311](https://doi.org/10.1145/2783258.2783311).
- Feng, R., Y. Yang, Y. Lyu, C. Tan, Y. Sun, and C. Wang. (2019). “Learning Fair Representations via an Adversarial Framework”. Apr. arXiv: [1904.13341](https://arxiv.org/abs/1904.13341) [cs.LG].
- Ferraro, A. (2019). “Music Cold-Start and Long-Tail Recommendation: Bias in Deep Representations”. In: *Proceedings of the 13th ACM Conference on Recommender Systems. RecSys '19*. Copenhagen, Denmark: Association for Computing Machinery. 586–590. DOI: [10.1145/3298689.3347052](https://doi.org/10.1145/3298689.3347052).
- Ferraro, A., X. Serra, and C. Bauer. (2021). “What is Fair? Exploring the Artists’ Perspective on the Fairness of Music Streaming Platforms”. In: *Proceedings of the 18th IFIP International Conference on Human-Computer Interaction (INTERACT 2021)*. URL: <http://arxiv.org/abs/2106.02415>.
- Ferro, N., N. Fuhr, G. Grefenstette, J. A. Konstan, P. Castells, E. M. Daly, T. Declerck, M. D. Ekstrand, W. Geyer, J. Gonzalo, T. Kuflik, K. Lindn, B. Magnini, J.-Y. Nie, R. Perego, B. Shapira, I. Soboroff, N. Tintarev, K. Verspoor, M. C. Willemsen, and J. Zobel. (2018). “The Dagstuhl Perspectives Workshop on Performance Modeling and Prediction”. *SIGIR Forum*. 52(1): 91–101. DOI: [10.1145/3274784.3274789](https://doi.org/10.1145/3274784.3274789).

- Fetterly, D. and Z. Gyöngyi, eds. (2009). *Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web*. Madrid, Spain: Association for Computing Machinery. URL: <https://dl.acm.org/doi/proceedings/10.1145/1531914>.
- Finn, M. and Q. DuPont. (2020). “From closed world discourse to digital utopianism: the changing face of responsible computing at Computer Professionals for Social Responsibility (1981–1992)”. *Internet Histories*. 4(1): 6–31.
- Fisher, E. (2022). “Do algorithms have a right to the city? Waze and algorithmic spatiality”. en. *Cultural Studies of Science Education*. 36(1): 74–95. DOI: [10.1080/09502386.2020.1755711](https://doi.org/10.1080/09502386.2020.1755711).
- Fleder, D. M. and K. Hosanagar. (2009). “Blockbuster Culture’s Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity”. *Management Science*. 55(5): 697–712. DOI: [10.1287/mnsc.1080.0974](https://doi.org/10.1287/mnsc.1080.0974).
- Foulds, J. R., R. Islam, K. N. Keya, and S. Pan. (2020). “An Intersectional Definition of Fairness”. In: *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. 1918–1921. DOI: [10.1109/ICDE48307.2020.00203](https://doi.org/10.1109/ICDE48307.2020.00203).
- Friedler, S. A., C. Scheidegger, and S. Venkatasubramanian. (2016). “On the (im)possibility of fairness”. Sept. arXiv: [1609.07236](https://arxiv.org/abs/1609.07236) [cs.CY].
- Friedler, S. A., C. Scheidegger, and S. Venkatasubramanian. (2021). “The (Im)possibility of Fairness”. en. *Communications of the ACM*. 64(4): 136–143. DOI: [10.1145/3433949](https://doi.org/10.1145/3433949).
- Friedler, S. A., C. Scheidegger, S. Venkatasubramanian, S. Choudhary, E. P. Hamilton, and D. Roth. (2019). “A Comparative Study of Fairness-Enhancing Interventions in Machine Learning”. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* ’19*. Atlanta, GA, USA: Association for Computing Machinery. 329–338. DOI: [10.1145/3287560.3287589](https://doi.org/10.1145/3287560.3287589).
- Friedman, A., B. P. Knijnenburg, K. Vanhecke, L. Martens, and S. Berkovsky. (2015). “Privacy Aspects of Recommender Systems”. In: *Recommender Systems Handbook*. Boston, MA: Springer US. 649–688. DOI: [10.1007/978-1-4899-7637-6\\_19](https://doi.org/10.1007/978-1-4899-7637-6_19).

- Friedman, B. and H. Nissenbaum. (1996). “Bias in Computer Systems”. *ACM Transactions on Information Systems*. 14(3): 330–347. DOI: [10.1145/230538.230561](https://doi.org/10.1145/230538.230561).
- Funk, S. (2006). “Netflix Update: Try This at Home”. URL: <http://sifter.org/~simon/journal/20061211.html>.
- García-Soriano, D. and F. Bonchi. (2021). “Maxmin-Fair Ranking: Individual Fairness under Group-Fairness Constraints”. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. KDD '21*. Virtual Event, Singapore: Association for Computing Machinery. 436–446. DOI: [10.1145/3447548.3467349](https://doi.org/10.1145/3447548.3467349).
- Gebru, T., J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumé III, and K. Crawford. (2018). “Datasheets for Datasets”. Mar. arXiv: [1803.09010](https://arxiv.org/abs/1803.09010) [cs.DB].
- Geyik, S. C., S. Ambler, and K. 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. KDD '19*. Anchorage, AK, USA: Association for Computing Machinery. 2221–2231. DOI: [10.1145/3292500.3330691](https://doi.org/10.1145/3292500.3330691).
- Geyik, S. C. and K. Kenthapadi. (2018). “Building Representative Talent Search at LinkedIn”. URL: <https://engineering.linkedin.com/blog/2018/10/building-representative-talent-search-at-linkedin>.
- Ghosh, A., R. Dutt, and C. Wilson. (2021). “When Fair Ranking Meets Uncertain Inference”. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '21*. Virtual Event, Canada: Association for Computing Machinery. 1033–1043. DOI: [10.1145/3404835.3462850](https://doi.org/10.1145/3404835.3462850).
- Glowacka, D. (2019). “Bandit Algorithms in Information Retrieval”. *Foundations and Trends® in Information Retrieval*. 13(4): 299–424. DOI: [10.1561/15000000067](https://doi.org/10.1561/15000000067).
- Goecks, J., A. Volda, S. Volda, and E. D. Mynatt. (2008). “Charitable Technologies: Opportunities for Collaborative Computing in Non-profit Fundraising”. In: *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work. CSCW '08*. New York, NY, USA: ACM. 689–698. DOI: [10.1145/1460563.1460669](https://doi.org/10.1145/1460563.1460669).

- Goffman, W. (1964). “A Searching Procedure for Information Retrieval”. *Information Storage and Retrieval*. 2(2): 73–78. DOI: [10.1016/0020-0271\(64\)90006-3](https://doi.org/10.1016/0020-0271(64)90006-3).
- Goldman, E. (2005). “Search Engine Bias and the Demise of Search Engine Utopianism”. *Yale Journal of Law & Technology*. 8: 188+. URL: <https://go.gale.com/ps/i.do?p=AONE&sw=w&issn=&v=2.1&it=r&id=GALE%7CA182079911&sid=googleScholar&linkaccess=fulltext&userGroupName=anon%7Ee0e23234>.
- Gómez, E., C. Shui Zhang, L. Boratto, M. Salamó, and M. Marras. (2021). “The Winner Takes it All: Geographic Imbalance and Provider (Un)fairness in Educational Recommender Systems”. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '21*. Virtual Event, Canada: Association for Computing Machinery. 1808–1812. DOI: [10.1145/3404835.3463235](https://doi.org/10.1145/3404835.3463235).
- Gosepath, S. (2011). “Equality”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2011. Metaphysics Research Lab, Stanford University. URL: <https://plato.stanford.edu/archives/spr2011/entries/equality/>.
- Green, B. and Y. Chen. (2019). “Disparate Interactions: An Algorithm-in-the-Loop Analysis of Fairness in Risk Assessments”. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* '19*. Atlanta, GA, USA: ACM. 90–99. DOI: [10.1145/3287560.3287563](https://doi.org/10.1145/3287560.3287563).
- Guo, R., X. Zhao, A. Henderson, L. Hong, and H. Liu. (2020). “Debiasing Grid-Based Product Search in E-Commerce”. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '20*. New York, NY, USA: Association for Computing Machinery. 2852–2860.
- Hamidi, F., M. K. Scheuerman, and S. M. Branham. (2018). “Gender Recognition or Gender Reductionism?: The Social Implications of Embedded Gender Recognition Systems”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. CHI '18*. ACM. 8. DOI: [10.1145/3173574.3173582](https://doi.org/10.1145/3173574.3173582).
- Hammer, M. C. (2021). “You bore us...” URL: <https://twitter.com/MCHammer/status/1363908982289559553>.

- Hanna, A., E. Denton, A. Smart, and J. Smith-Loud. (2020). “Towards a Critical Race Methodology in Algorithmic Fairness”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 501–512. DOI: [10.1145/3351095.3372826](https://doi.org/10.1145/3351095.3372826).
- Harambam, J., D. Bountouridis, M. Makhortykh, and J. van Hoboken. (2019). “. Designing for the Better by Taking Users into Account: A Qualitative Evaluation of User Control Mechanisms in (News) Recommender Systems”. In: *Proceedings of the 13th ACM Conference on Recommender Systems. RecSys '19*. Copenhagen, Denmark: Association for Computing Machinery. 69–77. DOI: [10.1145/3298689.3347014](https://doi.org/10.1145/3298689.3347014).
- Hardt, M., E. Price, and N. Srebro. (2016). “Equality of Opportunity in Supervised Learning”. In: *Advances in Neural Information Processing Systems*. papers.nips.cc. 3315–3323. URL: <http://papers.nips.cc/paper/6373-equality-of-opportunity-in-supervised-learning>.
- Harman, J. L., J. O'Donovan, T. Abdelzaher, and C. Gonzalez. (2014). “Dynamics of Human Trust in Recommender Systems”. In: *Proceedings of the 8th ACM Conference on Recommender Systems. RecSys '14*. Foster City, Silicon Valley, California, USA: Association for Computing Machinery. 305–308. DOI: [10.1145/2645710.2645761](https://doi.org/10.1145/2645710.2645761).
- Harper, F. M. and J. A. Konstan. (2015). “The MovieLens Datasets: History and Context”. *ACM Transactions on Interactive Intelligent Systems*. 5(4): 19:1–19:19. DOI: [10.1145/2827872](https://doi.org/10.1145/2827872).
- Hashimoto, T., M. Srivastava, H. Namkoong, and P. Liang. (2018). “Fairness Without Demographics in Repeated Loss Minimization”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. *Proceedings of Machine Learning Research*. Stockholm, Sweden: PMLR. 1929–1938. URL: <http://proceedings.mlr.press/v80/hashimoto18a.html>.
- Herlocker, J., J. A. Konstan, and J. Riedl. (2002). “An Empirical Analysis of Design Choices in Neighborhood-Based Collaborative Filtering Algorithms”. In: *Information Retrieval*. 5(4): 287–310. DOI: [10.1023/A:1020443909834](https://doi.org/10.1023/A:1020443909834).



- Hoffmann, A. L. (2019). “Where Fairness Fails: Data, Algorithms, and the Limits of Antidiscrimination Discourse”. *Information, Communication and Society*. 22(7): 900–915. DOI: [10.1080/1369118X.2019.1573912](https://doi.org/10.1080/1369118X.2019.1573912).
- HUD. (2020). “HUD’s Implementation of the Fair Housing Act’s Disparate Impact Standard”. *Federal Register*. 85(186): 60288–60333. URL: <https://www.federalregister.gov/documents/2020/09/24/2020-19887/huds-implementation-of-the-fair-housing-acts-disparate-impact-standard>.
- Hurley, N. and M. Zhang. (2011). “Novelty and Diversity in Top-N Recommendation – Analysis and Evaluation”. *ACM Transactions on Internet Technology*. 10(4): 14:1–14:30. DOI: [10.1145/1944339.1944341](https://doi.org/10.1145/1944339.1944341).
- Hutchinson, B. and M. Mitchell. (2019). “50 Years of Test (Un)fairness: Lessons for Machine Learning”. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* ’19*. Atlanta, GA, USA: Association for Computing Machinery. 49–58. DOI: [10.1145/3287560.3287600](https://doi.org/10.1145/3287560.3287600).
- Hutchinson, B., A. Smart, A. Hanna, E. Denton, C. Greer, O. Kjartansson, P. Barnes, and M. Mitchell. (2021). “Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure”. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT ’21*. Virtual Event, Canada: Association for Computing Machinery. 560–575. DOI: [10.1145/3442188.3445918](https://doi.org/10.1145/3442188.3445918).
- IFLA Governing Board. (2012). “Code of Ethics for Librarians and Other Information Workers”. *Tech. rep.* International Federation of Library Associations and Institutions. URL: <https://repository.ifla.org/handle/123456789/1850>.
- Introna, L. D. and H. Nissenbaum. (2000). “Shaping the Web: Why the Politics of Search Engines Matters”. *The Information Society*. 16(3): 169–185.
- Jacobs, A. Z. and H. Wallach. (2021). “Measurement and Fairness”. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT ’21*. Virtual Event, Canada: Association for Computing Machinery. 375–385. DOI: [10.1145/3442188.3445901](https://doi.org/10.1145/3442188.3445901).

- Jambor, T., J. Wang, and N. Lathia. (2012). “Using Control Theory for Stable and Efficient Recommender Systems”. In: *Proceedings of the 21st International Conference on World Wide Web. WWW '12*. 11–20. DOI: [10.1145/2187836.2187839](https://doi.org/10.1145/2187836.2187839).
- Jannach, D., L. Lerche, I. Kamehkhosh, and M. Jugovac. (2015). “What Recommenders Recommend: An Analysis of Recommendation Biases and Possible Countermeasures”. en. *User Modeling and User-Adapted Interaction*. 25(5): 427–491. DOI: [10.1007/s11257-015-9165-3](https://doi.org/10.1007/s11257-015-9165-3).
- Jiang, R., S. Chiappa, T. Lattimore, A. György, and P. Kohli. (2019). “Degenerate feedback loops in recommender systems”. In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 383–390.
- Joachims, T., L. Granka, B. Pan, H. Hembrooke, and G. Gay. (2017). “Accurately Interpreting Clickthrough Data as Implicit Feedback”. *SIGIR Forum*. 51(1): 4–11. DOI: [10.1145/3130332.3130334](https://doi.org/10.1145/3130332.3130334).
- Johnson, I., J. Henderson, C. Perry, J. Schöning, and B. Hecht. (2017). “Beautiful... but at What Cost? An Examination of Externalities in Geographic Vehicle Routing”. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*. 1(2): 1–21. DOI: [10.1145/3090080](https://doi.org/10.1145/3090080).
- Jones, R. and K. L. Klinkner. (2008). “Beyond the Session Timeout: Automatic Hierarchical Segmentation of Search Topics in Query Logs”. In: *Proceedings of the 17th ACM Conference on Information and Knowledge Management. CIKM '08*. Napa Valley, California, USA: Association for Computing Machinery. 699–708. DOI: [10.1145/1458082.1458176](https://doi.org/10.1145/1458082.1458176).
- Joseph, M., M. Kearns, J. Morgenstern, S. Neel, and A. Roth. (2018). “Meritocratic Fairness for Infinite and Contextual Bandits”. In: *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society. AIES '18*. New Orleans, LA, USA: Association for Computing Machinery. 158–163. DOI: [10.1145/3278721.3278764](https://doi.org/10.1145/3278721.3278764).

- Joseph, M., M. Kearns, J. H. Morgenstern, and A. Roth. (2016). “Fairness in Learning: Classic and Contextual Bandits”. In: *Advances in Neural Information Processing Systems 29*. Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett. Curran Associates, Inc. 325–333. URL: <http://papers.nips.cc/paper/6355-fairness-in-learning-classic-and-contextual-bandits.pdf>.
- Kallus, N., X. Mao, and A. Zhou. (2020). “Assessing Algorithmic Fairness with Unobserved Protected Class Using Data Combination”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 110. DOI: [10.1145/3351095.3373154](https://doi.org/10.1145/3351095.3373154).
- Kamiran, F. and T. Calders. (2009). “Classifying Without Discriminating”. In: *2009 2nd International Conference on Computer, Control and Communication*. [ieeexplore.ieee.org](http://ieeexplore.ieee.org). 1–6. DOI: [10.1109/IC4.2009.4909197](https://doi.org/10.1109/IC4.2009.4909197).
- Kamiran, F. and T. Calders. (2012). “Data Preprocessing Techniques for Classification Without Discrimination”. en. *Knowledge and Information Systems*. 33(1): 1–33. DOI: [10.1007/s10115-011-0463-8](https://doi.org/10.1007/s10115-011-0463-8).
- Kamishima, T. and S. Akaho. (2017). “Considerations on Recommendation Independence for a Find-Good-Items Task”. In: *Workshop on Fairness, Accountability and Transparency in Recommender Systems at RecSys 2017*. URL: <http://scholarworks.boisestate.edu/fatrec/2017/1/11/>.
- Kamishima, T., S. Akaho, H. Asoh, and J. Sakuma. (2018). “Recommendation Independence”. In: *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Ed. by S. A. Friedler and C. Wilson. Vol. 81. *Proceedings of Machine Learning Research*. New York, NY, USA: PMLR. 187–201. URL: <http://proceedings.mlr.press/v81/kamishima18a.html>.
- Karako, C. and P. Manggala. (2018). “Using Image Fairness Representations in Diversity-Based Re-ranking for Recommendations”. In: *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization. UMAP '18*. Singapore, Singapore: Association for Computing Machinery. 23–28. DOI: [10.1145/3213586.3226206](https://doi.org/10.1145/3213586.3226206).

- Karimi, M., D. Jannach, and M. Jugovac. (2018). “News Recommender Systems — Survey and Roads Ahead”. *Information Processing & Management*. 54(6): 1203–1227. DOI: [10.1016/j.ipm.2018.04.008](https://doi.org/10.1016/j.ipm.2018.04.008).
- Katell, M., M. Young, D. Dailey, B. Herman, V. Guetler, A. Tam, C. Bintz, D. Raz, and P. M. Krafft. (2020). “Toward Situated Interventions for Algorithmic Equity: Lessons from the Field”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 45–55. DOI: [10.1145/3351095.3372874](https://doi.org/10.1145/3351095.3372874).
- Kay, J., B. Kummerfeld, and P. Lauder. (2002). “Personis: A Server for User Models”. en. In: *Adaptive Hypermedia and Adaptive Web-Based Systems*. Ed. by P. D. Bra, P. Brusilovsky, and R. Conejo. *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 203–212. URL: [http://link.springer.com/chapter/10.1007/3-540-47952-X\\_22](http://link.springer.com/chapter/10.1007/3-540-47952-X_22).
- Kay, M., C. Matuszek, and S. 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. CHI '15*. Seoul, Republic of Korea: ACM. 3819–3828. DOI: [10.1145/2702123.2702520](https://doi.org/10.1145/2702123.2702520).
- Kaya, M., D. Bridge, and N. Tintarev. (2020). “Ensuring Fairness in Group Recommendations by Rank-Sensitive Balancing of Relevance”. In: *Fourteenth ACM Conference on Recommender Systems. RecSys '20*. Virtual Event, Brazil: Association for Computing Machinery. 101–110. DOI: [10.1145/3383313.3412232](https://doi.org/10.1145/3383313.3412232).
- Kearns, M., S. Neel, A. Roth, and Z. S. Wu. (2017). “Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness”. *arXiv preprint arXiv:1711.05144*.
- Kearns, M., S. Neel, A. Roth, and Z. S. Wu. (2019). “An Empirical Study of Rich Subgroup Fairness for Machine Learning”. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* '19*. Atlanta, GA, USA: Association for Computing Machinery. 100–109. DOI: [10.1145/3287560.3287592](https://doi.org/10.1145/3287560.3287592).
- Kelly, D. and J. Teevan. (2003). “Implicit Feedback for Inferring User Preference: A Bibliography”. *SIGIR Forum*. 37(2): 18–28. DOI: [10.1145/959258.959260](https://doi.org/10.1145/959258.959260).

- Khenissi, S., B. Mariem, and O. Nasraoui. (2020). “Theoretical Modeling of the Iterative Properties of User Discovery in a Collaborative Filtering Recommender System”. In: *Fourteenth ACM Conference on Recommender Systems. RecSys '20*. Virtual Event, Brazil: Association for Computing Machinery. 348–357. DOI: [10.1145/3383313.3412260](https://doi.org/10.1145/3383313.3412260).
- Kırnap, Ö., F. Diaz, A. Biega, M. Ekstrand, B. Carterette, and E. Yilmaz. (2021). “Estimation of Fair Ranking Metrics with Incomplete Judgments”. In: *Proceedings of the Web Conference 2021. WWW '21*. Ljubljana, Slovenia: Association for Computing Machinery. 1065–1075. DOI: [10.1145/3442381.3450080](https://doi.org/10.1145/3442381.3450080).
- Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan. (2018). “Human Decisions and Machine Predictions”. *The Quarterly Journal of Economics*. 133(1): 237–293.
- Kleinberg, J., S. Mullainathan, and M. Raghavan. (2017). “Inherent Trade-Offs in the Fair Determination of Risk Scores”. In: *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Ed. by C. H. Papadimitriou. Vol. 67. *Leibniz International Proceedings in Informatics (LIPIcs)*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany. DOI: [10.4230/LIPICS.ITCS.2017.43](https://doi.org/10.4230/LIPICS.ITCS.2017.43).
- Kohavi, R., D. Tang, and Y. Xu. (2020). *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press. DOI: [10.1017/9781108653985](https://doi.org/10.1017/9781108653985).
- Kohler-Hausmann, I. (2019). “Eddie Murphy and the Dangers of Counterfactual Causal Thinking About Detecting Racial Discrimination”. *Northwestern Law Review*. 113(5): 1163–1228. DOI: [10.2139/ssrn.3050650](https://doi.org/10.2139/ssrn.3050650).
- Koren, Y., R. Bell, and C. Volinsky. (2009). “Matrix Factorization Techniques for Recommender Systems”. *Computer*. 42(8): 30–37. DOI: [10.1109/MC.2009.263](https://doi.org/10.1109/MC.2009.263).
- Koren, Y. (2010). “Collaborative filtering with temporal dynamics”. *Communications of the ACM*. 53(4): 89–97. DOI: [10.1145/1721654.1721677](https://doi.org/10.1145/1721654.1721677).

- Kouki, P., I. Fountalis, N. Vasiloglou, X. Cui, E. Liberty, and K. Al Jadda. (2020). “From the Lab to Production: A Case Study of Session-Based Recommendations in the Home-Improvement Domain”. In: *Fourteenth ACM Conference on Recommender Systems. RecSys '20*. Virtual Event Brazil: ACM. 140–149. DOI: [10.1145/3383313.3412235](https://doi.org/10.1145/3383313.3412235).
- Kuhlman, C., W. Gerych, and E. Rundensteiner. (2021). “Measuring Group Advantage: A Comparative Study of Fair Ranking Metrics”. In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society. AIES '21*. Virtual Event, USA: Association for Computing Machinery. 674–682. DOI: [10.1145/3461702.3462588](https://doi.org/10.1145/3461702.3462588).
- Kuhlthau, C. C. (1993). “A Principle of Uncertainty for Information Seeking”. *Journal of Documentation*. 49(4): 339–355. DOI: [10.1108/eb026918](https://doi.org/10.1108/eb026918).
- Kulshrestha, J., M. Eslami, J. Messias, M. B. Zafar, S. Ghosh, K. P. Gummadi, and K. Karahalios. (2017). “Quantifying Search Bias: Investigating Sources of Bias for Political Searches in Social Media”. In: *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing. CSCW '17*. Portland, Oregon, USA: Association for Computing Machinery. 417–432. DOI: [10.1145/2998181.2998321](https://doi.org/10.1145/2998181.2998321).
- Kulynych, B., R. Overdorf, C. Troncoso, and S. Gürses. (2020). “POTs: Protective Optimization Technologies”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 177–188. DOI: [10.1145/3351095.3372853](https://doi.org/10.1145/3351095.3372853).
- Lahoti, P., K. P. Gummadi, and G. Weikum. (2019). “iFair: Learning Individually Fair Data Representations for Algorithmic Decision Making”. In: *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. 1334–1345. DOI: [10.1109/ICDE.2019.00121](https://doi.org/10.1109/ICDE.2019.00121).
- Lathia, N., S. Hailes, and L. Capra. (2009). “Evaluating Collaborative Filtering Over Time”. In: *SIGIR '09 Workshop on the Future of IR Evaluation*. URL: [http://www.cs.ucl.ac.uk/staff/n.lathia/papers/lathia\\_ireval09.pdf](http://www.cs.ucl.ac.uk/staff/n.lathia/papers/lathia_ireval09.pdf).

- Lee, M. K., A. Jain, H. J. Cha, S. Ojha, and D. Kusbit. (2019). “Procedural Justice in Algorithmic Fairness: Leveraging Transparency and Outcome Control for Fair Algorithmic Mediation”. *Proceedings of the ACM on Human-Computer Interaction*. 3(CSCW): 1–26. DOI: [10.1145/3359284](https://doi.org/10.1145/3359284).
- Lerner, M. (2021). “Fannie Mae to include rent payments in mortgage applicants’ credit history review”. *Washington Post*. URL: <https://www.washingtonpost.com/business/2021/09/08/fannie-mae-include-rent-payments-mortgage-applicants-credit-history-review/>.
- Leuski, A. and J. Allan. (2000). “Lighthouse: Showing the Way to Relevant Information”. In: *Proceedings of the IEEE Symposium on Information Visualization*. 125–130.
- Li, L., W. Chu, J. Langford, and R. E. Schapire. (2010). “A Contextual-Bandit Approach to Personalized News Article Recommendation”. In: *Proceedings of the 19th International Conference on World Wide Web*. 661–670.
- Li, Y., H. Chen, S. Xu, Y. Ge, and Y. Zhang. (2021). “Towards Personalized Fairness based on Causal Notion”. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’21*. Virtual Event, Canada: Association for Computing Machinery. 1054–1063. DOI: [10.1145/3404835.3462966](https://doi.org/10.1145/3404835.3462966).
- Lian, J. W., N. Mattei, R. Noble, and T. Walsh. (2018). “The Conference Paper Assignment Problem: Using Order Weighted Averages to Assign Indivisible Goods”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. No. 1. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/11484>.
- Lipton, Z., J. McAuley, and A. Chouldechova. (2018). “Does Mitigating ML’s Impact Disparity Require Treatment Disparity?” In: *Advances in Neural Information Processing Systems 31*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Curran Associates, Inc. 8125–8135. URL: <http://papers.nips.cc/paper/8035-does-mitigating-mls-impact-disparity-require-treatment-disparity.pdf>.

- Liu, J., C. Liu, and N. J. Belkin. (2020). “Personalization in Text Information Retrieval: A Survey”. *Journal of the Association for Information Science and Technology*. 71(3): 349–369. DOI: [10.1002/asi.24234](https://doi.org/10.1002/asi.24234).
- Liu, T.-Y. (2007). “Learning to Rank for Information Retrieval”. en. *Foundations and Trends® in Information Retrieval*. 3(3): 225–331. DOI: [10.1561/15000000016](https://doi.org/10.1561/15000000016).
- Longo, B. (2015). *Edmund Berkeley and the Social Responsibility of Computer Professionals*. Association for Computing Machinery and Morgan & Claypool.
- Lucherini, E., M. Sun, A. Winecoff, and A. Narayanan. (2021). “T-RECS: A simulation tool to study the societal impact of recommender systems”. July. arXiv: [2107.08959](https://arxiv.org/abs/2107.08959) [cs.CY].
- Luhn, H. P. (1960). “Key Word-in-Context Index for Technical Literature (KWIC Index)”. *American Documentation*. 11(4): 288–295. DOI: [10.1002/asi.5090110403](https://doi.org/10.1002/asi.5090110403).
- Madras, D., E. Creager, T. Pitassi, and R. Zemel. (2018). “Learning Adversarially Fair and Transferable Representations”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. *Proceedings of Machine Learning Research*. PMLR. 3384–3393. URL: <https://proceedings.mlr.press/v80/madras18a.html>.
- Mann, T. A., S. Gowal, A. Gyorgy, H. Hu, R. Jiang, B. Lakshminarayanan, and P. Srinivasan. (2019). “Learning from Delayed Outcomes via Proxies with Applications to Recommender Systems”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by K. Chaudhuri and R. Salakhutdinov. Vol. 97. *Proceedings of Machine Learning Research*. PMLR. 4324–4332.
- Manning, C. D., P. Raghavan, and H. Schütze. (2008). *Introduction to Information Retrieval*. Cambridge University Press.
- Mansoury, M., H. Abdollahpouri, M. Pechenizkiy, B. Mobasher, and R. Burke. (2020). “Feedback Loop and Bias Amplification in Recommender Systems”. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. Virtual Event Ireland: ACM. 2145–2148. DOI: [10.1145/3340531.3412152](https://doi.org/10.1145/3340531.3412152).



- McDonald, G. and I. Ounis. (2020). “University of Glasgow Terrier Team at the TREC 2020 Fair Ranking Track”. In: *The Twenty-Ninth Text REtrieval Conference (TREC 2020) Proceedings*. Vol. 1266. NIST SP. URL: <https://trec.nist.gov/pubs/trec29/papers/uogTr.FR.pdf>.
- Mehrabi, N., F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. (2019). “A survey on bias and fairness in machine learning”. *arXiv preprint arXiv:1908.09635*.
- Mehrotra, R., A. Anderson, F. Diaz, A. Sharma, H. Wallach, and E. Yilmaz. (2017). “Auditing Search Engines for Differential Satisfaction Across Demographics”. In: *Proceedings of the 26th International Conference on World Wide Web Companion. WWW '17 Companion*. Perth, Australia: International World Wide Web Conferences Steering Committee. 626–633. DOI: [10.1145/3041021.3054197](https://doi.org/10.1145/3041021.3054197).
- Mehrotra, R., J. McInerney, H. Bouchard, M. Lalmas, and F. Diaz. (2018). “Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems”. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management. CIKM '18*. Torino, Italy: Association for Computing Machinery. 2243–2251. DOI: [10.1145/3269206.3272027](https://doi.org/10.1145/3269206.3272027).
- Metaxa, D., M. A. Gan, S. Goh, J. Hancock, and J. A. Landay. (2021). “An Image of Society: Gender and Racial Representation and Impact in Image Search Results for Occupations”. *Proceedings of the ACM on Human-Computer Interaction*. 5(CSCW1): 1–23. DOI: [10.1145/3449100](https://doi.org/10.1145/3449100).
- Metzler, D. and T. Kanungo. (2008). “Machine Learned Sentence Selection Strategies for Query-Biased Summarization”. In: *SIGIR Learning to Rank Workshop*.
- Mitchell, S., E. Potash, S. Barocas, A. D’Amour, and K. Lum. (2020). “Algorithmic Fairness: Choices, Assumptions, and Definitions”. en. *Annual Review of Statistics and Its Application*. 8(Nov.). DOI: [10.1146/annurev-statistics-042720-125902](https://doi.org/10.1146/annurev-statistics-042720-125902).

- Mitchell, M., S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru. (2019). “Model Cards for Model Reporting”. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* '19*. Atlanta, GA, USA: Association for Computing Machinery. 220–229. DOI: [10.1145/3287560.3287596](https://doi.org/10.1145/3287560.3287596).
- Moffat, A. and J. Zobel. (2008). “Rank-biased Precision for Measurement of Retrieval Effectiveness”. *ACM Transactions on Information Systems*. 27(1): 2:1–2:27. DOI: [10.1145/1416950.1416952](https://doi.org/10.1145/1416950.1416952).
- Morik, M., A. Singh, J. Hong, and T. Joachims. (2020). “Controlling Fairness and Bias in Dynamic Learning-to-Rank”. In: *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York, NY, USA: Association for Computing Machinery. 429–438. DOI: [10.1145/3397271.3401100](https://doi.org/10.1145/3397271.3401100).
- Moulin, H. (2004). *Fair Division and Collective Welfare*. en. MIT Press. URL: <https://play.google.com/store/books/details?id=qQXtEnb2B2cC>.
- Mowshowitz, A. and A. Kawaguchi. (2002). “Assessing Bias in Search Engines”. en. *Information Processing & Management*. 38(1): 141–156. DOI: [10.1016/s0306-4573\(01\)00020-6](https://doi.org/10.1016/s0306-4573(01)00020-6).
- Nagel, R. (2021). “Went to Google (first mistake) to see...” URL: <https://twitter.com/rebeccanagle/status/1371535405942734849>.
- Narasimhan, H., A. Cotter, M. Gupta, and S. Wang. (2020). “Pairwise Fairness for Ranking and Regression”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 5248–5255. DOI: [10.1609/aaai.v34i04.5970](https://doi.org/10.1609/aaai.v34i04.5970).
- Nasr, M. and M. C. Tschantz. (2020). “Bidding Strategies with Gender Nondiscrimination Constraints for Online Ad Auctions”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 337–347. DOI: [10.1145/3351095.3375783](https://doi.org/10.1145/3351095.3375783).
- Nielek, R., A. Wierzbicki, A. Jatowt, and K. Tanaka. (2016). “Report on the WebQuality 2015 Workshop”. In: *ACM SIGIR Forum*. Vol. 50. 83–85.

- Ning, X. and G. Karypis. (2011). “SLIM: Sparse Linear Methods for Top-N Recommender Systems”. In: *Proceedings of the 2011 IEEE 11th International Conference on Data Mining. ICDM '11*. Washington, DC, USA: IEEE Computer Society. 497–506. DOI: [10.1109/ICDM.2011.134](https://doi.org/10.1109/ICDM.2011.134).
- Noble, S. U. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. en. NYU Press. URL: <https://market.android.com/details?id=book--ThDDwAAQBAJ>.
- Noriega-Campero, A., M. A. Bakker, B. Garcia-Bulle, and A. ? Pentland. (2019). “Active Fairness in Algorithmic Decision Making”. In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society. AIES '19*. Honolulu, HI, USA: Association for Computing Machinery. 77–83. DOI: [10.1145/3306618.3314277](https://doi.org/10.1145/3306618.3314277).
- O’Neil, C. (2017). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. en. Penguin Books. URL: <http://www.worldcat.org/oclc/1015602855>.
- Okonofua, J. A. and J. L. Eberhardt. (2015). “Two strikes: Race and the disciplining of young students”. *Psychological science*. 26(5): 617–624.
- Okonofua, J. A., G. M. Walton, and J. L. Eberhardt. (2016). “A Vicious Cycle: A Social-Psychological Account of Extreme Racial Disparities in School Discipline”. *Perspectives on psychological science: a journal of the Association for Psychological Science*. 11(3): 381–398.
- Olteanu, A., C. Castillo, F. Diaz, and E. Kiciman. (2019). “Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries”. en. *Frontiers in Big Data*. 2(July): 13. DOI: [10.3389/fdata.2019.00013](https://doi.org/10.3389/fdata.2019.00013).
- Otterbacher, J., J. Bates, and P. 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. CHI '17*. Denver, Colorado, USA: Association for Computing Machinery. 6620–6631. DOI: [10.1145/3025453.3025727](https://doi.org/10.1145/3025453.3025727).
- Pandey, S. and C. Olston. (2008). “Crawl Ordering by Search Impact”. In: *Proceedings of the 2008 International Conference on Web Search and Data Mining. WSDM '08*. Palo Alto, California, USA: Association for Computing Machinery. 3–14. DOI: [10.1145/1341531.1341535](https://doi.org/10.1145/1341531.1341535).

- Pariser, E. (2011). *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. en. Penguin.
- Pasquale, F. (2006). “Rankings, Reductionism, and Responsibility”. *Cleveland State Law Review*. URL: [https://heinonline.org/hol-cgi-bin/get\\_pdf.cgi?handle=hein.journals/clevslr54&section=10&casa\\_token=BILTyPvQ06wAAAAA:8Xa\\_ROPfsnjS9x52QtvwoCH1adCMbyENgLbwKrnWyscAtulNRpk9oKwAaegjm0oWDxiFFcVhr-0](https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/clevslr54&section=10&casa_token=BILTyPvQ06wAAAAA:8Xa_ROPfsnjS9x52QtvwoCH1adCMbyENgLbwKrnWyscAtulNRpk9oKwAaegjm0oWDxiFFcVhr-0).
- Patro, G. K., A. Chakraborty, A. Banerjee, and N. Ganguly. (2020). “Towards Safety and Sustainability: Designing Local Recommendations for Post-pandemic World”. In: *Fourteenth ACM Conference on Recommender Systems. RecSys '20*. Virtual Event, Brazil: Association for Computing Machinery. 358–367. DOI: [10.1145/3383313.3412251](https://doi.org/10.1145/3383313.3412251).
- Pelly, L. (2018). “Discover Weakly”. *The Baffler*. June.
- Pizzato, L., T. Rej, T. Chung, I. Koprinska, and J. Kay. (2010). “RECON: a reciprocal recommender for online dating”. In: *Proceedings of the fourth ACM conference on Recommender systems. RecSys '10*. Barcelona, Spain: Association for Computing Machinery. 207–214. DOI: [10.1145/1864708.1864747](https://doi.org/10.1145/1864708.1864747).
- Pleiss, G., M. Raghavan, F. Wu, J. Kleinberg, and K. Q. Weinberger. (2017). “On Fairness and Calibration”. In: *Proceedings of the 31st Conference on Neural Information Processing Systems*. URL: <https://proceedings.neurips.cc/paper/2017/file/b8b9c74ac526fffb2d39ab038d1cd7-Paper.pdf>.
- Raghavan, M., S. Barocas, J. Kleinberg, and K. Levy. (2020). “Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* '20*. Barcelona, Spain: Association for Computing Machinery. 469–481. DOI: [10.1145/3351095.3372828](https://doi.org/10.1145/3351095.3372828).
- Raj, A. and M. D. Ekstrand. (2022). “Measuring Fairness in Ranked Results: An Analytical and Empirical Comparison”. In: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM Press. DOI: [10.1145/3477495.3532018](https://doi.org/10.1145/3477495.3532018).

- Raj, A., A. Milton, and M. D. Ekstrand. (2021). “Pink for Princesses, Blue for Superheroes: The Need to Examine Gender Stereotypes in Kid’s Products in Search and Recommendations”. In: *Proceedings of the 5th International and Interdisciplinary Perspectives on Children and Recommender and Information Retrieval Systems: Search and Recommendation Technology through the Lens of a Teacher (KidRec 2021)*, co-located with ACM IDC 2021. URL: <http://arxiv.org/abs/2105.09296>.
- Raj, A., C. Wood, A. Montoly, and M. D. Ekstrand. (2020). “Comparing Fair Ranking Metrics”. Sept. arXiv: [2009.01311](https://arxiv.org/abs/2009.01311) [cs.IR].
- Rekabsaz, N., S. Kopeinik, and M. 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. SIGIR ’21*. Virtual Event, Canada: Association for Computing Machinery. 306–316. DOI: [10.1145/3404835.3462949](https://doi.org/10.1145/3404835.3462949).
- Rendle, S., C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. (2009). “BPR: Bayesian Personalized Ranking from Implicit Feedback”. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence. UAI ’09*. Arlington, Virginia, United States: AUAI Press. 452–461. URL: <http://dl.acm.org/citation.cfm?id=1795114.1795167>.
- Rich, E. (1979). “User Modeling via Stereotypes”. *Cognitive Science*. 3(4): 329–354. URL: <http://www.sciencedirect.com/science/article/B6W48-4FWF9GC-9/2/f924f793eb153d455893e8d39982ef45>.
- Rijsbergen, C. J. van. (1979). *Information Retrieval*. Butterworths.
- Robertson, S. E. (1977). “The Probability Ranking Principle in IR”. *Journal of Documentation*. 33(4): 294–304. DOI: [10.1108/eb026647](https://doi.org/10.1108/eb026647).
- Rochet, J.-C. and J. Tirole. (2003). “Platform Competition in Two-Sided Markets”. *Journal of the European Economic Association*. 1(4): 990–1029.

- Roegiest, A., A. Lipani, A. Beutel, A. Olteanu, A. Lucic, A.-A. Stoica, A. Das, A. Biega, B. Voorn, C. Hauff, D. Spina, D. Lewis, D. W. Oard, E. Yilmaz, F. Hasibi, G. Kazai, G. McDonald, H. Haned, I. Ounis, I. van der Linden, J. Garcia-Gathright, J. Baan, K. N. Lau, K. Balog, M. de Rijke, M. Sayed, M. Panteli, M. Sanderson, M. Lease, M. D. Ekstrand, P. Lahoti, and T. Kamishima. (2019). “FACTS-IR: Fairness, Accountability, Confidentiality, Transparency, and Safety in Information Retrieval”. *SIGIR Forum*. 53(2): 20–43. URL: <http://sigir.org/wp-content/uploads/2019/december/p020.pdf>.
- Rolf, E., T. T. Worledge, B. Recht, and M. Jordan. (2021). “Representation Matters: Assessing the Importance of Subgroup Allocations in Training Data”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by M. Meila and T. Zhang. Vol. 139. *Proceedings of Machine Learning Research*. PMLR. 9040–9051. URL: <https://proceedings.mlr.press/v139/rolf21a.html>.
- Rossetti, M., F. Stella, and M. Zanker. (2016). “Contrasting Offline and Online Results when Evaluating Recommendation Algorithms”. In: *Proceedings of the 10th ACM Conference on Recommender Systems. RecSys '16*. New York, NY, USA: ACM. 31–34. DOI: [10.1145/2959100.2959176](https://doi.org/10.1145/2959100.2959176).
- Roth, A. E. (2015). *Who Gets What and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt.
- Rothstein, R. (2017). *The Color of Law: A Forgotten History of How Our Government Segregated America*. en. Liveright Publishing. URL: <https://play.google.com/store/books/details?id=SdtDDQAAQBAJ>.
- Salimi, B., L. Rodriguez, B. Howe, and D. Suciu. (2019). “Interventional Fairness: Causal Database Repair for Algorithmic Fairness”. In: *Proceedings of the 2019 International Conference on Management of Data. SIGMOD '19*. Amsterdam, Netherlands: Association for Computing Machinery. 793–810. DOI: [10.1145/3299869.3319901](https://doi.org/10.1145/3299869.3319901).
- Salton, G., A. Wong, and C. S. Yang. (1975). “A Vector Space Model for Automatic Indexing”. *Communications of the ACM*. 18(11): 613–620. DOI: [10.1145/361219.361220](https://doi.org/10.1145/361219.361220).

- Sambasivan, N., E. Arnesen, B. Hutchinson, and V. Prabhakaran. (2020). “Non-portability of Algorithmic Fairness in India”. Dec. arXiv: [2012.03659 \[cs.CY\]](https://arxiv.org/abs/2012.03659).
- Sánchez-Monedero, J., L. Dencik, and L. Edwards. (2020). “What does it mean to ‘solve’ the problem of discrimination in hiring? social, technical and legal perspectives from the UK on automated hiring systems”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT\* ’20*. Barcelona, Spain: Association for Computing Machinery. 458–468. DOI: [10.1145/3351095.3372849](https://doi.org/10.1145/3351095.3372849).
- Santos, R. L. T., C. Macdonald, and I. Ounis. (2015). *Search Result Diversification*. Vol. 9. *Foundations and Trends in Information Retrieval*. Hanover, MA, USA: Now Publishers Inc.
- Santos, R. L. T., J. Peng, C. Macdonald, and I. Ounis. (2010). “Explicit Search Result Diversification through Sub-queries”. In: *ECIR 2010: Advances in Information Retrieval*. Vol. 5993. *LNCS*. Springer Berlin Heidelberg. 87–99. DOI: [10.1007/978-3-642-12275-0\\_11](https://doi.org/10.1007/978-3-642-12275-0_11).
- Sapiezynski, P., W. Zeng, R. E Robertson, A. Mislove, and C. Wilson. (2019). “Quantifying the Impact of User Attention on Fair Group Representation in Ranked Lists”. In: *Companion Proceedings of The 2019 World Wide Web Conference. WWW ’19*. San Francisco, USA: Association for Computing Machinery. 553–562. DOI: [10.1145/3308560.3317595](https://doi.org/10.1145/3308560.3317595).
- Schein, A. I., A. Popescul, L. H. Ungar, and D. M. Pennock. (2002). “Methods and Metrics for Cold-start Recommendations”. In: *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’02*. Tampere, Finland: ACM. 253–260. DOI: [10.1145/564376.564421](https://doi.org/10.1145/564376.564421).
- Schuler, D. and A. Namioka, eds. (1993). *Participatory Design: Principles and Practices*. English. Hillsdale, N.J: CRC / Lawrence Erlbaum Associates. URL: [http://www.amazon.com/Participatory-Design-Principles-Douglas-Schuler/dp/0805809511?ie=UTF8&keywords=participatory%20design&qid=1460985157&ref\\_=sr\\_1\\_2&sr=8-2](http://www.amazon.com/Participatory-Design-Principles-Douglas-Schuler/dp/0805809511?ie=UTF8&keywords=participatory%20design&qid=1460985157&ref_=sr_1_2&sr=8-2).

- Selbst, A. D., D. Boyd, S. A. Friedler, S. Venkatasubramanian, and J. Vertesi. (2019). “Fairness and Abstraction in Sociotechnical Systems”. en. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency. FAT\* '19*. Atlanta, GA, USA: Association for Computing Machinery. 59–68. DOI: [10.1145/3287560.3287598](https://doi.org/10.1145/3287560.3287598).
- Singh, A. and T. Joachims. (2018). “Fairness of Exposure in Rankings”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. London, United Kingdom: ACM. 2219–2228. DOI: [10.1145/3219819.3220088](https://doi.org/10.1145/3219819.3220088).
- Singh, A. and T. Joachims. (2019). “Policy Learning for Fairness in Ranking”. In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc. 5426–5436. URL: <http://papers.nips.cc/paper/8782-policy-learning-for-fairness-in-ranking.pdf>.
- Singh, V. K., M. Chayko, R. Inamdar, and D. Floegel. (2020). “Female librarians and male computer programmers? Gender bias in occupational images on digital media platforms”. en. *Journal of the Association for Information Science and Technology*. 71(11): 1281–1294. DOI: [10.1002/asi.24335](https://doi.org/10.1002/asi.24335).
- Singhal, A., C. Buckley, and M. Mitra. (2017). “Pivoted Document Length Normalization”. *SIGIR Forum*. 51(2): 176–184. DOI: [10.1145/3130348.3130365](https://doi.org/10.1145/3130348.3130365).
- Smith, J., N. Sonbolil, C. Fiesler, and R. Burke. (2020). “Exploring User Opinions of Fairness in Recommender Systems”. In: *Fair & Responsible AI Workshop @ CHI 2020*. URL: <https://fair-ai.owlstown.net/publications/1459>.
- Smucker, M. D. and J. Allan. (2006). “Find-Similar: Similarity Browsing as a Search Tool”. In: *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '06*. Seattle, Washington, USA: Association for Computing Machinery. 461–468. DOI: [10.1145/1148170.1148250](https://doi.org/10.1145/1148170.1148250).
- Sonboli, N. (2022). “Controlling the Fairness / Accuracy Tradeoff in Recommender Systems”. *PhD thesis*. University of Colorado.



- Sonboli, N., R. Burke, N. Mattei, F. Eskandanian, and T. Gao. (2020a). ““And the Winner Is...”: Dynamic Lotteries for Multi-group Fairness-Aware Recommendation”. arXiv: [2009.02590](https://arxiv.org/abs/2009.02590) [cs.IR].
- Sonboli, N., F. Eskandanian, R. Burke, W. Liu, and B. Mobasher. (2020b). “Opportunistic Multi-aspect Fairness through Personalized Re-ranking”. In: *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization. UMAP '20*. Genoa, Italy: Association for Computing Machinery. 239–247. DOI: [10.1145/3340631.3394846](https://doi.org/10.1145/3340631.3394846).
- Sonboli, N., J. J. Smith, F. Cabral Berenfus, R. Burke, and C. Fiesler. (2021). “Fairness and Transparency in Recommendation: The Users’ Perspective”. In: *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization. UMAP '21*. Utrecht, Netherlands: Association for Computing Machinery. 274–279. DOI: [10.1145/3450613.3456835](https://doi.org/10.1145/3450613.3456835).
- Speer, R. (2017). “ConceptNet Numberbatch 17.04: better, less-stereotyped word vectors”. URL: <https://blog.conceptnet.io/2017/04/24/conceptnet-numberbatch-17-04-better-less-stereotyped-word-vectors/>.
- Srivastava, M., H. Heidari, and A. Krause. (2019). “Mathematical Notions vs. Human Perception of Fairness: A Descriptive Approach to Fairness for Machine Learning”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '19*. Anchorage, AK, USA: Association for Computing Machinery. 2459–2468. DOI: [10.1145/3292500.3330664](https://doi.org/10.1145/3292500.3330664).
- Steck, H. (2018). “Calibrated Recommendations”. In: *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM. 154–162. DOI: [10.1145/3240323.3240372](https://doi.org/10.1145/3240323.3240372).
- Steenkamp, J.-B. E. M. and H. Baumgartner. (1998). “Assessing Measurement Invariance in Cross-National Consumer Research”. en. *The Journal of consumer research*. 25(1): 78–90. DOI: [10.1086/209528](https://doi.org/10.1086/209528).
- Stray, J. (2020). “Aligning AI Optimization to Community Well-Being”. *International Journal of Community Well-Being*. 3(4): 443–463. DOI: [10.1007/s42413-020-00086-3](https://doi.org/10.1007/s42413-020-00086-3).

- Sun, W., O. Nasraoui, and P. Shafto. (2018). “Iterated Algorithmic Bias in the Interactive Machine Learning Process of Information Filtering”. In: *Proceedings of the 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*. Seville, Spain: SCITEPRESS - Science and Technology Publications. 108–116. DOI: [10.5220/0006938301100118](https://doi.org/10.5220/0006938301100118).
- Sun, Z., D. Yu, H. Fang, J. Yang, X. Qu, J. Zhang, and C. Geng. (2020). “Are We Evaluating Rigorously? Benchmarking Recommendation for Reproducible Evaluation and Fair Comparison”. In: *Fourteenth ACM Conference on Recommender Systems. RecSys '20*. Virtual Event, Brazil: Association for Computing Machinery. 23–32. DOI: [10.1145/3383313.3412489](https://doi.org/10.1145/3383313.3412489).
- Suresh, H. and J. V. Gutttag. (2019). “A Framework for Understanding Unintended Consequences of Machine Learning”. Jan. arXiv: [1901.10002 \[cs.LG\]](https://arxiv.org/abs/1901.10002).
- Takács, G. and D. Tikk. (2012). “Alternating Least Squares for Personalized Ranking”. In: *Proceedings of the Sixth ACM conference on Recommender Systems. RecSys '12*. Dublin, Ireland: Association for Computing Machinery. 83–90. DOI: [10.1145/2365952.2365972](https://doi.org/10.1145/2365952.2365972).
- Thelwall, M. (2019). “Reader and Author Gender and Genre in GoodReads”. *Journal of Librarianship and Information Science*. 51(2): 403–430. DOI: [10.1177/0961000617709061](https://doi.org/10.1177/0961000617709061).
- Thomson, W. (2016). “Introduction to the Theory of Fair Allocation”. In: *Handbook of Computational Social Choice*. Ed. by F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. Cambridge University Press. 261–283. DOI: [10.1017/CBO9781107446984.012](https://doi.org/10.1017/CBO9781107446984.012).
- Thune, J. (2019). “S.2763 Filter Bubble Transparency Act”. URL: <https://www.congress.gov/bill/116th-congress/senate-bill/2763>.
- Tintarev, N. and J. Masthoff. (2007). “A Survey of Explanations in Recommender Systems”. In: *IEEE ICDE Workshop*. IEEE Computer Society. 801–810. DOI: [10.1109/ICDEW.2007.4401070](https://doi.org/10.1109/ICDEW.2007.4401070).

- Tombros, A. and M. Sanderson. (1998). “Advantages of Query Biased Summaries in Information Retrieval”. In: *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '98*. Melbourne, Australia: Association for Computing Machinery. 2–10. DOI: [10.1145/290941.290947](https://doi.org/10.1145/290941.290947).
- Vaughan, L. and M. Thelwall. (2004). “Search Engine Coverage Bias: Evidence and Possible Causes”. *Information Processing & Management*. 40(4): 693–707. DOI: [10.1016/S0306-4573\(03\)00063-3](https://doi.org/10.1016/S0306-4573(03)00063-3).
- Vaughan, L. and Y. Zhang. (2007). “Equal Representation by Search Engines? A Comparison of Websites across Countries and Domains”. *Journal of computer-mediated communication: JCMC*. 12(3): 888–909. DOI: [10.1111/j.1083-6101.2007.00355.x](https://doi.org/10.1111/j.1083-6101.2007.00355.x).
- Voorhees, E. M. (2001). “The Philosophy of Information Retrieval Evaluation”. en. In: *Evaluation of Cross-Language Information Retrieval Systems*. Ed. by C. Peters, M. Braschler, J. Gonzalo, and M. Kluck. *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 355–370. URL: [http://link.springer.com/chapter/10.1007/3-540-45691-0\\_34](http://link.springer.com/chapter/10.1007/3-540-45691-0_34).
- Wu, Y., J. Cao, G. Xu, and Y. Tan. (2021). “TFROM: A Two-sided Fairness-Aware Recommendation Model for Both Customers and Providers”. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '21*. Virtual Event, Canada: Association for Computing Machinery. 1013–1022. DOI: [10.1145/3404835.3462882](https://doi.org/10.1145/3404835.3462882).
- Wyden, R. (2019). “S.1108 Algorithmic Accountability Act of 2019”. URL: <https://www.congress.gov/bill/116th-congress/senate-bill/1108>.
- Xiang, A. and I. D. Raji. (2019). “On the Legal Compatibility of Fairness Definitions”. Nov. arXiv: [1912.00761 \[cs.CY\]](https://arxiv.org/abs/1912.00761).
- Xie, X., J. Mao, Y. Liu, M. de Rijke, Y. Shao, Z. Ye, M. Zhang, and S. Ma. (2019). “Grid-Based Evaluation Metrics for Web Image Search”. In: *The World Wide Web Conference. WWW '19*. San Francisco, CA, USA: Association for Computing Machinery. 2103–2114. DOI: [10.1145/3308558.3313514](https://doi.org/10.1145/3308558.3313514).

- Xu, D., Y. Wu, S. Yuan, L. Zhang, and X. Wu. (2019). “Achieving Causal Fairness Through Generative Adversarial Networks”. In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- Yang, K., J. R. Loftus, and J. Stoyanovich. (2020). “Causal Intersectionality for Fair Ranking”. June. arXiv: [2006.08688](https://arxiv.org/abs/2006.08688) [cs.LG].
- Yang, K. and J. Stoyanovich. (2017). “Measuring Fairness in Ranked Outputs”. In: *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. No. Article 22. Chicago, IL, USA: ACM. 1–6. DOI: [10.1145/3085504.3085526](https://doi.org/10.1145/3085504.3085526).
- Yang, K., J. Stoyanovich, A. Asudeh, B. Howe, H. V. Jagadish, and G. Miklau. (2018). “A Nutritional Label for Rankings”. en. In: *Proceedings of the 2018 International Conference on Management of Data - SIGMOD '18*. Houston, TX, USA: ACM Press. 1773–1776. DOI: [10.1145/3183713.3193568](https://doi.org/10.1145/3183713.3193568).
- Yao, S. and B. Huang. (2017). “Beyond Parity: Fairness Objectives for Collaborative Filtering”. In: *Advances in Neural Information Processing Systems 30*. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Curran Associates, Inc. 2925–2934. URL: <http://papers.nips.cc/paper/6885-beyond-parity-fairness-objectives-for-collaborative-filtering.pdf>.
- Yue, Y., R. Patel, and H. Roehrig. (2010). “Beyond Position Bias: Examining Result Attractiveness as a Source of Presentation Bias in Clickthrough Data”. In: *Proceedings of the 19th International Conference on World Wide Web. WWW '10*. Raleigh, North Carolina, USA: Association for Computing Machinery. 1011–1018. DOI: [10.1145/1772690.1772793](https://doi.org/10.1145/1772690.1772793).
- Zafar, M. B., I. Valera, M. Gomez Rodriguez, and K. P. Gummadi. (2017). “Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment”. In: *Proceedings of the 26th International Conference on World Wide Web. WWW '17*. Perth, Australia: International World Wide Web Conferences Steering Committee. 1171–1180. DOI: [10.1145/3038912.3052660](https://doi.org/10.1145/3038912.3052660).

- Zehlike, M., F. Bonchi, C. Castillo, S. Hajian, M. Megahed, and R. Baeza-Yates. (2017). “FA\*IR: A Fair Top-k Ranking Algorithm”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management. CIKM '17*. ACM. 1569–1578. DOI: [10.1145/3132847.3132938](https://doi.org/10.1145/3132847.3132938).
- Zehlike, M., K. Yang, and J. Stoyanovich. (2022). “Fairness in Ranking, Part I: Score-based Ranking”. *ACM Computing Surveys*. Apr. DOI: [10.1145/3533379](https://doi.org/10.1145/3533379).
- Zemel, R., Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. (2013). “Learning Fair Representations”. In: *Proceedings of the 30th International Conference on Machine Learning*. Ed. by S. Dasgupta and D. McAllester. Vol. 28. *Proceedings of Machine Learning Research*. Atlanta, Georgia, USA: PMLR. 325–333. URL: <https://proceedings.mlr.press/v28/zemel13.html>.
- Zhang, J., G. Adomavicius, A. Gupta, and W. Ketter. (2020). “Consumption and Performance: Understanding Longitudinal Dynamics of Recommender Systems via an Agent-Based Simulation Framework”. *Information Systems Research*. 31(1): 76–101. DOI: [10.1287/isre.2019.0876](https://doi.org/10.1287/isre.2019.0876).
- Zhang, X., M. Khaliligarekani, C. Tekin, and M. Liu. (2019). “Group Retention when Using Machine Learning in Sequential Decision Making: the Interplay between User Dynamics and Fairness”. In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc. 15269–15278. URL: <http://papers.nips.cc/paper/9662-group-retention-when-using-machine-learning-in-sequential-decision-making-the-interplay-between-user-dynamics-and-fairness.pdf>.
- Zhao, X., Z. Niu, and W. Chen. (2013). “Opinion-Based Collaborative Filtering to Solve Popularity Bias in Recommender Systems”. In: *Database and Expert Systems Applications*. Springer Berlin Heidelberg. 426–433. DOI: [10.1007/978-3-642-40173-2\\_35](https://doi.org/10.1007/978-3-642-40173-2_35).

- Zheng, G., F. Zhang, Z. Zheng, Y. Xiang, N. J. Yuan, X. Xie, and Z. Li. (2018). “DRN: A Deep Reinforcement Learning Framework for News Recommendation”. In: *Proceedings of the 2018 World Wide Web Conference. WWW '18*. Lyon, France: International World Wide Web Conferences Steering Committee. 167–176. DOI: [10.1145/3178876.3185994](https://doi.org/10.1145/3178876.3185994).
- Zhu, Z., J. Kim, T. Nguyen, A. Fenton, and J. Caverlee. (2021). “Fairness among New Items in Cold Start Recommender Systems”. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '21*. Virtual Event, Canada: Association for Computing Machinery. 767–776. DOI: [10.1145/3404835.3462948](https://doi.org/10.1145/3404835.3462948).
- Ziegler, C.-N., S. M. McNee, J. A. Konstan, and G. Lausen. (2005). “Improving Recommendation Lists Through Topic Diversification”. In: *Proceedings of the 14th international conference on World Wide Web. WWW '05*. Chiba, Japan: Association for Computing Machinery. 22–32. DOI: [10.1145/1060745.1060754](https://doi.org/10.1145/1060745.1060754).

## Index

---

- accountability, 10, 69
- anti-classification, **55**
- anti-subordination, **55**, 112
- bias, **9**
  - defined for this work, 9
  - societal, **45**
  - statistical, **45**
- browsing model, **30**, 113
- calibrated fairness, 109, 110
- candidate search, 11
- Chouldechova-Kleinberg theorem, 54
- cold start, 83, 95
- collaborative filter, **32**
- collection, 23
- consumer, 21, 22, **25**, 66, 67, 71–73
- consumer fairness, **71**, 93
- corpus, *see* repository
- creation, 23
- discount function, **31**, 114
- discovery
  - job candidates, 11
  - jobs, 11
  - music, 12
  - news, 12
- disparate effectiveness, **97**
- disparate impact, 8, **51**, 117
- disparate treatment, **50**, 117
- diversity, **37**, 75, 90, 105, 121
- document, *see* item
- dominant group, *see* group, dominant
- embedding
  - item, **34**
  - user, **36**
- ethics, 11
- evaluation, **28**, 88, 95
- exposure, 103, **112**
- fairness, **9**, 15

- defined for this work, 9
- feedback loop, 86, 124–125
- four-fifths rule, 51
- group, **50**
  - dominant, **50**
  - majority, *see* dominant
  - protected, **50**, 77, 107
  - unprotected, *see* dominant
- group fairness, 41, **50**, 78, 95, 116
- harm
  - distributional, 39, **48**, 78, 106
  - representational, 39, **48**, 78, 106
- individual fairness, 41, **48**, 55, 78, 94, 113
- information access system, **7**, 21, 23, 61
- information need, 21, 22, **25**, 65
- intersectionality, **8**, 18, 102
- item, 21, 22, **23**, 64, 114
- job search, 11, 77
- limitations, 18
- majority group, *see* group, dominant
- matching, 98
- maximum marginal relevance, **37**, 122
- measurement invariance, 98
- meritocratic fairness, **53**, 118
- metadata, 24
- multisided platform, **70**
- music discovery, 12
- news discovery, 12
- parity, 51–54
  - calibration, **53**
  - error, **52**, 55
  - predictive value, **53**
  - recall, **53**
  - statistical, **51**, 55, 116
- philanthropic giving, 13
- policy, **32**
- privacy, 10
- process, 100
- protected group, *see* group, protected
- provider, **23**, 67, 73–74, 103
- provider fairness, **73**, 75, 90, 103
- query, **25**, 65
- ranking, **27**, 32, 64, 83, 105
- re-ranking, **37**, 99, 111, 118
- repository, 21, **23**
- representation
  - item, **24**, 87
  - user, **26**, 88
- results, 21, **26**
- rivalrous, *see* subtractable
- safety, 10
- satisfaction, 21, **28**
- scoring function, **32**, 82
- sensitive attribute, 8, **41**
- separate, simultaneous, and symmetric, 64



- session, 25, 26, **27**
- social constructionism, 17
- societal bias, *see* bias, societal
- stakeholders, 70–78
- statistical bias, *see* bias, statistical
- stochastic ranking, 113
- subject, 15, **75**, 121
- subtractability, 83
- subtractable, 64, 100
- system, **6**
- time, 86
- transparency, 10
- turn, **26**
- usage, 24
- user model, **35**
- utility, **28**, 64
- we’re all equal (WAE), **49**, 51, 54, 55
- what you see is what you get (WYSIWYG), **49**, 52, 54, 112