

Algorithmic Harms in Child Welfare: Uncertainties in Practice, Organization, and Street-level Decision-Making

DEVANSH SAXENA, Carnegie Mellon University, USA

SHION GUHA, University of Toronto, Canada

Algorithms in public services such as child welfare, criminal justice, and education are increasingly being used to make high-stakes decisions about human lives. Drawing upon findings from a two-year ethnography conducted at a child welfare agency, we highlight how algorithmic systems are embedded within a complex decision-making ecosystem at critical points of the child welfare process. Caseworkers interact with algorithms in their daily lives where they must collect information about families and feed it to algorithms to make critical decisions. We show how the interplay between systemic mechanics and algorithmic decision-making can adversely impact the fairness of the decision-making process itself. We show how functionality issues in algorithmic systems can lead to process-oriented harms where they adversely affect the nature of professional practice, and administration at the agency, and lead to inconsistent and unreliable decisions at the street level. In addition, caseworkers are compelled to undertake additional labor in the form of repair work to restore disrupted administrative processes and decision-making, all while facing organizational pressures and time and resource constraints. Finally, we share the case study of a simple algorithmic tool that centers caseworkers' decision-making within a trauma-informed framework and leads to better outcomes, however, required a significant amount of investments on the agency's part in creating the ecosystem for its proper use.

CCS Concepts: • **Human-centered computing** → **Human-computer interaction (HCI)**; *Empirical studies in HCI*; • **Applied computing** → Computing in government.

Additional Key Words and Phrases: algorithmic harms, AI functionality, human-AI decision-making, repair work, child welfare system

ACM Reference Format:

Devansh Saxena and Shion Guha. 2018. Algorithmic Harms in Child Welfare: Uncertainties in Practice, Organization, and Street-level Decision-Making. In . ACM, New York, NY, USA, 30 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Algorithmic decision-making systems are increasingly being deployed in government agencies such as child welfare [31, 100], criminal justice [116, 117], public education [81], unemployment services [27] and homeless services [70] in an effort to improve public services delivery and target resources toward cases that most need them. Algorithms are being deployed across the public sector where they are informing decisions regarding which child maltreatment cases to investigate [66, 101], which families should be offered preventive services [106], which neighborhoods require more policing [17, 55], and who are offered public housing and unemployment benefits [43, 70], among others. However, audits of these systems show that they are achieving worse outcomes [29, 97], embedding human biases in decision-making [107], and exacerbating racial and socio-economic biases [43, 68, 94]. Algorithmic systems are causing real harm to vulnerable people who are witnessing cuts to their unemployment and healthcare benefits [36, 43], loss of public housing [125], as well as experiencing unfair investigations from the child welfare system [59, 114].

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

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

Child welfare (CW) agencies in the United States are increasingly facing severely limited resources, burdensome workloads, and high staff turnover and have turned towards algorithmic systems as they promise to allocate resources more efficiently and fairly [43, 75, 104] and offer them a means to "do more with less" [67, 94]. These agencies also rely on federal and philanthropic funding to support their operations with more grant funding opportunities being made available to agencies that invest in innovative evidence-based programs and systems reform and integration [113]. Recently, more federal funding has been made available to CW agencies that implement the new Comprehensive Child Welfare Information System (CCWIS) data system that further improves data collection mechanisms, integrates data from previously siloed departments, and supports bi-directional data exchange between courts, education systems, and Medicaid [56]. In addition, child welfare agencies have also been the center of public and media scrutiny because of harm caused to children who are removed from the care of their parents [24] but they also receive severe criticism and media attention for child abuse tragedies where the system failed to remove and protect a child [46, 54]. This has further mounted the pressure on CW agencies to employ algorithmic systems that can assist with high-stakes decisions (e.g., assessing the risk of future child maltreatment events) even though evidence points to such systems disproportionately targeting low-income and minority communities [29, 59].

Critics of AI systems across academia (e.g., [92, 111]) and the press (e.g., [36, 59, 64, 125]) have accumulated significant evidence regarding outcome-oriented harms caused by haphazardly deployed AI systems in the public sector, i.e., disparate harms that are caused to affected communities. Consequently, scholars have begun to synthesize a taxonomy of sociotechnical harms [111] as well as drawn attention to the fundamental misconceptions about the expected performance of AI systems [92]. Raji and Kumar et al. [92] refer to this as the *fallacy of AI functionality* and developed a taxonomy of AI functionality issues [92]. We build upon this body of work and further draw attention towards *process-oriented harms*, i.e. - harms caused to the fairness of the decision-making process itself as a result of algorithmic decision-making [51, 52]. In this study, we deepen this discussion about algorithmic harms by examining the impact of algorithmic tools on professional practices, administration at the agency, and street-level decision-making in child welfare. Digital technologies initially altered the workings of public agencies from 'street-level' to 'screen-level' bureaucracies [16]; a transformation referred to as digital era governance [20, 123]. However, despite promises of significant improvements in efficiencies and cost-effectiveness, several digital tools have fused onto existing street-level discretionary practices without altering public management at a deeper organizational behavior level [80, 123]. A similar exploration needs to take place to assess the impact of algorithmic decision-making on the nature of practice (i.e., street-level human discretionary work), administration (i.e., bureaucratic processes) as well as assess how street-level decision-making is changing and whether algorithms are living up to the promises of cost-effective, consistent, and fair decision-making. In this study, we make the following contributions –

- We show how functionality issues in algorithmic systems can lead to process-oriented harms where they adversely impact the nature of professional practice, and administration at the agency, and lead to inconsistent and unreliable street-level decision-making.
- We show how caseworkers are forced to assume the added labor resulting from these algorithmic harms and must conduct *repair work* [63] to address the disruption caused to administrative processes and street-level decision-making. This repair work is conducted within the bounds of organizational pressures, time constraints, and limited resources.
- We highlight the case study of a simple algorithmic tool that centers caseworkers' decision-making processes within a trauma-informed framework and leads to better outcomes for families, however, required a significant amount of investments on the agency's part in creating the ecosystem for its proper use.

In the following sections, we first discuss related scholarly work on public sector sociotechnical systems, the current landscape of algorithmic decision-making in the U.S. child welfare system, and the theoretical framework within which we ground our qualitative analysis.

2 BACKGROUND

In this section, we first explain our conceptualization of *algorithms* or *AI* since there is no agreed-upon definition of AI within the AI research community. Next, we discuss related scholarly work regarding digital technologies and algorithmic decision-making in the public sector.

2.1 What is an "Algorithm"?

Different definitions of algorithms have been adopted across different disciplines due to the intersection between statistical modeling and machine learning as well as ongoing innovations in advanced neural networks [58, 108, 120]. From a social perspective, we define algorithms through the lens of *street-level algorithms* [6] - computational predictive tools used to make on-the-ground decisions about human lives and welfare. That is, algorithms that directly affect families utilizing public services and not second or third-removed algorithms that might be internally used by government analytics teams. Technically, we define an algorithm as a computational system that takes in input data, computationally processes data, and produces an output. This output can be a predicted outcome of interest or change in output over time when some statistical property of variables change as is the case with change point detection algorithms [7].

2.2 Algorithmic Decision-Making in the Public Sector

Digital technologies have allowed public services to collect a vast amount of data about citizens in the course of their daily operations in managing and delivering public services [43, 44, 84, 124]. This has changed governance practices in two distinct ways. First, these systems have improved data sharing practices between different sectors of the government and as Veale and Brass [123] note, - these systems have allowed for the digital re-integration of siloed services, data sharing practices aimed at creating a 'one-stop-shop' and 'end-to-end' service delivery with minimal repeated information gathering. This comprehensive cross-sector data has allowed policymakers, technologists, and academics to narrow their focus on improving high-stakes decision-making by developing data-driven practices that purportedly provide consistent, transparent, objective, and defensible decisions to citizens [19, 43, 102]. Algorithmic decision-making in the public sector has generally been adopted in the form of risk assessment algorithms with their primary purpose being the preemptive recognition and mitigation of 'risk'; a core organizing concept of this shift in governance and of neoliberal economics [10]. That is, improving productivity, accountability, and efficiency by proactively identifying clients in the riskiest circumstances and targeting services towards them using algorithms [43, 94]. This further shift in *digital era governance* that embeds neoliberal principles within public administration has been called *New Public Analytics* [129] and focuses more on risk management based on individual client characteristics while driving attention away from structural and societal problems [43, 67, 94].

The breadth of research at the intersection of public technology and governance has been wide-ranging, including studies that examine how emergent technologies are shaping collaborative work [15, 83], designing technologies that empower affected communities [18, 37, 103, 115], studying issues concerning civic engagement [38, 49], understanding the social and ethical implications of datafied public services [61, 93, 94], as well as uncovering how gig work platforms exploit historical inequities which lead to poor working conditions for data workers and limit their agency. [85, 91]. Researchers have also unpacked the forms, limits, and complexities of participatory design within the public sector where

newer technologies are being designed for the governance of smart cities [57, 110, 128]. Several of these technologies are being developed through public-private partnerships [38, 75] where the expectation is that these entities are able to transform data into knowledge and inform decisions that are centered in the efficient allocation of resources [61, 86, 93]. As Holten Møller et al. [61] note, "*here, data becomes the promise of future bureaucratic efficiencies*".

Specific to affairs of public administration, researchers have studied how the shifting decision-making latitude has impacted the work of street-level bureaucrats.¹ Studies have found that value conflicts arise when the logics embedded within the government's digital platforms do not align with street-level bureaucrats' discretion when they tried enacting the same shared values in practice [37, 61, 66, 126]. Research within public administration, science and technology studies, and human-computer interaction have recently drawn attention to how the digitization of public services is leading to distinct changes in street-level bureaucrats' discretion [21], power asymmetries between public officials and citizens [106, 107, 114], need for re-skilling of public officials [8, 61, 66], and actor transparency in government decision-making [48]. Lindgren et al. [73] go a step further and argue that "public officials can no longer be understood as merely human". They call for a reinterpretation of citizens' trust in their government in regard to the legitimacy and accountability of e-governments. Busch et al. [22] explore public service workers' digitized discretionary practices as they balanced conflicting demands of market-oriented goals and norms of professional practice and found that workers responded positively to digitization when it supported professional aspects of their work. In a similar vein of work, Giest et al. [48] highlight that a disconnect between bureaucratic processes and digital tools magnified in street-level decision-making where workers' discretionary power was obfuscated and led to more complications and time-consumption in accomplishing daily tasks. In the context of smart cities, Meijer [82] further argues that building technologies for governance is centered in political, strategic, and value-laden choices made between three key actors: state, market, and civil society. This requires a re-conceptualization of sociotechnical structures that result from their interactions. Public administration scholars have also foreboded a "digital sclerosis" characterized by the stiffening of governmental processes and lowering of innovation feedback from workers [9]. They predict "decreased bargaining and discretionary power of governmental workers" as one of the early warning signs of this phenomenon. These changes in professional work practices through the adoption of digital tools have similar, yet more serious implications for the adoption of algorithmic tools which further shift discretion away from public service workers; instances of which are already being captured by recent studies in the public sector [66, 105].

As a result, researchers have started investigating the intersection of human discretionary work conducted at the street-level and algorithmic decision-making in public services [6, 8, 29, 61, 66, 101]. Alkhatib and Bernstein introduced the theory of *street-level algorithms* to distinctly highlight the gaps in algorithmic decision-making that human discretion needed to address [6]. Unlike street-level bureaucrats who used discretion to reflexively make decisions about novel cases, street-level algorithms produced illogical decisions that offered no recourse and could only be addressed by 'learning' through new data in the future. Pääkkönen et al. expand upon this theory to highlight that algorithm design needed to identify and cultivate important sources of uncertainty because it was at these locations that human discretion was most needed [88]. Additionally, recent work [8, 66, 101] has highlighted the collaborative nature of caseworkers' decision-making processes and the impact of bureaucratic structures that algorithm design needed to account for. In sum, researchers have reached a general consensus that any algorithmic interventions in the public sector needed to understand the complexities of human discretion carried out at the street level when implementing day-to-day bureaucratic processes and legislative policies. Recently, Saxena et al. [101] synthesized prior work conducted on

¹A street-level bureaucrat is a professional service worker (e.g., social worker, police officer, teacher) who operates in the frontline of public service provision. They interact closely with clients and make decisions about them based on how they interpret policies relating to the situations at hand [74].

algorithmic governance systems into a framework for algorithmic decision-making for the public sector which accounts for the complex interdependencies between human discretion, algorithmic decision-making, and bureaucratic processes. Next, we offer a brief overview of this framework and explicate its utility in unpacking the impact of algorithmic decision-making on the nature of professional practice, administration at the agency, and street-level decision-making.

2.3 Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS) Framework

Saxena et al. [101] developed a theoretical framework for algorithmic decision-making in the public sector that accounts for the complex inter-dependencies between *human discretion*, *bureaucratic processes*, and *algorithmic decision-making*. Below, we provide a brief summary and explain the three elements of the framework and their utility for our analysis -

Human Discretion refers to the decision-making process model that practitioners in the public sector engage in where they use their professional expertise, make value judgments, and engage in heuristic decision-making based on the available information. This is especially important in the public sector because government officials must make decisions within the bounds of policies and organizational and resource constraints. Moreover, they must use discretion in resolving missing or conflicting sources of information.

Bureaucratic Processes refer to critical governance characteristics that include the systemic constraints within which all decisions must be made, day-to-day protocols, and the legislation that the organization is legally mandated to follow. Bureaucratic processes directly impact practitioners' training and nature of practice (i.e., human discretion) at the agency as well as how well an algorithm is integrated into the day-to-day workflows and decision-making processes.

Algorithmic Decision-Making adopts *street-level algorithms* [6] as a theoretical lens to identify algorithmic systems that are used to make on-the-ground decisions about clients and welfare in the public sector. It draws attention to the relevant data needed for decision-making, the degree of uncertainty associated with predicted decisions, as well as the decision-making latitude (i.e., predictive or prescriptive) that is allocated to algorithmic systems. Authors refer to *algorithmic decision-making* as the most flexible element of the framework that researchers can directly affect by designing systems that balance the other two elements (i.e., human discretion and bureaucratic processes). In addition, the authors further highlight the need to understand human discretionary work because algorithm design needs to make space for (and be preceded by) human discretion to fill in the gaps in data as well as make sense of organizational and legislative protocols.

In sum, the ADMAPS framework showcases how practical trade-offs must be made to manage the cross-dependencies at both the macro- and micro-levels of the algorithmic model to offer autonomy to practitioners and improve human discretionary work. The framework also draws attention to the high degree of uncertainty inherent in the administrative data which consequently means unreliable predictions. Therefore, the goal of algorithms in the public sector must be re-evaluated to support the decision-making processes of stakeholders instead of providing predicted outcomes. In their study, Saxena et al. [101] focused on the micro-interactions between the dimensions of these three elements to understand why each algorithm failed (or succeeded) to offer utility to child welfare staff and their impact on human discretionary work. However, in this study, instead of focusing on singular algorithmic tools and their impact on human discretion, we draw attention to the broader decision-making ecosystem and critically investigate the macro-interactions between these three elements to assess the impact of algorithmic systems on the nature of practice (i.e., interaction between human discretion and algorithmic decision-making), the organization (i.e., interaction between bureaucratic processes and algorithmic decision-making), as well as three-way interactions between the three elements to understand how the nature of street-level decision-making is changing in child welfare.

2.4 Current Landscape OF Algorithms Used in the U.S. Child Welfare System

Child welfare (CW) agencies in the United States have increasingly adopted algorithms for high-stakes decisions as they promise to improve decision-making, lower costs, and provide better outcomes to citizens [43, 96]. A nationwide survey on predictive analytics in child welfare conducted by the American Civil Liberties Union (ACLU) in 2021 revealed that 26 states have considered employing predictive analytics in child welfare [99]. Of these 26 states, 11 are currently using them [99]. Several states in the United States continue to experiment with predictive analytics, however, audits of these systems reveal that they are achieving worse outcomes for families and exacerbating racial biases [45, 60, 87, 119]. In the past, Los Angeles County and the state of Illinois have shut down their predictive analytics programs for these reasons [45, 119] with Oregon recently joining their ranks in June 2022 [87]. A recent study conducted by Cheng and Stapleton et al. [29] on the Allegheny Family Screening Tool (AFST) found that AFST-predicted decisions were racially biased and workers reduced these biases by overriding erroneous decisions. AFST algorithm was designed to mitigate call screeners' biases and subjective decisions and augment decision-making by making it more objective through data. Ironically, AFST has introduced more complexities in decision-making and the call screeners are the ones mitigating algorithmic biases. A comprehensive literature review of algorithms in CWS revealed other sources of biases embedded in the predictors, outcomes, and computational methods being used to develop these systems [102]. Kawakami et al. [66] found that call screeners were offered minimal information about the workings of AFST, considered it to be unreliable due to unexpected model behaviors, and even engaged in a collaborative "game" to learn more about the tool. Gerchick et al. [47] unpack how design decisions made by AFST's developers effectively act as policy choices that sit outside the purview of democratic oversight. They show how families are labeled "risky by association" because designers aggregated risk scores at the household level and how systemic discrimination is amplified by the inclusion of static variables that capture families' involvement with the criminal legal system or public benefits programs. Similarly, Saxena et al. [105] draw attention to the multiplicity and temporality of risk factors that are experienced by families in the system, how the system itself poses significant risk where families experience over-surveillance and are subjected to several arbitrary decisions, and how these street-level structural biases become embedded in algorithmic systems such as the Allegheny Family Screening Tool.

Federal initiatives such as improved data infrastructures for CWS [56] have paved the way for tech startups to develop and pitch algorithmic systems to human services agencies across different states [1–4]. However, there is a need to critically examine the current points of failures in algorithm design as well as understand how workers engage with algorithms as they make critical high-stakes decisions about families. Critical to the conversation about predictive analytics or predictive risk models (PRMs) is also the underlying principle of "risk" and how its understanding has shifted in response to the restructuring of public services to be economically efficient, productive, and accountable [10, 23, 94]. Traditionally, child welfare services have focused on risks and protective factors within families to be able to provide them with individualized care. However, with a shift towards an economic understanding of risk and the introduction of PRMs, risk has now become a function of client characteristics as existing in prior cases (and not individual family circumstances) and their impact on a predicted outcome (i.e., risk of maltreatment) which Gerrick et al. [47] refer to families being labeled "risky by association". That is, the risk is estimated based on historical administrative data and is being used to identify the "deserving poor" who pose the most risk to governmental apparatus [43]. Redden et al. [94] refer to this as the embedded logic of actuarialism that also obfuscates and drives attention away from social and structural issues that bring poor and vulnerable communities under the attention of public services such as child welfare, housing authority, and public assistance [67].

3 METHODS

We partnered with a child welfare agency that serves a metropolitan area in the midwestern United States. This private, non-profit agency is contracted by the state's Department of Children and Families (DCF) to provide child welfare and family services and must comply with all DCF standards, including the use of mandated decision-making algorithms. DCF's Initial Assessment (IA) workers investigate allegations of child maltreatment, and if abuse/neglect is substantiated, the child(ren) may be removed from the care of their parents or an in-home safety plan might be developed. At this point, the case is referred to the agency to provide services. These services are negotiated between the parents' attorney, the district attorney's office, and the judge after caseworkers have conducted initial structured assessments and provided their recommendations to the court. A case manager is assigned to each case and is supported by a multi-disciplinary child welfare team that brings in domain expertise from social work, family psychology, medicine, and law. The agency is mandated to use some algorithms at different stages of a child welfare case per DCF standards but has also developed an algorithm in-house to improve its decision-making processes.

3.1 Study Overview

We conducted a two-year extensive ethnography to understand how caseworkers interacted with algorithmic systems, and their perspectives on these systems, as well as unpack how decisions were made at the intersection of child welfare practice, regulations, and algorithmic decision-making. This study is conducted within the same broader child welfare system that was also the ethnographic site for Saxena et al. [101]. However in this study, instead of focusing on singular algorithmic tools and assessing their impact on human discretion, we draw attention to the broader decision-making ecosystem and assess the impact of algorithms on the nature of practice, administration at the organization, and the changing nature of street-level decision-making. Before conducting observations or recruiting participants for interviews, we obtained Institutional Review Board (IRB) approval at our research institution to conduct our study. Before the interviews, we emailed the participants an IRB-approved consent form and obtained their verbal consent to participate in the study before beginning the interviews. The first author observed 80 child welfare team meetings and then conducted 20 semi-structured interviews with key stakeholders at the agency to better contextualize interactions that occurred at these meetings and gather caseworkers' perspectives. Below, we first highlight our scoping criteria for algorithms followed by a description of observations and interviews, and the qualitative data analysis process.

3.2 45-Day Staff Meetings, Permanency Consultations, and Concurrent Planning

45-day staff meetings occur within the first 45 days of a case coming into the care of the agency and are attended by child welfare staff involved at the front end of case management and planning. These meetings facilitate information sharing such that consensus decisions can be made regarding child well-being. Each meeting is scheduled for 90 minutes and the first author observed 25 such meetings. These meetings are typically attended by child welfare staff that work in case management, permanency planning, family preservation, and licensing. The goal of these meetings is to develop a deeper understanding of a family's circumstances from a trauma-informed perspective, identify parents' support system that can help with child care, identify systemic barriers that may inhibit reunification, and develop action steps for child welfare staff. These meetings facilitate collaborative decision-making and ensure congruence in case planning.

Permanency consultation meetings are specialized meetings designed to expedite permanency for children placed in out-of-home care by employing collaborative practices as well as actively addressing any systemic or policy-related barriers. These meetings are facilitated by permanency consultants and are staffed with many of the child welfare team

members that attend 45-day planning meetings. These meetings regularly occur at the 5, 10, and 15+ month marks for every case until the case is closed. These ongoing meetings tended to be more informative than the 45-day planning meetings because not enough information is available at the onset of a case. Moreover, permanency consultations involved cases that had been with the agency for several months (if not years) and revealed the complicated interactions between policies, systemic barriers (legal, resource, administrative), social work practice, and algorithmic systems being used at the agency. The first author attended 55 permanency consultations each of which was an hour long.

Critical to a discussion about collaborative meetings and the decision-making ecosystem is the *Adoption and Safe Families Act (1997)* which introduced some of the most sweeping changes to the child welfare system and shifted the focus primarily towards child safety concerns and away from the policy of reuniting children with parents regardless of prior neglect/abuse. The Act introduced federal funding to assist states with foster care, adoption, and guardianship assistance and expanded family preservation services. In addition, it also introduced a **15-month timeline** where the State must proceed with the termination of parental rights if the child has been in foster care for 15 out of the last 22 months [50]. This speedy termination of parental rights has received widespread criticism but still establishes the constraints within which caseworkers must conduct their work [53, 90, 118]. To ensure expedited permanency² for foster children, this child welfare agency employs concurrent planning such that two simultaneous plans begin when a child enters foster care: a plan for reunification with the birth parents and a plan for guardianship and/or adoption if reunification is not possible.

3.3 Semi-Structured Interviews

Next, we used the knowledge gathered from these observations to develop our interview protocol and recruit participants who consistently attended these meetings. We conducted these interviews to delve deeper into caseworkers' understanding of these algorithms as well as the benefits and challenges as perceived by them. We asked participants a series of questions about the nature of child welfare work, how algorithms in use impacted their practice, and the organizational support made available to them to mitigate conflicts or challenges associated with these systems. We also asked them to expand upon any interactions we had observed during the meetings such as the participant's dislike or appreciation for a certain algorithm or feature or their frustration with the misuse of algorithmic tools. During observations, we also noted any benefits or frustrations that the participant (or their team) experienced regarding algorithmic decision tools and brought them up during the interview to further expand upon and better understand their outlook. We conducted 20 interviews with participants that included program directors, child welfare supervisors, permanency consultants, caseworkers, data specialists, and clinical therapists.

3.4 Study Participants

We interviewed child welfare workers who attended the collaborative meetings mentioned above and had more experience working with algorithmic tools and assessments in their day-to-day work. The case manager position experiences high turnover with most case managers quitting within the first two years. Therefore, we focused our attention on the more experienced members of the child welfare staff who could provide deeper insights into the workings of the system as well as systemic barriers that impede their work. Below, we provide job descriptions of various positions and participant information in Table 1.

²Permanency is defined as reunification with birth parents, adoption or legal guardianship.

Participant	Sex	Job Title	Experience (years)
P1	F	Permanency Consultation Supervisor	22
P2	F	Permanency Consultant	20
P3	F	Permanency Consultant	9
P4	F	Permanency Consultant	8
P5	F	Permanency Consultant	12
P6	F	Permanency Consultant	3
P7	F	Child Welfare Program Director	20
P8	F	Child Welfare Supervisor	19
P9	F	Child Welfare Supervisor	13
P10	F	Child Welfare Supervisor	9
P11	F	Child Welfare Supervisor	9
P12	F	Child Welfare Supervisor	7
P13	M	Child Welfare Supervisor	12
P14	M	Child Welfare Supervisor	30
P15	F	Child Welfare Supervisor	9
P16	F	Clinical Therapist	5
P17	F	Child Welfare Case Manager	8
P18	M	Child Welfare Case Manager	2
P19	F	Data Specialist (Program Director)	17
P20	M	Data Specialist	17

Table 1. Interview Study Participants

Program Directors (n=1): Agency leadership responsible for professional development programs, trainings, research and policy initiatives, and grant writing. They supervise child welfare supervisors and the agency’s trauma-responsive service model.

child welfare Supervisors (n=8): Supervisors manage a case management team comprising 6-8 case managers and oversee about 140 cases. They are responsible for the professional development of case managers and provide additional support when interacting with families and legal parties. They also facilitate 45-day staffing meetings.

child welfare Case Managers (n=2): Frontline workers that directly interact with families and act as mediators between parents, foster parents, relatives, legal parties, and child welfare staff. They conduct home visits, safety assessments, and psychometric assessments, and transport foster children to supervised visits and medical appointments. On average, they manage about 20 cases.

Permanency Consultation Supervisor (n=1): The supervisor in charge of the permanency consultation program that is designed to address systemic barriers that pose obstacles in the way of achieving permanency for foster children. They manage and supervise permanency consultants.

Permanency Consultants (n=5): They are responsible for managing the legal process underscoring permanency where they prepare documentation for court, focus on recruitment and licensing of foster homes, and manage post-guardianship and post-adoption services, among other tasks that help expedite permanency for foster children. They provide consultations on about 150 cases each.

Data Specialists (n=2): Responsible for tracking federally-mandated performance benchmarks for the agency as well as case-level data. They also analyze data and manage algorithmic systems implemented at the agency and present results to agency leadership.

Clinical Therapist (n=1): A licensed clinical social worker who conducts mental health assessments and has a deeper understanding of psychometric assessments used in child welfare.

3.5 Qualitative Data Analysis

The first author took detailed observational notes during each team meeting and compiled a debriefing document with their initial insights. They also noted questions that needed further clarification and conferred with team leaders (i.e., supervisors or permanency consultants) after the meetings and noted their responses. All authors read through the observational notes and collectively discussed meeting notes that uncovered pertinent aspects of the decision-making ecosystem in regard to the nature of practice, policies and regulations, and/or the use of algorithmic tools. Interviews were audio-recorded and transcribed verbatim for analysis. After carefully reading through the transcripts, we conducted several rounds of iterative coding to identify patterns and converge on appropriate themes associated with the three elements of the ADMAPS framework (i.e., human discretion, bureaucratic processes, and algorithmic decision-making). We performed thematic analysis [33] to create these initial codes, form a consensus around the codes, as well as resolve any ambiguous codes. In our results, we also use our observational notes to augment the insights we gained from the interviews and note potential discrepancies and nuances from the holistic insights gained from our site observations. These initial codes were grouped into three high-level themes of *impact on practice*, *impact on organization*, and *impact on street-level decision-making*. Codes that highlighted pertinent issues regarding how human discretionary work has changed under algorithmic systems were grouped under *impact on practice*. For instance, 80% of the participants shared that the CANS algorithm had negatively impacted social work practice due to its lack of understanding of trauma. Similarly, codes that highlighted issues such as an algorithmic system's inability to account for organizational constraints or legislation were grouped under *impact on organization*. For instance, 75% of the participants shared that all decision-making processes were first situated in the 7ei algorithm which is central to the agency's trauma-responsive service model. Finally, codes that highlighted how street-level decision-making is changing on the ground where bureaucrats are mandated to use algorithmic decisions (but do so while employing their professional expertise and operating under organizational and legislative constraints) were grouped under *impact on street-level decision-making*. For instance, 80% of participants shared that they were mandated to use the foster care placement as recommended by CANS, however, foster children were more likely to achieve stability and well-being when placed with relatives rather than a more restrictive foster care setting as recommended by CANS. In addition, CANS did not account for the limited number of foster homes available in the system; an organizational constraint that rendered the algorithmic recommendation nonsensical.

4 RESULTS

In the following subsections, we first share some overarching findings regarding the decision-making ecosystem in child welfare that sits at the intersection of policies, social work practice, and algorithmic systems. Next, we discuss findings from two algorithms that are in use at the child welfare agency and how their adoption has impacted the nature of practice, administration at the agency, and street-level decision-making.

4.1 Decision-making Ecosystem in Wisconsin's Child Welfare System

In order to understand the role of algorithms in decision-making, it is necessary to map out the complexities within the broader decision-making ecosystem to be able to assess the utility and scope of algorithmic tools. In addition, it is imperative to understand the systemic constraints within which such systems must operate. This helps us to better contextualize how caseworkers interact with algorithms, their perspectives on these systems, as well as how decisions are made at the intersection of policies, social work practice, and algorithms, i.e. - how algorithmic decisions

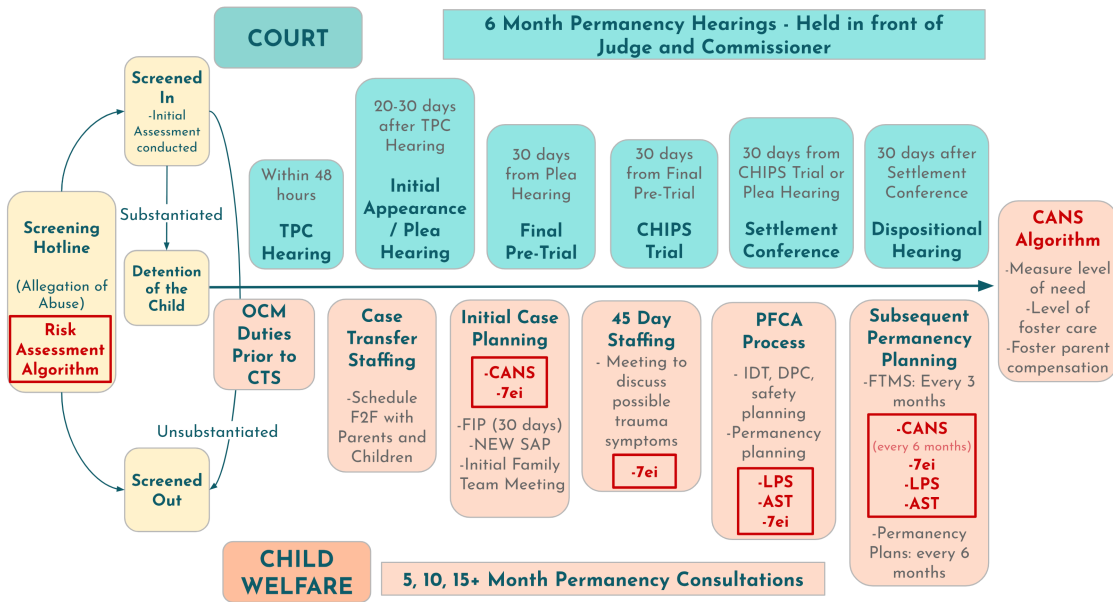


Fig. 1. Life of a child welfare Case in Wisconsin

IDT: Interdisciplinary team, PFCA: Private Foster Care Agencies, OCM: Ongoing Case Manager, CTS: Caretaker Supplement, TPC: Temporary Physical Custody, F2F: Face to Face, FIP: Family Investment Program, SAP: Substance Abuse Prevention, FTM: Family Team Meetings, DPC: Direct Primary Care

were used by child welfare staff working under legislative and organizational pressures. Figure 1 describes “Life of a Case” in child welfare and highlights all the critical decision-making steps and algorithms (red boxes) that are embedded throughout the child welfare process. We co-designed this diagram with caseworkers and it is generally used at the agency to depict the customary legislative process, however, this diagram can significantly vary based on case circumstances. The yellow section of the diagram represents the initial allegation of maltreatment and investigation that is conducted by Initial Assessment (IA) caseworkers at the Department of Children and Families (DCF). If maltreatment is substantiated, the case is officially opened and referred to the child welfare agency to provide ongoing services to the family. This is represented by the orange section of the figure. Finally, the blue section of the figure represents the court system. Contrary to popular belief, critical decision-making power in regard to permanency (i.e., reunification, adoption, or transfer of guardianship) for foster children sits with the court system and child welfare staff only makes recommendations to the district attorney’s office. Tensions between child welfare staff and the court system are well captured in social work literature [26, 39, 42]. This ethnography is conducted within the orange section of the figure (i.e., the child welfare agency) where algorithmic tools are being used by caseworkers to make several day-to-day street-level decisions (as opposed to critical permanency decisions made in courts) about children and families. Underscoring the use of algorithmic decision-making in child welfare are two dominant concerns that frequently arose in this ethnography -

- Caseworkers are not trained in “thinking statistically” about data, algorithms, and uncertainties but are legally mandated to input data, interact with algorithms, and make critical decisions.
- All algorithmic decisions in the public sector must be made within the bounds of policies, current practice, and organizational constraints.

As depicted in Figure 1, when an allegation of abuse is made at the hotline, a risk assessment algorithm helps call screeners to decide if the call should be screened for an investigation. We did not have access to this decision step since it is conducted at DCF before the case is referred to the agency where the ethnography was conducted. If the call is screened in and the investigation substantiates abuse and/or neglect, the child(ren) is removed from the care of their parent(s), or an in-home safety plan is put in place. The case is then referred to this child welfare agency. Throughout the life of the case, the CANS algorithm is used to assess the mental health needs of a child, the level of foster care (level 1 to 5) the child should be placed in, as well as calculate the compensation that foster parents should be paid by the state. 7ei (Seven essential ingredients) algorithm is used to develop trauma-responsive services. LPS (Legal Permanency Status) algorithm is used to track federal benchmarks such as placement stability, as well as policy and systemic barriers. And finally, AST (Anti Sex-Trafficking Response) algorithm is used to assess the risk of sex trafficking for a foster child. These algorithms are being used to make everyday decisions about foster kids but there are also other algorithms that are being used internally by the Department of Children and Families (DCF). For instance, an algorithm is used to assess the risk of re-entry into the system for every case. In this study, we only focus on the CANS and 7ei algorithms.

4.2 Caseworkers' Perspectives: Assessments or Algorithms?

Through the observations of child welfare meetings, we learned about several different sources of data that are collected by the agency as a means to provide consistent information about all cases. The agency uses several psychometric assessments for this purpose (see for e.g., [78, 79, 95]) as they provide a structured framework to conduct risk assessments. Child welfare initially adopted psychometric risk assessment instruments (RAI) as a means to standardize the process of assessing children's and parents' risks and needs and allow for a more consistent decision-making process. Consequently, RAIs have facilitated the collection of data about children and families for the past three decades. Over the last decade, psychometric and administrative data from prior substantiated cases of maltreatment is now being used to train algorithms to make predictions about current cases [102]. The algorithms depicted in Figure 1 have all been developed using their RAI counterparts. RAIs have been widely adopted in social work practice, and as a result, all participants recognized these assessments but did not actively recognize their algorithmic components (i.e., automated aspects) that routinely frustrated them. However, once we focused more on these frustrations and expectation violations, we uncovered several pertinent issues at the intersection of human discretion, bureaucracy, and algorithmic decision-making. In the following sections, we discuss the CANS and the 7ei algorithm and their impact on the nature of social work practice, administration at the agency, and street-level decision-making.

4.3 CANS Algorithm

CANS (Child and Adolescent Needs and Strengths) algorithm is constructed using the CANS communimetric assessment that consists of 104 psychometric items organized across eight domains that address child needs and strengths (see [30] for more details). It was designed to assess the *level of need* of a foster child and utilize this assessment to develop an individualized care plan. CANS offers the child welfare team a structured way to assess a case as well as share this information with other parties. With its primary purpose being communication, CANS is designed based on communication theory rather than psychometric theories centered in measurement development [76]. That is, the tool was not designed to explicitly measure any variables or predict outcomes based on these variables. CANS was designed to be the *expected outcome* of the decision-making process such that it facilitates the linkage between the assessment process and the design of individualized service plans. That is, CANS was designed to support care planning, facilitate quality improvement initiatives, and monitor the outcomes of new evidence-based practices.

CANS algorithm, on the other hand, uses the risks and needs data about children collected by the assessment and has been re-purposed to explicitly measure three outcomes - **1)** mental health services to be offered to foster children, **2)** level of foster care they should be placed in, and **3)** generate subsidized guardianship rate offered to foster parents. Based on a child's risks and needs, the algorithm recommends mental health services that must be offered. Next, based on this level of need, it predicts the level of foster care the child should be placed in. Foster homes in this state range from Level 1 to Level 5. Higher-level foster parents are more trained and certified to take care of higher-needs children. Therefore, the higher the level of need, the higher the level of foster care that the child should be placed in. Finally, based on the level of need and foster home setting, CANS generates the subsidized guardianship rate that foster parents are paid by the state for the costs associated with having a child placed in their care. CANS is conducted within the first 30 days of a child entering the child welfare system or moving to a new placement. Subsequently, it is then conducted every six months. CANS algorithmic assessment is completed by caseworkers with information (about child behaviors and needs) provided by foster parents.

This re-appropriation of CANS (and therefore, CANS data) to predict these outcomes has led to several unintended consequences that frustrate caseworkers and impedes theory-driven practice centered in trauma-informed care. All caseworkers at the agency are trained and certified in conducting the CANS assessment, however, they are not trained in managing conflicts that arise due to the re-appropriation of CANS. This further obfuscates the boundary between the role of the **assessment** versus the **algorithm**. Below, we discuss some unintended consequences on the nature of practice, administration at the agency, and street-level decision-making.

4.3.1 CANS Diminishes the Nature of Social Work Practice.

In this section, we show how the CANS algorithm adversely affects the nature of social work practice because it does not account for the principles of trauma-informed care, misses important details about the child's mental health, and offers no context regarding worsening behaviors.

CANS contradicts trauma-informed care. Caseworkers at the agency are trained in trauma-informed care (TIC) and apply principles and practices from TIC while employing the agency's trauma-responsive service model to cases. However, the CANS algorithm only offers a live snapshot of the child's mental health based on exhibited behaviors (over the past 30 days) and does not account for underlying trauma (or traumatic triggers) that can cause serious emotional dysregulation from time to time. Here, caseworkers co-opt CANS to account for trauma-informed care and anticipate the child's needs. One permanency consultant explained -

"We have a child that goes into manic depression every year around the holiday season. At this point, one of the foster parents has to quit their job and care for the child full time. We know that this traumatic trigger is coming up but there is no way to account for that in CANS. So we edit scores in anticipation for these upcoming needs" - P3, Permanency Consultant

CANS misses important context about the child's mental health. Algorithms attempt to generalize child and family characteristics and place them in certain categories to be able to make a determination. However, this inadvertently leads to a loss of information or context since all the information cannot be accounted for. CANS focuses on child behaviors, risks, and needs but it does not account for the child's interactions and relationships with other people. Caseworkers shared that CANS conducted the child's assessment in an isolated manner and did not account for the quality and impact of relationships in their lives which are often more important in determining their long-term well-being. For instance, a permanency consultant shared -

"A child's behaviors are often a result of what is happening in their environment. What or who is really triggering them and where is that coming from.. there is no way to put that in CANS" - P4, Permanency Consultant

CANS offers no context regarding worsening behaviors. A clinical therapist shared that it was more important for them to be able to understand the context around worsening behaviors. They shared that worsening behaviors were not always a bad thing since children's behaviors initially get worse (before getting better) when they started therapy. This is simply because children are finally starting to address the underlying trauma in their lives in therapy which might lead to worsened behaviors in the short-term but is necessary for their mental health and well-being in the long run. The clinical therapist also shared the following concern about the short-term focus on mental health -

"If CANS is used as a communication tool, that is, just as a structured way to talk about children and families and then to make decisions on it, then that's fine! But that's not how it's being used.. A lot of it is just entering [data] into a computer and seeing whether or not there's improvement from one point to another. And if there's no improvement then making a decision." - P16, Clinical Therapist

These concerns regarding the re-appropriation of CANS were inconsistently shared by child welfare workers. That is, not all caseworkers were equally aware of other concerns and only shared issues that affected their day-to-day practice. For instance, caseworkers were unaware of the clinical therapist's concerns regarding *worsening behaviors*. A supervisor shared that children can have severe underlying trauma that is not captured by CANS because of its focus on exhibited behaviors. At this point, we shared the clinical therapist's concerns with the supervisor and they further elaborated by saying that underlying trauma did not always manifest in terms of exhibit behaviors and that CANS may be measuring the wrong indicators altogether. This also raises questions regarding which explanations associated with the tool are deemed more useful - caseworkers were more invested in explanations of underlying trauma (based on their training in trauma-informed care), whereas the clinical therapist wanted better explanations about the context surrounding worsening behaviors. That is, different stakeholders, have different needs in regard to explainability and a "one-size-fits-all" approach with respect to algorithm design may not be feasible here.

Such interactions with the algorithm over time led to cumulative distrust in algorithmic decision-making. The implementation of CANS also further obfuscated the difference between the algorithm and the assessment. For instance, caseworkers generally score CANS on the paper-based assessment when speaking with foster parents. However, at the agency, they input this data into the CANS algorithm to predict outcomes. Re-appropriation of CANS is also a case in point regarding why it is imperative to understand the context within which (and how) data is collected about clients in the public sector.

In sum, CANS diminishes the nature of social work practice centered in trauma-informed care because it misses important context about children's mental health, does not account for the interactions in their ecosystem, and offers no context regarding worsening behaviors. Caseworkers' ongoing interactions with the algorithm where these conflicts continually arise have led to a cumulative distrust in the algorithm.

4.3.2 CANS has Introduced More Barriers to Administrative Processes at the Agency.

Below we discuss the impact of the CANS algorithm and its re-appropriation on the day-to-day bureaucratic processes carried out at the agency. CANS has added more barriers to day-to-day administration where caseworkers are forced to provide added labor to redress these barriers.

Subsidized guardianship rate has become the primary outcome of interest. Both caseworkers and foster parents are actively aware that the subsidized guardianship rate is directly tied to mental health needs. This leads to the gamification of the algorithm for three main reasons -

- Maintaining a stable placement - There is a lack of good foster homes in the system and caseworkers try to support foster parents by any means necessary to ensure that the placement is not disrupted and the child does not need to be placed elsewhere. Multiple placement moves for foster children are associated with poor well-being outcomes where they are unable to develop any meaningful relationships with foster parents or other caregivers [14]. Therefore, gaming the algorithm is the only way in which placements can be continually supported.
- Base rate is low - Caseworkers shared that the subsidized rate for most foster placements was too low and there was no financial incentive for foster parents to be doing this work. Here, gaming the algorithm is a convenient way to generate a higher compensation by exaggerating the mental health needs on the assessment.
- Maximizing financial incentive - Caseworkers also shared (see quote below) that even though most foster parents did their best to care for foster children, there were still other foster parents who accepted multiple placements in their home and ran foster care like a business. They routinely exaggerated child behaviors to receive higher compensations and renegotiated to what they considered to be their "standard" rate. One supervisor explained -

"Case managers are being pressured into scoring children higher. Foster parents will match [the rate] with the previous kids. They think that their previous kid got 1500 dollars, so now that is their standard rate, and will demand 1500 dollars for the next kid. They don't have an understanding of CANS or the child's strengths or needs. Money is the key part of these decisions." -P13, Child Welfare Supervisor

We refer to caseworkers manipulating the CANS data as *gamification* to draw attention to the fact that inexperienced caseworkers socially learn from seasoned caseworkers how to play with CANS data in order to generate higher compensations without raising any red flags. However, *repair work* [63] provides a better conceptual lens for capturing the shift in the nature of practice because caseworkers must provide the added labor to redress the disruption in decision-making caused by the CANS algorithm. Moreover, the intent behind conducting this repair work is to make the algorithm work for the clients and not just meet the allocation demands of policymakers.

4.3.3 CANS has Added Inconsistencies to Street-level Decision-Making.

In this section, we discuss CANS' impact on street-level decision-making. Instead of making decision-making more evidence-based, CANS has introduced arbitrariness to the process where children are being sent to unnecessary services and the data no longer reflects the true mental health needs of children. We discuss the impact of the CANS algorithm on street-level decisions that caseworkers make in their work relationships with foster parents.

CANS punishes good foster parents. Even if there were safeguards in place that prevented the gamification of CANS, the core incentive structure of the algorithm is problematic. Foster parents offer stability and support so that kids can develop good coping skills. They also help address mental health needs and help kids stabilize by taking them to therapy and all their activities. However, CANS scores are recalculated every six months and as the child supposedly exhibits less overt behaviors, their needs are lowered (per CANS assessment) and so is the compensation offered to foster parents. One supervisor explained (see quote below) that by lowering the rate, foster parents were being punished for being actively involved and caring for foster children. The CANS algorithm is set up to incentivize worsening behaviors.

"There are an enormous amount of foster parents that do an exceptional job. And so, it is hard, how do you reward them? Because in my opinion, it needs to be a rewarding system... like you're putting in the energy.. you're doing what you're supposed to do... There are foster parents getting their kid into all these activities and making sure this kid has a "normal social life", like normalcy, just part of parenting law. And they are getting half the rate compared to foster parents who are putting in no energy." - P9, Child Welfare Supervisor

Inconsistent placement decisions due to a lack of foster homes. All participants shared that the majority of placement decisions (i.e., where to house a foster child) come down to the availability of good foster homes. For instance, a child might have severe mental health or medical needs and the algorithm might recommend placing the child in a residential care center. However, most of these centers have limited openings, and the child welfare staff is forced to manipulate CANS scores and place the child in a group home or foster care setting that is ill-equipped to manage their needs. In addition, policy dictates placement decisions regarding relatives willing to accept guardianship for a child. For instance, a supervisor shared that relatives' homes must meet all legal and safety requirements and go through a licensing process to become foster parents (see quote below). However, if relatives fail to meet these requirements, CANS is used again to recommend a placement setting where the foster child can at least be temporarily placed until better placements become available.

Houses need to be legally compliant with all safety codes before a child can be placed there. These rules make sense on paper but in practice, we lose so many good placement options because relatives can't afford to move or fix everything in the house to be legally compliant" – P10, Child Welfare Supervisor

In sum, the re-appropriation of CANS has become a source of frustration for child welfare staff due to inconsistent street-level decisions where they must continually find ways to address policy and systemic barriers as well as ways to make the algorithm work for their clients.

4.4 7ei (Seven Essential Ingredients) Algorithm

This child welfare agency's core service delivery model is centered in trauma-informed care (TIC) where the staff assesses every case from a trauma-informed perspective and develops trauma-responsive interventions. However, program directors at the agency were concerned that decisions were still being made in an arbitrary manner where only some child welfare teams were employing TIC practices. To address these gaps in practice, agency leadership has made significant investments towards developing a comprehensive four-part evidence-based program. The program is based on core concepts of the Neurosequential Model of Therapeutics [89]. Recently, principles of trauma-informed care have also been adopted within computational research where Chen et al. [28] have proposed a framework for trauma-informed computing that translates these principles into practical design guidelines.

The 7ei algorithm (commonly referred to as the 7ei tool) is one of the four parts of this program. It is designed using TIC principles where the child welfare team discusses and scores variables across seven domains. We provide some high-level explanations of the seven domains in Table 1 but do not have the agency leadership's permission to share the complete tool. Instead of predicting a singular outcome of interest, 7ei is used to assess the trajectory of child welfare cases, i.e. - *change in 7ei domains over time*. Supervisors and program directors use 7ei to monitor trends in the seven domains to ensure progress is being made in cases using a seven-pronged TIC approach. For instance, if supervisors notice a downward trend for *Regulation* for a given case, it triggers consultations with a clinical supervisor and the development of consequent action steps that must be followed. This integrated approach allows the team to also focus on the child's and parent's ecosystem and their social support system such that parents have more caregiver support in the future. In addition, casenotes from the child welfare team are also collected and linked to the 7ei quantitative scores which provide program directors with more contextual information about the scoring of variables at any point in time. This has allowed the agency to collect pertinent data about their practices in TIC and informs future improvements towards their TIC service delivery model. Below, we discuss how 7ei has impacted the nature of social work practice, administration at the agency, and street-level decision-making.

7ei Domain	Explanation
Prevalence	Exposure to and difficulty adjusting to adverse life experiences, i.e. - what happened to the child regarding trauma exposure?
Impact	Trauma occurs when a person's ability to cope with an adverse event is overwhelmed and contributes to difficulties in functioning
Perspective Shift	Caregivers' understanding of child's trauma exposure and its relationship with child's behaviors and/or impairments
Regulation Relationships	Traumatic triggers that cause serious dysregulation and action steps to address it Current progress towards facilitating strong relationships in a child's life. Strong relationships help create resilience and mitigate effects of trauma
Reason to Be	Child's sense of identity and purpose and their sense to connectedness to their family, community, and culture
Caregiver Capacity	Caregivers' understanding of their importance to this process and whether they have additional social support

Table 2. 7ei Algorithm: Explanation of Seven Domains.

4.4.1 7ei Augments the Nature of Social Work Practice.

The first and second components of the four-part program involve extensive ongoing TIC trainings where the child welfare staff (e.g., caseworkers, supervisors, family preservation team, permanency consultants) is introduced to the complexity of trauma, frameworks for understanding the effects of trauma, and practices and principles of TIC. This is accompanied by specialized supervision and consultation with a clinical supervisor, caregiver support specialist, program administrator, and a national expert (Dr. Bruce Perry) who provide varying degrees of support to frontline caseworkers based on family circumstances. These two components ensure that the nature of practice is centered in the TIC model and allows for deeper engagement with 7ei. One supervisor explained -

"Every case is centered in the trauma-responsive model using 7ei. It ensures that caseworkers are always thinking through TIC" - P14, Child Welfare Supervisor

Even though most participants shared that they appreciated 7ei and the trauma-responsive model, some participants drew attention to deeper systemic issues in child welfare and the nature of social work practice that need redressing for tools like 7ei to offer more utility. One program director explained -

"During home visits, caseworkers need to ask the right questions and read situations to be able to derive meaningful information. But we are not hiring people who have been doing this for 5-10 years or are highly qualified. The only new hires available are newly graduated social workers." -P7, Program Director

In sum, 7ei was not designed to address such issues in child welfare, however, it is pertinent to note that any algorithmic tool can only offer limited utility when the workers' decision-making latitude (and social work practice) are inherently impacted by deeper systemic and structural issues. That is, the tool augments the quality of social work practice itself but it must still operate within the restrictive legislative framework of the child welfare system [105].

4.4.2 7ei Required the Development of New Administrative Processes at the Agency.

The third component introduces a staffing protocol in the form of specialized collaborative meetings (commonly referred to as *7ei meetings*) where experienced members of child welfare staff (e.g., family preservation, program director, permanency consultants) share their expertise and provide support to frontline staff members (i.e., caseworkers and supervisors). One supervisor explained -

"I get the most out of having those conversations. What usually happens is that we end up talking about other things, but then it [7ei] brings us back around then too. So, we may be talking about *Prevalence* and *Impact* but when we are on *Relationships* or *Regulation*, we can start tying those to *Prevalence* and how we may be able to help with an intervention or explain why *Regulation* is off. So, I like connecting the dots through TIC and having conversations and processing it [7ei] with my staff." - P8, Child Welfare Supervisor

Implementation of 7ei (i.e., the fourth component) has also helped the agency address some problems in CWS. Lack of supervisory support is one of the main reasons for high turnover in child welfare [25]. 7ei meetings ensure that new caseworkers are receiving adequate supervision and support from experienced members where they are also consistently using the trauma-informed framework. Agency also decided to train only one member of the child welfare team (permanency consultant or supervisor) on the definitions and scoring criteria for variables where they facilitate the meetings and utilization of the tool. This ensures that the rest of the team is able to freely brainstorm as they work through TIC principles. Moreover, any decision-support tool runs the risk of becoming a safe default in a high-stakes environment where the workers carry high workloads and shifts decision-making accountability away from the workers. Here, the 7ei meetings provide the adequate time and space for the tool's collaborative and ethical use.

4.4.3 7ei Improves Street-level Decision-Making within Systemic Constraints.

7ei algorithm is constructed using TIC principles and is embedded within the agency's service delivery model that child welfare staff must follow. One supervisor shared that the caseworker position experienced high turnover such that the agency was always understaffed. Here, it becomes imperative that the agency is creating mechanisms to train new caseworkers where they are employing evidence-based best practices in street-level decision-making when they work with families. 7ei algorithm and 7ei meetings have helped create this training process -

"7ei is embedded in the agency. Everything is based in the trauma-responsive model and often caseworkers are working through 7ei domains without even realizing it" - P9, Child Welfare Supervisor

However, participants also shared that even though 7ei is leading to better outcomes for some cases, it also ends up being impractical for cases where critical decisions come down to systemic or legal barriers. For instance, a family may have completed trauma-informed interventions and made changes to their household to provide a safe environment for children, however, the district attorney's office might still advocate against reunification because of the family's past (e.g., criminal history, domestic violence, drug use). One supervisor explained -

"7ei is useful but so many decisions come down to systemic problems. Like..for example..the biggest barriers can come from legal parties. They are not in the [parent's] home every month. They are not talking to these parents, day in and day out. So, sometimes it can be challenging to try and fight for reunification." - P12, Child Welfare Supervisor

Here, the supervisor alludes to ongoing systemic tensions between child welfare staff and the court system that are well-documented in social work literature [26, 39, 42]. 7ei was not designed to address these systemic issues, however, it still adds to the frustrations of caseworkers who continually employ trauma-informed practices using 7ei but are unable to receive a favorable decision in court for their clients.

5 DISCUSSION

In this section, we discuss how algorithmic tools at the agency are inadvertently causing harm to the nature of practice, harm to the administration at the agency, as well as uncertainties in street-level decision-making. Next, we discuss implications for human-AI decision-making and implications for responsible data science practices in the public sector.

5.1 Algorithms Harms to Social Work Practice

The child welfare system has traditionally suffered from inconsistent decision-making with respect to child safety and family maintenance [24]. The problem of inconsistent decision-making is further aggravated by the fact that CWS experiences chronic turnover with the majority of caseworkers quitting within the first two years [12, 112]. However, research in evidence-based social work suggests that it takes caseworkers about two years to learn how to do the job

and adeptly navigate and interact with the legal parties, service providers, children, and families, and learn the nature of child welfare practice [40]. Consequently, inexperienced caseworkers who use risk assessment algorithms such as CANS are led to believe that they are acting in an unbiased and objective manner. But as our results in this study and prior work [107] indicate, bias can be embedded within these assessments and the underlying data collected as a result of the practices of similarly inexperienced caseworkers. For instance, two of the most significant predictors of risk of maltreatment per the WARM risk assessment are *parents' cooperation with the agency* and *stress of caretaker* [107]. These variables are singly scored by caseworkers and encapsulate their impressions of the family with no input from the parents themselves. In addition, such variables capture a parent's response to the agency's intervention in their lives rather than the effectiveness or means of the intervention itself (i.e., how the caseworker approached and engaged with the family). child welfare is supposed to transition towards a "families as partners" model, however, algorithmic tools are creating a layer of obfuscation where such power asymmetries can become embedded within these tools. Ironically, risk assessments have worsened the turnover problem in CWS. Caseworkers leave due to their frustrations with practices and working conditions that have been exacerbated by the adversarial nature of risk assessments [23]. In addition, what remains missing from the conversations about algorithms in child welfare, and the public sector in general, is the significant amount of data labor that is required of caseworkers as well as the repair work [32, 63, 69] necessary to bring these systems to work for families and not just the governmental apparatus.

In addition, caseworkers at the agency were frustrated that the CANS algorithm required them to provide data labor where they must collect information about children and feed it to the system, however, the system also stripped them of discretionary power in regard to these decisions. As highlighted by recent work [62], predictive systems in high-stakes domains are *extractive by design* where they lead to a systemic extraction of discretionary power such that probabilistic outcomes are being used to supplant workers' contextual knowledge. Moreover, this interaction between caseworkers and algorithmic tools is further problematic because *human-in-the-loop* is the often proposed solution to erroneous decisions made by algorithms [35]. However, in this scenario, humans may be just as likely to make mistakes. Power asymmetries and incentive structures also directly impact how data is collected about children and how these assessments are scored. For instance, the CANS algorithm is conducted by both clinical therapists and caseworkers. However, a prior study [77] conducted on CANS found that clinical therapists were more likely to detect mental health needs because they are medically trained to detect these needs but also because they provide services when needs are detected, and consequently, they are paid for providing these services. That is, there was a clear financial incentive to exaggerate CANS scores. On the contrary, caseworkers very significantly less likely to detect mental health needs because 'detecting needs' inadvertently created more work for them. Once needs are detected, caseworkers must reach out to service providers and secure appointments for their clients.

5.2 Algorithmic Harms to the Child Welfare Agency

child welfare agencies in several states in the United States have continued to rely on counsel from the federal government in the form of initiatives, regulations, and evidence-based approaches to improve their practice model. As previously noted, recent federal initiatives have laid the groundwork for more algorithmic interventions in CWS. However, federal directives have continually focused on the need for CWS agencies to adopt data-driven practices without providing adequate guidelines that focus on the *why* and *how* CWS agencies could employ evidence-based data-driven practices within their day-to-day processes in addition to training caseworkers on these practices. Consequently, CWS agencies in several states have rushed to adopt "something" in order to prove that they are employing scientific and evidence-based practices without ensuring that child welfare stakeholders have a strong understanding of how the model works, how

to assure fidelity, and how to assess the model for issues of ethics and equity. The Allegheny Family Screening Tool (AFST) is a case in point of this scenario. A recent study conducted by Kawakami et al. [66] found that call screeners were offered minimal information about the working of AFST, considered it to be unreliable due to unexpected model behaviors, and even engaged in a collaborative "game" to learn more about the tool. Moreover, Cheng and Stapleton et al. [29] found that AFST-only decisions were racially biased and workers mitigated these biases by overriding erroneous decisions. Yet, these (sometimes hastily adopted) data-driven risk assessment models have become a central activity in many child welfare organizations [23]. As child welfare agencies across the United States begin to adopt the new CCWIS data model [56], it has also paved the way for tech startups [1–4] to start pitching CCWIS-based algorithmic tools to agencies to help them meet their federal and state accountability requirements [34].

At this agency, there are serious data provenance concerns about data collected about children through the CANS algorithm since the data is so heavily manipulated by both caseworkers and foster parents. CANS was reappropriated to calculate foster parent compensations because policymakers believed that it would offer a fair and unbiased means for allocating resources. In addition, it would reduce costs over time since the improvement in care would lower the mental health needs and also the resources required to provide care. However, data specialists at the agency shared with us several cases where the compensation continually increased over time because the algorithm was being gamed. Another unintended consequence of exaggerating CANS scores is that foster children are now being sent to services (e.g., individual therapy) that they don't necessarily need. This is an added financial burden on an underfunded system that must pay for these unnecessary services. Here, CANS has added more barriers to consistent and evidence-based decision-making and introduced more constraints within the process. Moreover, caseworkers are mandated to provide the data labor that runs AI systems, however, they have no agency over these data production or future data utilization processes. This non-profit agency is contracted by the state's Department of Children and Families to provide child welfare services and the agency must use DCF's centralized data system (i.e., eWiSACWIS) to record all case data. However, they are unable to download this data from their cases for the purpose of critiquing and improving their own practices. Skeptics of AI have invoked ethical frames where engineers and designers of algorithms must refuse to design systems that may raise social and ethical concerns [13]. However, such concerns might not arise at the onset, and as Selbst et al. note - "repurposing algorithmic solutions designed for one social context may do harm when applied to a different context" [109]. CANS algorithm is a case in point of this scenario which was developed by academics with all the best intentions and as a medium to facilitate sharing of information about children's needs. Refusal to design may also not be an option for most public sector agencies who must continually prove that they are employing innovative data-driven practices and may also be under political pressure to adopt algorithmic tools that seek to automate public service delivery as well as measure performance standards that impact future contracts.

5.3 Algorithmic Harms to Street-Level Decision-Making

"All government policy and regulation is contradictory. That's why we exist, to fill in these gaps and make sense of policy in implementing it. This CANS is basically policy, right? So we interpret it and bend it however necessary to make it work for the people. I don't think we are gaming the algorithm. Because that would mean we are gaming all policy." – P7, Program Director

With inadequate federal guidelines on how to adopt, implement, and use algorithmic tools, there is significant misunderstanding regarding the roles of these tools in implementing rules or standards. Rules can be thought of as triggering criteria, i.e. - authoritative conditions that can be algorithmically coded for. On the other hand, street-level bureaucrats have traditionally exercised a significant amount of autonomy and flexibility in applying professional

standards. CANS algorithm offers a case in point where it is expected to be implemented as a rule, however, caseworkers engage with it as if it were a standard. In addition, the program director in the above quote alludes to the *repair work* [63] that caseworkers undertake in order to make the algorithm work for their clients. Caseworkers are expected to use algorithmic tools that shift discretion away from them, however, they are also expected to assume responsibility when automated decisions lead to poor outcomes for families. Recently, public sector algorithms have received criticism when they lead to poor outcomes and exacerbate racial disparities [43, 121, 127]. However, finding the sources of harm can be difficult when we witness the distributed use of smaller algorithmic tools as opposed to a larger and more visible algorithmic system. Distributed nature of several of these tools also shifts accountability from one group to another within the agency, and therefore, can be hard to assess. For instance, as highlighted in prior work [101], the Anti Sex-Trafficking algorithm refers cases to the Human Anti-Trafficking Response Team (HART) based on a set of risk predictors. However, HART is now receiving an influx of cases with no added resources provided to them. Consequently, cases that really need their expertise are receiving inadequate attention; an indirect harm that is hard to measure. As illustrated in this case study, assessments such as CANS are transformed into algorithmic decision-supports which are now further being legally mandated to ensure caseworkers' compliance. The AFST algorithm implemented at a child welfare agency in Allegheny County, Pennsylvania is another example of this scenario where it was initially implemented in a voluntary capacity, however, its use was mandated when it did not receive enough engagement from caseworkers. AFST is still characterized as an algorithmic decision-support but offers an interesting example of how such tools are increasingly limiting human discretionary work instead of findings ways to improve it.

6 IMPLICATIONS FOR HUMAN-AI DECISION-MAKING

Public sector agencies such as the child welfare agency examined in this study offer an important setting for studying human-AI interaction for the following reasons - 1) full automation in such settings is not desirable because of ethical, legal, and safety concerns associated with high-stakes decision-making [71]. This leads to observable ongoing interactions between practitioners and algorithmic systems where practitioners learn to engage with such systems, assess their utility in informing decision-making, as well as develop perspectives regarding critical factors that lead to the development of (dis)trust of such systems, and 2) decisions are collaboratively made by the child welfare staff which allows us to inspect how a team of caseworkers with varying levels of expertise interact and reason with an algorithmic system and incorporate algorithmic decisions within their decision-making processes.

As Lai et al. [72] note, the underlying assumption for AI-assisted decision-making is that human reasoning can benefit from counter-intuitive patterns derived from empirical models. Such prediction tasks are also generally referred to as *AI for discovery* tasks [71]. For instance, ascertaining the risk of maltreatment as a function of prior involvement with public services (e.g., drug and alcohol services, housing authority, etc.) might help uncover patterns that aid decision-making and prevent future child maltreatment events. However, **we found that this led to a cumulative distrust of such systems because the decision task was no longer aligned with the caseworkers' practice model and did not help inform their day-to-day decisions.** For instance, caseworkers across jurisdictions [29, 101] shared that the knowledge of predicted long-term risk (i.e., the likelihood of being re-referred or placed in foster care within two years) did not help address risk factors that the families experience in the system (e.g., unstable housing or employment, lack of access to medical services, risk posed by procedural factors) and did not help them acquire services for their clients [105]. That is, there is a significant difference between the *target variable* that the human decision-makers care about and the *target proxy* being predicted. Guerdan et al. [52] use causal diagrams to unpack the relationship between predictors, decisions, target variables, and their proxies in the human-AI decision space and show

how outcome measurement error occurs due to this gap between the target variable and its proxy. In this study, **we show how outcome measurement error can result in process-oriented algorithmic harms in practice.**

In addition, for empirical predictions to provide any utility, it is imperative that caseworkers are provided with explanations of predictions that especially highlight these counter-intuitive factors that do not traditionally align with their practice. However, these explanations can differ based on the practitioner and their decision-making goals. Different practitioners had different needs in regard to explainability and a “one-size-fits-all” approach with respect to algorithm design may not be feasible here. For instance, the clinical therapist was more interested in explanations about the context surrounding worsening behaviors whereas the caseworkers wanted explanations about underlying trauma. As Ehsan et al. [41] note, explainability is a human factor and not just a model-inherent property that requires a careful examination of the sociotechnical gap, i.e. - the divide between the technical affordances and the social needs [5]. We direct the readers to Ehsan et al.’s work [41] where they operationalize a framework for explainable AI (XAI) and highlight how the framework can help researchers better understand the sociotechnical gap in different contexts and practical considerations for how we can begin to bridge this gap.

Most of the *AI for discovery* tasks require domain expertise but domain experts are seldom included in the design of AI systems [71]. The majority of studies conducted on human-AI interaction are conducted as isolated experimental trials of decision tasks in lab studies or on crowdsourcing platforms. Consequently, as is evident from the case study of the CANS algorithm, we found that such predicted decision tasks developed without domain expertise did not generalize to decision tasks carried out by caseworkers within the bounds of policies and systemic resources. In addition, seemingly cogent and well-aligned decision tasks can be at odds with one another in practice. In theory, it makes sense to predict the foster care placement setting and the subsidized guardianship rate based on the mental health assessment of foster children. However, in practice, we learned that due to a lack of foster homes in the system, the second-removed outcome (i.e., subsidized guardianship rate) became the primary outcome of interest and impacted how mental health assessments were conducted. **That is, a systemic issue is the critical driving factor that determines how decision tasks are carried out.** Here, ethnographic engagement is necessary to fully understand the decision-making ecosystem and assess how human-AI decision-making may occur within the bounds of organizational pressures, time constraints, and limited resources and which disruptions are likely to occur in a resource-constrained environment.

In high-stakes domains where improving human-AI interaction is paramount, AI research must be encapsulated within HCI methods because it helps us better understand stakeholders’ interactions with the system and assess critical human factors such as trust, reliance, intelligibility, and explainability of the system. In addition, as noted above, AI systems designed for high-stakes domains often seek to replicate or replace human labor without an adequate understanding of the nature of practice and the reality of street-level decision-making. Here, HCI methods can help us understand organizational workflows, day-to-day protocols that workers are expected to follow, systemic factors (e.g., policies, resources, etc.), and the bottlenecks in decision-making. Consequently, this helps us outline the systemic mechanics as well as the insertion points, scope, and utility of algorithmic decision-making. That is, designing systems that algorithmically support the practitioners’ cognitive environment with or without predicting decisions. The 7ei algorithm provides an interesting case study of a simple algorithmic tool specifically designed to overcome the obstacles described above. Having faced several frustrations with empirical models, the agency leadership engaged in theory-driven design and developed 7ei using the principles of trauma-informed care (TIC). This further allowed them to ensure that decisions made using 7ei are centered in the theory of practice. **The agency also ensured that child welfare staff were trained in TIC, created a collaborative setting to ensure the proper use of the tool, and decided to track multiple TIC outcomes over the life of cases instead of predicting them.** Consequently, what

we observed over an extended ethnography was that 7ei received collective buy-in and significantly better engagement from the caseworkers because the tool was centered in social work practice.

As highlighted by Vaughan and Wallach [122], it is imperative to understand the needs of relevant stakeholders to be able to design for intelligibility techniques that meet these needs. It is imperative to note here that 7ei was designed to cultivate the human factors of trust, reliance, and intelligibility. It facilitates the child welfare staff's intelligibility needs by 1) providing relevant information, 2) centering this information in a theoretical framework and 3) demonstrating compliance with evidence-based practice. The tool was also designed with the intent to decompose the algorithm and turn it into an open-ended and transparent process that promotes trust and reliance. The collaborative use of 7ei also helped establish informed mistrust of the tool for cases where decisions were dictated by policy and/or systemic constraints with less utility for TIC principles. However, it is necessary to note that too many interactions where the tool did not inform decision-making (due to persisting systemic factors) may lead caseworkers to distrust and abandon the tool altogether. Here, from a human-AI interaction standpoint, it is imperative to note that **designing for intelligibility needs preceded designing for explainability or transparency and impacted the caseworkers' engagement and reliance on the tool**. Transparency and explainability are often endorsed as human factors that might facilitate engagement and improve decision-making (i.e., making the workings of the algorithm transparent and providing explanations for predictions). Here, the case study of the CANS algorithm provides a cautionary tale against prioritizing these human factors over the intelligibility needs of stakeholders. Even though CANS promotes transparency where all the predictors are accessible to caseworkers, it frustrates the child welfare staff because it does not account for the principles of TIC and only offers an isolated view of the child. This further problematizes explainability because caseworkers were more invested in explanations of underlying trauma than exhibited behaviors. Consequently, we learned that transparency in the case of CANS promoted the manipulation of data, led to the diminishing of trust and reliance, and forced caseworkers to engage in repair work to make the tool work for their clients.

7 IMPLICATIONS FOR RESPONSIBLE DATA SCIENCE IN THE PUBLIC SECTOR

Crucial to the discussion of responsible data science practices [11] in the public sector are the following points. First, it is imperative to recognize that decision-making is a complex process where information needs to be shared among several parties (e.g., child welfare staff, district attorney's office, parents' attorneys, service providers, and judges) and decisions are collaboratively made. As illustrated by Figure 2, **algorithmic tools must be designed for collaborative use** for them to offer higher utility to practitioners and improve their decision-making processes. In addition, practitioners must be able to explain decisions made to other involved parties. Therefore, it is necessary to **design algorithms that facilitate explanations**. Here, Ehsan et al's framework for operationalizing explainable AI [41] offers important practical considerations for designing for explanations when the nature of practice is highly collaborative. Designing for collaborative use and explanations will aid deeper integration between human discretion and algorithmic decision-making, and consequently, lead to improved human-AI decision-making over time. In addition, **algorithm design must occur within organizational constraints**, i.e.- it is imperative to account for the resource and legislative constraints within which all decisions (human or algorithmic) must be made. In the absence of these design requirements, algorithms frustrate practitioners who are unable to follow the 'ideal algorithmic decision'. For instance, CANS might recommend that a foster child be placed in a Level 4 foster home, however, with a lack of good foster homes in the system, the caseworker must manipulate data to produce an outcome that points to an available placement [100]. Consequently, algorithmic tools that account for such constraints are more likely to be integrated as an organizational process as well as enhance the tool's trustworthiness among practitioners who must interact with it collaboratively.

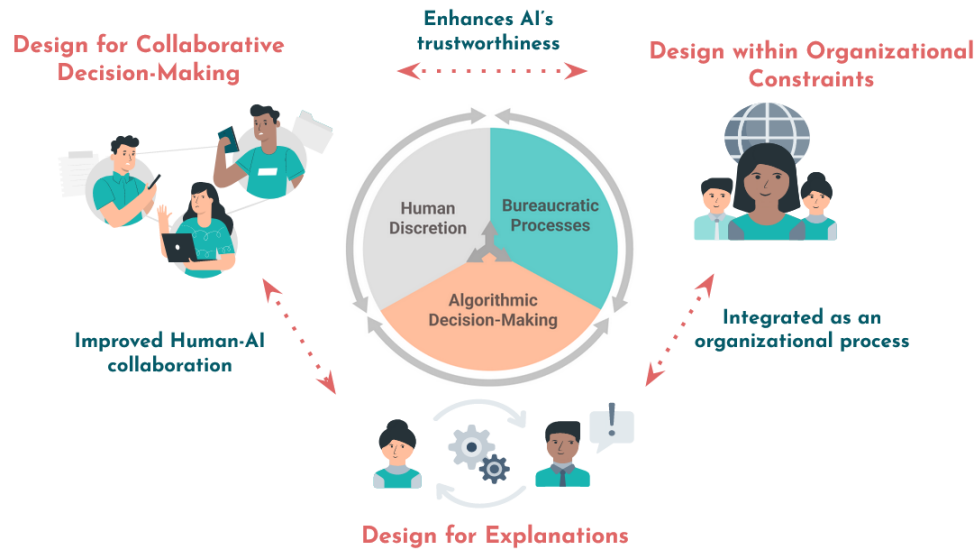


Fig. 2. Designing Algorithms for the Public Sector through the Lens of ADMAPS Framework

Second, as highlighted by the adoption and use of the 7ei algorithm, agency leaders had to invest a significant amount of resources in terms of trainings, specialized consultations, hiring experts, and creating the time and space (in terms of collaborative 7ei meetings) to ensure proper utilization of the algorithmic tool. That is, there is a significant amount of human labor and agency resources that went into the integration of 7ei into daily work routines and decision-making processes. It is imperative to note that these investments, especially in the public sector, must be made in order to deconstruct and rebuild decision-making processes that utilize algorithmic systems.

Third, as highlighted by the development of 7ei, the tool was designed with the intent to decompose the algorithm and turn it into an open-ended and transparent process such that it tracked outcomes over time instead of predicting an outcome of interest. As highlighted by prior work, there is an irreducible degree of uncertainty associated with each predicted outcome [35, 88] and this problem is further exacerbated in the public sector where all the relevant information is not always available and there may also be contradicting sources of information [19, 106]. In such scenarios, predicted decisions led to frustrations on the part of caseworkers because there is a significant gap between the *unobserved target variable* and the *observed proxy variable* of interest to the human decision-makers [52]. In sum, there is a need to prescribe away from the black-box model of *input-computation-output* and build tools that algorithmically support the practitioners' cognitive environment without predicting decisions [105, 106]. Designing to support decision-making processes will also lead to fair, transparent, and accountable decisions for clients as well as responsible AI practices from the data science community. Furthermore, as highlighted by prior work [122], there is no need to over-engineer solutions using complex neural networks when simple models have been shown to be just as accurate in high-stakes domains [65, 98, 122]. In addition, caseworkers' frustrations with CANS and the repair work they conducted highlighted key aspects of decision-making that the algorithm failed to capture. This further highlights the need to conduct ethnographic work as well as the need to co-design with stakeholders such that their needs and concerns are addressed. Here, we direct the readers to Yildirim et al.'s work [130] on how to effectively ideate AI concepts and explore AI's problem-solution space with domain experts.

8 LIMITATIONS

We conducted an extensive ethnographic study that highlights the messy interactions between social work practice, policies, and algorithmic decision-making at a child welfare agency. However, this study has some limitations that create opportunities for researchers to further expand upon this body of work. First, this study only focuses on the perspectives of on-the-ground caseworkers and on their interactions with algorithmic systems. It is important to understand the perspective of affected communities (i.e., foster children, parents, and foster parents) about whom decisions are being made through algorithms. For instance, a recent study conducted by Stapleton et al. [114] found that parents considered CWS algorithms to be punitive and unsupportive. Parents instead wanted systems that would help them fight against CPS as well as evaluate CPS and the caseworkers themselves. Future research should continue to focus on uncovering street-level complexities within this complicated sociotechnical environment to understand how different labor practices, systemic constraints, and power asymmetries might amplify algorithmic harms experienced by citizens on the street level [105, 106]. Finally, this ethnographic study only uncovers complexities at one child welfare agency, however, child welfare practices and policies can significantly vary from one state to another. Therefore, we recommend that researchers conduct similar qualitative explorations in other jurisdictions.

9 CONCLUSION

We conducted an in-depth ethnographic study to understand the daily algorithmic practices of caseworkers at a child welfare agency. We qualitatively coded our data from the ethnography to the dimensions of the ADMAPS framework to reveal the complex interdependencies between human discretion, algorithmic decision-making, and bureaucratic processes. We focused on the macro-interactions between the three core elements of ADMAPS to understand how the interplay between systemic mechanics and algorithmic decision-making impacts the fairness of the decision-making process itself. Our findings highlight how functionality issues in algorithmic systems can lead to process-oriented harms where they adversely impact the nature of professional practice, and administration at the agency, and lead to inconsistent and unreliable street-level decision-making. Here, caseworkers are compelled to undertake additional labor in the form of repair work to restore disrupted administrative processes and decision-making at the street level, all while facing organizational pressures and time and resource constraints. We learned that there is a need to focus on the proper implementation and integration of algorithmic tools into decision-making processes and not just the initial development and deployment of the algorithmic model. That is, there is a need to rethink the amount of investment required to ensure the proper adoption of algorithms in complex sociotechnical environments. In addition, we show how a simple algorithmic tool that tracks variables over time instead of predicting an outcome offered higher utility to caseworkers. Here, algorithms need to be designed to support explanations and collaborative use such that they augment human discretionary work without seeking to replace or replicate it. In addition, algorithmic tools need to be fully supported by bureaucratic processes by allocating necessary resources and accounting for organizational constraints to ensure that they are integrated as an organizational process. As a result of this study, we also offer implications for responsible data science practices in the public sector.

ACKNOWLEDGMENTS

This work was initially presented at Data & Society's "The Social Life of Algorithmic Harms" academic workshop held in March 2022. We sincerely thank the organizing team at Data & Society as well as the participants and reviewers for their valuable suggestions, detailed feedback, and most of all for their intellectual generosity.

REFERENCES

- [1] 2022. Analytics for Child Well-Being. *SAS Institute* (Accessed: March 2022). https://www.sas.com/en_us/software/analytics-for-child-well-being.html
- [2] 2022. ECKERD RAPID SAFETY FEEDBACK. *SAS Institute* (Accessed: March 2022). <https://eckerd.org/family-children-services/ersf/>
- [3] 2022. Improving Outcomes using data you already have. *MindShare Technology* (Accessed: March 2022). <https://mindshare-technology.com/analytics/>
- [4] 2022. Stopping Child Maltreatment Before it Happens. *Predict-Align-Prevent* (Accessed: March 2022). <https://papreports.org/little-rock-ar/>
- [5] Mark S. Ackerman. 2000. The Intellectual Challenge of CSCW: The Gap between Social Requirements and Technical Feasibility. *Hum.-Comput. Interact.* 15, 2 (Sept. 2000), 179–203.
- [6] 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*. 1–13.
- [7] Samaneh Aminikhanghahi and Diane J Cook. 2017. A survey of methods for time series change point detection. *Knowledge and information systems* 51, 2 (2017), 339–367.
- [8] Asbjørn Ammitzbøll Flügge, Thomas Hildebrandt, and Naja Holten Møller. 2021. Street-level algorithms and AI in bureaucratic decision-making: A caseworker perspective. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–23.
- [9] Kim Normann Andersen, Jungwoo Lee, and Helle Zinner Henriksen. 2020. Digital sclerosis? Wind of change for government and the employees. *Digital Government: Research and Practice* 1, 1 (2020), 1–14.
- [10] Mark Andrejevic. 2019. Automating surveillance. *Surveillance & Society* 17, 1/2 (2019), 7–13.
- [11] Cecilia Aragon, Shion Guha, Marina Kogan, Michael Muller, and Gina Neff. 2022. *Human-centered data science: an introduction*. MIT Press.
- [12] Michàlle E Mor Barak, Amy Levin, Jan A Nissly, and Christianne J Lane. 2006. Why do they leave? Modeling child welfare workers' turnover intentions. *Children and Youth Services Review* 28, 5 (2006), 548–577.
- [13] Solon Barocas, Asia J Biega, Benjamin Fish, Jędrzej Niklas, and Luke Stark. 2020. When not to design, build, or deploy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 695–695.
- [14] Joan M Blakey, Sonya J Leathers, Michelle Lawler, Tyreasa Washington, Chiralaine Natschke, Tonya Strand, and Quenette Walton. 2012. A review of how states are addressing placement stability. *Children and Youth Services Review* 34, 2 (2012), 369–378.
- [15] Susanne Bødker, Christian Dindler, and Ole Sejer Iversen. 2017. Tying knots: Participatory infrastructuring at work. *Computer Supported Cooperative Work (CSCW)* 26, 1-2 (2017), 245–273.
- [16] Mark Bovens and Stavros Zouridis. 2002. From street-level to system-level bureaucracies: how information and communication technology is transforming administrative discretion and constitutional control. *Public administration review* 62, 2 (2002), 174–184.
- [17] Sarah Brayne. 2017. Big data surveillance: The case of policing. *American sociological review* 82, 5 (2017), 977–1008.
- [18] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-making in Child Welfare Services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 41.
- [19] Justin B Bullock. 2019. Artificial intelligence, discretion, and bureaucracy. *The American Review of Public Administration* 49, 7 (2019), 751–761.
- [20] Justin B Bullock, Jesper Rosenberg Hansen, and David J Houston. 2018. Sector Differences in Employee's Perceived Importance of Income and Job Security: Can These be Found Across the Contexts of Countries, Cultures, and Occupations? *International Public Management Journal* 21, 2 (2018), 243–271.
- [21] Peter André Busch. 2020. Crafting or mass-producing decisions: Technology as professional or managerial imperative in public policy implementation. *Information Polity* 25, 1 (2020), 111–128.
- [22] Peter André Busch and Helle Zinner Henriksen. 2018. Digital discretion: A systematic literature review of ICT and street-level discretion. *Information Polity* 23, 1 (2018), 3–28.
- [23] Marilyn Callahan and Karen Swift. 2018. The paradox of risk assessment, child safety and empowerment in child welfare. In *Revitalising communities in a globalising world*. Routledge, 67–77.
- [24] Michael J Camasso and Radha Jagannathan. 2013. Decision making in child protective services: A risky business? *Risk analysis* 33, 9 (2013), 1636–1649.
- [25] Sarah Carnochan, Megan Moore, and Michael J Austin. 2013. Achieving placement stability. *Journal of evidence-based social work* 10, 3 (2013), 235–253.
- [26] Sarah Carnochan, Sarah Taylor, Anne Abramson-Madden, Meekyung Han, Sonja Rashid, Jennifer Maney, Sarah Teuwen, and Michael J Austin. 2006. Child welfare and the courts: An exploratory study of the relationship between two complex systems. *Journal of Public Child Welfare* 1, 1 (2006), 117–136.
- [27] Robert N Charette. 2018. Michigan's MiDAS unemployment system: Algorithm alchemy created lead, not gold. *IEEE Spectrum* 24 (2018).
- [28] Janet X Chen, Allison McDonald, Yixin Zou, Emily Tseng, Kevin Roundy, Acar Tamersoy, Florian Schaub, Thomas Ristenpart, and Nicola Dell. 2022. Trauma-Informed Computing: Towards Safer Technology Experiences for All. *Under review* (2022).
- [29] 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. 1–22.

- [30] Ka Ho Brian Chor, Gary M McClelland, Dana A Weiner, Neil Jordan, and John S Lyons. 2015. Out-of-home placement decision-making and outcomes in child welfare: A longitudinal study. *Administration and Policy in Mental Health and Mental Health Services Research* 42, 1 (2015), 70–86.
- [31] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. PMLR, 134–148.
- [32] Madeleine Clare Elish and Elizabeth Anne Watkins. 2020. Repairing Innovation: A Study of Integrating AI in Clinical Care. *Data and Society* (2020).
- [33] Victoria Clarke, Virginia Braun, and Nikki Hayfield. 2015. Thematic analysis. *Qualitative psychology: A practical guide to research methods* (2015), 222–248.
- [34] Michael Corrigan. 2019. Using Algorithms and Artificial Intelligence in Child Welfare. *The Imprint* (Jan 2019). <https://imprintnews.org/child-welfare-2/using-algorithms-artificial-intelligence-in-child-welfare/33429>
- [35] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 2020. A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [36] Alejandro De La Garza. 2020. States’ Automated Systems Are Trapping Citizens in Bureaucratic Nightmares With Their Lives on the Line. *TIME* (May 2020). <https://time.com/5840609/algorithm-unemployment/>
- [37] Lynn Dombrowski, Gillian R Hayes, Melissa Mazmanian, and Amy Volda. 2014. E-government intermediaries and the challenges of access and trust. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 2 (2014), 1–22.
- [38] Andy Dow, Rob Comber, and John Vines. 2018. Between grassroots and the hierarchy: Lessons learned from the design of a public services directory. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [39] Joe Duffy and Mary Elizabeth Collins. 2010. Macro impacts on caseworker decision-making in child welfare: A cross-national comparison. *European Journal of Social Work* 13, 1 (2010), 35–54.
- [40] Frank Edwards and Christopher Wildeman. 2018. Characteristics of the front-line child welfare workforce. *Children and Youth Services Review* 89 (2018), 13–26.
- [41] Upol Ehsan, Koustuv Saha, Mummun De Choudhury, and Mark O Riedl. 2023. Charting the Sociotechnical Gap in Explainable AI: A Framework to Address the Gap in XAI. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–32.
- [42] Raquel T Ellis. 2010. Child welfare workers’ perceptions of juvenile court influence on child welfare practices. *Journal of Public Child Welfare* 4, 2 (2010), 158–173.
- [43] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press.
- [44] Sergio Fernandez and Hal G Rainey. 2017. Managing successful organizational change in the public sector. In *Debating public administration*. Routledge, 7–26.
- [45] NCCPR: National Coalition for Child Protection Reform. 2017. Los Angeles County quietly drops its first child welfare predictive analytics experiment. Link: <https://www.nccprblog.org/2017/05/los-angeles-county-quietly-drops-its.html>.
- [46] Mahita Gajanan. 2020. The Heartbreaking Story Behind Netflix’s Documentary Series The Trials of Gabriel Fernandez. *Time.com* (Mar 2020). <https://time.com/5790549/gabriel-fernandez-netflix-documentary/>
- [47] 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*. 1292–1310.
- [48] Sarah Giest and Nadine Raaphorst. 2018. Unraveling the hindering factors of digital public service delivery at street-level: the case of electronic health records. *Policy Design and Practice* 1, 2 (2018), 141–154.
- [49] Connie Golsteijn, Sarah Gallacher, Licia Capra, and Yvonne Rogers. 2016. Sens-Us: Designing innovative civic technology for the public good. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. 39–49.
- [50] DeLeith Duke Gossett. 2017. The Client: How States Are Profiting from the Child’s Right to Protection. *U. Mem. L. Rev.* 48 (2017), 753.
- [51] Nina Grgić-Hlača, Muhammad Bilal Zafar, Krishna P Gummedi, and Adrian Weller. 2018. Beyond distributive fairness in algorithmic decision making: Feature selection for procedurally fair learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [52] Luke Guerdan, Amanda Coston, Zhiwei Steven Wu, and Kenneth Holstein. 2023. Ground (less) Truth: A Causal Framework for Proxy Labels in Human-Algorithm Decision-Making. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 688–704.
- [53] Martin Guggenheim. 2021. How Racial Politics Led Directly to the Enactment of the Adoption and Safe Families Act of 1997—the Worst Law Affecting Families Ever Enacted by Congress. *Columbia Journal of Race and Law* 11, 3 (2021).
- [54] Melissa Hanson. 2021. After release of scathing report on 14-year-old David Almond’s death, Massachusetts DCF says worker fired and changes are underway. *MASSLIVE* (Mar 2021). <https://www.masslive.com/boston/2021/03/after-release-of-scathing-report-on-14-year-old-david-almonds-death-massachusetts-dcf-says-worker-fired-and-changes-are-underway.html>
- [55] MD Romael Haque, Katherine Weathington, Joseph Chudzik, and Shion Guha. 2020. Understanding Law Enforcement and Common Peoples’ Perspectives on Designing Explainable Crime Mapping Algorithms. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 269–273.
- [56] Teresa M Harrison, Donna Canestraro, Theresa Pardo, Martha Avila-Marilla, Nicolas Soto, Megan Sutherland, Brian Burke, and Mila Gasco. 2018. A tale of two information systems: transitioning to a data-centric information system for child welfare. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age*. 1–2.

- [57] Sara Heitlinger, Nick Bryan-Kinns, and Rob Comber. 2019. The right to the sustainable smart city. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [58] Robin K Hill. 2016. What an algorithm is. *Philosophy & Technology* 29, 1 (2016), 35–59.
- [59] Sally Ho and Garance Burke. 2022. An algorithm that screens for child neglect raises concerns. *Associated Press* (Apr 2022). <https://apnews.com/article/child-welfare-algorithm-investigation-9497ee937e0053ad4144a86c68241ef1>
- [60] Sally Ho and Garance Burke. Associated Press. April 2022. An algorithm that screens for child neglect raises concerns. Link: <https://apnews.com/article/child-welfare-algorithm-investigation-9497ee937e0053ad4144a86c68241ef1>.
- [61] Naja Holten Møller, Irina Shklovski, and Thomas T Hildebrandt. 2020. Shifting concepts of value: Designing algorithmic decision-support systems for public services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–12.
- [62] Sun-ha Hong. 2022. Prediction as Extraction of Discretion. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 925–934.
- [63] Steven J Jackson. 2013. Rethinking repair: breakdown, maintenance and repair in media and technology studies today. *Media meets technology, MIT Press, forthcoming* 68 (2013).
- [64] Khari Johnson. 2023. Algorithms Allegedly Penalized Black Renters. The US Government Is Watching. *WIRED* (Jan 2023). <https://www.wired.com/story/algorithms-allegedly-penalized-black-renters-the-us-government-is-watching/>
- [65] Jongbin Jung, Connor Concannon, Ravi Shroff, Sharad Goel, and Daniel G Goldstein. 2017. Simple rules for complex decisions. *arXiv preprint arXiv:1702.04690* (2017).
- [66] Anna Kawakami, Venkatesh Sivaraman, Hao-Fei Cheng, Logan Stapleton, Yanghui Cheng, Diana Qing, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. Improving Human-AI Partnerships in Child Welfare: Understanding Worker Practices, Challenges, and Desires for Algorithmic Decision Support. In *CHI Conference on Human Factors in Computing Systems*. 1–18.
- [67] Emily Keddell. 2015. The ethics of predictive risk modelling in the Aotearoa/New Zealand child welfare context: Child abuse prevention or neo-liberal tool? *Critical Social Policy* 35, 1 (2015), 69–88.
- [68] Emily Keddell. 2019. Algorithmic justice in child protection: Statistical fairness, social justice and the implications for practice. *Social Sciences* 8, 10 (2019), 281.
- [69] Katherine C Kellogg, Melissa A Valentine, and Angele Christin. 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14, 1 (2020), 366–410.
- [70] Tzu-Sheng Kuo, Hong Shen, Jisoo Geum, Nev Jones, Jason I Hong, Haiyi Zhu, and Kenneth Holstein. 2023. Understanding Frontline Workers’ and Unhoused Individuals’ Perspectives on AI Used in Homeless Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [71] Vivian Lai, Chacha Chen, Q Vera Liao, Alison Smith-Renner, and Chenhao Tan. 2021. Towards a science of human-ai decision making: a survey of empirical studies. *arXiv preprint arXiv:2112.11471* (2021).
- [72] Vivian Lai, Chacha Chen, Alison Smith-Renner, Q Vera Liao, and Chenhao Tan. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1369–1385.
- [73] Ida Lindgren, Christian Østergaard Madsen, Sara Hofmann, and Ulf Melin. 2019. Close encounters of the digital kind: A research agenda for the digitalization of public services. *Government Information Quarterly* 36, 3 (2019), 427–436.
- [74] Michael Lipsky. 2010. *Street-level bureaucracy: Dilemmas of the individual in public service*. Russell Sage Foundation.
- [75] Thomas Lodato and Carl DiSalvo. 2018. Institutional constraints: the forms and limits of participatory design in the public realm. In *Proceedings of the 15th Participatory Design Conference: Full Papers-Volume 1*. 1–12.
- [76] John S Lyons. 2009. *Communitetrics: A communication theory of measurement in human service settings*. Springer Science & Business Media.
- [77] John S Lyons, Dana Aron Weiner, and Melanie Buddin Lyons. 2004. Measurement as communication in outcomes management: The child and adolescent needs and strengths (CANS). *The Use of Psychological Testing for Treatment Planning and Outcomes Assessment. Volume 2: Instruments for Children and Adolescents* (2004).
- [78] John S Lyons, Dana Aron Weiner, and Melanie Buddin Lyons. 2014. Measurement as communication in outcomes management: The child and adolescent needs and strengths (CANS). In *The use of psychological testing for treatment planning and outcomes assessment*. Routledge, 461–476.
- [79] Peter Lyons, Howard J Doueck, and John S Wodarski. 1996. Risk assessment for child protective services: A review of the empirical literature on instrument performance. *Social Work Research* 20, 3 (1996), 143–155.
- [80] Helen Margetts. 2012. *Information technology in government: Britain and America*. Routledge.
- [81] Kelly McConvey, Shion Guha, and Anastasia Kuzminykh. 2023. A Human-Centered Review of Algorithms in Decision-Making in Higher Education. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [82] Albert Meijer. 2018. Datapolis: A public governance perspective on “smart cities”. *Perspectives on Public Management and Governance* 1, 3 (2018), 195–206.
- [83] Amanda Meng, Carl DiSalvo, and Ellen Zegura. 2019. Collaborative data work towards a caring democracy. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.
- [84] Ines Mergel, R Karl Rethemeyer, and Kimberley Isett. 2016. Big data in public affairs. *Public Administration Review* 76, 6 (2016), 928–937.
- [85] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–37.

- [86] Trine Rask Nielsen, Maria Menendez-Blanco, and Naja Holten Møller. 2023. Who Cares About Data? Ambivalence, Translation, and Attentiveness in Asylum Casework. *Computer Supported Cooperative Work (CSCW)* (2023), 1–50.
- [87] NPR. June 2022. Oregon is dropping an artificial intelligence tool used in child welfare system. Link: <https://www.npr.org/2022/06/02/1102661376/oregon-drops-artificial-intelligence-child-abuse-cases>.
- [88] Juho Pääkkönen, Matti Nelimarkka, Jesse Haapoja, and Airi Lampinen. 2020. Bureaucracy as a Lens for Analyzing and Designing Algorithmic Systems. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [89] Bruce D Perry. 2009. Examining child maltreatment through a neurodevelopmental lens: Clinical applications of the neurosequential model of therapeutics. *Journal of Loss and Trauma* 14, 4 (2009), 240–255.
- [90] Nancy D Polikoff and Jane M Spinak. 2021. Strengthened Bonds: Abolishing the Child Welfare System and Re-Envisioning Child Well-Being. *Colum. J. Race & L.* 11 (2021), 427.
- [91] Julian Posada. 2022. Embedded reproduction in platform data work. *Information, Communication & Society* 25, 6 (2022), 816–834.
- [92] 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*. 959–972.
- [93] Trine Rask Nielsen and Naja Holten Møller. 2022. Data as a Lens for Understanding what Constitutes Credibility in Asylum Decision-making. *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (2022), 1–23.
- [94] Joanna Redden, Lina Dencik, and Harry Warne. 2020. Datafied child welfare services: unpacking politics, economics and power. *Policy Studies* (2020), 1–20.
- [95] Kellie Reed-Ashcraft, Raymond S Kirk, and Mark W Fraser. 2001. The reliability and validity of the North Carolina family assessment scale. *Research on Social Work Practice* 11, 4 (2001), 503–520.
- [96] Jeanne S Ringel, Dana Schultz, Joshua Mendelsohn, Stephanie Brooks Holliday, Katharine Sieck, Ifeanyi Edochie, and Lauren Davis. 2018. Improving child welfare outcomes: balancing investments in prevention and treatment. *Rand health quarterly* 7, 4 (2018).
- [97] Samantha Robertson, Tonya Nguyen, and Niloufar Salehi. 2021. Modeling Assumptions Clash with the Real World: Transparency, Equity, and Community Challenges for Student Assignment Algorithms. *arXiv preprint arXiv:2101.10367* (2021).
- [98] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (2019), 206–215.
- [99] Anjana Samant, Aaron Horowitz, Kath Xu, and Sophie Beiers. 2021. Family surveillance by algorithm: The rapidly spreading tools few have heard of. *American Civil Liberties Union (ACLU)*(2021).
- [100] Devansh Saxena, Karla Badillo-Urquiola, Pamela Wisniewski, and Shion Guha. 2020. Child Welfare System: Interaction of Policy, Practice and Algorithms. In *Companion of the 2020 ACM International Conference on Supporting Group Work*. 119–122.
- [101] Devansh Saxena, Karla Badillo-Urquiola, Pamela Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021).
- [102] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the US Child Welfare System. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [103] Devansh Saxena, Erhardt Graeff, Shion Guha, EunJeong Cheon, Pedro Reynolds-Cuellar, Dawn Walker, Christoph Becker, and Kenneth R Fleischmann. 2020. Collective Organizing and Social Responsibility at CSCW. In *conference companion publication of the 2020 on computer supported cooperative work and social computing*. 503–509.
- [104] Devansh Saxena and Shion Guha. 2020. Conducting Participatory Design to Improve Algorithms in Public Services: Lessons and Challenges. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 383–388.
- [105] Devansh Saxena, Erina Seh-Young Moon, Aryan Chaurasia, Yixin Guan, and Shion Guha. 2023. Rethinking “Risk” in Algorithmic Systems Through A Computational Narrative Analysis of Casenotes in Child-Welfare. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [106] Devansh Saxena, Seh Young Moon, Dahlia Shehata, and Shion Guha. 2022. Unpacking Invisible Work Practices, Constraints, and Latent Power Relationships in Child Welfare through Casenote Analysis. In *CHI Conference on Human Factors in Computing Systems*. 1–22.
- [107] Devansh Saxena, Charles Repaci, Melanie D Sage, and Shion Guha. 2022. How to Train a (Bad) Algorithmic Caseworker: A Quantitative Deconstruction of Risk Assessments in Child Welfare. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–7.
- [108] Nick Seaver. 2017. Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society* 4, 2 (2017), 2053951717738104.
- [109] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.
- [110] Nicole Shadowen, Thomas Lodato, and Daria Loi. 2020. Participatory governance in smart cities: Future scenarios and opportunities. In *International Conference on Human-Computer Interaction*. Springer, 443–463.
- [111] Renee Shelby, Shalaleh Rismani, Kathryn Henne, Ajung Moon, Negar Rostamzadeh, Paul Nicholas, YILLA-AKBARI N’MAH, Jess Gallegos, Andrew Smart, and GURLEEN VIRK. 2022. Identifying Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. *arXiv preprint arXiv:2210.05791* (2022).
- [112] Miseung Shim. 2010. Factors influencing child welfare employee’s turnover: Focusing on organizational culture and climate. *Children and Youth Services Review* 32, 6 (2010), 847–856.

- [113] June Simon, Jeeyeon Hong, Peter J Pecora, Geoffrey Cloepfil, Troy Gillespie, Elizabeth Parker, Reed Klein, Thalia Garcia, and Mariah Hardy. 2023. An examination of recent federal and philanthropic funding for child welfare research in the United States. *Children and Youth Services Review* 148 (2023), 106779.
- [114] Logan Stapleton, Min Hun Lee, Diana Qing, Marya Wright, Alexandra Chouldechova, 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 *2022 ACM Conference on Fairness, Accountability, and Transparency*. 1162–1177.
- [115] Logan Stapleton, Devansh Saxena, Anna Kawakami, Tonya Nguyen, Asbjørn Ammitzbøll Flügge, Motahhare Eslami, Naja Holten Møller, Min Kyung Lee, Shion Guha, Kenneth Holstein, et al. 2022. Who Has an Interest in “Public Interest Technology”? Critical Questions for Working with Local Governments & Impacted Communities. In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing*. 282–286.
- [116] Megan Stevenson. 2018. Assessing risk assessment in action. *Minn. L. Rev.* 103 (2018), 303.
- [117] Megan T Stevenson and Jennifer L Doleac. 2022. Algorithmic risk assessment in the hands of humans. *Available at SSRN 3489440* (2022).
- [118] Kele Stewart. 2022. Re-envisioning Child Well-being: Dismantling the Inequitable Intersections Among Child Welfare, Juvenile Justice, and Education. *Columbia Journal of Race and Law* 12, 1 (2022).
- [119] Chicago Tribune. 2017. Data mining program designed to predict child abuse proves unreliable, DCFS says. Link: <https://www.chicagotribune.com/investigations/ct-dcfs-eckerd-met-20171206-story.html>.
- [120] Rebecca Uliasz. 2020. Seeing like an algorithm: operative images and emergent subjects. *AI & SOCIETY* (2020), 1–9.
- [121] Sarah Valentine. 2019. Impoverished algorithms: Misguided governments, flawed technologies, and social control. *Fordham Urb. LJ* 46 (2019), 364.
- [122] Jennifer Wortman Vaughan and Hanna Wallach. 2020. A human-centered agenda for intelligible machine learning. *Machines We Trust: Getting Along with Artificial Intelligence* (2020).
- [123] Michael Veale and Irina Brass. 2019. Administration by algorithm? Public management meets public sector machine learning. *Public management meets public sector machine learning* (2019).
- [124] Eran Vigoda-Gadot and Itai Beeri. 2011. Change-oriented organizational citizenship behavior in public administration: The power of leadership and the cost of organizational politics. *Journal of Public Administration Research and Theory* 22, 3 (2011), 573–596.
- [125] John Villasenor and Virginia Foggo. 2021. Algorithms and housing discrimination: Rethinking HUD’s new disparate impact rule. *Brookings Institution* (Mar 2021). <https://www.brookings.edu/articles/algorithms-and-housing-discrimination-rethinking-huds-new-disparate-impact-rule/>
- [126] Amy Voida, Lynn Dombrowski, Gillian R Hayes, and Melissa Mazmanian. 2014. Shared values/conflicting logics: working around e-government systems. In *Proceedings of the sigchi conference on human factors in computing systems*. 3583–3592.
- [127] Anne L Washington. 2018. How to argue with an algorithm: Lessons from the COMPAS-ProPublica debate. *Colo. Tech. LJ* 17 (2018), 131.
- [128] 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*. 1–15.
- [129] Karen Yeung. 2018. Algorithmic regulation: a critical interrogation. *Regulation & Governance* 12, 4 (2018), 505–523.
- [130] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supriya Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, et al. 2023. Creating Design Resources to Scaffold the Ideation of AI Concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2326–2346.