

# Search and Society: Reimagining Information Access for Radical Futures

**Bhaskar Mitra**

*Microsoft Research*

*Montréal, Canada*

BMITRA@MICROSOFT.COM

**Editor:** Monica Lestari Paramita

## Abstract

Information retrieval (IR) research must understand and contend with the social implications of the technology it produces. Instead of adopting a reactionary strategy of trying to mitigate potential social harms from emerging technologies, the community should aim to proactively set the research agenda for the kinds of systems we *should* build inspired by diverse explicitly stated sociotechnical imaginaries. The sociotechnical imaginaries that underpin the design and development of information access technologies needs to be explicitly articulated, and we need to develop theories of change in context of these diverse perspectives. Our guiding future imaginaries must be informed by other academic fields, such as human-computer interaction, information sciences, media studies, design, science and technology studies, social sciences, humanities, democratic theory, and critical theory, as well as legal and policy experts, civil rights and social justice activists, and artists, among others. In this perspective paper, we motivate why the community must consider this radical shift in how we do research and what we work on, and sketch a path forward towards this transformation.

**Keywords:** IR and society, Sociotechnical imaginaries, Theory of change, Technology and power

## 1 Introduction

Information retrieval (IR) research must understand and contend with the social implications of the technology it produces. Nearly half a century ago, Belkin and Robertson (1976) concluded that IR research should acknowledge its responsibility to society and “must become both theoretically self-conscious and self-consciously based upon a social ideology”. This perspective has gained traction in the IR community in recent years. Researchers in attendance at the third Strategic Workshop in Information Retrieval in Lorne (SWIRL) (Culpepper et al., 2018) identified fairness, accountability, confidentiality, and transparency in IR (“FACT IR”) as socially consequential and strategically important research directions for the field. The following year, the FACTS-IR workshop (Olteanu et al., 2019) added “safety” as a fifth pillar. Subsequently, a large body of recent IR literature has grappled with questions of fairness, transparency, and explainability in the context of information access.

However, it is our perspective that this growing focus on fairness and ethics in IR—despite having played a critical role in bringing much-needed attention to the societal implications of IR systems and advancing the conversation about the IR community’s responsibility to broader society—operates within a severely constrained frame that leaves the many

underlying values, politics, and socioeconomic incentives that guide IR research largely unchallenged. For example, faced with the applications of generative artificial intelligence (AI) for information access the IR community has focused on concerns of fair ranking and representation and limiting model “hallucinations”<sup>1</sup> for good reasons, but have largely ignored other significant consequences of these technologies on society, such as for the information ecosystem and how these systems concentrate power and control (Mitra et al., 2024). In machine learning (ML), there has been similar recent perspectives (Blodgett et al., 2020; Miceli et al., 2022) that, for example, calls for shifting the lens from fairness and bias to the power differentials that exists between those who build technology, those who use it, and those who are subject to it. Others in the ML community have brought attention to the questions of how these technologies shift power (Kalluri et al., 2020) and simultaneously constrain ethics interventions in practice (Widder et al., 2023) and shape our collective sociotechnical futures. Even by accepting the frame that we should develop fairness and transparency mechanisms for certain systems, we may inadvertently ignore the alternative perspective that some of these technologies should be dismantled, not made fairer, nor more transparent (Barocas et al., 2020; Wilkinson et al., 2023; Merchant, 2023). Consequently, at the recent fourth SWIRL workshop<sup>2</sup>, researchers in attendance called for expanding the “FACTS-IR” framing to center IR research on societal, democratic, and emancipatory values. Similar sentiments, including re-centering IR on societal needs and informing IR research with democratic theories, were also discussed at the first Search Futures workshop (Azzopardi et al., 2024b).

**What do we propose?** In this paper, we argue that IR research, instead of adopting a reactionary strategy of trying to mitigate potential social harms from emerging technologies by developing new fairness and transparency interventions, should aim to proactively set the research agenda for the kinds of systems we *should* build inspired by diverse explicitly stated sociotechnical imaginaries. Towards that goal, IR research needs to explicitly articulate the sociotechnical imaginaries (Jasanoff and Kim, 2009, 2015) that underpin the design and development of information access technologies, and develop theories of change (Weiss, 1995; Brest, 2010; Taplin and Clark, 2012; Wikipedia contributors, 2013). Jasanoff and Kim (2015) define *sociotechnical imaginaries* as “collectively held, institutionally stabilized, and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and supportive of, advances in science and technology”. These shared visions do not only imagine but also co-produce our futures through development and government of digital technologies (Mager and Katzenbach, 2021). Diverse imaginaries promoted by different corporations, professional communities, political organizations, and social movements can coexist “in tension or in a productive dialectical relationship” (Jasanoff and Kim, 2015).

In technological research and development, these diverse, and often diverging, perspectives and visions that guide the community frequently remain implicit and unstated (Wilkinson et al., 2023) despite the significant influence they exert on what the community focuses on and produces. Because of the consequential role that access to information plays in political participation by citizens in democratic societies and social transformation (Higgins and

---

1. We acknowledge that the term “hallucination” anthropomorphizes AI models and its usage should be discouraged. However, given the popular usage of that term in the IR community, we make an exception here for clarity.

2. <https://sites.google.com/view/swirl2025>

Gregory, 2013; Polizzi, 2020; Goldstein, 2020; Correia, 2002; González, 2021), and as a social determinant of health (Moretti et al., 2012) and economic progress (Yu, 2002; Mutula, 2008), it is even more important to critically reflect on the values and motivations that guide the design and deployment of popular information access systems. *What sociotechnical futures do IR researchers and system designers envision and how do they influence the design of current and future IR systems? Whose sociotechnical imaginaries are granted normative status and what myriad of radically alternative futures are we overlooking?* For example, what are the implications of the reliance of popular search and social media platforms on advertising as the primary source of revenue generation (Ang, 2022) and how does Big Tech’s (Oremus, 2017) increasing dominance over academic research (Whittaker, 2021) influences and/or homogenizes the kinds of IR systems we build? *What is our role, as IR researchers, to safeguard communities from falling victim to crisis of imaginations (Haiven, 2014) and how do we become more open to welcoming influences from radically new sociotechnical imaginaries?* For example, what would IR systems look like if designed for futures informed by feminist, queer, decolonial, anti-racist, anti-casteist, anti-ableist, and abolitionist thoughts, and if the focus of IR research was not to prop up colonial cisheteropatriarchal capitalist structures but to dismantle them? We believe that explicating, critiquing, and consciously choosing the values and sociotechnical imaginaries that shape IR research is critical to realizing positive social outcomes through IR research. As Benjamin (2024) argues, exercising our imagination is “an invitation to rid our mental and social structures from the tyranny of dominant imaginaries”, or as Le Guin put it:

“The exercise of imagination is dangerous to those who profit from the way things are because it has the power to show that the way things are is not permanent, not universal, not necessary.”

– Ursula K. Le Guin

*The Wave in the Mind: Talks and Essays on the Writer, the Reader, and the Imagination* (Le Guin, 2004)

For IR to concretely support diverse sociotechnical imaginaries the research community also needs to develop their own theories of change. Theory of change (Weiss, 1995; Brest, 2010; Taplin and Clark, 2012; Wikipedia contributors, 2013) can be defined as a participatory process whereby groups and stakeholders articulate their long-term goals and identify necessary preconditions in a planning process. Consequently, in IR adjacent fields—*e.g.*, in the Fairness, Accountability, and Transparency (FAcCT)<sup>3</sup> community—there has been similar recent calls to reimagine our sociotechnical futures (da Hora et al., 2024) and develop theories of change (Wilkinson et al., 2023) to make explicit the visions for desired futures of responsible computing and the strategic pathways that lead to those desired futures. Here, we argue that IR research similarly needs to both explicitly articulate and support diverse imagined futures and develop corresponding theories of change for how new information access technologies can take us towards these desired worlds. Theories of change in this context may benefit IR research by encouraging community members to make their goals

---

3. <https://facctconference.org>

and assumptions explicit, making it feasible to test stated theories, encouraging the community to work towards building consensus, and even aid in developing potential means of evaluation of desired progress (Weiss, 1995).

Maximizing social good and minimizing harm in this context should not just be the concerns of the few in our community working on fairness, transparency, ethics, and related areas, but the domain of all IR research that should be guided by theories of change towards these envisioned futures. In this context, we largely agree with the perspective of Belkin and Robertson (1976) but diverge on a critical point which is that we believe IR research should not seek any singular notions of “universal” social ideology but explicitly adopt pluralistic humanistic and emancipatory values and make space for diverse visions and perspectives.

Consequently, the task for IR researchers here is not to put themselves in positions to pick the guiding social ideologies nor push technosolutionism to address today’s social problems. Instead, our guiding future imaginaries must be co-developed with scholars from diverse fields such as human-computer interaction (HCI), information sciences, media studies, design, science and technology studies (STS), social sciences, and political sciences, as well as legal and policy experts, civil rights and social justice activists, and artist, to name a few. Not all sociotechnical imaginaries are equal in this respect, and what futures guide our research must be informed by the values and ethics of our community, which should be constantly discussed, debated, and challenged by the community as part of our sociotechnical research and be open to external critique. To summarize, as a research community we should invest our energies and resources to: (i) Nurture digital spaces where radical visions and projects for human emancipation, social progress, and equity and justice can take shape, (ii) encourage experimentation within our research community with new approaches to information access informed by new sociotechnical imaginaries cross-pollinating through interdisciplinary scholarship, and (iii) ensure that the tools and artefacts we produce as a community do not uphold systems of oppression nor contribute towards other systemic social harms.

So far, we have argued that IR research should explicate the sociotechnical futures we want to realize and develop theories of change towards these desired futures. In Section 2 we present a brief overview of existing literature on fairness and ethics in IR, and share our critical perspectives on it to motivate our work. Then, in Section 3 we draw from relevant movements in IR-adjacent fields, specifically those with explicit values and prefigurative politics; we share some ideas on how the community can get started on this journey; and who should be doing this work. In Section 4 we motivate why now is an appropriate time for the community to consider this shift. We conclude in Section 5 with some final remarks on potential pitfalls and desired outcomes. Our goal with this paper is to raise sociopolitical consciousness in the IR research community so that we *all* see our research embedded in projects of future world making and recognize our collective responsibility to affect social good; and to dismantle the artificial separation between the work on fairness and ethics in IR and the rest of IR research.

## 2 Background

IR systems act as intermediaries between information seekers and information artefacts. These artefacts may represent: (i) economic and other opportunities for consumers, (ii) mon-

etization opportunities for content creators and publishers, (iii) specific sociopolitical frames and ideologies, and (iv) lenses to view individuals and groups that are subjects of representation by the content. These systems infer the information need from highly incomplete expressions of interests (*e.g.*, short keyword-based search queries for web search) or implicit signals (*e.g.*, history of previously accessed artefacts in case of recommender systems), and make subjective estimates of an artefact’s relevance to the information need. Consequently, these systems are not neutral tools for lookup (Noble, 2018) and the choices these systems make exert systemic influence over what information is exposed and consumed at scale. These systems bear a responsibility to society to not only mitigate potential harms, like allocative and representational harms (Crawford, 2017), but also to maximize social good.

Representational harms may happen due to reinforcement of negative stereotypes (*e.g.*, by disproportionately suggesting arrest record searches in ads corresponding to searches for black-identifying first names (Sweeney, 2013) or suggesting racist stereotypes in query autocompletion (Noble, 2018)), by pandering to the white male gaze (*e.g.*, by sexualizing women of color in search results (Noble, 2018; Urman and Makhortykh, 2024)), and through erasure (*e.g.*, underrepresenting women and other historically marginalized peoples in image search for occupational roles (Kay et al., 2015)). Allocative harms may manifest from disparate exposure in search and recommendation results (Singh and Joachims, 2018)—*e.g.*, when women are recommended lower-paying jobs in ads (Datta et al., 2014) or by influencing traffic to websites that depend on ad-monetization. Beyond direct representational and allocative harms, these systems also hold tremendous power to shape political discourse and culture (Grimmelmann, 2008; Gillespie, 2019; Hallinan and Striplas, 2016).

In light of these, there has been multiple calls (Culpepper et al., 2018; Olteanu et al., 2019) in the IR community to study and address these potential harms. Bernard and Balog (2023) report a significant rise in publications in this area after 2016, with fairness and transparency receiving the most attention. The increasing focus on these sociotechnical aspects of information access has been at least partly in response to recent advances in foundation models (Bommasani et al., 2021) and their implications for the future of information access (Shah and Bender, 2022; Mitra et al., 2024).

Fairness in ranking has garnered so much interests that there are numerous recent surveys (Ekstrand et al., 2021; Zehlike et al., 2021, 2022a,b; Pitoura et al., 2022; Dinnissen and Bauer, 2022; Aalam et al., 2022; Patro et al., 2022; Wang et al., 2023; Li et al., 2023; Deldjoo et al., 2023) and tutorials (Ekstrand et al., 2019; Gao and Shah, 2020; Li et al., 2021; Fang et al., 2022; Bigdeli et al., 2022) summarizing this emerging body of work, as well as shared tasks (Biega et al., 2020, 2021; Ekstrand et al., 2023). The fairness questions have typically been framed around disparate quality-of-service—*e.g.*, (Mehrotra et al., 2017, 2018; Neophytou et al., 2022; Wu et al., 2024)—and disparate exposure—*e.g.*, (Biega et al., 2018; Singh and Joachims, 2018, 2019; Diaz et al., 2020; Zehlike and Castillo, 2020; Patro et al., 2020; Wu et al., 2022b). Several recent works (Smith and Beattie, 2022; Raj and Ekstrand, 2020, 2022; Boratto et al., 2022, 2023) have also systemically compared various fairness metrics proposed in the literature.

Beyond fairness, there has been renewed interests in questions of transparency and explainability (Zhang et al., 2020; Anand et al., 2022), addressing misinformation (Zhou and Zafarani, 2018; Kumar and Shah, 2018; Sharma et al., 2019; Collins et al., 2021; Zhou and Zafarani, 2020; Saracco et al., 2021; Guo et al., 2022), and broader ethical concerns in

IR (Schedl et al., 2022). Transparency in IR covers a broad range of scenarios and concerns. Examples include transparency about how the system behaves (Singh and Anand, 2019; Verma and Ganguly, 2019; Singh and Anand, 2018; Zhuang et al., 2020) and how data subjects are represented in search results (Biega et al., 2017; Li et al., 2022). Transparency needs may specifically arise in the context of how information and knowledge access systems modulate *what* and *who* get exposure and influence how we see ourselves and others (Cortiñas-Lorenzo et al., 2024). Different notions of transparency may be relevant here, including but not limited to: System transparency (*i.e.*, how does the system work?), procedural transparency (*i.e.*, in what social norms and processes is the system use embedded?), or transparency of outcomes (*i.e.*, what is the impact of the system’s use on individuals and society?). New transparency needs (Liao and Vaughan, 2023) may also arise in the context of emerging technologies, such as large language models (LLMs) (OpenAI, 2023; Thoppilan et al., 2022; Chowdhery et al., 2022; Touvron et al., 2023).

## 2.1 Critical perspectives on fairness and ethics research in IR

It is important that we critically assess how the body of fairness and ethics research in IR translates to real world impact. But on this, there is little in the published literature to go on. It is plausible, for example, that the research on ranking fairness has been operationalized in, or at least has influenced the designs of, popular IR systems in recent years, but institutions who build these systems have rarely publicly disclosed any information in that regard, may be to obfuscate details of system design from bad actors or for competitive reasons. Alternatively, it is also plausible that in fact many of these approaches have *not* been operationalized, or only been operationalized in very constrained settings, because industry adoption is lagging behind research or that existing fairness research is built on abstractions and assumptions that are incompatible with real world deployment.

Studies of search logs (Jiang et al., 2013; Chuklin et al., 2015) have historically served an important role in IR research for understanding user and system behaviors and how the two interact. Similarly, online experimentation (Kohavi et al., 2007, 2009, 2020) has been key to validating the outputs of IR research in the real world. In contrast, to the best of our knowledge, there are very few fairness studies—*e.g.*, (Mehrotra et al., 2017, 2018; Raj et al., 2023)—in IR that make use of search logs, while it may be argued that the user interaction data in these logs are exactly where we *should* be looking to identify which social harms are common in practice and understand how exactly they manifest. Similarly, there is an urgent need for validating proposed fairness interventions from the literature through online experimentation, involving real users and real information needs, to ensure that fairness research is grounded in actual needs of the society and does not amount to just academic intellectual exercises. While the lack of log-based studies and online experimentation in fairness research is likely due to the lack of access to commercially-deployed systems and corresponding log data for academic research, we must critically enquire why this manifests so much more severely in fairness research compared to other areas of IR, such as in ranking. To do so, we must expand the frame beyond questions of algorithmic fairness, and examine the very sociopolitical context in which this research is being conducted, the economic incentives and risks that shape it, and the power differentials between institutions and individuals that determine what research is allowed and who is allowed to do it (Whittaker,

2021; Widder et al., 2023). This lack of access to data and systems not only makes it difficult to reproduce, validate, and challenge the claims in existing fairness studies,<sup>4</sup> but also limits what fairness questions the community is allowed to investigate.

We should also critically reflect on what questions should or should not be framed as fairness problems, and the societal consequences of doing so. For example, one of the motivating scenarios described by Morik et al. (2020) is exposure fairness for search results across different ends of the political spectrum. A similar question has also been studied by Kulshrestha et al. (2017). Casting this as a fairness issue, however, has several problematic implications and consequences. Firstly, this assumes an overly-simplistic frame in which complicated intersecting political ideologies are mapped to a linear spectrum (*e.g.*, left *vs.* right and Democrats *vs.* Republicans) and holds the two ends static rather than a continuously shifting window of acceptable discourse (Lehman, 2014; Giridharadas, 2019a). Furthermore, it also assumes that two ends of a political discourse have equal merit and deserve equal exposure, which amounts to *algorithmic bothsism*. Finally, in arguing that builders of IR systems should shift exposure out of fairness concerns, it inadvertently normalizes the idea that it is acceptable for institutions and individuals who own these systems to exert enormous influence over public discourse of social and moral import. Instead, it is our perspective that IR needs a fundamentally different and cross-disciplinary approach to these questions, one that centers on engaging and co-producing with other academic sub-fields, such as STS (Hackett et al., 2008) and critical theory of technology (Feenberg et al., 1991; Feenberg, 2002).

We must also assess the validity of constructs that we employ in fairness research. For example, Jacobs and Wallach (2021) point out that race and gender, that are often the focus of group fairness research, are contested constructs, and indeed so is the construct of fairness itself. While several papers on fairness have employed gender as a group variable, Pinney et al. (2023) caution us that much more care should be taken in this practice, for example, to ensure that we respect everyone’s right to self-identify their gender and recognize the fairness concerns of non-binary peoples. Patro et al. (2022) encourage us to move beyond fairness definitions that are grounded in discrete moments and to consider the long-term impact of fairness interventions. Fairness research itself may contribute towards certain negative externalities in the long-term, such as encouraging more pervasive collection of protected demographic attributes and further intensification of data surveillance (Zuboff, 2023) of already marginalized groups in a misguided attempt to bridge the data gaps that may be responsible for the system’s disparate quality of service across groups.

Similarly, an important question that all transparency research must contend with is: *transparency towards what end?* Some works (Polley et al., 2021; Schmitt et al., 2022) motivate transparency as a means to increase user trust in the system. However, we should be critical of whether that trust is warranted, or whether transparency mechanisms could in

---

4. For example, some studies on log-based fairness audits depend on the availability of user provided demographic attributes. In our experience, such attributes are often available only for a small subset of users—such as for signed-in users—who typically are more loyal users of the product and on average much better satisfied with system performance than the general population. Audits based on this sub-population may significantly under-report bias and unfairness issues faced by users. Better understanding of such practical challenges can motivate the community to work on specific research questions, such as fairness considerations under distributional shifts and with noisy (Ghazimatin et al., 2022; Mehrotra and Vishnoi, 2022) or missing demographic attributes (Lazovich et al., 2022; Do and Usunier, 2022).

fact draw users into a false sense of safety and distract them from noticing how the system surveils them and subtly manipulates their behavior (Ravenscraft, 2020; Morrison, 2021). Indeed, Hollanek (2023) argue that “only the sort of transparency that arises from critique—a method of theoretical examination that, by revealing pre-existing power structures, aims to challenge them—can help us produce technological systems that are less deceptive and more just”. It is that kind of critical reflection that we believe should inform transparency research in IR and lead the community towards explicit goals to challenge power (*e.g.*, the power that system owners hold over users) and safeguard user agency.

On the same lines, misinformation research should be motivated by sociotechnical visions for the future of democratic societies, public health, and knowledge production. When we broaden that frame, it becomes apparent that the community must not only focus on automated fact checking, an important research problem, but also understand the social, political, and economic conditions under which misinformation and disinformation is produced and disseminated. The focus of misinformation research then should include identifying, understanding, and addressing the structural mechanisms of misinformation—*e.g.*, data voids (Golebiewski and Boyd, 2019)—as well as ground itself in the articulation of IR’s role in online knowledge production, public health education, and information literacy.

Our perspective here on ongoing research on fairness, transparency, and ethics in IR should **not** be misconstrued as an argument for doing less of this kind of work. Instead, we believe that the community should be explicit and more ambitious about the changes it wants to affect in broader society and conduct research with a clear mapping between the research goals and the desired social impact. Just as Johnson (2014) challenges the notion that open data directly leads to information justice, we want the IR community to be cautious in their assumption that working on narrowly defined questions of fairness, transparency, and ethics necessarily contributes towards practical social good. To be effective in that endeavour, we believe that we should be explicit in articulating our collective visions for our sociotechnical futures, the changes we want to affect in society, and how we envision our research can bring about those changes.

### 3 Towards sociotechnical change

Sociotechnical imaginaries are not born in vacuum. They are moulded and shaped by our values and our politics. Deliberation over what futures we want to bring into being *is* essentially political, and challenges us to critically reflect on our community’s shared, and yet pluralistic, political values. While the call for explicit political reflection in IR may come as a surprise to some, we need to recognize that our research and the artefacts we produce do not exist outside of the current sociopolitical order (Friedman and Kahn Jr, 2007; Flanagan et al., 2008; Miller, 2021) but as essential cogs in the system, and the absence of political reflection does not imply an absence of politics in our work, but rather translates to implicit complicity in propping up the status quo and neoliberalism (Dourish, 2010; Feltwell et al., 2018; Keyes et al., 2019). Instead, we can learn from how some of our neighboring fields, *e.g.*, HCI and AI, have engaged with these questions, and reflect on how politics shapes and intersects with our own research agendas.



### 3.1 Prefigurative politics in other IR-adjacent fields

There are several strands of research in IR-adjacent fields that explicate prefigurative politics (Asad, 2019) and ground research in humanistic (Bardzell and Bardzell, 2016a, 2015; Werthner et al., 2024), anti-oppressive and emancipatory (Smyth and Dimond, 2014; Bardzell and Bardzell, 2016a; Kane et al., 2021; Monroe-White, 2021; Saxena et al., 2023), feminist (Wajcman, 2004, 2010; Bardzell, 2010; Bardzell and Bardzell, 2011; Bardzell et al., 2011; Bardzell and Bardzell, 2016b; Bardzell, 2018; D’ignazio and Klein, 2020), queer (Light, 2011; Klippahhn-Karge et al., 2024; Guyan, 2022), postcolonial and decolonial (Irani et al., 2010; Philip et al., 2012; Dourish and Mainwaring, 2012; Sun, 2013; Ali, 2014, 2016; Akama et al., 2016; Irani and Silberman, 2016; Adams, 2021; Mohamed et al., 2020), anti-racist (Abebe et al., 2022), anti-casteist (Kalyanakrishnan et al., 2018; Sambasivan et al., 2021; Vaghela et al., 2022a,b; Shubham, 2022; Kanjilal, 2023), anti-ableist (Williams et al., 2021; Sum et al., 2024), anti-fascist (McQuillan, 2022), abolitionist (Benjamin, 2019; Barabas, 2020; Earl, 2021; Jones and Melo, 2021; Williams and Haring, 2023), post-capitalistic (Feltwell et al., 2018; Browne and Green, 2022), and anarchist (Keyes et al., 2019; Linehan and Kirman, 2014; Asad et al., 2017) epistemologies. Reviewing this full body of literature is out-of-scope of this work but we briefly present a sample to draw from and motivate new IR research agendas for sociotechnical change.

Bardzell and Bardzell (2016a) define humanistic HCI as “any HCI research or practice that deploys humanistic epistemologies (e.g., theories and conceptual systems) and methodologies (e.g., critical analysis of designs, processes, and implementations; historical genealogies; conceptual analysis; emancipatory criticism) in service of HCI processes, theories, methods, agenda-setting, and practices”, and include emancipatory HCI as an aspiration of humanistic HCI. Kane et al. (2021) propose to incorporate emancipatory pedagogy (Freire, 2020) that does “not advocate the oppressed simply rise and overthrow their oppressors. Instead, [...] the oppressors and oppressed create new educational processes that would allow them to work together to create a new type of society that was emancipatory for all”.

In STS, there is a body of work (Longino, 1987; Wajcman, 1991; Hubbard, 2001; Turkle, 2004; Herring et al., 2006; von Hellens et al., 2007; Haraway, 2013; Michelfelder et al., 2017) on gendered inequities caused by technology, and how technology and gender relations mutually shape each other (Wajcman, 2004, 2010). In HCI, Bardzell et al. (Bardzell, 2010; Bardzell and Bardzell, 2011; Bardzell et al., 2011; Bardzell and Bardzell, 2016b; Bardzell, 2018) propose to incorporate feminist theories (Kolmar and Bartkowski, 1999; Friedan, 2010) into research and practice. Bardzell (2010) posits that feminist theories can contribute to interaction design both by critiquing and by generating new insights that inform and shape designs and design processes. We can see feminist HCI in practice in the works of Dimond et al. (Dimond, 2012; Dimond et al., 2013). In the context of speculative design (Auger, 2013), Martins (2014) emphasize the need for intersectional (Crenshaw, 1989; McCall, 2005; Crenshaw, 2013) feminist lens in critiquing and dismantling structures of oppression. Feminist theory, methods, and epistemologies have also influenced AI research (Adam, 1995, 2013; Wellner and Rothman, 2020; Browne et al., 2023; Toupin, 2024). Erscoi et al. (2023) highlight how women are erased from and by AI technologies. Leavy et al. (2021) propose ethical data curation approaches grounded in feminist principles. Using feminist epistemology, Huang et al. (2022) critique existing practices of explainable AI, and Varon and Peña

(2021) critique practices for obtaining digital consent in data extractivist practices in AI. Gender theory have also been employed in these fields in the forms of Queer HCI (Light, 2011) and Queer AI (Klipphahn-Karge et al., 2024). Both highlight *queering* (Sta, 1997; Brooks et al., 2021) as a tactic to challenge the basis on which categories are constructed.

Irani et al. (2010) define postcolonial computing as one that is “centered on the questions of power, authority, legitimacy, participation, and intelligibility in the contexts of cultural encounter, particularly in the context of contemporary globalization. [...] It asserts a series of questions and concerns inspired by the conditions of postcoloniality”. Avle et al. (2017) criticize the “colonizing impulse” to valorize “universal methods” that are supposedly appropriate across cultural and geopolitical boundaries; instead we can draw from works (Alsheikh et al., 2011; Wong, 2012; Winschiers-Theophilus and Bidwell, 2013; Shaw et al., 2014; Fox and Le Dantec, 2014; Ahmed et al., 2015; Akama et al., 2016; Sambasivan et al., 2021) that center on and incorporate indigenous and non-western values and ethics in the critique and development of technologies. Ali (2014) argue for decolonial computing over postcolonial which he criticizes as “Eurocentric critique of Eurocentrism” that “tends to privilege cultural issues over political-economic concerns” and “is noticeably silent about past injustices and does not engage with the matter of reparations”.

Kaba (2021) define abolition as “a long-term project and a practice around creating the conditions that would allow for the dismantling of prisons, policing, and surveillance and the creation of new institutions that actually work to keep us safe and are not fundamentally oppressive”. The movement challenges us to move beyond the default assumptions and world views of the carceral state and to dismantle the prison-industrial complex. Incorporating abolitionist values in computing requires us to oppose carceral technologies, surveillance technologies, and military applications (Earl, 2021).

Post-capitalist computing assumes “a socio-economic model that completely replaces capital as the primary method of organising society” (Feltwell et al., 2018). Among other subjects, research in this area contends with, “the racialized dynamics of labor competition” (Irani, 2018), dismantling Big Tech’s concentration of power (Verdegem, 2022; Srnicek, 2017), and imagining post-work futures (Browne and Green, 2022; Butler, 2018; Srnicek and Williams, 2015). Perhaps, the challenges in this area are best summed up in the words of Jameson (2003): “it is easier to imagine the end of the world than to imagine the end of capitalism”.

These different political lens lend to imagining new futures of computing but there are some themes that cut across them. Firstly, they all recognize technology and society as mutually shaping, and reject both technological determinism (Greene et al., 2019) and the frame in which technology exists in, what Pfaffenberger (1988) calls, a fetishised form (Marx, 1867) where technology is disembodied and disconnected from social relations. Secondly, they recognize that the perspectives, goals, and approaches across this spectrum while sometime distinct are also intersecting. Finally, they all call for structural changes and progress towards alternative futures for society and computing. Perhaps, these aspirations are best articulated by Keyes et al. (2019): “radically reorienting the field towards creating prefigurative counterpower—systems and spaces that exemplify the world we wish to see, as we go about building the revolution in increment”. To affect said changes we need to both recognize the politics of our work and ground it in broader context of political actions (Wickenden, 2018; Moore, 2020; Green, 2021; Widder et al., 2023; Young et al., 2021).

### 3.2 Proposals for IR

The survey of works presented in Section 3.1 hopefully provides some seeds of ideas for how IR research can be driven by radical new sociotechnical imaginaries. This is not to imply that these other IR-adjacent fields have achieved the desired success from these approaches, in fact there are some evidence (Chivukula and Gray, 2020) that point otherwise. Rather, we should recognize that how values and politics can inform computing research is still an open question, and they may apply differently to IR than these other fields. The challenge then for the community is to collectively engage and push towards sociotechnical change. In the remaining of this section, we discuss how we imagine some of these frames and values can guide us towards open challenges in information access. However, these examples should be interpreted as just that, as examples, not our recommendation for specific research questions the community should focus on. The actual research agenda should be developed through participatory processes that simultaneously focuses on both identifying technical research questions and building diverse communities with shared understanding of these challenges and shared commitments to address them.

Through the lens of feminist, queer, and anti-racist IR, we could critique existing approaches to ranking fairness, not only in terms their use of socially constructed categories, such as race and gender (Pinney et al., 2023), but question if it is the appropriate framing at all for the problems it purports to solve. For example, instead of trying to algorithmically fix under-representation of women and people of color in image search results for occupational roles, we could reclaim that digital space as a site of resistance and emancipatory pedagogy by allowing feminist, queer, and anti-racist scholars, activists, and artists to create experiences that teach the history of these movements and struggles.

In context of decolonial IR, ongoing fairness research may co-develop relevant local intervention strategies with legal scholars in recognition of significant differences in legal treatment of topics such as *affirmative action* across geographies. This shifts fairness research away from abstract universal notions of bias and fairness towards locally-significant societal impact (*i.e.*, *think local, act local*).

Anti-oppressive IR research may concern itself with questions such as: *Can we translate Freire’s (Freire, 2020) anti-oppressive pedagogy to strategies for anti-oppressive information access? Can search result pages support dialogical interactions between searchers that allows for communities of searchers to add context to the search results, as an alternative to centralized moderation?* Unlike conversation search, that is framed as interaction between the user and the system, the idea of dialogical search interfaces challenges us to build sophisticated sociotechnical solutions to support dialog between searchers in context of specific search intents in ways that leads to knowledge production and better digital literacy. Anti-oppressive and anti-capitalist perspectives may also motivate us to reimagine search and recommender systems as decentralized and federated.

IR research may also employ these lenses as instruments of critique. For example, in the enterprise context, Gausen et al. (2023) adopt decolonial and anti-capitalist lens to expose how information and knowledge access systems may commodify and appropriate knowledge from workers. We should also critically challenge the employment of what Gray and Suri (2019) calls Ghost Work in IR research both as a labor issue and through the lens of decolonization. In abolitionist IR, we must ensure that the techonologies we build cannot

be used for surveillance or any other military or carceral applications. The community may also consider more radical direct actions such as developing critical theories of information access, or collectively organizing to abolish Big Tech (Kwet, 2020).

### 3.3 Where do we start?

We are calling for not only a significant shift in what the IR community works on, but fundamentally changing the arrangements within our community that determine on an ongoing basis our research agendas. Beyond explicating our values and sociotechnical imaginaries, we need to develop frameworks that help us appropriately prioritize societal needs against the needs of the user, the publisher, and the platform owners. We also need new research that reimagines how IR can be informed by different epistemologies and political theories. Finally, we must also critically reexamine the arrangements within our community and create spaces for shared sense-making in collaboration with those outside of our field. We elaborate on these further in this section.

In both academic research and industrial deployment, IR places a strong emphasis on the needs of the user (consumer). This focus motivates various lines of research including: understanding user needs (through lab studies, log analysis, surveys, *etc.*), optimizing the search system towards those needs (*e.g.*, relevance optimization, personalization, and improving response time), and validating that proposed system changes indeed benefit the user (through online experimentations and further lab studies). Salient in industry settings are the needs of the system owners—*e.g.*, revenue, market share, and brand—that drive significant decisions for system design and deployment, but have historically been of lesser concern to academic IR research. Real-world IR system deployments also engage with content producers and publishers, *e.g.*, web publishers and the search engine optimization community for web search engines, and artists for music recommender systems); although how their needs are weighed against the needs of system owners and users may vary, *e.g.*, (Guttenberg, 2012; Plausic, 2015). IR research has considered questions of fairness between producers, but have rarely focused on the power differentials between system owners and producers, and its implications for producers.

Societal concerns in both IR research and industry settings have commonly been framed through a narrow lens of harm mitigation, such as “*how do we make the IR system more fair?*” and “*how do we reduce misinformation in search results?*”, without fundamentally challenging the frames in which these systems are designed and deployed, *e.g.*, centralized control and profit incentives (Mager, 2012; Taplin, 2017). IR systems are deeply embedded in sociopolitical and organization context. However, instead of grounding IR research in questions around its role in online and institutional knowledge production, literacy and informed citizenry, public health education, and social justice, the community has typically constrained themselves to improving measurable system attributes like relevance and efficiency.

Articulating different stakeholder concerns is a prerequisite to any conversation about reprioritizing our research agendas and recentering IR research on societal needs. In Figure 1, we propose a hierarchy of stakeholder needs that IR research should concern with. Contrary to the status quo, we believe that IR system design and research should explicitly reflect how these systems should contribute to knowledge production, public education, and social

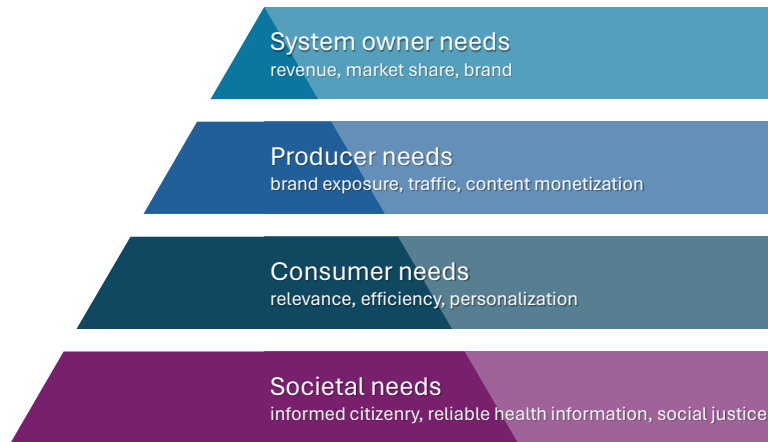


Figure 1: Hierarchy of IR stakeholder needs. More fundamental and critical needs are at the bottom of the pyramid. This figure is inspired by Maslow’s Hierarchy of Needs (Maslow, 1958) and Siksika (Blackfoot) way of life (Ravilochan, 2021).

movements, and that broader framing of societal concerns should be the most fundamental stakeholder need that should inform and shape IR research. This should then be followed by concerns of the consumer and producer needs, and lastly the needs of the system owners themselves. The needs of the consumers, producers, and system owners should not override the need of the collective society, just as the needs of the system owners should not be prioritized over the needs of the consumer and the producers. This is the shift in our research thinking, agendas, and impact that we are pushing for in this paper.

To realize structural changes in IR and meaningfully challenge dominant imaginaries we must also invest in research that specifically explores and builds on the connections between IR and different epistemologies and social and political theories. Examples of such work may consider how IR system design can support different models of democracy (Vrijenhoek et al., 2021; Helberger, 2021) and emancipatory aspirations (Mitra, 2025). We must also ensure that such research does not happen in isolation but is grounded in a collective effort to build a movement within the IR community. This requires us to acknowledge and value community building as an important part of IR research. In this context, it is important to emphasize that our call for re-centering societal needs in IR research must not be confused for a call for technosolutionism. Quite the contrary, we believe we need safe spaces where IR researchers can engage with scholars from other diverse fields, legal and policy experts, activists, and artists, in a recognition of a collective struggle to develop shared understandings of the core challenges and what IR research can offer to this process. These spaces are critical for cross-pollination of ideas and shared sense-making, which are vital for realizing structural changes. As Tsing aptly puts it:

“We are contaminated by our encounters; they change who we are as we make way for others. As contamination changes world-making projects, mutual worlds—and new directions—may emerge. Everyone carries a history of contamination; pu-

urity is not an option. [...] staying alive—for every species—requires livable collaborations. Collaboration means working across difference, which leads to contamination. Without collaborations, we all die.”

– Anna Lowenhaupt Tsing

*The mushroom at the end of the world:*

*On the possibility of life in capitalist ruins* (Tsing, 2015)

Through these exchanges we must build relations of solidarity; work together to articulate pluralistic desirable sociotechnical futures, co-develop theories of change and new research agendas to support our aspirations; vigilantly critique our own assumptions and the structures that we exist in and conduct our research in; and critically assess the impact of our work not by publications or scholarly metrics but in terms of affecting real social change. In other words, before we can transform our platforms and systems we need to transform our communities and how we conduct research. And it is vital that we approach these spaces with curiosity and humility; in recognition of our own incomplete understanding of the world; open to change and be changed by these encounters. In this context, it would do IR researchers good to keep in mind the words of activist Lilla Watson.

“If you have come here to help me you are wasting your time, but if you have come because your liberation is bound up with mine, then let us work together.”

– Lilla Watson and other members of an  
Aboriginal Rights group in Queensland<sup>5</sup>

This praxis will be unfamiliar and the learning curve frustrating; but hopefully the mutual shaping of IR technology and society will be ultimately rewarding for us all. And we must embark on this transformative journey fully acknowledging that challenging dominant imaginaries is challenging power itself. To affect counter-imaginaries we must therefore actualize counter-structures and alternative funding mechanisms that can sustain this research in the face of likely reprisals from those whose power and visions are threatened by our proposed transformations.

### 3.4 Who should do this work?

In his highly influential work “Pedagogy of the oppressed”, Freire (2020) criticizes the “banking” concept of education: “Education thus becomes an act of depositing, in which the students are the depositories and the teacher is the depositor. [...] knowledge is a gift bestowed by those who consider themselves knowledgeable upon those whom they consider to know nothing. [...] But the humanist, revolutionary educator cannot wait for this possibility to materialize. From the outset, her efforts must coincide with those of the students to engage in critical thinking and the quest for mutual humanization. His efforts must be imbued with a profound trust in people and their creative power. To achieve this, they must be partners of the students in their relations with them.” These profound words present a

---

5. While often credited for this quote, Watson explained (Reblog of report from Northland Poster) that this came out of a collective process by an Aboriginal rights group in Queensland that she was part of.

critique that we believe is also relevant to how we conduct fairness and ethics research in the IR community today.

Our default modes of doing research, like pedagogy, takes for granted the validity of *experts* and *expertise* as integral to knowledge production. While this may be effective when our research involves improving ranking or developing new evaluation measures, we posit it is the wrong approach when the goal of the research is to affect sociotechnical change. It would be narcissistic to imagine that sociotechnical research of moral import can be conducted by the few in our community and then “deposited” to the rest. Instead these concerns should be central to all IR research and we should collectively engage in dialogical praxis. Ultimately, research that attempts to affect sociotechnical change does not just transform technology, but also the researcher, and both are necessary for progress.

To challenge the homogeneity of the future imaginaries—saliently bound by colonial, cisheteropatriarchal, and capitalist ways of knowing the world—that shape our research, we need broad and diverse participation from our community. But it is also in that very context that we must critically reflect on the topic of membership in our community itself. ACM SIGIR has a commendable emphasis on the topic of Diversity, Equity, and Inclusivity (DEI) (Verberne et al., 2024; Kobayashi, 2017; Goharian and Bast, 2022; Goharian et al., 2023, 2021). For our sociotechnical imaginaries to be informed by pluralistic social, cultural, and political perspectives we not only need significantly improved representation from historically underrepresented and marginalized groups in our community, but also be inclusive of their politics and world views. Inclusion of people without inclusion of their history and struggles is simply tokenism and epistemic injustice (Fricker, 2007). That is why we believe that we, as a community, should go beyond just diversity and inclusivity (D&I), and enshrine as our goal Justice, Equity, and Diversity & Inclusivity (JEDI)—in which context D&I is both a means towards justice and equity, and also an end in itself. As Keyes et al. (2019) put it: “This must be about more than just bodies: it is not diversity if we only accept marginalised people who are stripped of the epistemic models that underpin experiences of being Other, or have the work they draw from those models held to an unequal standard of legitimacy”.

Lastly, we reiterate the important role of industry researchers in this process. They should take advantage of their proximity to large-scale systems to identify, understand, and communicate concerns of societal import and partner with academia to work on those challenges. The spaces they occupy are also sites for resistance (Wickenden, 2018; Widder et al., 2023).

## 4 Why now?

The arguments we present in this paper to reimagine our sociotechnical futures and center IR research on societal needs have always been relevant to the field. However, there is a confluence of several factors that makes this discourse particularly relevant in the present moment. Research communities constantly evolve, shaped by ideas and developments from both within the field and adjacent communities, and in response to real world events and changing societal needs. In case of IR, we believe we are seeing significant developments in both context at present: (i) A fast-changing landscape across IR and adjacent fields, such as natural language processing (NLP), HCI, ML, and AI, spurred by recent progress in LLMs

and other generative AI approaches, and (ii) an increasing recognition of the role of technology, and the communities that build it, in determining our collective futures. Consequently, there is a shared sense in the community that right now both IR technologies and IR research have been made malleable and are undergoing transformative changes under these forces of emerging new computational capabilities and evolving societal needs (Azzopardi et al., 2024a). This presents a timely opportunity for the field to consciously, collectively, and ambitiously engage in purposeful dialog about the future of the field while metaphorically the “IR(on)” is hot and before it is irrevocably shaped by unexamined imaginaries of those with power and influence over present day IR research. In doing so, we must also critically reflect on “*where do we want to go?*” (*i.e.*, our sociotechnical imaginaries), “*how do we get there?*” (*i.e.*, our theories of change), and “*who will we go there with?*” (*i.e.*, our relationships, and that of our work, with other disciplines, governments, industry, and society). These considerations should drive future IR research as a whole, and we should accept this opportunity to re-center our research agendas on societal needs while dismantling the artificial separation between the work on fairness and ethics in IR and the rest of IR research.

**LLMs are changing how we access information.** The natural language generation capabilities of LLMs are having a profound effect on how we access information and in what context. Conversational search interfaces have gone from being aspirational (Anand et al., 2021; Metzler et al., 2021) to being deployed at web-scale (*e.g.*, Bing Chat<sup>6</sup> and Google Bard<sup>7</sup>) in a span of two years. While the long-term social implications of inserting an LLM between a retrieval system and the information seeker should rightly be met with rigorous skepticism (Shah and Bender, 2022), natural language interfaces are already impacting how we interact with IR systems. The ubiquitous search box is being challenged as IR becomes more *context-driven* than *user-driven* as a consequence of LLMs increasingly embedding themselves in user’s work processes—*e.g.*, Microsoft Copilot for M365 (Mehdi, 2024; Warren, 2024)—and interacting with the IR system on the user’s behalf, under retrieval-augmentation (Lewis et al., 2020; Zamani et al., 2022).

While we should be excited with the new prospects that these emerging AI technologies unlock and recognize that they will shape how we access and interact with information in the future, we must not be duped by AI techno-determinism into believing that there is a single pre-determined path forward. Instead, we must hold pluralistic views of what IR’s future, one that is yet to be determined, looks like and how these technologies will take us there. In a study of top-cited AI papers, many of which are coauthored by researchers affiliated with industry or elite universities, Birhane et al. (2022) find that the dominant values expressed and operationalized support concentration of power. So, we must ask: *in what new ways can we imagine accessing and interacting with information, aided by LLMs, if large-scale IR systems were not just a purview of Big Tech? How would LLMs empower us to reimagine IR systems whose explicit goal is to dismantle hierarchies and redistribute power, not to centralize it? What role would AI technologies play in information access that is built explicitly to facilitate dialogical social processes for knowledge production, world building, and our collective struggles for universal emancipation?* It is critical that we have

6. <https://chat.bing.com/> (now Microsoft Copilot (Mehdi, 2023))

7. <https://bard.google.com/>



these conversations now in the face of ongoing massive technology-driven power shifts in favor of dominant established platforms that grants their visions of the future normative status and shrinks the space for any critique, resistance, or counter-imaginaries.

**LLMs are shifting priorities of IR research.** Over the last decade, deep learning technologies became the new hammer in the toolbox for IR research (Mittra and Craswell, 2018; Lin et al., 2020; Fan et al., 2022), dominating IR publications with nearly four out of five papers at the ACM SIGIR 2020 conference being related to deep learning by some estimates (Mittra, 2021). One particular focus of neural IR has been on estimating relevance of information artefacts (*e.g.*, documents) to an information intent (*e.g.*, as expressed by a search query) for ranking, a central problem in IR. Curiously, many of the key ingredients for this research, such as the Transformer architecture (Vaswani et al., 2017) and the idea of pretrained LLMs, like BERT (Devlin et al., 2019), came from fields adjacent to IR; correspondingly, shifting the focus within the field more towards adapting these models for the relevance estimation task—*e.g.*, (Nogueira and Cho, 2019)—and making them more efficient (Fröbe et al., 2024).

More recently, Thomas et al. (2024) demonstrated that LLMs, like GPT-4 (OpenAI, 2023), are able to estimate the actual searcher’s preference for documents, given their query, better than several populations of human relevance assessors. This technology has already been deployed in production at Bing.<sup>8</sup> Putting it bluntly, these LLMs may be getting close to the best we can expect with machine learned general purpose relevance estimators. If these claims stand the test of time, it may mark a watershed moment for IR research. Speculatively, we may see the IR community further shifting towards: (i) Improving efficiency of these models, (ii) focusing on more specialized information needs—*e.g.*, tip-of-the-tongue information needs (Arguello et al., 2021), and (iii) increasing investments in measurement and evaluation—*e.g.*, for emerging new IR scenarios, such as retrieval-enhanced machine learning (Zamani et al., 2022). Alternatively, we may ask: *How can the IR community meet this moment, not with apprehension nor with unchallenged exuberance for progress happening in adjacent fields, but truly grasp this opportunity to redefine what it means to work on IR research? Can we be unabashedly discontent with imagining the future of IR based wholly on what new AI progress makes plausible, and instead reimagine our field as a place where knowledge, culture, and radical aspirations meet to demand of technology to make new futures possible?* Alternatively, if we fail to articulate an aspiring vision for IR research, we risk as a field being reduced to just an application of AI.

**LLMs are raising new sociotechnical concerns.** It is well-known that language models reproduce, and even amplify, harmful stereotypes and biases of moral import (Friedman and Nissenbaum, 1996) that are present in their training data (Bolukbasi et al., 2016; Caliskan et al., 2017; Gonen and Goldberg, 2019; Blodgett et al., 2020; Bender et al., 2021; Abid et al., 2021). One particular mitigation strategy involves using ML approaches that learn from human preferences, such as reinforcement learning from human feedback (RLHF) (Christiano et al., 2017; Ziegler et al., 2019), to align with “human values” (Kasirzadeh and Gabriel, 2023; Tamkin et al., 2021). While RLHF has been quite effective in constraining LLMs from producing certain types of offensive and harmful content, we must be wary of any framing of AI ethics, such as *AI alignment* (Russell et al., 2015; Gabriel, 2020; Gabriel and Ghazavi,

8. [https://twitter.com/IR\\_oldie/status/1659413086007328768](https://twitter.com/IR_oldie/status/1659413086007328768)

2021), that presupposes the existence of *universal values* but that assumption does not hold true in reality (Prabhakaran et al., 2022; Birhane and Cummins, 2019; Jobin et al., 2019; Png, 2022; Sambasivan et al., 2021). This is particularly concerning if we look at this in the context of power asymmetries that exist between powerful private corporations, who have outsized influence over what values these models are optimized for, and those who use these models or are represented in some fashion in the model outputs. This is further compounded by the lack of appropriate mechanisms for civil society to participate in and challenge these choices. Indeed, by placing these controls in the hands of the privileged few, we risk further concentration of power. The concerns of biases in what these models produce, and even biases in what context they refuse to generate (Urman and Makhortykh, 2023), and who gets to influence those decisions have serious implications for information access and society. Finally, the development of LLMs themselves may involve potential harms to authors (Davis, 2023; Lawler, 2023; Browne, 2023; Shetler, 2024; Milmo, 2024), crowdworkers (Gray and Suri, 2017; Ekbja and Nardi, 2017; Gray and Suri, 2019; Roberts, 2019; Jones-Imhotep, 2020; Roberts, 2021; Williams et al., 2022; Perrigo, 2023; Dzieza, 2023), and even the environment (Bender et al., 2021; Patterson et al., 2021; Bommasani et al., 2021; Wu et al., 2022a; Dodge et al., 2022; Patterson et al., 2022).

In a critical perspective, Shah and Bender (2022) recommend that IR research should focus on developing appropriate guardrails in anticipation of the social implications of these emerging technologies and not be constrained by a singular LLM-powered conversational search vision for IR. With this we agree, we should be excited by the new capabilities unlocked by recent progress in LLMs but must not limit our imaginations and aspirations by only what LLMs make plausible. We must also consider the broad sociotechnical implications of deploying these emerging technologies in the context of information access (Mitra et al., 2024), and their systemic consequences and risks.

**Our relationships with adjacent communities are changing and so is how we do research.** Not long ago, many in the IR community would tout IR, and specifically web search, as a rare success story of real-world application of ML, AI, and NLP technologies. A sea change in these adjacent communities in the last decade have shifted this balance. Now, many see IR as just another NLP task, sometimes included in NLP benchmarks (*e.g.*, HELM (Liang et al., 2022)) for evaluating ML and AI models. With retrieval-augmented LLMs, IR is auditioning for a new role, as a tool for AI models, curiously inverting the relationship between these technologies, where AI was one of many in the IR toolbox.

These communities are also undergoing significant changes in research culture, often influencing each other on the way. One particular trend in NLP, and the broader ML and AI communities, that has influenced IR, is *leaderboard-driven research*. Several NLP leaderboards (Bajaj et al., 2016; Goyal et al., 2017; Joshi et al., 2017; Rajpurkar et al., 2018; Wang et al., 2018; Talmor et al., 2018; Wang et al., 2019; Kwiatkowski et al., 2019; Yang et al., 2018; Liang et al., 2020) have been instrumental in encouraging progress on specific tasks. IR has a long history of focus on shared tasks and benchmarks, notably TREC (Voorhees et al., 2005) that has been a venue for developing new tasks and benchmarks, as well as building research communities with shared interests around them. What differentiates IR benchmarking in venues such as TREC from NLP leaderboards is that the former is framed not as a *competition*, but as a *coopetition*. In a competition, the goal of the participant is to

outperform others, while in a coopetition the participants share a collective goal to develop a better understanding of both the task and the models in the spirit of scientific enquiry. In the words of IR researcher Ian Soboroff:<sup>9</sup> “The datasets were not built to be solved. They were built as tools to understand the problem and the systems we build to ‘solve’ them.”; or as Voorhees (2021) put it: “Coopetition is defined as competitors cooperating for the common good... While competition can give one a bigger piece of the pie, cooperation makes the whole pie bigger.” By emphasizing the goals of community development and understanding of the tasks and the models, these evaluation effort try to promote scientific enquiry over sportive competition. Even the MS MARCO ranking task that initially started as a competition later reframed itself as a coopetitive evaluation effort (Craswell et al., 2021; Lin et al., 2022).

As a community we should cautiously embrace insights and trends from our neighboring fields. However, we should not let IR be minimized to just an evaluation task which undermines the critical responsibility that IR researchers owe to the broader society. Similarly, while leaderboards and competitions may be effective in creating excitement and increasing participation in certain tasks, we must be mindful of the implications of potentially a large section of the IR community being driven predominantly by these practices. When the goal is to win, then scientific inquiry takes a back seat, and the ones with the most compute and data resources take the metaphorical steering wheel. It risks, what Gausen et al. (2023) call, albeit in a different context, shifting from *praxis*—*i.e.*, “reflection and action directed at the structures to be transformed” (Freire, 2020)—to *proxies*, *i.e.*, optimizing towards proxy quantitative measures of outcomes. Actions in this context may refer to any research activity, including but not limited to: formalization, design, experimentation, publishing, artefact creation, open sourcing, and community building; and examples of proxies include state-of-the-art (SOTA) performance on benchmarks and leaderboard rankings that do not translate to better scientific understanding or positive impact on people.<sup>10</sup>

Yet another relationship that we must critically examine is the one between industry and academia. Whittaker (2021) point out that the concentration of data and compute resources, two key ingredients in recent advances in AI, in the hands of few large tech corporations is giving these same institutions tremendous power to shape academic research agenda. Big tech also shapes academic research agendas in various other ways, including academic engagements and employments. In IR, the MS MARCO dataset (Bajaj et al., 2016) and leaderboard (Craswell et al., 2021; Lin et al., 2021, 2022), and the TREC Deep Learning track (Craswell et al., 2020), that has been broadly adopted for benchmarking deep ranking models were produced and is currently maintained by industry researchers. Indeed, the organizers behind these efforts themselves recognized (Craswell et al., 2021; Lin et al., 2022) the critical responsibility that comes with defining critical research tasks for the community—effectively playing “the Pied Piper guiding a significant section of the community down specific lanes of research”—and recommend all benchmark developers to engage in open and inclusive discussions with the rest of the community to critically examine

9. [https://twitter.com/ian\\_soboroff/status/1426901262369439751](https://twitter.com/ian_soboroff/status/1426901262369439751)

10. This is a case of Goodhart’s law (Chrystal and Mizen, 2001; Goodhart, 1975; Hoskin, 1996; Thomas and Uminsky, 2022) whereby improvements on benchmarks and corresponding metrics do not translate to progress on the problem the benchmark was created for, as has been argued for example by Hsia et al. (2023).

their impact. While academia-industry collaboration is critically important for the field to ground our research in real large-scale systems and see our research outputs materialize into real impact on system users, we must also resist the homogenization of our research agendas towards a singular world view put forth by Big Tech capitalism.

**The world is changing and so is our relationship to that world.** Our world at large is experiencing a confluence of many simultaneous, and mutually reinforcing, forces that are increasingly pushing us towards precarity, including but not limited to: increasing global wealth and income inequality (Chancel et al., 2022), rising global conflicts (Taylor, 2023; United Nations Meetings Coverage and Press Releases), pandemics (Scientist, 2021; Taylor, 2022; for Disease Control et al., 2022), and impending climate catastrophes (Parmesan et al., 2022; IPCC, 2013; Poynting and Rivault, 2024). At a moment when the world needs global solidarity built on trust and consensus, and informed citizenry with robust access to reliable information, online disinformation and misinformation are undermining both (Turrentine, 2022; Treen et al., 2020; Kata, 2010; Allcott and Gentzkow, 2017; Doubek, 2017; Beaumont et al., 2020; Zadrozny, 2024; Swenson and Fernando, 2023). While these complex global challenges require sophisticated and multifaceted response that spans across the political, legal, economic, and technological realms, one thing is for certain that information access research has a role to play. *So, will we answer the call?*

## 5 Conclusion

**Despite good intentions.** We must be vigilant and reflexively critique our impact, whether under the model of existing fairness and ethics research in IR or under the proposed shift. The call for pluralistic sociotechnical imaginaries must not in this context be confused with uncritical acceptance of all possible futures as equally valid or desirable. Instead, this is a call for critical examination of our community’s existing normative values and future aspirations. This work not only involves explicating our sociotechnical imaginaries but also engaging critically with the history of technology—*e.g.*, (Merchant, 2023)—and challenging harmful silicon valley ideologies that are counter to the goals of universal emancipation and justice—*e.g.*, (Gebu and Torres, 2023). Above all, we should be wary of any promises of the future that further concentrates wealth and power, or advances any notion of altruism in place of structural change (Giridharadas, 2019b).

**Desired outcomes.** Having emphasized the importance of theories of change and ensuring that our research has the desired societal impact and not merely constitute an intellectual exercise, it is only fair that we explicate our own desired outcome of this particular work. We authored this paper because we sincerely believe that information access has a critical role to play in determining our collective futures; and that real change can not be realized by fairness and ethics research happening in silos but only when combined with raising social consciousness, organizing, and movement building. We would consider it a failure if this paper is only cited in future IR papers as a passing remark on social responsibility of IR research. Instead, we hope this work sparks many passionate conversations and debates within the community, and radicalizes us to work on issues of social import in collaboration with other disciplines and civil society. But above all, we hope this paper serves as a clarion call to all IR researchers to reflect on why we do what we do. Personally, we hope that the

community continues to build technology not just because we love technology itself, but as an act of radical love for all peoples and the worlds we share. So, we conclude with one final quote for our readers.

“Another world is not only possible, she is on her way. On a quiet day, I can hear her breathing.”

– Arundhati Roy  
*War talk* (Roy, 2003)

### Positionality statement

The author of this paper works at a large technology company in the global north. However, the perspectives presented in this work is intended to challenge Big Tech and global north’s view of technology and our collective futures.

### Acknowledgments and Disclosure of Funding


The author gratefully acknowledges feedback from Asia Biega, Michael D. Ekstrand, and Ida Larsen-Ledet on the various drafts of this paper. No external funding was received in support of this work.

### References

- Syed Wajid Aalam, Abdul Basit Ahanger, Muzafar Rasool Bhat, and Assif Assad. Evaluation of fairness in recommender systems: A review. In *International Conference on Emerging Technologies in Computer Engineering*, pages 456–465. Springer, 2022.
- Veronica Abebe, Gagik Amaryan, Marina Beshai, Ilene, Ali Ekin Gurgun, Wendy Ho, Naaji R Hylton, Daniel Kim, Christy Lee, Carina Lewandowski, et al. Anti-racist HCI: notes on an emerging critical technical practice. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–12, 2022.
- Abubakar Abid, Maheen Farooqi, and James Zou. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 298–306, 2021.
- Alison Adam. A feminist critique of artificial intelligence. *European Journal of Women’s Studies*, 2(3):355–377, 1995.
- Alison Adam. Feminist AI projects and cyberfutures. In *The Gendered Cyborg*, pages 276–290. Routledge, 2013.
- Rachel Adams. Can artificial intelligence be decolonized? *Interdisciplinary Science Reviews*, 46(1-2):176–197, 2021.
- Syed Ishtiaque Ahmed, Nusrat Jahan Mim, and Steven J. Jackson. Residual mobilities: infrastructural displacement and post-colonial computing in bangladesh. In *Proceedings*

- of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pages 437–446, 2015.
- Yoko Akama, Seth Keen, and Peter West. Speculative design and heterogeneity in indigenous nation building. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, pages 895–899, 2016.
- Syed Mustafa Ali. Towards a decolonial computing. In *In Ambiguous Technologies: Philosophical Issues, Practical Solutions, Human Nature*, pages 28–35, 2014.
- Syed Mustafa Ali. A brief introduction to decolonial computing. *XRDS: Crossroads, The ACM Magazine for Students*, 22(4):16–21, 2016.
- Hunt Allcott and Matthew Gentzkow. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–236, 2017.
- Tamara Alsheikh, Jennifer A Rode, and Siân E Lindley. (Whose) value-sensitive design: a study of long-distance relationships in an Arabic cultural context. In *Proceedings of the ACM conference on Computer supported cooperative work*, pages 75–84, 2011.
- Avishek Anand, Lawrence Cavedon, Matthias Hagen, Hideo Joho, Mark Sanderson, and Benno Stein. Dagstuhl seminar 19461 on conversational search: seminar goals and working group outcomes. In *ACM SIGIR Forum*, volume 54, pages 1–11. ACM, 2021.
- Avishek Anand, Lijun Lyu, Maximilian Idahl, Yumeng Wang, Jonas Wallat, and Zijian Zhang. Explainable information retrieval: A survey. *arXiv preprint arXiv:2211.02405*, 2022.
- Carmen Ang. How do big tech giants make their billions? *Visual Capitalist*, 25 April 2022. URL <https://www.visualcapitalist.com/how-big-tech-makes-their-billions-2022/>.
- Jaime Arguello, Adam Ferguson, Emery Fine, Bhaskar Mitra, Hamed Zamani, and Fernando Diaz. Tip of the tongue known-item retrieval: A case study in movie identification. In *Proceedings of the Conference on Human Information Interaction and Retrieval*, 2021.
- Mariam Asad. Prefigurative design as a method for research justice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–18, 2019.
- Mariam Asad, Christopher A Le Dantec, Becky Nielsen, and Kate Diedrick. Creating a sociotechnical api: Designing city-scale community engagement. In *Proceedings of the CHI conference on human factors in computing systems*, pages 2295–2306, 2017.
- James Auger. Speculative design: Crafting the speculation. *Digital Creativity*, 24(1):11–35, 2013.
- Seyram Avle, Silvia Lindtner, and Kaiton Williams. How methods make designers. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 472–483, 2017.

- Leif Azzopardi, Charles LA Clarke, Paul Kantor, Bhaskar Mitra, Johanne Trippas, Zhaochun Ren, Mohammad Aliannejadi, Negar Arabzadeh, Raman Chandrasekar, Maarten de Rijke, et al. Report on the search futures workshop at ECIR 2024. In *ACM SIGIR Forum*, volume 58, pages 1–41. ACM, 2024a.
- Leif Azzopardi, Charlie Clarke, Paul Kantor, Bhaskar Mitra, Johanne Trippas, and Zhaochun Ren. The search futures workshop (ecir2024), 2024b. URL <https://searchfutures.github.io/>.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. MS MARCO: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*, 2016.
- Chelsea Barabas. Beyond bias: Re-imagining the terms of “ethical AI” in criminal law. *Georgetown Journal of Law and Modern Critical Race Perspectives*, 12:83, 2020.
- Jeffrey Bardzell and Shaowen Bardzell. What is humanistic HCI? In *Humanistic HCI*, pages 13–32. Springer, 2015.
- Jeffrey Bardzell and Shaowen Bardzell. Humanistic HCI. *Interactions*, 23(2):20–29, 2016a.
- Shaowen Bardzell. Feminist HCI: taking stock and outlining an agenda for design. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1301–1310, 2010.
- Shaowen Bardzell. Utopias of participation: Feminism, design, and the futures. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 25(1):1–24, 2018.
- Shaowen Bardzell and Jeffrey Bardzell. Towards a feminist HCI methodology: social science, feminism, and HCI. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 675–684, 2011.
- Shaowen Bardzell and Jeffrey Bardzell. Feminist design in computing. *The Wiley Blackwell Encyclopedia of Gender and Sexuality Studies*, pages 1–7, 2016b.
- Shaowen Bardzell, Elizabeth Churchill, Jeffrey Bardzell, Jodi Forlizzi, Rebecca Grinter, and Deborah Tatar. Feminism and interaction design. In *CHI’11 Extended Abstracts on Human Factors in Computing Systems*, pages 1–4. 2011.
- Solon Barocas, Asia J Biega, Benjamin Fish, Jędrzej Niklas, and Luke Stark. When not to design, build, or deploy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 695–695, 2020.
- Peter Beaumont, Julian Borger, and Daniel Boffey. Malicious forces creating “perfect storm” of coronavirus disinformation. *The Guardian*, 24, 2020.
- Nicolas Belkin and Stephen Robertson. Some ethical and political implications of theoretical research in information science. In *Proceedings of the ASIS Annual Meeting*, 1976.

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? . In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 2021.
- Ruha Benjamin. Race after technology: Abolitionist tools for the new jim code. *Social Forces*, 2019.
- Ruha Benjamin. *Imagination: A manifesto*. WW Norton & Company, 2024.
- Nolwenn Bernard and Krisztian Balog. A systematic review of fairness, accountability, transparency and ethics in information retrieval. *ACM Computing Surveys*, 2023.
- Asia J Biega, Azin Ghazimatin, Hakan Ferhatosmanoglu, Krishna P Gummadi, and Gerhard Weikum. Learning to un-rank: quantifying search exposure for users in online communities. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 267–276, 2017.
- Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR ’18, pages 405–414, New York, NY, USA, 2018. ACM.
- Asia J Biega, Fernando Diaz, Michael D Ekstrand, and Sebastian Kohlmeier. Overview of the TREC 2019 fair ranking track. *arXiv preprint arXiv:2003.11650*, 2020.
- Asia J Biega, Fernando Diaz, Michael D Ekstrand, Sergey Feldman, and Sebastian Kohlmeier. Overview of the TREC 2020 fair ranking track. *arXiv preprint arXiv:2108.05135*, 2021.
- Amin Bigdeli, Negar Arabzadeh, Shirin SeyedSalehi, Morteza Zihayat, and Ebrahim Bagheri. Gender fairness in information retrieval systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3436–3439, 2022.
- Abeba Birhane and Fred Cummins. Algorithmic injustices: Towards a relational ethics. *arXiv preprint arXiv:1912.07376*, 2019.
- Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. The values encoded in machine learning research. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, pages 173–184, 2022.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. Language (technology) is power: A critical survey of “bias” in NLP. *arXiv preprint arXiv:2005.14050*, 2020.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29, 2016.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.



- Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. Consumer fairness in recommender systems: Contextualizing definitions and mitigations. In *European Conference on Information Retrieval*, pages 552–566. Springer, 2022.
- Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. Consumer fairness benchmark in recommendation. In *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023). Pisa, Italy*, pages 60–65, 2023.
- Paul Brest. The power of theories of change, *Stanford Social Innovation Review*, 2010.
- Lonny J Avi Brooks, Jason Tester, Eli Kosminsky, and Anthony D Weeks. Queering: Liberation futures with afrofuturism. In *Routledge Handbook of Social Futures*, pages 260–274. Routledge, 2021.
- Jacob Browne and Laurel Green. The future of work is no work: A call to action for designers in the abolition of work. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–8, 2022.
- Jude Browne, Stephen Cave, Eleanor Drage, and Kerry McInerney. *Feminist AI: Critical Perspectives on Algorithms, Data, and Intelligent Machines*. Oxford University Press, 2023.
- Ryan Browne. New York Times sues Microsoft, ChatGPT maker OpenAI over copyright infringement. 2023. URL <https://www.cnbc.com/2023/12/27/new-york-times-sue-s-microsoft-chatgpt-maker-openai-over-copyright-infringement.html>.
- Lise Butler. Interview: Technology, capitalism, and the future of the left. *Renewal: A Journal of Social Democracy*, 26(1):18–31, 2018.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, 2017.
- Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman. *World inequality report 2022*. Harvard University Press, 2022.
- Shruthi Sai Chivukula and Colin M Gray. Bardzell’s “feminist HCI” legacy: Analyzing citational patterns. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–8, 2020.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- K. Alec Chrystal and Paul D. Mizen. Goodhart’s law: Its origins, meaning and implications for monetary policy. Prepared for the Festschrift in honour of Charles Goodhart, 2001.

- Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. Click models for web search. *Synthesis lectures on information concepts, retrieval, and services*, 7(3):1–115, 2015.
- Botambu Collins, Dinh Tuyen Hoang, Ngoc Thanh Nguyen, and Dosam Hwang. Trends in combating fake news on social media—a survey. *Journal of Information and Telecommunication*, 5(2):247–266, 2021.
- Ana Maria Ramalho Correia. Information literacy for an active and effective citizenship. In *White Paper prepared for UNESCO, the US National Commission on Libraries and Information Science, and the National Forum on Information Literacy, for use at the Information Literacy Meeting of Experts, Prague, The Czech Republic*, 2002.
- Karina Cortiñas-Lorenzo, Siân Lindley, Ida Larsen-Ledet, and Bhaskar Mitra. Through the looking-glass: Transparency implications and challenges in enterprise AI knowledge systems. *arXiv preprint arXiv:2401.09410*, 2024.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. Overview of the TREC 2019 deep learning track. In *Proceedings of the Text REtrieval Conference (TREC)*, 2020.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Jimmy Lin. Ms marco: Benchmarking ranking models in the large-data regime. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2021.
- Kate Crawford. The trouble with bias. In *Conference on Neural Information Processing Systems, invited speaker*, 2017.
- Kimberlé Williams Crenshaw. Mapping the margins: Intersectionality, identity politics, and violence against women of color. In *The public nature of private violence*, pages 93–118. Routledge, 2013.
- KW Crenshaw. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine (pp. 139–168). In *University of Chicago legal forum*, 1989.
- J Shane Culpepper, Fernando Diaz, and Mark D Smucker. Research frontiers in information retrieval: Report from the third strategic workshop on information retrieval in lorne (SWIRL 2018). In *ACM SIGIR Forum*, volume 52, pages 34–90. ACM New York, NY, USA, 2018.
- Nina da Hora, Joana Varon, and Annette Zimmermann. Better utopias: resisting silicon valley ideology and decolonizing our imaginaries of the future, 2024. URL <https://faccconference.org/2024/acceptedcraft>.
- Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *arXiv preprint arXiv:1408.6491*, 2014.
- Wes Davis. Sarah silverman is suing OpenAI and Meta for copyright infringement. *The Verge*, 9 July 2023. URL <https://www.theverge.com/2023/7/9/23788741/sarah-sil>

verman-openai-meta-chatgpt-llama-copyright-infringement-chatbots-artificial-intelligence-ai.

Yashar Deldjoo, Dietmar Jannach, Alejandro Bellogin, Alessandro Difonzo, and Dario Zanzonelli. Fairness in recommender systems: research landscape and future directions. *User Modeling and User-Adapted Interaction*, pages 1–50, 2023.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.

Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. Evaluating stochastic rankings with expected exposure. In *Proc. CIKM*, pages 275–284, 2020.

Catherine D’ignazio and Lauren F Klein. *Data feminism*. MIT press, 2020.

Jill P Dimond. *Feminist HCI for real: Designing technology in support of a social movement*. Georgia Institute of Technology, 2012.

Jill P Dimond, Michaelanne Dye, Daphne LaRose, and Amy S Bruckman. Hollaback! the role of storytelling online in a social movement organization. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 477–490, 2013.

Karlijn Dinnissen and Christine Bauer. Fairness in music recommender systems: A stakeholder-centered mini review. *Frontiers in big Data*, 5:913608, 2022.

Virginie Do and Nicolas Usunier. Optimizing generalized gini indices for fairness in rankings. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 737–747, 2022.

Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A Smith, Nicole DeCario, and Will Buchanan. Measuring the carbon intensity of AI in cloud instances. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, pages 1877–1894, 2022.

James Doubek. How disinformation and distortions on social media affected elections worldwide. *NPR*, 16 November 2017.

Paul Dourish. HCI and environmental sustainability: the politics of design and the design of politics. In *Proceedings of the 8th ACM conference on designing interactive systems*, pages 1–10, 2010.

Paul Dourish and Scott D Mainwaring. Ubicomp’s colonial impulse. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 133–142, 2012.

- Josh Dzieza. AI is a lot of work. *The Verge*, 20 July 2023. URL <https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots>.
- Charles C Earl. Towards an abolitionist ai: the role of historically black colleges and universities. *arXiv preprint arXiv:2101.02011*, 2021.
- Hamid R Ekbia and Bonnie A Nardi. *Heteromation, and other stories of computing and capitalism*. MIT Press, 2017.
- Michael D Ekstrand, Robin Burke, and Fernando Diaz. Fairness and discrimination in retrieval and recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1403–1404, 2019.
- Michael D Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. Fairness and discrimination in information access systems. *arXiv preprint arXiv:2105.05779*, 2021.
- Michael D Ekstrand, Graham McDonald, Amifa Raj, and Isaac Johnson. Overview of the TREC 2022 fair ranking track. *arXiv preprint arXiv:2302.05558*, 2023.
- Lelia Erscoi, Annelies V Kleinherenbrink, and Olivia Guest. Pygmalion displacement: When humanising AI dehumanises women. *SocArXiv. February*, 11, 2023.
- Yixing Fan, Xiaohui Xie, Yinqiong Cai, Jia Chen, Xinyu Ma, Xiangsheng Li, Ruqing Zhang, Jiafeng Guo, et al. Pre-training methods in information retrieval. *Foundations and Trends in Information Retrieval*, 16(3):178–317, 2022.
- Yi Fang, Hongfu Liu, Zhiqiang Tao, and Mikhail Yurochkin. Fairness of machine learning in search engines. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 5132–5135, 2022.
- Andrew Feenberg. *Transforming technology: A critical theory revisited*. Oxford University Press, 2002.
- Andrew Feenberg et al. *Critical theory of technology*, volume 5. Oxford University Press New York, 1991.
- Tom Feltwell, Shaun Lawson, Enrique Encinas, Conor Linehan, Ben Kirman, Deborah Maxwell, Tom Jenkins, and Stacey Kuznetsov. “Grand visions” for post-capitalist human-computer interaction. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–8, 2018.
- Mary Flanagan, Daniel C Howe, and Helen Nissenbaum. Embodying values in technology: Theory and practice. In *Information Technology and Moral Philosophy*, pages 322–353. Cambridge University Press, 2008.
- Centers for Disease Control, Prevention. CDC museum COVID-19 timeline. 2022. URL <https://www.cdc.gov/museum/timeline/covid19.html>

- Sarah Fox and Christopher Le Dantec. Community historians: scaffolding community engagement through culture and heritage. In *Proceedings of the 2014 conference on Designing interactive systems*, pages 785–794, 2014.
- Paulo Freire. Pedagogy of the oppressed. In *Toward a sociology of education*, pages 374–386. Routledge, 2020.
- Miranda Fricker. *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press, 2007.
- Betty Friedan. *The feminine mystique*. WW Norton & Company, 2010.
- Batya Friedman and Peter H Kahn Jr. Human values, ethics, and design. In *The human-computer interaction handbook*, pages 1267–1292. CRC press, 2007.
- Batya Friedman and Helen Nissenbaum. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)*, 14(3):330–347, 1996.
- Maik Fröbe, Joel Mackenzie, Bhaskar Mitra, Franco Maria Nardini, and Martin Potthast. ReNeuIR at SIGIR 2024: The third workshop on reaching efficiency in neural information retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3051–3054, 2024.
- Iason Gabriel. Artificial intelligence, values, and alignment. *Minds and machines*, 30(3): 411–437, 2020.
- Iason Gabriel and Vafa Ghazavi. The challenge of value alignment: From fairer algorithms to AI safety. *arXiv preprint arXiv:2101.06060*, 2021.
- Ruoyuan Gao and Chirag Shah. Counteracting bias and increasing fairness in search and recommender systems. In *Proceedings of the 14th ACM Conference on Recommender Systems*, pages 745–747, 2020.
- Anna Gausen, Bhaskar Mitra, and Siân Lindley. A framework for exploring the consequences of ai-mediated enterprise knowledge access and identifying risks to workers. *arXiv preprint arXiv:2312.10076*, 2023.
- Timnit Gebru and Émile P Torres. Eugenics and the promise of utopia through artificial general intelligence. *First Monday*, 2023.
- Azin Ghazimatin, Matthaus Kleindessner, Chris Russell, Ziawasch Abedjan, and Jacek Golebiowski. Measuring fairness of rankings under noisy sensitive information. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, pages 2263–2279, 2022.
- Tarleton Gillespie. Algorithmically recognizable: Santorum’s Google problem, and Google’s santorum problem. In *The Social Power of Algorithms*, pages 63–80. Routledge, 2019.
- A. Giridharadas. How America’s elites lost their grip. *Time Magazine*, 2019a.

- Anand Giridharadas. *Winners take all: The elite charade of changing the world*. Vintage, 2019b.
- Nazli Goharian and Hannah Bast. Report on women in ir (wir 2021) at sigir 2021. In *ACM SIGIR Forum*, volume 55, pages 1–3. ACM, 2022.
- Nazli Goharian, Xin Ma, and Suzan Verberne. Women and disparities in leadership and wages. In *ACM SIGIR Forum*, volume 54, pages 1–3. ACM New York, NY, USA, 2021.
- Nazli Goharian, Faegheh Hasibi, Maria Maistro, and Suzan Verberne. Report on the SIGIR 2022 session on Women in IR (WIR). In *ACM SIGIR Forum*, volume 56, pages 1–2. ACM New York, NY, USA, 2023.
- Stéphane Goldstein. *Informed Societies*. Facet publishing, 2020.
- Michael Golebiewski and Danah Boyd. Data voids: Where missing data can easily be exploited. *Data & Society*, 2019.
- Hila Gonen and Yoav Goldberg. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In *Proc. NAACL*, pages 609–614, 2019.
- M González. A better-informed society is a freer society. 2021. URL <https://www.unesco.org/en/articles/better-informed-society-freer-society>.
- Charles A E Goodhart. Problems of monetary management: The UK experience. In *Papers in Monetary Economics*, volume 1. Reserve Bank of Australia, 1975.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- Mary L Gray and Siddharth Suri. The humans working behind the AI curtain. *Harvard Business Review*, 9(1):2–5, 2017. URL <https://hbr.org/2017/01/the-humans-working-behind-the-ai-curtain>.
- Mary L Gray and Siddharth Suri. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books, 2019.
- Ben Green. Data science as political action: Grounding data science in a politics of justice. *Journal of Social Computing*, 2(3):249–265, 2021.
- Daniel Greene, Anna Lauren Hoffmann, and Luke Stark. Better, nicer, clearer, fairer: A critical assessment of the movement for ethical artificial intelligence and machine learning. 2019.
- James Grimmelman. The Google dilemma. *New York Law School, Law Review*, 53:939, 2008.

- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206, 2022.
- Steve Guttenberg. Is Spotify unfair to musicians?, 2012. URL <https://www.cnet.com/tech/home-entertainment/is-spotify-unfair-to-musicians/>.
- Kevin Guyan. Queer data. 2022.
- Edward J Hackett, Olga Amsterdamska, Michael Lynch, Judy Wajcman, et al. *The handbook of science and technology studies*. Mit Press Cambridge, 2008.
- Max Haiven. *Crises of imagination, crises of power: Capitalism, creativity and the commons*. Bloomsbury Publishing, 2014.
- Blake Hallinan and Ted Striphas. Recommended for you: The Netflix prize and the production of algorithmic culture. *New media & society*, 18(1):117–137, 2016.
- Donna Haraway. A cyborg manifesto: Science, technology, and socialist-feminism in the late twentieth century. In *The transgender studies reader*, pages 103–118. Routledge, 2013.
- Natali Helberger. On the democratic role of news recommenders. In *Algorithms, automation, and news*, pages 14–33. Routledge, 2021.
- Susan C Herring, Christine Ogan, Manju Ahuja, and Jean C Robinson. Gender and the culture of computing in applied it education. In *Encyclopedia of gender and information technology*, pages 474–481. IGI Global, 2006.
- Shana Higgins and Lua Gregory. *Information literacy and social justice: Radical professional praxis*. Library Juice Press, 2013.
- Tomasz Hollanek. AI transparency: a matter of reconciling design with critique. *AI & Society*, 38(5):2071–2079, 2023.
- Keith Hoskin. The ‘awful’ idea of accountability: Inscribing people into the measurement of objects. In R Munro and J Mouritsen, editors, *Accountability: Power, ethos and technologies of managing*. International Thompson Business Press, London, 1996.
- Jennifer Hsia, Danish Pruthi, Aarti Singh, and Zachary C Lipton. Goodhart’s law applies to NLP’s explanation benchmarks. *arXiv preprint arXiv:2308.14272*, 2023.
- Linus Ta-Lun Huang, Hsiang-Yun Chen, Ying-Tung Lin, Tsung-Ren Huang, and Tzu-Wei Hung. Ameliorating algorithmic bias, or why explainable AI needs feminist philosophy. *Feminist Philosophy Quarterly*, 8(3/4), 2022.
- Ruth Hubbard. Science, facts, and feminism. *Women, science, and technology: A reader in feminist science studies*, pages 148–154, 2001.
- IPCC. The physical science basis, Working group I contribution to the UN IPCC’s fifth assessment report (WG1 AR5), 2013. URL <https://www.ipcc.ch/report/ar5/wg1/>

- Lilly Irani. “Design thinking”: Defending silicon valley at the apex of global labor hierarchies. *Catalyst: Feminism, Theory, Technoscience*, 4(1):1–19, 2018.
- Lilly Irani, Janet Vertesi, Paul Dourish, Kavita Philip, and Rebecca E Grinter. Postcolonial computing: a lens on design and development. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1311–1320, 2010.
- Lilly C Irani and M Six Silberman. Stories we tell about labor: Turkopticon and the trouble with “design”. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 4573–4586, 2016.
- Abigail Z Jacobs and Hanna Wallach. Measurement and fairness. In *Proceedings of the ACM conference on fairness, accountability, and transparency*, pages 375–385, 2021.
- Fredric Jameson. Future city. *New left review*, 21:65, 2003.
- Sheila Jasanoff and Sang-Hyun Kim. Containing the atom: Sociotechnical imaginaries and nuclear power in the United States and South Korea. *Minerva*, 47:119–146, 2009.
- Sheila Jasanoff and Sang-Hyun Kim. *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*. University of Chicago Press, 2015.
- Daxin Jiang, Jian Pei, and Hang Li. Mining search and browse logs for web search: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 4(4):1–37, 2013.
- Anna Jobin, Marcello Ienca, and Effy Vayena. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9):389–399, 2019.
- Jeffrey Alan Johnson. From open data to information justice. *Ethics and Information Technology*, 16:263–274, 2014.
- Stephanie T Jones and Natalie Araujo Melo. We tell these stories to survive: Towards abolition in computer science education. *Canadian Journal of Science, Mathematics and Technology Education*, 21:290–308, 2021.
- Edward Jones-Imhotep. The ghost factories: histories of automata and artificial life. *History and technology*, 36(1):3–29, 2020.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.
- Mariame Kaba. *We do this ’til we free us: Abolitionist organizing and transforming justice*, volume 1. Haymarket Books, 2021.
- Pratyusha Kalluri et al. Don’t ask if artificial intelligence is good or fair, ask how it shifts power. *Nature*, 583(7815):169–169, 2020.
- Shivaram Kalyanakrishnan, Rahul Alex Panicker, Sarayu Natarajan, and Shreya Rao. Opportunities and challenges for artificial intelligence in India. In *Proceedings of the 2018 AAAI/ACM conference on AI, Ethics, and Society*, pages 164–170, 2018.



- Gerald C Kane, Amber G Young, Ann Majchrzak, and Sam Ransbotham. Avoiding an oppressive future of machine learning: A design theory for emancipatory assistants. *Mis Quarterly*, 45(1):371–396, 2021.
- Sucharita Kanjilal. The digital life of caste: Affect, synesthesia and the social body online. *Feminist Media Studies*, pages 1–16, 2023.
- Atoosa Kasirzadeh and Iason Gabriel. In conversation with artificial intelligence: aligning language models with human values. *Philosophy & Technology*, 36(2):1–24, 2023.
- Anna Kata. A postmodern pandora’s box: anti-vaccination misinformation on the internet. *Vaccine*, 28(7):1709–1716, 2010.
- Matthew Kay, Cynthia Matuszek, and Sean A Munson. 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*, pages 3819–3828, 2015.
- Os Keyes, Josephine Hoy, and Margaret Drouhard. Human-computer insurrection: Notes on an anarchist HCI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–13, 2019.
- Michael Klippbahn-Karge, Ann-Kathrin Koster, and Sara Morais dos Santos Bruss. Introduction: Queer AI. In *Queer Reflections on AI*, pages 1–19. Routledge, 2024.
- Mei Kobayashi. Opportunities for women, minorities in information retrieval. *Communications of the ACM*, 60(11):10–11, 2017.
- Ron Kohavi, Randal M Henne, and Dan Sommerfield. Practical guide to controlled experiments on the web: listen to your customers not to the hippo. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 959–967, 2007.
- Ron Kohavi, Diane Tang, and Ya Xu. *Trustworthy online controlled experiments: A practical guide to a/b testing*. Cambridge University Press, 2020.
- Ronny Kohavi, Thomas Crook, Roger Longbotham, Brian Frasca, Randy Henne, Juan Lavista Ferres, and Tamir Melamed. Online experimentation at microsoft. *Data Mining Case Studies*, 11(2009):39, 2009.
- Wendy K Kolmar and Frances Bartkowski. *Feminist theory: A reader*. 1999.
- Juhi Kulshrestha, Motahhare Eslami, Johnnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. Quantifying search bias: Investigating sources of bias for political searches in social media. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, pages 417–432, 2017.
- Srijan Kumar and Neil Shah. False information on web and social media: A survey. *arXiv preprint arXiv:1804.08559*, 2018.

- Michael Kwet. A digital tech deal: Digital socialism, decolonization, and reparations for a sustainable global economy. *Global Information Society Watch*, 2020.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- Kelly Lawler. AI already is making its mark in Hollywood. *USA Today*, pages 01D–01D, 2023.
- Tomo Lazovich, Luca Belli, Aaron Gonzales, Amanda Bower, Uthaipon Tantipongpipat, Kristian Lum, Ferenc Huszar, and Rumman Chowdhury. Measuring disparate outcomes of content recommendation algorithms with distributional inequality metrics. *Patterns*, 3(8), 2022.
- Ursula K Le Guin. *The Wave in the Mind: Talks and Essays on the Writer, the Reader, and the Imagination*. Shambhala Publications, 2004.
- Susan Leavy, Eugenia Siapera, and Barry O’Sullivan. Ethical data curation for AI: An approach based on feminist epistemology and critical theories of race. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 695–703, 2021.
- Joseph Lehman. A brief explanation of the Overton window. *Mackinac Center for Public Policy*, 2014.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.
- Ruohan Li, Jianxiang Li, Bhaskar Mitra, Fernando Diaz, and Asia J Biega. Exposing query identification for search transparency. In *Proceedings of the ACM Web Conference 2022*, pages 3662–3672, 2022.
- Yunqi Li, Yingqiang Ge, and Yongfeng Zhang. Tutorial on fairness of machine learning in recommender systems. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2654–2657, 2021.
- Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. Fairness in recommendation: Foundations, methods, and applications. *ACM Transactions on Intelligent Systems and Technology*, 14(5):1–48, 2023.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation. *arXiv preprint arXiv:2004.01401*, 2020.

- Q Vera Liao and Jennifer Wortman Vaughan. AI transparency in the age of LLMs: A human-centered research roadmap. *arXiv preprint arXiv:2306.01941*, 2023.
- Ann Light. HCI as heterodoxy: Technologies of identity and the queering of interaction with computers. *Interacting with computers*, 23(5):430–438, 2011.
- Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. Pretrained transformers for text ranking: BERT and beyond. *arXiv preprint arXiv:2010.06467*, 2020.
- Jimmy Lin, Daniel Campos, Nick Craswell, Bhaskar Mitra, and Emine Yilmaz. Significant improvements over the state of the art? A case study of the MS MARCO document ranking leaderboard. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2283–2287. ACM, 2021.
- Jimmy Lin, Daniel Campos, Nick Craswell, Bhaskar Mitra, and Emine Yilmaz. Fostering coopetition while plugging leaks: The design and implementation of the MS MARCO leaderboards. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022.
- Conor Linehan and Ben Kirman. Never mind the bollocks, I wanna be anarchi: a manifesto for punk HCI. In *CHI’14 Extended Abstracts on Human Factors in Computing Systems*, pages 741–748. 2014.
- Helen E Longino. Can there be a feminist science? *Hypatia*, 2(3):51–64, 1987.
- Astrid Mager. Algorithmic ideology: How capitalist society shapes search engines. *Information, Communication & Society*, 15(5):769–787, 2012.
- Astrid Mager and Christian Katzenbach. Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified, *New Media & Society* 23(2), Sage Journals, 2021.
- Luiza Prado de O Martins. Privilege and oppression: Towards a feminist speculative design. *Design’s Big Debates DRS International Conference*, 2014.
- Karl Marx. The fetishism of commodities and the secret thereof. *Capital: A Critique of Political Economy*, 1:71–83, 1867.
- Abraham Harold Maslow. A dynamic theory of human motivation. 1958.
- Leslie McCall. The complexity of intersectionality. *Signs: Journal of women in culture and society*, 30(3):1771–1800, 2005.
- Dan McQuillan. Anti-fascist AI. In *Resisting AI*, pages 135–148. Bristol University Press, 2022.
- Yusuf Mehdi. Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web, 2023. URL <https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/>.

- Yusuf Mehdi. Bringing the full power of Copilot to more people and businesses, 2024. URL <https://blogs.microsoft.com/blog/2024/01/15/bringing-the-full-power-of-copilot-to-more-people-and-businesses/>.
- Anay Mehrotra and Nisheeth Vishnoi. Fair ranking with noisy protected attributes. *Advances in Neural Information Processing Systems*, 35:31711–31725, 2022.
- Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. Auditing search engines for differential satisfaction across demographics. In *Proc. WWW*, pages 626–633, 2017.
- Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th acm international conference on information and knowledge management*, pages 2243–2251, 2018.
- Brian Merchant. *Blood in the Machine: The Origins of the Rebellion Against Big Tech*. Little, Brown, 2023.
- Donald Metzler, Yi Tay, Dara Bahri, and Marc Najork. Rethinking search: making domain experts out of dilettantes. In *Acm sigir forum*, volume 55, pages 1–27. ACM New York, NY, USA, 2021.
- Milagros Miceli, Julian Posada, and Tianling Yang. Studying up machine learning data: Why talk about bias when we mean power? *Proceedings of the ACM on Human-Computer Interaction*, 6(GROUP):1–14, 2022.
- Diane P Michelfelder, Galit Wellner, and Heather Wiltse. Designing differently: Toward a methodology for an ethics of feminist technology design. *The ethics of technology: methods and approaches*, pages 193–218, 2017.
- Boaz Miller. Is technology value-neutral? *Science, Technology, & Human Values*, 46(1): 53–80, 2021.
- Dan Milmo. ‘Impossible’ to create AI tools like ChatGPT without copyrighted material, OpenAI says. *The Guardian*, 8 January 2024. URL <https://www.theguardian.com/technology/2024/jan/08/ai-tools-chatgpt-copyrighted-material-openai>.
- Bhaskar Mitra. *Neural Methods for Effective, Efficient, and Exposure-Aware Information Retrieval*. PhD thesis, UCL (University College London), 2021.
- Bhaskar Mitra. Emancipatory information retrieval. *arXiv preprint arXiv:2501.19241*, 2025.
- Bhaskar Mitra and Nick Craswell. An introduction to neural information retrieval. *Foundations and Trends in Information Retrieval*, 2018.
- Bhaskar Mitra, Henriette Cramer, and Olya Gurevich. Sociotechnical implications of generative artificial intelligence for information access. *arXiv preprint arXiv:2405.11612*, 2024.

- Shakir Mohamed, Marie-Therese Png, and William Isaac. Decolonial ai: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33:659–684, 2020.
- Thema Monroe-White. Emancipatory data science. 2021.
- Jared Moore. Towards a more representative politics in the ethics of computer science. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 414–424, 2020.
- Felipe Azevedo Moretti, Vanessa Elias de Oliveira, and Edina Mariko Koga da Silva. Access to health information on the internet: a public health issue? *Revista da Associação Médica Brasileira*, 58:650–658, 2012.
- Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. Controlling fairness and bias in dynamic learning-to-rank. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 429–438, 2020.
- Sara Morrison. Dark patterns, the tricks websites use to make you say yes, explained. *Vox*, 1 April 2021.
- Stephen M Mutula. Digital divide and economic development: Case study of sub-saharan africa. *The Electronic Library*, 26(4):468–489, 2008.
- Nicola Neophytou, Bhaskar Mitra, and Catherine Stinson. Revisiting popularity and demographic biases in recommender evaluation and effectiveness. In *Proceedings of the European Conference on Information Retrieval (ECIR)*, 2022.
- Safiya Umoja Noble. Algorithms of oppression. In *Algorithms of oppression*. New York university press, 2018.
- Rodrigo Nogueira and Kyunghyun Cho. Passage re-ranking with BERT. *CoRR*, abs/1901.04085, 2019. URL <http://arxiv.org/abs/1901.04085>.
- Alexandra Olteanu, Jean Garcia-Gathright, Maarten de Rijke, Michael D Ekstrand, Adam Roegiest, Aldo Lipani, Alex Beutel, Alexandra Olteanu, Ana Lucic, Ana-Andreea Stoica, et al. Facts-ir: fairness, accountability, confidentiality, transparency, and safety in information retrieval. In *ACM SIGIR Forum*, volume 53, pages 20–43. ACM New York, NY, USA, 2019.
- OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Will Oremus. Big tobacco. big pharma. big tech? 2017. URL <https://slate.com/technology/2017/11/how-silicon-valley-became-big-tech.html>.
- Camille Parmesan, Mike D Morecroft, and Yongyut Trisurat. Climate change 2022: Impacts, adaptation and vulnerability, 2022.

- Gourab K. Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *WWW*, pages 1194–1204. ACM / IW3C2, 2020.
- Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. Fair ranking: a critical review, challenges, and future directions. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1929–1942, 2022.
- David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. 2021.
- David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David R So, Maud Texier, and Jeff Dean. The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7):18–28, 2022.
- Billy Perrigo. Exclusive: OpenAI used Kenyan workers on less than \$2 per hour to make chatgpt less toxic. *Time*, 18 January, 2023. URL <https://time.com/6247678/openai-chatgpt-kenya-workers/>.
- Bryan Pfaffenberger. Fetishised objects and humanised nature: towards an anthropology of technology. *Man*, pages 236–252, 1988.
- Kavita Philip, Lilly Irani, and Paul Dourish. Postcolonial computing: A tactical survey. *Science, Technology, & Human Values*, 37(1):3–29, 2012.
- Christine Pinney, Amifa Raj, Alex Hanna, and Michael D Ekstrand. Much ado about gender: Current practices and future recommendations for appropriate gender-aware information access. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval*, pages 269–279, 2023.
- Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. Fairness in rankings and recommendations: an overview. *The VLDB Journal*, pages 1–28, 2022.
- Lizzie Plausic. Spotify’s year in music shows just how little we pay artists for their music. *The Verge*, 2015.
- Marie-Therese Png. At the tensions of south and north: Critical roles of global south stakeholders in AI governance. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1434–1445, 2022.
- Gianfranco Polizzi. Information literacy in the digital age: Why critical digital literacy matters for democracy. *Informed Societies: Why information literacy matters for citizenship, participation and democracy*, pages 1–23, 2020.
- Sayantana Polley, Rashmi Raju Koparde, Akshaya Bindu Gowri, Maneendra Perera, and Andreas Nuernberger. Towards trustworthiness in the context of explainable search. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2580–2584, 2021.

- Mark Poynting and Erwan Rivault. 2023 confirmed as world’s hottest year on record, *BBC*, 9 January 2024.
- Vinodkumar Prabhakaran, Margaret Mitchell, Timnit Gebru, and Iason Gabriel. A human rights-based approach to responsible AI. *arXiv preprint arXiv:2210.02667*, 2022.
- Amifa Raj and Michael D Ekstrand. Comparing fair ranking metrics. *arXiv preprint arXiv:2009.01311*, 2020.
- Amifa Raj and Michael D Ekstrand. Measuring fairness in ranked results: An analytical and empirical comparison. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 726–736, 2022.
- Amifa Raj, Bhaskar Mitra, Nick Craswell, and Michael Ekstrand. Patterns of gender-specializing query reformulation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2241–2245, 2023.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don’t know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*, 2018.
- Eric Ravenscraft. How to spot—and avoid—dark patterns on the web. *Wired*, 29 July 2020.
- Teju Ravilochan. The Blackfoot wisdom that inspired Maslow’s hierarchy. *Resilience*. June, 18:2021, 2021.
- Reblog of report from Northland Poster. Attributing words, 2006. URL <https://unnecessaryevils.blogspot.com/2008/11/attributing-words.html>.
- Sarah T Roberts. *Behind the screen*. Yale University Press, 2019.
- Sarah T. Roberts. Your AI is a human. *Your Computer is on fire*, MIT Press, 2021.
- Arundhati Roy. *War talk*. South End Press, 2003.
- Stuart Russell, Daniel Dewey, and Max Tegmark. Research priorities for robust and beneficial artificial intelligence. *AI magazine*, 36(4):105–114, 2015.
- Nithya Sambasivan, Erin Arnesen, Ben Hutchinson, Tulsee Doshi, and Vinodkumar Prabhakaran. Re-imagining algorithmic fairness in India and beyond. In *Proceedings of the ACM conference on fairness, accountability, and transparency*, pages 315–328, 2021.
- F. Saracco, M. Viviani, et al. Overview of ROMCIR 2021: workshop on reducing online misinformation through credible information retrieval. In *ROMCIR 2021 CEUR Workshop Proceedings*, volume 2838, 2021.
- Deepak Saxena, P.J. Wall, and Dave Lewis. Artificial intelligence (AI) ethics: A critical realist emancipatory approach. In *2023 IEEE International Symposium on Technology and Society (ISTAS)*, pages 1–5. IEEE, 2023.

- Markus Schedl, Emilia Gómez, and Elisabeth Lex. Retrieval and recommendation systems at the crossroads of artificial intelligence, ethics, and regulation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3420–3424, 2022.
- Anuschka Schmitt, Thiemo Wambsganss, and Andreas Janson. Designing for conversational system trustworthiness: the impact of model transparency on trust and task performance. *ECIS 2022 Research Papers*, 2022.
- New Scientist. Covid-19: the story of a pandemic. *New Scientist*, 10, 2021.
- Chirag Shah and Emily M Bender. Situating search. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*, pages 221–232, 2022.
- Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. Combating fake news: A survey on identification and mitigation techniques. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(3):1–42, 2019.
- Grace Shaw, Margot Brereton, and Paul Roe. Mobile phone use in Australian indigenous communities: future pathways for HCI4D. In *Proceedings of the 26th Australian computer-human interaction conference on designing futures: the future of design*, pages 480–483, 2014.
- Konnor Shetler. AI and consent: What the SAG-AFTRA and WGA agreements tell us about the future of generative AI. *Seton Hall University, Student Works*, 2024.
- Shreeti Shubham. Caste and the digital sphere. *Shuddhashar Youth Journal*, 2022.
- Ashudeep Singh and Thorsten Joachims. Fairness of exposure in rankings. In *KDD*, pages 2219–2228. ACM, 2018.
- Ashudeep Singh and Thorsten Joachims. Policy learning for fairness in ranking. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5427–5437. Curran Associates, Inc., 2019.
- Jaspreet Singh and Avishek Anand. Posthoc interpretability of learning to rank models using secondary training data. *arXiv preprint arXiv:1806.11330*, 2018.
- Jaspreet Singh and Avishek Anand. Exs: Explainable search using local model agnostic interpretability. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 770–773, 2019.
- Jessie J Smith and Lex Beattie. Recsys fairness metrics: Many to use but which one to choose? *arXiv preprint arXiv:2209.04011*, 2022.
- Thomas Smyth and Jill Dimond. Anti-oppressive design. *Interactions*, 21(6):68–71, 2014.
- Nick Srnicek. *Platform capitalism*. John Wiley & Sons, 2017.



- Nick Srnicek and Alex Williams. *Inventing the future: Postcapitalism and a world without work*. Verso Books, 2015.
- Susan Leigh Sta. The politics question in feminist science and technology projects: the queering of infrastructure. Talk presented at the Technology and Democracy – Comparative Perspectives Conference, University of Oslo, Norway, 1997. URL <http://www.drury.edu/faculty/Ess/Technology/starr.htm>.
- Cella M Sum, Franchesca Spektor, Rahaf Alharbi, Leya Breanna Baltaxe-Admony, Erika Devine, Hazel Anneke Dixon, Jared Duval, Tessa Eagle, Frank Elavsky, Kim Fernandes, et al. Challenging ableism: A critical turn toward disability justice in HCI. *XRDS: Crossroads, The ACM Magazine for Students*, 30(4):50–55, 2024.
- Huatong Sun. Critical design sensibility in postcolonial conditions. *AoIR Selected Papers of Internet Research*, 2013.
- Latanya Sweeney. Discrimination in online ad delivery. *Communications of the ACM*, 56(5):44–54, 2013.
- Ali Swenson and Christine Fernando. As social media guardrails fade and AI deepfakes go mainstream, experts warn of impact on elections, 2023. URL <https://apnews.com/article/election-2024-misinformation-ai-social-media-trump-6119ee6f498db10603b3664e9ad3e87e>.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.
- Alex Tamkin, Miles Brundage, Jack Clark, and Deep Ganguli. Understanding the capabilities, limitations, and societal impact of large language models. *arXiv preprint arXiv:2102.02503*, 2021.
- Dana H Taplin and Hélène Clark. Theory of change basics: A primer on theory of change. *New York: ActKnowledge*, page 72, 2012.
- Jonathan Taplin. *Move fast and break things: How Facebook, Google, and Amazon have cornered culture and what it means for all of us*. Pan Macmillan, 2017.
- Adam Taylor. A historic rise in global conflict deaths suggests a violent new era. 2023. URL <https://www.washingtonpost.com/world/2023/06/29/conflict-war-deaths-global-peace-rise-casualty/>.
- Luke Taylor. Covid-19: True global death toll from pandemic is almost 15 million, says who. *BMJ: British Medical Journal (Online)*, 377:ø1144, 2022.
- Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. Large language models can accurately predict searcher preferences. In *Proc. SIGIR*, 2024.
- Rachel L. Thomas and David Uminsky. Reliance on metrics is a fundamental challenge for AI. *Patterns*, 3(5), 2022.

- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.
- Sophie Toupin. Shaping feminist artificial intelligence. *New Media & Society*, 26(1):580–595, 2024.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Kathie Treen, Hywel Williams, and Saffron O’Neill. Guest post: How climate change misinformation spreads online, *Carbon Brief*, 2020.
- Anna Lowenhaupt Tsing. *The mushroom at the end of the world: On the possibility of life in capitalist ruins*. Princeton University Press, 2015.
- Sherry Turkle. Computational reticence: Why women fear the intimate machine. In *Technology and women’s voices*, pages 44–60. Routledge, 2004.
- J. Turrentine. Climate misinformation on social media is undermining climate action, 2022. URL <https://www.nrdc.org/stories/climate-misinformation-social-media-undermining-climate-action>.
- United Nations Meetings Coverage and Press Releases. With highest number of violent conflicts since second world war, united nations must rethink efforts to achieve, sustain peace, speakers tell security council. <https://press.un.org/en/2023/sc15184.doc.htm>, 2023.
- Aleksandra Urman and Mykola Makhortykh. The silence of the LLMs: Cross-lingual analysis of political bias and false information prevalence in Chatgpt, Google Bard, and Bing chat. *Telematics and Informatics 96(C)*, 2023.
- Aleksandra Urman and Mykola Makhortykh. “foreign beauties want to meet you”: The sexualization of women in google’s organic and sponsored text search results. *New media & society*, 26(5):2932–2953, 2024.
- Palashi Vaghela, Steven J Jackson, and Phoebe Sengers. Interrupting merit, subverting legibility: Navigating caste in ‘casteless’ worlds of computing. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–20, 2022a.
- Palashi Vaghela, Ramaravind Kommiya Mothilal, Daniel Romero, and Joyojeet Pal. Caste capital on twitter: A formal network analysis of caste relations among indian politicians. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1):1–29, 2022b.
- Joana Varon and Paz Peña. Artificial intelligence and consent: A feminist anti-colonial critique. *Internet Policy Review*, 10(4):1–25, 2021.

- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Proc. of NeurIPS*, 2017.
- Suzan Verberne, Hussein Suleman, Luca Soldaini, and Avijit Ghosh. Report on the SIGIR 2023 session on diversity, equity and inclusivity. In *ACM SIGIR Forum*, volume 57, pages 1–2. ACM, 2024.
- Pieter Verdegem. Dismantling AI capitalism: the commons as an alternative to the power concentration of big tech. *AI & society*, pages 1–11, 2022.
- Manisha Verma and Debasis Ganguly. LIRME: locally interpretable ranking model explanation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1281–1284, 2019.
- Liisa von Hellens, Sue Nielsen, Kaylene Clayton, and Jenine Beekhuyzen. Conceptualising gender and it: Australians taking action in germany. *Proceedings of QualIT2007- Qualitative Research in IT & IT in Qualitative Research*, 2017.
- Ellen M Voorhees. Coopetition in ir research. In *ACM SIGIR Forum*, volume 54, pages 1–3. ACM, 2021.
- Ellen M Voorhees, Donna K Harman, et al. *TREC: Experiment and evaluation in information retrieval*, volume 63. MIT press Cambridge, 2005.
- Sanne Vrijenhoek, Mesut Kaya, Nadia Metoui, Judith Möller, Daan Odijk, and Natali Helberger. Recommenders with a mission: assessing diversity in news recommendations. In *Proceedings of the 2021 conference on human information interaction and retrieval*, pages 173–183, 2021.
- Judy Wajcman. *Feminism confronts technology*. Penn State Press, 1991.
- Judy Wajcman. *Technofeminism*. Cambridge: Polity. 2004.
- Judy Wajcman. Feminist theories of technology. *Cambridge journal of economics*, 34(1): 143–152, 2010.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32, 2019.
- Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. A survey on the fairness of recommender systems. *ACM Transactions on Information Systems*, 41(3):1–43, 2023.
- Tom Warren. Microsoft’s new Copilot pro brings AI-powered office features to the rest of us, 2024. URL <https://www.theverge.com/2024/1/15/24038711/microsoft-copilot-pro-office-ai-apps>.

- Carol H Weiss. Nothing as practical as good theory: Exploring theory-based evaluation for comprehensive community initiatives for children and families. *New approaches to evaluating community initiatives: Concepts, methods, and contexts*, 1:65–92, 1995.
- Galit Wellner and Tiran Rothman. Feminist AI: Can we expect our AI systems to become feminist? *Philosophy & Technology*, 33(2):191–205, 2020.
- Hannes Werthner, Carlo Ghezzi, Jeff Kramer, Julian Nida-Rümelin, Bashar Nuseibeh, Erich Prem, and Allison Stanger. *Introduction to Digital Humanism: A Textbook*. Springer Nature, 2024.
- Meredith Whittaker. The steep cost of capture. *Interactions*, 28(6):50–55, 2021.
- Dorothy Wickenden. A reckoning at Facebook, *The New Yorker*, 19 February 2018. URL <https://www.newyorker.com/podcast/political-scene/a-reckoning-at-facebook>.
- David Gray Widder, Derrick Zhen, Laura Dabbish, and James Herbsleb. It’s about power: What ethical concerns do software engineers have, and what do they (feel they can) do about them? In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 467–479, 2023.
- Wikipedia contributors. Theory of change — Wikipedia, the free encyclopedia, 2013. URL [https://en.wikipedia.org/wiki/Theory\\_of\\_Change](https://en.wikipedia.org/wiki/Theory_of_Change).
- Darcia Wilkinson, Michael Ekstrand, Janet A. Vertesi, and Alexandra Olteanu. Theories of change in responsible AI. Craft session at the ACM conference on fairness, accountability, and transparency, 2023. URL <https://facctconference.org/2023/acceptedcraft#theor>.
- Adrienne Williams, Milagros Miceli, and Timnit Gebru. The exploited labor behind artificial intelligence. *Noema Magazine*, 13, 2022. URL <https://www.noemamag.com/the-exploited-labor-behind-artificial-intelligence/>.
- Rua M Williams, Kathryn Ringland, Amelia Gibson, Mahender Mandala, Arne Maibaum, and Tiago Guerreiro. Articulations toward a crip HCI. *Interactions*, 28(3):28–37, 2021.
- Tom Williams and Kerstin Sophie Haring. No justice, no robots: From the dispositions of policing to an abolitionist robotics. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, pages 566–575, 2023.
- Heike Winschiers-Theophilus and Nicola J Bidwell. Toward an Afro-centric indigenous HCI paradigm. *International Journal of Human-Computer Interaction*, 29(4):243–255, 2013.
- Pak-Hang Wong. Dao, harmony and personhood: Towards a confucian ethics of technology. *Philosophy & technology*, 25(1):67–86, 2012.
- Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. Sustainable AI: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems*, 4:795–813, 2022a.

- Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. Joint multisided exposure fairness for recommendation. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*, 2022b.
- Haolun Wu, Bhaskar Mitra, and Nick Craswell. Towards group-aware search success. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*, pages 123–131, 2024.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*, 2018.
- Meg Young, PM Krafft, and Michael A Katell. A call for scholar activism: A response to power and technology. *AI Activism*, 28:43, 2021.
- Peter K Yu. Bridging the digital divide: Equality in the information age. *Cardozo Arts & Ent. LJ*, 20:1, 2002.
- Brandy Zadrozny. Disinformation poses an unprecedented threat in 2024 — and the U.S. is less ready than ever, 2024. URL <https://www.nbcnews.com/tech/misinformation/disinformation-unprecedented-threat-2024-election-rcna134290>.
- Hamed Zamani, Fernando Diaz, Mostafa Dehghani, Donald Metzler, and Michael Bendersky. Retrieval-enhanced machine learning. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2875–2886, 2022.
- Meike Zehlike and Carlos Castillo. Reducing disparate exposure in ranking: A learning to rank approach. In *Proceedings of the web conference 2020*, pages 2849–2855, 2020.
- Meike Zehlike, Ke Yang, and Julia Stoyanovich. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000*, 2021.
- Meike Zehlike, Ke Yang, and Julia Stoyanovich. Fairness in ranking, part i: Score-based ranking. *ACM Computing Surveys*, 55(6):1–36, 2022a.
- Meike Zehlike, Ke Yang, and Julia Stoyanovich. Fairness in ranking, part II: Learning-to-rank and recommender systems. *ACM Computing Surveys*, 55(6):1–41, 2022b.
- Yongfeng Zhang, Xu Chen, et al. Explainable recommendation: A survey and new perspectives. *Foundations and Trends in Information Retrieval*, 14(1):1–101, 2020.
- Xinyi Zhou and Reza Zafarani. Fake news: A survey of research, detection methods, and opportunities. *arXiv preprint arXiv:1812.00315*, 2, 2018.
- Xinyi Zhou and Reza Zafarani. A survey of fake news: Fundamental theories, detection methods, and opportunities. *ACM Computing Surveys (CSUR)*, 53(5):1–40, 2020.
- Honglei Zhuang, Xuanhui Wang, Michael Bendersky, Alexander Grushetsky, Yonghui Wu, Petr Mitrichev, Ethan Sterling, Nathan Bell, Walker Ravina, and Hai Qian. Interpretable learning-to-rank with generalized additive models. *arXiv preprint arXiv:2005.02553*, 2020.

Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.

Shoshana Zuboff. The age of surveillance capitalism. In *Social Theory Re-Wired*, pages 203–213. Routledge, 2023.