

Towards Critical IR Theories and Practices

Bhaskar Mitra

Independent Researcher

Tiohtià:ke / Mooniyang / Montréal, Canada

`bhaskar.mitra@acm.org`

Abstract

Belkin and Robertson urged us, half a century ago, to develop a theoretical foundation for understanding what constitutes *societal good* that can inform information retrieval (IR) research and serve as a basis for determining when we should limit our scientific inquiry in the face of demands that are contradictory to societal good. In this article, I argue that to achieve this, IR should embrace critical theories and practices in our work, and shift away from the dominant liberal frame through which much of the IR community today view societal concerns in context of our research. Unlike the liberal frame, the critical frame explicitly adopts nondomination as its stated goal which can clarify our conceptualization of societal good within the field, provide necessary theoretical underpinning that Belkin and Robertson urged the community to develop, and serve as a basis for critical appraisals of our progress in enacting desired societal change.

1 Introduction

Half a century ago, [Belkin and Robertson \[1976\]](#) called on the information science (IS) community to recognize the field’s salient responsibilities to society. They urged the community to develop a theoretical foundation for understanding what constitutes societal good that can inform information access research and technology development. In particular, they wrote:

“Thus, information science must become theoretically self-conscious, and self-consciously based upon a social ideology. Inevitably such a position will lead to conflicts between ethical and politically expedient imperatives. Not only must we be willing to limit our scientific inquiry according to public interest, we must have a concept of the public interest and be able to argue our position when faced with contradictory demands. And in order to do this successfully, we must have developed our theoretical position and our limitations before being faced with such conflicting demands (otherwise, will we realize the conflict before it is too late?).”

Unfortunately, in spite of their prescient call-to-action, five decades later the field of information retrieval (IR), which could be situated at the intersection of computing and IS, finds itself in exactly the troubling situation that Belkin and Robertson warned us about. IR research and practices are often motivated by grand ambitions like “organiz[ing] the world’s

information and mak[ing] it universally accessible and useful”¹ or supporting human cognition with information artifacts [Zobel, 2018]. But the reality of dominant search and recommendation platforms is that they codify oppression [Noble, 2018] and have materialized massive infrastructures for (i) large-scale user surveillance [Zuboff, 2023], (ii) commodification and harvesting of user attention [Wu, 2017], and increasingly (iii) manipulating user’s political and commercial preferences via generation of persuasive language and visualizations using generative AI [Mitra et al., 2024]. IR platforms have also increasingly become targets of ideological capture by authoritarian state and private actors because they represent sites of struggle between oppression and justice [Mitra et al., 2026b]. These concerns are particularly serious and urgent today in light of the rising levels of democratic erosion worldwide and the increasing concentration of economic and political power in the hands of Big Tech.

Simultaneously, there has been a significant uptick in IR scholarship around fairness and explainability since 2018, following the third Strategic Workshop in Information Retrieval in Lorne (SWIRL) [Culpepper et al., 2018] where IR researchers in attendance recognized fairness, accountability, confidentiality, and transparency in IR (“FACT IR”) as consequential research directions for the field. However, how much of this research has in practice meaningfully translated into societally positive outcomes is unclear [Mitra, 2025b]. The current frames through which the IR community approaches societal concerns have also failed to thwart the community from engaging in research that support user surveillance and manipulation, and serious concerns about how proprietary platforms have concentrated power in the hands of few individuals and institutions making our information ecosystem vulnerable to authoritarian capture have also hitherto received insufficient attention.

The under-theorization of “societal good” in IR is a key contributing factor here. After all, how can we appraise our progress in affecting positive societal change—or limit our scientific inquiry when faced with demands contradictory to societal good—if we cannot articulate what societal good we are aiming for? This has led to several recent calls for the IR community to explicate its normative positions [Vrijenhoek et al., 2023], values [Trippas et al., 2025], and sociotechnical imaginaries [Mitra, 2025b]. Consequently, the IR-for-Good track at the 2026 edition of the European Conference on Information Retrieval (ECIR) adopted as its theme the question of “*what is IR-for-Good?*” [Mitra and Heuss, 2025] which again attempts to get at this exact underlying tension in the field.

Against this backdrop, this article invites the IR community to engage in thoughtful deliberation over how to develop a shared normative position on societal good. I believe this would require IR to embrace *critical* theories and practices in our work, and shift away from the dominant *liberal* frame through which much of the IR community today view societal concerns in context of our ongoing research. The liberal frame concerns itself with designing a *good system* without explicating what *societal good* it wants to realize in the process. In contrast, the critical frame adopts realizing a *nondominative society* as its stated goal which can clarify our conceptualization of societal good within the field, provide necessary theoretical underpinning that Belkin and Robertson urged the community to develop, and serve as a basis for critical appraisals of our progress in enacting desired societal change. Furthermore, the critical frame grounds itself in structural understanding of intersecting systems of oppression, and extends our purview of concerns beyond how the system behaves to also include the sociopolitical context in which the system is embedded and who gets to exert

¹Source: Google’s mission statement (<https://about.google/company-info/>)

Table 1. Contrasting the liberal and critical frames with respect to their goals, their methods and practices, and their underlying values, politics, and epistemologies.

	Liberal frame	Critical frame
Goals	Desired system attributes, <i>e.g.</i> , fairness, confidentiality, transparency, and sustainability.	Desired societal outcomes, <i>e.g.</i> , social justice, emancipation, and democracy.
Methods and practices	Algorithmic harm mitigation, <i>e.g.</i> , fair ranking, fact-checking, privacy mechanisms, transparency, and explainability.	Sociotechnical mitigation, repair, refusal, building counter-technologies; foregrounds praxis and movement building.
Values, politics, and epistemologies	Sometimes (mis-)represented as apolitical; politically espouses individual rights, equality, freedom of speech, and market economies; philosophically positivist and tied to technological and instrumental rationality.	Explicitly political and politically decolonial, anti-racist, feminist, queer, pro-labor, anti-capitalist, anti-casteist, anti-ableist, and abolitionist; philosophically social constructionist and principally nondominative.

power and control over the system. The critical approach is also not constrained to exploring algorithmic mitigation exclusively and includes in its arsenal the right of refusal [Barocas et al., 2020] and to build radical alternatives. I implore the IR community to evolve our current research practices to rigorously incorporate critical theories and perspectives, and hold ourselves accountable for our impact on society.

2 The two frames of IR & Society

Within the relatively short period between 2018-2026, we observe that IR has employed two distinct frames—the *liberal* and the *critical*—to motivate societal concerns in IR research. A frame, in this context, may be loosely defined as comprising of (i) a set of desired goals, (ii) a collection of methods and practices, and (iii) a body of underlying values, politics, and epistemologies that distinguish it from other frames. Table 1 summarizes some of the key distinctions between the liberal and the critical frames along these dimensions.

2.1 The liberal frame

The liberal frame emphasizes desired attributes of the IR system—*e.g.*, fairness, confidentiality, transparency, and sustainability—and consequently focuses on algorithmic mechanisms to achieve these objectives. We can see this reflected in the call for “FACT IR” at SWIRL 2018 [Culpepper et al., 2018], and in the technical research agenda proposed for Responsible IR by the FACTS-IR workshop [Olteanu et al., 2021] organized the following year. Subsequently, there has been a broad body of work on fairness [Ekstrand et al., 2022; Zehlike et al.,

2021; Patro et al., 2022; Wang et al., 2023; Li et al., 2023], explainability [Zhang et al., 2020; Anand et al., 2022], and other related topics in IR.

The liberal frame largely distances itself from any explicit political alignments. But the absence of political expression does not imply an absence of politics in the research conducted through this frame. Instead, the liberal frame maybe better understood as valuing individual rights, equality, freedom of speech, and market economies. IR has historically held positivist notions about the “neutrality” of information and its research practices, and operated within the frames of technological and instrumental rationality [Pyati, 2006]. The liberal frame also draws from these IR traditions and is grounded in related epistemologies.

By defining its goals in terms of what constitutes a *good system*, the liberal frame subtly sidesteps the question of what constitutes a *good society*. The latter questions is saliently more political and I speculate that posing it might be a source of much discordance within the IR community whose members reflect vastly different values, cultures, politics, identities, privileges, lived experiences, and understandings of history. From a cursory look, this may seem like a clever strategy to allow the community to start building a foundation for IR research that prioritizes societal good without waiting on the community to reach reasonable political consensus on what societal good we are aiming for. However, after several years of working on fair ranking and related topics through this lens, I have reached the conclusion that any attempt to bypass or depoliticize the question of what constitutes societal good is ultimately self-defeating and cannot lead us towards ensuring positive societal impact.

The challenges of defining a good system without defining a good society is multifold. Let’s take fairness in IR, for example. It is reasonable to assume that a majority of the IR community prefer that our IR systems be fair. However, how we can move fairness from an ideal to a practical implementation in the context of historical dimensions of marginalization—such as race and gender—is still unclear in spite of the substantive body of existing research in this area. A fundamental challenge is that racial and gender categories are multidimensional, relational, and socially-constructed. However, as Hanna et al. [2020] and Pinney et al. [2023] point out, typical fairness frameworks represent these categories as single-dimensional variables that can take one of few discrete values and treat groups as interchangeable erasing the context of historical harms and the distinct nature of marginalization faced by different groups. As Hanna et al. [2020] summarize: “In short, group fairness approaches try to achieve sameness across groups without regard for the difference between the groups. Group fairness offers an incomplete version of what a race-conscious policy would be.” Practical deployment of fairness interventions also raises concerns of further intensification of data collection on marginalized groups and creation of population-scale databases containing categorical data on algorithmic subjects that risk being instrumentalized by state actors and other powerful entities to target historically marginalized groups.

Fairness constructs in IR also typically treat information artifacts as interchangeable commodities valued through the lens of relevance, source authority, factuality, and other similar factors with little regard to their complex relationship with societal concerns. This is evident if we consider fairness principles in IR like *equity of attention*² [Biega et al., 2018] and *equal expected exposure*³ [Diaz et al., 2020]. However, information artifacts are in fact *not* commodities but rather objects that are defined not just by their relevance to the query but by

²“A ranking offers equity of attention if each subject receives attention proportional to its relevance.”

³“Given a fixed information need, no item should be exposed (in expectation) more or less than any other item of the same relevance”

their complex relationships with societal concerns and their role in our collective sense-making of our place in this world. For example, say, there are two articles both deemed topically “relevant” to the topic of the civil rights movement in the US but one article is grounded in anti-racist epistemology while the other espouses white supremacist point-of-view; surely these two articles do not deserve equal exposure. This example fits into a broader pattern of glaring misapplication of the bias and fairness construct in IR in the context of political ideologies. There are several examples of IR papers⁴ that study whether specific IR methods give disparate representation to different ends of the political spectrum (*e.g.*, between political parties, or left *vs.* right) or different stances on politically charged topics. However, these papers rarely consider how we can make normative judgments on how much different perspectives deserve exposure, the absence of which reinforces the idea that maybe exposure should be equally distributed over them. Such misguided scholarly position fails to adequately recognize the continuously shifting Overton window of public discourse [Lehman, 2014] and disregards the fact that different political beliefs do not hold equal merit—to treat them otherwise would amount to *algorithmic bothsiding* [Mitra, 2025b].

We should not think of this notion of *societal value* (or its corollary, *societal harmfulness*) as an attribute of the commodified information artifact removed from its relational context that can be determined autocratically by the platform and incorporated into our existing notions of relevance or utility of the information—because that would elevate those who wield power over the platform to a position of gatekeeping what constitutes acceptable speech. Instead, this calls for acknowledging the complex processes of social construction of information and the necessity to reimagine how we can encode spaces for community participation, deliberation, and negotiation within our IR systems that is informed by structural understanding of power relations between dominant and oppressed groups. The liberal temptation to cast concerns of social justice as a fairness problem also discourages exploration of alternative interventions. For example, instead of algorithmically fixing the under-representation of women and people of color in image search results for the query “CEO”, we could reclaim space on the search result page to teach the history of erasure and undervaluing of gendered and racial labor [Mitra, 2025b]. The choice here depends on whether our goal is to defend the system against accusations of bias or to materially support the causes of social justice.

Given these concerns, it is unsurprising that there is little evidence that the existing body of fairness research has found meaningful application in dominant search platforms in the context of social justice concerns [Mitra, 2025b]; and in the absence of meaningful societal impact, the primary outcome of this body of work—which includes significant contributions from industry—seems largely to convince ourselves of our own altruism and ethics-wash the tech industry. Herein, lies one of the more serious concerns with the liberal frame that in the absence of a clear articulation of our societal goals, any specification of desired system attributes, like fairness, will be ultimately coopted and systemically stripped of its critical interpretation by those who hold social, economic, and political power over our information access platforms to ensure that this research does not in turn threaten their dominance in society. This is evident in how large corporations editorialize “Responsible AI” research coming out of their institutions to ensure that the critical eye of scholarship is always pointed outwards in search for the “bad actor” and must never be turned inwards to challenge their concentration of wealth and power or their complicity in the military-industrial complex [Mi-

⁴I intentionally do not list specific papers as examples here because my goal is not to shame the authors, but to grow our collective sociopolitical sensitivity to these concerns.

tra, 2026]. Putting it differently, only the type of research that does not challenge oppressive power—but instead enshrines it—is tolerated, supported, funded, and publicized by the powers that be. Even if the IR research community produced effective algorithmic interventions, we must reconcile with the reality that we do not hold the power to enforce the deployment of these interventions in real-world systems; in fact we should recognize that our work is happening in the backdrop of state actors and high-net-worth individuals trying to actively capture our information access platforms for their own ideological goals [Mitra et al., 2026b]. The questions of power concentration over our platforms and ecosystems and large power asymmetries within our community spaces are not orthogonal to our research. They should be intentionally centered in our work to encourage honest critical discourse within the community and development of power-aware research agendas and theories of change.

While I have so far argued my position in the context of fairness research, similar critique is relevant to other concerns of the liberal frame, such as transparency and sustainability. Transparency in IR has come to denote a body of research that is largely predicated on the goal of convincing users to trust the IR system. Contrast this with the more critical notion of transparency that Hollanek [2023] espouses when he writes “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”. Similarly, concerns of environmental impact in IR have largely been met with proposals for algorithmic and procedural interventions to bring down the cost of training and inference of models—*e.g.*, Scells et al. [2022]—without adequate consideration of power relationships between different real-world actors, their incentives, and the well-known phenomenon of Jevons paradox⁵ that are likely to render these well-intentioned strategies ineffectual. Another relevant example is the emerging literature on model alignment approaches like reinforcement learning from human feedback (RLHF) [Christiano et al., 2017; Ziegler et al., 2019] that seems untethered from the concerns of who holds the power to choose the values that models should be aligned with [Mitra et al., 2024]. Lastly, in the EU, the work on “open search” has begun [CERN, 2025] in response to the increasing recognition of US Big Tech’s oligopolistic control over search platforms. However, these initiatives—motivated through the nationalist framing of “digital sovereignty” [FitzGerald, 2025; Djuric, 2025]—point to saliently different notions of openness than what is required to meet the needs of the global south and marginalized communities [Neophytou and Mitra, 2026].

To be clear, it is not my intention here to claim that all work in fairness, explainability, and related topics strictly adhere to the liberal frame. I believe that the reality is much more complicated and much of this work has roots in critical scholarship. But nonetheless it would be accurate to say that this “first wave” of IR & Society research, if we may call it so, has predominantly operationalized liberal framings and engaged inadequately with critical theories and practices.

I also want to be clear that my strong critique of the liberal frame does not imply a wholesale dismissal on my part of the research that has been produced via this frame and its *theoretical* impact. Quite the contrary, the seeds that were sown at SWIRL 2018 [Culpepper et al., 2018] and the FACTS-IR workshop [Olteanu et al., 2021] brought together a community of societally-motivated researchers and sparked a body of work that serves as the foundation on which we stand today discussing concerns at the intersection of IR & society. My own transition from neural IR research to IR & Society research was against this backdrop and

⁵https://en.wikipedia.org/wiki/Jevons_paradox

I personally hold immense gratitude towards these scholars and their scholarship that foregrounded societal concerns in IR research. However, I also steadfastly believe that the liberal frame is saliently incomplete because it is not grounded in critical analysis, structural understanding of systems of oppression, and theories of power. The time is now ripe to transition to a critical frame if we want to see our work translate to practical societal impact.

2.2 The critical frame

After 2018, the next SWIRL workshop was organized in 2025, where another focus group convened to discuss IR’s responsibility to society. I had the enormous privilege to be invited to SWIRL’25 and the opportunity to participate in this particular focus group. A central theme of the discussion centered on explicating the values that guide IR research and centering concerns of social justice, emancipation, and democracy in IR [Trippas et al., 2025]. Related conversations were already ongoing within the IR community, including on the topics of normative design and evaluation of IR systems [Vrijenhoek et al., 2023] and reasserting IR’s interdisciplinary roots [Zangerle et al., 2025]. At that time, I had a couple of manuscripts in the preprint stage that also informed some of the conversations; one on explicating sociotechnical imaginaries in IR [Mitra, 2025b] and another conceptualizing an emancipatory IR framework [Mitra, 2025a], both of which were later published in the Information Retrieval Journal⁶ (IRRJ) the same year. Other recent work that bridges IR with political science and democratic theory—*e.g.*, Vrijenhoek et al. [2021]—were also discussed. Those crucial conversations were the spark that prompted Dana McKay, Michael D. Ekstrand, and myself to join with Sanne Vrijenhoek and Maria Murray to organize the workshop titled *Justice, Emancipation, Democracy, and Information Access (JEDI)*⁷ [Mitra et al., 2026a] at the ACM SIGIR 2026 conference in NAARM / Melbourne, Australia.

In hindsight, the emerging consensus among a subset of the IR community that the field should acknowledge the saliently political nature of our work and re-theorize our goal for “societal good” went hand in hand with making the case that IR research needs critical theories and practices. If SWIRL 2018 can be said to have kicked off the “first wave” of IR & Society research within the liberal frame, then SWIRL 2025 may turn out to be the catalyst for the “second wave” that grounds itself in the critical frame.

The critical frame distinguishes itself from the liberal frame in several key aspects, including:

1. **Structural analysis:** It grounds itself in critical theories and structural critiques of systems of oppression.
2. **Nondomination:** It adopts a normative position of nondomination and solidarity with the oppressed.
3. **Social construction of technology:** It rejects technological determinism and rationality in favor of reimagining and co-constructing anti-oppressive sociotechnical futures.
4. **Social construction of knowledge:** It renounces positivist views of information and the traditional role ascribed to IR within capitalist logics of consumption in favor of reframing IR’s role within social processes of sense-making and knowledge production.
5. **Praxis:** It affirms emancipatory praxis and movement building as research practices.

⁶<https://irrj.org/>

⁷<https://jedi.inertial.science/sigir2026/>

Next, I discuss each of these aspects in more details.

2.2.1 Structural analysis

Critical theory is a western tradition of thought with roots in the Frankfurt School in the early 20th century [Fuchs, 2021]. Early critical theorists include Herbert Marcuse, Max Horkheimer, and Theodor W. Adorno. Critical theory analyzes, with an intent to challenge, the systemic power relations in society. It posits that knowledge and social structures are saliently shaped by power relations between dominant and oppressed groups, and emphasizes the need to examine intersecting systems of oppression through historical and structural lenses. Critical information theory [Fuchs, 2009] further combines critical theory with information studies, and situates the study of information in context of relational power asymmetries, domination, and oppression. Pyati [2006] argues that Herbert Marcuse’s critique of ‘technological rationality’ as a tool of domination is particularly relevant to a critical theoretic analysis of modern information technologies.

Critical IR must be founded on a theoretical basis for situating and studying IR in the context of power relations and systems of oppression. As the name suggests, critical IR should be grounded in critical theory, but here I use ‘critical theory’ a bit more loosely than alluding to just the Frankfurt School of scholarship. I am also including other liberatory epistemologies—such as feminist and decolonial theories—as well as other forms of knowledge from activism, political organizing, and movement building. I am including the many strands of justice-oriented scholarships in information science and in computing—such as from science and technology studies (STS), Fairness, Accountability, and Transparency (FAccT), and human-computer interaction (HCI)—that engage with liberatory epistemologies, including humanistic [Bardzell and Bardzell, 2016], feminist [Hannigan and Crew, 1993; Wajcman, 2004; Bardzell, 2010; D’ignazio and Klein, 2020], queer [Light, 2011; Kitzie et al., 2022], postcolonial and decolonial [Irani et al., 2010; Ali, 2016; Mohamed et al., 2020; Jimenez et al., 2023], anti-racist [Ogbonnaya-Ogburu et al., 2020], anti-casteist [Vaghela et al., 2022], anti-ableist [Williams et al., 2021], anti-fascist [McQuillan, 2022], post-capitalist [Feltwell et al., 2018], and anarchist [Keyes et al., 2019] thought. The common thread shared between these diverse bodies of scholarship is that they recognize systems of oppression encoded within our social, political, economic, cultural, and technological realms, and how they co-constitute broader matrix of domination [Collins, 2000].

The challenge for the IR community is to develop a critical IR research agenda that engages with this rich body of scholarship; and to develop our own theoretical foundation and reflexive muscles equipped with structural analysis to visibilize how IR interacts and intersects with structural oppression. For IR to take this leap, we must both reach back to reconnect with our roots in library and information sciences as well as reach beyond our disciplinary boundaries to engage with critical scholarship from both within and outside of computing.

2.2.2 Nondomination

A key distinctive feature of critical theory is its normative commitment to nondomination in explicit solidarity with the oppressed. Horkheimer [1972] echoes this ethos by clarifying that the aim of critical theory is not simply to produce knowledge, but to progress the cause

of emancipation. Fuchs [2009] asserts that critical theory “*struggles for a classless, nondominative, co-operative, participatory democracy*”. Consequently, with respect to information studies, Fuchs [2009] postulate:

“Critical information theory therefore must study not just the role of information and information concepts in society, academia, nature, culture, etc, but how it is related to processes of oppression, exploitation, and domination, which implies a normative judgment in solidarity with the dominated and for the abolishment of domination.”

It is this commitment to nondomination and participatory democracy that I suggest should be the tenet around which we theorize the notion of “societal good” within IR. System behaviors that are desired by the liberal frame—such as fairness and transparency—must then be appraised with respect to whether they help us to sufficiently challenge systems of oppression and move us towards universal emancipation and participatory democracy. This goal broadens our purview of concern beyond algorithmic mitigations and asks of us to situate our work in broader sociopolitics and power relationships of our world. It both clarifies the societal change we aspire to affect through our work as well as calls for the development of relevant theories of change for how we envision our work to achieve its desired outcomes.

2.2.3 Social construction of technology

Technological determinism is a reductionist theory that asserts that technology develops following its own internal logic of efficiency and acts as an exogenous force on society and culture causing them to adapt [Héder, 2021]. We can observe such a worldview dominant within computing today that presupposes the inevitability of a future shaped by “AI” technologies. This technological deterministic thinking is also reflected by the liberal frame that is reluctant to challenge industry narratives and holds the capitalist logic of market-based decision-making as both legitimate and outside the realm of critique; therefore relegating itself strictly to the role of trying to moderate or delay the harms and inequities caused by these so-claimed “inevitable” technologies.

The critical frame rejects technological determinism in favor of recognizing that technology and society mutually shape each other. It encourages us to challenge technological and instrumental rationality and reclaim our agency to intentionally choose our desired sociotechnical futures. These visions of desirable futures constitute our sociotechnical imaginaries [Jasanoff and Kim, 2015] and are unsurprisingly shaped by power relations between social groups and capitalist logic of efficiency. The IR community must reflect on whose sociotechnical imaginaries are granted normative status in our field and intentionally reimagine a liberatory role of information access technologies in realizing nondominative futures. In doing so, we must also go beyond trying to make oppressive technologies more transparent or fair, and instead redirect our energy to radically redesigning our sociotechnical information systems in service of society’s liberatory aspirations.

2.2.4 Social construction of knowledge

The critical frame challenges positivist views on information and the role ascribed to IR within a capitalist model of information consumption. In IR, it is typical to assume the existence of a corpus of information artifacts (*e.g.*, documents or items) and a set of information

needs of users (represented as queries or their current context) that exist independently of the IR system. The prescribed goal of the IR system is to match the information artifacts (“supply”) with the information needs (“demand”) based on some notion of relevance that the system should estimate automatically. Relevance here is generally defined in terms of topical match—*i.e.*, does the information artifact satisfy the information need—but may also consider other attributes such as authority and reputation of the information source and factuality of its claims. IR tasks, in this context, are often represented as pre-curated datasets with a corpus, a set of information needs, and static relevance judgments that guide the development, training, and evaluation of models for relevance estimation.

The social constructionist view of information refutes this perspective and instead sees the corpus, the information needs, the IR system, and society in complex mutually-shaping relations with each other. Today’s IR systems are no longer simple keyword-based lookup tools; they actively model users and information, and shape consumer behavior, political discourse, and culture [Grimmelmann, 2008; Gillespie, 2019; Hallinan and Striphas, 2016]. They are integral to our social processes of sense-making and knowledge production across communities and populations. The business models of commercial IR platforms and the sociopolitical values of those who control these platforms wield enormous influence over global information production, not just consumption. In light of this, the social constructionist perspective challenges the reductionist view of IR as divorced from historical and current sociopolitical context, and instead encourages us to see IR systems as sociotechnical components embedded within dynamic knowledge ecosystems. This motivates very different views of IR tasks, those that cannot be captured by static datasets in isolation. Instead, it requires modeling complex information processes in which information is appraised not just based on their relevance or factuality but in the context of power relations within society. The role of IR then becomes to actively affect certain process outcomes, including growing our collective understanding of our physical and social environments—and our relations with them—and realizing nondominative futures.

2.2.5 Praxis

Another distinctive feature of the critical IR frame is that it challenges the boundaries of what we traditionally consider as research activities within the field. Our desired social transformations cannot be realized by theorization and research being conducted in isolation. Instead, it calls for recognizing praxis (liberatory research + liberatory action) as legitimate and necessary activity within our work. Examples of such liberatory actions includes:

1. **Building consensus on red lines in our research**, *e.g.*, prohibiting IR research on generative ads and other related technologies whose core objective is to manipulate public opinion at scale.
2. **Building counter-structures**, *e.g.*, building public-service IR platforms as alternatives to for-profit proprietary systems, exploring worker-owned cooperatives as alternatives to for-profit corporations, and building alternatives⁸ to the for-profit academic publishing ecosystem.
3. **Collective organizing and movement building**, *e.g.*, building collectives within the research community, organizing issue-based protests, supporting labor unioniza-

⁸*E.g.*, <https://irrj.org/about>

tion within academic and industry institutions, and participating in broader political movements and collective actions.

4. **Sociopolitical consciousness raising within the research community**, *e.g.*, creating spaces where IR community members can engage in critical dialog as a form of liberatory pedagogy [Freire, 1970], initiating book clubs and reading groups to engage with critical scholarship and build shared understanding, and leveraging different artistic medium like zines to educate and mobilize the community.
5. **Public education**, *e.g.*, creating materials and resources that inform and mobilize civil society against oppressive actors who instrumentalize new technologies to entrench their wealth and power and reinforce historical systems of oppression.
6. **Diversity, Equity, and Inclusion**, *e.g.*, ensuring that our communities are welcoming and safe spaces for researchers and practitioners from historically underrepresented and marginalized groups and their radical liberatory politics, and actively countering systemic discrimination and marginalization in our community spaces.
7. **Community care**, *e.g.*, creating community infrastructure to ensure psychological and financial safety of community members in the face of likely reprisals from those whose power is threatened by our work, and standing in solidarity with our community members who themselves (or their loved ones) are impacted by or are at risk from ongoing human rights violations, immigration violence, displacement, apartheid, war, and genocide.

3 Concluding thoughts: From *here* to *there*

Critical IR presents a high-potential direction for theoretical and applied IR research. The serious risks to our information ecosystems in the emerging era of technofascism [Coeckelbergh, 2026] makes investing in this research critically urgent. The call to center societal, democratic, and emancipatory values in IR research at SWIRL 2025 [Trippas et al., 2025] was in recognition of the several ongoing strands of IR research that are already working to incorporate theories from social and political sciences as well as developing our own critical theoretical foundation for the field. This includes structural critiques of our current approaches to gender-aware information access [Pinney et al., 2023], critical analysis of the implications of generative AI for IR [Mitra et al., 2024], and questions of how IR can support different forms of democratic governance [Vrijenhoek et al., 2021].

In my own work [Mitra, 2025b], I have argued the need for IR to reject technological determinism and critically examine the sociotechnical imaginaries that shape our research. I have also advocated for adopting nondomination as a central tenet in our field, offering the following operative definition for emancipatory IR [Mitra, 2025a] towards that goal:

“Emancipatory IR is the study and development of information access methods that challenge all forms of human oppression and situates its activities within broader collective emancipatory praxis.”

I theorized about emancipatory IR intentionally to encourage my IR colleagues to center the concerns of universal humanization and liberation from structural forms of oppression—*e.g.*, colonialism, racism, patriarchy, casteism, transphobia, religious persecution, and ableism—in our work; and to discourage the non-performative academic gaze. But these concerns

should not be simply the focus of a subset of the IR community. At ECIR 2026 IR-for-Good track [Mitra and Heuss, 2025], we posed the following rhetorical question:

“Does a special track on IR-for-Good imply that other contributions to the conference are IR-for-Bad?”.

This is an intentional prompt to the IR community to reaffirm that *all* IR research should aspire to realize societal good, and any demands for scientific inquiry that stands contradictory to the wellbeing of society should be vigorously rejected as Belkin and Robertson [1976] recommended fifty years ago. To reemphasize, not only should we inject the critical lens in ongoing research on fairness, explainability, and related topics, we should also subject IR research that do not generally make explicit societal arguments—*e.g.*, work on indexing data structures and AI aided information access—to the same critical demands. Even sub-topics of IR like data structures that may superficially seem detached from societal concerns, may in fact conceal underlying assumptions about centralized platform control and ownership. My argument here mirrors the perspective of Fuchs [2009] that our ultimate goal should not be to realize the critical frame as one alternative that coexists with a plurality of other frames, but to realize a unified critical approach to information studies.

To shift the IR community towards critical theories and practices will require pulling together the current strands of ongoing work into a coherent research area. But more importantly, it will also require raising the sociopolitical consciousness of the community that has historically taken comfort in being portrayed as apolitical. Any desired transformation to IR technologies must be preceded by transformation of the IR community itself [Mitra, 2025b]; and we cannot prescribe this transformation top-down—nor from theorists to practitioners—but instead we must build the scaffolding that others can follow for their own liberatory education and development of their conscientização [Freire, 1970]. We need community spaces that are explicitly politicized and amenable to open and critical conversations about the role of IR in supporting social justice, emancipation, and democracy—as is our goal with the SIGIR 2026 JEDI workshop [Mitra et al., 2026a].

Our work in support of this transition will also require developing new community practices. For example, ECIR 2026 IR-for-Good track [Mitra and Heuss, 2025] introduced a new requirement for papers to articulate their theories of change and underlying assumptions that must hold true for their desired goals to be realized in practice. The motivation was that making our theories of change explicit in our scholarship allows for their necessary critique and refinement over time. Other relevant practices may include requiring papers to provide positionality statements and elaborate on their ethical considerations. IR would benefit from looking at adjacent communities—such as IS, FAccT, and HCI—for lessons learned in those fields and actively experiment with new practices to continue evolving our research norms.

While throughout this article, I have called on the IR community to center justice and emancipation in our work, I must forewarn that as a community we must reject any false notions of technosolutionism and vigorously challenge the technologists-as-liberator framing [Mitra et al., 2026b]. Instead, we must do the challenging and radical task of reimagining all IR technologies as sociotechnical systems and encode spaces for community participation, collective deliberation, and democratic negotiation within our platforms, and intentionally allow our platforms to be co-opted by the oppressed in service of their liberatory struggles.

My call for incorporating critical theories and practices in IR is also a call for more rigor in our scholarly scientific activities. Rigorous IR scholarship cannot emerge from a reactionary

understanding of our world and the role of information in it. Rigor in IR, similar to in adjacent disciplines like AI [Olteanu et al., 2025], must be understood as not just methodological rigor but also epistemic, normative, conceptual, reporting, and interpretative rigor. In summary, the goal of critical IR research is to employ rigorous science in service of realizing Fuchs’s “classless, nondominative, co-operative, participatory democracy” [Fuchs, 2009].

Positionality statement

I am a middle-class Southasian Bengali immigrant and a cisgender man, who grew up in India, got a doctorate in Computer Science in the UK, and is currently working as an independent researcher in Canada. My lived experiences in India and as an immigrant worker in the USA, the UK, and Canada continue to shape my sociopolitical values and my research. I worked as an industry researcher in Big Tech for nearly two decades but last year I decided to leave in objection to their harmful impact on society.⁹ I participate in social justice movement organizing which has deepened my understanding of anti-oppression struggles and traditions of community resistance and solidarity. I see positionality as an evolving process. I continue to reflect on how my identity and lived experiences shape my perspectives, values, and research.

Territorial acknowledgment

I acknowledge that I authored this article while located in Tiohtià:ke / Mooniyang situated on unceded Indigenous lands and traditional territory of the Kanien’kehá:ka and Anishnaabe Nations.

Generative AI usage statement

I hereby declare that no generative AI technologies were used in the writing of this paper.

Acknowledgments

I gratefully acknowledge feedback from Dana McKay, Michael D. Ekstrand, and Nicola Neophytou on earlier drafts of this article.

References

Syed Mustafa Ali. A brief introduction to decolonial computing. *XRDS: Crossroads, The ACM Magazine for Students*, 22(4):16–21, 2016.

⁹<https://disjunctionsmag.com/articles/why-leaving-big-tech/>

-
- Avishek Anand, Lijun Lyu, Maximilian Idahl, Yumeng Wang, Jonas Wallat, and Zijian Zhang. Explainable information retrieval: A survey. *arXiv preprint arXiv:2211.02405*, 2022.
- Jeffrey Bardzell and Shaowen Bardzell. Humanistic hci. *Interactions*, 23(2):20–29, 2016.
- 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.
- 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.
- NJ Belkin and SE Robertson. Some ethical and political implications of theoretical research in information science. In *Proceedings of the ASIS Annual Meeting*, 1976.
- Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *Proc. SIGIR*, pages 405–414, New York, NY, USA, 2018. ACM.
- CERN. European project to make web search more open and ethical, 2025. URL <https://home.cern/news/news/computing/european-project-make-web-search-more-open-and-ethical>.
- 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.
- Mark Coeckelbergh. Technofascism: Ai, big tech, and the new authoritarianism. *AI & SOCIETY*, pages 1–14, 2026.
- Patricia Hill Collins. *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*. routledge, 2000.
- 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.
- 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.
- Mickey Djuric. Canada’s digital sovereignty push challenges america’s reach. *Politico Pro*, 2025.
- Michael D Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. Fairness in information access systems. *Foundations and Trends® in Information Retrieval*, 16(1-2):1–177, 2022.

-
- 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 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–8, 2018.
- Michael FitzGerald. ‘digital sovereignty’: why the eu may be shifting from internet regulation to building homegrown tech. *The Conversation*, 2025.
- Paulo Freire. *Pedagogy of the Oppressed*. Herder and Herder New York, 1970.
- Christian Fuchs. Towards a critical theory of information. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society*, 7(2): 243–292, 2009.
- Christian Fuchs. What is critical theory? In *Foundations of critical theory*, pages 17–51. Routledge, 2021.
- Tarleton Gillespie. Algorithmically recognizable: Santorum’s google problem, and google’s santorum problem. In *The Social Power of Algorithms*, pages 63–80. Routledge, 2019.
- James Grimmelman. The google dilemma. *NYL Sch. L. Rev.*, 53:939, 2008.
- 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.
- Alex Hanna, Remi Denton, Andrew Smart, and Jamila Smith-Loud. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 501–512, 2020.
- Jane Anne Hannigan and Hilary Crew. A feminist paradigm for library and information science. *Wilson Library Bulletin*, 68(2):28–32, 1993.
- Mihály Héder. Ai and the resurrection of technological determinism. *Információs Társadalom: Társadalomtudományi Folyóirat*, 21(2):119–130, 2021.
- Tomasz Hollanek. Ai transparency: a matter of reconciling design with critique. *Ai & Society*, 38(5):2071–2079, 2023.
- Max Horkheimer. Traditional and critical theory. *Critical theory: Selected essays*, 188(243): 1–11, 1972.
- 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.
- Sheila Jasanoff and Sang-Hyun Kim. *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*. University of Chicago Press, 2015.
- Andrea Jimenez, Sara Vannini, and Andrew Cox. A holistic decolonial lens for library and information studies. *Journal of Documentation*, 79(1):224–244, 2023.

-
- 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.
- Vanessa L Kitzie, Travis L Wagner, Valerie Lookingbill, and Nicolas Vera. Advancing information practices theoretical discourses centered on marginality, community, and embodiment: Learning from the experiences of lesbian, gay, bisexual, transgender, queer, intersex, and asexual (lgbtqia+) communities. *Journal of the Association for Information Science and Technology*, 73(4):494–510, 2022.
- Joseph Lehman. A brief explanation of the overton window. *Mackinac Center for Public Policy*, 2014.
- 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.
- Ann Light. Hci as heterodoxy: Technologies of identity and the queering of interaction with computers. *Interacting with computers*, 23(5):430–438, 2011.
- Dan McQuillan. Anti-fascist ai. In *Resisting AI*, pages 135–148. Bristol University Press, 2022.
- Bhaskar Mitra. Emancipatory information retrieval. *Information Retrieval Research*, 1(2): 313–339, 2025a.
- Bhaskar Mitra. Search and society: Reimagining information access for radical futures. *Information Retrieval Research*, 1(1):47–92, 2025b.
- Bhaskar Mitra. Why i’m leaving big tech: A tech worker’s reflections. *Disjunctions*, (1), 2026.
- Bhaskar Mitra and Maria Heuss. What is ir-for-good? <https://bhaskar-mitra.github.io/posts/2025/09/01/what-is-ir-for-good/>, 2025.
- Bhaskar Mitra, Henriette Cramer, and Olya Gurevich. Sociotechnical implications of generative artificial intelligence for information access. In *Information Access in the Era of Generative AI*, pages 161–200. Springer, 2024.
- Bhaskar Mitra, Dana Mckay, Michael D. Ekstrand, Sanne Vrijenhoek, and Maria Murray. Justice, emancipation, democracy, and information access (jedi): The sigir workshop on resisting corporate and authoritarian capture of information access platforms. In *sigir*, 2026a.
- Bhaskar Mitra, Nicola Neophytou, and Sireesh Gururaja. Information access of the oppressed: Freirean design for emancipatory information access. *arXiv preprint arXiv:2601.09600*, 2026b.
- Shakir Mohamed, Marie-Therese Png, and William Isaac. Decolonial ai: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33:659–684, 2020.

-
- Nicola Neophytou and Bhaskar Mitra. Open, to what end? a capability-theoretic perspective on open search. *arXiv preprint arXiv:2603.14584*, 2026.
- Safiya Umoja Noble. Algorithms of oppression. In *Algorithms of oppression*. New York university press, 2018.
- Ihudiya Finda Ogbonnaya-Ogburu, Angela DR Smith, Alexandra To, and Kentaro Toyama. Critical race theory for hci. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pages 1–16, 2020.
- 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, 2021.
- Alexandra Olteanu, Su Lin Blodgett, Agathe Balayn, Angelina Wang, Fernando Diaz, Flavio du Pin Calmon, Margaret Mitchell, Michael Ekstrand, Reuben Binns, and Solon Barocas. Rigor in ai: Doing rigorous ai work requires a broader, responsible ai-informed conception of rigor. *arXiv preprint arXiv:2506.14652*, 2025.
- Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. Fair ranking: a critical review, challenges, and future directions. In *Proc. FAccT*, pages 1929–1942, 2022.
- 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.
- Ajit K Pyati. Critical theory and information studies: A marcusean infusion. *Policy Futures in Education*, 4(1):83–89, 2006.
- Harrison Scells, Shengyao Zhuang, and Guido Zuccon. Reduce, reuse, recycle: Green information retrieval research. In *sigir*, pages 2825–2837, 2022.
- Johanne R Trippas, J Shane Culpepper, et al. Report from the fourth strategic workshop on information retrieval in lorne (swirl 2025). In *ACM SIGIR Forum*, volume 59, page 68, 2025.
- 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, 2022.
- Sanne Vrijenhoek, Mesut Kaya, Nadia Metoui, Judith Möller, Daan Odijk, and Natali Helberger. Recommenders with a mission: assessing diversity in news recommendations. In *chiir*, pages 173–183, 2021.
- Sanne Vrijenhoek, Lien Michiels, Johannes Kruse, Alain Starke, Jordi Viader Guerrero, and Nava Tintarev. Report on normalize: The first workshop on the normative design and

-
- evaluation of recommender systems. In *CEUR Workshop Proceedings*, volume 3639. CEUR-WS, 2023.
- Judy Wajcman. *Technofeminism*. cambridge: Polity. 2004.
- 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.
- Rua M Williams, Kathryn Ringland, Amelia Gibson, Mahender Mandala, Arne Maibaum, and Tiago Guerreiro. Articulations toward a crisp hci. *Interactions*, 28(3):28–37, 2021.
- Tim Wu. *The attention merchants: The epic scramble to get inside our heads*. Vintage, 2017.
- Eva Zangerle, Alan Said, and Christine Bauer. Beyond algorithms: Reclaiming the interdisciplinary roots of recommender systems (beyond 2025). In *Proc. RecSys*, pages 1360–1361, 2025.
- Meike Zehlike, Ke Yang, and Julia Stoyanovich. Fairness in ranking: A survey. *arXiv preprint arXiv:2103.14000*, 2021.
- Yongfeng Zhang, Xu Chen, et al. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval*, 14(1):1–101, 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.
- Justin Zobel. What we talk about when we talk about information retrieval. In *ACM SIGIR Forum*, volume 51, pages 18–26. ACM New York, NY, USA, 2018.
- Shoshana Zuboff. The age of surveillance capitalism. In *Social Theory Re-Wired*, pages 203–213. Routledge, 2023.