

Information Access of the Oppressed: A Problem-Posing Framework for Envisioning Emancipatory Information Access Platforms 🙌🙌🙌

BHASKAR MITRA, Independent Researcher, Canada

NICOLA NEOPHYTOU, Independent Researcher, Canada

SIREESH GURURAJA, Carnegie Mellon University, USA

Online information access (IA) platforms, like other forms of media in our history, are targets of authoritarian capture. These concerns are particularly serious and urgent today in light of the rising levels of democratic erosion worldwide, the emerging capabilities of generative AI technologies such as *AI persuasion*, and the increasing concentration of economic and political power in the hands of Big Tech. This raises the question of what alternative IA infrastructure we must reimagine and build to mitigate the risks of authoritarian capture of our information ecosystems. We explore this question through the lens of Paulo Freire’s theories of emancipatory pedagogy. Freire’s theories provide a radically different lens for exploring IA’s sociotechnical concerns relative to the current dominating frames of fairness, accountability, confidentiality, transparency, and safety. We make explicit, with the intention to challenge, the dichotomy of how we relate to technology as either technologists—who envision and build technology—and its users. We posit that this mirrors the teacher-student relationship in Freire’s analysis. By extending Freire’s analysis to IA, we challenge the notion that it is the burden of the (altruistic) technologists to come up with interventions to mitigate the risks that emerging technologies pose to marginalized communities. Instead, we advocate that the first task for the technologists is to pose these as problems to the marginalized communities, to encourage them to make and unmake the technology as part of their material struggle against oppression. Their second task is to redesign our online technology stacks to structurally expose spaces for community members to co-opt and co-construct the technology in aid of their emancipatory struggles. We operationalize Freire’s theories to develop a problem-posing framework for envisioning emancipatory IA platforms of the future, and call on technologists to rid ourselves of our complicity in the global rise of authoritarianism and intentionally engage in building the infrastructure for resistance.

CCS Concepts: • **Information systems** → **Information retrieval**; **World Wide Web**; • **Human-centered computing**;

Additional Key Words and Phrases: Emancipatory Information Access, Search and Society, Sociotechnical Information Systems

ACM Reference Format:

Bhaskar Mitra, Nicola Neophytou, and Sireesh Gururaja. 2026. Information Access of the Oppressed: A Problem-Posing Framework for Envisioning Emancipatory Information Access Platforms 🙌🙌🙌. In *Proceedings of Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency (FAcT ’26)*. ACM, New York, NY, USA, 18 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Motivation

On July 23, 2025, U.S. president Donald J. Trump issued an executive order titled “Preventing Woke AI in the Federal Government” [88] prohibiting federal procurement of AI models that incorporate concepts like “critical race theory,

Authors’ Contact Information: [Bhaskar Mitra](#), bhaskar.mitra@acm.org, Independent Researcher, Tiohtià:ke / Montréal, Canada; Nicola Neophytou, neophytounicola@gmail.com, Independent Researcher, Tiohtià:ke / Montréal, Canada; [Sireesh Gururaja](#), sgururaj@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

transgenderism, unconscious bias, intersectionality, and systemic racism; and discrimination on the basis of race or sex”. The executive order, while claiming to safeguard AI models from purported “ideological biases”, seeks to strong-arm AI companies to conform with the administration’s ideological agenda [72]. Likewise, China’s National Technical Committee 260 on Cybersecurity of SAC published a technical document¹ mandating all Chinese generative AI models to not generate content that violates the country’s “core socialist values”. Consequently, the Chinese AI chatbot DeepSeek refuses to answer certain political questions, like on the topics of Tiananmen Square and Taiwanese independence [61]. Such acts of ideological imposition on AI model outputs is also not limited only to state actors. In May 2025, xAI’s chatbot Grok promoted conspiracy theories about “white genocide” in South Africa [55] after said conspiracy theories were mainstreamed by xAI’s owner Elon Musk [46]. On another occasion the same year, Grok responded to users with antisemitic and pro-Nazi content [86], which followed only a few months after Musk himself was embroiled in a controversy for making hand gestures at Trump’s presidential inauguration event that appeared to be Nazi salutes [47].

As AI becomes a widely used technology for information access (IA) [13], these instances of attempted capture can be seen to follow historical examples of regimes trying to capture and manipulate different forms of media [24], and fit into a growing trend of *digital authoritarianism* [78]—defined as “the use of digital information technology by authoritarian regimes to surveil, repress, and manipulate domestic and foreign populations” [76]. Traditional media, social media, and IA platforms have historically been sites of conflict between oppression and emancipation. However, these concerns are particularly serious today in the context of: (i) Rising levels of democratic erosion [38] and geopolitical conflicts [85, 90], (ii) the increasing concentration of economic and political power in the hands of Big Tech [30, 51, 74, 96] which increasingly places itself in alignment with militarized and authoritarian actors [3, 28, 52–54, 57, 60], and (iii) emerging capabilities of generative AI technologies, such as *AI persuasion* [9, 10, 27, 75], and the potential for its weaponization within IA platforms for mass manipulation [69] and generative propaganda [18]. The confluence of these factors calls for increased critical scrutiny from IA researchers on the evolving relationship between users and IA platforms [67], to quote: “Imagine every time you searched online or accessed information via your digital assistant, the information was presented to you exactly in the form most likely to alter your consumer preferences or political opinions. Or, consider AI-generated digital characters in ads and videos that appropriate marginalized identities to say or act in ways that real members of that community may be strongly opposed to—a new form of ‘Digital Blackface’ [49].”

Even if we were to assume that IA companies take seriously their role as a pillar of civic society, the existence of such capabilities and infrastructure incentivizes authoritarian capture of platforms. There may be a temptation to search for algorithmic or policy interventions that can act as guardrails around AI technologies. But given the lack of transparency around how these models are trained by private corporations, identifying and enforcing such interventions become challenging. Also, many of the harms must be understood in the context of the IA platforms where these AI models are deployed and in context of their complex interactions with people searching for information. We must also not take it for granted that these AI technologies necessarily have a beneficial role in IA and must leave room for refusal of their deployment [7]. Finally, while AI capabilities may make IA platforms more attractive for capture, the risk does not go away simply by their exclusion from the platform. Therefore, we posit that the IA platforms are more meaningful sites for systemic mitigation. We ask: “What is the alternative IA infrastructure that we must reimagine and realize in practice as a form of mitigation for authoritarian capture of IA platforms?”.

¹<https://www.tc260.org.cn/upload/2024-03-01/1709282398070082466.pdf>

In the social media landscape, decentralization has emerged as a dominant strategy to safeguard against capture [15]. Decentralization may also hold the key to protecting IA platforms—e.g., search engines. However, while the decentralization of technological infrastructure may be necessary, it is not sufficient to ensure that platforms are emancipatory [44]. Our pursuit for alternative information infrastructure is not value-neutral. In fact, our motivation stems from the recognition of IA as sites of political struggles between oppressors and oppressed. We reject any ambivalence or false neutrality with respect to oppressor-oppressed relations and situate our work in recent calls to explicitly align IA research with humanistic and emancipatory goals [6, 67, 87]. To put it plainly, it is our emancipatory aspirations that prompt us to recognize the grave risks of state and corporate actors capturing our information platforms to manipulate mass opinions, and therefore it is necessary that we go beyond focusing on platform decentralization. This is in contrast to other frames of critique of information platform capture, like *digital sovereignty* [23, 29] whose critique isn't targeted at the practices of capture itself but focused on whose nationalistic interests are served by the platforms. Instead, in this work we develop a theoretical framework to envision IA platforms that foregrounds how oppressive power relations subverts knowledge production and consensus building, uplifts marginalized and oppressed voices, seeks epistemic justice [35], and aims to realize humanistic and emancipatory futures for all globally.

Practices of futuring [36] and envisioning [77] shape technology development. Our collectively held sociotechnical imaginaries [48] do not only imagine but also co-construct our digital futures [62]. In IA, these practices commonly manifest as scholarly publications (e.g., [65, 81]) and strategic workshops (e.g., [6, 41, 87]). However, the dominant underlying values and epistemologies that shape these practices and their outcomes in the field often remain implicit [68, 94]. Envisioning emancipatory IA poses two distinct challenges in this context. The first more obvious requirement is that the envisioning process must produce figurative blueprint for IA platforms that safeguard against authoritarian capture and contribute towards emancipatory societal outcomes. A second less obvious but nonetheless critical requirement is that the envisioning process must involve reflective participation of marginalized communities who should get to shape the values of the platform. This second requirement is motivated by insights drawn from critical scholarship on emancipatory pedagogy, in particular Paulo Freire's canonical work "*Pedagogy of the Oppressed*" [34]. In the context of emancipatory pedagogy, Freire emphasizes, that the oppressed are not simply targets but rather the executors of liberatory actions. He critiques what he refers to as the "*banking concept of education*" in which emancipatory knowledge is "gifted" by those committing the act of liberating to those who are purportedly being liberated. He writes:

"Attempting to liberate the oppressed without their reflective participation in the act of liberation is to treat them as objects which must be saved from a burning building; it is to lead them into the populist pitfall and transform them into masses which can be manipulated. ... The conviction of the oppressed that they must fight for their liberation is not a gift bestowed by the revolutionary leadership, but the result of their own conscientização."

When we embark to envision emancipatory IA platforms, we must vigorously reject the notion that it is the burden of a few individuals or institutions to determine the design that makes platforms emancipatory. Fundamentally, we must reject *technosolutionism* and the idea that there are some key elusive design tricks and system features that are waiting to be discovered that once incorporated renders IA platforms universally emancipatory and safe from capture in all diverse cultural and sociopolitical context. Rather the basis of emancipatory IA lies in holding space for reflective participation of community members in the process of envisioning, design, implementation, maintenance, and governance of platforms. This concept is central to the framework we propose in this work. The framework adopts Freire's problem-posing approach by trying to articulate fundamental questions that emancipatory IA platforms must address. The framework

intentionally stops there to make space for diverse reflection, dialogue, participation, collaboration, and negotiation by community members in the rest of the envisioning and actualization of emancipatory IA that we hope will follow our current work. Envisioned IA platforms must further create new opportunities for ongoing community participation in extending, maintaining, and governing platforms. Community participation is also necessary to critique and revise the problems posed by our framework. By foregrounding community participation this way, we want to move away from framings of IA where people relate to platforms as either technologists (*those who envision, build, and maintain platforms*) or users (*those who use platforms*); in favor of framings where we all see ourselves engaged in epistemic activities of knowledge production and consensus building through the making and unmaking of IA technologies.

The rest of this article is organized as follows: We next discuss some background concepts relevant to our framework in Section 2. As recent calls [67, 87] for aligning IA research with the goals of universal humanization and emancipation serves as the starting point of our work, we provide a brief introduction to emancipatory IA (§2.1) and then discuss other initiatives in the fields of IA and social media that are relevant to our work (§2.2). Next, we situate our work in Freire’s theories and analyses of emancipatory pedagogy in Section 3 to develop a critical perspective on emancipatory IA. In Section 4, we present our problem-posing framework for envisioning Emancipatory IA. In Section 5, we speculate about some of the inherent tensions involved in this work before finally concluding in Section 6.

2 Background

2.1 Emancipatory IA

Emancipatory IA [67] is defined² as “*as the study and development of information access methods that challenge all forms of human oppression, and situates its activities within broader collective emancipatory praxis*”. The original formulation clarifies that emancipatory IA aspires for universal humanization and the elimination of all structural forms of oppression including colonialism, racism, patriarchy, casteism, transphobia, religious persecution, and ableism. This framing rejects the premise of *technological determinism* in favor of recognizing the agency and responsibilities of researchers and developers to affect positive societal transformations via IA technologies. It proposes a framework of practices, projects, and provocations for emancipatory IA. The practices—adapted from Wright’s emancipatory social science framework [98]—include: (i) *Diagnosing and critiquing* to identify the ways in which IA approaches and practices may contribute systemic harms and impede emancipatory struggles, (ii) *imagining viable alternative futures* informed by feminist, queer, decolonial, anti-racist, anti-casteist, anti-ableist, abolitionist, and other emancipatory epistemologies, and (iii) *elaborating theories of change* that makes explicit the path to realizing alternative desirable futures recognizing current sociopolitical realities. The framework encourages the field to proactively seek out provocations (e.g., through cross-disciplinary engagements and practices of reflexivity) to challenge current ways of thinking and reimagine radical alternatives. Finally, the framework proposes a nonexhaustive list of projects including: safeguarding against capture, surveillance, manipulation, dehumanization, and inequitable outcomes, and promoting emancipatory pedagogy.

Our current work intersects with the emancipatory IA framework in several ways. Our work is closely relevant to the practices of imagining viable alternatives. The problem-posing framework detailed in Section 4 aims to provoke new imaginaries of IA systems as well as identify concrete projects for actualizing said visions. While the emancipatory IA framework enumerates several largely-independent projects motivated by broad concerns of emancipation, our work is predicated on the assumption that we want to explore these questions in the context of a single IA platform that may be part of a broader ecosystem. This assumption not only prompts us to envision more concrete artifacts and

²While the original definition [67] was applied in relation to information *retrieval*, we extend the scope to information *access* as a whole.

future directions, but also stresses the importance of considering how these different concerns and projects relate and intersect with each other in complex and significant ways. In doing so, it attempts to provide a concrete starting point for the field of IA to begin putting the theoretical framework of emancipatory IA to actual practice.

2.2 Related Initiatives

Open Web Search. The Open Search Foundation³ is engaged in projects that are particularly relevant to our work. This includes the OpenWebSearch.EU project⁴ and the #FreeWebSearch Charter [31]. The OpenWebSearch.EU project started in response to the growing concerns of the concentration of power over IA platforms in the hands of a few large corporations bound by specific national interests. According to the project website: “*The project has its origins in concerns over the imbalance of the search engine market. Despite being a backbone of our digital economy, web search is dominated and limited by a few gatekeepers like Google, Microsoft, Baidu or Yandex. Thus, information as public good, with free, unbiased and transparent access is not under public control anymore. This imbalance endangers democracy and limits the innovative potential of Europe’s research landscape and economy.*” They furthermore clarify their intention to contribute towards “*Europe’s digital sovereignty as well as promoting an open human-centered search engine market*”.

Consequently, the #FreeWebSearch Charter enumerates ten fundamental principles and demands, namely: (i) Guarantee Transparency and Explainability, (ii) foster competition and independent infrastructure, (iii) strengthen privacy and enable search without tracking and profiling, (iv) give users, content creators and advertisers more control, (v) secure equal access and non-discrimination, (vi) support information plurality and multiple perspectives, (vii) take responsibility for environmental and social impact, (viii) safeguard integrity and protect against disinformation, (ix) build search competences and critical awareness, and (x) enforce democratic control and binding rules across borders.

These Open Search projects and charters significantly overlap with calls for emancipatory IA and presents critical opportunity for collaborations across proponents of both approaches. However, there are also critical divergences in motivations and aspirations between the two perspectives. The Open Search Foundation is focused on Europe’s digital sovereignty and is centered on European values, principles, legislation, and standards. While, it represents important initiatives to build counter-structures to challenge US Big Tech’s outsized dominance over global IA platforms, it does not speak for the values and needs of the global south that represents 85% of the world’s population⁵. We must recognize that *Eurocentrism* is not a meaningful antidote to *Amercan Exceptionalism*, and no IA platform is truly open nor reflects the internet in all its diversity if it systemically erases the voices of the global south. This model of challenging US Big Tech reflects the values of multipolarism in which influence over our information ecosystems is concentrated in the hands of few powerful state and supranational actors within the international order (e.g., USA, EU, China, and Russia); at the cost of further disenfranchising other often-marginalized global populations.

In contrast, emancipatory IA seeks epistemic justice for both the global south as well as for marginalized populations within powerful nation-states, and is saliently decolonial in its aspirations. These differences are not simply philosophical divergences, but have material consequences for platform design and research directions. For example, the absence of trustworthy centralized institution(s) to steer and sponsor the platform development and governance, exacerbates the need for decentralized technological and social infrastructure in emancipatory IA. Emancipatory IA also requires reflective participation of marginalized populations in platform development and governance in ways that an open search platform for EU does not. Emancipatory IA also espouses different approaches to addressing certain societal

³<https://opensearchfoundation.org/about-open-search-foundation/>

⁴<https://openwebsearch.eu/the-project/>

⁵The global south’s population estimate is based on Verón [92].

concerns compared to OpenWebSearch.EU and the #FreeWebSearch Charter. For example, the #FreeWebSearch Charter promotes transparency and diversity to tackle political bias and dominant narratives. However, that risks encoding false equivalence between oppressive and emancipatory politics which may contribute to *algorithmic bothsism* [68]. Emancipatory IA, in contrast, pushes for structural ways to counter systemic power differentials and uplift marginalized voices. That is, while the #FreeWebSearch vision promotes what may be considered a liberal model of democracy, emancipatory IA aligns with a more critical model of democracy [93].

However, these differences notwithstanding, we believe that both these approaches can benefit from one another. The digital sovereignty framing is likely to attract significantly more funding and investments that may enable projects like OpenWebSearch.EU to prototype some the technologies that may turn out to be critical for emancipatory IA. Emancipatory IA research may in turn unlock new innovations in the space of sociotechnical design of IA platforms that provides structural mechanisms to ensure compliance with the demands of #FreeWebSearch.

Decentralized social media. Adjacent to IA, in the social media landscape, there have been longstanding concerns about the degree that platforms’ unequal application of moderation rules, creating unfair outcomes for minoritized groups [26], and the degree to which platforms make decisions at the expense of their users [66]. These concerns expanded to include explicit ideological capture after Elon Musk’s takeover of Twitter in 2022 [2, 42, 58, 64]. In response, protocol-based decentralized social media like the ActivityPub-based Mastodon⁶ and ATProto-based Bluesky⁷, have emerged as alternatives that explicitly center resistance to elite and corporate capture, and individual and community agency in processes like feed curation and moderation. The novelty of the decentralized paradigm means that many issues are yet to be fully solved; in particular, two issues are particularly relevant to our discussion here: the degree to which nominally decentralized networks actually facilitate decentralized operation [100] and effectively avoid capture, and the tension between democratic control of sections of federated networks and the access to a (relatively) unified discourse on a centralized platform [33, 44]. However, decentralized social media remains under active development, and their protocol-based nature structurally encourages community development. The ATProto ecosystem, for example, now features independently developed interoperable services that focus on blogging⁸, scientific preprints⁹, photo sharing¹⁰, and even source control¹¹. Emancipatory IA should draw insights from these ongoing efforts and in turn contribute back new insights and sociotechnical innovations.

2.3 Participatory design

Central to our framework is the notion that communities and marginalized groups are not only essential in the design, building, and governance of emancipatory IA platforms, but that the technology cannot truly be emancipatory without their reflective participation. Existing models for collaboration with wider communities—such as participatory design (also referred to as co-design)—have been popularized particularly in the field of human-computer interaction (HCI). We make explicit how our framework diverges from existing notions of participatory design, and argue that these approaches maintain a notion of centralized power and responsibility, reinforcing the same structures that make technologies susceptible perpetual de-prioritization of marginalized communities, and to authoritarian capture.

⁶<https://www.w3.org/TR/2018/REC-activitypub-20180123/>

⁷<https://bsky.social/about/blog/10-18-2022-the-at-protocol>

⁸<https://leaflet.pub>, <https://pckt.blog/>

⁹<https://chive.pub/>

¹⁰<https://www.flashes.blue/>

¹¹<https://tangled.org/>

Participatory design refers to building technologies with the involvement of relevant stakeholders—often end-users—such that the final design outcome reflects and incorporates their opinions, desires, and expertise. Participatory design is therefore inherently visionary; it exists for researchers and participants to cooperatively envision and build idealized systems. In practice, participatory design has been leveraged in the development of initiatives across multiple fields including tech and computing [25, 79], social programs and community projects [5, 56], and urban and environmental design [8]. Several of these initiatives, such as the Decidim platform, to facilitate citizen participation in the allocation of public funds by Barcelona’s City Council [1], include the wider community in not just the initial construction but also the ongoing governance of the program.

Costanza-Chock [17] describes *design justice* as a conscious redistribution of the “benefits and burdens” in design, cognizant of community-led and Indigenous traditional practices, and an investigation into how design can reinforce or challenge the “matrix of domination (white supremacy, heteropatriarchy, capitalism, ableism, settler colonialism, and other forms of structural inequality)”. The author highlights several pitfalls of participatory design to explain the continual de-prioritization of marginalized voices in design. Common shortcomings include selected participants being unrepresentative of all potential users, or resembling the researchers themselves, the general lack of diversity in ‘default’ end users, and the economic pressures to appeal to the most profitable users. Additionally, naive implementations of participatory design—“Add Diverse Stakeholders and Stir” [20]—can lead to de-politicizing of well-intentioned participation, whereby a focus on the needs of individual stakeholders obscures the consideration of potential power structures or designing solutions for collective good [19]. Further, seemingly inclusive and democratic enactments for participation (in a neoliberal approach) can incidentally reinforce existing power structures [17]. The ways in which participants are selected and included to shape technology are therefore political, and contribute to an ongoing feedback-loop which progressively alienates marginalized groups through disproportionately worse usability and relevance. The consequences of worse utility can manifest—in the context of information access—as disproportionately higher censorship and filtering for marginalized groups [40, 89], among other forms of systematic exclusion.

While authoritarian capture implies the intentional “hijacking” of technology or media to manipulate mass opinion, the above limitations of participatory design discussed so far can lead to the unintentional skewing of social and political opinions, due to the structural suppression of presence and influence of marginalized groups. However, realized examples of participatory design still tend towards centralized power structures, since project owners maintain executive authority in deciding the parameters of participation, i.e., the selection and role of stakeholders in the design process [19], and ultimately have the power to block proposals made by stakeholders in favor of their own ideas [20]. We argue that this upholding of central power is a crucial vulnerability of participatory design, that, in the context of an emancipatory IA system, would render it susceptible to authoritarian capture.

The principles of participatory design as outlined by Spinuzzi [82] underscores a separation between the two groups of researchers and participants, with participants being elected to contribute by the researchers, given their relevance and potential insight to the problem. This two-camp distinction between researchers and participants is the most significant way in which our framework for emancipatory IA deviates from existing approaches of participatory design. While the two approaches are aligned in harnessing the ‘tacit’ perspectives of the wider community, our work does not distinguish between the designers and the users, nor does it promote a structure in which feedback is collected from one group and implemented by another. Rather, the emancipatory IA framework embeds a structure of collaboration which removes the primary onus for a particular group or individual, but ensures equal agency and influence of members of the community in building future IA technologies. Costanza-Chock [17] emphasizes the necessity to ground design

justice in Freire’s theory of *popular education* (or critical pedagogy). We expand on possible alternative collaborative structures, inspired by Freire’s work, in the next section.

3 Information Access of the Oppressed

“*Pedagogy of the Oppressed*” (original Portuguese title: “*Pedagogia do Oprimido*”) [34] authored by Brazilian Marxist educator Paulo Freire was first published in 1968. The book presents a detailed Marxist class analysis exploring oppressor-oppressed relations in the context of pedagogy. In his canonical work, Freire critiques the traditional “banking model of education” in which it is the task of the teacher to be the arbiter of what constitutes valid knowledge and to “deposit” their teachings, while it is the duty of the students to “*patiently receive, memorize, and repeat*”. Instead, Freire advocates for a problem-posing dialog between teacher and student that engages the latter in critical reflection and sensemaking of their own state of oppression. In Freire’s approach, teachers and students co-construct knowledge and collectively engage in *praxis*—i.e., “*reflection and action directed at the structures to be transformed*”.

While Freire’s work centers on pedagogy and learning, we apply his theories to modern information access, where information seekers become ‘learners’ via the search of digital information. We argue that information access platforms can and should serve as sites to enact emancipatory pedagogy: to encourage information seekers to co-opt and co-construct the IA platform, and in turn shape how information is experienced, as part of critically examining their realities within oppressive structures, and taking action to transform them. Our work begins with the observation that we typically relate to technology either as technologists (who envision and build technology) or as users (who seek information via the platform). Although developers and researchers may sometime use the technologies they build, most users of technology do not typically contribute directly to their envisioning or development. This dichotomy between technologists and users of technology mirrors the teacher-student relationship that Freire analyzes. By delegating the tasks of envisioning, designing, and governing information access platforms exclusively to technologists, we in turn elevate them to the position of gatekeepers of knowledge and hand them tremendous influence over our collective sensemaking of our place and relationships in this world; while the rest of the population is designated to be the *objects* of technological influence—their existence digitized for monitoring, modeling, and manipulation. As Freire writes:

“*More and more, the oppressors are using science and technology as unquestionably powerful instruments for their purpose: the maintenance of the oppressive order through manipulation and repression. The oppressed, as objects, as ‘things,’ have no purposes except those their oppressors prescribe for them.*”

This designation of information seekers as agency-free “*objects*” normalizes the platform monitoring their mouse clicks [14] (and even mouse movements [21]) to farm signals and training data for relevance models. The reduction of individual users to their interactions with the system incentivizes their manipulation even by non-platform actors—e.g., the growing practices of clickbaiting [73] and rage-baiting [45] by content producers. Furthermore, even in the face of serious societal concerns this frame restricts our imagination to only be able to conceive the same technologists—or a subset of their altruistic converts—as those that must bear the burden of mitigating the risks to society and its marginalized populations. Seeing through this frame one comes to the inevitable (but false) conclusion that all mitigative research must then focus solely on inventing mechanisms that IA platforms can operationalize to address concerns of bias, fairness, misinformation and disinformation, privacy, and safety; and we must ensure platforms employ said mechanisms by appealing to their better nature, through public pressure, or via enforcement of legal regulations. But what remains elusive just beyond the reach of this frame is the idea that may be the control over these platforms should

never have been centralized in the hands of large corporations, because that threatens Big Tech’s grand project of empire building [43]. To those well-intentioned “converts”, we quote the following text by Freire:

“Our converts, on the other hand, truly desire to transform the unjust order; but because of their background they believe that they must be the executors of the transformation. They talk about the people, but they do not trust them; and trusting the people is the indispensable precondition for revolutionary change. A real humanist can be identified more by his trust in the people, which engages him in their struggle, than by a thousand actions in their favor without that trust.”

This is not an argument against research (or researchers) that develop mitigative intervention mechanisms, but rather a call for us to recognize that those efforts must be grounded in broader collective acts of movement-building to challenge the concentration of power in the hands of large (often multinational) platforms and to ultimately *abolish Big Tech*. Otherwise, any positive impact of these efforts will be temporary and potentially contribute to *ethics washing* [91, 95] and further entrenchment of corporate hold over our IA platforms. Instead, we advocate that societally-motivated technologists should join the struggle in solidarity with the people who are most at risk and harmed by these technologies. With this re-framing, the role of the technologists becomes two-fold. They must present the material risks of emerging technologies for the marginalized as an object of study to the people (users), to encourage them to see themselves as the necessary makers and unmakers of technology in support of their liberation. Simultaneously, technologists must re-envision the structure of our online technology stacks, such that they facilitate marginalized communities co-opting and co-constructing the technology to build pedagogical information experiences that encourage information seekers to critically examine their relations with oppressive structures, and to take action in support of universal emancipation. The inclusion of those at the cross hairs of technology is therefore much more than symbolic. It is *only* through their participation that we can aim for the necessary social, political, and cultural sensitivity in our pursuit for universal emancipation. Their participation is a radical act of connecting their struggles with the making and unmaking of technology, that must be shaped by intra-community and inter-community negotiations. This co-construction of technology is both informed by social movements and itself becomes an integral part of liberatory movement-building.

When designing new modes of participation to center the needs of the marginalized, we must counter cynicism about the likelihood of broad public engagement by reminding ourselves that communities have long been resisting and appropriating everyday technologies to alleviate their oppression [71]. It is true that the ability to appropriate and rework technology is considerably more limited in the age of online services, since the “code” and infrastructure is not easily accessible to or modifiable by users, relative to the ease of appropriating physical devices. Therefore, we cannot assume a reluctance to engage, but must rethink online technology stacks to re-expose those cracks, transforming them into legitimate spaces for not just participation, but appropriation in support of everyday struggles against oppression. These spaces must be structurally guaranteed in our design to support people marginalized by these technologies to reflect, negotiate, and collaborate with each other to co-opt and co-construct the technology to suit their own needs.

To re-construct technology in support of emancipation, we must work together to unequivocally reject the premise of technological determinism [12] and confront the repeating cycles of technodeterministic hype coming from our own technology communities—e.g., just in recent years: Blockchain, Non-fungible tokens (NFTs), Virtual Reality, and now AI. This is not to deny that many of these technologies have their material usefulness,¹² but the ask is to recognize that each of them drove a frenzied hype that tried to convince the world that the ubiquity of the technology was inevitable. Consequently, we must also break the ever-repeating cycles of critiquing emerging technologies and

¹²OK, may be not NFTs. [99]

developing incomplete means of mitigation when our resistance is better directed at those who wield power over these technologies to serve their own need for wealth and power. Instead of only obsessing about mitigations, we must build for something better, something for our collective selves, something hopeful and emancipatory. To quote Freire:

“In order for the oppressed to be able to wage the struggle for their liberation, they must perceive the reality of oppression not as a closed world from which there is no exit, but as a limiting situation which they can transform. This perception is a necessary but not a sufficient condition for liberation; it must become the motivating force for liberating action.”

So, to operationalize Freire’s principles we adopt a problem-posing approach to rethink IA platforms. As the instigators of this work—and as recovering “converts” ourselves who previously approached the sociotechnical concerns of IA through the traditional lenses of fairness, transparency, and related frames—we see our new role as that of problem-posing, to build collective movements for universal emancipation, to free our IA platforms from authoritarian control, and in turn to liberate ourselves from the complicity in the global rise in authoritarianism, following Freire’s guidance:

“Problem-posing education, as a humanist and liberating praxis, posits as fundamental that the people subjected to domination must fight for their emancipation.”

4 Problem-Posing Framework for Emancipatory IA

Our framework comprises of a series of problem-posing questions selected to prompt participants to critically reflect, ideate, prototype, experiment, and study different types of emancipatory IA approaches and capabilities. The problems are chosen to reflect some of the primary concerns of emancipatory IA while minimizing assumptions and prescriptions about *how* these problems are to be addressed. We intend for each of these problems to solicit a diverse range of candidate sociotechnical design ideas and interventions that foregrounds our emancipatory aspirations. Each problem may be considered in isolation, but ultimately we seek visions for IA platforms that can jointly address as many of these concerns as possible. We aim for this framework of problems to expose a rich design space of alternative approaches that communities can explore to challenge the status quo IA platforms. Next, we enumerate our framework of questions. We group these questions into two thematic groups, focusing on (1) the operations and governance of an envisioned emancipatory IA platform, and (2) safeguarding such a platform from capture and co-option.

Theme: Platform operation and governance. We pose a provocation to IA researchers and platform developers to rethink the design of our technological stacks to create the conditions for community co-option and co-construction of IA technologies. In addition, the platform must also explore models of democratic governance and effective mechanisms of policy enforcements. We also need research and prototyping—informed by multiple disciplines including HCI, information sciences, and critical theory—to explore the design space of pedagogical information experiences. To summarize, this theme concerns with models for community co-option and co-construction, democratic governance, policy enforcement, and design of pedagogical information experiences.

Prob1. *How do we design and develop the platform from the margins by legitimizing community co-option and co-construction of the platform?*

Desiderata. The platform must legitimize its own co-option and co-construction through reflective participation of community members. Community members should be empowered to build new IA artifacts (e.g., information retrieval components, moderation services, and information experiences). Community members should be able

to leverage artifacts built by other communities, and may also build competing artifacts if existing ones are inadequate or deemed harmful to community members. Through these participatory interactions the platform must allow for intra- and inter-community negotiations. The platform must mediate these interactions to ensure an equal playing-field for community members. Participation in co-construction must not be biased by who has access to more resources and labor. The model of engaging communities must be informed by theories of power and strive for epistemic justice. Finally, to ensure the platform does not simply extract free labor from already marginalized communities, it must also consider mechanisms for their compensation and long-term sustainability of this work.

Prob2. *How do we develop a critical democratic governance framework?*

Desiderata. Community participation is not a miracle-cure for all our sociotechnical challenges with IA. The opportunity to co-opt and co-construct may not only be leveraged by the oppressed but also the oppressor. We see these tensions already in context of app stores where apps that enable racial surveillance are allowed to exist [70] but those that oppose authoritarian violence are pulled down under government pressure [4]. We must recognize the limits of community negotiations and need for platform governance to bend the arc of platform development towards emancipation. All platforms must also moderate [37] and platform governance must deal with critical moderation policies and mediate challenging cases for moderation. Content moderation and governance are not at odds with emancipatory aspirations; mitigating real-world harms is a core tenet of anti-oppression and emancipatory praxis. The challenge is to develop a critical democratic framework for decision-making with respect to platform governance that is also robust to capture. Any appropriate governance structures must ensure effective and timely removal of critically harmful content from circulation—e.g., child sexual abuse material (CSAM)—as well as to exercise sociopolitical sensitivity in decision-making for challenging cases like the sharing of the ‘Napalm girl’ photo [59].

Prob3. *How do we develop robust enforcement mechanisms for platform policies?*

Desiderata. Democratic decision-making must be supported by robust enforcement of its governing policies and decisions. For critical cases, like that of CSAM, the model for empowering communities to co-construct the platform must not be a loophole to circumvent moderation. However, there is a natural tension here between developing infrastructure for the enforcement of policies and recreating infrastructure of censorship and centralized ideological control. To address this tension, we need innovations in sociotechnical infrastructure; and radical proposals for harm mitigation and justice that are restorative, reparative, and transformative [50].

Prob4. *How do we design information experiences to promote emancipatory pedagogy?*

Desiderata. Emancipatory IA poses the challenge to imagine new ways for us to experience information that are pedagogical and foreground critical scholarship. Pedagogy in this context must not replicate the “banking concept of education”—nor act as tools for indoctrination—but encourage reflection, critical thinking, and dialog. This work must be informed by multidisciplinary perspectives including HCI, design, critical theory, social and political sciences, humanities, and arts.

Theme: Safeguards from capture and co-option. No struggle for liberation has the luxury to step outside of its own cage to architect a new liberated reality; the work begins from within the bounds of the oppressive system to transform

the state of oppression to one of emancipation. The work on emancipatory IA must likewise begin situated within present-day structures of oppression. This requires our work to prioritize safeguarding the platform from capture, blocking, political pressure, economic competition, and co-option. We consider these problems under this theme.

Prob5. *How do we safeguard the IA platform against authoritarian capture?*

Desiderata. The platform must ensure that it cannot be pressured to take an action (or be blocked from taking an action) via state regulations or orders, through corporate or other institutional mandates, using economic and other incentives, by threats and acts of force or violence, and other forms of coercion. This requires that the infrastructure and its management, the governance of the platform, its sources of funding, and other critical aspects be decentralized to avoid exposing critical points of weakness that oppressive-actors can put pressure on to capture the platform. Any serious attempt at building an emancipatory IA platform must strategically consider adversarial strategies of platform capture and develop necessary mitigation.

Prob6. *How do we ensure that access to the platform cannot be blocked by authoritarian forces?*

Desiderata. The platform should support reasonably convenient mechanisms to circumvent attempts by governments or Internet service providers to block access to the service. There are precedents for online platforms—e.g., the messaging app Signal [97]—that provide mechanisms for circumventing authoritarian block on the app. Emancipatory IA must support similar mechanisms to ensure continued access to information under authoritarian censorship.

Prob7. *How do we make the platform sustainable long-term in the face of political pressure and economic competition?*

Desiderata. Any platform that succeeds in, or even attempts to, aid liberatory struggles would likely come under immediate political pressure. Beyond safeguarding against capture, the project of emancipatory IA must also consider how to ensure the physical and economic safety of those conducting the research and actively participating in development of the platform. This includes the safety of community members who participate in the co-construction of the platform. Safeguards may range from funding support, internal policies, community norms, technological safeguards, and other mechanisms. The platform would also face economic competition from other platforms, and capitalist corporations might be particularly motivated make the economics of the emancipatory IA platform untenable for its survival. Therefore, the economics of the platform should be a critical focus of the project. The platform strategy must adopt robust strategies to overcome these external socioeconomic and political forces to be sustainable for the long-term.

Prob8. *How do we safeguard the platform from co-option and extractivism?*

Desiderata. If the platform were to succeed, an immediate concern would be that it might get co-opted by authoritarian and corporate actors. Other existing IA platforms may also attempt to extract economic value from the emancipatory IA platform without reciprocating. The platform design must be intentional in safeguarding against these risks and codify the need for reciprocity through technological design, and licensing and other legal means—e.g., Copyleft licensing [32]. Emancipatory IA should draw insights from the successes and failures of past and current initiatives like the Open Source movement.

5 The Salient Tensions of Emancipatory IA

Emancipatory IA situates itself within the tension between the oppressor and oppressed; and resolves that tension by making a normative commitment to side with the oppressed. However, attempting to simultaneously resolve several of the problems posed in Section 4 evokes other critical tensions—spanning the technological, social, political, economic, epistemic, legal, and personal. These tensions will manifest as we collectively put the theories to practice. In this section, we briefly discuss a short selection of such tensions that we must negotiate.

Platform vs. Protocol. In response to the growing concentration of power of a handful of platforms over public speech and information access, there is an emerging debate over the platform vs. protocol approach [63]. This is particularly dominant in the social media discourse where protocols like ActivityPub and ATProto have gained popularity as alternatives to billionaire-owned social media platforms. In this paper, we have liberally used the term “platform” to refer to the object of our envisioning in the context of emancipatory IA. This is intentional, partly to sidestep this debate earlier in the paper and use “platform” as a shorthand to refer to whatever solution that emerges as appropriate but also to reflect the realities of building for users who may not understand or care about the mechanics of platforms vs protocols, as we discuss below. While protocols present distinct advantages over platforms for user ownership of data, competition, and resistance to single points of capture, as Hendrix and Flowers [44] argue, protocols in lieu of platforms do not inherently challenge the hierarchical social structures that they emerge from; the benefits that protocols do provide often rest on a neoliberal belief that market competition will drive better outcomes. However, Emancipatory IA nudges us to consider democratic governance designed explicitly to drive towards emancipatory outcomes.

Governance vs. Control. Platform governance [39] is a critical challenge for online platforms. Emancipatory IA not only inherits these challenges but must also be hardened against the risk of capture of any governance mechanism by authoritarian actors. It therefore requires additional safeguards and safety protocols, as well as contingency plans for when members of the governing body may be targeted or coerced. Policy enforcement mechanisms must navigate a similarly challenging landscape, balancing effectiveness of enforcement and the risks of implementing the tools for its own capture. This is one of the core tensions that future research on emancipatory IA must resolve.

Top-down vs. Bottom-up Design. Design-related tensions emerge from opening up the platform to community co-option and co-construction. When large masses of Twitter users abandoned the social media site in 2022 to flock to Mastodon, they experienced significant growing pains onboarding onto the new platform [22]. New users had to shoulder the burden to familiarize themselves with a new platform that posed significant friction, especially for non-tech-savvy users, and exercise trust-based actions (such as selecting a Mastodon instance) that many felt ill-equipped to do [83]. In enabling, co-option and co-construction, emancipatory IA must still ensure not to offload responsibilities to the user that might overwhelm them, and continue to prioritize the ease of usability of the platform. It must also avoid fragmentation and drive towards consistency of user experience. Fundamentally, it must drive a balance between the flexibility of design it affords to contributing communities and the overall usability of the platform.

Pedagogy vs. Indoctrination. The pedagogical capabilities of the platform also require necessary safeguards to not recreate the same risks associated with AI persuasion. These safeguards will need to combine technological guardrails with policy interventions, and make space for social negotiations around their design and deployment.

Justice vs. Legality. Lastly, we must embark on this mission clear-eyed about challenges that emerge when the law diverges from justice. We opened this article with examples of authoritarian attempts at capture using executive orders

and mandates. Under the conditions of oppression, the law disproportionately sides with the oppressor, which may create situations where the mandate of Emancipatory IA conflict with existing laws. Emancipatory IA platforms must find a robust protocol for decision-making to navigate these tensions while respecting legal frameworks where possible, human rights, and our social commitment to do no harm.

6 Conclusion

Information access is too critical to our collective sensemaking and sociopolitical discourse to be left in the hands of powerful corporations and economic elites. In a globalized world, our IA platforms must represent the interconnectedness of different communities and cultures worldwide, and not privilege the perspectives of a handful of powerful nations nor continue to reproduce the epistemic injustices of our past. These injustices are the poison fruits of the structures of oppression that have their roots in our history and continue their domination [16] into the present, including colonialism, cisheteropatriarchy, classism, casteism and ableism among others; and so it is the explicit goal of emancipatory IA to aid the dismantling of these structures.

Our approach to emancipatory IA is predicated on a rejection of technosolutionism in that it is not the design of the technology that confronts oppression; but rather it is the creation of spaces for community participation, collective deliberation, and democratic negotiation within the platform, that is so often missing in our modern technology design, that we believe holds the key to realizing emancipatory access to information. The role of design is to structurally encode these spaces within the IA platform, which is necessary but in itself not sufficient for the platform to be emancipatory. Emancipation after all is a political project and so it is the social arrangements within these spaces that are key to determining whether the platform will be emancipatory or oppressive. Emancipatory IA does not brush this political reality under the rug but rather foregrounds it. To bend the arc of online information access towards emancipation we need to experiment with different structures of democratic participation, governance, and decision-making. But beyond structures and processes, we also need to ground these discourses and negotiations in critical scholarship, theories, and practices that distinguishes structural understanding of systems of oppression from reactionary narratives. For this, we must lean on relevant disciplines such as critical theory, information and library sciences, science and technology studies, and others; while being clear-eyed about the growing phenomenon of reactionary co-option and weaponization of social justice language to prop up false victimhood—e.g., [11, 80, 84]. Our “neutrality” must manifest through equal treatment of all groups with respect to humanization and freedom from oppression; not drawing of false equivalences between oppressors and oppressed.

Our problem-posing framework, as the name explicitly suggests, does not point to miracle-cures for the ailing growth of global authoritarianism and oppression. Instead, it responds to the call for building an emancipatory movement within the IA community [67] by identifying some of the key challenges that the movement should focus on. This problem-posing framework is intended to be a concrete starting point for envisioning radically new alternatives. Envisioning in this context must not be idle conjectures but *praxis*; we must use these questions to reflect and act on the problems collectively and in the process co-construct the social arrangements of individuals and institutions committed to the causes of universal humanization and emancipation.

Our movement begins here. We are moved not by the opportunity to imagine a different sociotechnical future but by the abject necessity of it in the face of growing authoritarianism and oppression. We do not move alone but in solidarity with others because our liberation, and the struggle for it, are bound together. We move not in fear but with courage and out of love for our own humanity and that of others. Where the oppressor attempts to capture and control,

we move to co-opt, co-construct, and liberate. We move not just to negate oppression but to create something new, something joyful, something emancipatory. We conclude with one last quote from Freire:

“Because love is an act of courage, not of fear, love is commitment to others. No matter where the oppressed are found, the act of love is commitment to their cause—the cause of liberation. And this commitment, because it is loving, is dialogical.”

Generative AI usage statement

The authors hereby declare that no generative AI technologies were used in the writing of this paper nor at any stage of the research process.

References

- [1] 2025. *Decidim Docs*. Retrieved January 12, 2026 from <https://docs.decidim.org/en/develop/index.html>
- [2] 2025. Free Our Feeds: Save Social Media From Billionaire Capture. <https://freeourfeeds.com/>
- [3] Azadeh Akbari. 2025. Big Tech authoritarianism: Political synergies of an emerging power. *Dialogues on Digital Society* (2025), 29768640251370968.
- [4] Bobby Allyn. 2025. Legal experts condemn Apple bowing to White House’s request to remove ICE tracking app. *NPR* (2025).
- [5] Mariam Asad, Christopher A Le Dantec, Becky Nielsen, and Kate Diedrick. 2017. Creating a sociotechnical API: Designing city-scale community engagement. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 2295–2306.
- [6] Leif Azzopardi, Charlie Clarke, Paul Kantor, Bhaskar Mitra, Johanne Trippas, and Zhaochun Ren. 2024. The Search Futures Workshop (ECIR2024). <https://searchfutures.github.io/>
- [7] Solon Barocas, Asia J Biega, Benjamin Fish, Jędrzej Niklas, and Luke Stark. 2020. When not to design, build, or deploy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 695–695.
- [8] Karl Baumann, Benjamin Stokes, François Bar, and Ben Caldwell. 2017. Infrastructures of the Imagination: Community Design for Speculative Urban Technologies. In *Proceedings of the 8th International Conference on Communities and Technologies* (Troyes, France) (C&T ’17). Association for Computing Machinery, New York, NY, USA, 266–269. doi:10.1145/3083671.3083700
- [9] Matthew Burtell and Thomas Woodside. 2023. Artificial influence: An analysis of AI-driven persuasion. *arXiv preprint arXiv:2303.08721* (2023).
- [10] Micah Carroll, Alan Chan, Henry Ashton, and David Krueger. 2023. Characterizing manipulation from AI systems. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–13.
- [11] Southern Poverty Law Center. [n. d.]. Men’s Rights Activists. <https://www.splcenter.org/resources/extremist-files/mens-rights-activists/>
- [12] Daniel Chandler. 1995. Technological or media determinism.
- [13] Aaron Chatterji, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. *How people use chatgpt*. Technical Report. National Bureau of Economic Research.
- [14] Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. 2015. Click models for web search. *Synthesis lectures on information concepts, retrieval, and services* 7, 3 (2015), 1–115.
- [15] Raffaele F Ciriello, Julian Marx, Marc Cheong, Christoph Mueller-Bloch, and Lars Mathiassen. 2025. Decentralized Social Media. *Business & Information Systems Engineering* (2025), 1–16.
- [16] Patricia Hill Collins. 1990. Black feminist thought in the matrix of domination. *Black feminist thought: Knowledge, consciousness, and the politics of empowerment* 138, 1990 (1990), 221–238.
- [17] Sasha Costanza-Chock. 2020. *Design Justice: Community-Led Practices to Build the Worlds We Need*. The MIT Press.
- [18] Madeleine IG Daepf, Alejandro Cuevas, Robert Osazuwa Ness, Vickie Yu-Ping Wang, Bharat Kumar Nayak, Dibyendu Mishra, Ti-Chung Cheng, Shaily Desai, and Joyojeet Pal. 2025. Generative Propaganda. *arXiv preprint arXiv:2509.19147* (2025).
- [19] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (Boston, MA, USA) (EAAMO ’23). Association for Computing Machinery, New York, NY, USA, Article 37, 23 pages. doi:10.1145/3617694.3623261
- [20] Fernando Pedro Delgado, Stephen Yang, Michael A. Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond “Add Diverse Stakeholders and Stir”. *ArXiv abs/2111.01122* (2021). <https://arxiv.org/pdf/2111.01122>
- [21] Fernando Diaz, Ryen White, Georg Buscher, and Dan Liebling. 2013. Robust models of mouse movement on dynamic web search results pages. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. 1451–1460.
- [22] Jesus Diaz. 2022. Using Mastodon is way too complicated to ever topple Twitter. *Fast Company* (2022).
- [23] Mickey Djuric. 2025. Canada’s digital sovereignty push challenges America’s reach. *Politico Pro* (2025).
- [24] Marius Dragomir and Minna Aslama Horowitz. 2025. Old Tricks, New Tech: How Legacy Media Capture Fuels Today’s Digital Authoritarianism. *Tech Policy Press* (2025).

- [25] Ana Maria Bustamante Duarte, Nina Brendel, Auriol Degbelo, and Christian Kray. 2018. Participatory Design and Participatory Research: An HCI Case Study with Young Forced Migrants. *ACM Trans. Comput.-Hum. Interact.* 25, 1, Article 3 (Feb. 2018), 39 pages. doi:10.1145/3145472
- [26] Ángel Díaz and Laura Hecht-Felella. 2021. *Double Standards in Social Media Content Moderation*. Report. Brennan Center for Justice at New York University School of Law. https://www.brennancenter.org/sites/default/files/2021-08/Double_Standards_Content_Moderation.pdf
- [27] Seliem El-Sayed, Canfer Akbulut, Amanda McCroskery, Geoff Keeling, Zachary Kenton, Zaria Jalan, Nahema Marchal, Arianna Manzini, Toby Shevlane, Shannon Vallor, et al. 2024. A Mechanism-Based Approach to Mitigating Harms from Persuasive Generative AI. *arXiv preprint arXiv:2404.15058* (2024).
- [28] Marwa Fatafta. 2024. Big Tech and the risk of genocide in Gaza: what are companies doing? *Access Now* (2024).
- [29] Michael FitzGerald. 2025. 'Digital sovereignty': why the EU may be shifting from internet regulation to building homegrown tech. *The Conversation* (2025).
- [30] World Economic Forum. 2024. World Economic Forum Global Risks Report 2024. <https://www.weforum.org/publications/global-risks-report-2024/>
- [31] Open Search Foundation. 2025. #FreeWebSearch Charter: Search for the Common Good. <https://freewebsearch.org/en/charter/>
- [32] Dean A Frantsvog. 2012. All Rights Reversed: A Study of Copyleft, Open-Source, and Open-Content Licensing. *Contemporary Issues In Education Research* 5, 1 (2012), 15–22.
- [33] Rudy Fraser. 2025. Beyond Horseless Carriages: Building Communities for the Decentralized Era. Conference talk at ATmosphere 2025. <https://www.youtube.com/watch?v=OZN8y8kVFFI> YouTube video.
- [34] Paulo Freire. 1970. *Pedagogy of the Oppressed*. Herder and Herder New York.
- [35] Miranda Fricker. 2007. *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- [36] Tony Fry. 2009. Design futuring. *University of New South Wales Press, Sydney* (2009), 71–77.
- [37] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
- [38] Yana Gorokhovskaia and Cathryn Grothe. 2024. The mounting damage of flawed elections and armed conflict. *Freedom House* (2024).
- [39] Robert Gorwa. 2019. What is platform governance? *Information, communication & society* 22, 6 (2019), 854–871.
- [40] Rachel Griffin. 2025. The Politics of Algorithmic Censorship: Automated Moderation and its Regulation. In *Music and the Politics of Censorship: From the Fascist Era to the Digital Age*, James Garratt (Ed.). Brepols. <https://sciencespo.hal.science/hal-04325979>
- [41] Matthias Hagen, Josiane Mothe, Smaranda Muresan, Martin Potthast, Min Zhang, Benno Stein, et al. 2026. Retrieval-Augmented Generation—The Future of Search? (Dagstuhl Seminar 25391). *Dagstuhl Reports* (2026).
- [42] Blayne Haggart, Courtney C. Radsch, Chris Riley, Natasha Tusikov, Heidi Tworek, and Samuel Woolley. 2022. As Elon Musk Upends Twitter, Whither the Digital Town Square? *Centre for International Governance Innovation (CIGI)* (2022). <https://www.cigionline.org/articles/as-elon-musk-upends-twitter-whither-the-digital-town-square/>
- [43] Karen Hao. 2025. *Empire of AI: Dreams and nightmares in Sam Altman's OpenAI*. Penguin Group.
- [44] Justin Hendrix and J Flowers. 2022. The whiteness of Mastodon. *Tech Policy Press, November 23* (2022), 2022.
- [45] Kyra-lin Hom. 2015. Rage baiting. *Westside Seattle* (2015). <https://www.westside-seattle.com/west-seattle-herald/2015/05/25/rage-baiting>
- [46] Gerald Inray. 2025. South Africa president says persecution of whites a 'false narrative' as Musk repeats genocide claim. *Associated Press News* (2025).
- [47] Lenz Jacobsen. 2025. Ein Hitlergruß ist ein Hitlergruß ist ein Hitlergruß. *Die Zeit* (2025).
- [48] Sheila Jasanoff and Sang-Hyun Kim. 2015. *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*. University of Chicago Press.
- [49] Lauren Johnson and Ryan Joe. 2023. Generative AI has a Digital Blackface problem. *Business Insider* (2023).
- [50] Mariame Kaba. 2021. *We do this' til we free us: Abolitionist organizing and transforming justice*. Haymarket Books.
- [51] Amba Kak and Sarah Myers West. 2023. AI Now 2023 Landscape: Confronting Tech Power. <https://ainowinstitute.org/2023-landscape> AI Now Institute.
- [52] Dake Kang and Yael Grauer. 2025. Silicon Valley enabled brutal mass detention and surveillance in China, internal documents show. *AP* (2025).
- [53] Jay Caspian Kang. 2025. The Big Tech takeover of American politics. <https://www.newyorker.com/news/fault-lines/the-big-tech-takeover-of-american-politics>
- [54] Jamal Kanj. 2025. Big Tech and the Architecture of the Israeli Genocide Against Palestinians. *CounterPunch* (2025).
- [55] Dara Kerr. 2025. Musk's AI Grok bot rants about 'white genocide' in South Africa in unrelated chats. *The Guardian* (2025).
- [56] Creative Reaction Lab. 2018. *Equity-Centered Community Design Field Guide*.
- [57] Adrienne LaFrance. 2024. The Rise of Techno-Authoritarianism. *The Atlantic* (2024).
- [58] Paddy Leerssen. 2025. From Murdoch to Musk: Platform ownership and the political economy of online content governance. *Platforms & Society* 2 (2025), 29768624251386260.
- [59] Sam Levin, Julia Carrie Wong, and Luke Harding. 2016. Facebook backs down from 'napalm girl' censorship and reinstates photo. *The Guardian* (2016).
- [60] Becca Lewis. 2025. 'Headed for technofascism': the rightwing roots of Silicon Valley. *The Guardian* (2025).
- [61] Donna Lu. 2025. We tried out DeepSeek. It worked well, until we asked it about Tiananmen Square and Taiwan. *The Guardian* (2025).

- [62] Astrid Mager and Christian Katzenbach. 2021. Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified. 223–236 pages.
- [63] Mike Masnick. 2019. Protocols, not platforms: A technological approach to free speech. *Knight First Amendment Institute* 19, 5 (2019).
- [64] Mastodon. 2025. The people should own the town square. <https://blog.joinmastodon.org/2025/01/the-people-should-own-the-town-square/>
- [65] Donald Metzler, Yi Tay, Dara Bahri, and Marc Najork. 2021. Rethinking search: making domain experts out of dilettantes. In *Acm sigir forum*, Vol. 55. ACM New York, NY, USA, 1–27.
- [66] Milmo, Daniel. 2022. Instagram Rolls Back Some Changes to App. *The Guardian* (28 July 2022). <https://www.theguardian.com/technology/2022/jul/28/instagram-rolls-back-some-changes-to-app> Technology section.
- [67] Bhaskar Mitra. 2025. Emancipatory Information Retrieval. *Information Retrieval Research* 1, 2 (2025), 313–339.
- [68] Bhaskar Mitra. 2025. Search and Society: Reimagining Information Access for Radical Futures. *Information Retrieval Research* 1, 1 (2025), 47–92.
- [69] Bhaskar Mitra, Henriette Cramer, and Olya Gurevich. 2024. Sociotechnical Implications of Generative Artificial Intelligence for Information Access. In *Information Access in the Era of Generative AI*. Springer, 161–200.
- [70] Rani Molla. 2019. The rise of fear-based social media like Nextdoor, Citizen, and now Amazon’s Neighbors. *Vox* (2019).
- [71] David Nemer. 2022. *Technology of the oppressed: Inequity and the digital mundane in favelas of Brazil*. MIT Press.
- [72] Tori Noble and Kit Walsh. 2025. President Trump’s War on “Woke AI” Is a Civil Liberties Nightmare. *Electronic Frontier Foundation Deeplinks Blog* (2025).
- [73] Caroline O’Donovan. 2014. What is clickbait? <https://www.niemanlab.org/2014/08/what-is-clickbait/>
- [74] Will Oremus. 2017. Big Tobacco. Big Pharma. Big Tech? (2017). <https://slate.com/technology/2017/11/how-silicon-valley-became-big-tech.html>
- [75] Peter S Park, Simon Goldstein, Aidan O’Gara, Michael Chen, and Dan Hendrycks. 2023. AI deception: A survey of examples, risks, and potential solutions. *arXiv preprint arXiv:2308.14752* (2023).
- [76] Alina Polyakova and Chris Meserole. 2019. Exporting digital authoritarianism: The Russian and Chinese models. *Policy brief, democracy and disorder series* August 2019 (2019), 1–22.
- [77] Stuart Reeves. 2012. Envisioning ubiquitous computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1573–1582.
- [78] Tony Roberts and Marjoke Oosterom. 2025. Digital authoritarianism: A systematic literature review. *Information Technology for Development* 31, 4 (2025), 860–884.
- [79] Amalia G. Sabiescu, Salomão David, Izak van Zyl, and Lorenzo Cantoni. 2014. Emerging spaces in community-based participatory design: reflections from two case studies. In *Proceedings of the 13th Participatory Design Conference: Research Papers - Volume 1* (Windhoek, Namibia) (PDC ’14). Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/2661435.2661446
- [80] Savera. 2024. HAF Way To Supremacy: How the Hindu American Foundation Rebrands Bigotry As Minority Rights. <https://www.wearesavera.org/wp-content/uploads/2024/10/HAFWayToSupremacy.pdf>
- [81] Chirag Shah and Emily M Bender. 2024. Envisioning information access systems: What makes for good tools and a healthy Web? *ACM Transactions on the Web* 18, 3 (2024), 1–24.
- [82] Clay Spinuzzi. 2005. The Methodology of Participatory Design. *Technical Communication* 52, 2 (2005), 163–174.
- [83] Ana Pop Stefanija. 2023. Solid, mastodon, and the risk of overburdening the user. *Internet Policy Review* (2023).
- [84] Patrick Strickland. 2016. White lives matter: A new US hate group shows its face. *Al Jazeera* (2016).
- [85] Adam Taylor. 2023. A historic rise in global conflict deaths suggests a violent new era. (2023). <https://www.washingtonpost.com/world/2023/06/29/conflict-war-deaths-global-peace-rise-casualty/>
- [86] Josh Taylor. 2025. Musk’s AI firm forced to delete posts praising Hitler from Grok chatbot. *The Guardian* (2025).
- [87] Johanne R Trippas, J Shane Culpepper, et al. 2025. Report from the fourth strategic workshop on information retrieval in lorne (swirl 2025). In *ACM SIGIR Forum*, Vol. 59. 68.
- [88] Donald J. Trump. 2023. Preventing Woke AI in the Federal Government. <https://www.whitehouse.gov/presidential-actions/2025/07/preventing-woke-ai-in-the-federal-government/>
- [89] Eddie L. Ungless, Nina Markl, and Björn Ross. 2025. Experiences of Censorship on TikTok Across Marginalised Identities. *Proceedings of the International AAAI Conference on Web and Social Media* 19, 1 (Jun. 2025), 1952–1965. doi:10.1609/icwsm.v19i1.35912
- [90] United Nations Meetings Coverage and Press Releases 2023. 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>.
- [91] Gijs Van Maanen. 2022. AI ethics, ethics washing, and the need to politicize data ethics. *Digital Society* 1, 2 (2022), 9.
- [92] Nicolas Verón. 2023. Much of the Global South is on Ukraine’s Side. *Peterson Institute for International Economics* (2023).
- [93] Sanne Vrijenhoek, Mesut Kaya, Nadia Metoui, Judith Möller, Daan Odijk, and Natali Helberger. 2021. Recommenders with a mission: assessing diversity in news recommendations. In *chiir*. 173–183.
- [94] Sanne Vrijenhoek, Lien Michiels, Johannes Kruse, Alain Starke, Jordi Viader Guerrero, and Nava Tintarev. 2023. Report on normalize: The first workshop on the normative design and evaluation of recommender systems. In *CEUR Workshop Proceedings*, Vol. 3639. CEUR-WS.
- [95] Ben Wagner. 2018. Ethics as an escape from regulation. From “ethics-washing” to ethics-shopping? (2018).
- [96] Meredith Whittaker. 2021. The steep cost of capture. *Interactions* 28, 6 (2021), 50–55.
- [97] Meredith Whittaker. 2024. Proxy Please: Help People Connect to Signal. <https://signal.org/blog/proxy-please/>
- [98] Erik Olin Wright. 2020. *Envisioning real utopias*. Verso Books.

- [99] Maya Yang. 2023. The vast majority of NFTs are now worthless, new report shows. *The Guardian* (2023).
- [100] Tianyuan Yu, Adam Robert Thieme, and Lixia Zhang. 2025. Identifying Roadblocks in Building Decentralized Apps. *Proceedings of the 2025 Applied Networking Research Workshop* (2025). <https://api.semanticscholar.org/CorpusID:280059772>