



FAIRS: a framework for rethinking algorithmic fairness in the context of information access

Emmanuel Sebastian Udoh, Xiaojun Yuan, Abebe Rorissa

DOI: <https://doi.org/10.47989/ir31262019>

Abstract

Introduction. Artificial intelligence and machine learning increasingly shape information access, enhancing efficiency but also amplifying biases that affect equity in access, exposure, and opportunity. This study introduces the FAIRS (‘Fair Algorithms for Information Retrieval and Seeking’) framework to address the pressing need for a comprehensive conceptualisation of algorithmic fairness in the context of information access.

Method. We build on a PRISMA-based literature review on scholarship relevant to algorithmic fairness and information access published between 2015 and 2025. The intersecting themes from over 100 peer-reviewed articles were triangulated to develop a conceptual framework.

Analysis. Thematic synthesis iteratively mapped themes connecting algorithmic fairness and information access across user roles, technologies, fairness issues, dimensions, solutions, metrics, bias, and contextual factors. These themes form the foundation of the FAIRS conceptual design.

Results. Existing research focuses on fairness in classification and rank-based personalisation. We determined that a comprehensive model of algorithmic fairness in information access must integrate metrics, context, barriers, technology, bias, fairness dimensions, and user perspectives—core components reflected in FAIRS.

Conclusions. FAIRS offers a novel, context-sensitive approach for defining, assessing, and operationalising fairness in information access. It provides a foundation for new fairness models, clarifies trade-offs, and supports the creation of holistic, equitable, and ethically grounded information systems.

Introduction

Artificial intelligence (AI) and machine learning (ML) increasingly shape online and public sector domains by enhancing efficiency, accuracy, and productivity while potentially mitigating human bias while reproducing and amplifying biases embedded in their training data, resulting in inequitable outcomes such as biased opportunity allocation, skewed search rankings, and information disparities (Gao & Shah, 2021; Joachims, 2021). Consequently, issues of privacy, accountability, equity, inclusion, and informed consent in automated decision-making become more complex, underscoring the urgent need to ensure algorithmic fairness and minimise bias.

Adoption of AI and ML also varies across regions and communities (Conn, 2010), shaped by global and national geopolitical dynamics (Yang & Chen, 2015). Internet governance policies and the inherent structures of AI/ML systems can perpetuate digital divides and other inequalities in information access (Yang & Chen, 2015). Thus, algorithmic fairness has emerged as a central concern in aligning technological advancement with ethical governance (Sonnenberg, 2020).

Despite growing attention, the literature lacks consensus on what constitutes ‘algorithmic fairness’ in the context of information access. This paper fills that gap by addressing the following research questions (RQs):

RQ1: What does algorithmic fairness mean in the context of information access?

RQ2: What themes and concepts should form the core of a comprehensive conceptualisation of algorithmic fairness in the context of information access?

Using a systematic (PRISMA-based) literature review (Page et al., 2021), we identify and synthesise key themes linking algorithmic fairness and information access to form the foundation of a conceptual framework. Ekstrand et al. (2022) reviewed fairness issues in information access systems, focusing primarily on user-item trade-offs despite acknowledging multiple stakeholders.

This paper proposes FAIRS (‘Fair Algorithms for Information Retrieval and Seeking’), a conceptual framework that integrates fairness at the design stage of information systems, rather than assessing it post-deployment. FAIRS also offers guiding questions for embedding fairness principles throughout system development and evaluation while emphasising that fairness is normative and context dependent.

Literature review

Literature search

Building on our earlier PRISMA review (Udoh, Yuan, & Rorissa, 2022), we revised our initial search query and expanded our search to include additional databases, including those that have become key venues for publishing on information access and algorithmic fairness. We used predetermined keywords with the query: Information AND (retriev* OR system*) AND (search* OR engine) AND Fair* AND “Information access.”

We searched eight databases in three rounds, limiting our inclusion to peer-reviewed full-text articles published from 2015 to 2025, written in English, and relevant to algorithmic fairness and information access. Titles and abstracts were screened using the predetermined inclusion/exclusion criteria, and the full texts of selected articles were reviewed to ensure they met these criteria.

From the initial collection of 635 articles, we removed 164 duplicates, resulting in 471 unique sources. The authors independently screened the titles and abstracts of these articles, removing 359 that were irrelevant to the review, selecting 112 articles for further screening, and including

103 in the final sample. Figure 1 summarises the search results and workflow for the literature search.




Figure 1. PRISMA literature search workflow.

Algorithmic fairness in information access has emerged as a central concern for information research, while remaining conceptually fragmented across technical, legal, and socio-cultural domains (Udoh et al., 2024). This fragmentation undermines the field's ability to evaluate and design systems that equitably mediate access to information, opportunities, and social participation (Ekstrand et al., 2019; Udoh et al., 2022). We synthesise recent PRISMA-based scholarship on algorithmic fairness and information access to examine key concepts, stakeholders, metrics, solutions, and their application in core technologies. These strands reveal a pressing need for a comprehensive, information-access-specific definition of algorithmic fairness that can support cumulative research and practice.

Dimensions of fairness

The literature converges on the view that AI/ML fairness is a multidimensional construct, though the dimensions are treated in isolation. A foundational distinction is between individual (that seeks to ensure that similar individuals are treated similarly) and group (focuses on comparable outcomes for socially salient groups defined by protected attributes such as gender, race, or disability) fairness (Barocas & Selbst, 2016; Dwork et al., 2012). Another recurring distinction is between fairness in treatment (intentional differentiation) and fairness in impact (which captures unintended but systematically unequal outcomes) (Barocas & Selbst, 2016). In algorithmic contexts, these are typically operationalised via statistical parity and related measures, which assess whether decisions or exposure to information are independent of group membership (Friedler et al., 2021). These concepts highlight that fairness cannot be reduced to the absence of explicit discrimination; instead, the distributional effects of ranking and recommendation policies should be considered.

The literature also identifies several further dimensions relevant to information research. Internet or ‘net’ neutrality concerns equitable access to online content and services, requiring that providers avoid blocking, throttling, or prioritising specific sites or applications (Conn, 2010). Network fairness examines whether users or applications receive a fair share of system resources and how network position can confer self-reinforcing advantages (Venkatasubramanian et al., 2021). Fairness in ranking and exposure focuses on how search and recommendation systems allocate visibility and attention, explicitly recognising that exposure is a resource (Balagopalan et al., 2023; Diaz et al., 2020). Finally, legal and policy dimensions of fairness are anchored in accessibility protections, such as Section 508 in the United States (U.S.), and cross-national regulatory comparisons (Sonnenberg, 2020; Yang & Chen, 2015).

These dimensions point to fairness as a layered construct encompassing equal treatment and impact, exposure and resource allocation, infrastructural conditions (net and network neutrality), and compliance with accessibility and anti-discrimination regimes. However, these layers rarely are integrated into a unified conceptualisation specific to information access.

Fairness issues and contexts involved

Algorithmic fairness spans multiple, interrelated issues across information systems, including bias, inclusion, network neutrality, and the broader harms of automated decision-making such as stereotyping, misrepresentation, and disenfranchisement. These concerns arise in contexts like internet access, search and recommendation systems, information quality, digital literacy, and metadata practices.

Inclusion remains a central challenge, particularly in equitable access to digital resources. Algorithmic filtering and prioritisation can undermine network neutrality and restrict access (Conn, 2010; Yang & Chen, 2015). Persistent disparities affect older adults, people with disabilities, and marginalised populations, reflecting gaps in both access and digital literacy, especially in developing regions (Gebremichael & Jackson, 2006; Yang & Chen, 2015).

Result diversity highlights how systems often reproduce dominant perspectives due to biased data and design. For instance, although women comprise 28% of U.S. chief executive officers (CEOs), they appear in only 10% of Google Image results for “CEO” (Kay et al., 2015). Efforts such as diversity-aware summarisation aim to mitigate this (Celis & Keswani, 2020), yet disparities persist, including in multilingual retrieval, where Bantu-language speakers may experience worse performance than English speakers (Chavula & Hussein, 2016).

Network fairness and net neutrality address how infrastructure and policy shape equitable access. Unequal network structures can amplify disparities through feedback loops, while policy

debates, such as U.S. net neutrality, demonstrate how access to essential information, including healthcare, can be constrained (Conn, 2010; Venkatasubramanian et al., 2021).

Algorithmic systems also shape visibility and opportunity through ranking and recommendation. Biases may stem from training data (exogenous) or broader systemic conditions (endogenous), producing allocational and representational harms (Blodgett et al., 2020; Rekabsaz et al., 2021). These dynamics align with concerns about algorithmic oppression, where systems reproduce inequality through mechanisms like stereotyping and redlining.

Bias and stereotyping persist due to unrealistic assumptions, such as completely protected-attribute data, and flawed modelling choices that skew exposure (Gao & Shah, 2021; Joachims, 2021; Kirnap et al., 2021). Inequities are compounded by differences in readability and information quality, disproportionately affecting users with lower digital literacy, disabilities, or age-related barriers (Hamwela et al., 2018; McKearney et al., 2018; Yang & Chen, 2015). Privacy-preserving behaviour may also lead to less dependable or more extreme recommendations (Spinelli & Crovella, 2020).

Accountability and transparency remain difficult to standardise. Fairness is measured through computational metrics, while accountability and transparency rely on audits and user studies, prompting efforts to develop clearer evaluative frameworks (Bernard & Balog, 2023). Interpretability adds another dimension: users frequently disagree with system outputs, for example, 81% disagreed with Alexa's top result in one study, highlighting the gap between system logic and user expectations (Dash et al., 2022).

Overall, fairness in information systems extends beyond accuracy to include equitable access, representation, interpretability, and the distribution of informational benefits and harms.

Stakeholders involved

The reviewed literature increasingly adopts stakeholder and multi-stakeholder perspectives, emphasising that those most affected by algorithmic systems should participate in their governance (Costa et al., 2020; Webb et al., 2018). Studies of patient-facing health information highlight misalignments between the content and presentation of online resources and the preferences and capacities of people with chronic conditions, such as low back pain, signalling the need for participatory design and evaluation (Costa et al., 2020; McKearney et al., 2018).

Stakeholders identified include technology firms, regulators, researchers, educators, librarians, civil society organisations, and end users, particularly marginalised communities such as people with disabilities, older adults, and users with lower literacy (Webb et al., 2018; Yang & Chen, 2015). In recommender systems, Sonboli et al. (2022) articulate a multi-stakeholder framework that distinguishes between consumer fairness (users), item fairness (content providers), and platform fairness (system operators), and argue that fairness objectives must be negotiated among these groups rather than specified unilaterally.

Within the information research community, librarianship and media ecosystems are positioned as active stakeholders in fairness. Critical librarianship initiatives call for exposing subject-heading bias, teaching search engine bias, and curating social-justice-oriented guides, while credibility models for news communities leverage interaction graphs to identify trustworthy sources and 'citizen journalists' (Bains, 2020; Mukherjee & Weikum, 2015). These strands suggest that algorithmic fairness in information access is not a purely technical property of systems, but also involves professional practice, governance, and collective responsibility.

Fairness metrics and frameworks used

Fairness metrics in information access have evolved from simple decision-level measures to more nuanced, exposure-oriented, and context-sensitive frameworks. In ranking contexts, Diaz

et al. (2020) define fairness as equal expected exposure for items, (or groups, of equal relevance, explicitly accounting for position bias and user attention. Kirnap et al. (2021) introduce unbiased estimators for fair ranking metrics that address incomplete judgments and logging bias, enabling more accurate evaluation of exposure fairness under realistic data constraints.

Beyond specific metrics, researchers propose frameworks for evaluating fairness in ranked lists. Sakai and colleagues advance group fairness and relevance (GFR) measures that manage nominal and ordinal groups, as well as soft group membership, and quantify how far exposure distributions depart from fairness targets (Sakai et al., 2023). Complementing this, Bernard and Balog (2023) provide a systematic review of fairness, accountability, transparency, and ethics in information retrieval and conclude that the field lacks standard definitions and calls for taxonomies of requirements that can be operationalised in evaluation (Bernard & Balog, 2023).

Where explicit protected attributes are unavailable, estimation and auditing approaches have been proposed. Quantification-based estimators approximate group composition and adjust outputs to promote fairer distributions, while counterfactual auditing evaluates whether individual-level outcomes change under hypothetical modifications to sensitive attributes (Cornacchia et al., 2023; Fabris et al., 2023). In recommender systems, multi-stakeholder fairness frameworks define distinct metrics for consumer, item, and platform fairness, and axiomatic models specify desired fairness properties for ranking systems, including search and job recommendations (Giner, 2023; Sonboli et al., 2022).

These developments underscore that fairness is increasingly measured at the level of exposure and attention, but that metric design is shaped by assumptions about context, group definitions, and stakeholder priorities. This reinforces the need for an overarching conceptualisation that clarifies how different metric families relate to each other in information access settings.

Algorithmic fairness solutions proffered

The identified algorithmic harms centre on inclusion and equity across protected characteristics such as race, gender, age, religion, national origin, marital status, socioeconomic status, and proxy variables like location.

Proposed solutions address fairness in search, retrieval, and recommendation systems through agnostic ranking models (Ai et al., 2023; Cornacchia et al., 2023; Rekabsaz et al., 2021; Rieger et al., 2024), distribution-aware recommender systems (Ekstrand et al., 2024; Diaz et al., 2020), diversity and multi-sided exposure methods (Balagopalan et al., 2023; Dhaliwal et al., 2021; Gao et al., 2022; Gao & Shah, 2020; Sonboli et al., 2022; Wu et al., 2022), image summarisation diversity (Celis & Keswani, 2020), and pairwise fairness measures (Fabris et al., 2023).

Model-based approaches address bias directly. For example, AdvBERT uses adversarial learning to remove protected attributes while preserving relevance, though proxy variables, such as location for race, can still reintroduce bias (Rekabsaz et al., 2021). Counterfactual auditing methods, such as PRET, instead assess fairness by testing outcome changes under altered sensitive attributes (Cornacchia et al., 2023).

Several fairness interventions engage librarianship and media ecosystems to preserve fairness and facts in the face of misinformation and disinformation. Critical librarianship calls for naming subject-heading bias, teaching search-engine bias, and curating social-justice guides, while credibility models for news communities infer credible articles, trustworthy sources, and expert 'citizen journalists' from interaction graphs (Bains, 2020; Mukherjee & Weikum, 2015).

Preprocessing decisions and evaluation strategies further shape fairness outcomes. Metadata decisions significantly affect information retrieval (IR) performance, while randomised evaluation and optimisation frameworks promote equal expected exposure and parity of attention (Diaz et

al., 2020; Roy et al., 2018). Several solutions focus on preserving fairness at the resource ranking level in search engines, including uncertainty-aware retrieval and unbiased estimation of fairness metrics from sparse labels, and improved tag recommendation using topic models (Guiver & Snelson, 2008; Kirnap et al., 2021; Krestel et al., 2009).

Accountability and transparency are advanced through system design and auditing. For instance, IBM's CLEVER (Continuous Learning for Evolving and Verifiable Embodied Representations) search system provides visibility into ranking pipelines, building on Kleinberg's (1999) HITS (or Hypertext Induced Topic Search) algorithm framework, while broader efforts combine audits and user studies to evaluate fairness (Kumar et al., 2006).

Another group of solutions that aim to ensure the fairness of algorithms within search engines includes cluster-based retrieval for improved access (Lamprier et al., 2010), spam-resilient ranking (Du et al., 2007), and fairness in advertising mechanisms such as keyword bidding (Rey & Kannan, 2010).

Context-aware and user-centred approaches further refine fairness, including models that incorporate task context, can better infer user intent to retrieve relevant and valuable information (Koskela et al., 2018), while search behaviour (e.g., literature search and data seeking) is more closely intertwined than previous literature suggests (Kramer et al., 2021).

Several sources frame fairness in terms of inclusion and accessibility to resources for people with disabilities. These accessibility solutions range from early communications stations for blind users (Rahimi & Eulenberg, 1974) to modern legal and policy audits of government compliance with Section 508 (Sonnenberg, 2020), and cross-national comparisons of the rules and laws of Internet accessibility protections, and calls for stronger private-sector accessibility mandates (Yang & Chen, 2015). Mobile and geospatial systems also enhance equitable access through integrated, location-based services (Sun et al., 2017).

Sociotechnical interventions also address misinformation and bias through critical librarianship and credibility modelling in news ecosystems (Bains, 2020; Mukherjee & Weikum, 2015). Fairness-aware rank fusion methods, such as WISE Fusion, aim to ensure proportional or equal representation in high-stakes contexts like hiring (Cachel & Rundensteiner, 2024), much more than the fairness and utility performance of state-of-the-art methods.

These interventions show that achieving fairness in information access in practice requires coordinated technical, evaluative, and institutional strategies.

Fairness and access technologies involved

Search engines, information retrieval systems, recommender systems, ranking algorithms and mechanisms, and websites are among information access technologies relevant to the discussion of algorithmic fairness.

In information access systems, fairness concerns whether benefits and resources are distributed equitably across separate groups (Ekstrand, Burke, & Diaz, 2019; Jeong, 2006). Fairness concerns could be raised regarding the algorithm, the data they utilise to produce results, the output, or the search results (Makri et al., 2021), and interfaces such as graphical displays used by visually impaired or blind users (Jeong, 2006). The potential for bias in search engines is significant, from the introduction or embedding of 'social, political, and moral values in their ranking functions' (Ekstrand, Burke, & Diaz, 2019, p. 7) to biased crawling and indexing, where webpages in specific languages could be under-indexed. Therefore, any solutions to algorithmic fairness issues in information access must directly address factors relevant to these technologies, including developing and applying fairness metrics that assess web usability, readability, and quality

(Hamwela, Ahmed, & Bath, 2018; McKearney et al., 2018), and are pivotal in objectively evaluating the fairness of search system ranking decisions (Kirnap et al., 2021; Rekabsaz et al., 2021).

The two major sectors directly impacted are online retail and streaming services, where the influence of issues such as biased training data on the fairness of these technologies is substantial (Joachims, 2021). The heightened scrutiny around search engines and fairness issues has prompted some to propose fairness evaluation metrics that consider protected attributes like gender or ethnicity (Kirnap et al., 2021), while others (Shah & Bender, 2024) advocate for reassessing the core information retrieval research focus to adequately address the tremendous impact algorithmic decision systems have on information access, individuals, and society. It suggests that algorithmic fairness cannot be fully understood without considering the interplay of data, algorithms, interfaces, and user characteristics across the entire information access pipeline.

Information access fairness

The literature explicitly connects algorithmic fairness to broader notions of information access as a human right rooted in equality, justice, and non-discrimination (Udoh et al., 2022, 2024). As AI and ML systems increasingly shape how individuals and communities search for, encounter, and interpret information, fairness in these systems becomes a prerequisite for equitable participation in knowledge production, public discourse, and everyday life (Ekstrand et al., 2019; Li et al., 2014).

However, fairness is currently conceptualised through multiple, partially overlapping lenses: statistical parity and disparate impact; exposure equality and diversity; accessibility and disability rights; stakeholder participation and commonsense expectations; and socio-technical accountability and transparency (Bauer & Schiele, 2024; Bernard & Balog, 2023). While this diversity reflects the multidimensional nature of fairness, it also creates a fragmented landscape in which definitions, metrics, and interventions are difficult to compare and aggregate.

Thus, while information research has developed rich work on specific fairness dimensions, metrics, and solutions, it lacks a comprehensive, field-specific definition of or a conceptual framework for algorithmic fairness in the context of information access (Udoh et al., 2024). Developing and debating such a framework within the Information Research community would support more coherent evaluation, design, and governance of information access systems in an increasingly algorithmic world.

FAIRS: A framework for rethinking algorithmic fairness in the context of information access

Several overlapping and intersecting themes emerged from the previous section, and we triangulated those findings with current themes not captured in the review. Again, the overlapping/emerging themes form the foundation for FAIRS ('Fair Algorithms for Information Retrieval and Seeking') that conceptualises algorithmic fairness in the context of information access (Figure 2). We present the framework before discussing crucial concepts that should constitute any conceptualisation of algorithmic fairness in the context of information access.




Figure 2. The FAIRS framework for conceptualising algorithmic fairness in the context of information access. Source: Authors' work.

Bias

Bias appears in the literature as a pervasive, overlapping theme that links data practices, modelling choices, and power relations in information access systems. It manifests in stereotyping, redlining, disenfranchisement, and 'dirty' or biased data, often through imbalances in user and item data, content production, and interaction logs that feed retrieval, ranking, and evaluation models (Ferguson, 2017; Noble, 2018; Richardson et al., 2019). Scholars argue that algorithmic systems rely on flawed assumptions, such as that data accurately reflects reality, that the future mirrors the past, and that algorithms are neutral, despite significant human judgment shaping their design and outcomes (Barocas & Selbst, 2016; Moses & Chan, 2016; Selbst et al., 2019).

This work converges on the view of bias as socio-technical: preexisting in institutions and hierarchies, technical in design and implementation, and emergent in use (Crawford, 2016; Friedman & Nissenbaum, 1996; Kroll et al., 2017). Moreover, algorithms align with and advance ideological worldviews (Beer, 2009, 2013; Danaher, 2016; Fuller & Goffey, 2012; Mackenzie, 2015).

Algorithmic systems both personalise and sort while also enabling discrimination, exclusion, and the reinforcement of racialised and gendered orders, as documented in search engines, social media feeds, and recommender systems (Beer, 2009; Gillespie, 2014, 2019; Noble, 2018; Pasquale, 2015; Tufekci, 2015). In information access specifically, this includes representational harms such as misrepresentation and stereotyping, as well as allocative harms tied to unequal exposure and opportunity, including in criminal justice and credit scoring (Ferguson, 2017; Pariser, 2012; Richardson et al., 2019; Wu et al., 2022). For example, the same search query entered by three

users, each from a different continent at the same time, tends to produce different sets or orders of results.

Across these strands, concerns about ‘algocracy’ foreground how hiddenness and opacity in data collection, modelling, and deployment undermine accountability, enabling algorithmic systems to act as de facto arbiters of public relevance and legitimacy (Danaher, 2016; Kroll et al., 2017; Peter, 2017; Selbst et al., 2019). The resulting picture, which FAIRS adopts, treats bias not as a purely statistical artifact but as a structuring condition of contemporary information ecosystems that must be confronted through contextual, multi-level interventions in data, models, interfaces, and institutions (Crawford, 2016; Friedman & Nissenbaum, 1996; Noble, 2018; Pasquale, 2015).

The user

The user is a key stakeholder in evaluating algorithmic fairness. In the context of information access, the question is ‘Fair to whom?’ To whom should the algorithmic decision system be fair? And how much is fair enough?

In different application domains of algorithmic systems, users face different fairness challenges. In job and candidate seeking search, whether the user receives a fair set of job opportunities in the recommendations or ads in their feed, whether the match or fit score is a fair score, does not redline or overvalue some candidates, whether the candidates have fair visibility when recruiters filter for top candidates (Geyik & Kenthapadi, 2018), and what fairness concerns come from regulatory requirements or procedures. Similarly, fairness in discoverability also directly impacts users’ news and music discoverability. For example, do the music and news items receive fair exposure and visibility across topics and genres, are users exposed to diverse content regardless of demographics or location, and are recommendations and search results free from biases affecting both users and creators (Epps-Darling, Bouyer, & Cramer, 2020)? The point is, ‘News search and recommendation influences user exposure to news articles on social media, news aggregation applications, and search engines’ (Ekstrand et al., 2022, p. 12). This tends to place users in filter bubbles, exposing them primarily to news items that reinforce their beliefs and deepen ideological polarisation (Alstyne & Brynjolfsson, 2005; Pariser, 2011). The bottom line is that information access systems ought to meet various objectives for multiple stakeholders.

The information access technology involved

Fairness discussions in information access also need to consider the various technologies involved, including search engines, information retrieval systems, recommender systems, ranking algorithms, and web-based interfaces as the core infrastructures through which algorithmic fairness issues materialise (Ekstrand et al., 2019; Makri et al., 2021). These systems are consistently framed as socio-technical pipelines, spanning crawling and indexing, learning-to-rank, personalisation, and interface design, whose design choices embed social, political, and moral values and can systematically skew exposure and opportunity across groups (Ekstrand et al., 2019; Joachims, 2021; Kirnap et al., 2021; Rekabsaz et al., 2021). A substantial subset of studies highlights modality- and context-specific access technologies, including assistive interfaces, mobile platforms, and large language model (LLM)-based systems, demonstrating how accessibility constraints, biased training data, and under-indexing of particular languages or regions exacerbate digital divides and representational harms (Chavula & Suleman, 2016; Jeong, 2006; Schmidt et al., 2010; Shah & Bender, 2024).

Collectively, this work motivates FAIRS’ treatment of information access technologies as a unifying axis in the literature, linking concrete systems and interfaces to fairness metrics, evaluation methods, and governance interventions that seek to ensure equitable exposure, usability, and information quality for diverse users and content providers (Diaz et al., 2020; Hamwela et al., 2018; McKearney et al., 2018; Sonboli et al., 2022; Udoh et al., 2022, 2024).

Fairness

Algorithmic fairness debates still hinge on which conception of fairness is assumed and what counts as ‘fair enough’ in applied settings, especially given the dominance of Euro-American legal and statistical traditions (Sambasivan et al., 2020; Udoh et al., 2024). Recent work argues for treating fairness as a realistic sufficiency threshold rather than an ideal of perfection, emphasising context-sensitive standards against which complex information access systems can be evaluated and governed (Ekstrand, 2022, 2024; Sambasivan et al., 2020). Drawing on FAIRS, we conceptualise algorithmic fairness as a multidimensional responsibility structure in which accountability, transparency, safety, privacy, and ethics specify interrelated but analytically distinct conditions for responsible information access (Bernard & Balog, 2023; Udoh et al., 2022, 2024).

Fairness as accountability links measured disparities in exposure, error, or relevance to concrete mechanisms for traceability, audit, and redress, making fairness metrics and audits actionable governance tools rather than purely diagnostic instruments (Bernard & Balog, 2023; Ekstrand, 2022; Selbst et al., 2019). Fairness as transparency concerns the scrutability of system purposes, data flows, and optimisation logics for affected stakeholders, including intelligible explanations of decisions, objectives, and trade-offs; without such transparency, neither informed consent nor meaningful contestation of unfair outcomes is possible (Bernard & Balog, 2023; Guo et al., 2023; Udoh et al., 2024).

Fairness as safety foregrounds obligations to anticipate, mitigate, and monitor informational harms, such as exposure to misinformation, discriminatory content, or unsafe recommendations, by integrating harm-sensitive objectives and pessimistic or worst-case evaluation methods that emphasise the most vulnerable users and providers rather than aggregate utility (Diaz, 2024; White & Hassan, 2014). Fairness as privacy highlights how unequal exposure to data collection, profiling, and inference constitutes an unfair burden, positioning privacy both as a constraint on fairness interventions and as a dimension of fairness concerned with whose data, queries, and interaction histories are most heavily exploited (Akter et al., 2020; Udoh et al., 2024).

Finally, fairness as ethics provides the normative grounding that connects these sub concepts to broader commitments to non-discrimination, equity, and respect for persons, extending beyond mere compliance toward critical reflection on how information access systems reproduce or resist structural injustices and representational harms (Gotterbarn et al., 2018; Noble, 2018; Shah & Bender, 2024; Sonboli et al., 2022).

Synthesising these strands, we define algorithmic fairness in the information access domain as:

Algorithmic fairness in the domain of information access is the context-sensitive responsibility of information systems to allocate exposure, relevance, and opportunities in ways that (1) avoid unjust, systematic disadvantages to individuals or groups, and (2) satisfy interrelated requirements of accountability, transparency, safety, privacy, and ethics across the system’s data, models, interfaces, and governance.

---1

Barriers to fair access

Some barriers to algorithmic fairness include historical and persistent discrimination, censorship, content divide, disproportionate access to infrastructure, deliberate digital misrepresentation, disparate distribution of content, insufficient legislation and regulation, and lack of net neutrality. Discussions on algorithmic fairness in information access must pay special attention to the fact that some groups are more susceptible to these barriers than others. In the United States, there is a long history of discrimination against black and minority groups dating back to the county poorhouses of the 19th century (Eubanks, 2018). These surreptitious redlining

activities persist online through the opaque algorithmic ranking, sorting, and recommendation systems.

Fairness metrics

The metrics used to determine what and how much is fair in the context of information access also need to be considered. There needs to be greater transparency about the sampling strategies used to determine who is included in the dataset, accompanied by detailed codebooks that not only define the variables but also specify how each variable is recorded, particularly for categorical variables. The selection and definition of variables are crucial, as observations or proxies must be valid and unbiased measures of the target construct to avoid measurement bias. A codebook reflects specific perspectives on how observations should be recorded or coded, which is extremely important in contexts where protected and personal variables are used.

Considering group fairness, group size, or the amount of representation can also contribute to unfairness (Rolf et al., 2021), because naive modelling is likely to be more accurate at predicting the majority group than the minority groups. This is often the root of representational or distributional harms online, where more data and examples from dominant geographical and cultural environments are used to train algorithms for filtering, recommendation, and ranking, and to determine how much exposure a resource gets and who has more access to it.

The context of implementation

Technologies do not exist in a vacuum; they are deployed in domains and environments coloured by culture, politics, economics, legal and policy framework, and social norms, among others. Therefore, fairness discourses in the context of information access need to focus on implementation contexts as well. Despite the pervasiveness of the Internet, net neutrality is not a common feature across all geopolitical zones, as policies, laws, and guidelines govern Internet use, digital information searches, and access to information. With censorship, for instance, entire groups and countries are denied unfettered access to Internet resources. In contrast, some geopolitical regions have no say in how digital resources are distributed and must contend not only with a digital divide but also a content divide.

Similarly, input data and the algorithmic decision results need to be collected and interpreted, considering the data's social and cultural context to avoid unfairly disregarding local knowledge, perspectives, societal norms, mores, and ethos, as well as bias. A vital direction would be to focus on the entire data pipeline, as each point in the pipeline is a potential vulnerability to unfairness and unfair access practices.

Discussion

Drawing on our previous work, we propose FAIRS, a framework for rethinking and studying the intersection of algorithmic fairness and information access.

Algorithmic fairness in information access is best understood as a multidimensional, context-sensitive responsibility of sociotechnical systems rather than a single technical property or metric (Ekstrand et al., 2019; Udoh et al., 2022, 2024). Building on our PRISMA-based review and the conceptual synthesis presented earlier, this section interprets what FAIRS contributes, where its boundaries lie, and which tensions and open questions it surfaces for future work (Bernard & Balog, 2023; Page et al., 2021).

Interpreting FAIRS as a responsibility structure

Most work on algorithmic fairness in information access has evolved around classification and rank-based personalisation, typically operationalised through group and individual fairness metrics, exposure constraints, and multi-stakeholder trade-offs (Barocas & Selbst, 2016; Diaz et al., 2020; Dwork et al., 2012; Giner, 2023; Sonboli et al., 2022). FAIRS reframes fairness not merely

as a distributional outcome, but as a responsibility structure that spans data, models, interfaces, and governance (Bernard & Balog, 2023; Ekstrand et al., 2022; Shah & Bender, 2024). This shift echoes philosophical accounts that treat algorithmic fairness as an institutional obligation to avoid unjustified differential treatment and impact, rather than as a purely statistical criterion (Wachter & Mittelstadt, 2019).

First, FAIRS integrates bias, user roles, technologies, barriers, metrics, and implementation context into a single conceptual map, thereby countering the fragmentation that separates ‘fairness in ranking,’ ‘accessibility,’ and ‘policy/regulation’ into parallel but rarely connected discourses (Conn, 2010; Ekstrand et al., 2019; Sonnenberg, 2020; Yang & Chen, 2015). Ekstrand et al. (2022) emphasise stakeholder and exposure trade-offs in information access; Selbst et al. (2019) foreground sociotechnical abstraction; and Burke (2017) distinguishes consumer- and provider-side fairness in recommender systems. FAIRS extends these contributions by insisting that any serious account of fairness in information access must treat bias, stakeholders, access technologies, barriers, metrics, and context as jointly constitutive, not optional add-ons (Friedman & Nissenbaum, 1996; Nissenbaum, 2001; Noble, 2018; Udoh et al., 2022, 2024).

Second, FAIRS makes explicit that accountability, transparency, safety, privacy, and ethics are analytically distinct but interdependent fairness sub concepts that define what responsible information access requires (Bernard & Balog, 2023; Makri et al., 2021; Rekabsaz et al., 2021). This aligns with recent work that maps fairness, accountability, transparency, and ethics (FATE) in information retrieval as overlapping but under-specified notions in need of clearer taxonomies and operational requirements (Bernard & Balog, 2023; Visier, 2023). By embedding these sub concepts at the core of the framework, FAIRS links fairness metrics and audits to concrete governance demands, such as traceability of decisions, meaningful contestation, harm mitigation, and ethical orientation, rather than treating them as external constraints (Danaher, 2016; Gao & Shah, 2021; Joachims, 2021).

Third, FAIRS foregrounds the user not as a passive recipient of rankings and recommendations, but as a central stakeholder whose informational autonomy, exposure, and opportunities are directly shaped by algorithmic decisions across search, recommendation, and filtering systems (Ekstrand et al., 2022; Geyik & Kenthapadi, 2018; Kramer et al., 2021). This perspective resonates with work on filter bubbles, confirmation bias, and the homogenisation of recommendations, which documents how iterative feedback loops can entrench exposure patterns and worldviews (Alstyne & Brynjolfsson, 2005; Pariser, 2011; Tufekci, 2015). FAIRS situates these phenomena within a broader picture of distributive and representational harms, thus connecting experiential user harms to structural issues such as digital divides, accessibility failures, and network or net neutrality constraints (Gebremichael & Jackson, 2006; Hamwela et al., 2018; Venkatasubramanian et al., 2021; Yang & Chen, 2015).

These elements motivate a more formal, comprehensive definition of algorithmic fairness tailored to information access:

‘Algorithmic fairness in the information access domain may be defined as the context-sensitive and multi-stakeholder responsibility of sociotechnical systems that search, rank, recommend, filter, or present information to ensure equitable access, exposure, and opportunity, while minimising allocative and representational harms and upholding accountability, transparency, safety, privacy, ethics, accessibility, and appropriate governance across data, models, interfaces, and implementation contexts.’

---2

This definition builds on and extends existing accounts that emphasise fair opportunities for exposure and quality for both consumers and providers in information access systems (Burke, 2017; Ekstrand et al., 2019; Ekstrand et al., 2022).

Core tensions and trade-offs in FAIRS

While FAIRS provides a more holistic lens, it also foregrounds tensions that are not easily resolved within existing technical or policy toolkits and reveal why a responsibility-based conception of fairness is normatively attractive but practically demanding (Barocas & Selbst, 2016; Blodgett et al., 2020; Sambasivan et al., 2021).

Fair for whom versus fair how

The multi-stakeholder nature of information access systems means that ‘Fair for whom?’ and ‘Fair how?’ cannot be answered independently (Costa et al., 2020; Sonboli et al., 2022; Webb et al., 2018). Provider-side fairness, consumer-side fairness, and platform-side objectives (e.g., engagement, revenue, regulatory compliance) often come into conflict (Burke, 2017; Ekstrand et al., 2019). For instance, equalising exposure for underrepresented artists or news outlets may reduce short-term click-based relevance for some users, while optimising for maximum utility may exacerbate existing inequalities in visibility and opportunity (Ekstrand et al., 2022; Gao et al., 2022; Wu et al., 2022).

FAIRS do not resolve these conflicts; instead, it requires that stakeholders explicitly surface and govern them (Costa et al., 2020; Webb et al., 2018). The framework positions fairness metrics (e.g., equal expected exposure, pairwise fairness, demographic parity, calibration) as tools for articulating trade-offs rather than as stand-alone solutions (Diaz et al., 2020; Fabris et al., 2023; Giner, 2023; Kirnap et al., 2021). This stance directly acknowledges impossibility results showing that different fairness criteria cannot simultaneously be satisfied under realistic conditions and that value-laden choices among these criteria are unavoidable (Friedler et al., 2021; Wachter & Mittelstadt, 2019). In doing so, FAIRS invites designers and regulators to treat these conflicts as sites of negotiation shaped by legal, cultural, and ethical commitments, not purely technical optimisation problems (MacCarthy, 2019; Sonnenberg, 2020).

Fairness versus relevance and utility

Another central tension lies between fairness constraints and traditional information retrieval goals of relevance and utility (Guiver & Snelson, 2008; Lewandowski & Spree, 2011; Li et al., 2014; Menczer et al., 2004). Fair ranking approaches that constrain exposure or prioritise diversity can degrade conventional effectiveness metrics, while ranking systems optimised solely for relevance risk reinforcing dominant perspectives and undeserving marginalised users or content (Celis & Keswani, 2020; Chavula & Suleman, 2016; Kay et al., 2015). Studies of search engines and scientific information show that highly ranked results often privilege specific framings or communities, with lower-ranked results offering systematically different perspectives that remain effectively invisible to most users (Ekstrand et al., 2019; Li et al., 2014).

Within FAIRS, this tension is reframed as a normative design choice: whose utility, over what time horizon, and under which constraints count as relevant (Bauer & Schiele, 2024; Ekstrand et al., 2019). Distributionally informed evaluation and stochastic ranking approaches already move in this direction by considering exposure distributions and worst-case group performance, but they still leave open fundamental questions about how to integrate long-term social welfare, epistemic diversity, and harm reduction into “relevance” (Diaz et al., 2020; Ekstrand et al., 2024; Gao & Shah, 2020). FAIRS thus highlights the need to reconceive relevance itself as a fairness-aware construct, especially in high-stakes domains such as health, employment, and civic information, where unfair access can have material consequences (Joachims, 2021; McKearney et al., 2018).

Fairness versus privacy and data minimisation

FAIRS also illuminates a persistent friction between fairness interventions and privacy protections (Cornacchia et al., 2023; Fabris et al., 2023; Spinelli & Crovella, 2020). Group fairness metrics often require knowledge of protected attributes, yet many legal and ethical regimes encourage or mandate minimisation, obfuscation, or non-collection of such data (Barocas & Selbst, 2016; Joachims, 2021). Proxy-based or quantification-based approaches attempt to approximate group composition without direct identifiers, but they introduce their own risks of misclassification, opacity, and potential reidentification (Cornacchia et al., 2023; Fabris et al., 2023).

By treating privacy as both a constraint on and a dimension of fairness, FAIRS underscores that fairness interventions can themselves produce unfair burdens when they require invasive or asymmetric data collection from specific groups (Research Agenda for Algorithmic Fairness, 2022; Wachter & Mittelstadt, 2019). At the same time, privacy-preserving practices that indiscriminately suppress sensitive information may obscure systematic harms and impede targeted remediation (Blodgett et al., 2020; Noble, 2018). This creates a genuine governance dilemma: How to design accountability and audit mechanisms that can detect and address unfairness without deepening surveillance or eroding informational autonomy for already vulnerable communities (Danaher, 2016; Tufekci, 2015; Udoh et al., 2022, 2024).

Local context versus global portability

A further tension arises from the non-portability of fairness notions across cultural, legal, and infrastructural contexts. Much of the fairness discourse is grounded in Western, Euro-American understandings of discrimination, legal protections, and institutional capacities, which do not straightforwardly apply to contexts with different histories, regulatory environments, or social norms (Gebremichael & Jackson, 2006; Sambasivan et al., 2021). Work on the non-portability of algorithmic fairness in India, as well as cross-national comparisons of accessibility laws and net neutrality regimes, illustrates how seemingly neutral technical criteria can fail or even backfire when transplanted across domains (Conn, 2010; Sambasivan et al., 2021; Yang & Chen, 2015).

FAIRS directly incorporates implementation context as a constitutive element, raising an open question: How far can a single framework travel before it becomes an instrument of epistemic or regulatory colonialism (Noble, 2018; Udoh et al., 2022, 2024)? The emphasis on context-sensitivity, sociotechnical analysis, and multi-stakeholder governance is intended to mitigate this risk, yet it also implies that instantiations of FAIRS will, and must, diverge across settings (Costa et al., 2020; Webb et al., 2018). This variability complicates efforts to standardise fairness metrics or benchmarks globally and demands further work on reconciling local adaptation with cross-context comparability and accountability (Ekstrand et al., 2019).

Limitations of the FAIRS framework

The FAIRS framework, while more comprehensive than prevailing models, has several limitations that must be acknowledged.

First, FAIRS is primarily conceptual and synthesising: It identifies the necessary components and relations for thinking about fairness in information access, but does not specify operational metrics, algorithms, or governance procedures (Bernard & Balog, 2023; Ekstrand et al., 2019). It offers guiding questions rather than prescriptive answers, and its efficacy depends on how practitioners, policymakers, and communities translate its categories into concrete practices and constraints (Costa et al., 2020; MacCarthy, 2019). This leaves a gap between high-level responsibility claims and the everyday realities of engineering and product decisions, where competing incentives and resource constraints are acute (Gao & Shah, 2021; Shah & Bender, 2024).

Second, the framework centres algorithmic fairness in the information access domain, which may limit its immediate applicability to adjacent but distinct contexts such as credit scoring, criminal justice, or hiring platforms that are less obviously 'information access' systems yet still rely heavily on ranking and classification (Barocas & Selbst, 2016; Wachter & Mittelstadt, 2019). While many FAIRS components (e.g., bias, stakeholders, metrics, context) generalise, the specific emphasis on exposure, visibility, and access as primary goods may underemphasise other forms of harm or benefit in those domains (Pasquale, 2015; Zarsky, 2016).

Third, FAIRS is grounded in a particular scholarly canon and methodological approach, a PRISMA-based review of largely peer-reviewed, English-language literature from 2015 to 2025, plus additional triangulated sources (Page et al., 2021; Udoh et al., 2022, 2024). This corpus inevitably reflects selection biases, including the relative underrepresentation of non-English scholarship, practitioner know-how, and community-based or activist perspectives on algorithmic harms (Crawford, 2016; Noble, 2018). As such, FAIRS risks systematically underweighting certain experiences of unfairness (e.g., those documented primarily in grey literature or local advocacy) and overweighting concerns prominent in well-resourced academic and industrial research communities (Kontinen et al., 2022; Sambasivan et al., 2021).

Finally, the framework presupposes the feasibility and desirability of embedding fairness considerations at the design stage and throughout the system life cycle (Ekstrand et al., 2019; Shah & Bender, 2024). In practice, information access systems are often legacy, opaque, and tightly coupled with business models or infrastructural dependencies that resist substantial redesign (Danaher, 2016; Joachims, 2021). Retrofitting FAIRS into such environments may require political and economic changes that extend well beyond the control of technical teams or even individual organisations (Ekstrand et al., 2022; McCarthy, 2019). This limitation underscores that FAIRS is as much a normative agenda for governance and institutional reform as it is a conceptual framework for system design (Costa et al., 2020; Webb et al., 2018).

Open questions and future directions

The FAIRS framework raises a set of open questions that define a research and practice agenda for algorithmic fairness in information access.

1. Operationalising multi-dimensional fairness. How can accountability, transparency, safety, privacy, and ethics be jointly operationalised in ways that are measurable yet sensitive to context (Bernard & Balog, 2023; Ekstrand et al., 2019)? Existing taxonomies and audits typically focus on one or two dimensions at a time; FAIRS calls for integrated evaluation pipelines that can surface cross-dimensional trade-offs and conflicts (Makri et al., 2021; Research Agenda for Algorithmic Fairness, 2022).
2. Fairness-aware relevance modelling. What would it mean to build relevance models that incorporate fairness, diversity, and harm minimisation as first-class objectives, rather than post hoc constraints on ranking outputs (Diaz et al., 2020; Ekstrand et al., 2024)? This requires new theoretical formulations of utility and loss functions that reflect both user satisfaction and social equity, as well as new benchmarks that weight performance for historically marginalised groups more heavily than aggregate metrics (Chavula & Hussein, 2016; Kay et al., 2015; Wu et al., 2022).
3. Context-sensitive portability and governance. How can frameworks like FAIRS be adapted to local contexts without losing their critical edge, and how can cross-context accountability be maintained when fairness standards legitimately diverge (Sambasivan et al., 2021; Yang & Chen, 2015)? This question points to the need for participatory, multi-stakeholder governance structures capable of articulating and revising context-specific fairness norms over time (Costa et al., 2020; Udoh et al., 2022, 2024; Webb et al., 2018).

4. Privacy-preserving fairness audits. What technical and institutional arrangements can support robust fairness auditing under strong privacy and data minimisation constraints (Cornacchia et al., 2023; Fabris et al., 2023)? Promising directions include privacy-preserving statistics, federated audits, and the use of synthetic or counterfactual data, but these approaches raise their own questions about validity, interpretability, and governance (Blodgett et al., 2020; MacCarthy, 2019).

5. Lifecycle integration and institutional change. Finally, how can the responsibility-based view of fairness be integrated into organisational and regulatory processes across the system lifecycle, from data collection and model design to deployment, monitoring, and decommissioning (Ekstrand et al., 2019; Shah & Bender, 2024)? FAIRS implies that fairness cannot be relegated to a single team or phase; instead, it must be woven into incentives, accountability structures, and professional norms, echoing calls from ethical computing codes and sociotechnical analyses of AI governance (Danaher, 2016; Noble, 2018; Tufekci, 2015).

In sum, FAIRS advances the field by offering a domain-specific, responsibility-oriented definition and framework for algorithmic fairness in information access, while openly acknowledging that any serious pursuit of fairness must grapple with deep trade-offs, context dependencies, and structural constraints (Ekstrand et al., 2019; Sambasivan et al., 2021; Udoh et al., 2022, 2024). The framework is intended to provide shared language and structure for more rigorous, context-aware, and ethically grounded conversations among researchers, practitioners, regulators, and affected communities about what fairness in information access should mean and how it might be pursued in practice (Costa et al., 2020; Research Agenda for Algorithmic Fairness, 2022; Webb et al., 2018).

FAIRS answers the concern raised by Selbst et al. (2019) of the lack of a universally accepted and contextually sensitive definition of algorithmic fairness by giving a context-sensitive definition of algorithmic fairness for information access and by specifying the substantive elements that must be considered together: Bias, users, technologies, barriers, metrics, and implementation context. Unlike Burke (2017), which focuses on multisided fairness in recommender systems and especially provider- and consumer-side fairness, FAIRS expands the unit of analysis to the broader information access domain, including search, retrieval, ranking, recommender systems, accessibility, and governance. Most importantly, FAIRS defines fairness not only as exposure or allocation across stakeholders but as an interrelated responsibility structure spanning accountability, transparency, safety, privacy, and ethics across data, models, interfaces, and governance.

Conclusion

Algorithmic fairness in information access is best understood as a context-sensitive responsibility of information systems, rather than as a single metric or purely technical property. Existing work tends to fragment fairness across classification, ranking, and stakeholder perspectives, and demonstrates that this fragmentation obscures how bias, users, technologies, barriers, metrics, and implementation contexts interact to shape inequitable exposure and opportunity online.

In response, we introduce the FAIRS framework, which systematically integrates these elements into a single conceptual model designed specifically for information access systems. FAIRS reconceptualises fairness as a multidimensional responsibility structure—spanning accountability, transparency, safety, privacy, and ethics, that must be embedded from the earliest stages of system design through deployment, evaluation, and governance, rather than treated as an after-the-fact constraint. By making explicit the trade-offs among competing fairness notions and metrics and by centring multi-stakeholder and socio-technical contexts, the framework

provides a theoretically grounded foundation for defining, assessing, and operationalising fairness in search, retrieval, ranking, and recommender systems.

Debates about ‘Fair for whom?’ and ‘Fair how?’ cannot be resolved within narrow, metric-only discussions but require attention to structural barriers, historical discrimination, and the specific domains in which information systems are implemented. FAIRS therefore positions algorithmic fairness in information access as an ongoing governance project: One that demands transparent metrics, careful attention to protected and proxy attributes, and sensitivity to context to avoid unjust, systematic disadvantages for individuals and groups. In doing so, the framework establishes a basis for future domain-specific fairness models and for more holistic, ethically robust approaches to designing and regulating information access technologies.

Future work

Our Framework aims to establish the core constitutive concepts in the algorithmic fairness discourse, enriching and balancing the discourse and providing a better foundation for other research efforts, especially in the information access domain. With this foundational work, new fairness models applicable to and appropriate for the information access context become a viable research agenda. It also provides a reference point for AI, ML, and information technology practitioners, vendors, and researchers on the ethical implications and challenges of integrating AI and ML in information access spaces.

About the authors

Emmanuel Sebastian Udoh is a Visiting Assistant Professor in the Massry School of Business, University at Albany, State University of New York, USA. He received his Ph.D. from the University at Albany, and his research interests include algorethics, responsible artificial intelligence, machine learning, the Internet of Things, smart cities, data fairness, algorithmic fairness, and the value-sensitive design of algorithmic decision systems. He can be contacted at eudoh@albany.edu.

Xiaojun (Jenny) Yuan is an Associate Professor in the College of Emergency Preparedness, Homeland Security and Cybersecurity, University at Albany, State University of New York. She received her Ph.D. from Rutgers University, the State University of New Jersey, New Brunswick, NJ, USA, and her research interests include human information behaviour, user interface design and evaluation, voice-enabled intelligent agent, and socially vulnerable populations. She can be contacted at xyuan@albany.edu.

Abebe Rorissa is a Professor & Rene and Clara Said Director of the School of Information Sciences, University of Tennessee, Knoxville, USA. He received his Ph.D. from the University of North Texas and his research interests include multimedia information organization and retrieval; scaling of users’ information needs; use, acceptance, adoption, and impact of information and communication technologies, and data analytics. He can be contacted at arorissa@utk.edu.

References

- Ai, Q., Wang, X., & Bendersky, M. (2023, July). Metric-agnostic ranking optimization. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2669-2680). <https://doi.org/10.1145/3539618.3591935>

- Akter, T., Ahmed, T., Kapadia, A., & Swaminathan, S. M. (2020, October). Privacy considerations of the visually impaired with camera based assistive technologies: Misrepresentation, impropriety, and fairness. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (pp. 1-14). <https://doi.org/10.1145/3373625.3417003>
- Alstyne, M. V., & Brynjolfsson, E. (2005). Global village or cyber-Balkans? Modeling and measuring the integration of electronic communities. *Management Science*, 51(6) (June 2005), 851–868. <https://doi.org/10.1287/mnsc.1050.0363>
- Balogopalan, A., Jacobs, A. Z., & Biega, A. J. (2023, July). The role of relevance in fair ranking. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2650-2660). <https://doi.org/10.1145/3539618.3591933>
- Bains, J. (2020). Search engines and algorithms are biased. Here's why that matters: Librarians can take steps to address inequality and bias that result from using data historically weighted in favor of white men. *Information Outlook*, 24(1). https://scholarworks.sjsu.edu/sla_io_2020/1/ (Archived at https://web.archive.org/web/20240709201850/https://scholarworks.sjsu.edu/sla_io_2020/1/)
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732. <http://dx.doi.org/10.15779/Z38BG31>
- Bauer, M. W., & Schiele, B. (Eds.). (2024). *AI and common sense: Ambitions and frictions*. Taylor & Francis.
- Beer, D. (2009). Power through the algorithm? Participatory web cultures and the technological unconscious. *New Media and Society*, 11(6), 985–1002. <https://doi.org/10.1177/1461444809336551>
- Beer, D. (2013). Algorithms: Shaping tastes and manipulating the Circulations of popular culture. In *Popular culture and new media: The politics of circulation* (pp. 63–100). Palgrave Macmillan. https://doi.org/10.1057/9781137270061_4
- Bernard, N., & Balog, K. (2025). A systematic review of fairness, accountability, transparency, and ethics in information retrieval. *ACM Computing Surveys*, 57(6). <https://doi.org/10.1145/3637211>
- Blodgett, S. L., Barocas, S., Daumé III, H., & Wallach, H. (2020). Language (technology) is power: A critical survey of "bias" in NLP. *arXiv*. <https://doi.org/10.48550/arXiv.2005.14050>
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 29). Curran Associates, Inc.
- Burke, R. (2017, July). Multisided fairness for recommendation. *arXiv*. <https://doi.org/10.48550/arXiv.1707.00093>
- Cachel, K., & Rundensteiner, E. (2024, October). Wise fusion: Group fairness enhanced rank fusion. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management* (pp. 163-174). <https://doi.org/10.1145/3627673.3679649>
- Celis, L. E., & Keswani, V. (2020). Implicit diversity in image summarization. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), 1-28. <https://doi.org/10.1145/3415210>

- Chaney, A. J. B., Stewart, B. M., & Engelhardt, B. E. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems* (pp. 224–232). Association for Computing Machinery. <https://doi.org/10.1145/3240323.3240370>
- Chavula, C., & Suleman, H. (2016, December). Assessing the impact of vocabulary similarity on multilingual information retrieval for Bantu languages. In *Proceedings of the 8th Annual Meeting of the Forum for Information Retrieval Evaluation* (pp. 16–23). <https://doi.org/10.1145/3015157.3015160>
- Conn, J. (2010). Still stuck on neutrality. No legislative remedy likely over Internet fair play. *Modern Healthcare*, 40(41), 30–31. <https://pubmed.ncbi.nlm.nih.gov/21033002/> (Archived at <https://web.archive.org/web/20260422193955/https://pubmed.ncbi.nlm.nih.gov/21033002/>)
- Cornacchia, G., Anelli, V. W., Biancofiore, G. M., Narducci, F., Pomo, C., Ragone, A., & Di Sciascio, E. (2023). Auditing fairness under unawareness through counterfactual reasoning. *Information Processing & Management*, 60(2), 103224. <https://doi.org/10.1016/j.ipm.2022.103224>
- Costa, N., Nielsen, M., Jull, G., Claus, A. P., & Hodges, P. W. (2020). Low back pain websites do not meet the needs of consumers: A study of online resources at three time points. *Health Information Management Journal*, 49(2–3), 137–149. <https://doi.org/10.1177/1833358319857354>
- Crawford, K. (2016). Can an algorithm be agonistic? Ten scenes from life in calculated publics. *Science, Technology & Human Values*, 41(1), 77–92. <https://doi.org/10.1177/0162243915589635>
- Crawford, K. (2017, December). *The trouble with bias* [Video]. YouTube. https://youtu.be/fMym_BKWQzk (Archived at https://web.archive.org/web/20251220102224/https://www.youtube.com/watch?v=fMym_BKWQzk)
- Danaher, J. (2016). The threat of algocracy: Reality, resistance and accommodation. *Philosophy & Technology*, 29(3), 245–268. <https://doi.org/10.1007/s13347-015-0211-1>
- Dash, A., Chakraborty, A., Ghosh, S., Mukherjee, A., & Gummadi, K. P. (2022). Alexa, in you, I trust! Fairness and interpretability issues in e-commerce search through smart speakers. In *Proceedings of the ACM Web Conference 2022* (pp. 3695–3705). Association for Computing Machinery. <https://doi.org/10.1145/3485447.3512265>
- Dhaliwal, M. P., Muskaan, & Seth, A. (2021). Fairness and diversity in the recommendation and ranking of participatory media content. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* (pp. 43–47). Association for Computing Machinery. <https://doi.org/10.1145/3487351.3488363>
- Diaz, F., Mitra, B., Ekstrand, M. D., Biega, A. J., & Carterette, B. (2020). Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management* (pp. 275–284). Association for Computing Machinery. <https://doi.org/10.1145/3340531.3411962>
- Diaz, F. (2024). Pessimistic evaluation. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*

- (pp. 115–124). Association for Computing Machinery.
<https://doi.org/10.1145/3673791.3698428>
- Du, Y., Shi, Y., & Zhao, X. (2007). Using spam farm to boost PageRank. In *Proceedings of the 3rd International Workshop on Adversarial Information Retrieval on the Web* (pp. 29–36). Association for Computing Machinery. <https://doi.org/10.1145/1244408.1244415>
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *Proceedings of the Innovations in Theoretical Computer Science Conference* (pp. 214–226). Association for Computing Machinery. <https://doi.org/10.1145/2090236.2090255>
- Ekstrand, M. D., Burke, R., & Diaz, F. (2019). Fairness and discrimination in recommendation and retrieval. In *Proceedings of the 13th ACM Conference on Recommender Systems* (pp. 576–577). Association for Computing Machinery. <https://doi.org/10.1145/3298689.3346964>
- Ekstrand, M. D., Das, A., Burke, R., & Diaz, F. (2022). Fairness in information access systems. *Foundations and Trends in Information Retrieval*, 16(1-2), 1-177.
<https://doi.org/10.48550/arXiv.2105.05779>
- Ekstrand, M. D., Carterette, B., & Diaz, F. (2024). Distributionally-informed recommender system evaluation. *ACM Transactions on Recommender Systems*, 2(1), 1-27.
<https://doi.org/10.1145/3613455>
- Epps-Darling, A., Bouyer, R. T., & Cramer, H. (2020). Artist gender representation in music streaming. In *Proceedings of the 21st International Society for Music Information Retrieval Conference* (pp. 248–254). <https://archives.ismir.net/ismir2020/paper/000148.pdf>
 (Archived at <https://web.archive.org/web/20260210055209/http://archives.ismir.net/ismir2020/paper/000148.pdf>)
- Eubanks, V. (2018, January). The digital poorhouse. *Harper's Magazine*.
<https://harpers.org/archive/2018/01/the-digital-poorhouse/> (Archived at <https://web.archive.org/web/20260210075441/https://harpers.org/archive/2018/01/the-digital-poorhouse/>)
- Fabris, A., Silvello, G., Susto, G. A., & Biega, A. J. (2023). Pairwise fairness in ranking as a dissatisfaction measure. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining* (pp. 931–939). Association for Computing Machinery.
<https://doi.org/10.1145/3539597.3570459>
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., & Venkatasubramanian, S. (2015). Certifying and removing disparate impact. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 259–268). Association for Computing Machinery. <https://doi.org/10.1145/2783258.2783311>
- Ferguson, A. G. (2017). *The rise of big data policing: Surveillance, race, and the future of law enforcement*. New York University Press.
- Friedler, S., Scheidegger, C., & Venkatasubramanian, S. (2021). The (im)possibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4), 136–143. <https://doi.org/10.1145/3433949>
- Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, 14(3), 330–347. <https://doi.org/10.1145/230538.230561>
- Fuller, M., & Goffey, A. (2012). *Algorithms*. In *Evil media* (pp. 69–82). MIT Press.

- Gao, R., & Shah, C. (2021). Addressing bias and fairness in search systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2643–2646). Association for Computing Machinery. <https://doi.org/10.1145/3404835.3462807>
- Gao, R., Ge, Y., & Shah, C. (2022). FAIR: Fairness-aware information retrieval evaluation. *Journal of the Association for Information Science and Technology*, 73(10), 1461–1473. <https://doi.org/10.1002/asi.24648>
- Gao, R., & Shah, C. (2020). Toward creating a fairer ranking in search engine results. *Information Processing & Management*, 57(1), 102138. <https://doi.org/10.1016/j.ipm.2019.102138>
- Gebremichael, M. D., & Jackson, J. W. (2006). Bridging the gap in Sub-Saharan Africa: A holistic look at information poverty and the region's digital divide. *Government Information Quarterly*, 23(2), 267–280. <https://doi.org/10.1016/j.giq.2006.02.011>
- Geyik, S. C., & Kenthapadi, K. (2018, October). Building representative talent search at LinkedIn. *LinkedIn Engineering Blog*. <https://engineering.linkedin.com/blog/2018/10/building-representative-talent-search-at-linkedin> (Archived at <https://web.archive.org/web/20260123093645/https://www.linkedin.com/blog/engineering/hiring/building-representative-talent-search-at-linkedin>)
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–193). MIT Press.
- Gillespie, T. (2016, February). #trendingistrending: When algorithms become culture. *Culture Digitally*. <http://culturedigitally.org/2016/02/trendingistrending/> (Archived at <https://web.archive.org/web/20260218143639/https://culturedigitally.org/2016/02/trendingistrending/>)
- Gillespie, T. (2019). Algorithmically recognizable: Santorum's Google problem, and Google's Santorum problem. In *The social power of algorithms* (pp. 63–80). Routledge.
- Giner, F. (2023). Information retrieval evaluation measures defined on some axiomatic models of preferences. *ACM Transactions on Information Systems*, 42(3), 1–35. <https://doi.org/10.1145/3632171>
- Gotterbarn, D. W., Brinkman, B., Flick, C., Kirkpatrick, M. S., Miller, K., Vazansky, K., & Wolf, M. J. (2018). ACM code of ethics and professional conduct. <https://dora.dmu.ac.uk/server/api/core/bitstreams/1e5b3cb8-2d77-4ab4-885b-d5996b74605f/content> (Archived at <https://web.archive.org/web/20250501011834/https://dora.dmu.ac.uk/server/api/core/bitstreams/1e5b3cb8-2d77-4ab4-885b-d5996b74605f/content>)
- Guiver, J., & Snelson, E. (2008). Learning to rank with Softrank and Gaussian processes. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 259–266). Association for Computing Machinery. <https://doi.org/10.1145/1390334.1390380>
- Guo, S., Zhang, S., Sun, W., Ren, P., Chen, Z., & Ren, Z. (2023). Towards explainable conversational recommender systems. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2786–2795). Association for Computing Machinery. <https://doi.org/10.1145/3539618.3591884>

- Hamwela, V., Ahmed, W., & Bath, P. A. (2018). Evaluation of websites that contain information relating to malaria in pregnancy. *Public Health*, 157, 50–52.
<https://doi.org/10.1016/j.puhe.2018.01.001>
- Hargittai, E. (2007). The social, political, economic, and cultural dimensions of search engines: An introduction. *Journal of Computer-Mediated Communication*, 12(3), 769–777.
<https://doi.org/10.1111/j.1083-6101.2007.00349.x>
- Introna, L. D., & Nissenbaum, H. (2000). Shaping the web: Why the politics of search engines matters. *The Information Society*, 16(3), 169–185.
<https://doi.org/10.1080/01972240050133634>
- Jeong, W. (2006). Force feedback textual and graphic displays for the blind. *Proceedings of the American Society for Information Science and Technology*, 43(1), 1–11.
<https://doi.org/10.1002/meet.14504301156>
- Joachims, T. (2021). Fairness and control of exposure in two-sided markets. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval* (pp. 1–1). Association for Computing Machinery. <https://doi.org/10.1145/3471158.3472226>
- Joachims, T., Granka, L., Pan, B., Hembrooke, H., & Gay, G. (2017). Accurately interpreting clickthrough data as implicit feedback. *SIGIR Forum*, 51(1), 4–11.
<https://doi.org/10.1145/3130332.313033>
- Kay, M., Matuszek, C., & Munson, S. A. (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* (pp. 3819–3828). Association for Computing Machinery. <https://doi.org/10.1145/2702123.270252>
- Kearns, M., & Roth, A. (2020). *The ethical algorithm: The science of socially aware algorithm design*. Oxford University Press.
- Kirnap, Ö., Diaz, F., Biega, A., Ekstrand, M., Carterette, B., & Yilmaz, E. (2021). Estimation of fair ranking metrics with incomplete judgments. In *Proceedings of the Web Conference 2021* (pp. 1065–1075). Association for Computing Machinery.
<https://doi.org/10.1145/3442381.345008>
- Kleinberg, J. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5), 604–632. <https://doi.org/10.1145/324133.324140>
- Kontinen, L., Koulu, R., & Sankari, S. (2022). Research agenda for algorithmic fairness studies: Access to justice lessons for interdisciplinary research. *Frontiers in Artificial Intelligence*, 5, 882134. <https://doi.org/10.3389/frai.2022.882134>
- Koskela, M., Luukkonen, P., Ruotsalo, T., Sjöberg, M., & Floréen, P. (2018). Proactive information retrieval by capturing search intent from primary task context. *ACM Transactions on Interactive Intelligent Systems*, 8(3), 1–25. <https://doi.org/10.1145/3150975>
- Krämer, T., Papenmeier, A., Carevic, Z., Kern, D., & Mathiak, B. (2021). Data-seeking behaviour in the social sciences. *International Journal on Digital Libraries*, 22, 175–195.
<https://doi.org/10.1007/s00799-021-00303-0>
- Krestel, R., Fankhauser, P., & Nejdl, W. (2009). Latent Dirichlet allocation for tag recommendation. In *Proceedings of the Third ACM Conference on Recommender Systems* (pp. 61–68). Association for Computing Machinery. <https://doi.org/10.1145/1639714.1639726>

- Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165, 633–705. <https://repository.law.upenn.edu/Documents/Detail/accountable-algorithms/155735> (Archived at <https://web.archive.org/web/20260218210455/https://repository.law.upenn.edu/Documents/Detail/accountable-algorithms/155735>)
- Kumar, R., Raghavan, P., Rajagopalan, S., & Tomkins, A. (2006). Core algorithms in the CLEVER system. *ACM Transactions on Internet Technology*, 6(2), 131–152. <https://doi.org/10.1145/1149121.1149123>
- Lamprier, S., Amghar, T., Saubion, F., & Levrat, B. (2010). Traveling among clusters: A way to reconsider the benefits of the cluster hypothesis. In *Proceedings of the 2010 ACM Symposium on Applied Computing* (pp. 1774–1780). Association for Computing Machinery. <https://doi.org/10.1145/1774088.1774465>
- Latzer, M., Hollnbuchner, K., Just, N., & Saurwein, F. (2014). *The economics of algorithmic selection on the internet* (Working Paper – Media Change & Innovation Division). University of Zurich. <https://doi.org/10.4337/9780857939852.00028>
- Lederman, N. G., & Lederman, J. S. (2015). What is a theoretical framework? A practical answer. *Journal of Science Teacher Education*, 26(7), 593–597. <https://doi.org/10.1007/s10972-015-9443-2>
- Lewandowski, D., & Spree, U. (2011). Ranking of Wikipedia articles in search engines revisited: Fair ranking for reasonable quality? *Journal of the American Society for Information Science and Technology*, 62(1), 117–132. <https://doi.org/10.1002/asi.21423>
- Li, N., Anderson, A. A., Brossard, D., & Scheufele, D. A. (2014). Channeling science information seekers' attention? A content analysis of top-ranked vs. lower-ranked sites in Google. *Journal of Computer-Mediated Communication*, 19(3), 562–575. <https://doi.org/10.1111/jcc4.12043>
- MacCarthy, M. (2019). *Fairness in algorithmic decision-making*. Brookings Institution. <https://www.brookings.edu/articles/fairness-in-algorithmic-decision-making/> (Archived at <https://web.archive.org/web/20260322175848/https://www.brookings.edu/articles/fairness-in-algorithmic-decision-making/>)
- Mackenzie, A. (2015). The production of prediction: What does machine learning want? *European Journal of Cultural Studies*, 18(4–5), 429–445. <https://doi.org/10.1177/1367549415577384>
- Makri, S., McKay, D., Buchanan, G., Chang, S., Lewandowski, D., MacFarlane, A., Cole, L., Vrijenhoek, S., & Ferraro, A. (2021). Search a great leveler? Ensuring more equitable information acquisition. *Proceedings of the Association for Information Science and Technology*, 58(1), 613–618. <https://doi.org/10.1002/pra2.511>
- Mansoury, M., Abdollahpouri, H., Pechenizkiy, M., Mobasher, B., & Burke, R. (2020). Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management* (pp. 2145–2148). Association for Computing Machinery. <https://doi.org/10.1145/3340531.3412152>
- McKearney, R. M., MacKinnon, R. C., Smith, M., & Baker, R. (2018). Tinnitus information online—does it ring true? *The Journal of Laryngology & Otology*, 132(11), 984–989. <https://doi.org/10.1017/S0022215118001792>

- Menczer, F., Pant, G., & Srinivasan, P. (2004). Topical web crawlers: Evaluating adaptive algorithms. *ACM Transactions on Internet Technology*, 4(4), 378–419.
<https://doi.org/10.1145/1031114.1031117>
- Moses, L. B., & Chan, J. (2018). Algorithmic prediction in policing: Assumptions, evaluation, and accountability. *Policing and Society*, 28(7), 806–822.
<https://doi.org/10.1080/10439463.2016.1253695>
- Moulin, H. (2004). *Fair division and collective welfare*. MIT press.
- Mukherjee, S., & Weikum, G. (2015). Leveraging joint interactions for credibility analysis in news communities. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management* (pp. 353–362). Association for Computing Machinery.
<https://doi.org/10.1145/2806416.2806537>
- Nahon, K. (2016). Where there is social media, there is politics. In A. Bruns, E. Skogerbø, C. Christensen, A. O. Larsson, & G. S. Enli (Eds.), *The Routledge companion to social media and politics* (pp. 39–55). Routledge. <https://doi.org/10.4324/9781315716299>
- Nissenbaum, H. (2001). How computer systems embody values. *Computer*, 34(3), 118–119.
<https://doi.org/10.1109/2.910905>
- Noble, S. (2012). Missed Connections: What Search Engines Say about Women. *Bitch magazine*, 12(4): 37–41.
https://www.academia.edu/1975319/Missed_Connections_What_Search_Engines_Say_About_Women
- Noble, S. (2018). *Algorithms of oppression: Race, gender, and power in the digital age*. NYU Press.
- O’Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Penguin Books.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., McGuinness, L. A., & Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Pariser, E. (2012). *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin Books.
- Park, Y. J., & Jang, S. M. (2016). African American Internet use for information search and privacy protection tasks. *Social Science Computer Review*, 34(5), 618–630.
<https://doi.org/10.1177/08944393155974>
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Peter, F. (2017). Political legitimacy. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*.
<https://plato.stanford.edu/archives/sum2017/entries/legitimacy/>
- Rahimi, M. A., & Eulenberg, J. B. (1974). A computing environment for the blind. In *Proceedings of the National Computer Conference and Exposition* (pp. 121–124). Association for Computing Machinery. <https://doi.org/10.1145/1500175.1500197>
- Rekabsaz, N., Kopeinik, S., & Schedl, M. (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 (pp. 306–316). Association for Computing Machinery.
<https://doi.org/10.1145/3404835.3462949>
- Rey, B., & Kannan, A. (2010). Conversion rate based bid adjustment for sponsored search. In *Proceedings of the 19th International Conference on World Wide Web* (pp. 1173–1174). Association for Computing Machinery. <https://doi.org/10.1145/1772690.1772860>
- Richardson, R., Schultz, J. M., & Crawford, K. (2019). Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *New York University Law Review*, 94(1), 15–55. <https://nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice/> (Archived at <https://web.archive.org/web/20260416223004/https://nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice/>)
- Rieger, A., Draws, T., Theune, M., & Tintarev, N. (2024). Nudges to mitigate confirmation bias during web search on debated topics: Support vs. manipulation. *ACM Transactions on the Web*, 18(2), 1–27. <https://doi.org/10.1145/3635034>
- Rolf, E., Worledge, T., Recht, B., & Jordan, M. I. (2021). Representation matters: Assessing the importance of subgroup allocations in training data. *arXiv preprint arXiv:2103.03399*. <https://doi.org/10.48550/arXiv.2103.03399>
- Roy, D., Mitra, M., & Ganguly, D. (2018). To clean or not to clean: Document preprocessing and reproducibility. *Journal of Data and Information Quality*, 10(4), Article 18, 1–25. <https://doi.org/10.1145/3242180>
- Sakai, T., Kim, J. Y., & Kang, I. (2023). A versatile framework for evaluating ranked lists in terms of group fairness and relevance. *ACM Transactions on Information Systems*, 42(1), 1–36. <https://doi.org/10.1145/3589763>
- Sambasivan, N., Arnesen, E., Hutchinson, B., & Prabhakaran, V. (2020). Non-portability of algorithmic fairness in India. *arXiv preprint arXiv:2012.03659*. <https://doi.org/10.48550/arXiv.2012.03659>
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68). Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287598>
- Schmidt, C., Gorman, T. J., Bayor, A. A., & Gary, M. S. (2010). Impact of low-cost, on-demand information access in a remote Ghanaian village. In *Proceedings of the 4th ACM/IEEE International Conference on Information and Communication Technologies and Development* (pp. 1–10). Association for Computing Machinery. <https://dl.acm.org/doi/10.1145/2369220.2369261>
- Shah, C., & Bender, E. M. (2024). Envisioning information access systems: What makes for good tools and a healthy Web? *ACM Transactions on the Web*, 18(3), 1–24. <https://doi.org/10.1145/3649468>
- Sonboli, N., Burke, R., Ekstrand, M., & Mehrotra, R. (2022). The multisided complexity of fairness in recommender systems. *AI Magazine*, 43(2), 164–176. <https://doi.org/10.1002/aaai.12054>
- Sonnenberg, C. (2020). E-government and social media: The impact on accessibility. *Journal of Disability Policy Studies*, 31(3), 181–191. <https://doi.org/10.1177/1044207320906521>

- Spinelli, L., & Crovella, M. (2020). How YouTube leads privacy-seeking users away from reliable information. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization* (pp. 244–251). Association for Computing Machinery. <https://doi.org/10.1145/3386392.3399566>
- Sun, Z., Di, L., Heo, G., Zhang, C., Fang, H., Yue, P., Wang, X., & Lin, L. (2017). GeoFairy: Towards a one-stop and location-based service for geospatial information retrieval. *Computers, Environment and Urban Systems*, 62, 156–167. <https://doi.org/10.1016/j.compenvurbsys.2016.11.007>
- Sun, W., Nasraoui, O., & Shafto, P. (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* (pp. 108–116). SCITEPRESS. <https://doi.org/10.5220/0006938301100118>
- Tufekci, Z. (2015). Algorithmic harms beyond Facebook and Google: Emergent challenges of computational agency. *Journal on Telecommunications & High Technology Law*, 13(2), 203–216. <https://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf> (Archived at <https://web.archive.org/web/20260421173754/https://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf>)
- Udoh, E. S., Yuan, X., & Rorissa, A. (2022). Rethinking algorithmic fairness in the context of information access. *Proceedings of the Association for Information Science and Technology*, 59(1), 815–817. <https://doi.org/10.1002/pra2.736>
- Udoh, E., Yuan, X., & Rorissa, A. (2024). A framework for defining algorithmic fairness in the context of information access. *Proceedings of the Association for Information Science and Technology*, 61(1), 667–672. <https://doi.org/10.1002/pra2.1077>
- Venkatasubramanian, S., Scheidegger, C., Friedler, S., & Clusett, A. (2021). Fairness in networks: Social capital, information access, and interventions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (pp. 4078–4079). Association for Computing Machinery. <https://doi.org/10.1145/3447548.3470821>
- Wachter, S., & Mittelstadt, B. (2019). A right to reasonable inferences: Re-thinking data protection law in the age of big data and AI. *Columbia Business Law Review*, 2019(2), 494–620. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248829 (Archived at https://web.archive.org/web/20260414121218/https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248829)
- Webb, H., Koene, A., Patel, M., & Vallejos, E. P. (2018). Multi-stakeholder dialogue for policy recommendations on algorithmic fairness. In *Proceedings of the 9th International Conference on Social Media and Society* (pp. 395–399). Association for Computing Machinery. <https://doi.org/10.1145/3217804.3217952>
- White, R. W., & Hassan, A. (2014). Content bias in online health search. *ACM Transactions on the Web*, 8(4), 1–33. <https://doi.org/10.1145/2663355>
- Wu, H., Mitra, B., Ma, C., Diaz, F., & Liu, X. (2022). Joint multisided exposure fairness for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 703–714). Association for Computing Machinery. <https://doi.org/10.1145/3477495.3532007>
- Yang, Y. T., & Chen, B. (2015). Web accessibility for older adults: A comparative analysis of disability laws. *The Gerontologist*, 55(5), 854–865. <https://doi.org/10.1093/geront/gnv057>

- Yue, Y., Patel, R., & Roehrig, H. (2010). Beyond position bias: Examining result attractiveness as a source of presentation bias in clickthrough data. In *Proceedings of the 19th International Conference on the World Wide Web* (pp. 1011–1018). Association for Computing Machinery. <https://doi.org/10.1145/1772690.1772793>
- Zambonelli, F., Salim, F., Loke, S. W., De Meuter, W., & Kanhere, S. (2018). Algorithmic governance in smart cities: The conundrum and the potential of pervasive computing solutions. *IEEE Technology and Society Magazine*, 37(2), 80–87. <https://doi.org/10.1109/MTS.2018.2826080>
- Zarsky, T. (2016). The trouble with algorithmic decisions: An analytic road map to examine efficiency and fairness in automated and opaque decision making. *Science, Technology, & Human Values*, 41(1), 118–132. <https://www.jstor.org/stable/43671285>

Copyright

Authors contributing to *Information Research* agree to publish their articles under a [Creative Commons CC BY-NC 4.0 license](https://creativecommons.org/licenses/by-nc/4.0/), which gives third parties the right to copy and redistribute the material in any medium or format. It also gives third parties the right to remix, transform and build upon the material for any purpose, except commercial, on the condition that clear acknowledgment is given to the author(s) of the work, that a link to the license is provided and that it is made clear if changes have been made to the work. This must be done in a reasonable manner, and must not imply that the licensor endorses the use of the work by third parties. The author(s) retain copyright to the work. You can also read more at: <https://publicera.kb.se/ir/openaccess>