

# The Fall of an Algorithm: Characterizing the Dynamics Toward Abandonment

Nari Johnson  
narij@andrew.cmu.edu  
Carnegie Mellon University  
USA

Sanika Moharana  
smoharan@andrew.cmu.edu  
Carnegie Mellon University  
USA

Christina N. Harrington  
charring@andrew.cmu.edu  
Carnegie Mellon University  
USA

Nazanin Andalibi  
andalibi@umich.edu  
University of Michigan  
USA

Hoda Heidari\*  
hheidari@cmu.edu  
Carnegie Mellon University  
USA

Motahhare Eslami\*  
meslami@andrew.cmu.edu  
Carnegie Mellon University  
USA

## ABSTRACT

As more algorithmic systems have come under scrutiny for their potential to inflict societal harms, an increasing number of organizations that hold power over harmful algorithms have chosen, or were required under the law, to *abandon* them. While social movements and calls to abandon harmful algorithms have emerged across application domains, little academic attention has been paid to studying abandonment as a means to mitigate algorithmic harms. In this paper, we take a first step towards conceptualizing “*algorithm abandonment*” as an organization’s decision to stop designing, developing, or using an algorithmic system due to its (potential) harms. We conduct a thematic analysis of real-world cases of algorithm abandonment to characterize the dynamics leading to this outcome. Our analysis of 40 cases reveals that campaigns to abandon an algorithm follow a common process of six iterative phases: *discovery*, *diagnosis*, *dissemination*, *dialogue*, *decision*, and *death*, which we term the *6 D’s of abandonment*. In addition, we highlight key factors that facilitate (or prohibit) abandonment, which include characteristics of both the technical and social systems that the algorithm is embedded within. We discuss implications for several stakeholders, including proprietors and technologists who have the power to influence an algorithm’s (dis)continued use, FAccT researchers, and policymakers.

## CCS CONCEPTS

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

## KEYWORDS

abandonment, contestation, refusal, accountability

### ACM Reference Format:

Nari Johnson, Sanika Moharana, Christina N. Harrington, Nazanin Andalibi, Hoda Heidari, and Motahhare Eslami. 2024. The Fall of an Algorithm: Characterizing the Dynamics Toward Abandonment. In *The 2024 ACM*

\*Both senior authors contributed equally to this research.



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

FAccT ’24, June 03–06, 2024, Rio de Janeiro, Brazil  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0450-5/24/06  
<https://doi.org/10.1145/3630106.3658910>

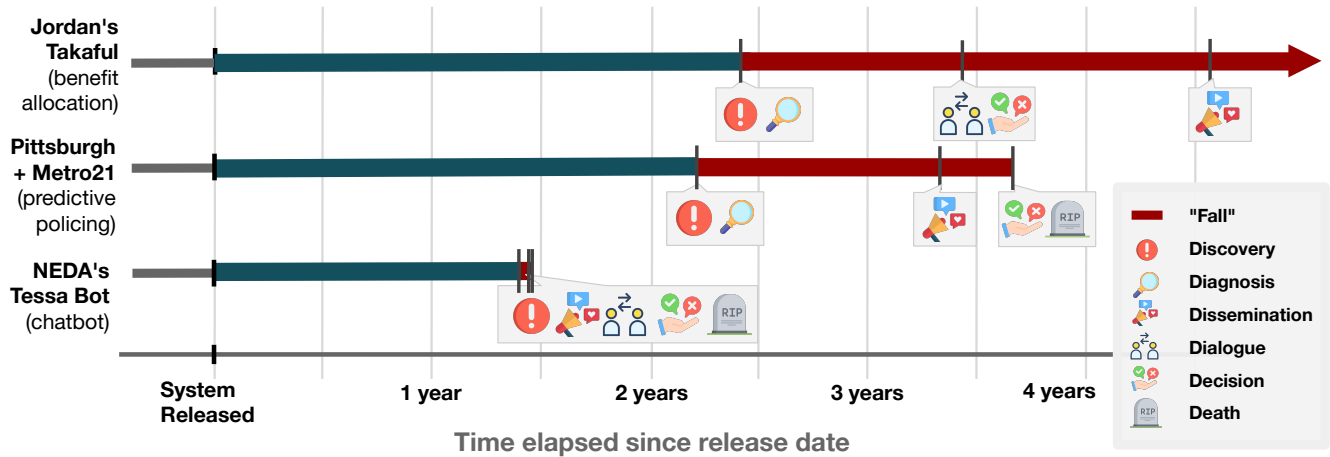
*Conference on Fairness, Accountability, and Transparency (FAccT ’24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 22 pages.  
<https://doi.org/10.1145/3630106.3658910>

## 1 INTRODUCTION

An increasing number of incidents have demonstrated how algorithmic systems can inflict harm, by (for example) upholding oppressive power structures [28, 123], reinforcing stereotypes [20, 165], or making life-altering decisions through inscrutable logic [59, 64, 104, 148]. In response to these harms, a growing body of research has explored how harmful algorithms might be *repaired* through technical or procedural interventions [26, 130, 145]. However, scholars and activists alike have argued that not all algorithmic harms can be effectively addressed or mitigated through these kinds of repairs [73, 119, 167]. This stance often arises from a recognition of the deeply ingrained nature of the corresponding harm, where mere adjustments fail to resolve fundamentally systemic issues. In some instances, resistance from impacted communities has precipitated the *fall* of an algorithmic system, in which its proprietors might choose, or be forced to, *abandon* its use.

Take, for example, a case from Arkansas: when a new resource allocation algorithm resulted in significant cuts to the number of in-home care hours that Medicaid beneficiaries would receive [87], one beneficiary wrote in with a clear request: to turn it off, and return to the old system of human-based assessment. No one in the state could offer an explanation for why this beneficiary’s in-home care hours were reduced, cuts that for many beneficiaries with chronic disabilities resulted in drastic changes to their quality of life. While three years later, the algorithm (called RUGS) was abandoned, the road to this outcome was long and tedious. Efforts to abandon the algorithm were met with resistance from state officials [173]. The eventual decommissioning of RUGS after it was invalidated by a state court [87] was a result of concerted advocacy efforts by impacted individuals, legal advocates, and media attention.

Although some campaigns to abandon harmful algorithms have been successful, other harmful algorithms continue to remain in use today. Furthermore, while some algorithms were abandoned in early ideation phases *before* they were deployed, and some were abandoned within mere hours of harm discovery, other algorithms continue(d) to inflict harm for years. Thus, open questions remain when considering the possibility of abandoning harmful algorithms as a mechanism of harm mitigation and accountability: What types



**Figure 1: Event timelines for three example algorithms [83, 91, 171].** We visualize the time elapsed to each event since each algorithm’s release date (x-axis), and label important events by tagging them with one (or multiple) phase(s) from the 6 D’s of abandonment (see legend, bottom right). We provide detailed descriptions of each event in Appendix B. Two of the pictured algorithms (the city of Pittsburgh’s predictive policing pilot and the National Eating Disorder Association’s “Tessa” chatbot) were abandoned, and had a “Death” phase. Takaful, a cash transfer allocation algorithm deployed in Jordan, was *not* abandoned and is still in use at the time of writing (denoted by the red arrow into the present). These timelines illustrate how the time elapsed between events (e.g., for the algorithm to be discovered, or be abandoned after discovery) varied considerably between the three cases shown.

of events preceded and precipitated the decision to abandon an algorithm? How did impacted communities and other actors work together to contest harmful algorithms? What technical, political, or social characteristics of real-world algorithmic deployments posed barriers for contestation, accountability, and abandonment? How can future calls for abandonment be bolstered and supported?

In this paper, we take a first step toward answering these questions by exploring the concept of *algorithm abandonment*. We define algorithm abandonment as a decision made by actors with jurisdiction over the system, to discontinue the process of developing, deploying, or using the algorithm due to its (potential) harms. The decision to abandon an algorithm can be made at any point in the algorithm’s development lifecycle [90, 157], and includes both decisions not to build a proposed technology, and decisions to stop using tools that are already built. The algorithmic systems that we explore span diverse ways of coming into being, including hand-coded rules and machine learning models. Our study of algorithm abandonment draws from literature from Science and Technology Studies which examines algorithms not just as technical artifacts mapping inputs to outputs, but rather as socio-technical “assemblages” of social relationships that enable and sustain technologies [145]. Thus, understanding abandonment requires examining the social, political, and economic systems that algorithms are embedded within [147].

In our thematic analysis of 40 different campaigns to abandon harmful algorithms in recent years, we observe that the fall of an algorithm (or “*dynamics*” leading to abandonment) often follows a consistent pattern of six iterative, non-linear phases which we term the 6 “D”s of abandonment: (1) *discovery*, (2) *diagnosis*, (3) *dissemination*, (4) *dialogue*, (5) *decision*, and (6) *death*, which we visualize for

three example cases in Figure 1. We review each of these phases and highlight both tactics that algorithm owners used to avoid scrutiny (e.g., discrediting or ignoring criticism), and tactics that other actors such as impacted individuals, advocacy organizations, journalists, or regulators used to shift power to advocates of abandonment. We additionally identify key *socio-technical factors* relevant to the outcome of whether an algorithm was ultimately abandoned (and how long it took to reach this stage), such as public awareness of and access to the algorithm, characteristics of the algorithm’s implementation, visibility of criticism, the regulatory environment, and for whom (and how) the algorithm is marketed. We conclude by discussing implications for technologists designing algorithmic systems, regulators and oversight entities, and the FAcCT community in identifying barriers to abandonment and interventions to support future calls to abandon harmful algorithms. We release our annotated case set online<sup>1</sup> as a resource to aid future research and advocacy efforts.

## 2 BACKGROUND & RELATED WORK

We contextualize our study of abandonment using scholarship that studies how different stakeholders respond to algorithmic harm. “Algorithmic harm,” as described by Shelby et al. [150], refers to “*the adverse lived experiences resulting from a system’s deployment and operation in the world*” which occurs through the interplay of technical components and societal power dynamics [28, 112, 123]. We group approaches taken to address algorithmic harm into two broad categories: (1) actions taken “from above” by organizations in positions of power (e.g., those that develop, oversee, or govern the algorithm’s use), and (2) actions taken “from below” by other

<sup>1</sup><https://njohnson99.github.io/fall-of-algorithm-database/>

stakeholders (e.g., those impacted by algorithmic harm). We adopt the language of “above” and “below” from Whitney et al. [175], who uses these terms to articulate the “*tensions between the wants, needs, and knowledge of unequally empowered stakeholders*”.

We situate our study of abandonment as an organizational decision made “from above” by algorithm owners or government oversight entities. We use the term “algorithm owner” to broadly refer to the organizations and people within them that have the power to influence an algorithm’s (dis)continued use on the ground. Our analysis of the events, or “dynamics”, that precede abandonment in practice draws from the growing body of scholarship that examines how actors “from below” organize to contest harmful algorithmic systems [52, 185]. We use the term “impacted communities” to describe the people who are affected by the use of an algorithm. Those impacted may include the direct consumers of a technology or any other people who are subject to an algorithm’s decisions [65, 121].

**Repair from above.** Much work concerning algorithmic harm explores the actions that algorithm owners might take to repair flawed technologies. The AI/ML and FAccT community have proposed numerous technical methods to promote values of “fairness” and “explainability” [55, 103, 148]. This work has led to the development of metrics [23, 133, 154, 156], methods [56, 74, 140, 168], and toolkits [22, 25, 51, 172] that practitioners can use to measure and fix specific harmful behaviors. Beyond technical repairs, human-computer interaction researchers have documented how organizational users and decision-makers engage in their own ad-hoc or improvisational repair activities when they encounter flawed technologies [61, 84, 145]. For example, social workers using AI tools to screen calls can override them to mitigate disparities in final outcomes [41, 42, 50]. Such methods can be used to correct the algorithm’s flaws, often defined in terms of narrowly scoped measures [93], while maintaining the algorithm’s operation.

In both research and policy proposals, repair has emerged as a dominant approach to address algorithmic harm from above [26, 49, 100]. Both governments and private-sector organizations have instituted policies requiring mitigation strategies to repair potential algorithmic harms [118, 126]. However, not all algorithmic harms are readily fixable, particularly in cases where critics take issue with the very existence of the system itself. Critical scholarship has scrutinized the use of algorithms in specific domains (e.g., predictive policing [95, 119] or criminal sentencing [64]), and questioned if such algorithms can ever be repaired. Instead of examining how flawed technologies might be repaired, we focus our inquiry on abandonment, which is predicated on questioning if an algorithm should exist in the first place.

**Refusal from below.** The concept of refusal, defined simply as the act of saying “no” [17], has been studied in feminist and de-colonial discourses [82, 111, 152, 159]. Today, scholars and activists continue to draw from these traditions in studying refusal of emerging technologies [44, 52, 185]. Among theories of anti-techno-solutionism, Benjamin [27] refers to indigenous refusal to analyze vignettes where research subjects refuse to engage with new biotechnologies. A group of scholars drafted the Feminist Data “Manifest-No”, a declaration of refusal to harmful data regimes and exploitative research practices [44]. There is a rich history of grassroots social

movements formed to refuse technologies, tracing back hundreds of years to the British Luddites during the Industrial Revolution [45], sparking modern adaptations of “Luddism” by workers organizing to oppose harmful algorithmic management practices [40, 115].

Refusal of harmful technologies may take the form of individual users choosing to no longer interact with harmful algorithms [52, 185]. For example, in response to Amazon’s introduction of AI-powered cameras for worker surveillance, some drivers engaged in an act of refusal by obscuring the cameras with stickers, rendering them unable to record [70]. But refusal also extends beyond individual acts into collective calls for a system to be abandoned. Understanding calls to abandon an algorithm through the lens of refusal sheds light on the possibilities that abandonment entails to “*open up possibilities for power shifts*” [38].

**Conceptualizing & contextualizing abandonment.** In this paper, we conceptualize algorithm abandonment as “a decision made *from above* (e.g., by algorithm proprietors, vendors, purchasers, or governance bodies), to discontinue developing, deploying, or using a harmful algorithmic system, usually due to calls *from below*”. We focus our study on abandonment as an *organizational* decision (rather than that of individual user decisions) as many systems are implemented in a way where opting out is difficult or impossible [4, 99, 158]. We differentiate abandonment from repair in that while repair considers how one might improve an existing algorithmic system, abandonment entails doing away with specific types of automation altogether. This conceptualization of abandonment builds upon scholarship [24, 54, 81, 85, 155] focused on “*when the implication is not to design ([a] technology)*” and advocates for technologists to adopt “*a reflective awareness for when [automation] is inappropriate or harmful*” [24].

### 3 METHODOLOGY

We adopt a case study approach to analyze documented cases where stakeholders called for abandonment in recent years. Our goal was to understand the dynamics of algorithm abandonment by studying the events that occurred throughout these calls. As such, we first gathered real-world cases of algorithms that were (called to be) abandoned, and then conducted an iterative thematic analysis.

**Case collection.** To develop our dataset, we conducted an extensive search of various sources for instances where a harmful algorithm was either (a) already abandoned, or (b) critics *called for* the algorithm to be abandoned, regardless of whether it was abandoned or not at the time of writing. Our search methodology was a dynamic, iterative process that unfolded concurrently with our analysis. This approach involved consulting a diverse group of resources, detailed further below:

- *Initial dataset:* We began with an initial set of cases that the authors were familiar with from research articles and reading lists focused on resisting and refusing algorithmic harm [35, 52, 150, 151, 185]. We also conducted web searches with keywords such as “algorithm abandonment”, “algorithmic bias”, and “algorithmic harm”.
- *AI incident databases:* We reviewed all of the reported incidents to date<sup>2</sup> in three incident databases that compiled and documented

<sup>2</sup>We conducted our review from July to December 2023, and reviewed all incidents that had been posted in each database before January 1, 2024.

real-world harms caused by algorithmic systems: (a) the AI Incident Database (AIID) [112], (b) the AI, Algorithmic, and Automation Incidents and Controversies (AIAAIC) repository [18], and (c) Not My AI [164] which focuses specifically on mapping cases of algorithmic harm in Latin America.

- *Expert consultations:* We also identified additional cases through conversations with experts who had academic or lived experiences relating to algorithm abandonment. Specifically, our research team met with members of an independent U.S.-based advocacy group founded to help impacted communities resist oppressive surveillance technologies. We also met with academic colleagues who have done research on topics relating to resisting and responding to algorithmic harm. These experts pointed us to relevant cases that they had encountered in their work.
- *International cases:* Given that U.S. and European-based incidents were over-represented in our initial case set, we extended our search to be intentionally inclusive of other regions globally. We conducted additional keyword searches on Google to identify relevant cases that occurred in countries outside of Western contexts, such as Korea [98], Argentina [21], Jordan [143], India [4], among others.

*Case set.* From the cases that met the abandonment criteria, we selected a subset of 40 as our case studies for our thematic analysis. We chose this subset based on several criteria, including (a) coverage of different domains, (b) access to enough sources to gain more information about a case to do a thematic analysis, and (c) reaching saturation in finding new themes in the analysis process. The selected cases included algorithms used in social media, chatbots, hiring, family and social services, benefits allocation, healthcare, education, property tech, and law enforcement. We chose among cases within the same domain based on the amount of public information and documentation that was available for us to construct accurate event timelines. We performed our thematic analysis in parallel to identifying relevant cases, and stopped adding new cases to the subset once we had reached saturation in finding new themes (*i.e.*, the types of events and factors relevant to abandonment). Our final case set contains 40 different algorithmic systems, 24 of which had been abandoned at the time of writing<sup>3</sup>, while the remainder were still in use. 16 out of the 40 total cases involved campaigns that took place in countries outside of the U.S. To provide context, we describe three example cases in Table 1 and a complete list in Appendix C. These 40 cases are not intended to be a comprehensive list of all algorithms that were called to be abandoned, but rather serve as a starting point to conceptualize algorithm abandonment, illustrate its dynamics, and map out its implications.

**Thematic analysis.** We gathered and reviewed public resources about each case, including news articles, research publications, social media posts, and any other relevant primary sources (*e.g.*, algorithm source code or court documents). For each case, we constructed a timeline by noting the month and year of important

events. Our goal in noting important events was not to be comprehensive, but rather to provide “good enough” coverage of the types of events leading to abandonment, a standard used in past qualitative research [184]. We reviewed sources to code metadata about the algorithm, such as its application domain, the type(s) of algorithmic harm [150], and other information relevant to whether it was ultimately abandoned (more details in Appendix A). We walk through the set of sources and resulting codes and event timeline for an example case in Appendix A.1.

We conducted a bottom-up thematic analysis [30] of case event timelines and metadata in a series of meetings among research team members. In our thematic analysis of event timelines, we paid particular attention to the sequencing of key events, the time lapsed between them, and how each event contributed to the final decision of whether the system was abandoned. Our analysis yielded six high-level types of events (or “phases”) as the dynamics of the algorithm abandonment that we describe in Section 4. To identify factors relevant to abandonment presented in Section 5, we noted differences among algorithms that were abandoned almost immediately (*e.g.*, within days of their release), versus algorithms that took months or even years to be abandoned. We organized these differences into higher-level categories that resulted in the final seven factors. We also noted different barriers that advocates of abandonment faced when interacting with algorithm owners (*i.e.*, asserting their role in this accountability relationship [117]).

**Limitations.** In our study of the events that influenced abandonment, we analyzed *publicly available* materials such as news articles or information obtained through public records requests. We speculate, however, that there were other reasons why organizations chose to abandon the algorithm beyond those named in public sources, such as pressure from company shareholders [96, 176] or financial projections [120, 136]. Future work that engages more directly with people involved in making the decision to abandon an algorithm may shed light on their motives.

## 4 THE DYNAMICS TOWARD ALGORITHM ABANDONMENT

We explore the *dynamics* of abandonment by proposing a process-oriented view of the events leading to abandonment. We visualized these dynamics for individual cases by creating a *timeline of events* (shown in Figure 1) that occurred until a system was abandoned (or until the time of writing for systems that were still in use). In our thematic analysis, we found that most algorithms in our case set were abandoned *after* they were deployed. However, our case set also includes cases of algorithms abandoned earlier in their development lifecycle. These cases provide meaningful comparison points to consider how algorithmic harms might be anticipated and prevented before a system is released.

**The 6 D’s of abandonment.** In our analysis of event timelines for the 40 identified cases, the process leading to abandonment involved events that fell into six iterative phases of (1) **discovery**, which initiated critique and calls to abandon the system by the public, impacted communities, or their advocates, (2) **diagnosis**, where these actors gathered further evidence about the system and its impacts, (3) **dissemination**, where these actors amplified their concerns to a wider audience, (4) **dialogue** between critics of the

<sup>3</sup>Some algorithms that we study were deployed in multiple different contexts (*e.g.*, PredPol was used by dozens of U.S. police departments [75]), and were abandoned in some, but not others. We count products like PredPol as a single “algorithmic system”, and count them as being “abandoned” if they were abandoned in at least one such context.

Example Cases (domain)	Description
<b>U.S. National Eating Disorder Association (NEDA)’s Tessa Wellness Bot</b> [174] (chatbot)	Shortly after NEDA fired their hotline staff, activists posted on social media about how Tessa (NEDA’s chatbot) gave harmful and triggering advice.
<b>Facial Recognition in NYC Atlantic Plaza Towers</b> [66] (facial recognition)	Residents of a Brooklyn apartment complex partnered with a legal group to resist their landlord’s plan to install facial recognition technology to “secure” access to the building.
<b>The Spanish Ministry of the Interior’s VioGén Gender Violence Recidivism Risk Assessment</b> [5] (family and social services)	The Spanish government’s risk assessment algorithm underestimated risks posed by perpetrators, withholding necessary protection services that led to adverse outcomes in domestic violence cases.

**Table 1: Three example cases from our list of 40 total case studies of real-world examples of campaigns to abandon an algorithm. See the complete list of 40 cases in Appendix C.**

system and algorithm owners, until they reached a (5) **decision** about whether to repair or abandon the system. We discuss how many abandoned systems still continued to impact peoples’ lives in the final (6) **death** phase, in which algorithm owners transitioned away from the algorithm and those impacted sought justice for past harms. Note that all six phases may occur in any stage of an algorithm’s lifecycle (*e.g.*, design, development, usage), and are still relevant for algorithms that are not yet developed or deployed. In practice, the “fall” of an algorithm often proceeds *iteratively* and *non-linearly* throughout these six phases. Similarly, not all six phases may occur for every case. Regardless, this process-oriented view provides a useful lens to examine what types of events were influential across a range of cases.

#### 4.1 Discovery

Many cases had a catalyzing event where information that was previously unknown about the algorithm was made public and initiated wide-scale critique. We group these discoveries into two categories: (1) discovering the *existence* of a previously obscured algorithmic system; and (2) discovering potential or realized algorithmic *harms*.

**Discovery of existence.** In some cases, criticism began once the algorithm’s mere existence was made public. While algorithm owners could in theory provide public notice of their intentions to develop or use an algorithm in early stages of its lifecycle, many did not provide notice until the system was already in use. For example, the City of Pittsburgh “*experimented with*” using a predictive policing algorithm for over a year before the public was made aware of its existence through the release of a public report [91]. Several other algorithms such as the New York Police Department’s contract with facial recognition software vendor Clearview AI [76] evaded critique for months due to a lack of public notice. Discoveries of an algorithms’ existence often sparked criticism about a lack of community engagement and oversight [91] and the values they encoded (*e.g.*, through their task formulation or deployers’ stated goals) [2, 77, 95].

**Discovery of algorithmic harm.** In other cases, initial critiques began after a discovery of potential or realized algorithmic harms. To understand what types of harms motivated calls to abandonment, we mapped all mentions of harms in case materials to the five

major types of sociotechnical harms in Shelby et al. [150]: representational, allocative, quality-of-service, interpersonal, and social system harms. We provide complete definitions and examples of each harm type Appendix Table 2. In summary, the algorithms in our case set spanned all five of Shelby et al. [150]’s harm categories. Algorithms deployed in similar domains resulted in similar harms, *e.g.*, chatbots could demean and stereotype [97, 98], and resource allocation algorithms could result in opportunity loss [87, 143]. Several different types of harm often co-occurred for each single algorithmic system.<sup>4</sup>

Importantly, critics identified harms that implicated more than just the algorithm’s *functionality*, or ability to work as intended [138]. In other words, simply improving the algorithm to be more accurate, *i.e.*, to reduce what Shelby et al. [150] deem “Quality-of-Service” performance disparities, would not redress all harms. A more effective surveillance technology, for example, could still contribute to civil rights violations or interpersonal harms, such as accelerating incarceration or family separation at scale [158, 169]. While many cases we considered did involve dysfunctional algorithms, the remaining taxonomy categories remind us that other harms exist too.

#### 4.2 Diagnosis

As the number of people aware of an algorithms’ existence and potential harms grew, so did their capacity to *diagnose* its potential to cause (further) harm. We group critics’ investigations into two categories: audits that analyzed the algorithms’ *outputs*, and investigations into *human* decisions (such as design choices) surrounding the algorithm.

**Examining the output: algorithm auditing.** Critics attempted to better understand the behavior of an algorithm by interacting with or probing it directly to observe its output. Many such audits involved querying the algorithm on several (real or simulated) inputs, and measuring properties of its output. Large-scale targeted audits were a highly effective tool that precipitated decisions to

<sup>4</sup>For example, San Diego’s TACIDS algorithm [158, 178] matched photographs of suspects against mugshot databases, where people with previous law enforcement encounters were more likely to be (falsely) flagged (representational harm). TACIDS was used to aid immigration enforcement efforts that separated families (interpersonal harm), and had implications for residents’ constitutional rights (social system harms).

abandon an algorithm. For example, many organizations referenced the Gender Shades audit (which measured racial and gender accuracy disparities of commercial facial recognition algorithms [33]) when they later instituted moratorium policies on the sale or usage of facial recognition technology [34].

While large-scale audits were highly effective, in several cases, critics faced significant obstacles to conduct such audits. One dimension that significantly influenced the auditability of an algorithm was access to the model, as algorithms that were administered by a decision-maker internal to the organization (*e.g.*, a social worker assessing child welfare cases [79]) could not be queried at will. Independent auditors also often lacked access to realistic and representative data from the population of interest. In some cases, the public was even unaware of what information the algorithm took as input to make decisions. For example, in Jordan, government officials refuse to release the “socio-economic indicators” of poverty that an algorithm which allocated cash assistance to citizens took into consideration [171].

To overcome these obstacles, in some cases critics worked together to collectively develop work-around strategies. One strategy to obtain access to realistic input data was to design coordinated data collection efforts from consenting people who would be or were already impacted by an algorithmic system. In Missouri, a coalition of legal aid organizations and third-party researchers partnered with home care nurses to fill out assessments to estimate a proposed home care allocation algorithm’s impact for 1200 patients [173]. In Spain, an independent nonprofit conducted phone interviews to understand the experiences of women who completed an algorithmic gender-based violence risk assessment that determined the police and legal resources allocated to them [5]. Successful data donation initiatives gave auditors access to population-level data that was critical to understanding the algorithm’s behavior.

**Examining the process: design decisions.** Algorithmic systems do not just “come about”, but are rather the product of a series of *decisions made by designers, developers, and deployers* [117, 149]. Critics attempted to understand algorithmic systems by examining the important decisions that brought them into being. For example, concerned Carnegie Mellon University students who opposed Pittsburgh’s use of a predictive policing model filed a Right to Know request that revealed that the algorithm was rolled out without any meaningful consultation with the communities it was used to police [91]. Several advocacy groups have created resources to help people collect information about human decisions and processes surrounding algorithmic systems. For algorithms used by public agencies, organizations including the Electronic Frontier Foundation (“EFF”) [9] and Lucy Parsons Labs [15] have developed resources to help people navigate filing public records requests, which in some cases still only resulted in partial information being released [91, 110].

### 4.3 Dissemination

Broader dissemination was a common strategy employed by critics to raise public awareness and put public pressure on algorithm owners. Dissemination sometimes caught the attention of elected officials or regulators who could apply pressure or exercise authority from above [79, 80, 170]. Some campaigns led by advocates of abandonment took place *in-person* (offline), whereas in other cases campaigns took place *online*. Other cases were supported by global

media campaigns, such as the #BanTheScan coalition supporting grassroots activist groups’ efforts to ban government use of facial recognition technologies in NYC, Jerusalem, and Hyderabad [4]. In many instances, only after the visibility of criticism increased did algorithm owners respond to and acknowledge the critique.

**Online vs. Offline.** While less common, a small number of cases involved in-person *public demonstrations*. All of the public demonstrations in our case set were organized to oppose systems that were intended to be deployed in a specific *locality* (*e.g.*, in a particular city [39]). Many such demonstrations were organized and attended by people who were directly impacted by the algorithm, such as student-organized protests against algorithmically assigned exam grades in the UK [135], or a resident-led protest of a proposed facial recognition algorithm held on the sidewalk outside of their apartment [58]. In contrast, criticism of tools with spatially diffuse user bases took place *online*. Several cases gained visibility through high engagement on online platforms, such as Twitter, Instagram, or YouTube by going “viral”, such as users’ interactions with NEDA’s Tessa Bot [174], or the Korean chatbot Lee-Luda [98]. We found that highly visible cases (as approximated through engagement metrics, such as high reaction or re-posts counts) were significantly more likely to receive a timely response from the algorithm owner.

**Getting press.** In several cases, media coverage of criticism played a critical role in the fall of an algorithm. As one example, Oregon’s DHS decided just two weeks after an Associated Press story [79] to halt a pilot on child maltreatment risk prediction that had been in progress for four years. US Senator Ron Wyden “reached out” to the department immediately after the story to “ask questions about [the tool’s] racial bias” [7]. Similar reporting efforts helped activist efforts gain momentum, and sometimes caught the attention of those in positions of power who applied pressure from above.

### 4.4 Dialogue

When criticism about an algorithm became more publicly visible, actors who held power over the algorithm were called to respond in a public *dialogue* with their critics. These dialogues occurred in sites ranging from those that are more formal or compulsory for algorithm owners (*e.g.*, in court [87, 88, 109, 134, 163]), to those that were more informal (*e.g.*, on social media [53, 86, 165, 174, 183]). Some algorithm owners took weeks, months, or even years to acknowledge their critics and respond *at all*. In this section, we examine how algorithm owners (dis)engaged with criticism about the algorithm. While some algorithms were primarily criticized by the public, other algorithms found support from other actors beyond their owners, and sparked a broader discourse about whether or not they should exist.

**Denial vs. acceptance of algorithmic harm.** When presented with evidence of algorithmic harm, some algorithm owners responded by *discrediting* the evidence presented to them, either by creating counter-evidence [62] or ad-hominem attempts to discredit the critic [72, 182]. For example, when an Instagram user shared an infographic that described a harmful and triggering interaction with NEDA’s Tessa Wellness chatbot<sup>5</sup>, NEDA’s VP of Communications accused her of lying in the post’s comment section [182]. Interestingly, in many cases [72, 182], initial denials were later retracted

<sup>5</sup><https://www.instagram.com/p/Cs1jp1pPkOs/>



or withdrawn due to public scrutiny. After the NEDA VP's comment sparked outrage from other users, she deleted her comment and posted that she "*retracted her statement*" [182]. This pattern of denial followed by retraction is a broader trend in responses to attempts for algorithmic accountability [137].

In contrast, rather than resist, some algorithm owners immediately acknowledged the validity of the critiques. For example, when Twitter users began to investigate if its image cropping algorithm was racially biased [183], Twitter immediately issued a public statement that they were "*prioritizing work*" to remedy bias in their algorithms [16]. Cooperative responses that immediately acknowledged critiques allowed organizations to take actions (including abandonment) to redress harms more quickly.

**The public vs. private dialogue.** One dimension that influenced whether an algorithm was abandoned was whether algorithm owners responded publicly. Cases that were resolved interpersonally rather than in public were significantly less likely to result in critiques being resolved or the system being abandoned. For example, Zoom users observed in their personal experiences [151] that Zoom's virtual background feature tended to be more dysfunctional for users with darker skin, prompting some to tweet about it, but Zoom did not respond until one of these tweets<sup>6</sup> went viral. Zoom corporate replied to the tweet to invite its author to a call. Later that week, the author of the tweet posted that he had "*had a great convo with [Zoom's Chief DEI Officer] [...] I have no reason to doubt that they are serious about fixing problems with virtual backgrounds [...] the problem is deep and nuanced*". Yet, such private resolutions still leave the broader public in the dark about what commitments the algorithm owner has made.

**Opposing the abandonment/supporting the algorithm.** As discourse surrounding an algorithm became increasingly visible, not all people believed that the algorithm should be abandoned. For example, calls to ban police use of facial recognition technologies worldwide were met with concern from law enforcement officials. In the U.S., police unions and conservative advocacy groups invested resources in lobbying against bans of facial recognition technologies [67, 144]. Similarly, in Uganda, local police and members of the majority party lauded the facial recognition systems they purchased from Chinese vendor Huawei for "*boosting security*" [179]. Beyond law enforcement officials, even citizens support the use of facial recognition in some contexts: for example, 60% of surveyed Indian citizens supported the use of facial recognition to identify protesters [14]. These examples demonstrate how the debate to abandon an algorithm often implicates the larger surrounding political contexts and historic ideals (e.g., surveillance, nationalism, and criminalization) that animate new technologies.

#### 4.5 Decision

Those who held power over the algorithm could *make decisions* to intervene on its design, development, or usage. We group these decisions about the algorithm into 4 categories based on their outcome: whether the algorithm was *not abandoned*, in which the algorithm owners could (1) *repair* the algorithm, or (2) make no changes; or the algorithm was *abandoned* (3) by the algorithm owner's choice, or (4) by law or policy. In some cases, multiple decisions were made,

e.g., many algorithm owners first attempted to repair an algorithm before later deciding to abandon it.

**Repair.** Algorithm owners took several actions to repair or remediate potential harms, including fixing or updating dysfunctional technology [104, 142, 161] or creating better infrastructure or policies around its use [48, 142, 153]. For example, when independent investigations revealed flaws in Epic's sepsis prediction model (which was deployed in several hundred U.S. hospitals) [180], the company completely "overhauled" its original algorithm and released a second version that used completely different input and target variables, and recommended that the model be trained on a hospital's *own* data rather than being applied off-the-shelf [142]. However, repairs that improved the technology by making it more functional or less biased did not redress harms that resulted from the algorithms' existence.

**Abandoned by the algorithm owner's choice.** In our analysis, the majority of systems that were abandoned were abandoned by choice, *i.e.*, by a decision made by an organization that developed or used the algorithm. To name a few, Korean start-up ScatterLab disabled their Lee-Luda chatbot [98], Twitter decided to stop using an algorithm to crop users' images [43], and leadership at Oregon's Department of Human Services decided to stop using a child welfare call screening algorithm [7]. Relatively few algorithm owners released public statements with explanations or named reasons why they decided to abandon the algorithm. Public statements that we found often referenced algorithmic harms as a primary reason why the system was abandoned [43, 83, 98, 101].

**Abandoned through policy.** Algorithms overseen by uncooperative algorithm owners could still be abandoned through another governing body creating or enforcing applicable laws. Several algorithms in our case set were deemed illegal through the interpretation of *existing laws* (including laws written without an algorithm in mind) in a court of law. For example, a Dutch court ruled that their government must stop using an algorithm that falsely predicted that thousands of low-income residents had committed benefits fraud, after finding that the algorithm violated citizens' right to privacy in the European Convention on Human Rights [1]. U.S. courts ruled Medicaid benefits allocation algorithms in violation of existing laws (*i.e.*, due process law and the Administrative Procedure Act) in both Arkansas and Idaho [36, 59]. Other algorithms were banned by existing governing bodies founded to protect citizens' rights, such as the U.S. Federal Trade Commission issuing a five-year ban on Rite Aid's use of facial recognition because of the technology's risks to consumers [13].

However, all past cases of systems abandoned through executive or judicial enforcement took a *significantly longer time* (e.g., several years [59, 104]) before the decision to abandon the algorithm was reached. These cases also required significant investment: e.g., the legal team representing Idaho Medicaid beneficiaries spent over \$40,000 on expert consultations and 2,000 attorney hours to secure a settled agreement [59].

Algorithmic systems were also abandoned through the *creation of new policies* by a governing body (e.g., a state legislature or city council) that had jurisdiction over the algorithm's design, development, or usage. New laws that target algorithmic systems have been rapidly growing in countries across the world [10]. For example,

<sup>6</sup><https://twitter.com/colinmadland/status/1307111816250748933>

several U.S. state legislatures have proposed bans on algorithms developed for specific purposes, including facial recognition [144], betting or casino algorithms [71], employee monitoring [177], or providing mental health counseling [12]. All relevant policies that we identified for our cases targeted banning the *usage* (rather than the *development*) of technology.

#### 4.6 Death

The *death* phase explores what happens *after* an algorithm is abandoned. Because algorithm development often entails creating infrastructure (such as data collection pipelines or portable products) [57], we discuss how abandoned algorithms' impacts can "live on" by influencing future processes [57], or in "reincarnated" products deployed elsewhere. Our understanding of the afterlife of abandoned algorithmic systems builds upon Ehsan et al. [57]'s conceptualization of the *algorithmic imprint*, which describes how an algorithm's impact can extend beyond its deployment, by influencing infrastructural conditions, organizational practices, and societal norms that continue on in its "afterlife".

**Overriding past decisions.** Past decisions issued by an algorithm can continue to negatively impact people who received adverse outcomes. In the majority of cases of algorithms that were abandoned after their deployment, we found that the algorithm owners did *not* re-consider or override past (harmful) decisions made using the algorithm. For example, to our knowledge, rejected resumes were not re-screened [47] and closed child welfare investigations were not re-opened [7]. In a few notable cases, however, decisions were overturned and re-issued en masse. For example, in response to global protests against algorithmically-assigned A-level grades, U.K. exam boards globally rescinded all algorithm-assigned grades and revised them based on teacher-assessed grades (without the algorithmic standardization under criticism [57]). Similarly, a legal settlement in Idaho guaranteed that people whose home care benefits were cut by an algorithm would continue to receive their original budgets [88].

**Making reparations.** Beyond the rare re-issuing of decisions, some algorithm owners provided other kinds of reparations to symbolically or materially mend past harms. Material reparations, for example, may be especially appropriate for systems that caused allocative harm that made individuals incur additional financial costs, e.g., by withholding access to healthcare benefits or employment. The majority of algorithm owners did not provide financial reparations for the harm caused, with a few notable exceptions. The Dutch government paid €30,000 to all families falsely accused of committing benefits fraud by an algorithm [60]. The Arkansas DHS was required to pay \$460,000 in a court settlement to three plaintiffs whose benefits were severely cut, but many other hundreds of beneficiaries who were similarly impacted by the algorithm were not compensated [108]. No algorithm owners in our case set provided reparations for other types of non-material harms (e.g., representational, psychological, or interpersonal harms).

**Reincarnation.** A growing number of investigations uncover how abandoned technologies can re-appear in different contexts [94, 106, 146]. We observed that when a city cancelled a contract with a technology vendor, the company often remained and continued to sell their product elsewhere [75, 76]. Even as cities such as

Palo Alto decided to drop their contracts with the predictive policing company PredPol, public records requests revealed that PredPol still had (many previously undisclosed) contracts with dozens of U.S. cities [75]. In response to bad press, some algorithm owners re-branded [3]: e.g., one month after the LAPD stopped using PredPol, the company changed its name to Geolitica and its slogan from "The Predictive Policing Company" to "Data-Driven Community Policing" [29]. In 2023, SoundThinking (the company that created ShotSpotter, another prominent law enforcement algorithm) acquired IP from Geolitica, and still maintains versions of its predictive policing algorithms [114]. The case of PredPol exemplifies how built technologies are valued as infrastructure that remains and can be recycled even after they are abandoned.

## 5 FACTORS RELEVANT TO ALGORITHM ABANDONMENT

Our analysis of historic cases across a range of domains reveals several patterns, including common phases that took place toward abandonment. However, beyond identifying commonalities, we also noticed key *differences* relevant to the outcome of whether the system was ultimately abandoned, and how long it took to get to that stage. In this section, we single out these *differentiating socio-technical factors*, which span characteristics of the technical system and the social systems that the algorithm was embedded within. We highlight the implications of each factor for several stakeholders. For example, because abandonment is an important step that algorithm owners can take to redress (potential) algorithmic harms, we discuss ways that algorithm owners can design algorithmic systems that *can be more easily abandoned* at different stages of their lifecycle. We also discuss implications of each factor for other external stakeholders, including FAcCT researchers and oversight entities.

**Users vs. Subjects.** *Is the algorithm a consumer product, or used on (unaware) subjects?* Algorithms in consumer products (e.g., online platforms) "fell" more quickly than algorithms used by organizations (e.g., governments) to analyze or make decisions about subjects who were potentially unaware of the algorithm's existence. Direct users or consumers who interacted with technologies were more likely to be aware that there was an algorithm in place, and had more agency to resist harmful interactions with it. In cases where the direct user who consumed the algorithm's outputs used it to make decisions about someone else (e.g., a policing agency using facial recognition on city residents), impacted individuals were often unaware of its use or existence, which prolonged the process leading to abandonment. This distinction draws attention to the importance of understanding power differentials in who purchases and uses, versus who is most (negatively) impacted by algorithmic systems. Algorithm owners and critics alike should aim to empower and amplify the perspectives of those most impacted by algorithmic systems, even if they are not direct users.

**Value proposition.** *What is the "value proposition" behind the use of the algorithm?* A critical factor relevant to cases' outcomes was understanding algorithm owners' motivation for creating or adopting the algorithm, i.e., its value proposition. Publicly available value propositions often articulated a clear set of political values, e.g., to reduce Medicaid spending (by cutting care) [173], to maximize



returns to landlords (by raising rents) [166], or to enable constant and wide-scale surveillance [48]. In some cases, algorithm owners valued that automation allowed them to de-personalize decisions formerly made by humans, *e.g.*, that leasing agents or nurse assessors had “*too much empathy*” [166] to raise rents or cut hours. When an algorithm’s value proposition was clearly stated, in some cases critics organized successful campaigns that countered them (*e.g.*, challenging the notion that surveillance makes us feel safe [2]). Algorithm owners should make an algorithm’s value proposition clear to the public. Critics should search for value propositions (*i.e.*, how is the technology being marketed to its buyers? [141]) and make them visible so that they can be contested.

**The “domino effect” caused by deep dependencies.** *How embedded is the algorithm into other systems? What other practices and processes “depend” on it?* Stand-alone algorithms released as independent products, like chatbots, were much easier to abandon than those embedded into larger social or technical systems, like algorithms used to automate important bureaucratic decisions [19]. Abandoning an algorithm that other components rely on might have a domino effect that requires changing its surrounding socio-technical components. For example, when courts ruled algorithms that assessed beneficiaries’ eligibility illegal, public officials scrambled to develop a new process (*e.g.*, to train human assessors) to allocate care resources [87]. Algorithm owners should aim to minimize dependencies by creating and budgeting for back-up plans in a scenario where the algorithm is no longer in use, a proposal that has become popular in recent AI governance frameworks [126, 127].

**Procedural transparency and public notice.** *Is the public aware that an algorithm exists? Do we know how important decisions were made?* We found that many harmful algorithms were used for months (or even years) before they were discovered and the general public was made aware of their existence. In contrast, when algorithm owners made their plans to use an algorithm public *before* it was deployed (like Missouri’s NF LOC algorithm [89]), in some cases they decided to abandon the algorithm before it could inflict harm in the real world. Thus, the extent to which algorithm owners made information about the algorithm public, and when, played a pivotal role in the time it took for harmful systems to be abandoned. Algorithm owners should provide public notice and publicly release information about the algorithm’s design at the earliest possible stages of the algorithm’s lifecycle. FAccT researchers and advocacy groups can contribute to the small (yet growing) number of tools identified by Ojewale et al. [129] designed to aid “target identification”, *i.e.*, to “*uncover deployed systems by making them visible to external stakeholders*”. Finally, policymakers or organizations can institute transparency requirements for algorithm owners (*e.g.*, those proposed in [68, 121, 139]).

**Access and auditability.** *What access do independent actors, e.g., members of the public, have to the algorithm?* As discussed in Section 4.2, while independent audits were an extremely effective tool that triggered the decision to abandon an algorithm, independent actors had varying degrees of access to different algorithmic systems. FAccT researchers can contribute to infrastructure and build tools to support independent audits, such as those taxonomized by [129]. One promising trend is that coalitions of different actors have started collaborations to share knowledge and data about their

interactions with algorithmic systems, such as academic or non-profit research groups supporting public defenders and worker’s unions [63, 89, 173]. Policymakers and advocacy groups can consider funding these independent audits [122].

**Visibility of criticism and timely media coverage.** *How visible is the critique to other stakeholders?* As discussed in Sections 4.3 and 4.4, critics’ efforts to abandon an algorithm were often intertwined with a broader struggle for visibility. Getting media coverage that uplifted critics’ concerns and hopes for the future was often a turning point for accelerating an algorithm’s fall. Other cases that went viral from organic engagement (such as amplification from influential users) were also significantly more likely to result in algorithm owners taking action to address algorithmic harms [83, 183]. Thus, journalists, media organizations, and “everyday users” [151] can play an important role in advocating for abandonment.

**Oversight and regulatory environment.** *Are policymakers receptive to creating new policies? Who is held liable for algorithmic harms?* As discussed in Section 4.5, instituting new “bright-line rules” [8] that explicitly ban specific types of algorithms takes far less time (and material investment) compared to interpreting existing laws in court. Policymakers can consider instituting such bans *proactively*, even before the banned systems are in use, like the City of Somerville’s ban on employees’ use of facial recognition [181]. Another trend we observed is a lack of accountability under the law for vendors that develop algorithmic systems and sell them to organizations that oversee their deployment. For example, when courts ruled Arkansas’ roll-out of an algorithm developed by interRAI (an independent vendor) illegal, the state of Arkansas, not interRAI, was required to pay the full settlement amount to plaintiffs who were harmed [108]. This accountability structure means that even when one organization decides to stop using a harmful algorithm, it can continue to be used elsewhere with little consequence to the vendor who would continue to benefit from the algorithm [105]. Policymakers should create and interpret existing laws to promote what Metcalf [116] calls “*distributed accountability*”, where vendors are also held liable for remedying algorithmic harms.

## 6 CONCLUSION & FUTURE WORK

Grassroots movements, moratorium bans, and boardroom decisions to abandon harmful algorithmic systems continue to grow across the world. Yet, comparatively little academic attention has been paid to understanding the phenomena of algorithm abandonment. In our analysis of real-world cases involving (calls for) algorithm abandonment, we identified common phases of events that took place and key socio-technical factors that enabled (or prolonged) abandonment efforts. Our analysis reveals power imbalances and information asymmetries that prevent critics from holding algorithm owners accountable. Yet, victories from below in our case set illustrate how these power differentials might be challenged through policy interventions and collective action. In our discussion of relevant factors, we highlight steps that algorithm owners can take to create algorithms that can be more easily be contested and abandoned. Regulators can consider making these steps mandatory, such as requiring algorithm owners to make important information public at early stages of design and development [139].

Our preliminary work points to several exciting directions for future work on algorithm abandonment. Researchers can take inspiration from successful collaborations that have built tooling and infrastructure to enable collective contestation, such as partnerships between academic researchers and worker’s unions [63, 89, 173]. Future research efforts can conduct more in-depth studies, such as interviews, to understand the experiences of actors involved in making the decision to abandon an algorithm, as well as other important stakeholders such as impacted communities, advocacy groups, and policymakers. Future work can also center campaigns that go beyond calls to *abandon*, and instead fight to *abolish* harmful technologies. In her conceptualization of “*abolitionist technologies*”, Ruha Benjamin writes that “*calls for abolition are never simply about bringing harmful systems to an end but also about envisioning new ones. After all, the etymology of abolition includes Latin root words for ‘destroy’ (abolere) and ‘grow’ (olere)*” [28]. In our analysis, we found that algorithms in our case set implicate oppressive social systems or politics that informed their design. How can the FAcCT community can better support the many advocacy initiatives [2, 3, 28, 119] that work to contest the values that animate harmful algorithmic systems, and imagine more just futures?

## 7 IMPACT STATEMENTS

**Researcher Positionality.** We acknowledge that our lived experiences shape our positionality and perspectives [102] in the study of algorithmic harms, which are often disproportionately faced by members of marginalized groups. As researchers who were trained and work in the U.S., we recognize that our analysis, particularly our discussions of algorithmic accountability, contestation (e.g., through public protests), and governance was influenced by the present political context of the U.S. In our initial search for cases, we noticed that the majority of cases cited in FAcCT research articles or reported in crowd-sourced AI incident databases occurred in the U.S. Further, because our search was conducted in English, we likely missed relevant cases that occurred in non-English-speaking regions. We attempted to address this regional bias by performing additional keyword searches and including an additional Latin America-based incident database [164]. We hope that our exploratory research raises academic awareness of algorithm abandonment as a *global* phenomena, and inspires future work that more deeply interrogates similarities and differences relevant to how the dynamics of abandonment play out across regional contexts.

While all authors have experience studying algorithmic harm as academic researchers, it was important for us to better understand the needs and obstacles faced by activists taking action to resist oppressive technologies on the ground. Thus, as we conducted this research, we consulted experts who had lived experiences organizing to resist algorithmic harm. Specifically, our research team met with members of the Surveillance Resistance Lab, a U.S.-based advocacy group founded to resist state and corporate surveillance technologies, who directed us to cases and offered feedback. Looking forward, our research team plans to translate parts of this work to create resources that can directly benefit advocacy groups.

**Ethical Considerations.** In proposing a process-oriented view of the dynamics of algorithm abandonment, our analysis reveals commonalities that place seemingly dissimilar cases (of chatbots

[83], B2B SAAS products [166], and carceral technologies [158]) in conversation with one another. These similarities enable us to formulate recommendations for future research and policy interventions related to algorithm abandonment.

However, our analysis has several limitations. In focusing on what is *shared* between a large and diverse set of algorithmic systems that are deployed in a wide variety of political contexts, we exclude important details that illustrate the uniqueness and complexity of each case. Further, given the rapidly evolving nature of algorithmic capabilities and harms, we acknowledge that our process-oriented view may fail to describe the events leading to the fall of all algorithmic systems (including recent generative, multi-modal, and multi-purpose AI systems). Because of these concerns, the 6 D’s of abandonment are *not* meant to serve as a model of steps for critics to follow, but rather provide descriptive context to understand similar historic struggles.

**Adverse Impact.** Our aim in identifying factors that permit (or hinder) algorithm abandonment is to support actors *from below* in contesting harmful algorithmic systems. To do so, we name and draw attention to common actions that actors *from above* took to dodge and discredit their critics. However, this work also has the potential to be misused by algorithm owners who model their actions after historic cases that successfully reinforced oppressive power structures (e.g., by intentionally obscuring an algorithm). Despite this risk, we advocate for the value of naming and identifying these tactics as a tool to strengthen the eco-system of accountability around algorithmic systems.

## ACKNOWLEDGMENTS

We thank Mizue Aizeki and Ed Vogel from the Surveillance Resistance Lab for offering feedback on this work. We thank Ryan Steed, Meg Young, Upol Ehsan, and Sohini Upadhyay for directing us to related scholarship and cases. We thank Alana Rogers and Ayush Tripathi for conducting preliminary research that informed this project. HH and NJ acknowledge support from NSF (IIS2040929 and IIS2229881), PwC (through the Digital Transformation and Innovation Center), and the Block Center for Technology and Society at CMU. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not reflect the views of the National Science Foundation and other funding agencies.

## REFERENCES

- [1] AlgorithmWatch 2020. *How Dutch activists got an invasive fraud detection algorithm banned*. AlgorithmWatch. Retrieved January 16, 2024 from <https://algorithmwatch.org/en/syri-netherlands-algorithm/> Email newsletter.
- [2] Logic(s) Magazine 2020. *Safe or Just Surveilled?: Tawana Petty on the Fight Against Facial Recognition Surveillance*. Logic(s) Magazine. Retrieved January 16, 2024 from <https://logicmag.io/security/safe-or-just-surveilled-tawana-petty-on-facial-recognition/> Interview.
- [3] Stop LAPD Spying Coalition (Ed.). 2021. *The Ghosts of White Supremacy in AI Reform*. <https://ainowinstitute.org/publication/a-new-ai-lexicon-surveillance>
- [4] Amnesty International (Ed.). 2022. *Ban the Scan*. Retrieved January 16, 2024 from <https://banthescan.amnesty.org/>
- [5] Eticas 2022. *The External Audit of the VioGen System*. Eticas. <https://eticasfoundation.org/wp-content/uploads/2022/03/ETICAS-FND-The-External-Audit-of-the-VioGen-System.pdf>
- [6] 2022. Galactica Demo — galactica.org. <https://galactica.org/>. [Accessed 22-01-2024].
- [7] The Associated Press (Ed.). 2022. *Oregon is dropping an artificial intelligence tool used in child welfare system*. Retrieved January 16,

- 2024 from <https://www.npr.org/2022/06/02/1102661376/oregon-drops-artificial-intelligence-child-abuse-cases>
- [8] AI Now Institute (Ed.). 2023. *Algorithmic Accountability: Moving Beyond Audits*. Retrieved January 16, 2024 from <https://ainowinstitute.org/publication/algorithmic-accountability>
- [9] Electronic Frontier Foundation 2023. *FOIA How To*. Electronic Frontier Foundation. Retrieved January 16, 2024 from <https://www.eff.org/issues/transparency/foia-how-to>
- [10] International Association of Privacy Associates 2023. *Global AI Legislation Tracker*. International Association of Privacy Associates. Retrieved January 16, 2024 from <https://iapp.org/resources/article/global-ai-legislation-tracker/>
- [11] 2023. New liver transplant rules yield winners, losers as wasted organs reach record high — [washingtonpost.com](https://www.washingtonpost.com/business/2023/03/21/liver-transplants-acuity-circle-policy/). [Accessed 22-01-2024].
- [12] Duane Morris Government Strategies (Ed.). 2023. *Regulating Artificial Intelligence in Mental Health: States' Attempts To Revolutionize Mental Health Services*. Retrieved January 16, 2024 from <https://statecapitallobbyist.com/artificial-intelligence-ai/regulating-artificial-intelligence-in-mental-health-services-states-attempts-to-revolutionize-mental-health-services/>
- [13] The U.S. Federal Trade Commission 2023. *Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology without Reasonable Safeguards*. The U.S. Federal Trade Commission. Retrieved January 16, 2024 from <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without-press-release>
- [14] Common Cause & Lokniti – Centre for the Study Developing Societies (CSDS) 2023. *Status of Policing in India Report 2023: Surveillance and the Question of Privacy*. Common Cause & Lokniti – Centre for the Study Developing Societies (CSDS). Retrieved January 16, 2024 from [https://www.commoncause.in/wotadmin/upload/REPORT\\_2023.pdf](https://www.commoncause.in/wotadmin/upload/REPORT_2023.pdf)
- [15] Lucy Parsons Labs 2023. *Who We Are*. Lucy Parsons Labs. Retrieved January 16, 2024 from <https://lucyparsonslabs.com/about/>
- [16] Parag Agrawal and Dantley Davis. 2020. *Transparency around image cropping and changes to come*. [https://blog.twitter.com/en\\_us/topics/product/2020/transparency-image-cropping](https://blog.twitter.com/en_us/topics/product/2020/transparency-image-cropping)
- [17] Sarah Ahmed. 2017. *No*. Retrieved January 16, 2024 from <https://feministkilljoys.com/2017/06/30/no/>
- [18] AIAAIC. 2023. *AI, Algorithmic, and Automation Incidents and Controversies (AIAAIC) Repository*. Retrieved January 18, 2024 from <https://www.aiaaic.org/aiaaic-repository>
- [19] Ali Alkhatib and Michael Bernstein. 2019. *Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions*. 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, 1–13. <https://doi.org/10.1145/3290605.3300760>
- [20] Nazanin Andalibi, Cassidy Pyle, Kristen Barta, Lu Xian, Abigail Z. Jacobs, and Mark S. Ackerman. 2023. *Conceptualizing Algorithmic Stigmatization*. 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 373, 18 pages. <https://doi.org/10.1145/3544548.3580970>
- [21] Laboratorio de Inteligencia Artificial Aplicada. 2018. *Sobre la predicción automática de embarazos adolescentes*.
- [22] Vijay Arya, Rachel K. E. Bellamy, Pin-Yu Chen, Amit Dhurandhar, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Q. Vera Liao, Ronny Luss, Aleksandra Mojsilović, Sami Mourad, Pablo Pedemonte, Ramya Raghavendra, John Richards, Prasanna Sattigeri, Karthikeyan Shanmugam, Moninder Singh, Kush R. Varshney, Dennis Wei, and Yunfeng Zhang. 2019. *One Explanation Does Not Fit All: A Toolkit and Taxonomy of AI Explainability Techniques*. arXiv:1909.03012 [cs.AI]
- [23] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.
- [24] Eric P.S. Baumer and M. Six Silberman. 2011. *When the Implication is Not to Design (Technology)*. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (, Vancouver, BC, Canada.) (CHI '11). Association for Computing Machinery, New York, NY, USA, 2271–2274. <https://doi.org/10.1145/1978942.1979275>
- [25] Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. 2018. *AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias*. arXiv:1810.01943 [cs.AI]
- [26] Emily Bembeneck, Rebecca Nissan, , and Ziad Obermeyer. 2021. *Algorithmic Bias Playbook*. [https://www.ftc.gov/system/files/documents/public\\_events/1582978/algorithmic-bias-playbook.pdf](https://www.ftc.gov/system/files/documents/public_events/1582978/algorithmic-bias-playbook.pdf)
- [27] Ruha Benjamin. 2016. *Informed Refusal: Toward a Justice-based Bioethics*. *Science, Technology, & Human Values* 41, 6 (2016), 967–990. <https://doi.org/10.1177/0162243916656059>
- [28] Ruha Benjamin. 2019. *Race after technology: Abolitionist tools for the new Jim code*. Polity.
- [29] Johana Bhuiyan. 2021. *LAPD ended predictive policing programs amid public outcry. A new effort shares many of their flaws*. Retrieved January 16, 2024 from <https://www.theguardian.com/us-news/2021/nov/07/lapd-predictive-policing-surveillance-reform>
- [30] Virginia Braun and Victoria Clarke. 2006. *Using thematic analysis in psychology*. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. <https://doi.org/10.1191/1478088706qp0630a> arXiv:<https://www.tandfonline.com/doi/pdf/10.1191/1478088706qp0630a>
- [31] Rachel Bukowitz and Tim O'Loughlin. 2022. *Governing by Algorithm? Child Protection in Aotearoa New Zealand*. <https://anzsog.edu.au/research-insights-and-resources/research/governing-by-algorithm-child-protection-in-aotearoa-new-zealand/>
- [32] Bundesverfassungsgericht. 2023. *Bundesverfassungsgericht - Press - Legislation in Hesse and Hamburg regarding automated data analysis for the prevention of criminal acts is unconstitutional*. <https://www.bundesverfassungsgericht.de/SharedDocs/Pressemittelungen/EN/2023/bvg23-018.html>
- [33] Joy Buolamwini and Timnit Gebru. 2018. *Gender shades: Intersectional accuracy disparities in commercial gender classification*. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [34] Joy Buolamwini and Timnit Gebru. 2023. *Algorithmic Justice League Gender Shades 5th Anniversary Celebration*. (2023). <https://www.youtube.com/watch?v=8JSxbZyivvE> Virtual discussion (recorded).
- [35] Jenna Burrell and Deirdre Mulligan. 2020. *The Berkeley Algorithmic Fairness and Opacity Group Refusal Conference*. Retrieved January 16, 2024 from <https://afog.berkeley.edu/programs/the-refusal-conference>
- [36] Ryan Calo and Danielle Keats Citron. 2021. *The Automated Administrative State: A Crisis of Legitimacy*. *Emory Law Journal* (2021). [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3553590](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3553590)
- [37] Zeph Grunspan and Janet Bass. 2017. *Federal Suit Settlement: End of Value-Added Measures for Teacher Termination in Houston* | American Federation of Teachers. <https://www.aft.org/press-release/federal-suit-settlement-end-value-added-measures-teacher-termination-houston>
- [38] Nathalie Casemajor, Stéphane Couture, Mauricio Delfin, Matthew Goerzen, and Alessandro Delfanti. 2015. *Non-participation in digital media: toward a framework of mediated political action*. *Media, Culture & Society* 37, 6 (2015), 850–866. <https://doi.org/10.1177/0163443715584098> arXiv:<https://doi.org/10.1177/0163443715584098>
- [39] Cindy Chang. 2018. *LAPD officials defend predictive policing as activists call for its end*. Retrieved January 16, 2024 from <https://www.latimes.com/local/lanow/la-me-lapd-data-policing-20180724-story.html>
- [40] Kyle Chayka. 2023. *Rethinking the Luddites in the Age of AI*. Retrieved January 16, 2024 from <https://www.newyorker.com/books/page-turner/rethinking-the-luddites-in-the-age-of-ai>
- [41] Hao-Fei Cheng, Logan Stapleton, Anna Kawakami, Venkatesh Sivaraman, Yanghui Cheng, Diana Qing, Adam Perer, Kenneth Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. *How Child Welfare Workers Reduce Racial Disparities in Algorithmic Decisions*. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (, New Orleans, LA, USA.) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 162, 22 pages. <https://doi.org/10.1145/3491102.3501831>
- [42] Lingwei Cheng and Alexandra Choudhachova. 2022. *Heterogeneity in Algorithm-Assisted Decision-Making: A Case Study in Child Abuse Hotline Screening*. 6, CSCW2, Article 376 (nov 2022), 33 pages. <https://doi.org/10.1145/3555101>
- [43] Rumman Chowdhury. 2021. *Sharing learnings about our image cropping algorithm*. [https://blog.twitter.com/engineering/en\\_us/topics/insights/2021/sharing-learnings-about-our-image-cropping-algorithm](https://blog.twitter.com/engineering/en_us/topics/insights/2021/sharing-learnings-about-our-image-cropping-algorithm)
- [44] M. Cifor, P. Garcia, T.L. Cowan, J. Rault, T. Sutherland, A. Chan, J. Rode, A.L. Hoffmann, N. Salehi, and L. Nakamura. 2019. *Feminist Data Manifest-No*. Retrieved January 16, 2024 from <https://www.manifestno.com/>
- [45] Richard Conniff. 2011. *What the Luddites Really Fought Against*. Retrieved January 16, 2024 from <https://www.smithsonianmag.com/history/what-the-luddites-really-fought-against-264412/>
- [46] Jeffrey Dastin. [n. d.]. *Insight - Amazon scraps secret AI recruiting tool that showed bias against women*. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scrap-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G/>. [Accessed 22-01-2024].
- [47] Jeffrey Dastin. 2018. *Insight - Amazon scraps secret AI recruiting tool that showed bias against women*. Retrieved January 16, 2024 from <https://www.reuters.com/article/idUSKCN1MK0AG/>
- [48] Jeffrey Dastin. 2020. *Rite Aid deployed facial recognition systems in hundreds of U.S. stores*. Retrieved January 16, 2024 from <https://www.reuters.com/investigates/special-report/usa-riteaid-software/>
- [49] Jenny L. Davis, Apryl Williams, and Michael W. Yang. 2021. *Algorithmic repatriation*. *Big Data & Society* 8, 2 (2021), 20539517211044808. <https://doi.org/10.1177/20539517211044808> arXiv:<https://doi.org/10.1177/20539517211044808>

- [50] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 2020. A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020). <https://api.semanticscholar.org/CorpusID:211171909>
- [51] Wesley Hanwen Deng, Manish Nagireddy, Michelle Seng Ah Lee, Jatinder Singh, Zhiwei Steven Wu, Kenneth Holstein, and Haiyi Zhu. 2022. Exploring How Machine Learning Practitioners (Try To) Use Fairness Toolkits. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAcCT '22). Association for Computing Machinery, New York, NY, USA, 473–484. <https://doi.org/10.1145/3531146.3533113>
- [52] Alicia DeVrio, Motahare Eslami, and Kenneth Holstein. 2023. Building, Shifting, & Employing Power: A Taxonomy of Responses From Below to Algorithmic Harm. *In submission* (2023).
- [53] Megan Rose Dickey. 2020. Twitter and Zoom’s algorithmic bias issues. <https://techcrunch.com/2020/09/21/twitter-and-zoom-algorithmic-bias-issues/>
- [54] Lynn Dombrowski, Ellie Harmon, and Sarah Fox. 2016. Social Justice-Oriented Interaction Design: Outlining Key Design Strategies and Commitments. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) (DIS '16). Association for Computing Machinery, New York, NY, USA, 656–671. <https://doi.org/10.1145/2901790.2901861>
- [55] Finale Doshi-Velez and Been Kim. 2015. Towards A Rigorous Science of Interpretable Machine Learning. arXiv:1702.08608 [stat.ML]
- [56] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (Cambridge, Massachusetts) (ITCS '12). Association for Computing Machinery, New York, NY, USA, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [57] Upol Ehsan, Ranjit Singh, Jacob Metcalf, and Mark Riedl. 2022. The Algorithmic Imprint. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAcCT '22). Association for Computing Machinery, New York, NY, USA, 1305–1317. <https://doi.org/10.1145/3531146.3533186>
- [58] Elizabeth Elizalde and Michael Gartland. 2019. *Brooklyn tenants in rent-regulated apartments push state to nix landlord’s facial recognition software*. Retrieved January 16, 2024 from <https://www.nydailynews.com/2019/05/01/brooklyn-tenants-in-rent-regulated-apartments-push-state-to-nix-landlords-facial-recognition-software/>
- [59] Ritchie Eppink. 2023. *Testimony of Ritchie Eppink: AI in Government United States Senate Committee on Homeland Security & Government Affairs*. Retrieved January 16, 2024 from <https://www.hsgac.senate.gov/wp-content/uploads/Testimony-Eppink-2023-05-16-1.pdf>
- [60] Thomas Erdbrink. 2021. *Government in Netherlands Resigns After Benefit Scandal*. Retrieved January 16, 2024 from <https://www.nytimes.com/2021/01/15/world/europe/dutch-government-resignation-rutte-netherlands.html>
- [61] Motahare Eslami, Kristen Vaccaro, Karrie Karahalios, and Kevin Hamilton. 2017. “Be Careful; Things Can Be Worse than They Appear”: Understanding Biased Algorithms and Users’ Behavior Around Them in Rating Platforms. *Proceedings of the International AAAI Conference on Web and Social Media* 11, 1 (May 2017), 62–71. <https://doi.org/10.1609/icwsm.v11i1.14898>
- [62] Virginia Eubanks. 2018. *A response to Allegheny County DHS*. <https://virginia-eubanks.com/2018/02/16/a-response-to-allegheeny-county-dhs/>
- [63] FairFare. 2023. *FairFare: Unveiling Ridehail Fairness*. Retrieved January 16, 2024 from <https://getfairfare.org/>
- [64] Katherine B. Forrest. 2021. *When Machines Can Be Judge, Jury, and Executioner: Justice in the Age of Artificial Intelligence*. World Scientific.
- [65] Batya Friedman, Peter Kahn, Alan Borning, Ping Zhang, and Dennis Galletta. 2006. *Value Sensitive Design and Information Systems*. [https://doi.org/10.1007/978-94-007-7844-3\\_4](https://doi.org/10.1007/978-94-007-7844-3_4)
- [66] Yasmin Gagne. 2019. *How we fought our landlord’s secretive plan for facial recognition—and won*. Retrieved January 16, 2024 from <https://www.fastcompany.com/90431686/our-landlord-wants-to-install-facial-recognition-in-our-homes-but-were-fighting-back>
- [67] William Gavin. 2022. *As privacy concerns arise, organizations using facial recognition technology spend on lobbying*. <https://www.opensecrets.org/news/2022/03/as-privacy-concerns-arise-organizations-using-facial-recognition-technology-continue-to-employ-lobbyists/>
- [68] Albert Gehami and Leila Doty. 2023. When the Rubber Meets the Road: Experience Implementing AI Governance in a Public Agency with the City of San José. (2023). [https://www.youtube.com/watch?v=Bif3fwI\\_d20](https://www.youtube.com/watch?v=Bif3fwI_d20) ACM FAcCT Tutorial.
- [69] Marissa Gerchick, Tobi Jegede, Tarak Shah, Ana Gutierrez, Sophie Beiers, Noam Shemtov, Kath Xu, Anjana Samant, and Aaron Horowitz. 2023. The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool. 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, 1292–1310. <https://doi.org/10.1145/3593013.3594081>
- [70] Lauren Kaori Gurley. 2021. *Amazon’s AI Cameras Are Punishing Drivers for Mistakes They Didn’t Make*. Retrieved January 16, 2024 from <https://www.vice.com/en/article/88npjv/amazons-ai-cameras-are-punishing-drivers-for-mistakes-they-didnt-make>
- [71] Rebecca Hanchett. 2022. *Rhode Island Looks At Banning Facial Recognition By Sports Betting Apps*. Retrieved January 16, 2024 from <https://www.gamingtoday.com/news/rhode-island-banning-facial-recognition-sports-betting-apps/>
- [72] Karen Hao. 2020. *The two-year fight to stop Amazon from selling face recognition to the police*. <https://www.technologyreview.com/2020/06/12/1003482/amazon-stopped-selling-police-face-recognition-fight/>
- [73] Bernard E. Harcourt. 2006. *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age*. The Chicago University Press.
- [74] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2016/file/9d2682367c3935defcb1f9e247a97c0d-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/9d2682367c3935defcb1f9e247a97c0d-Paper.pdf)
- [75] Caroline Haskins. 2019. *Dozens of Cities Have Secretly Experimented With Predictive Policing Software*. <https://www.vice.com/en/article/d3m7jq/dozens-of-cities-have-secretly-experimented-with-predictive-policing-software>
- [76] Caroline Haskins. 2021. *The NYPD Has Misled The Public About Its Use Of Facial Recognition Tool Clearview AI*. <https://www.buzzfeednews.com/article/carolinehaskins1/nypd-has-misled-public-about-clearview-ai-use>
- [77] Kashmir Hill. 2020. *The Secretive Company That Might End Privacy as We Know It*. Retrieved January 16, 2024 from <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- [78] Kashmir Hill. 2020. *Wrongfully Accused by an Algorithm*. *The New York Times* (Jun 2020). <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>
- [79] Sally Ho and Garance Burke. 2022. *An algorithm that screens for child neglect raises concerns*. Retrieved January 16, 2024 from <https://apnews.com/article/child-welfare-algorithm-investigation-9497ee937e0053ad414a86c68241ef1>
- [80] Sally Ho and Garnace Burke. 2023. *Child welfare algorithm faces Justice Department scrutiny*. <https://apnews.com/article/justice-scrutinizes-pittsburgh-child-welfare-ai-tool-4f61f45bfc3245fd256e886c2da988b,lastaccessed=January16,2024>
- [81] Sarah Homewood. 2019. Inaction as a Design Decision: Reflections on Not Designing Self-Tracking Tools for Menopause. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290607.3310430>
- [82] Bonnie Honig. 2021. *A Feminist Theory of Refusal*. Harvard University Press.
- [83] Amanda Hoover. 2023. *An Eating Disorder Chatbot Is Suspended for Giving Harmful Advice*. Retrieved January 16, 2024 from <https://www.wired.com/story/tessa-chatbot-suspended/>
- [84] Lara Houston, Steven J. Jackson, Daniela K. Rosner, Syed Ishiaque Ahmed, Meg Young, and Laewoo Kang. 2016. Values in Repair. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1403–1414. <https://doi.org/10.1145/2858036.2858470>
- [85] Noura Howell, Audrey Desjardins, and Sarah Fox. 2021. Cracks in the Success Narrative: Rethinking Failure in Design Research through a Retrospective Triethnography. *ACM Trans. Comput.-Hum. Interact.* 28, 6, Article 42 (nov 2021), 31 pages. <https://doi.org/10.1145/3462447>
- [86] <https://www.theguardian.com/profile/justinmccurry>. [n. d.]. South Korean AI chatbot pulled from Facebook after hate speech towards minorities – theguardian.com. <https://www.theguardian.com/world/2021/jan/14/time-to-properly-socialise-hate-speech-ai-chatbot-pulled-from-facebook>. [Accessed 22-01-2024].
- [87] Benefits Tech Advocacy Hub. 2022. *Arkansas Medicaid Home and Community Based Services Hours Cuts*. Retrieved January 16, 2024 from <https://www.btah.org/case-study/arkansas-medicaid-home-and-community-based-services-hours-cuts.html>
- [88] Benefits Tech Advocacy Hub. 2022. *Idaho Medicaid Home and Community Based Services Care Cuts*. Retrieved January 16, 2024 from <https://www.btah.org/case-study/idaho-medicaid-home-and-community-based-services-care-cuts.html>
- [89] Benefits Tech Advocacy Hub. 2022. *Missouri Medicaid Home and Community Based Services Eligibility Issues*. Retrieved January 16, 2024 from <https://www.btah.org/case-study/missouri-medicaid-home-and-community-based-services-eligibility-issues.html>
- [90] Benefits Tech Advocacy Hub. 2022. *Understanding the Lifecycle of Benefits Technology*. Retrieved January 16, 2024 from <https://www.btah.org/lifecycle.html>
- [91] The Coalition Against Predictive Policing in Pittsburgh. 2020. *Predictive Policing in Pittsburgh: A Primer*. Retrieved January 16, 2024 from [https://capp-pgh.com/files/Primer\\_v1.pdf](https://capp-pgh.com/files/Primer_v1.pdf)
- [92] The Coalition Against Predictive Policing in Pittsburgh. 2020. *Responding to the “Completion” of Predictive Policing in Pittsburgh*. Retrieved January 16, 2024 from <https://capp-pgh.com/files/Metro21%20Counter-Statement.pdf>

- [93] Abigail Z. Jacobs and Hanna Wallach. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. ACM. <https://doi.org/10.1145/3442188.3445901>
- [94] Pratyusha Ria Kalluri, William Agnew, Myra Cheng, Kentrell Owens, Luca Soldaini, and Abeba Birhane. 2023. The Surveillance AI Pipeline. arXiv:2309.15084 [cs.CV]
- [95] Hamid Khan and Pete White. 2021. *Police Surveillance Can't Be Reformed. It Must Be Abolished*. Retrieved January 16, 2024 from <https://www.vice.com/en/article/xgzj7n/police-surveillance-cant-be-reformed-it-must-be-abolished>
- [96] Rebecca Klar. 2023. *Exclusive: Meta faces pressure to support independent audit of risk oversight committee*. Retrieved January 16, 2024 from <https://news.yahoo.com/exclusive-meta-faces-pressure-support-164507051.html>
- [97] Amy Kraft. 2016. *Microsoft shuts down AI chatbot after it turned into a Nazi*. Retrieved January 16, 2024 from <https://www.cbsnews.com/news/microsoft-shuts-down-ai-chatbot-after-it-turned-into-racist-nazi/>
- [98] Junhyup Kwon and Hyeong Yun. 2021. *AI Chatbot Shut Down After Learning To Talk Like a Racist Asshole*. Retrieved January 16, 2024 from <https://www.vice.com/en/article/akd4g5/ai-chatbot-shut-down-after-learning-to-talk-like-a-racist-asshole>
- [99] Colin Lecher. 2018. *What happens when an algorithm cuts your health care*. <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy>
- [100] Nicol Turner Lee, Paul Resnick, and Genie Barton. 2019. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms. <https://www.brookings.edu/articles/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>
- [101] Peter Lee. 2016. *Learning from Tay's introduction*. <https://blogs.microsoft.com/blog/2016/03/25/learning-tays-introduction/>
- [102] Calvin Liang. 2021. *Reflexivity, positionality, and disclosure in HCI*. Retrieved January 16, 2024 from <https://medium.com/@caliang/reflexivity-positionality-and-disclosure-in-hci-3d95007e9916>
- [103] Q. Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. arXiv:2306.01941 [cs.HC]
- [104] Kevin De Liban. 2018. *Comments re: Notice of Rule-Making for ARChoices Program*. Retrieved January 21, 2024 from <https://www.arkleg.state.ar.us/Home/FTPDocument?path=%2FAssembly%2FMeeting+Attachments%2F430%2F663%2FHandout+1+-+Legal+Aid-Kevin+De+Liban.pdf> Email providing public comment on proposed revisions to ARChoices program..
- [105] Gabriel Lima, Nina Grgic-Hlaca, Jin Keun Jeong, and Meeyoung Cha. 2023. Who Should Pay When Machines Cause Harm? Laypeople's Expectations of Legal Damages for Machine-Caused Harm. 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, 236–246. <https://doi.org/10.1145/3593013.3593992>
- [106] Antony Loewenstein. 2023. *The Palestine Laboratory: How Israel Exports the Technology of Occupation Around the World*. Verso. <https://www.versobooks.com/products/2684-the-palestine-laboratory>
- [107] Lynn Lomibao. 2021. Stop LAPD Spying Sues LAPD to Uncover Communications with UCLA Professor Who Founded PredPol, Inc. <https://stoplapdspying.org/stop-lapd-spying-sues-lapd-for-communications-with-widely-condemned-ucla-professor-who-founded-predpol-inc/>
- [108] My Ly. 2023. *Arkansas DHS agrees to pay \$460,000 to settle case over in-home care cuts*. Retrieved January 16, 2024 from <https://www.arkansasonline.com/news/2023/aug/09/arkansas-dhs-agrees-to-pay-460000-to-settle-case/>
- [109] Angelica Mari. 2022. São Paulo subway ordered to suspend use of facial recognition. <https://www.zdnet.com/article/sao-paulo-subway-ordered-to-suspend-use-of-facial-recognition/>
- [110] Jesse Marx and Lilly Irani. 2021. *Redacted*. Taller California. <https://www.printedmatter.org/catalog/58378/>
- [111] Carole McGranahan. 2016. Theorizing Refusal: An Introduction. *Cultural Anthropology* 31 (08 2016), 319–325. <https://doi.org/10.14506/ca31.3.01>
- [112] Sean McGregor. 2020. Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. arXiv:2011.08512 [cs.CY]
- [113] James Meadway. 2020. "Fuck the Algorithm": How A-Level Students Have Shown the Future of Protest. <https://novaramedia.com/2020/08/17/fuck-the-algorithm-how-a-level-students-have-shown-future-of-protest/>
- [114] Dhruv Mehrotra and Dell Cameron. 2023. *The Maker of ShotSpotter Is Buying the World's Most Infamous Predictive Policing Tech*. Retrieved January 16, 2024 from <https://www.wired.com/story/soundthinking-geolitic-acquisition-predictive-policing/>
- [115] Brian Merchant. 2023. *Blood in the Machine: The Origins of the Rebellion Against Big Tech*. Little, Brown and Company.
- [116] Jacob Metcalf. 2023. *What federal agencies can learn from NYC's AI Hiring Law*. Retrieved January 16, 2024 from <https://thehill.com/opinion/technology/4360523-what-federal-agencies-can-learn-from-new-york-citys-ai-hiring-law/>
- [117] Jacob Metcalf, Ranjit Singh, Emanuel Moss, Emnet 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. <https://doi.org/10.1145/3593013.3594092>
- [118] Microsoft. 2022. *Microsoft Responsible AI Standard, v2 (General Requirements)*. Retrieved July 21, 2023 from <https://blogs.microsoft.com/wp-content/uploads/prod/sites/5/2022/06/Microsoft-Responsible-AI-Standard-v2-General-Requirements-3.pdf>
- [119] Yeshimabeit Milner. 2020. *Abolish Big Data*. <https://medium.com/@YESHICAN/abolish-big-data-ad0871579a41>
- [120] Aaron Mok. 2023. *AI is expensive. A search on Google's chatbot Bard costs the company 10 times more than a regular one, which could amount to several billion dollars*. Retrieved January 16, 2024 from <https://www.businessinsider.com/ai-expensive-google-chatbot-bard-may-cost-company-billions-dollars-2023-2>
- [121] Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish, and Jacob Metcalf. 2021. *Assembling accountability: algorithmic impact assessment for the public interest*. Retrieved July 21, 2023 from <https://datasociety.net/wp-content/uploads/2021/06/Assembling-Accountability.pdf>
- [122] Mozilla. 2023. *Auditing AI: Announcing the 2023 Mozilla Technology Fund Cohort*. Retrieved January 16, 2024 from <https://foundation.mozilla.org/en/blog/auditing-ai-announcing-the-2023-mozilla-technology-fund-cohort/>
- [123] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press. <http://www.jstor.org/stable/j.ctt1pw9w5>
- [124] Oryem Nyeko. 2023. Uganda: Rights Concerns Over License Plate Tracking | Human Rights Watch. <https://www.hrw.org/news/2023/11/14/uganda-rights-concerns-over-license-plate-tracking>
- [125] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [126] The Government of Canada. 2023. *Algorithmic Impact Assessment tool*. Retrieved July 21, 2023 from <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html>
- [127] The White House Office of Science and Technology Policy. 2022. *Blueprint for an AI Bill of Rights: A Vision for Protecting Our Civil Rights in the Algorithmic Age*. Retrieved October 2, 2023 from <https://www.whitehouse.gov/ostp/news-updates/2022/10/04/blueprint-for-an-ai-bill-of-rights-a-vision-for-protecting-our-civil-rights-in-the-algorithmic-age/>
- [128] UN OHCHR. 2022. OHCHR Assessment of Human Rights Concerns in the Xinjiang Uyghur Autonomous Region, People's Republic of China.
- [129] Victor Ojewale, Ryan Steed, Briana Vecchione, Abeba Birhane, and Inioluwa Deborah Raji. 2024. Towards AI Accountability Infrastructure: Gaps and Opportunities in AI Audit Tooling. arXiv:2402.17861 [cs.CY]
- [130] Kalia Orphanou, Jahna Otterbacher, Styliani Kleanthous, Khuyagbaatar Batsuren, Fausto Giunchiglia, Veronika Bogina, Avital Shulner Tal, Alan Hartman, and Tsvi Kuflik. 2022. Mitigating Bias in Algorithmic Systems—A Fish-eye View. *ACM Comput. Surv.* 55, 5, Article 87 (dec 2022), 37 pages. <https://doi.org/10.1145/3527152>
- [131] Phil Pennington. 2023. Facial recognition: Officials yet to meet obligation to seek views of Māori - documents. <https://www.rnz.co.nz/news/national/501761/facial-recognition-officials-yet-to-meet-obligation-to-seek-views-of-maori-documents>
- [132] Nick Perry. 2023. New Zealand debates whether ethnicity should be a factor for surgery waitlists. <https://apnews.com/article/new-zealand-surgery-ethnicity-algorithm-maori-1b44026f2661772a7eb3bd4444619446>
- [133] Dana Pessach and Erez Shmueli. 2022. A Review on Fairness in Machine Learning. *ACM Comput. Surv.* 55, 3, Article 51 (feb 2022), 44 pages. <https://doi.org/10.1145/3494672>
- [134] Brendan Pierson and Brendan Pierson. 2023. Lawsuit claims UnitedHealth AI wrongfully denies elderly extended care. *Reuters* (Nov 2023). <https://www.reuters.com/legal/lawsuit-claims-unitedhealth-ai-wrongfully-denies-elderly-extended-care-2023-11-14/>
- [135] Jon Porter. 2020. *UK ditches exam results generated by biased algorithm after student protests*. Retrieved January 16, 2024 from <https://www.theverge.com/2020/8/17/21372045/uk-a-level-results-algorithm-biased-coronavirus-covid-19-pandemic-university-applications>
- [136] Gabriel Puron-Cid and J. Ramon Gil-Garcia. 2022. Are Smart Cities Too Expensive in the Long Term? Analyzing the Effects of ICT Infrastructure on Municipal Financial Sustainability. *Sustainability* 14, 10 (2022). <https://doi.org/10.3390/su14106055>
- [137] 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 (dec 2022), 101–108. <https://doi.org/10.1145/3571151>
- [138] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *Proceedings of the 2022 ACM Conference*

- on Fairness, Accountability, and Transparency (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 959–972. <https://doi.org/10.1145/3531146.3533158>
- [139] Dillon Reisman, Jason Schultz, Kate Crawford, and Meredith Whittaker. 2018. *Algorithmic Impact Assessments: A Practical Framework for Public Agency*. Retrieved July 21, 2023 from <https://www.nist.gov/system/files/documents/2021/10/04/aiareport2018.pdf>
- [140] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. arXiv:1602.04938 [cs.LG]
- [141] Kat Roemmich, Tillie Rosenberg, Serena Fan, and Nazanin Andalibi. 2023. Values in Emotion Artificial Intelligence Hiring Services: Technosolutions to Organizational Problems. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 109 (apr 2023), 28 pages. <https://doi.org/10.1145/3579543>
- [142] Casey Ross. 2021. *Epic's AI algorithms, shielded from scrutiny by a corporate firewall, are delivering inaccurate information on seriously ill patients*. Retrieved January 21, 2024 from <https://www.statnews.com/2021/07/26/epic-hospital-algorithms-sepsis-investigation/>
- [143] Tate Ryan-Mosley. 2023. *An algorithm intended to reduce poverty might disqualify people in need*. Retrieved January 16, 2024 from <https://www.technologyreview.com/2023/06/13/1074551/an-algorithm-intended-to-reduce-poverty-in-jordan-disqualifies-people-in-need/>
- [144] Tate Ryan-Mosley. 2023. *How face recognition rules in the US got stuck in political gridlock*. <https://www.technologyreview.com/2023/07/24/1076668/how-face-recognition-rules-in-the-us-got-stuck-in-political-gridlock/>
- [145] Nete Schwennesen. 2019. Algorithmic assemblages of care: imaginaries, epistemologies and repair work. *Sociology of Health & Illness* 41 (10 2019), 176–192. <https://doi.org/10.1111/1467-9566.12900>
- [146] Joey Scott. 2023. *LAPD Is Using Israeli Surveillance Software That Can Track Your Phone and Social Media*. Retrieved January 16, 2024 from <https://knock-la.com/lapd-is-using-israeli-surveillance-software-that-can-track-your-phone-and-social-media/#:~:text=During%20a%202014%20trip%20to,would%20be%20using%20all%20three>
- [147] Nick Seaver. 2019. *Knowing Algorithms*. Princeton University Press, Princeton, 412–422. <https://doi.org/doi:10.1515/9780691190600-028>
- [148] Andrew D. Selbst and Solon Barocas. 2018. The Intuitive Appeal of Explainable Machines. *Fordham Law Review* 87 (2018), 1085. <https://api.semanticscholar.org/CorpusID:59548063>
- [149] Andrew D. Selbst, Suresh Venkatasubramanian, and I. Elizabeth Kumar. 2024. Deconstructing Design Decisions: Why Courts Must Interrogate Machine Learning and Other Technologies. *Ohio State Law Journal* 85 (2024). <https://ssrn.com/abstract=4564304>
- [150] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Roshtamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (, Montréal, QC, Canada), (AIES '23). Association for Computing Machinery, New York, NY, USA, 723–741. <https://doi.org/10.1145/3600211.3604673>
- [151] 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. <https://doi.org/10.1145/3479577>
- [152] Audra Simpson. 2007. On Ethnographic Refusal: Indigeneity, 'Voice' and Colonial Citizenship. *Junctures* (2007).
- [153] Michael Sisitzky and Ben Schaefer. 2021. *The NYPD Published its Arsenal of Surveillance Tech. Here's What We Learned*. Retrieved January 16, 2024 from <https://www.nyclu.org/en/news/nypd-published-its-arsenal-surveillance-tech-heres-what-we-learned>
- [154] Till Speicher, Hoda Heidari, Nina Grgic-Hlaca, Krishna P. Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. 2018. A Unified Approach to Quantifying Algorithmic Fairness: Measuring Individual & Group Unfairness via Inequality Indices. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (London, United Kingdom) (KDD '18). Association for Computing Machinery, New York, NY, USA, 2239–2248. <https://doi.org/10.1145/3219819.3220046>
- [155] Logan Stapleton, Min Hun Lee, Diana Qing, Marya Wright, Alexandra Chouldchova, Ken Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. Imagining New Futures beyond Predictive Systems in Child Welfare: A Qualitative Study with Impacted Stakeholders. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 1162–1177. <https://doi.org/10.1145/3531146.3533177>
- [156] Varshini Subhash, Zixi Chen, Marton Havasi, Weiwei Pan, and Finale Doshi-Velez. 2022. What Makes a Good Explanation?: A Harmonized View of Properties of Explanations. In *Progress and Challenges in Building Trustworthy Embodied AI*. <https://openreview.net/forum?id=YDyLZWwpBK2>
- [157] Harini Suresh and John Gutttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (, NY, USA) (EAAAMO '21). Association for Computing Machinery, New York, NY, USA, Article 17, 9 pages. <https://doi.org/10.1145/3465416.3483305>
- [158] The Automated Regional Justice Information System. 2021. *TACIDS: Tactical Identification System Using Facial Recognition*. Retrieved January 16, 2024 from <https://voiceofsandiego.org/wp-content/uploads/2021/04/TACIDS-Final-Report-FINAL.pdf>
- [159] Kim TallBear. 2013. *Native American DNA: Tribal Belonging and the False Promise of Genetic Science*. U of Minnesota Press.
- [160] A. Toh and Human Rights Watch. 2023. *Automated Neglect: How the World Bank's Push to Allocate Cash Assistance Using Algorithms Threatens Rights*. Human Rights Watch. <https://books.google.com/books?id=rNMG0AEACAAJ>
- [161] Rhema Vaithianathan, Emily Kulick, Emily Putnam-Hornstein, and Diana Benavides Prado. 2019. *Allegheny Family Screening Tool: Methodology, Version 2*. Retrieved January 21, 2024 from [https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/Methodology-V2-from-16-ACDHS-26\\_PredictiveRisk\\_Package\\_050119\\_FINAL-7.pdf](https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/Methodology-V2-from-16-ACDHS-26_PredictiveRisk_Package_050119_FINAL-7.pdf)
- [162] Matias Valderrama. 2021. Sistema Alerta Niños y la predicción del riesgo de vulneración de derechos de la infancia. *Derechos Digitales* (2021).
- [163] C van Veen. 2020. Landmark judgment from The Netherlands on digital welfare states and human rights. *Open Global Rights* (2020).
- [164] Joana Varon and Paz Peña. 2022. Not My A.I.: Towards Critical Feminist Frameworks to Resist Oppressive A.I. Systems. *Carr Center Discussion Paper Series* (2022). <https://carrcenter.hks.harvard.edu/publications/notmyai>
- [165] James Vincent. [n. d.]. Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day — theverge.com. <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>. [Accessed 22-01-2024].
- [166] Heather Vogell, Haru Coryne, and Ryan Little. 2022. *Rent Going Up? One Company's Algorithm Could Be Why*. Retrieved January 16, 2024 from <https://www.propublica.org/article/yieldstar-rent-increase-realtor-rent>
- [167] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2023. Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms that Optimize Predictive Accuracy. 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, 626. <https://doi.org/10.1145/3593013.3594030>
- [168] Fulton Wang and Cynthia Rudin. 2015. Falling Rule Lists. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 38)*, Guy Lebanon and S. V. N. Vishwanathan (Eds.). PMLR, San Diego, California, USA, 1013–1022. <https://proceedings.mlr.press/v38/wang15a.html>
- [169] Nina Wang, Allison McDonald, Daniel Bateyko, and Emily Tucker. 2022. American Dragnet: Data-Driven Deportation in the 21st Century. (2022). Retrieved January 16, 2024 from <https://americandragnet.org/>
- [170] Elizabeth Warren. 2023. *Warren, Lawmakers Urge Justice Department to Review YieldStar, Warn of De-Facto Price Setting and Collusion After Senate Investigation*. Retrieved January 21, 2024 from <https://www.warren.senate.gov/oversight/letters/warren-lawmakers-urge-justice-department-to-review-yieldstar-warn-of-de-facto-price-setting-and-collusion-after-senate-investigation> Press Release.
- [171] Human Rights Watch. 2023. *Automated Neglect: How the World Bank's Push to Allocate Cash Assistance Using Algorithms Threatens Rights*. Retrieved January 16, 2024 from <https://www.hrw.org/report/2023/06/13/automated-neglect/how-world-banks-push-allocate-cash-assistance-using-algorithms>
- [172] Hilde Weerts, Miroslav Dudik, Richard Edgar, Adrin Jalali, Roman Lutz, and Michael Madaio. 2023. Fairlearn: Assessing and Improving Fairness of AI Systems. arXiv:2303.16626 [cs.LG]
- [173] Emma Weil and Elizabeth Edwards. 2023. Using Technical Skills to Fight Actual Public Benefits Cuts and Austerity Policies, with the Benefits Tech Advocacy Hub. (2023). <https://www.youtube.com/watch?v=ZELeRp74PE> ACM FAccT Tutorial.
- [174] Kate Wells. 2023. *An eating disorders chatbot offered dieting advice, raising fears about AI in health*. Retrieved January 16, 2024 from <https://www.npr.org/sections/health-shots/2023/06/08/1180838096/an-eating-disorders-chatbot-offered-dieting-advice-raising-fears-about-ai-in-health>
- [175] 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. <https://doi.org/10.1145/3411764.3445314>
- [176] Zack Whittaker. 2019. *Amazon shareholders reject facial recognition sale ban to governments*. Retrieved January 16, 2024 from <https://techcrunch.com/2019/05/22/amazon-reject-facial-recognition-proposals/>
- [177] Christopher Wilkinson and Dorothy Lukens. 2023. *The Growing Regulation of AI-Based Employment Decision Tools*. Retrieved January 16, 2024

- from <https://www.perkinscoie.com/en/news-insights/the-growing-regulation-of-ai-based-employment-decision-tools.html>
- [178] Timothy Williams. 2015. *Facial Recognition Software Moves From Overseas Wars to Local Police*. Retrieved January 16, 2024 from [https://www.nytimes.com/2015/08/13/us/facial-recognition-software-moves-from-overseas-wars-to-local-police.html?\\_r=1](https://www.nytimes.com/2015/08/13/us/facial-recognition-software-moves-from-overseas-wars-to-local-police.html?_r=1)
- [179] Tom Wilson and Madhumita Murgia. 2019. *Uganda confirms use of Huawei facial recognition cameras*. Retrieved January 16, 2024 from <https://www.ft.com/content/e20580de-c35f-11e9-a8e9-296ca66511c9>
- [180] Andrew Wong, Erkin Otles, John P. Donnelly, Andrew Krumm, Jeffrey McCullough, Olivia DeTroyer-Cooley, Justin Pestru, Marie Phillips, Judy Konye, Carleen Penzo, Muhammad Ghous, and Karandeep Singh. 2021. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. *JAMA Internal Medicine* 181, 8 (08 2021), 1065–1070. <https://doi.org/10.1001/jamainternmed.2021.2626> arXiv:[https://jamanetwork.com/journals/jamainternalmedicine/articlepdf/2781307/jamainternal\\_wong\\_2021\\_oi\\_210027\\_1627674961.11707.pdf](https://jamanetwork.com/journals/jamainternalmedicine/articlepdf/2781307/jamainternal_wong_2021_oi_210027_1627674961.11707.pdf)
- [181] Sarah Wu. 2019. *Somerville City Council passes facial recognition ban*. Retrieved January 16, 2024 from <https://www.bostonglobe.com/metro/2019/06/27/somerville-city-council-passes-facial-recognition-ban/>
- [182] Chloe Xiang. 2023. *Eating Disorder Helpline Disables Chatbot for 'Harmful' Responses After Firing Human Staff*. Retrieved January 16, 2024 from <https://www.vice.com/en/article/qjvk97/eating-disorder-helpline-disables-chatbot-for-harmful-responses-after-firing-human-staff>
- [183] Kyra Yee, Uthaipon Tantipongpipat, and Shubhanshu Mishra. 2021. Image Cropping on Twitter: Fairness Metrics, their Limitations, and the Importance of Representation, Design, and Agency. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 450 (oct 2021), 24 pages. <https://doi.org/10.1145/3479594>
- [184] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supriya Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, and John Zimmerman. 2023. Creating Design Resources to Scaffold the Ideation of AI Concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (, Pittsburgh, PA, USA.) (*DIS '23*). Association for Computing Machinery, New York, NY, USA, 2326–2346. <https://doi.org/10.1145/3563657.3596058>
- [185] Jonathan Zong and J. Nathan Matias. 2023. Data Refusal From Below: A Framework for Understanding, Evaluating, and Envisioning Refusal as Design. *ACM J. Responsib. Comput.* (oct 2023). <https://doi.org/10.1145/3630107> Just Accepted.

## A METHODOLOGY (EXTENDED)

*Keyword searches for international cases.* To identify additional cases that met our criteria outside of the U.S., we conducted additional keyword searches on Google of names of regions/countries and our keywords from before, such as “*China algorithmic harm*”.

### A.1 Thematic analysis (continued)

*Closed metadata coding.* In addition to conducting a bottom-up thematic analysis where we performed *open coding* to create event timelines, we also did some *closed coding* where we recorded the below metadata for all cases:

- Case domain
- Name of the organization overseeing the algorithm’s deployment
- Name of any organizations involved with developing the algorithm (if different from deploying organization)
- Region(s) in which the algorithm is developed and deployed
- What does the algorithm take as input, and what does it predict? How are these predictions then used to make decisions?
- Was the algorithm ever deployed?
- Who was (potentially) impacted by the algorithm?
- What were the algorithms’ (potential) harms? (We later mapped these using the categories from Shelby et al. [150]’s taxonomy.)
- Who critiqued the algorithm, and why?
- Was the algorithm ever abandoned? For abandoned algorithms,
  - How was the algorithm abandoned? (*i.e.*, by the algorithm owners choice? through policy?)
  - Were there any caveats, *e.g.*, did the algorithm continue to be used in any other contexts?
  - When in the algorithm’s development lifecycle did abandonment occur?
  - Is there a publicly available statement of why the algorithm was abandoned?

We referred to answers to these questions explicitly when developing our process-oriented view (*e.g.*, in understanding the different harms that occurred, or mechanisms that algorithms were abandoned) and in identifying key socio-technical factors that differentiated algorithms that could be abandoned quickly.

*Example annotation.* For illustrative purposes, we walk-through how we coded an event timeline for an example case: the National Eating Disorder Association (NEDA)’s Tessa Wellness Bot (see the final timeline in Section B below). We became aware of this case by searching the AI Incident Database [112]. We began learning about the case by reading the news articles linked on its AIID webpage. While reading these articles, we closed-coded case metadata, and also noted any other important events that appeared to be relevant to the algorithmic harm(s) and decisions made by the model owners. For example, we learned from a Vice article [182] that shortly before the Tessa bot came under criticism for its potential to inflict harm, a news story broke that NEDA had decided to fire all of its hotline staffers following their decision to unionize, and intended to “*replace them*” by instead directing people to Tessa, a chatbot that had been in operation since February 2022. From this article, we also learned the name of the organization that developed Tessa, and gained insight into Tessa’s design (*e.g.*, how users interacted with Tessa online). Two more recent articles linked on the AIID



webpage [83, 174] interviewed users who took to social media to criticize the prominent NEDA firing and the Tessa chatbot. These articles directly linked to social media posts where we could directly engage with users' critiques and discussions about the algorithm, such as Instagram posts (link). We referenced these interviews and critiques to better understand the nature of the algorithm's potential for harm and who might be adversely impacted. Finally, we also referred to these articles and primary sources (e.g., NEDA's Instagram announcement, link) to understand the organization's decision to abandon the algorithm. We conducted a final confirmatory Google search to check if the Tessa Bot (or any similar service created by Cass) is still in operation today.

## B EVENT TIMELINE FIGURE (EXPANDED)

We provide a detailed description of each labeled event in the event timeline visualization.

### Jordan's Takaful Cash Assistance Program [143, 171]

- (1) May 2019, release. The Jordanian government launches the Takaful cash transfer program in partnership with funding from the World Bank.
- (2) October 2021, discovery of harm. Human Rights Watch (an international NGO) begins conducting investigative interviews with Jordan residents about their experiences with Takaful (diagnosis).
- (3) October 2022. Jordan's National Aid Fund (NAF), the social protection agency administering Takaful, participates in an interview with the Human Rights Watch (dialogue), in which they reveal that they changed the algorithm to stop predicting the benefit *amount* after complaints that this was "arbitrary and confusing" (decision to repair). The algorithm now predicts benefit *eligibility* instead.
- (4) June 2023. The Human Rights Watch publishes the results of their investigation, which gets picked up and covered by several other large media outlets (dissemination).

### Metro21's Predictive Policing Pilot in Pittsburgh [91, 92]

- (1) February 2017, release. Pittsburgh begins a hotspot-based predictive policing pilot with researchers at CMU (through the Metro21 Smart Cities Institute). They do not notify or involve the public in the development of the program.
- (2) October 2018, discovery of existence. CMU researchers publish a paper about an evaluation of their predictive policing pilot. The public is made aware of its existence. A concerned group of CMU students begin to organize and file records requests (diagnosis).
- (3) June 2020, In the wake of mass protests after the murder of George Floyd, CMU students form the Coalition Against Predictive Policing, where they hold virtual teach-ins and create a petition to the city (dissemination). The petition gathers 3.6k signatures. The City states that the "Metro21 pilot has ended" and that they are no longer using predictive policing.
- (4) September 2020, Pittsburgh City Council votes to ban predictive policing in perpetuity (decision, death).

### The National Eating Disorder Association (NEDA)'s Tessa Chatbot [83, 174, 182]

- (1) February 2022, release. NEDA releases Tessa, a free AI "wellness bot" to help users "discover coping skills" for eating and body concerns.
- (2) May 2023, discovery of potential harm. NPR runs an exposé on how NEDA fired all helpline staff following their decision to unionize, and intends to "replace them" with Tessa.
- (3) May 2023. Five days after the NPR exposé, a fat-positive activist shares screenshots of harmful interactions with Tessa, in which the bot gave advice that might trigger an eating disorder (e.g., to count calories and lose weight). Her post goes viral (dissemination). A NEDA VP comments on the author's post that she is lying and asks the author to provide her with more evidence that the interaction actually occurred (dialogue).
- (4) May 2023. Two days later, NEDA announces that Tessa will be shut down (decision, death).

## C CASE LIST

Interact with and download our case dataset online at this (link)!

<b>Harm Category</b>	<b>Definition</b>	<b>Examples</b>
<b>Representational</b>	Socially constructed beliefs and unjust hierarchies about social groups are reflected in system inputs and outputs.	Microsoft’s Tay Chatbot [165] began tweeting offensive content after interacting with trolls ( <i>demeaning, stereotyping</i> ).
<b>Allocative</b>	A system withholds information, opportunities, or resources from historically marginalized groups; often in domains that affect material wellbeing.	Arkansas’ RUGS algorithm [104] results in unexpected and severe cuts to beneficiaries’ home care hours ( <i>opportunity loss</i> ).
<b>Quality of Service</b>	The system is dysfunctional, or underperforms disproportionately for certain groups of people along social categories of difference like disability, ethnicity, etc.	Zoom’s Virtual Background feature [53] is dysfunctional for users with darker skin ( <i>alienation, increased labor</i> ).
<b>Interpersonal</b>	The system adversely shapes relations between people or communities.	The Allegheny County Family Screening Tool (AFST) [69] assigns higher risk scores to parents with developmental disabilities, which may influence Allegheny DHS to open an investigation or deem them negligent, resulting in separation ( <i>loss of agency/control, privacy violation</i> ).
<b>Social System</b>	The system has adverse macro-level effects to knowledge systems, cultural impacts, political or civic systems, economic systems, or the environment.	San Diego’s TACIDS algorithm [178] was used by cops or immigration officers to scan and instantly identify people, often without their consent ( <i>political and civic harms</i> ).

**Table 2: We map reports of potential or realized harms from the algorithmic systems we identified, to the five major types of sociotechnical harms from the taxonomy proposed by Shelby et al. [150]. We also note the relevant harm "sub-types" for each case in italics, after a brief description of the harm that occurred. Definitions of each category (middle column) are taken verbatim from Shelby et al. [150].**

	<i>Case Name</i>	<i>Case Description</i>
<b>Online platforms</b>		
1	Twitter Cropping Algorithm [183]	Users hypothesized that an algorithm that cropped images to show feed previews compatible with device aspect ratios was more likely to crop out Black people (or other minority groups, like women) in images with multiple people.
2	Zoom Virtual Background [53]	Zoom's virtual background feature is more likely to be dysfunctional ( <i>i.e.</i> , hide the body of or include parts of the background) for users with darker skin.
<b>Chatbot</b>		
3	Microsoft Tay Bot [165]	Microsoft's Tay chatbot began tweeting offensive content after interacting with users who wanted it to do so.
4	U.S. National Eating Disorder Association (NEDA)'s "Tessa" Wellness Bot [174]	Shortly NEDA fired all of their hotline staff (following their decision to unionize), activists posted on social media about how Tessa (NEDA's chatbot) gave harmful and triggering advice.
5	Meta's Galactica [6]	Within three days of releasing a public demo of their Galactica language model, Meta took down the public demo and re-branded Galactica as "available for researchers" only.
6	Scatter Lab's "Lee-Luda" Chatbot [86]	Seoul-based startup's AI-chatbot "Lee-Luda" was removed from Facebook messenger 20 days upon launch after users complained it was using offensive language, discriminatory and generated hate speech towards LGBTQ folks and people with disabilities.
<b>Employment</b>		
7	Amazon AI Recruiting Tool [46]	Amazon disbanded an internal team whose goal was to train ML to screen resumes due to algorithmic gender bias.
<b>Child Welfare</b>		
8	Oregon DHS's Safety at Screening Tool [7]	Oregon DHS decided to stop using an AI-based screening tool (which is modeled after the AFST), in the wake of an April AP news article reporting that these tools have the ability to "heighten racial disparities in the child welfare system".
9	Pittsburgh Allegheny County Family Screening Tool (AFST) [69]	External researchers, journalists, and advocacy groups have raised concerns about the AFST, including that it is more likely to investigate (and ultimately deem negligent) (1) poor families, (2) Black families, and (3) people with disabilities.

	<i>Case Name</i>	<i>Case Description</i>
<b>Child Welfare</b>		
10	Chile Local Childhood Office's Sistema Alerta Niñez ("Childhood Alert System") [162]	Chile's Ministry of Social Development and Family launched the program to strengthen childhood protection with an algorithm that provided risk scores used to determine necessary intervention, but rejected an independent external audit's results revealing the system wrongly analyzed variables associated with direct measures of poverty rather than safety, targeting and perpetuating over-intervention in poor neighborhoods.
11	New Zealand's "Vulnerable Children PRM" (Predictive Risk Modeling) [31]	New Zealand's Social Development Minister intervened and halted the deployment of the algorithm after seeing major concerns from additional research during the request for its ethics approval for further study.
<b>Family &amp; Social Services</b>		
12	Teen Pregnancy Risk Prediction in Argentina [21]	The Applied Artificial Intelligence Laboratory (LIAA) in the University of Buenos Aires published concerns in a detailed technical analysis and review of an algorithm developed through a partnership between Microsoft and province of Salta's Ministry of Early Childhood to predict teenage pregnancy.
13	The Spanish Ministry of the Interior's "VioGén" Gender Violence Recidivism Risk Assessment [5]	The Eticas Foundation partnered with victims to publish an external audit after the Spanish Ministry declined requests of auditing the algorithm that underestimated level of risk in domestic and intimate partner violence cases, leading to adverse outcomes.
<b>Benefits</b>		
14	Arkansas Department of Human Services' Homecare Medicaid Program "RUGS" Assessment System [87]	Arkansas switched to using a new algorithm called "RUGS" to assess the hours of in-home care that residents are eligible for. The algorithm resulted in severe cuts to the number of hours allocated to most residents.
15	Missouri "NF LOC" Home Care [89]	The State of Missouri released (in advance) their plan to switch to using an algorithm called "NF LOC" to assess eligibility for home-care services. They released the algorithm for public comment, and external researchers found that it would result in severe cuts for a large group of residents (and also had other logical errors).
16	Idaho Department of Health and Welfare's Care Recommendation Algorithm [88]	The Idaho Department of Health and Welfare used an algorithm to assess individual Medicaid budgets available to people with disabilities in their home-based services waiver program. A class-action lawsuit, <i>KW vs. Armstrong</i> , resulted in the algorithm being made public, and the algorithm was deemed unconstitutional and "arbitrary". The case was settled so that people whose budgets were cut would continue to get their original budgets.

	<i>Case Name</i>	<i>Case Description</i>
	<b>Benefits</b>	
17	World Bank's Automated Cash Transfer Program "Takaful" in Jordan [160]	Human Rights Watch published an investigation report that revealed how the World Bank's algorithm for distributing and ranking financial assistance through a cash transfer program spearheaded by Jordan government's social protection agency was wrongfully disqualifying those in need. The algorithm used 57 factors that they declined to publicly disclose.
18	Netherlands' System Risk Indication (SyRI) [163]	Local civil rights organizations filed a case against the state to stop using an algorithm to detect welfare benefits fraud. A Dutch court ruled that the algorithm violated the European Convention on Human Rights (ECHR) protections and halted its operation.
	<b>Healthcare</b>	
19	EPIC' Sepsis Prediction Tool [142]	Epic, an American healthcare software company (and prominent EHR vendor), launched the Epic Sepsis Model (ESM) in 2017, a proprietary linear model that alerts patients at high risk of sepsis. From 2021 onward, external and independent audits or investigations showed that the model was flawed; <i>e.g.</i> , it had lower AUC than reported, had a high false alarm rate, etc.
20	US HHS and UNOS (the United Network for Organ Sharing)'s "Acuity Circles Policy" [11]	A joint investigation by the Washington Post and The Markup exposed how a new shift in the UNOS allocation algorithm's criteria in removing distance as a factor for organ donation and prioritizing sickest patients regardless of location, drastically decreased the number of transplant surgeries and increased deaths in poorer states.
21	UnitedHealth & NaviHealth's "nH predict" Post-Acute Care Management [134]	Families of deceased beneficiaries filed a lawsuit in the U.S. District Court for the District of Minnesota against a insurance company's healthcare algorithm for determining post-acute care. The algorithm prematurely discharged patients from facilities against doctor recommendations and discontinued payment for necessary services and claims, targeting those on Medicaid.
22	U.S. Health Systems' High-risk Healthcare Management" Prediction [125]	Researchers at the UC Berkeley School of Public Health exploited a rich dataset that provided insight into a live, scaled algorithm that had been deployed nationwide, and found it perpetuated significant racial bias through using health care costs to determine who was most likely to benefit from care management programs.
23	New Zealand Health System's Equity Adjuster-Waitlist Tool [132]	The ACT New Zealand political party and other clinicians started a petition against Auckland Health Minister's efforts to make healthcare more equitable for low-income, rural, and indigenous Maori and Pasifika communities through an algorithm that adjusted patients' position in elective surgery waitlists, claiming it was racial discrimination against other New Zealanders.

	<i>Case Name</i>	<i>Case Description</i>
<b>Property Tech</b>		
24	RealPage's YieldStar Algorithm [166]	RealPage, a Texas-based company that provides property management software for landlords, developed an algorithm, YieldStar, that suggests daily prices for open rental units. Critics claim that the software brings anti-trust concerns and has resulted in rent inflation across the U.S.
<b>Education</b>		
25	UK Grading Tool [113]	In the 2020 pandemic, UK-based boards decided to use an algorithm to grade students' GCE A-level exams (international certificates which are important for university placements and scholarship opportunities). Worldwide protests broke out when the algorithmically calculated results were announced, which resulted in the grades being globally rescinded and re-assigned one week later.
26	SAS Institute's Educational Value-Added Assessment System (EVAAS) in Houston [37]	Houston Independent School District ruled in the settlement of a lawsuit brought by Houston and the American Federation of Teachers, to terminate the use of EVAAS, which used a student's performance on prior standardized tests to predict academic growth in the current year, as basis for evaluating teacher effectiveness.
<b>Facial Recognition</b>		
27	San Diego County Facial Recognition Tool (TACIDS) [158]	San Diego County had a regional database known as the Tactical Identification System (TACIDS), which contained 1.8million booking photos and was used by 30 law enforcement agencies (including ICE). Use of the database was ruled illegal by CA state AB 1215.
28	NYPD using Clearview AI [76]	Released public records revealed that the NYPD had a formal relationship with facial recognition company Clearview AI, including a "trial" vendor agreement, despite the PD's remarks that they had never had a relationship with the vendor. The NYPD released a formal facial recognition policy that banned use of Clearview in spring of 2020.
29	Detroit Project Green Light + DataWorks Plus [78]	The Detroit PD signed a contract with a for-profit vendor (DataWorks Plus) to begin running facial recognition on live video streams from surveillance cameras (from the "Project Green Light" initiative), and other surveillance cameras (e.g., drones).
30	City of Somerville's Face Surveillance Full Ban Ordinance [181]	A unanimous (11-0) vote by Somerville City Council passed an ordinance to ban all use of FR technology by city employees (including, but not limited to, the city of Somerville Police).
31	Facial Recognition in NYC Atlantic Plaza Towers [66]	Residents of a Brooklyn apartment complex partnered with a legal group to resist their landlord's plan to install facial recognition technology as a means to access the building.

	<i>Case Name</i>	<i>Case Description</i>
	<b>Facial Recognition</b>	
32	RiteAid's Enrollement Database for Surveillance Systems [48]	RiteAid, a large U.S. drugstore chain, deployed facial recognition surveillance systems in over 200 stores. Employees used the FR to surveil customers (e.g., "persons of interest" who might attempt to engage in criminal activity). RiteAid stopped using FR after a Reuters investigation broke in 2020, and in 2023 were issued a five-year ban on their use of FR by the FTC.
33	Hyderabad's #BantheScan Campaign [4]	India's Internet Freedom Foundation partnered with Amnesty International in kickstarting the #BantheScan campaign globally and locally to resist law enforcement use of facial recognition.
34	São Paulo Metro Intelligent Security System [109]	The city judge of São Paulo ruled that the proposed facial recognition system did not meet the legal requirements under Brazil's General Data Protection Law (LGPD) in a lawsuit brought by the public defender and multiple human rights and civil advocacy organizations.
35	Chinese Surveillance in the Xinjiang Uyghur Autonomous Region [128]	The Office of the United Nations High Commissioner for Human Rights published an assessment of concerns against police use of facial recognition technology to surveil and track Uyghur people across the Xinjiang region in China.
36	New Zealand Identity Check Program [131]	Data specialists and ethicists from New Zealand's indigenous Maori communities publicly commented and accused the government of ignoring their requests of involvement in improving the facial recognition technology with a 45% failure rate for not only security surveillance, but also fraud detection.
37	Uganda's "Intelligent Transport Monitoring System" [124]	Human right organizations, local politicians, and other activists spoke out against the multi-purpose technology that will monitor traffic and public spaces by electronic license plate recognition, other high-density camera surveillance, and facial recognition for tracking vehicle locations for criminal activity and investigations, political opponents, and protests.
	<b>Predictive Policing</b>	
38	UCLA PredPol [107]	The LAPD began a 2011 contract with PredPol, Inc (which began as a research project with a UCLA professor), which they used for 9 years until the contract was not renewed in 2019. A data leak showed that their algorithms disproportionately targeted Black, brown, and poor neighborhoods.
39	City of Pittsburgh [91]	Pittsburgh City Council voted to ban "obtaining, retaining, accessing, or using FR or predictive policing technology" after facing criticism from student activists for a predictive policing experiment in collaboration with researchers at Carnegie Mellon (Metro21).
	<i>Case Name</i>	<i>Case Description</i>
	<b>Predictive Policing</b>	
40	Germany's Predictive Policing Software Program "HessenDATA" [32]	In a case brought by the brought by the German Society for Civil Rights (GFF), the German Federal Constitutional Court ruled that a predictive policing algorithm developed by U.S. company Palantir was unconstitutional on the basis of the right to informational self-determination under the country's fundamental privacy rights.