

"It just requires so much more creativity": Barriers and Workarounds to Gathering Information for AI Contestation

Sohini Upadhyay
Harvard University
Cambridge, Massachusetts, USA
supadhyay@g.harvard.edu

Dasha Pruss
University of Illinois Chicago
Chicago, Illinois, USA
dpruss@uic.edu

Alicia DeVrio
Human-Computer Interaction
Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
adevos@andrew.cmu.edu

Krzysztof Z. Gajos
School of Engineering and Applied
Sciences
Harvard University
Allston, Massachusetts, USA
kgajos@g.harvard.edu

Naveena Karusala
School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia, USA
nkarusala3@gatech.edu

Abstract

Gathering information about AI systems is essential for contesting their use; it forms the basis of arguments about how and to what extent AI is causing harm. Information thus plays a central role for advocates like lawyers, journalists, and auditors contesting harmful AI systems. However, there is little systematic understanding of how these actors, many of whom are newly encountering AI in their advocacy work, access and use information effectively in this process. Understanding this information work can offer valuable insights for supporting effective contestation of harmful AI systems—work that is typically taken on by underresourced advocacy groups to begin with. To better understand information work in AI contestation, we interviewed 18 advocates in the United States (US) who have contested the use of AI in high-stakes domains, such as public benefits and housing. We characterize advocates' strategies for accessing information that is useful for contestation, including a range of creative yet resource-intensive and risky workarounds that they use to overcome opacity. We discuss implications of our findings for the effectiveness of popular transparency policy strategies in the US and offer additional ways to support the social fabric that makes advocates' information work effective.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**.

Keywords

contestation, AI, algorithmic decision-making, information work, advocacy, transparency

ACM Reference Format:

Sohini Upadhyay, Dasha Pruss, Alicia DeVrio, Krzysztof Z. Gajos, and Naveena Karusala. 2026. "It just requires so much more creativity": Barriers and Workarounds to Gathering Information for AI Contestation. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3772318.3791923>

1 Introduction

As governments and private companies integrate AI into high-stakes decision-making, evidence is mounting of the real and potential harms that ensue. In the United States (US), where our study is situated, algorithms have cut people off from care, such as in cases where public benefits were slashed for the elderly and people with disabilities [2]. AI has also been found to amplify bias in critical infrastructure, such as job search tools that exhibit gender bias [99] and tenant screening tools that are racially discriminatory [84]. Indeed, entire databases have been dedicated to tracking the unfolding harms of AI [1, 2].

With the proliferation of algorithmic tools across public life, people with limited prior relationship to AI have suddenly and unexpectedly found themselves contending with its harmful effects. Simultaneously, top-down policy approaches to mitigate these problems, such as US state legislation requiring deployers to implement proactive audits of AI tools, have had limited success [43, 44]. As a result, impacted communities are taking matters into their own hands. AI tools are increasingly being contested not only through dedicated professional audits but also through social, legal, and discursive means. This has motivated a nascent body of work in Human-Computer Interaction (HCI) and adjacent areas like Fairness, Accountability, and Transparency (FAccT) on how advocates from diverse backgrounds—lawyers, journalists, and community members—go about the work of contesting the use of AI tools [40, 49, 50, 98]. Still, there has been relatively more focus on the efforts of professional auditors, who have specialized skills and dedicated time to conduct AI audits [20, 44, 75, 94], or investigations of an algorithmic system for potential problematic behavior [16]. By contrast, advocates who are working to contest AI on the frontlines



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3791923>

draw on contextual knowledge of their respective domains—such as public benefits or housing—but may not have dedicated resources or technical expertise on algorithms or auditing more broadly.

Existing research has established that AI contestation can be challenging work for advocates and communities, with a major barrier being a lack of transparency [11]. Harmful algorithms are often used for months or years before being discovered [50], and proprietary barriers to technical information about the function of AI tools are common [21]. Yet access to information provides the foundation for contestation: information evidences harm, clarifies how and why it occurred, and points to where responsibility lies. Prior work has recommended technical and policy pathways to improving transparency, such as calls for ex-ante¹ AI accountability evaluations [68, 97]. These recommendations target ideal outcomes, but it is less clear whether and how they might function in practice.

A key aspect of AI contestation that could be better understood is *information work*, or the work to identify, access, understand, and use information [33]. Specifically, there is an urgent need to understand the perspectives of the diverse range of stakeholders on the frontlines of AI contestation, who may be engaging in forms of information work that differ or intersect with professional auditors. The stakes of successful and timely contestation are high—individuals appealing automated cuts to Medicaid assistance, for example, have passed away waiting for justice to be served [41, 42]. Understanding the information work underlying contestation allows us to visibilize power structures that seek to undermine it and to evaluate pathways to make it easier.

To explore the social, technical, and legal dimensions of information work, we conducted interviews with 18 advocates who have experience contesting AI tools. Participants were third-party auditors, journalists, lawyers, researchers, and policy advocates. We sought to understand their information work—and barriers they faced—throughout the process of contestation, from learning about harms caused by AI, to accessing and using information to contest them. Our findings indicate that as AI contestation has become a central part of more traditional forms of advocacy, advocates have started to face new challenges to the information work they have always needed to do to oppose social injustices; meanwhile advocates from technical and non-technical backgrounds alike encounter resource constraints and risk as actors opposing more powerful institutions. To effectively contest the harms of AI tools, advocates must overcome barriers including opacity, adversarial tactics, resource constraints, power dynamics advantaging AI deployers and developers, and risks to themselves and the communities they serve. In response, advocates drew from their professional backgrounds to devise a range of creative workarounds, relying on interdisciplinary strategies, sustained effort, and occasional luck to build productive relationships, interface with AI systems and developers, and forge effective open records strategies.

Drawing on our findings, we discuss why policy mandates for transparency are insufficient for accountability. Not only do developers and deployers find ways to obstruct information access, but there are also additional obstacles posed by the political economy of AI that add to the burden of information access and use for contestation. We conclude by highlighting the essential role of

community and collaboration in information work for contestation, particularly as AI contestation happens in diverse professions and organizations. The contributions of this work are thus three-fold. We offer an empirical understanding of the challenges and creativity involved in the information work of AI contestation. We provide an analysis of the political economy that contributes to the burdens of information work, discussing the limits of transparency as a mechanism for accountability without community enforcement. We also draw on insights from our findings to discuss implications for sociotechnical interventions to support information work.

2 Background and Related Work

AI contestation has been examined by researchers in the HCI, FAccT, and technology policy communities. We start by discussing existing research on information needs in AI contestation, particularly in auditing and litigation. We then summarize the policy landscape on access to information about AI systems. We then describe the lens of information work, which we draw on to situate information access and use in the social context of AI contestation.

2.1 Information Needs for AI Contestation

Facilitating the contestation of harmful AI systems has emerged as a crucial tenet of AI governance [12]. Some scholarship understands contestation as a technical problem that enables a response to specific decisions made by AI systems [52]. In other scholarship, it is viewed as a sociotechnical problem of supporting responsiveness to human intervention throughout the development lifecycle [9]. Still other work analyzes contestation of AI systems more broadly as a right that can be enabled through policy [15, 67]. Other research taking a more bottom-up perspective explores community-led, often oppositional, methods of contesting AI systems [13, 40]. Empirical research demonstrates that contestation is a challenging process for marginalized communities, emphasizing the importance of advocates and the care they put into supporting the end-to-end process [54]. Retrospective research has highlighted how surfacing and sharing information about negative impacts is important to the practice of contesting and abandoning harmful AI systems [50]. Within this body of work, research on information needs has focused on particular communities of advocates, like auditors and lawyers, which we detail below.

2.1.1 Information Needs in AI Auditing. AI auditing, or empirical investigations of an algorithmic system for potential problematic behavior [16], has grown in recent years into a robust field [20, 34, 63]. AI auditors aim to identify potential issues in AI systems and collect additional information to prove or disprove that those issues exist. For example, audits have evidenced bias in search engines [72], online ad delivery [93], resume search tools [30], and chatbots [46]. AI audits and related forms of data-based advocacy can be driven by a variety of different stakeholders, ranging from more standardized and centralized approaches typical in government or industry [34, 82, 94, 101] to more participatory ones often involving community groups [48, 60], labor unions [26, 36, 59, 83, 90], and end users [27, 38, 39, 62, 64, 89]. They can also be first-, second-, or third-party audits [34], which typically involve decreasing levels of access to AI systems but increasing public transparency of results.

¹Ex-ante is a Latin term for “before the event” employed in legal and policy contexts.

Investigations of the dynamics and impacts of second- and third-party AI auditing have highlighted that, although audits are typically initiated for the ultimate goal of accountability, their actual impacts can be minimized due to the lack of support for accessing, understanding, proving, and translating relevant information [44, 102]. In this vein, a recent analysis of AI audit tools and practices emphasizes the importance of disclosing use of AI, data access, and open records laws like the Freedom of Information Act (FOIA), in order to meet information needs for impactful audits [75]. That said, other work has demonstrated that open records data can be mired in issues of availability, usability, and comparability [85]. Prior work has also recognized that once audits are complete, simply sharing the results may do little to enable accountability [44, 75, 81]. Making results understandable and accessible to those involved in advocacy and litigation is a necessary next step [75]. Advocacy technologists who are more embedded in civil society organizations confirm that translational information work is important [92], and describe working with non-technical colleagues to translate needs into technical specifications and to support digital literacy [29]. In this study we focus on third-party auditing in nonprofit settings, as one of many methods for AI contestation and accountability more broadly. We seek a more in-depth understanding of information work not only in auditing but also in litigation, journalism, and other avenues for advocacy, in order to gain a broader view of how information feeds into these interconnected strategies.

2.1.2 Information Needs in Legal Contexts. Legal scholars have grappled with AI governance from multiple perspectives. Prior work has identified where due process—the assurance that fair procedures and other legal rights will be respected—fails in the face of automated-decision making systems [31]. Scholars have advocated for an individual right to contest AI [51]. They have also called for ex-ante AI accountability, where demonstrating that AI is not harmful is the legal onus of deployers and developers [68]. In parallel, a growing body of HCI research has studied litigation and other legal mechanisms to counteract harmful AI systems, emphasizing the importance of accessing, understanding, and translating information in legal spheres. For example, one study identified challenges that US public defenders faced with accessing information about computational forensic software used by the government to convict people and effectively convincing judges and jurors of the software’s flaws [49]. Another study similarly highlights barriers to litigation in cases where AI has caused harm, including accessing and framing information that makes harms legible to courts [70]. Taking a design approach, Rao et al. [83] built, deployed, and evaluated a system for estimating lost wages to support a rideshare labor union in legally contesting members’ deactivations. This process visibilized interconnected information workflows between frontline organizers and legal team members, sped up and reduced errors in the legal team’s workflows, and enabled the legal team to pursue new advocacy work, including crosschecking company-provided information and legal action beyond individual appeals. Our study aims to extend these lines of work, probing information access, use, and barriers in the process of legal advocacy, alongside interrelated forms of advocacy, across a wide range of AI harm domains.

2.2 Policy Landscape and Implementation Challenges

The information work driving AI contestation is situated in a policy landscape affording varying degrees of transparency and legal mechanisms of redress. Within the US context, federal AI policy has been scarce and shifting [8], often relying on common law, or law built on precedent from judicial decisions as a supplement. An example of the latter is *Sandvig v. Barr*, which established the legality of adversarial auditing [73]. As a result, an assortment of AI-related policy proposals has emerged at the state level [17, 28, 74]. A common theme amongst these efforts is a documentation disclosure-based model of transparency and accountability [88], which proposes that audits and impact or risk assessments are reported to agencies, the public, or internally; however, emerging evidence suggests this may have limited success [101]. Supplementing this strategy are nascent “burden shifting” proposals for civil legal procedure reform; these are intended to prevent AI lawsuits from being dismissed before legal discovery² by requiring deployers or developers to provide information demonstrating that their tools are not harmful at the onset [25]. However, there is not yet a thorough understanding of the types of information that should be made available to plaintiffs who have been harmed, which is needed for this reform to be effective. To this end, in the context of predictive optimization, Wang et al. [97] conducted a case analysis to build a rubric of questions developers must address to demonstrate that their tools are not harmful. In empirically studying the information work underlying AI contestation across a wider range of contexts, we aim to expand understanding of the kind of information that policy proposals could mandate developers and deployers to provide. Given that empirical work can uncover the conditions required to render policy effective [66, 87], we analyze policy proposals in relation to our findings in Section 5.

2.3 Critical Perspectives on Transparency

Critical perspectives on transparency contextualize information access within political context and meaningfulness to people. Questions of what to reveal and what to conceal are laden with power [19, 45], and new forms of control may arise alongside visibility of information [104]. A singular focus on mandating transparency, or “transparency-washing,” can also have potential negative effects, redirecting attention away from underlying questions about the concentration of power and leaving the status quo unchanged [103]. Recognizing this can ensure that institutions do not take credit for ineffective transparency, where actual accountability is never reached [103]. Corbett and Denton [32] argue that transparency must center people and their agency, needs for justice, and political context. Dalal [36] and Calacci et al. [26] demonstrate the limitations of transparency in their work on supporting a rideshare union with platform accountability, finding that in practice, data, data access, or transparency is meaningless without the infrastructure to do something with it. As we focus on the work of information access and use, we analyze the impacts and limitations of transparency for enabling AI contestation specifically.

²Discovery is the formal, pre-trial phase of a lawsuit where parties exchange information and evidence to prepare for trial or a settlement.

2.4 Information Work

We orient this research around the notion of information work, developed by medical sociologists Corbin and Strauss [33]. They define information work as “the quest for, the receiving of, and the passing of information” [33]. This definition has been expanded upon by HCI research on health informatics to highlight insights into health management and inform the design of digital health tools [24, 55–58, 78]. Information work has also been leveraged in other domains, such as understanding the labor of essential workers [14], disability service providers [7], and school liaisons for immigrant parents [100]. Taken together, this body of work provides us with the language to highlight how information relates to workers’ broader goals—in our case, contestation of AI. Research has conceptualized social dimensions of information work. Employing the vocabulary of the information journey—recognizing need, seeking, using, and interpreting information—can draw attention to access barriers, as well as to who or what facilitates access, highlighting quality of information and how people make sense of it [24]. It also conceptualizes the movement of information across contexts: information can be interpreted differently across different social worlds, and effectively translated (or not) across these boundaries [56]. We draw on these notions of information work to situate information access for contestation in social, political, and economic factors.

3 Methods

The goal of this study was to understand the information work of advocates seeking to contest the harms of AI in the US context and the challenges they face in accessing and using information. Below, we describe how we collected and analyzed data to answer this research question, ethical considerations, and positionalities.

3.1 Data Collection

3.1.1 Participants. We conducted interviews with 18 participants between April and September of 2025. Participants included advocates based in the US who had experience with discovering and pushing back against harms caused by AI. We initially recruited participants through a combination of convenience [91] and purposive [77] sampling. We advertised the study across our professional networks, which spanned academic, legal, policy, union, non-profit, and industry spaces. We leveraged these networks for participant introductions or ideas about who to contact. The remaining participants were recruited via snowball sampling [71]. Throughout, we prioritized diversity of domain expertise and advocate role. We recruited people with experience in multiple high-stakes domains, including housing, healthcare, hiring, education, hate speech, finance, and criminal justice. These efforts resulted in participants who were lawyers, journalists, researchers, auditors, and policy advocates. They typically had been involved in contesting AI tools that were used in a specific setting, such as a healthcare system or a state’s public service, but a few participants had also been involved in contesting AI tools used on platforms with global reach, such as social media sites. Participant roles and their domain expertise are summarized in Tables 1 and 2, respectively.

3.1.2 Interview Protocol. Interviews were one hour each, with one interview lasting an hour and a half, and another lasting 2 hours.

ID	Role
A1-A3	Auditor
J1-J2	Journalist
L1-L9	Lawyer
P1-P2	Policy Advocate
R1	Researcher, Union
R2	Researcher, Client Advocacy Startup

Table 1: Participant roles. Some participants held multiple professional roles (e.g. lawyer and policy advocate), but we list participants based on the role they highlighted in our discussion. To preserve anonymity, these roles do not reflect participants’ actual professional titles.

Domain	Number of Participants
Employment	6
Healthcare	5
Social Media	5
Public Benefits	3
Education	2
Housing	2
Finance	2
Criminal Justice	2

Table 2: Domains discussed by participants. Participants had expertise contesting harms of AI in a wide range of domains, but in this table we tally the domains they focused on in their interviews. These domains are overlapping (e.g., contesting platform-mediated employment ad discrimination is counted under both Employment and Social Media).

All interviews took place over Zoom and were conducted in English, which all participants were fluent in. We audio recorded participants’ interviews after asking for consent and explaining their right to speak off the record both pre-emptively and reactively. Interviews started by asking about participants’ professional background. We then asked them to choose a particular instance of contesting an AI harm to focus on for the interview (some participants chose to discuss multiple related experiences at a high level, due to privacy concerns). Based on this instance, we asked them about the timeline of events, the information required to achieve their goals around contestation, why and how this information was accessed and used, and barriers to doing so. Finally, we ended by asking participants to speculate about what might have been the ideal information that could have supported their case, as well as how other aspects of their work could have been made easier. Interviews were transcribed using Zoom’s software and then deidentified. We offered every participant a \$25 Visa e-gift card as recognition of their time spent participating in the study.

3.2 Data Analysis

We conducted an inductive interpretive analysis [69] of the interview data. Sohini, Dasha, and Naveena started by open coding one transcript each, and also all coding an additional fourth transcript. Throughout this period, we met to discuss our codes to understand how we each interpreted the data. In some cases, we updated the

codes to align with different interpretations that were agreed upon during discussions. Finally, we split up coding the rest of the transcripts, with Sohini coding the bulk of them. Codes focused on the underlying meaning of participants' words, rather than topics. Examples of codes include “political climate shaping risk of audit” and “resource scarcity driving FOIA noncompliance”. After coding the majority of transcripts, we met to begin constructing themes from the codes, at which point we noticed that information work was a concept that helped frame the data. We thus created themes such as “iterative nature of information work” and “barriers beyond information access”, which inform our findings below.

3.3 Ethical Considerations

This study was reviewed by our Institutional Review Boards (Harvard University #IRB24-1733). More broadly, it is important to acknowledge that participants in this study are engaged in advocacy work that is adversarial to powerful AI deployers and developers. Given the revealing, instructional, or critical nature of participants' reflections, anonymity is essential to their ability to continue this work without retaliation. To mitigate the risk of deductive disclosure, we limited the granularity of reported participant information. We use broad role categories rather than professional titles and report domains of expertise only in aggregate form. We omitted years of experience and demographic details. Participant requests to speak off the record were honored. Some cases discussed in interviews are part of the public record. In these instances, we did our best to reduce the risk of deductive disclosure by contextualizing quotes with minimal details. At the same time, we attended to the overall meaning participants sought to convey, in recognition of their agency and role as public activists, too.

3.4 Positionality

We acknowledge that the study design, data collection, and analysis were all influenced by our positionalities. We are researchers based in academic institutions in the US, with training in computer science, HCI, and philosophy, prior professional experience in the tech industry and legal nonprofit work, and prior personal experience in community, labor, and political organizing. We understand communities' contestation of AI systems as an expression of public voice regarding whether and how AI should be used in society, motivating our interest in conducting the study. The research focus on information needs was motivated by Sohini's proximity to burden shifting proposals advocated by the Surveillance Technology Oversight Project (STOP). We did not interview members of STOP but they did refer some participants. Our limited legal expertise likely impacted the granularity of responses shared by participants with a legal background as well as the lens through which we engaged with the findings.

4 Findings

Our findings describe different facets of the information work required of advocates for effective AI contestation. We begin by describing the life cycle of information work for contestation, highlighting the opacity imbued throughout and how participants' professional roles shaped their practices. Next, we dive deeper into those practices and describe creative workarounds participants

used to access the information they needed, despite opacity. We then describe how conditions in the broader political economy of AI advantaged deployers and developers of AI, amplifying and increasing the burden of participants' information work. We end by discussing opportunities that participants articulated for supporting information work. Our findings are summarized in Table 3.

4.1 Encountering Opacity in the AI Contestation Life Cycle

The information work of contestation typically began with harm discovery, followed by defining the goals of contestation, and then seeking and learning from information. We note, however, that this “life cycle” of information work for contestation was a non-linear and iterative process—the goals and information required for contestation were dependent on each other and could change throughout sensemaking. At each stage of this life cycle, advocates encountered opacity. Throughout, advocates' professional backgrounds shaped how they approached contestation and therefore information work. We found that auditors often sought data for evaluation (as described in prior work [20, 75]), but many participants who were client-facing advocates in high-stakes domains—long before the proliferation of AI—sought many types of information that could support their domain-specific advocacy work and client needs.

4.1.1 Harm Discovery. The initial decision to contest AI typically began with learning about harm being done to a community and suspecting AI was involved. Many participants were part of organizations that welcome contact from community members, such as legal nonprofits, client advocacy startups, and unions, so learning about harm was relatively straightforward. Discovering that AI was responsible, however, required more effort due to opacity, resource constraints, and limited awareness that AI could affect their respective domains at all. For example, L4 explained how clients called their organization to report unexplained cuts to public benefits:

“We probably had two dozen calls from clients who were experiencing cuts in their [benefits]... the common thread among the various clients calling with these stories was that there was no explanation, and they would tend to ask the [workers]... and the [workers] would say, ‘It’s not me. It’s the computer.’”

Similarly, R1, a researcher for a labor union, described how union members reported being displaced from their roles with little explanation, leading the union to file a complaint. However, it was only years later, when the practice spread to another work site, that R1 investigated further and discovered that an algorithm was involved in the work restructuring. R1 noted that numerous other violations obscured this particular issue: *“Unfortunately, there’s many, many problems [in the workplace]... and so there are a ton of violations that come before me and that I work on.”*

In contrast, participants who were auditors, policy advocates, journalists, or lawyers engaged in strategic civil rights litigation did not describe being community members' go-to contacts in the same way. While these participants were aware of potential harms of AI at scale, they often struggled to connect with affected individuals. For example, in policy advocacy around AI use in employment, P2 explained that *“we make sure we partner with labor unions”* but

Finding	Overview
4.1 Encountering Opacity in the AI Contestation Life Cycle	Participants' backgrounds and systemic opacity shaped their practices across the life cycle of information work, characterized by harm discovery, strategizing, and information-seeking and sensemaking. These phases were iterative and non-linear, requiring access to various kinds of technical and sociotechnical information.
4.2 Creative Workarounds for Information Access	Participants overcame opacity by leveraging relationships with affected individuals and other advocates, experimenting with AI systems and interacting with their developers, and using public records laws.
4.3 The Political Economy of Information Work	Participants encountered significant barriers to information work, including adversarial tactics, labyrinthine bureaucracy, flawed governance, resource constraints, power dynamics that shaped the nature of contestation, and risks for impacted individuals.
4.4 Opportunities for Supporting Information Work	Participants described possible remedies, ranging from unfettered access to proprietors' data, documentation requirements, and system-wide reforms. Participants expressed distrust towards deployers and skepticism about how realistic or impactful such changes would be.

Table 3: Summary of findings

could not directly engage with affected workers “*just because of [organization] size and bandwidth.*” These participants were also often investigating AI in contexts where harms were less visible compared to issues like benefits cuts or work restructuring. For instance, when discussing a job application screening tool, J1 explained that rejected applicants typically blame themselves:

“...often people don't even know that they're being harmed, right? They're getting rejected from a job. What do they know if it's, like, a resume parser that didn't work or not? They just assume that they're not the most qualified person.”

For auditors, engaging community members was not common in their profession, making it challenging to know what harms to even evaluate algorithms for. As A1 shared: “*You still like don't know if you're missing something. It's still like poking into a black box, and like hoping...*” because “*I [A1] only see these people as lines on a spreadsheet...it's easy to lose your grounding.*” Participants who were more embedded in community and aware of AI still could not avoid opacity entirely, since AI deployers obfuscated details. L5, who discovered that a tenant screening algorithm was denying housing to individuals, remarked, “*even if they [affected individuals] suspect it [AI use], they don't know which vendor it was... I don't know whose website to start digging around on.*”

4.1.2 Developing Contestation Strategies and Tailoring Information Work. Upon discovering that algorithms were causing harm, advocates then needed to decide on strategies for contestation, shaping the information sought and disseminated.

Contestation strategies were often aligned with the responsibilities and ethos of participants' professions. For many participants in client-facing roles, such as lawyers and researchers, contestation strategies prioritized the autonomy and agency of impacted individuals. For client-facing lawyers, this sometimes meant prioritizing

harm reduction over pursuing litigation that could strengthen regulation or lead to AI tool reform or abandonment. L7 described this tradeoff and its impact on information work in the context of contesting faulty algorithmic evidence used against people in criminal trials—they abruptly stopped gathering evidence when the client wanted to take a plea deal. They described how this was a pervasive tradeoff in their profession: “*We do risk mitigation on like so many different kinds of levels as lawyers. We cannot prevent harm from happening but we can make that harm less bad... That's the universe in which we're operating.*” Other participants, especially journalists, policy advocates, and auditors, were more focused on raising public or legislator awareness of AI-related harms. For these participants, there were other professional norms that shaped their information work. For example, J1 viewed their role as a journalist as going beyond “*theoretical harms*” to share “*what's happening on the ground.*” Accordingly, they shared how they were planning on a longform publication spanning a longer timeframe of incidents, but made sure to publish articles along the way because they “*felt like some of it had, like sort of, needed to be told about immediately.*”

Often, effective contestation strategies required working with expertise or people from different sectors, shaping contestation strategy, and how participants sought and shared information. For professional auditors, staying up to date on journalism or hearing from community-engaged advocates helped them decide how to approach audits in order to support legal claims. On the other hand, lawyers and policy advocates sometimes worked with civil society organizations to support lobbying and impact litigation, and with technical experts to support evaluation of AI systems they were contesting. We note that these collaboration networks were not always robust, however, and information-sharing may or may not find the right people at the right time. For example, some auditors shared their audit results through “*traditional PR*” as

A3 noted, using tactics like sending reports to companies and civil society groups who might act on the information. More tenuously, they also relied on *“an informal network of people”* and conference presentations to make themselves available to litigators who might need their expertise.

Beyond advocacy goals and collaborations, factors like opacity and target audience also impacted contestation strategy and associated information work. In some cases, the opacity of impacts on individual people led advocates to prioritize systemic change. For example, L1 shared their experience challenging ad discrimination: *“It’s very difficult to challenge the ad delivery system with an individual plaintiff, because people don’t know the ads they don’t get.”* Given the difficulty of establishing injury for individuals, they took a consumer protection approach, which required evidence of aggregate harms. Contestation strategy could also change after accessing information and assessing its quality. L8 described discarding a disparate impact approach to litigation in the context of predictive policing: *“We looked at the data, and we were like, if we can’t figure out who’s Hispanic and who’s white, it’s gonna be really hard to try and make out one of those claims.”* In other cases, the intended audience for a given contestation strategy shaped information work. For example, R2, who supported people with insurance denials, needed to gather information that spoke to the medical background of reviewers, who sometimes *“care most about the medical validity and need for the situation, and not what [insurance] contracts are relevant.”* In the same vein, L2 described how their data on ads discrimination against older women resonated with the judge, who was also an older woman, speculating that this helped them win the case. P2, also working on policy advocacy to address ads discrimination, similarly emphasized the impact of connecting with receptive lawmakers who could help with translation of advocacy goals; they shared how one strong sponsor of a bill was able to *“say to [their] colleagues, like, no, the internet is not going to break, no, like, small businesses will still be able to advertise online.”*

4.1.3 Information-seeking and Sensemaking. As participants developed contestation strategies, they sought a range of information and came up against numerous barriers to access. We found that different kinds of information were useful for different professional practices. For many participants, including lawyers, journalists, and the union researcher, documentation describing a tool’s function and intended purpose facilitated preliminary sensemaking around why the tool was being deployed and whether it complied with existing laws. For example, R1 explained how knowing more about how an algorithm was being used alerted R1 that it might violate mandates around having licensed workers involved in decision-making processes. However, there was little incentive for information about an algorithm’s exact function to be shared publicly by developers and deployers, and this had worsened with increasing attempts to contest AI. L5 noted the increasing opacity of vendors’ webpages: *“Oh, gee! A few years ago I was finding this level of information, and now I’m finding, on a lot of other sites, less information.”*

Some participants sought further insight into AI logic, such as explanations, features, and code, and ran into the same barriers. Advocates often sought features and predictive variables to make connections with their domain knowledge and to better understand

harm. For example, L5, an expert in housing, was able to contest algorithmic rental screening by making an argument against the use of housing vouchers as an algorithmic input. Particular technical strategies for contestation could warrant more extensive information about model structure and training data, as in the case of L6 and technical collaborators who wanted to demonstrate that there were less discriminatory alternatives to the tool being contested.

Navigating technical information needs required dealing with opaque model changes. L8 described how such changes made sense-making challenging: *“They were actively making changes to the program in real time, changing the curve criteria, changing the weights of the model... So it was very hard to kind of pin down exactly what we were looking at...”* To this end, L7, who was contesting algorithmic evidence in criminal court, required the tool’s code as well as previous versions of the code and build environments used when the purported evidence was generated. Reflecting on these technical needs, L7 emphasized that the absence of technical documentation is itself useful grounds for contesting an AI tool, explaining, *“...if you’re building a car, I want to make sure that you have the requisite number of screws and washers and other pieces... to make sure that I can go from what you started with into the functional version.”*

Many participants also needed to understand tool outputs and effects, such as logs of decisions. Combined with validation studies and relevant demographic data, this information helped advocates demonstrate discrimination or incorrect evaluations. Across professional roles, participants sought this information for audits at various levels of formality, scale, and statistical rigor. Auditors and lawyers found algorithmic outputs useful for formal and statistically rigorous evaluations in service of a broader legal strategy, while journalists used informal audits to make sense of tool behavior for investigations. Across the board, a lack of access to quality data was a barrier. L6, a lawyer who frequently collaborated with technical experts, explained, *“...a big pain point when you get to quantitative testing is having data, like clean data, that you actually know what it is and what it represents and where it came from.”*

In addition to technical details, participants wanted to understand the sociotechnical context of AI tools. This was often challenging since it required piecing together many disparate sources of data that were not always publicly accessible, but understanding the bureaucratic processes surrounding AI tools allowed for effective legal cases, complaints to regulatory bodies, reporting, and even new avenues for contestation. This information could include a tool’s design process and goals, company policies, governance structures, evidence of meeting compliance standards, funding sources, and records of evaluation and communication (such as emails) that indicated whether deployers or developers knew about errors. For example, in legal cases where liability needed to be established, advocates sought to understand the bureaucratic structure and chain of events through which employers became aware of AI tools, what kinds of information they were aware of, and when they reported errors (or not). In L8’s case, understanding funding sources of a predictive policing algorithm opened up avenues to contest the tool by pressuring the funder to cut off funding. Outside of understanding the development and deployment context, participants needed information about how harm was felt by affected individuals. This information was often used to establish injury or publicly share stories that illustrated the impact of AI tools.

4.2 Creative Workarounds for Information Access

Many layers of AI opacity made the information work life cycle for contestation challenging. A strength of the diversity of training and background of participants we spoke to was the corresponding diversity and innovation of their approaches. Across professions, participants drew from existing strategies and crossed disciplines to devise a wide range of workarounds that relied on creativity, extensive effort, or some amount of luck to achieve their informational and strategic goals. These included building relationships and solidarity networks—especially in ways that went beyond a baseline level of collaboration, directly engaging with AI systems and developers, and forging effective open records strategies. These workarounds were connected, with relationships and solidarity often driving the latter two.

4.2.1 Building Relationships and Solidarity. Effective contestation was often driven by participants going above and beyond norms of collaboration, characterized by effortful, creative, or atypical relational labor that strengthened their capacity for information work. Some examples of exceptional relational work centered around connecting with individuals harmed by AI at the individual level. While the amount of energy to devote to research versus interviews is a personal decision among journalists, J1 got creative in order to connect with individuals facing algorithmic employment termination, searching on social media support groups for potentially impacted gig workers. They invested considerable time and effort to conduct multiple, iterative interviews with these affected individuals, explaining, “I know that these things are very personal—I wouldn’t tell these sorts of very specific, harmful things that happened to me after five minutes of talking to someone I had never met.” Through these practices, J1 learned that affected individuals were thorough and capable data gatherers, explaining that “...people who feel like they’re being harmed, they usually have some sort of documentation, like the screenshots of things”, supporting harm discovery and prompting further information-seeking. A similar strategy to support information work unfolded when R1 contested the displacement of licensed healthcare providers with algorithm-assisted unlicensed clerks. In this case, R1 leveraged labor organizing networks to go beyond working with union members who were displaced to also connect with the clerks, who were represented by a different union. Building these relationships enabled R1 to solicit information about the tool’s purpose and bureaucratic context, like tool documentation and details about the human-in-the-loop process. These connections also supported sensemaking around the multiple harms occurring—reduction in the quality of patient care, displaced providers, and clerks’ anxieties around new and demanding responsibilities they felt were out of their purview.

Building connections with other community and advocacy stakeholders also supported participants’ information work. J2 described working at a media outlet in the Global South where a loose, informal network of local colleagues were able to support harm discovery around hate speech and presumed automation in reporting mechanisms. These colleagues were able to leverage their understanding of colloquial language and slang to identify violations of moderation rules that may have otherwise gone unnoticed by foreigners to the region, such as J2, using typical translation tools.

For lawyers and policy advocates, some level of multi-stakeholder collaboration was baked into the job, as described above. Efforts to engage in interdisciplinary collaboration past these standards differentiated participants like L8. At the start of contesting a predictive policing system deployed in a public school system, L8 and their colleagues invested heavily in community connections, explaining that for “*maybe about the first month or two, we were really just kind of following the school board meetings, following the local reporting, listening to, like, kind of following social media for some of the local civil society groups that we had relationships with.*” While this had a clear impact on their sensemaking around harm, these connections also came with advice on how to effectively navigate local politics and actors in the small, racially segregated municipality the predictive policing system was deployed in. These media, community, and civil society relationships were not just sources of information for L8; they became partners in a coalitional contestation strategy. L8 underscores the transformative, essential nature of this kind of interdisciplinary collaboration:

“...what I found from this work, it really changed the way that I think about legal advocacy. It changed, kind of, the course of my career to a certain extent, because I was very much gung-ho on, like, you know, litigation and thinking about courtroom-based advocacy. But this really showed me, especially if you’re thinking about legal advocacy in the tech justice context, it just requires so much more creativity, and really, like, all these interdisciplinary things, I’ve had to embrace for the rest of my career.”

The power of coalition had a similarly transformative effect on P1, a policy advocate trying to prevent harms like automated deplatforming of gig workers. P1 explained that “*my thinking is very much like I need to be evidence-based,*” but working with unions reframed the idea that top-down transparency is a strict prerequisite for justice: “*What I’ve learned from labor is like, it doesn’t matter... Like we know this is a problem, right? We’re just gonna ask for the moon, and we’ll land somewhere else.*”

Exceptional interdisciplinary relationship-building also involved proactively skill-sharing about AI contestation, which in turn could have positive downstream effects on participants’ information work. For example, A2 shared that when their manager attended a legal conference and presented information about algorithmic harms, there were “*many people in attendance who were like, Oh, my God! I’ve been dealing with something that is familiar like this. Seems like it would explain something that I’ve been dealing with.*” This led to A2 connecting with a lawyer contesting a specific instance of algorithmic allocation of state-administered health benefits. This opportunity was atypical of the cases discussed by our auditor participants, where audits would often focus on high profile platforms or genres of tools, arising from auditors’ knowledge of AI related harms at a societal level. Subsequent connections with local healthcare providers enabled A2 to collect data at scale and conduct a crowdsourced audit of the tool, demonstrating the harms in practice. Similarly, L7 describes driving a robust advice network for lawyers. In solidarity with broader AI contestation struggles, with no immediate benefit to their own cases, L7 would present at

public defenders' offices on information work strategies for effective AI contestation, including sharing template legal motions for challenging algorithmic information and checklists for information requests to make during the contestation of AI tools.

4.2.2 Direct Engagement with AI Tools and Deployers. Many participants embraced strategies that pushed the boundaries of information work typical of their profession, overcoming opacity of AI use and documentation by directly interfacing with AI tools and developers to meet their informational and strategic objectives. Matched-pair testing is a common practice in civil rights litigation—it was historically conducted with human “testers,” where individuals of different protected classes engage with a service and document their experiences. In evaluating platform-mediated ad discrimination, participants accessed information by purchasing tool access and acting as an employer, digitizing a historical practice to collect data at scale, which they could then analyze with technical experts. Although non-technical participants would often contract technical experts, some also described efforts to document the user experience of tools themselves. For example, to detect suspected discrimination in automated review on a social media platform, J2 took screenshots of different attempts to report online harassment. L3 used the same strategy to understand discrimination in the availability of agents to view real estate listings. These participants described using simple tools like spreadsheets to collate information and metadata from screenshots and identify trends or potentially important algorithm features.

In some cases, participants accessed information about AI systems by directly engaging with their developers, a form of relational labor that depended on atypical social capital and a willingness to place themselves in the line of fire of powerful AI developers. Although reaching out to investigation targets for comment is a common journalistic practice, the journalists we spoke to went above and beyond in doing so to overcome AI opacity and enact change. For example, while investigating hiring tools, J1 leveraged professional and academic networks to get in touch with tool proprietors. Through luck and perhaps a lack of due diligence by the developers, they were invited to a business-to-business seminar about the assessment tool, where they were able to learn about the algorithm's factors, such as which keywords imply particular characteristics of job applicants. While contesting suspected automated review of hate speech reporting, J2 was able to open a direct line of communication with social media platform executives, allowing them to learn about the specific language-based policies driving the platform's strategy and to share data in support of removing overlooked instances of hate speech. J2, a white American reporting from the Global South, speculated that this was made possible by their institutional background and identity-based privilege at an earlier time when social media was nascent and US-based platform executives were more receptive to addressing harm.

4.2.3 Accessing Open Records. FOIA or open records requests are a well-established and already challenging information-seeking practice in advocacy work, but we found that open records strategies had to get creative in the face of AI opacity. A long history of operating in antagonistic and obstructionist conditions has, in some ways, honed non-technical advocates' research strategies to take on these challenges. For example, while investigating a hiring tool,

J1 pulled from their research strategy repertoire, searching for public statements from AI developers, sifting through legal databases for any cases involving said developers, and searching patent and trademark filings to understand the purposes and logic underlying their AI tools. As a result, they determined that the developer of the tool had previously engaged with a public entity. J1 then submitted a FOIA request to that public entity for emails related to the tool, allowing J1 to deduce its hyperparameters and learn that a regulatory investigation of the tool was underway. In L8's case, effective open records requests relied on relationships and solidarity, which enabled iterative, coordinated efforts to uncover the interconnected data streams powering AI. For example, after discovery in an unrelated lawsuit suggested that a police department was using predictive policing, journalists leveraged FOIA to uncover and report on the existence of an adult predictive policing program and a data-sharing agreement with the local school district, raising concerns that a second predictive policing program was being deployed in schools, targeting minors. This discovery incited a coalitional, interdisciplinary effort to request additional records and contest the school's predictive policing program, with one organization targeting the school district, while another targeted the police department. Additional local partners were consulted to help design the requests. The resulting requests and division of labor were strategic, designed to anticipate the kinds of pushback received from the police department versus the school. L8 explained how this was a successful strategy:

“It's kind of almost like a... no pun intended, good cop, bad cop strategy here, where it's like, let's actually, you know, apply the carrot over here with the school district, let's apply the stick over here with the sheriff's office. And it ended up proving to be very valuable, because we got a lot of data from the school district that we did not get from the sheriff's office, and the open records litigation was actually against the sheriff's office, because they refused to provide any information without having to go to court.”

The information obtained through these requests revealed insights not only into the AI tools, but also uncovered the role of federal funding, opening new legal and policy advocacy avenues for contesting the tool. The information was used for a wide range of purposes, including assessing potential avenues for litigation, handing off to regulators, sharing for community education, and enabling responsible media reporting.

Although record requests offered an essential workaround to AI opacity, participants also cautioned that their outcomes could be uncertain. As discussed further in Section 4.3, inappropriately scoped requests could backfire, require iteration, and exacerbate resource limitations. As L4 shared, requests could be exceptionally effortful, such as cases of “*getting a million documents and having to review them all.*” Alternatively, scoping requests narrowly could result in too little information, as L4 shared:

“...sometimes you scope your request, and it's not effective, or what you get is still not responsive because they are playing games, and you don't always know that they haven't given you the full response...”

While these issues are characteristic of records requests in any domain, they are especially tricky in AI contestation, where data logs can be prolific, and scoping must navigate the opacity of AI use, purpose, oversight, logic, testing, and more. On the other hand, the prolific nature of record requests occasionally led to poor redaction. L4 shared a lucky instance where poor redaction revealed that AI deployers were previously aware of an error that caused their tool to cut benefits for certain populations. Media reporting of that error pressured the deployers into rectifying it.

4.3 The Political Economy of Information Work

The resource constraints and circumscribed social and political capital typical of advocacy meant that participants engaged in AI contestation and information work were at a significant structural disadvantage relative to developers and deployers. The conditions of information work included facing obstructionist tactics, limited resources for conducting information work, and governance structures that prioritized the interests of deployers and developers. These disadvantages meant that AI contestation carried significant risks for affected individuals and their advocates.

4.3.1 Navigating Bureaucracy and Adversarial Tactics. Information work involved confronting convoluted bureaucratic processes and intentionally obstructionist tactics. For example, when supporting clients contesting suspected automated insurance claim denials, R2 described having to navigate broken links and communicate extensively with call centers to access claim information that patients were entitled to. Additionally, due to a flaw in how insurance is governed, insurance companies—who issued the decisions being contested—were allowed to determine whether clients qualified for access to independent third-party appeals.

Similar adversarial scenarios unfolded in legal contestation of AI. When L9's clients faced cuts to public benefit funds for their health care, L9 inquired how funding amounts were determined. The government department administering the program refused to clarify, explaining that the funding algorithm was a trade secret. Through the subsequent legal challenges, it was revealed that the trade secret claim was “preposterous,” because the algorithm was simply regression analysis in an Excel spreadsheet developed by “a government bureaucrat who had, like, you know, some statistical training in college or something.” Discovery was also lengthy, labor-intensive, adversarial, and expensive. As L1 explained, developers or deployers who are supposed to provide information could attempt to obstruct the process:

“There’s going to be a lot of semantics, linguistics battles, like, ‘How did you actually define the thing you want me to turn over and what time frame?’ And blah, blah, and like, ‘Judge, they’re asking for too much. This is an unreasonable burden on us to ask us to produce everything from such and such period.’ And it’s [our response is] like ‘No judge, like, they’re being obstructionist, like, we need you to order them to comply,’ [then we] go back and forth 10 times over six months before they actually produce anything... You’re going to fight over whether every single production is actually complete and whether this thing should actually be privileged or

withheld. Or whether they’re even being honest about saying this is the universe of documents that exists.”

Such dynamics are typical of any discovery process, but when AI tools are involved, obstruction can be amplified by sharing highly technical and illegible information. For example, L7 described how a developer shared code:

“There were a couple of previous cases where they’ve been ordered to disclose the code, and in the previous cases they had delivered it—I’m like not making this up—on an iPad in a series of file directories...[when the expert] got the iPad, he’s like, well, first of all, this stuff isn’t run on an iPad. So, what a weird way to give me this information! But it was like, folders. It was like, folder trees where he had to keep opening each individual one, and then like, manually reviewing the code.”

Open records requests were met with similar obstructionist tactics, sometimes requiring litigation over noncompliance. Beyond the scoping issues discussed in Section 4.2.3, participants reported widespread practices that enabled noncompliance. For example, workers at public agencies might not take meeting notes or use private cell-phones to avoid creating records. L4 speculated that these issues stemmed from intentional defunding of state agencies that left them resource-constrained; evading FOIA may not be a sign of “acting with malign intent, they’re just like, ‘Look, we gotta serve people who need benefits, and you asking us for this FOIA request is somehow taking us away from doing that.’” An alternative explanation was active hostility to the notion of government transparency, where “You have elected officials sometimes putting pressure on agencies to be non-transparent” and more generally, “governments oftentimes see the people that they serve, or the advocates that represent the people they serve, as being antagonistic” (L4). Flaws in open records law also enabled obstruction. L8 encountered this while contesting predictive policing, explaining that the police department “had the race data... they just didn’t want to share it with us, and we couldn’t compel them to share it, [because of] the way that they categorize it on the back end.”

4.3.2 Information Work under Resource Constraints. Participants navigated lengthy and complex processes to conduct information work amid constraints around money, time, and expertise. Participants emphasized the financial cost of litigation, making it especially difficult to legally contest high-stakes algorithmic decisions that disproportionately affect low-income populations. By the time litigation makes it to discovery, information access increases, but as L4 explained, discovery requires significant time and money: “This is another factor that you have to think about in contestation is, how do you have the money to do discovery?” When L9 began litigating an algorithmic system cutting public benefit funds over a decade ago, they spent around \$60-70,000 contracting expert work from statisticians and domain experts in developmental disability. They emphasized how the cost of reaching the first legal settlement was outrageously high relative to the simplicity of the system being contested:

“When we think about the kind of systems that are out there being used, this is an Excel spreadsheet with some relatively simple formulas in them...like if you

compare it to some of the machine learning systems that are in use...it's a primitive automated system in the grand scheme of things, and the amount of time and money and expert analysis required to, you know, present to the court a cogent explanation as to what the problem was, three years of work and tens of thousands of dollars...[for] this first go around."

After reaching this initial settlement, the department delayed meeting the system reform terms and in parallel, procured a replacement system with indications of flaws similar to the original tool. L9 explained that details of the replacement system flew under the radar *"until several years into the development of the new thing, and so now everything is back, basically, to square one."* Contestation of this system is still ongoing—altogether, L9 expects that it's *"gonna end up being probably 19 years of litigation to get a new system."*

Outside of legal contestation, auditors also faced the prohibitive cost of tool access that limited the amount of data they could gather. For auditors who wanted to raise awareness of information that could be later acted on by legal or regulatory actors, publishing at an academic venue felt like a seal of legitimacy that made their work more likely to have downstream impact. However, A1 described how review standards were misaligned with the realities of resource constraints, *"So like, there are people [academic reviewers] who would say like, there wasn't enough data. ... It's like, well, shit, I don't have enough money [to pay for another month of tool access]. I don't. I'm a small nonprofit."* Financial constraints were compounded in settings like nonprofits with few employees, where there was often limited support for information work. Participants were overburdened, with A1 emphasizing: *"I'm just one fucking person poking from the outside. Like, this is not sustainable in many ways."* Policy advocates also faced asymmetric information sharing battles against disingenuous industry lobbyists who possessed far greater resources. As P2 explained: *"[lobbyists] just flat-out lie, it's unbelievable...like, they ran, the industry [lobby] ran ads during the [major sporting event], it was crazy."*

4.3.3 Power Dynamics in the Expectations of Information Work. In addition to shaping how information work was done, we found that power dynamics also narrowed the frame of what kind of information work was demanded in the first place, and to what end. That is, the standards of evidence that were required of participants were set by actors such as courts and public agencies, and often advantaged developers and deployers. For example, L4 saw how the introduction of AI tools to allocate state benefits changed the very nature of contesting benefits cuts. Typically, in administrative hearings, lawyers would demonstrate clients' medical needs, answering questions like *"How long does it take you to get out of bed? How long does it take you to prepare a meal?"* With the introduction of AI, argumentation had to be in terms of the algorithm's criteria, even if it was incorrect:

"...the only consideration became whether the assessment was completed accurately at the time that it was given, not how many hours of care you need...the judge early on would just say, 'No, it's irrelevant,' I think later, because he perceived that that could be a problem, on appeal would admit the evidence, but would just not

give it any weight, right, and say, 'No, the state has chosen to assess medical need through this process.'"

This captures an ongoing displacement of criteria related to people by criteria related to technology where *"The algorithm becomes so embedded in sort of the bureaucratic process that it displaces any other way of getting at the issue."*

In other cases, the information work expected of participants was so disproportionate relative to developers and deployers, that participants needed to demonstrate how tools worked even when the developers and deployers themselves could not do so. For example, when contesting public benefits cuts driven by a tool developed internally by a government bureaucrat, L9 *"took the testimony of everyone, including that guy [the developer], and, you know, asked people, 'How does this work?' And, you know, each person we asked said, well, you should ask this other person, until basically they were pointing around in circles. Nobody at the department could actually explain how their system worked in any coherent way."* L6 speculated about how widespread this phenomenon is among deployers:

"Do they even know what their models are? Many of them don't, frankly. And they're kind of just willy-nilly changing all the time. That's a huge pain point, right? Like, if you wanna even answer the question of like, 'Is this an appropriate use case for this model,' ...I'm not sure many institutions can answer that."

Participants were also frustrated about contesting tools under appeals systems that deployers did not actually have infrastructure or capacity for. For example, when contesting suspected automated denials of insurance claims, R2 encountered responses to appeals riddled with indications of more AI at play, like logic irrelevant to the claim, ostensibly hallucinated by LLMs. Similarly, L4 explained that contesting government AI deployments was plagued by a *"weaponization of the lack of government capacity."* Emblematic of this issue, L4 described how states would terminate unemployment benefits due to faulty fraud detection algorithms and issue requests for information to affected individuals, yet lacked the capacity to process the outpouring of resulting responses.

4.3.4 Information Work Risks. Accessing information, and contestation in general, required considering risks. For example, in line with past findings [75], one risk that multiple auditors shared was that scraping information could violate terms and conditions of use and even lead to legal issues. A3 explained the need to consider *"...how long could we do this until it kind of triggered policies that would kind of get us kicked off, or prevent the research?"* Information sharing could also come with risks due to the current political climate in the United States around equity issues. A1 explained that their nonprofit decided to shelve an audit investigating gender discrimination on hiring platforms because they did not want to *"point a bigger flashlight on it"* due to the *"anti-DEI crusade."*

Participants engaged in legal contestation also expressed concerns about client well-being. L4, for example, worried about unwanted media attention, where *"they could be dragging up things about your past that have no bearing on the case at all, but are just there to make you suffer,"* and depositions, *"which can be hours and hours of inconvenient and difficult, hard questions."* L4 also highlighted burdensome logistical overhead, where *"[clients] have to*

deal with a lot of paperwork, you can have to go to court hearings that take a lot of time and may be difficult for [clients] to reach.” L9 noticed similar risks when contesting cuts to public benefits funds, explaining that contesting *“adds a whole lot of anxiety and stress, and some people just give up.”* Sometimes taking on these risks did not pay off for impacted individuals—in L9’s case, *“of the... people who initially were plaintiffs that filed this lawsuit, about half of them have died in the process of waiting to get a fair system.”*

4.4 Opportunities to Support Information Work

We explicitly asked participants to speculate about what kinds of information they would ideally like access to in contestation. Many of the desires that participants articulated reiterated the information needs detailed in Section 4.1, emphasizing the importance of enshrining access to this information in policy. Other desires went beyond, indicating opportunities for additional support systems.

Some participants speculated that it would be ideal to mandate information sharing in a way that avoids filtering data and provides a more complete picture of it. For example, L1 desired statistics over the entirety of proprietors’ data, which is impossible for participants to compute given present barriers. A1 desired direct system access due to a lack of trust in deployers and developers and the slow nature of information access through open records or litigation. Due to this distrust in provided data, A1 explained that at present, *“I’m like, give me a copy of your entity relationship diagram of your databases,”* to ensure attributes and relationships are not omitted. In a perfect world, however, they would *“select star from this whole table and grab everything from you, like I just want you to tell me where it actually is.”*

Several participants advocated for documentation requirements, where deployers would be required to document certain information, whether publicly or privately, so that it can be requested when needed. This could, in theory, package and improve access to some of the salient technical and sociotechnical information needs participants articulated. L1 advocated for private assessments and public summary reporting. Both L1 and L6 emphasized that private documentation requirements would facilitate discovery, because there is a known tangible document to request, reducing obstructionist tactics to avoid sharing information at all.

Participants also desired infrastructural support or system reforms that would facilitate their work. L3 expressed the need for more technically sophisticated matched-pair testing practices given federal defunding of such initiatives:

“How to increase the technology of [matched-pair] testing, to address all the tenant background checks, the criminal background checks, all the automated kinds of checks and systems that particularly landlords and brokers are using is a challenge... because HUD is retracting its funding and its presence in the world I live in of housing discrimination.”

J2 desired mechanisms to report online harassment or hate speech at scale, as opposed to having to document and report each instance individually. As algorithm criteria increasingly displaced common sense notions of health needs, L4 sought procedural reforms that would allow appeals to health benefits cuts to center clients’ actual, physical needs as evidence, regardless of the accuracy of the AI

tool’s inputs. Also at the procedural level, L8 desired *“open source agencies,”* meaning that public agencies or departments would have expansive, standardized, public information access.

Some of the desired procedural reforms centered around AI tools’ initial conception. Through engaging in prolonged legal battles and witnessing the material harms endured by their clients, L9 came to the conclusion that *“the people who any of these systems are making decisions about have to be involved from the get-go. And the get-go means the decision whether or not to even use an automated system in the first place.”* A2 echoed a similar sentiment while explaining their desire for more meaningful avenues for public engagement with government AI deployments. They pointed to an example where government transparency in the form of sharing a benefits allocation algorithm for public comment missed the mark because *“it was transparent. But it also was like, sort of like, illegible to the people who needed to interact with it.”* A2 described the crux of the issue as an overly narrow focus on transparency of technical information, *“making a political discussion into a technical discussion.”* They advocated for prioritizing the former, where *“people interact with it [benefits allocation] as a political discussion and not just as a technical discussion.”*

We note, however, that participants expressed concerns about how these ideal proposals would work in practice. Participants worried about audit-washing and doubted that corporations would allow open-ended audits. To this end, J1 shared examples of instances where they witnessed audits buried because of corporate influence. P1 echoed these concerns, explaining that companies *“make this compliance theater. You’re not gonna get meaningful answers.”* Multiple participants suspected that mandating information access would be met with pushback from companies claiming that information is proprietary. L6 also expressed skepticism over the feasibility of public disclosure proposals due to both privacy and government capacity issues: *“Some of the state laws I’ve seen would require entities to submit information to an [attorney general] or some other state body. That does not sound effective to me. I mean, it sounds like a nightmare, because... the documentation, if you do it correctly, is like way too detailed and voluminous.”* Some participants also expressed pessimism around the future of contestation work in general due to the current political climate shifting commitments away from civil rights, combined with the underresourced nature of state regulatory agencies.

5 Discussion

Our findings highlight the ways in which opacity and the political economy of AI is making traditional forms of advocacy work more challenging and labor-intensive, while posing risks to advocates from technical and non-technical backgrounds alike. After further unpacking the interconnections between AI contestation within different professions and the political economy of information work, we discuss two main takeaways: the limitations of transparency policy approaches in supporting contestation, and the central role of relationships and coalition for overcoming opacity and the other burdens of contestation.

5.1 Situating AI Contestation in the Political Economy of Information Work

Many of our participants, namely journalists, lawyers, policy advocates, and researchers, came to AI contestation as a result of the introduction of algorithms into unjust practices they had already been working to combat. Our findings thus contribute to research drawing attention to community-based strategies for contestation [40, 50, 60], and affirm arguments that there is much to learn from traditional forms of advocacy for AI contestation in the present [59]. The efforts of our participants, who have long been doing advocacy work before algorithms proliferated in public life, reveal insights that we might otherwise miss through a focus on only technical approaches to the information work of contestation: successful contestation can be deeply relational, relying on situated knowledge and coalitions to determine which information is most compelling for a given contestation strategy, how to obtain it, and what actions can achieve accountability while mitigating risks.

We also analyze these advocates' perspectives alongside that of third-party auditors working in a nonprofit setting. Prior work typically has studied different kinds of advocacy work, especially litigation [49, 70] and auditing [34, 44, 75], separately. Prior work has also discussed how a sole focus on professional and technical auditing, especially that led by corporate actors, can narrow the type of expertise and objectives privileged in evaluations of AI systems [18, 20, 63, 94]. By analyzing diverse work practices together, we can see how there are many shared struggles because of the aligned goals of justice and community-centered contestation, rather than solely reporting results or improving AI tools.

Analyzing participants' shared struggles around information work, we observed that while there have always been obstacles navigating bureaucratic opacity and uncooperative agencies, there is something distinctly challenging about the information work of AI contestation in a world shaped by political and legal structures that privilege the interests of AI deployers and developers. The attributes that make AI so enticing to deployers and developers—the efficiency and perceived objectivity of automated decision-making—have had opposite effects on the process of contestation. From the so-called efficiency standpoint, AI adoption lowers the organizational costs of making austerity decisions without improving the organization's overall capacity to process information, thus making those decisions more opaque and labor-intensive to contest and appeal. Algorithms are adopted by organizations in part for their perceived objectivity [79], but in practice, automation simply displaces the values and discretion inherent to decision-making to a less visible (and thus more difficult to contest) location [23].

Further compounding this problem is the growth of public-private partnerships in AI development and deployment [80]. This growth means that information barriers related to private sector "trade secret" claims are infiltrating the public sector, making advocacy work there more challenging. Meanwhile, these claims may not even be legitimate, as we saw in our findings. The lack of open communication about these partnerships also means that opportunities to request open records about private AI deployments from the public entity's side, a strategy highlighted in our findings, are not immediately visible. On the other end of the equation, we found that algorithm deployers themselves may have limited understanding of

their own tools, making it difficult to determine who is responsible when an algorithm does not function as intended [86].

Similar to arguments about the disempowering nature of data-driven activism [10, 35], we also see how information use, not just access, ends up on the terms of those who hold power. Our findings suggest that at least within public benefits, AI is changing the target of contestation in traditional advocacy work, limiting the strategies that advocates can successfully invoke in challenging the status quo. In some cases, contestation is no longer about whether needs are being met or how resources should be allocated (a political question), but about whether AI assessments are accurate and completed correctly (a technical question).

Within this landscape, our findings on the burdens experienced by participants and the communities they served are a stark contrast to the relative power and privilege of AI deployers and developers. To this end, participants had to respond to obstructionist tactics, manage financial and time constraints, and accept personal risks in accessing data or pursuing litigation. Still, we also see how when connections among advocates are forged, contestation strategies become stronger. Below, we unpack what these dynamics mean for the effectiveness of transparency policies for enabling contestation, and how coalition-building is an important path forward.

5.2 Misalignment between Information Work Realities and Transparency Policy

By applying the lens of information work to our findings, we were able to disentangle the differences between information access and information use, and identify their associated social and emotional worlds. These findings, coupled with evidence that overcoming technical opacity was just one of the many challenges of information work, contribute to the body of literature pointing to the insufficiency of transparency [26, 32, 36, 103]. We highlight disconnects between popular transparency policy strategies in the US and information work realities, not to suggest that these strategies be discarded altogether, but to caution against a myopic focus.

In the US, federal AI policy has been unstable and patchy [8]. Recent policies have made access to government information through FOIA more difficult [65, 76] and, in line with previous findings [70] and our own, legal challenges to AI-related harms have often relied on non-algorithm-related laws in order to achieve legal standing. Amid gaps left by the federal government, several proposed, failed, or enacted AI-related bills have emerged at the state level [17, 28, 74]. Representative of these efforts, the California Consumer Privacy Act, passed in 2025 and set to be enforced in 2027, mandates submitting risk assessments for automated decision-making technology in high-stakes domains to the California Consumer Protection Agency [6]. Several other states [17, 28, 74] have considered versions of an assessment disclosure-based model [88], sometimes with limited success [101]. The theory of change behind these approaches is that assessment disclosure disincentivizes the deployment of harmful AI. A second underlying assumption is that if issues slip through the cracks, disclosures indicating these issues will prompt regulatory agencies with enforcement power or affected private actors to contest harms of AI [101], whether through non-legal pathways or through legal means when a private right of

action³ exists. Outside of disclosure-based approaches, in a separate effort, US-based policy advocates have proposed inverting civil legal procedure, or “burden shifting,” to place the burden on defendants (e.g., developers and deployers) in AI litigation to prove that their system did not harm plaintiffs. This aims to resolve the current paradox in which plaintiffs must first demonstrate how an AI system is unlawful before entering the legal discovery phase needed to access information supporting such a claim [25]. Our findings suggest that the assumptions underlying these proposals, regarding the impacts of information disclosure and the challenges of legal contestation, are misaligned with the realities of information work.

First, the premise that assessment disclosure disincentivizes the deployment of harmful AI is undermined by our findings. Participants in our study were intimately aware of adversarial tactics employed by deployers and developers. In situations where they were able to compel information disclosure, advocates encountered a level of incompetence or malicious compliance, such that the requested data was either low quality or non-existent. These experiences suggest that, at present, AI deployers and developers are far from having the capacity to engage in effective assessment. This is in line with findings from prior work on poor data quality in open records request assays [85] and null compliance with auditing mandates [101]. Harmful AI deployments cannot be disincentivized if there is no internal capacity to determine whether AI is harmful. Given these experiences, participants expressed suspicion towards information shared by companies or developers, undermining the legitimacy of assessment disclosure. They felt that companies were likely engaging in “performative transparency” [32], or sharing the minimum needed to legally comply. As Corbett and Denton [32] explain, power asymmetry is inherent to information-sharing; developers or deployers sharing information will do so in accordance with their priorities of self-preservation. Our findings affirm that meaningful disclosure is not just about the information being shared, but also about how, by whom, and under what social conditions. Similarly, Rao et al. [83] found that data access needs to be paired with trust to effectively contest AI, in the context of the crowdsourcing infrastructure they built. Our participants’ collaborative approaches to information-seeking (including through a crowdsourced audit) support Rao et al. [83]’s argument and design implications. The premise that *internal* assessment will disincentivize harm is also at odds with participants’ emphasis that the burdens of information work could be ameliorated if affected individuals—*external* stakeholders—could meaningfully participate in the design of AI tools in the first place. The value of multi-stakeholder involvement and participatory design is well established in HCI [37, 96], and this finding presents an important connection: supporting the uptake of participatory design can be an indirect way of reducing the labor of contestation and information work.

Second, the assumption that assessment disclosure would prompt contestation is also at odds with our findings. In practice, contestation was often prompted by material harm discovery. For participants engaged in client-facing roles, where AI was a new encroachment, the role of AI in these harms was rarely legible to clients, complicating participants’ ability to triage and strategize around situations. Furthermore, even when algorithm details were accessed

(e.g., through open records or discovery) or disclosed publicly (e.g., via public comment), they were not always legible to affected individuals or advocates, and often required technical collaboration to unpack. Combined with the likelihood of performative compliance discussed above, it is unlikely that the public disclosure of technical details would adequately address the legibility of AI-related harms. This implication directly aligns with Calacci et al. [26] and Dalal [36]’s work in the context of gig worker advocacy, which argues that public data disclosures are unlikely to prompt change due to illegibility of harms to advocates and gig workers. Accordingly, we echo their recommendation for building infrastructure for harm discovery, like data collection and analytics that affected individuals and advocates can execute themselves. Besides these disconnects to the realities of harm discovery, it was often sociotechnical information, like how harm was felt, that made for effective contestation—which is outside the purview of any possible company disclosure. Finally, assuming that assessment disclosure would prompt contestation ignores the risks of contestation. Contestation comes at risks to advocates and clients, and could politically backfire. Sometimes, there are not enough resources to make contestation viable, and client harm reduction must be prioritized. Participants were also skeptical that regulatory agencies would lead contestation due to underfunding and political priority shifts.

Relatedly, burden shifting to facilitate discovery access does not account for all the challenges of legal information work. Burden shifting would not address the aforementioned issues of trust in information, the challenges of harm discovery in the first place, or the costs of litigation that lead clients to choose harm reduction. Even if litigation is pursued, our findings show that discovery is lengthy and costly, met by common adversarial tactics, and uniquely amplified by the complexity of algorithms, like the delivery of highly illegible and unexecutable code. Once past discovery, over the course of legal contestation, participants were held to standards of information sensemaking about how tools worked that deployers and developers themselves could not meet. It is not obvious how burden shifting would change this double standard.

An important corollary to the notion that popular transparency proposals may not address the breadth of challenges advocates face is that many of the participants we spoke to employed creative workarounds to circumvent a lack of transparency to sensemake about AI logic, harms, and use. We draw attention to these workarounds not to suggest that opacity is acceptable or that these workarounds are not laborious, but to draw contrast with the adversarial tactics, resource constraints, power dynamics, and material risks these transparency proposals do not address, which participants found harder to navigate.

5.3 Strengthening Community and Collaboration for AI Contestation

The lens of information work calls for us to understand how information processes are situated in social processes. Drawing on this line of inquiry, we found that while multiple strategies addressed the opacity of AI, it was investing in connections, organizing, and solidarity that was often key to building power when facing *any* of the burdens of contestation we identified: opacity, adversarial

³A private right of action enables individuals or organizations to file civil lawsuits.

tactics, resource constraints, power dynamics, and risks. In recommending strengthening community and collaboration, we join calls from prior work advocating for greater cooperation, shared infrastructure, and organizing towards AI accountability [29, 40, 75] and point to strategies and models for doing so in the context of information work and contestation.

Relational labor and solidarity helped reduce the burdens of information work by redistributing them across a broader coalition. They also enabled knowledge sharing and collective strategizing that increased procedural efficacy and the likelihood of a positive outcome. Recall, for example, that in an open records campaign organized in coalition, advocates were able to pool resources, deriving benefits across several different aspects of the process. This pooling parallelized time costs while reducing the financial and labor costs to any individual organization. It also reduced the risks to individuals impacted by prolonged contestation. Finally, strategically dividing the requests across the coalition worked to minimize opportunities for typical adversarial encounters. In other instances of legal contestation, connecting across sectors to the media pressured deployers to abandon the tool, in line with findings by Johnson et al. [50]. This interdisciplinary approach could circumvent unfair power dynamics in the expectations of information work when pursuing a purely legal contestation strategy. Relationship building also facilitated harm discovery and sensemaking, overcoming opacity to build strong arguments, reducing the risk of an unsuccessful outcome. While some degree of collaboration was commonplace (such as contracting technical experts, though itself a resource cost), the aforementioned interdisciplinary, coalitional, and connection-powered strategies were atypical, indicative of advocates' creativity, effort, and some amount of luck. To better support the information work of contestation, we argue for making the atypical typical.

There are multiple strategies that could strengthen infrastructures for collaborative and interdisciplinary information work. Given the insufficiency of transparency, we previously echoed recommendations [26, 36, 83] to build trusted infrastructures for harm discovery, information-seeking, and sensemaking. These infrastructures could and should support interdisciplinary multi-stakeholder collaboration; the systems built by Calacci et al. [26], Dalal [36], and Rao et al. [83] are promising models for this work, supporting interconnected information work processes across affected workers, frontline union organizers, and union lawyers contesting algorithm-based pay and deactivation harms. It may be useful to think of these infrastructures as boundary objects [22, 66], which enable mutual understanding among diverse stakeholders; designing with this in mind could ensure that information is structured in a way that could be legible to advocates across domains and useful for multiple advocacy goals.

In designing interventions to strengthen solidarity, we must also avoid introducing additional epistemic burdens to already overstretched and underresourced advocate communities. One strategy could be to iterate on practices already integrated in these advocacy communities. Our findings highlight opportunities for improved resource-sharing in ways that are accessible, like sharing checklists of information needs in litigation, templates for open records requests, and risk-aware litigation pathways and strategies to compile more knowledge in this space. Successful models for accessible resource sharing include the Benefits Tech Advocacy Hub [2], which

collates resources on fighting AI-driven benefits cuts, The Markup, in their reporting on how they access data [95], and FOIA repositories [3]. Adapting lightweight “know your rights” strategies, like red cards in immigration justice [4], could support raising awareness towards AI harm discovery.

Beyond increasing support for the *processes* of interdisciplinary collaborative information work, our findings suggest a need to build community, encouraging opportunities to *act* in coalition. For example, having more spaces for convening and making sense of emerging technologies within legal and labor organizing networks may be beneficial. This also brings up the question of how computing researchers might make their expertise available as and when needed; models such as technology clinics within university settings [5], if built in coalition with community organizations, can offer pathways to linking with technical expertise.

5.4 Limitations

This study focused on the US context. Through this focus, we miss opportunities to draw lessons from comparing across the Global South and North [61]. The US is also not representative of policy landscapes across the Global North. Especially in the contexts of AI and policy, asymmetric scholarly attention can introduce or entrench first-mover advantages, promoting the uptake of policies or interventions ill-suited across geographies [47, 53]. To this end, we emphasize that different socio-political ecosystems could produce very different dynamics around information work and transparency policy. Within the geographic context of our study, our insights were also shaped by the scope and diversity of our participant population. We did not interview individuals directly impacted by the harms of AI, who likely engage in information work themselves. Additionally, our sampling strategy was shaped by our personal networks and participants' institutional visibility, meaning that the perspectives of advocates engaging in AI contestation in lower-resourced contexts may be underrepresented. Researching the information work of contestation in cross-cultural contexts and with wider participant populations are important future directions.

6 Conclusion

In this paper, we examined how advocates gather information to contest AI systems. We found that, with increasing automation in high-stakes domains, AI contestation has become a key part of social justice advocacy more broadly, and that the political economy of AI is complicating the work advocates are doing to address harm. We found that information access was crucial for effective contestation, and advocates used creative workarounds in response to the opacity of this information, including informal networks of collaboration, new forms of investigative work, and coordinated open records requests. However, advocates faced substantial additional barriers and risks to information work—risks of challenging the economic interests of the tech industry, obstruction posed by bureaucratic systems, limited resources common to advocacy work, and a hostile political climate for such work. We discussed why these issues are unlikely to be resolved merely by legally enshrining transparency, which is the target of many popular proposals for AI regulation in the US. By surfacing the social processes core to advocates' information work, we found that we must go beyond legal

mandates to support solidarity across the diverse advocates engaging in contestation. We ended by discussing supporting information work through amplifying collaboration among advocates.

Acknowledgments

We thank participants for taking the time to share their experiences and referring us to others. We thank Albert Fox Cahn and colleagues at the Surveillance Technology Oversight Project (STOP) for advising on early stages of this project and supporting initial recruitment. We thank Angela Jin and Nari Johnson for their contributions to early conceptions of this project. Finally, we thank the anonymous reviewers whose feedback enriched this work.

References

- [1] [n. d.]. Artificial Intelligence Incident Database. <https://incidentdatabase.ai>. Accessed: 2026-02-04.
- [2] [n. d.]. Benefits Tech Advocacy Hub. <https://www.btah.org/>. Accessed: 2026-02-04.
- [3] [n. d.]. The FOIA Project. <https://foiaproject.org>. Accessed: 2026-02-04.
- [4] [n. d.]. Red Cards / Tarjetas Rojas. <https://www.ilrc.org/redcards>. Accessed: 2026-02-04.
- [5] [n. d.]. Technology Law and Public Policy Clinic | UW School of Law. <https://www.law.uw.edu/academics/experiential-learning/clinics/technology-law>. Accessed: 2026-02-04.
- [6] 2018. California Consumer Privacy Act (CCPA). <https://oag.ca.gov/privacy/ccpa>. Accessed: 2026-02-04.
- [7] Latifah Abubakar, Tulsi Shrivastava, Samantha A Whitman, Kathleen H Pine, and Myeong Lee. 2024. Understanding Information Managers: A Thematic Analysis of Information Challenges of Disability Service Providers. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*. 629–634.
- [8] Sacha Alanoca, Shira Gur-Arieh, Tom Zick, and Kevin Klyman. 2025. Comparing Apples to Oranges: A Taxonomy for Navigating the Global Landscape of AI Regulation. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. ACM, 914–937. doi:10.1145/3715275.3732059
- [9] Kars Alfrink, Ianus Keller, Gerd Kortuem, and Neelke Doorn. 2023. Contestable AI by design: Towards a framework. *Minds and Machines* 33, 4 (2023), 613–639.
- [10] McKane Andrus, Sucheta Ghoshal, and Sayamindu Dasgupta. 2025. From Data Activism to Activism in a Time of Data-Centrism: Affirming Epistemological Heterogeneity in Social Movements. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–32.
- [11] McKane Andrus, Elena Spitzer, Jeffrey Brown, and Alice Xiang. 2021. What We Can't Measure, We Can't Understand: Challenges to Demographic Data Procurement in the Pursuit of Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (Virtual Event, Canada) (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 249–260. doi:10.1145/3442188.3445888
- [12] Blair Attard-Frost. 2023. AI Countergovernance. <https://www.midnightsummag.ca/ai-countergovernance/>. Accessed: 2026-02-04.
- [13] Blair Attard-Frost. 2025. Transfeminist AI governance. *First Monday* 30, 4 (Apr. 2025). doi:10.5210/fm.v30i4.14121
- [14] Marianne Aubin Le Quééré, Ting-Wei Chiang, Karen Levy, and Mor Naaman. 2022. Information needs of essential workers during the COVID-19 Pandemic. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–19.
- [15] Agathe Balayn, Yulu Pi, David Gray Widder, Kars Alfrink, Mireia Yurrita, Sohini Upadhyay, Naveena Karusala, Henrietta Lyons, Cagatay Turkyay, Christelle Tessono, Blair Attard-Frost, and Ujjwal Gadiraju. 2024. From Stem to Stern: Contestability Along AI Value Chains. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing (San Jose, Costa Rica) (CSCW Companion '24)*. Association for Computing Machinery, New York, NY, USA, 720–723. doi:10.1145/3678884.3681831
- [16] Jack Bandy. 2021. Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 74 (April 2021), 34 pages. doi:10.1145/3449148
- [17] BCLP. [n. d.]. US STATE-BY-STATE AI LEGISLATION SNAPSHOT. <https://infographics.bclplaw.marketing/ai-legislation-tracker-table/>. Accessed: 2026-02-04.
- [18] Juana Catalina Becerra Sandoval and Felicia S Jing. 2025. Historical methods for ai evaluations, assessments, and audits. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 1371–1386.
- [19] Clare Birchall. 2018. Interrupting transparency. In *Transparency, Society and Subjectivity: Critical Perspectives*. Springer, 343–367.
- [20] Abeba Birhane, Ryan Steed, Victor Ojewale, Briana Vecchione, and Inioluwa Deborah Raji. 2024. AI auditing: The broken bus on the road to AI accountability. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. IEEE, 612–643.
- [21] Hannah Bloch-Wehba. 2019. Access to algorithms. *Fordham L. Rev.* 88 (2019), 1265.
- [22] Geoffrey C Bowker and Susan Leigh Star. 1999. *Sorting things out: classification and its consequences*. MIT Press. Cambridge MA (1999).
- [23] Sarah Brayne and Angèle Christin. 2020. Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts. *Social Problems* (2020). doi:10.1093/socpro/spaa004
- [24] Eleanor R. Burgess, Madhu C. Reddy, Andrew Davenport, Paul Laboi, and Ann Blandford. 2019. "Tricky to get your head around": Information Work of People Managing Chronic Kidney Disease in the UK. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3290605.3300895
- [25] Albert Fox Cahn and Nina Loshkajian. 2024. Shifting AI Litigation Burden Would Help Plaintiffs in Court. <https://news.bloomberglaw.com/us-law-week/shifting-ai-litigation-burden-would-help-plaintiffs-in-court>. Accessed: 2026-02-04.
- [26] Dana Calacci, Varun Nagaraj Rao, Samantha Dalal, Catherine Di, Kok-Wei Pua, Andrew Schwartz, Danny Spitzberg, and Andrés Monroy-Hernández. 2025. FairFare: A Tool for Crowdsourcing Rideshare Data to Empower Labor Organizers. *ACM Trans. Comput.-Hum. Interact.* (Oct. 2025). doi:10.1145/3769680
- [27] Dana Calacci and Alex Pentland. 2022. Bargaining with the black-box: Designing and deploying worker-centric tools to audit algorithmic management. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–24.
- [28] Orrick AI Law Center. [n. d.]. U.S. State AI Law Tracker. <https://ai-law-center.orrick.com/us-ai-law-tracker-see-all-states/>. Accessed: 2025-11-12.
- [29] Lauren Marietta Chambers. 2025. Beyond Big Tech: Advocacy Technologists within Mission-Driven Civil Society Organizations. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (Oct. 2025), 1–31. doi:10.1145/3757559
- [30] Le Chen, Ruijun Ma, Anikó Hannák, and Christo Wilson. 2018. Investigating the Impact of Gender on Rank in Resume Search Engines. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3173574.3174225
- [31] Danielle Keats Citron. 2007. Technological due process. *Wash. UL Rev.* 85 (2007), 1249.
- [32] Eric Corbett and Remi Denton. 2023. Interrogating the T in FAccT. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1624–1634.
- [33] J.M. Corbin and A.L. Strauss. 1988. *Unending Work and Care: Managing Chronic Illness at Home*. Jossey-Bass Publishers. <https://books.google.com/books?id=RLH8V52HqwC>
- [34] Sasha Costanza-Chock, Inioluwa Deborah Raji, and Joy Buolamwini. 2022. Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1571–1583.
- [35] Roderic Crooks and Morgan Currie. 2021. Numbers will not save us: Agonistic data practices. *The Information Society* 37, 4 (2021), 201–213.
- [36] Samantha Dalal. 2025. *Going Beyond the Gig: An Ecological Approach to Studying and Intervening in the Gig Economy*. Ph. D. Dissertation. University of Colorado at Boulder.
- [37] 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*. 1–23.
- [38] Wesley Hanwen Deng, Boyuan Guo, Alicia DeVrio, Hong Shen, Motahhare Eslami, and Kenneth Holstein. 2023. Understanding practices, challenges, and opportunities for user-engaged algorithm auditing in industry practice. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [39] Alicia DeVos, Aditi Dhabalia, Hong Shen, Kenneth Holstein, and Motahhare Eslami. 2022. Toward User-Driven Algorithm Auditing: Investigating users' strategies for uncovering harmful algorithmic behavior. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–19.
- [40] Alicia DeVrio, Motahhare Eslami, and Kenneth Holstein. 2024. Building, Shifting, & Employing Power: A Taxonomy of Responses From Below to Algorithmic Harm. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (Rio de Janeiro, Brazil) (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 1093–1106. doi:10.1145/3630106.3658958
- [41] Alyssa Edes and Emma Bowman. 2018. "Automating Inequality": Algorithms in Public Services Often Fail The Most Vulnerable. <https://www.npr.org/sections/alltechconsidered/2018/02/19/586387119/automating-inequality-algorithms-in-public-services-often-fail-the-most-vulnerab>. Accessed: 2026-02-04.

- [42] Virginia Eubanks. 2025. *Automating inequality: How high-tech tools profile, police, and punish the poor*. Macmillan+ ORM.
- [43] Marissa Kumar Gerchick, Ro Encarnación, Cole Tanigawa-Lau, Lena Armstrong, Ana Gutiérrez, and Danaë Metaxa. 2025. Auditing the Audits: Lessons for Algorithmic Accountability from Local Law 144's Bias Audits. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 29–44.
- [44] Lara Groves, Jacob Metcalf, Alayna Kennedy, Briana Vecchione, and Andrew Strait. 2024. Auditing Work: Exploring the New York City algorithmic bias audit regime. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 1107–1120. doi:10.1145/3630106.3658959
- [45] Hans Krause Hansen, Lars Thøger Christensen, and Mikkel Flyverbom. 2015. Introduction: Logics of transparency in late modernity: Paradoxes, mediation and governance. 117–131 pages.
- [46] Emma Harvey, Rene F. Kizilcec, and Allison Koenecke. 2025. A Framework for Auditing Chatbots for Dialect-Based Quality-of-Service Harms. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '25). Association for Computing Machinery, New York, NY, USA, 2025–2039. doi:10.1145/3715275.3732137
- [47] Michael C Horowitz. 2018. Artificial intelligence, international competition, and the balance of power (May 2018). (2018).
- [48] Noura Howell, Watson F Hartsoe, Jacob Amin, and Vyshnavi Namani. 2024. Reflective Design for Informal Participatory Algorithm Auditing: A Case Study with Emotion AI. In *Proceedings of the 13th Nordic Conference on Human-Computer Interaction*. 1–17.
- [49] Angela Jin and Niloufar Salehi. 2024. (Beyond) Reasonable Doubt: Challenges that Public Defenders Face in Scrutinizing AI in Court. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 467, 19 pages. doi:10.1145/3613904.3641902
- [50] Nari Johnson, Sanika Moharana, Christina Harrington, Nazanin Andalibi, Hoda Heidari, and Motahhare Eslami. 2024. The Fall of an Algorithm: Characterizing the Dynamics Toward Abandonment. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 337–358. doi:10.1145/3630106.3658910
- [51] Margot E Kaminski and Jennifer M Urban. 2021. The right to contest AI. *Columbia Law Review* 121, 7 (2021), 1957–2048.
- [52] Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. 2022. A survey of algorithmic recourse: contrastive explanations and consequential recommendations. *Comput. Surveys* 55, 5 (2022), 1–29.
- [53] Panagiotis Karkatsoulis, Pantelis Capros, Panagiotis Fragkos, Leonidas Parousos, and Stella Tsani. 2016. First-mover advantages of the European Union's climate change mitigation strategy. *International Journal of Energy Research* 40, 6 (2016), 814–830.
- [54] Naveena Karusala, Sohini Upadhyay, Rajesh Veeraraghavan, and Krzysztof Z Gajos. 2024. Understanding Contestability on the Margins: Implications for the Design of Algorithmic Decision-making in Public Services. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [55] Elizabeth Kazianus, Mark S Ackerman, and Tiffany CE Veinot. 2013. Localizing chronic disease management: Information work and health translations. *Proceedings of the American Society for Information Science and Technology* 50, 1 (2013), 1–10.
- [56] Elizabeth Kazianus, Ayse G Buyuktur, Jasmine Jones, Sung W Choi, David A Hanauer, and Mark S. Ackerman. 2015. Transition and Reflection in the Use of Health Information: The Case of Pediatric Bone Marrow Transplant Caregivers. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) (CSCW '15). Association for Computing Machinery, New York, NY, USA, 1763–1774. doi:10.1145/2675133.2675276
- [57] Elizabeth Kazianus, Ayse G Buyuktur, Jasmine Jones, Sung W Choi, David A Hanauer, and Mark S Ackerman. 2015. Transition and reflection in the use of health information: the case of pediatric bone marrow transplant caregivers. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 1763–1774.
- [58] Elizabeth Kazianus, Silvia Lindtner, Mark S Ackerman, and Joyce M Lee. 2018. Lived data: tinkering with bodies, code, and care work. *Human-Computer Interaction* 33, 1 (2018), 49–92.
- [59] Vera Khovanskaya, Lynn Dombrowski, Jeffrey Rzeszotarski, and Phoebe Sengers. 2019. The Tools of Management: Adapting Historical Union Tactics to Platform-Mediated Labor. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–22. doi:10.1145/3359310
- [60] P. M. Krafft, Meg Young, Michael Katell, Jennifer E. Lee, Shankar Narayan, Micah Epstein, Dharma Dailey, Bernese Herman, Aaron Tam, Vivian Guetler, Corinne Bintz, Daniella Raz, Pa Ousman Jobe, Franziska Putz, Brian Robick, and Bissan Barghouti. 2021. An Action-Oriented AI Policy Toolkit for Technology Audits by Community Advocates and Activists. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 772–781. doi:10.1145/3442188.3445938
- [61] Neha Kumar, Naveena Karusala, Azra Ismail, Marisol Wong-Villacres, and Aditya Vishwanath. 2019. Engaging feminist solidarity for comparative research, design, and practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [62] Michelle S Lam, Mitchell L Gordon, Danaë Metaxa, Jeffrey T Hancock, James A Landay, and Michael S Bernstein. 2022. End-user audits: A system empowering communities to lead large-scale investigations of harmful algorithmic behavior. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–34.
- [63] Michelle S Lam, Ayush Pandit, Colin H Kalicki, Rachit Gupta, Poonam Sahoo, and Danaë Metaxa. 2023. Sociotechnical audits: Broadening the algorithm auditing lens to investigate targeted advertising. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–37.
- [64] Rena Li, Sara Kingsley, Chelsea Fan, Proteeti Sinha, Nora Wai, Jaimie Lee, Hong Shen, Motahhare Eslami, and Jason Hong. 2023. Participation and Division of Labor in User-Driven Algorithm Audits: How Do Everyday Users Work together to Surface Algorithmic Harms?. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [65] Sydney Lupkin. 2025. After promising transparency, RFK guts public records teams at HHS. NPR. <https://www.npr.org/sections/shots-health-news/2025/04/03/g-s1-57888/hhs-fda-rfk-foia-public-records>.
- [66] Jonathan Lynn, Rachel Y Kim, Sicun Gao, Daniel Schneider, Sachin S Pandya, and Min Kyung Lee. 2025. Regulating Algorithmic Management: A Multi-Stakeholder Study of Challenges in Aligning Software and the Law for Workplace Scheduling. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 547–572.
- [67] Henrietta Lyons, Eduardo Velloso, and Tim Miller. 2021. Conceptualising contestability: Perspectives on contesting algorithmic decisions. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–25.
- [68] Gianclaudio Malgieri and Frank Pasquale. 2022. From transparency to justification: toward ex ante accountability for AI. *Brooklyn Law School, Legal Studies Paper* 712 (2022).
- [69] Sharan B Merriam and Robin S Grenier. 2019. *Qualitative research in practice: Examples for discussion and analysis*. John Wiley & Sons.
- [70] Jacob Metcalf, Ranjit Singh, Emanuel Moss, Emmet Tafesse, and Elizabeth Anne Watkins. 2023. Taking Algorithms to Courts: A Relational Approach to Algorithmic Accountability. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1450–1462. doi:10.1145/3593013.3594092
- [71] Mahin Naderifar, Hamideh Goli, Fereshteh Ghaljaie, et al. 2017. Snowball sampling: A purposeful method of sampling in qualitative research. *Strides in development of medical education* 14, 3 (2017), 1–6.
- [72] Safiya Umoja Noble. 2018. *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- [73] ACLU of DC. 2025. Sandvig v. Barr - First Amendment Challenge to Federal Computer Fraud and Abuse Act. <https://www.aclu.org/cases/sandvig-v-barr-first-amendment-challenge-federal-computer-fraud-and-abuse-act/>
- [74] National Conference of State Legislatures. [n. d.]. Artificial Intelligence 2025 Legislation. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation>
- [75] Victor Ojewale, Ryan Steed, Briana Vecchione, Abeba Birhane, and Imoluwa Deborah Raji. 2025. Towards AI Accountability Infrastructure: Gaps and Opportunities in AI Audit Tooling. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 815, 29 pages. doi:10.1145/3706598.3713301
- [76] House Committee on Oversight and Government Reform. 2025. Oversight Democrats Expose Trump Administration for Undermining FOIA and Preventing Transparency. <https://oversightdemocrats.house.gov/news/press-releases/oversight-democrats-expose-trump-administration-undermining-foia-and-preventing>. Accessed: 2026-02-04.
- [77] Lawrence A. Palinkas, Sarah M. Horwitz, Carla A. Green, Jennifer P. Wisdom, Naihua Duan, and Kimberly Hoagwood. 2015. Purposeful Sampling for Qualitative Data Collection and Analysis in Mixed Method Implementation Research. *Administration and Policy in Mental Health and Mental Health Services Research* 42, 5 (09 2015), 533–544. <https://www.proquest.com/scholarly-journals/purposeful-sampling-qualitative-data-collection/docview/1705641789/se-2> Copyright - Springer Science+Business Media New York 2015; Last updated - 2024-03-21; Subject's Term:NotLitGenreText - Thousand Oaks California; United States-US; New York.
- [78] Adrienne Pichon, Kayla Schiffer, Emma Horan, Bria Massey, Suzanne Bakken, Lena Mamykina, and Noemie Elhadad. 2021. Divided we stand: the collaborative work of patients and providers in an enigmatic chronic disease. *Proceedings of the ACM on human-computer interaction* 4, CSCW3 (2021), 1–24.
- [79] Theodore M. Porter. 1995. *Trust in numbers: the pursuit of objectivity in science and public life*. Princeton University Press, Princeton, NJ.

- [80] Dasha Pruss, Hannah Pullen-Blasnik, Nikki Stevens, Shakeer Rahman, Clara Belitz, Logan Stapleton, Mallika G Dharmaraj, Mizue Aizeki, Petra Molnar, Annika Pinch, et al. 2024. Prediction and Punishment: Critical Report on Carceral AI. (2024).
- [81] Inioluwa Deborah Raji and Joy Buolamwini. 2022. Actionable auditing revisited: Investigating the impact of publicly naming biased performance results of commercial ai products. *Commun. ACM* 66, 1 (2022), 101–108.
- [82] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 33–44. doi:10.1145/3351095.3372873
- [83] Varun Nagaraj Rao, Samantha Dalal, Andrew Schwartz, Amna Liaqat, Dana Calacci, and Andrés Monroy-Hernández. 2025. FareShare: A Tool for Labor Organizers to Estimate Lost Wages and Contest Arbitrary AI and Algorithmic Deactivations. *arXiv preprint arXiv:2505.08904* (2025).
- [84] Eva Rosen, Philip ME Garboden, and Jennifer E Cossyleon. 2021. Racial discrimination in housing: How landlords use algorithms and home visits to screen tenants. *American Sociological Review* 86, 5 (2021), 787–822.
- [85] Amelia Roskin-Fraze and Christine Denise Head. 2025. The Bureaucratic Challenges of Unified Data: Examining Predictive Policing Technology Design Issues Posed by Federal and State US Criminal Legal System Data Sets. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–11.
- [86] Filippo Santoni de Sio and Giulio Mecacci. 2021. Four responsibility gaps with artificial intelligence: Why they matter and how to address them. *Philosophy & technology* 34, 4 (2021), 1057–1084.
- [87] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2021. A framework of high-stakes algorithmic decision-making for the public sector developed through a case study of child-welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–41.
- [88] Andrew D Selbst. 2021. An institutional view of algorithmic impact assessments. *Harv. JL & Tech.* 35 (2021), 117.
- [89] Hong Shen, Alicia DeVos, Motahhare Eslami, and Kenneth Holstein. 2021. Everyday Algorithm Auditing: Understanding the Power of Everyday Users in Surfacing Harmful Algorithmic Behaviors. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 433 (Oct. 2021), 29 pages. doi:10.1145/3479577
- [90] Franchesca Spektor, Vera Khovanskaya, Jodi Forlizzi, and Sarah Fox. 2025. Data Tactics in Worker Advocacy Research. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW283 (Oct. 2025), 29 pages. doi:10.1145/3757464
- [91] Samuel J Stratton. 2021. Population research: convenience sampling strategies. *Prehospital and disaster Medicine* 36, 4 (2021), 373–374.
- [92] Cella M Sum, Anh-Ton Tran, Jessica Lin, Rachel Kuo, Cynthia L Bennett, Christina Harrington, and Sarah E Fox. 2023. Translation as (Re)mediation: How Ethnic Community-Based Organizations Negotiate Legitimacy. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 603, 14 pages. doi:10.1145/3544548.3581280
- [93] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Queue* 11, 3 (2013), 10–29.
- [94] Petros Terzis, Michael Veale, and Noëlle Gaumann. 2024. Law and the emerging political economy of algorithmic audits. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1255–1267.
- [95] Lam Thuy Vo. 2024. How to Run a Public Records Audit with a Team of Students. <https://themarkup.org/levelup/2024/01/24/how-to-run-a-public-records-audit-with-a-team-of-students>. Accessed: 2026-02-04.
- [96] John Vines, Rachel Clarke, Peter Wright, John McCarthy, and Patrick Olivier. 2013. Configuring participation: on how we involve people in design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 429–438.
- [97] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2024. Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy. *ACM Journal on Responsible Computing* 1, 1 (2024), 1–45.
- [98] Cedric Deslandes Whitney, Teresa Naval, Elizabeth Quepons, Simrandeep Singh, Steven R Rick, and Lilly Irani. 2021. HCI Tactics for Politics from Below: Meeting the Challenges of Smart Cities. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 297, 15 pages. doi:10.1145/3411764.3445314
- [99] Kyra Wilson and Aylin Caliskan. 2024. Gender, race, and intersectional bias in resume screening via language model retrieval. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, Vol. 7. 1578–1590.
- [100] Marisol Wong-Villacres, Neha Kumar, and Betsy DiSalvo. 2019. The work of bilingual parent-education liaisons: Assembling information patchworks for immigrant parents. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [101] Lucas Wright, Roxana Mika Muenster, Briana Vecchione, Tianyao Qu, Pika (Senhuang) Cai, Alan Smith, Comm 2450 Student Investigators, Jacob Metcalf, and J. Nathan Matias. 2024. Null Compliance: NYC Local Law 144 and the challenges of algorithm accountability. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 1701–1713. doi:10.1145/3630106.3658998
- [102] Juliette Zaccour, Reuben Binns, and Luc Rocher. 2025. Access Denied: Meaningful Data Access for Quantitative Algorithm Audits. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–31.
- [103] Monika Zalnieriute. 2021. "Transparency Washing" in the Digital Age: A Corporate Agenda of Procedural Fetishism. *Critical Analysis L.* 8 (2021), 139.
- [104] Shoshana Zuboff. 2019. The age of surveillance capitalism: The fight for a human future at the new frontier of power, edn. *PublicAffairs, New York* (2019).