

A Human-Centered Review of the Algorithms used within the U.S. Child Welfare System

Devansh Saxena

Marquette University
Milwaukee, WI, USA

devansh.saxena@marquette.edu

Karla Badillo-Urquiola

University of Central Florida
Orlando, FL, USA

kbadillo@ist.ucf.edu

Pamela Wisniewski

University of Central Florida
Orlando, FL, USA

pamwis@ucf.edu

Shion Guha

Marquette University
Milwaukee, WI, USA

shion.guha@marquette.edu

ABSTRACT

The U.S. Child Welfare System (CWS) is charged with improving outcomes for foster youth; yet, they are overburdened and underfunded. To overcome this limitation, several states have turned towards algorithmic decision-making systems to reduce costs and determine better processes for improving CWS outcomes. Using a human-centered algorithmic design approach, we synthesize 50 peer-reviewed publications on computational systems used in CWS to assess how they were being developed, common characteristics of predictors used, as well as the target outcomes. We found that most of the literature has focused on risk assessment models but does not consider theoretical approaches (e.g., child-foster parent matching) nor the perspectives of caseworkers (e.g., case notes). Therefore, future algorithms should strive to be context-aware and theoretically robust by incorporating salient factors identified by past research. We provide the HCI community with research avenues for developing human-centered algorithms that redirect attention towards more equitable outcomes for CWS.

Author Keywords

Child Welfare System; Algorithmic Decision-Making; Human-centered Algorithm Design

CCS Concepts

•Applied computing → Computing in government;
•Information systems → Decision support systems;

INTRODUCTION

As of September 2016, there were 437,465 children in the child welfare system (CWS) in United States [87]. This is a significant (10%) rise in just 4 years since September 2012 [87], and this number is expected to keep rising unless significant efforts are made to improve youth outcomes [87]. Child abuse and neglect are severe issues that policymakers in the United States continue to battle with, and which is consistently at the foreground of public policy [37]. In recent years, CWS has been the center of public and media scrutiny [38] because of the potential damage done to the children who are

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

CHI '20, April 25–30, 2020, Honolulu, HI, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-6708-0/20/04...\$15.00

DOI: <https://doi.org/10.1145/3313831.3376229>

removed from the care of their parents [45]. Therefore, there is significant pressure on CWS to systematize the decision-making process and show that these decisions were unbiased and evidenced-based [85]. For most policymakers, algorithmic decisions are perceived to be the epitome of being unbiased, evidence-based, and objective [109, 2]. Thus, algorithms have been developed for almost every aspect of services provided by CWS in different states. For instance, models have been developed for predicting risk of future maltreatment event of a child [110], recommending appropriate placement settings [97] and matching children with foster parents who can meet the unique needs of every child [80]. Many of these algorithms have achieved various degrees of early successes and have shown to reduce costs [93] for CWS. However, they have also come under significant criticisms for being biased [34, 16], being opaque [109], complex and hard to explain [110, 33], being too reductive [36] and non-contextual [99] and for not incorporating factors that arise from relevant social science research literature [30].

The SIGCHI community is at the forefront of research on algorithmic bias [43, 20, 69], and has begun to examine some of the challenges of algorithmic decision-making within CWS. Brown et al. [25] studied community perspectives on algorithmic decision-making systems in CWS and found several aspects of algorithmic systems that bolstered distrust, perpetuated bias, concern over the lack of contextual understanding and ‘black-box’ nature of the algorithms, as well as concerns about how these algorithms may negatively impact child-welfare workers’ decisions. Moreover, scholars outside of HCI have discussed how algorithms impact decision-making in CWS [29, 99, 100, 51]. Engaging in research that helps people and organizations, such as CWS, is well-suited and important for the HCI community. Therefore, a critical step in building a strategic research agenda is to synthesize the breadth of work that has already been done to identify a pathway forward. To forge this path, we posed the following high-level research questions:

RQ1: *What methods have been used to build algorithms in the child welfare system?*

RQ2: *What factors (i.e., independent variables) have been shown to be salient in predicting CWS outcomes?*

RQ3: *What outcomes (i.e., dependent variables) have CWS organizations been predicting?*

To answer these questions, we conducted a comprehensive literature review (n=50) of algorithms used for decision-making

in CWS in the United States. We qualitatively analyzed these articles using the lens of human-centered algorithm design [13]. Overall, we found that majority of the algorithms in CWS are empirically constructed, even though the empirical knowledge is quite fragmented [55]. Our results also revealed considerable differences in the predictors currently being used and those found salient in the child-welfare literature. Finally, CWS has traditionally focused on ‘risk assessment,’ rather than positive outcomes that improve the lives of foster children. Based on Woobrock and Kientz’s encapsulation of research contributions in HCI [117], this paper is a survey of the existing literature and makes the following unique research contributions:

1. We apply a human-centered conceptual framework [13] to critically review the algorithms used within the U.S. child welfare system.
2. We introduce domain knowledge from the child welfare system to embed it within the SIGCHI community to allow for collaborative research between the two disciplines.
3. We identify the potential gaps in the existing literature and recommend future research opportunities with careful attention to the human-centered design of algorithms to benefit CWS.

In the following sections, we discuss Human-Centered Algorithm Design and how we used this framework to inform our literature review methodology. Next, we situate our research within the SIGCHI community.

A Human-Centered Approach to Algorithm Design

As algorithms begin to permeate through every aspect of social life, HCI researchers have begun to ask, "Where is the Human?", that is, recognizing that humans are a critical, if not the central component of many domains for which Artificial Intelligence (AI) systems are being developed. A workshop organized at CHI 2019 [60], tackled this topic to identify several pertinent issues in algorithmic design, such as the opaque and isolated development of algorithms and a lack of involvement of the human stakeholders, who use these systems and are most affected by them. To address these problems, Baumer proposed Human-Centered Algorithm Design (HCAD) [13]; a conceptual framework founded in practices derived from human-centered design [61]. It incorporates human and social interpretations through both the design and evaluation phases [13]. Baumer [13] lays out three strategies that help algorithm design become more human-centered, namely, 1) theoretical, 2) speculative and 3) participatory strategies. We draw from the theoretical perspective to frame our research questions and as the qualitative lens for our analysis. Human-centered theoretical design strategy informs algorithm design as follows:

- **Meaning-making:** Theoretical foundations provide a much-needed scaffolding for dealing with complexity, identifying and evaluating design opportunities [89]. Designers need to study the socio-cultural domain in which they intend to situate their work.
- **Design:** Theoretical approaches aim to incorporate concepts and theories from social sciences into data science [13].

- **Evaluation:** The stakeholders’ social interpretations of results can help ensure that the algorithm has higher utility and integrates well with practice.

CWS is one such domain that suffers from a complete lack of human perspectives through the design process. Therefore, our work focuses on how HCAD strategies can be employed to answer critical research questions in CWS.

BACKGROUND

We situate our research within the SIGCHI community and provide an overview of the work that has been done to develop integrated data systems for CWS.

SIGCHI Research to Support the Child Welfare System

The SIGCHI community has recognized the importance of conducting research with organizations that help disadvantaged communities, such as those experiencing homelessness [108, 118] or recovering from substance abuse [75]. For example, Strohmayer, Comber, and Balaam [108] partnered with a center for people of low social stability to understand homeless young adults’ perceptions of education. Similarly, Woelfer and Hendry [118] created a community technology center at a local service agency to work with homeless young people, case managers, and outreach workers. Similarly, SIGCHI researchers have started to engage with CWS to find ways to improve the lives of youth who have been displaced from their families. Some SIGCHI research has focused on foster youth and parents. For instance, Gray et al.’s [56] research with fostered and adopted children introduces a new digital memory box for creating and storing childhood memories. More recently, researchers have begun to study algorithmic decision-making systems within the child-welfare community. Badillo-Urquiola et al. [9] presented the challenges foster parents face mediating teens’ technology use within the home.

Most relevant to our current work, Brown et al. [25] engaged in a participatory design effort and conducted workshops with families involved in CWS, child-welfare workers, and service providers. They found that participants were uncomfortable with algorithmic systems. Participants felt that these systems used deficit-based frameworks to make decisions and questioned the bias present within the data. Based on their findings, the investigators provide recommendations for researchers and designers to work together with public service agencies to develop systems that provide a higher comfort level to the community. Our study builds upon this related work by critically investigating the algorithms used within CWS and highlighting opportunities for future research. We provide a foundation for implementing human-centered approaches in the design and development of algorithmic systems for CWS.

Sociotechnical Systems for Child-Welfare

In this section, we provide necessary background context about integrated data systems that laid the foundation for algorithmic work in CWS. In 1995, the federal government launched SACWIS (State Automated Child Welfare Information System) initiative to provide states with a federally funded and automated case management tool. These data systems allow states to collect and maintain data for program management and informing their decision making [66]. States that implement

SACWIS must also report their data to federal databases, such as NCANDS (National Child Abuse and Neglect Data System) [86] and AFCARS (Adoption and Foster Care Analysis and Reporting System) [87], to allow for the continual curation of comprehensive national databases. These data systems became the foundation for actuarial risk assessment tools, which have been mandated into practice, even though controversy still remains as to whether these tools should override the judgment of case workers who are most knowledgeable about a particular child’s case [99, 98, 100, 111, 29].

Past survey papers have analyzed algorithms in CWS from a macro perspective, focusing on their reliability and validity with respect to consensus-based or clinical risk assessment models [99, 29]. Yet, they do not examine the mathematical or human-centered construction of these algorithms, that is, the techniques, the variable sets, or the outcomes predicted. This is especially important in CWS because each case of child neglect or abuse is contextually different and cannot be evaluated using the same set of significant predictors derived empirically [29]. To this end, we conducted a systematic literature review and identify the potential gaps in the literature with careful attention to the development of algorithms across time, as well as the methods and variable sets used.

METHODS

We describe our scoping criteria, systematic literature search and data analysis process.

Scoping Criteria: Defining Algorithms

To understand how "algorithms" are used in CWS, we first need to contextualize what we mean by algorithms. We conceptualized "algorithms" through the lens of *Street-level Algorithms*, a term recently coined by Alkhatib and Bernstein [6] in the HCI community. Street-level algorithms are algorithmically based systems that directly interact with and make on-the-ground decisions about human lives and welfare in a sociotechnical system [6]. From a more technical perspective, we use recent inclusive definitions [48, 68] for a whole suite of computational methods from statistical modeling (for e.g., generalized linear models) and machine learning. This allowed us to take a holistic viewpoint toward most forms of quantitative data analysis in CWS. Statistical modeling and machine learning are not mutually exclusive but we differentiate between them based on assumptions made about the data as specified by Breiman [23].

Systematic Literature Search

This study has been undertaken as a systematic literature review based on the guidelines proposed by Webster and Watson [115]. The unit of analysis for this literature review was peer-reviewed articles. We wanted to examine not just the algorithms currently being used in CWS but also newer solutions (algorithms) being proposed by researchers to better assess the current state of research. We used the following search terms to find papers at the intersection of CWS and algorithms – "child protective services," "child welfare," "foster care," "child and family services," "algorithm," "computation," "regression," "machine learning," "neural network," "data-driven," "actuarial," "computer program," "application". We used the following inclusion criteria for the articles:

Code	n	Breakdown
Peer reviewed	43	40 (social science); 3 (computer science)
Agency report	7	—
Theory	5	1 (implemented); 4 (proposed)
Psychometric scales	30	—
Actual system	27	15 (RAs); 11 (PLs); 1 (MT)
Hypothetical system	23	13 (RAs); 4 (PLs); 1 (MT); 5 (S-PL)
Model performance	35	—

RA: Risk Assessment model

PL: Placement Recommendation model

MT: Child-Foster parent Matching model

S-PL: Characteristics of successful placements

Table 1: Descriptive Characteristics of the Data Set

- The paper was peer-reviewed, published work or a systems (or policy) report produced by a government agency.
- The study (or report) engaged in a technical discussion about the computational methods, predictors and outcomes.

Articles that did not meet these two criteria were considered irrelevant for this study and were not included in our review. We conducted a comprehensive search to identify relevant research across multiple disciplines. We searched a diverse set of digital libraries which included the ACM Digital Library, IEEE Xplore, Routledge, Elsevier, and Springer. We chose these libraries to take into account research published in multi-disciplinary conferences and journals. We then cross-referenced the citations of each article to identify additional articles or government reports that met our inclusion criteria. We did not place any constraints on our search based on the time period in which the papers were published. We identified 50 relevant articles that met our inclusion criteria.

Data Analysis Approach

To analyze our data, we conducted a structured qualitative analysis to answer our over-arching research questions. We used a grounded thematic process [22] to generate codes based on the data as shown in Table 2. We define theory in two ways – the system discussed in the study was developed using a theoretical framework or the system was developed theoretically based upon factors considered significant in evidence-based social work. The first author coded all of the articles, and co-authors were consulted to form a consensus around codes early in the coding process and again during coding to resolve ambiguous codes. We also coded for descriptive characteristics of the article set as shown in Table 1.

RESULTS

In this section, we present our key findings from our review of the literature. We begin by first discussing the descriptive characteristics of our data set. Next, we organize and present the results by our three research questions, as shown in Table 2. Finally, we explore the relationship between the computational methods, predictors, and outcomes identified in our analyses.

Descriptive Characteristics of the Data Set

The majority of the papers (n=40 or 80%) were published in social science venues with 3 papers (6%) published in computer science conferences, [4, 7, 33] all in 2018. We also included 7 reports (14%) from non-profit organizations, including the Children’s Research Center [3]. One study discussed an algorithm which was theoretically constructed based

Research Question	Dimension	Codes	Count	%	Example	
RQ1 (Computational Method)	Inferential Statistics	Generalized Linear Models (GLM)	28	56%	[110]	
	Machine Learning	Discriminant Analysis/Statistical tests (DAS)	6	12%	[96]	
		Supervised Learning (SUP)	13	26%	[33]	
		Unsupervised Learning (UNSUP)	3	6%	[78]	
RQ2 (Predictor variables)	Demographics	Child Demographics (C-DEM)	20	40%	[7]	
	Systemic Factors	Biological parents Demographics (P-DEM)	10	20%	[65]	
		Characteristics of Agency (AGENCY)	2	4%	[80]	
		Characteristics of Caseworker (WORKER)	1	2%	[80]	
	Child Strengths	Child Strengths (CHI-S)	11	22%	[32]	
		Child Needs	Functioning (CHI-F)	15	30%	[77]
	Child Risks	Child Behavioral/Emotional Needs (CHI-BE)	26	52%	[70]	
		Suicide Risk (CHI-SR)	9	18%	[39]	
		Child Risk Behaviors (CHI-BR)	20	40%	[93]	
		Traumatic Experiences (CHI-T)	30	60%	[10]	
		Child Involvement in CWS (CHI-CWS)	9	18%	[110]	
		Bio-Parent Risk/Needs Foster Parents	Needs and Risky behavior (PAR-NS)	26	52%	[32]
			Characteristics (income, occupation) (FP-CHAR)	4	8%	[7]
	Preferences (FP-PREF)		2	4%	[80]	
Past performance (FP-PAST)	1		2%	[80]		
RQ3 (Outcome Variables)	Outcome	Capabilities (training/certifications) (FP-CAPS)	1	2%	[80]	
		Risk of a future maltreatment event (RISK)	28	56%	[65]	
		Placement recommendation for a child (PLACE)	15	30%	[32]	
		Matching children with foster parents (MATCH)	2	4%	[80]	
		Characteristics of a successful placement (S-PLACE)	5	10%	[107]	

Table 2: Structured Codebook: Dimensions are mapped onto their respective research questions

on child-welfare research literature and four studies proposed theoretically-driven solutions. 30 papers (60%) employed psychometric scales [53] to assess the strengths, needs and risks associated with foster children and/or the biological parents. 27 papers (54%) discussed an actual algorithmic system and 23 papers (46%) proposed a new algorithmic system. Model performance was reported by 35 papers (70%).

Computational Methods used to build Algorithms (RQ1)

In this section, we discuss the computational methods used to develop algorithms and organize them into the *Inferential Statistics* and *Machine Learning* dimensions.

Inferential Statistics approaches

Inferential statistics account for computational methods used in the majority of papers (68%), with 28 papers (56%) using a form of a generalized linear model (GLM). In Figure 1, we see a dramatic rise in the use GLMs after 1995, i.e., the post-SACWIS era. GLMs are being used to develop mostly two types of models; actuarial risk assessment and placement recommendation models. There was a general trend around the use of GLMs for developing risk assessment models [110, 29, 49]. We also identified two major concern surrounding GLMs: their atheoretical and reductive nature and performance with respect to outliers.

Social scientists use validated psychometric scales [73, 27] to quantify the level of risk. GLMs have been developed using these psychometric scales, such as the CANS Algorithm [32] that only uses the most statistically significant items from the scale. This reductive and atheoretical model development has received criticism [29, 99, 100]. Each case of child neglect/abuse is contextually different and factors that are significant for one case might be peripheral to another. Moreover, GLMs do not account for the contextual factors that influence caseworker decisions leading to variable omission bias [29].

Outliers can significantly impact the performance of a regression model [88]. Traditionally, regression models seek to omit outliers as a means of improving predictive power and still account for majority of the variance explained by significant variables [88]. However, for CWS, cases of severe abuse and neglect are the statistical outliers [21]. Regression models that are designed to predict the most moderate (average) outcomes tend to perform poorly on outliers [106]. In terms of CWS, poor performance on outliers raises several ethical and accountability concerns [40].

Four papers (8%) used discriminant analysis to differentiate between the characteristics of foster children served by different placement settings. Figure 1 illustrates that discriminant analysis was a popular technique during 1985-1990, however, with the advent of regression techniques it gradually faded away. These papers were some of the earliest attempts at introducing algorithms to aid decision-making in CWS. However, the data was limited and its quality questionable because of the lack of standardized data collection processes [103].

Machine Learning (ML) approaches

Machine Learning methods in CWS gained some momentum as early as 1986 with the introduction of *PLACECON*, a system designed to assist CWS with placement decisions [95]. However, with the increasing popularity of risk assessment models and limited funding available, resources were directed towards traditional regression models. Figure 1 shows a resurgence of ML methods starting 2015 and a growing interest within the computer science communities towards studying the research problems in CWS starting 2018 [4, 7, 33, 59]. Thirteen papers (26%) utilized ML methods in the form of decision trees, Bayesian networks or inference trees. Decision-tree learning has been popular as a means of organizing large amount of factual and empirical knowledge in the form of rules [102]. The CART (Classification and Regression Trees)

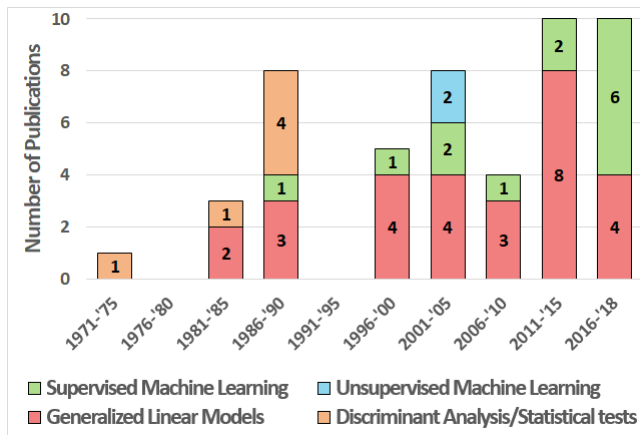


Figure 1: Methods used to build Algorithms (RQ1)

algorithm has been recently used to build a child-foster parent matching system [7]. It has also been used to identify the characteristics of the most troubled children in CWS [39] as well as to study trends in child abuse and neglect data [4]. However, with such a strong emphasis on risk assessment, Children's Research Center [3] used ML methods to develop the Structured Decision-Making (SDM) model.

SDM is a decision-making framework where a risk assessment tool is used in conjunction with clinical assessment [65]. SDM utilizes an array of ML tools such as decision, value and inference trees, and Bayesian networks [57] and has been adopted by CWS in several states [19]. However, several studies have also shown that SDM produces mixed results especially when accounting for race and ethnicity [42, 44, 64]. There is also an ongoing struggle between caseworker's theoretical assessments and the tool's empirical judgment [99, 100]. Three papers (6%) used unsupervised ML methods in the form of neural networks [77, 78] and natural language processing (NLP) [24]. Brindley et al. [24] propose a web platform that allows foster youth to create personalized goals and talk to a chat-bot that uses NLP to parse inputs and respond intelligently with recommendations about goals, finances, and housing. McDonald et al. [78] and Marshall et al. [77] propose the use of neural networks over regression techniques because their non-parametric approach performs better at modeling non-linear relationships and interactions.

One possible reason for the perpetual conflict between ML risk assessment tools and caseworkers' assessment might be at the core of Machine Learning itself and how it handles outliers. Statistical outliers in the case of child maltreatment are the most severe cases of child abuse and neglect [21]. Researchers [11] suggest that in the case of CWS, outliers are often more important for caseworkers and demands significant attention beyond the norm. Figure 1, depicts a significant dearth in the use of unsupervised learning methods with only two papers published in early 2000s [78, 77] and one paper published in 2018 [24]. Employing neural networks in social sciences comes with its own complexities because there needs to be transparency about the proposed decisions [33]. Vaithianathan et al. [110] explored several ML methods such as Naive Bayes and Random Forests for risk assessment and achieved higher

accuracies. However, they reverted to using a probit regression model because the outcomes were more explainable.

Predictors used in Algorithms (RQ2)

In this section, we examine the predictors that are being used in algorithms in CWS. Most algorithms are using over a hundred predictors so we systematically coded them and then grouped the emergent codes into seven dimensions (see Table 2).

Demographics and Systemic Factors

Child demographics were accounted for by 20 papers (40%) and biological parents demographics were accounted for by 10 papers (20%). Surprisingly, more than half the papers did not include child or parent demographics in their models even though racial and ethnic disparities in CWS have been recognized in social sciences [84, 91, 47]. Figure 2, illustrates that after 1990, there was a decline in the number of studies that used demographic variables in their algorithms. The *Systemic factors* dimension includes factors associated with CWS, such as, characteristics of the agency and caseworkers. Two papers (4%) use variables relating to characteristics of the agency, such as location and staffing vacancies and one paper (2%) accounted for the characteristics of the caseworker, such as, caseloads and the level of training. This is surprising because child-welfare literature acknowledges the impact caseworkers have on child outcomes [94, 30]. The caseworker is a child's primary contact between the biological parents, foster parents and CWS. They navigate through the system and find services for children and families. In fact, caseworker turnover is directly associated with placement instability [30]. Factors that lead to high caseworker turnover include low salary, high caseloads, administrative burdens, low levels of training and lack of supervisory support [30]. Systemic factors is one of the biggest reasons why children experience multiples placement moves in CWS [41]. This once again alludes to the atheoretical model construction that does not account for the salient factors well-established in evidence-based social work.

Foster-child related factors

Seven codes emerged out of the coding process and were grouped into three dimensions: child strengths, child needs, and child risks. 11 papers (22%) use variables that align with *Child Strengths*, such as, interpersonal skills, coping skills and level of optimism. Twenty-six papers (52%) took into account a child's emotional and behavioral needs and 15 papers (30%) recorded the child's day-to-day well-being and functioning, such as, their school attendance and behavior, personal hygiene and communication skills. We also coded for variables associated with risk factors that endanger child well-being. Suicide risk, risk behaviors, traumatic experiences, and child involvement with CWS were our four emergent codes that were grouped under the *Child Risks* dimension. 9 papers (18%) conducted a mental health screening to see if a child was suicidal or having suicidal thoughts. 20 papers (40%) accounted for risk behaviors such as self-harm, recklessness, social misbehavior, and 30 papers (60%) accounted for traumatic experiences such as neglect, physical/sexual abuse, history of family violence, and community violence. We noticed a trend here in that, almost all the risk assessment systems focused

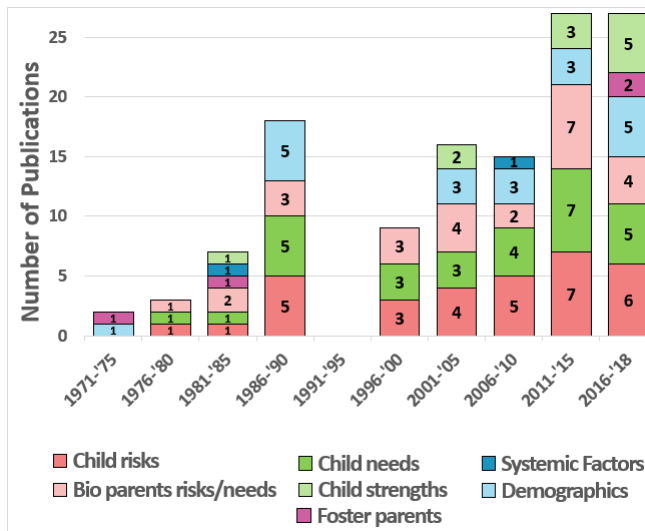


Figure 2: Predictors used in Algorithms (RQ2)

heavily on the *Child Risks* dimension, whereas, placement recommendation systems focused on the *Child Needs* dimension. Figure 2 depicts a rise in the number of studies that account for child strengths, child risk and child needs since 1995, that is, the post-SACWIS era. This alludes to the fact that these child characteristics are well-documented in SACWIS and are being used for modeling purposes.

All the studies we reviewed accounted for foster child related factors in terms of their needs and associated risks. However, only one study accounts for the child’s interactions with other people, such as siblings, relatives, and the system itself. Moore et al. [80] account for factors such as *Placement with a sibling*, *Proximity to child’s home/relatives*, and *Characteristics of the agency and caseworker*; factors well-studied in child-welfare literature [30]. Fluke et al. [49] found that placement decisions may be made as a result of interaction effects of non-case related factors such as characteristics of the agency and/or the caseworker. A study conducted in San Diego County found that 70% of the placement moves were a result of systemic or policy related factors [63].

Biological parents related factors

26 papers (52%) accounted for the biological parents’ risk behaviors and needs, such as, physical/mental health, substance abuse problems, residential stability and knowledge of child’s needs. We coded these variables into the *Bio-Parents Risks/Needs* dimension. Figure 2 shows that biological parent factors have been consistently used by several studies, however, we see a decline during 2005-2010. We also see a rise in the use of foster child related factors during the same time period. The introduction of CANS algorithm that focuses on the child’s level of need may be a plausible explanation for this trend. Different algorithms are using biological parent related variables differently. For example, risk assessment models quantify biological parents’ risky behavior so as to discern the risk of a future maltreatment. On the other hand, placement recommendation models are using this dimension to determine the level of trauma a child has experienced and

recommend a placement setting based on their level of need. Factors surrounding biological parents have been studied in great detail and accounted for by most algorithms.

Foster parents related factors

Four papers (8%) that we reviewed accounted for the characteristics of the foster parents, that is, their income level, occupation, demographics etc. Figure 2 shows that only 4 studies account for foster parent related factors with a significant gap between 1985 and 2016 where no study accounted for these factors. Two papers (8%) look at the preferences of foster parents and one paper (2%) accounts for the foster parents’ past performance and capabilities. Matching children with foster parents that are trained and prepared to meet their behavioral and emotional needs leads to increased stability for the children [30]. Matching children with foster parents that come from the same cultural background also leads to better outcomes because it leads to smoother transitions, lower stress and a feeling of security for the children [26]. These factors are well-studied in child-welfare literature [30, 92, 114], however, we see that very little research has been done from an algorithmic perspective. CWS has historically focused on ensuring safety and permanency rather than child well-being, that is, improving the quality of lives of foster children [16].

Target Outcomes of Algorithms (RQ3)

In this section, we examine the target outcomes of the algorithms used in CWS. Figure 3 depicts the trends in the target outcomes that algorithms have sought to model.

Risk Assessment

Predicting the risk of future maltreatment involves developing models using the empirical study of cases of child abuse/neglect [10]. The factors that show a strong association with abuse and/or neglect outcomes are selected to create an actuarial model which is then used to assess new cases of alleged abuse/neglect. Twenty-eight papers (56%) focused on predicting risk as their target outcome. Figure 3 illustrates that risk assessment has always received significantly more attention than any other outcome since the introduction of regression models in social sciences. The greatest criticism against these models is that they are not theoretically founded; these models are probabilistic in nature and not causal [10, 67, 99, 100]. Therefore, these models need to be empirically validated by follow-up studies to ensure their reliability. Direct comparison of any two actuarial models is a hard problem [10] and requires an in-depth understanding of the contexts in which the predictors were collected, measured and weighted in the models.

Studies conducted on risk assessment models show that these models are more accurate at predicting target events like child maltreatment than unaided judgment, however, they lack utility [99]. Seven papers (14%) discuss the Structured Decision-Making (SDM) model, a framework that integrates predictive and contextual assessments. CWS in several states have developed their own versions of SDM, however, there are significant enough differences to treat them independently as part of our review. Even though SDM is designed to assist caseworker decisions, studies have found that there are constant disagreements between the tool (empirically-driven) and caseworker




Figure 3: Target Outcomes of Algorithms (RQ3)

assessment (conceptually/theoretically-driven) to the point that caseworkers detest using these tools as they were intended [99]. However, caseworkers must continue to rely on these tools as a means of standardizing decisions in CWS, especially in cases of high uncertainty [100].

Placement Recommendations and Successful Placements

Models that focused on these two target outcomes were the precursors in the development of algorithms in CWS. Figure 3 depicts that these target outcomes were being studied during the time period of 1985-1990. However, no studies were published between 1991 and 2005 that focused on these target outcomes. A plausible explanation for this decline would be the increased focus on studying risk assessment during that period. 20 papers (40%) discussed recommendation systems for foster care placements.

The most prominent algorithm that determines the placement criteria based on a child's level of need is the CANS algorithm [32]. 6 papers (12%) discuss the CANS algorithm which is developed using the CANS psychometric scale [73]. CWS in a few states have developed their own versions of this algorithm, and therefore, were treated independently as part of our review. It makes a recommendation from six levels of care in the order of increasing severity – independent living, transitional living program, foster home, specialized foster care, group home, and residential treatment center. It is used in a hybrid approach in conjunction with a multi-disciplinary team which allows CWS to follow a standardized admission criteria for cases with lower levels of uncertainty [32]. This is a good initial approach to ensure child safety, however, it is a minimal approach and does not seek to improve the quality of a child's life.

Child-Foster Parent Matching

This approach seeks to match the specific needs of a child with the capabilities of foster parents. That is, placing children with foster parents who are trained and certified to manage their needs. It is a proactive approach towards improving the quality of lives of children and not just minimizing risk of maltreatment. Figure 3 shows that *Child-Foster parent matching* has only been implemented since 2015 (2 studies). This approach is different from the placement recommendation approach in that it addresses the specific needs of the child and preferences of the caregiver. For instance, matching with respect to child temperament, parent temperament, and parental

		Computational methods			
		Supervised Machine Learning	Unsupervised Machine Learning	Generalized Linear models	Discriminant analysis/statistical tests
Outcome Variables	RA	9	2	16	1
	PL	2	1	8	4
	MT	2	-	-	-
	S-PL	-	-	4	1

RA: Risk Assessment model
 PL: Placement Recommendation model
 MT: Child-Foster parent Matching model
 S-PL: Characteristics of successful placements

Table 3: Crosstabs between Methods and Outcomes

expectations leads to increased stability [92]. Placing children who have higher emotional needs with foster parents who prefer to be emotionally involved offers these children a better chance towards stability [113] than placing these children in a restrictive treatment setting. *Child-Foster parent matching* is well-studied in evidence-based social work and is known to improve stability and permanency outcomes [30, 92]. However, there is a dearth of information within CWS on how to guide this process [92]. This is a significant knowledge gap for both CWS and social scientists who seek to computationally model this approach. Moore et al. [80] recently validated a matching algorithm that was implemented by CWS in the state of Kansas for resulting in more stable placements.

Relationship between Methods, Predictors and Outcomes

Relationship between Algorithms (RQ1) and Outcomes (RQ3)

Table 3 depicts crosstabs between the computational methods used and the outcome from all the papers in our corpus. We saw that generalized linear models have mostly been used for developing risk assessment models (16 studies) followed by placement recommendation models (8 studies). Even with the emergence of newer machine learning methods, majority of the studies still continue to focus on risk assessment. 9 studies used supervised machine learning for risk assessment, 2 studies focused of placement recommendation and 2 studies focused on child-foster parent matching.

Relationship between Predictors (RQ2) and Outcomes (RQ3)

Table 4 depicts the crosstabs between predictors used by computational models and the outcome they seek to predict. First, we cross examine the risk assessment models with respect to the predictors that inform child characteristics. Majority of the models use a combination of predictors that assess *Child Behavioral/Emotional Needs* (7 studies), *Child Risk Behaviors* (7 studies), and *Traumatic Experiences* (15 studies). Several predictors coded under these dimensions (e.g., self-harm, recklessness, physical/sexual abuse) are assessed by a caseworker at an initial investigation and made available for predictive modeling. These predictors might already exist in the data if the family has previously come under the attention of CWS. This approach of aggregating the negative aspects of people's lives while ignoring the positive aspects has been criticized because of its deficit-based nature [25]. There is also an overlap between the *Traumatic Experiences* of a child and the *Bio-Parents Needs/Risk Behavior* because the same predictors (for

		Outcome Variables				Computational Methods			
		RA	PL	MT	S-PL	SUP	UNSUP	GLM	DAS
Predictors	Child demographics	8	6	2	4	7	1	7	5
	Bio-parents demographics	5	2	2	1	5	1	2	2
	Characteristics of Agency	-	-	-	2	-	-	2	-
	Characteristics of Caseworker	-	-	-	1	-	-	1	-
	Child Strengths	3	5	2	1	5	1	5	-
	Functioning	3	10	2	-	5	1	7	2
	Child Behavioral/Emotional Needs	7	13	2	4	7	-	13	6
	Suicide Risk	2	7	-	-	3	-	4	2
	Child Risk Behaviors	7	12	1	-	6	1	10	3
	Traumatic Experiences	15	10	2	3	9	1	16	4
	Child Involvement with CWS	2	6	1	-	3	-	3	3
	Bio-Parent Risk/Needs	15	8	-	3	7	1	16	3
	Foster parent characteristics	-	-	2	2	2	-	1	1
	Foster parent preferences	-	-	1	1	1	-	-	1
	Foster parent past performance	-	-	1	-	1	-	-	-
Foster parent capabilities	-	-	1	-	1	-	-	-	

RA: Risk Assessment model
PL: Placement Recommendation model
MT: Child-Foster parent Matching model
S-PL: Characteristics of successful placements
SUP: Supervised Machine Learning
UNSUP: Unsupervised Machine Learning
GLM: Generalized Linear models
DAS: Discriminant Analysis/Statistical tests

Table 4: Relationship between Computational Methods (RQ1), Predictors (RQ2) and Outcomes (RQ3)

e.g., history of physical/sexual abuse, medical trauma, parents' criminal activity) are used to conduct both *needs* assessment for a child and *risks* assessment for a parent.

Next, we cross examine the predictors used by placement recommendation models. These models are not employed at the onset of an investigation and are used by CWS when a child needs to be placed in a permanent foster care setting. These models are generally more equitable as compared to risk assessment models because they try to weigh-in the positive characteristics of a child, such as talents, interests, cultural identity, and school achievements to find an appropriate placement setting that meets their needs. Table 4 shows that placement recommendation models account for predictors around *Child Strengths* (5 studies), *Functioning* (10 studies), and *Child Behavioral/Emotional Needs* (13 studies) to weigh in the positive aspects and needs of a child and balance that with predictors around *Child Risk Behaviors* (12 studies) and *Traumatic Experiences* (10 studies) to find a suitable placement setting well equipped to meet their needs.

Relationship between Methods (RQ1) and Predictors (RQ3)

Table 4 depicts the crosstabs between predictors and computational methods. Most computational methods including both supervised machine learning and generalized linear models focused on *Child Behavioral/Emotional Needs*, *Child Risk Behaviors*, and *Traumatic Experiences* to assess the risk of a maltreatment event or the needs of a child. Some predictors that inform these three codes include traumatic events (e.g., physical/sexual abuse, medical trauma), child's conduct and anger management, and delinquent behavior. After an initial investigation is conducted by a caseworker and psychometric risk assessments completed, these predictors become available for modeling. However, quantifying risk from such a narrow set of predictors has been criticized because it fails to account for the wide range of risk factors that arise as a result of systemic issues in CWS itself [54].

DISCUSSION

Algorithms Need to be theoretical & context-aware (RQ1)

Overall, we found a lack of theoretically derived and validated algorithms that demonstrated that they took measures to integrate knowledge from the social sciences into their designs. Only one study [80] constructed their model based on the child-welfare literature. Four studies even discussed this lack of theory and proposed solutions in the form of cumulative risk models [74], causal models [98], and revised SDM models grounded in risk and resilience theory [99]. Yet, based on the published literature, such models have yet to be consistently implemented.

This finding is problematic because it shows that these algorithms ignore many factors that affect how decisions are made in CWS. For instance, the decisions are often constrained by current policies or the scarce resources [55]. Many current empirical models frustrate child-welfare workers because they do not account for such systemic factors. While some researchers have suggested [66] that empirical prediction is enough and that theory, context, or causal inferences are not always necessary in policy making when outcomes remain desirable, we argue that this is not a desirable stance to take in child-welfare contexts because there is significant debate on how and which types of data, models, and outcomes are to be used in predictive modeling (with or without theory). Empirical knowledge related to child-welfare practice is fragmented and social science theories must be used to fill the gaps [54].

Therefore, we recommend that human-centered theoretical approaches be used to incorporate factors arising from evidence-based social work [30] and understand the causal pathways that often dictate decision-making processes. Such algorithms that are informed by appropriate causal theory would also have a greater likelihood of utilization as compared to their a-theoretical counterparts [98]. Significant work has also found

disconnects between the functioning of algorithms and their social interpretations [13]. We see a similar phenomena in CWS where the caseworkers using Structured Decision Making (SDM) model must translate information from both forms of assessments (clinical and algorithmic) leading to uncertainty and unreliable decision making [100]. Therefore, algorithms that are meant to aid decision-making often become the source of frustration and force caseworkers to abandon their contextual judgments [99]. Human-centered theoretical approaches can help by placing the meaning-making process [89] at the center of the design process. It can help designers understand the theory of practice and uncover practitioners' sense-making processes (e.g., how they perceive quantified metrics [13]). Child-welfare workers are generally not trained in statistical thinking and make decisions based on experience, intuition, and individual heuristics [54]. Human-centered theoretical approaches can help us understand the mental models of child-welfare workers, inform feature selection (design), as well as interpret the results (evaluation).

Our results also indicate that several states adopted the SDM approach because it was supposed to integrate predictive and contextual assessments, however, it has fallen short of that goal [100, 99]. There are several factors at play in regards to any child-welfare case and it becomes critical to offer context to the case instead of focusing on a few broad factors without giving weight to important nuances [83]. For instance, understanding contextual knowledge with respect to an organization requires incorporating the organizational memory of the organization and its people [76, 5] which is inherently HCI research. Social workers are trained in writing detailed case notes by translating their context-specific experiences into text [35]. This unstructured, unanalyzed, textual data is added to SACWIS systems [87]. We hypothesize that valuable theoretical signals from these case notes can be considered within methodological approaches like topic modeling that can make good use of such unstructured data. Indeed, in recent years, HCI has developed a rich methodological tradition [14, 81, 31] of using signals from such unstructured data as predictors within algorithms to study complex, sociotechnical systems.

Going Beyond what is "Easily Quantifiable" (RQ2)

Our results suggest that majority of the algorithms used predictors around child and parent characteristics, such as their needs, strengths, and associated risks (see Table 2). The vast majority of these predictors that are used for predictive modeling are derived from information that is easily available and readily quantifiable. For example, child-welfare workers use psychometric scales [42] to assess child and parent associated risks and needs during an initial investigation which then become available for predictive modeling. Some of these predictors are found in almost every risk assessment model even though they have no predictive validity. For instance, severity of abuse is easily quantifiable and is found in several risk assessment models even though there is little to no indication that it is related to recurrence of abuse [28]. Moreover, several predictors (parenting skills, parent conflict etc.) have not been properly validated [54], and can lead to unreliable predictions [54]. Such issues led to the Illinois CWS (in 2017) to shut down their predictive analytics program [62]. In addition, none of

the predictors account for the temporality of risk assessment. After an allegation of abuse, the assumption of escalation is the baseline for risk assessment leading to inflated risk scores and excessive interventions from CWS [104].

Human-centered theoretical approaches can result in a rigorous feature selection process that relies on predictors that have been well-studied, understood, and validated in social sciences [13]. De Choudhury et al.'s [46] work in mental health is a good example, where the researchers validated constructs, focused on data biases and unobserved factors, as well as conducted sensitivity analysis. Moreover, it compels us to look towards sources of information that have been hereto hard to quantify. For instance, referring to our prior example around case notes, advances in natural language processing [112] now allows us to quantify and make holistic inferences about all the stakeholders involved in a child-welfare case. This can address persistent issues among cases which appear similar based on the empirical data but exhibit high variation in outcomes [54].

Furthermore, human-centered participatory design [13] allows for HCI researchers to actively engage with domain experts in child-welfare to understand how risk accumulates (and how to model it), as well as engage with other stakeholders to better understand the systemic factors around policies, laws and organizational culture [116]. Here, PD [82] can navigate the thorny, contextual differences between different legal and policy systems and needs/values of stakeholders. Lodato and DiSalvo [72] highlight the different forms and limitations of PD as well as how PD can be conducted within such institutional constraints. Advances in CWS data systems [50] can accommodate for the collection of several new predictors concerning child well-being and systemic factors. Here, the active consideration of needs and values of all stakeholders can help avoid the same reliability and validity pitfalls for the new predictors that exist for many of the current predictors.

Improve Lives and not just 'Minimize Risk'(RQ3)

One of the fundamental goals of CWS in the United States is to ensure positive outcomes for foster children [1], however, as our results confirm, majority of the efforts in computational modeling continues to be focused on risk assessment (see Table 2). Risk assessment models only seek to minimize the risk of future harm and not improve the quality of lives of foster children. The target outcome of "risk of maltreatment" is poorly defined [120]. Federal and State law dictate how child abuse and neglect are defined and the state definitions often vary and establish the grounds for intervention by CWS [1]. Algorithms are trained on cases of substantiation, that is, cases where CWS judged maltreatment to have occurred [54]. This judgment in itself is very subjective and depends on state laws, policies, and CWS intervention criteria which is often dictated by the level of funding and caseloads [30].

Human-centered approaches can help theoretically define not only the predictors but also the target outcomes with the help of stakeholders and domain experts to ensure these key ingredients needed for algorithm design are validated and reliable. Human-centered participatory design can also unravel concerns around the social interpretations of algorithmically-based systems. For instance, Brown et al [25] investigated the

community perspectives of risk assessment models in child-welfare. Child-welfare workers criticized these models because of their 'deficit-based' nature, that is, this approach only captures negative inputs to predict a negative outcome. There is growing concern that such an approach drives disproportionate negative caseworker perceptions that ultimately leads to negative actions [25]. Badillo-Urquiola et al. [9] and Pinter et al. [90] also recognized the problems with a deficit-based framing in that it creates a sense of moral panic and diverts attention away from positive outcomes. They suggested that researchers focus on "strength-based approaches" that focus on positive factors that help improve lives.

CWS should actively focus on approaches that disrupt the status quo [58] and seek to improve the lives of foster children, such as *Child-Foster Parent Matching* [30, 92, 113]. This requires an ongoing engagement with foster parents and foster children to understand their specific values and needs as well as their cultural and parental expectations. HCI can contribute here by drawing on its rich tradition of work in action research, participatory design, and value sensitive design to incorporate the values and needs of the stakeholders [52, 105, 8, 12, 18]. In addition, HCI researchers have developed methodological approaches that not only incorporate stakeholders into the design process but also the data analysis and interpretation processes [15, 119]. Moreover, advocating for foster children, a vulnerable and marginalized population, is inherently a social justice issue. HCI researchers have a long history of contending with social injustices and have developed theoretical and methodological approaches that seek "not so much to predict the future, but rather to imagine a radically better one [52]." Given the paucity of human-centered research into this domain and the richness of available social science literature [30, 92, 17, 26, 94, 41], this presents HCI researchers with a set of complex socio-technical challenges to study.

Recommendations for Future Research

Bridging AI and HCI Through Participatory Design

Our results indicate that there is a lack of theoretically-designed algorithms (see Table 1) which adds to the frustrations of child-welfare workers who are being pressured into using these algorithms as a means of standardizing decisions [99]. This situation is further exacerbated by a lack of PD leading to algorithmic systems that offer low utility [100]. Only one study in our corpus engaged with child-welfare workers to understand their concerns and needs [25]. PD [82] allows for the active inclusion of people most affected by a system. Engaging child-welfare workers in the design as well as evaluation processes ensures that their needs are met and that the system integrates well with child-welfare practice. Child-welfare workers who use algorithms on a daily basis strongly stress the need to be able to explain these models to each other and to policymakers [83]. Not only does this depend on which computational methods are used to construct an algorithm, how they are deployed but also how outcomes are defined and measured. This offers research pathways for HCI researchers who have increasingly started devoting attention to explaining outcomes and predictions [71, 79]. Moreover, it is imperative that researchers engage with the stakeholders because there are both, ethical and legal ramifications of using certain types

of data. For instance, legal requirements might not allow a juvenile's criminal record or history of physical and/or sexual abuse to be used for modeling [101].

Algorithmic Decision Making via Speculative Design

We found that 56% of studies took a deficit-based approach to mitigate risks even though child-welfare literature has discussed the significance of equitable outcomes (e.g., child-foster parent matching). Recent studies based in newer technologies still continue to focus on risk assessment and uncritically reproduce the status quo. Designing against the status quo means setting our goals beyond risk assessment, and moving more ambitiously toward design that challenges underlying problems [58]. Human-centered speculative design [13] can allow stakeholders to shift their focus away from algorithms and be truly innovative in how they imagine problems and their underlying causes without being constrained by what might be technologically feasible. This is especially important for algorithm design where the boundaries of possibility change every day [13]. For instance, child-foster parent matching is well-documented in child-welfare literature for almost two decades but it has only recently been explored in an algorithm [80] because of advances in decision-tree learning. Similarly, algorithmic advances also create novel avenues for studying the interactions and decision pathways resulting from different policies, practices and programs [116].

LIMITATIONS AND FUTURE WORK

We conducted a comprehensive and systematic literature review which was limited to the US-based child welfare system. We may have also missed algorithms used within CWS that are not publicly available for review. For instance, reports by non-profit organizations or state governments may have been distributed internally. Therefore, we plan to work directly with CWS agencies and conduct user interviews about the systems and algorithms being used within CWS to identify any other algorithms that have been implemented. To move towards using a human-centered approach to build new, evidence-based, and theoretically-driven algorithms, we plan to work with stakeholders in CWS to understand how different policies, practices, and programs create different decision pathways for child placements and services offered to families.

CONCLUSION

In conclusion, we recommend that the HCI community partners with CWS to do the following: **1)** A renewed focus on theoretically-designed algorithms with the active engagement of stakeholders through the design and evaluation phases, **2)** Develop algorithms for practice that incorporate a more comprehensive set of predictors well-studied in child-welfare literature, as well as predictors hard to quantify thus far, and **3)** Focus on equitable outcomes founded in evidence-based child-welfare research that improve the quality of lives of foster children instead of merely mitigating future risks.

ACKNOWLEDGMENTS

This research is funded in part by the Facebook Computational Social Science Methodology Research Award, William T. Grant Foundation (187941 and 190017), and the Northwestern Mutual Data Science Institute.

REFERENCES

- [1] 2013. *How the Child Welfare System Works*. Technical Report. Children's Bureau: Child Welfare Information Gateway.
- [2] 2018. *Using Data To Help Protect Children and Families Act*. 115th Congress, Senate of the United States.
- [3] 2019. Child Welfare Goals, Legislation, and Monitoring. (2019). Retrieved April 2, 2019 from <https://www.nccdglobal.org/what-we-do/children-s-research-center>
- [4] Abdurazzag A Aburas, Mohammad Hassan, Hilary Lin, and Shreshtha Batshu. 2018. Child Maltreatment Forecast Using Bigdata Intelligent Approaches. In *2018 Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*. IEEE, 302–308.
- [5] Mark S Ackerman and Christine Halverson. 2004. Organizational memory as objects, processes, and trajectories: An examination of organizational memory in use. *Computer Supported Cooperative Work (CSCW)* 13, 2 (2004), 155–189.
- [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*. ACM, 530.
- [7] Rachmadita Andreswari, Irfan Darmawan, and Warih Puspitasari. 2018. A Preliminary Study on Detection System for Assessing Children and Foster Parents Suitability. In *2018 6th International Conference on Information and Communication Technology (ICoICT)*. IEEE, 376–379.
- [8] Mariam Asad and Christopher A Le Dantec. 2015. Illegitimate civic participation: supporting community activists on the ground. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 1694–1703.
- [9] Karla Badillo-Urquiola, Xinru Page, and Pamela Wisniewski. 2019. Risk vs. Restriction: The Digital Divide between Providing a Sense of Normalcy and Keeping Foster Teens Safe Online. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM.
- [10] Christopher Baird, Dennis Wagner, Theresa Healy, and Kristen Johnson. 1999. Risk assessment in child protective services: Consensus and actuarial model reliability. *Child Welfare* 78, 6 (1999), 723.
- [11] Zuriana Abu Bakar, Rosmayati Mohamad, Akbar Ahmad, and Mustafa Mat Deris. 2006. A comparative study for outlier detection techniques in data mining. In *2006 IEEE conference on cybernetics and intelligent systems*. IEEE, 1–6.
- [12] Shaowen Bardzell. 2014. Utopias of participation: design, criticality, and emancipation. In *Proceedings of the 13th Participatory Design Conference: Short Papers, Industry Cases, Workshop Descriptions, Doctoral Consortium papers, and Keynote abstracts-Volume 2*. ACM, 189–190.
- [13] Eric PS Baumer. 2017. Toward human-centered algorithm design. *Big Data & Society* 4, 2 (2017), 2053951717718854.
- [14] Eric PS Baumer, David Mimno, Shion Guha, Emily Quan, and Geri K Gay. 2017a. Comparing grounded theory and topic modeling: Extreme divergence or unlikely convergence? *Journal of the Association for Information Science and Technology* 68, 6 (2017), 1397–1410.
- [15] Eric PS Baumer, Xiaotong Xu, Christine Chu, Shion Guha, and Geri K Gay. 2017b. When Subjects Interpret the Data: Social Media Non-use as a Case for Adapting the Delphi Method to CSCW. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 1527–1543.
- [16] Lawrence M Berger, Sarah K Bruch, Elizabeth I Johnson, Sigrid James, and David Rubin. 2009. Estimating the "impact" of out-of-home placement on child well-being: Approaching the problem of selection bias. *Child development* 80, 6 (2009), 1856–1876.
- [17] 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.
- [18] Alan Borning and Michael Muller. 2012. Next steps for value sensitive design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1125–1134.
- [19] Emily Adlin Bosk. 2018. What counts? quantification, worker judgment, and divergence in child welfare decision making. *Human Service Organizations: Management, Leadership & Governance* 42, 2 (2018), 205–224.
- [20] Engin Bozdag. 2013. Bias in algorithmic filtering and personalization. *Ethics and information technology* 15, 3 (2013), 209–227.
- [21] Ann E Brand and Paul M Brinich. 1999. Behavior problems and mental health contacts in adopted, foster, and nonadopted children. *The journal of child psychology and psychiatry and allied disciplines* 40, 8 (1999), 1221–1229.
- [22] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [23] Leo Breiman and others. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science* 16, 3 (2001), 199–231.
- [24] Meredith Brindley, James P Heyes, and Darrell Booker. 2018. Can Machine Learning Create an Advocate for Foster Youth? *Journal of Technology in Human Services* 36, 1 (2018), 31–36.

- [25] 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.
- [26] Jason D Brown, Natalie George, Jennifer Sintzel, and David St Arnault. 2009. Benefits of cultural matching in foster care. *Children and Youth Services Review* 31, 9 (2009), 1019–1024.
- [27] Michael J Camasso and Radha Jagannathan. 1995. Prediction accuracy of the Washington and Illinois risk assessment instruments: An application of receiver operating characteristic curve analysis. *Social Work Research* 19, 3 (1995), 174–183.
- [28] Michael J Camasso and Radha Jagannathan. 2000. Modeling the reliability and predictive validity of risk assessment in child protective services. *Children and Youth Services Review* 22, 11-12 (2000), 873–896.
- [29] Michael J Camasso and Radha Jagannathan. 2013. Decision making in child protective services: A risky business? *Risk analysis* 33, 9 (2013), 1636–1649.
- [30] Sarah Carnochan, Megan Moore, and Michael J Austin. 2013. Achieving placement stability. *Journal of Evidence-Based Social Work* 10, 3 (2013), 235–253.
- [31] Stevie Chancellor, Zhiyuan Lin, Erica L Goodman, Stephanie Zerwas, and Munmun De Choudhury. 2016. Quantifying and predicting mental illness severity in online pro-eating disorder communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 1171–1184.
- [32] Ka Ho Brian Chor, Gary M McClelland, Dana A Weiner, Neil Jordan, and John S Lyons. 2012. Predicting outcomes of children in residential treatment: A comparison of a decision support algorithm and a multidisciplinary team decision model. *Children and Youth Services Review* 34, 12 (2012), 2345–2352.
- [33] 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*. 134–148.
- [34] Christopher E Church and Amanda J Fairchild. 2017. In Search of a Silver Bullet: Child Welfare’s Embrace of Predictive Analytics. *Juvenile and Family Court Journal* 68, 1 (2017), 67–81.
- [35] James Clifford and George E Marcus. 1986. *Writing culture: The poetics and politics of ethnography*. Univ of California Press.
- [36] Patricia Cohen, Stephen G West, and Leona S Aiken. 2014. *Applied multiple regression/correlation analysis for the behavioral sciences*. Psychology Press.
- [37] US Congress. 2008. Fostering connections to success and increasing adoptions act of 2008. (2008).
- [38] Lindsay D Cooper. 2005. Implications of media scrutiny for a child protection agency. *J. Soc. & Soc. Welfare* 32 (2005), 107.
- [39] Katharan D Cordell, Lonnie R Snowden, and Laura Hosier. 2016. Patterns and priorities of service need identified through the Child and Adolescent Needs and Strengths (CANS) assessment. *Children and Youth Services Review* 60 (2016), 129–135.
- [40] Michael Corrigan. 2019. Building A Comprehensive Child Welfare Information System. (Jan 2019). <https://chronicleofsocialchange.org/child-welfare-2/building-comprehensive-child-welfare-information-system/33426>
- [41] Theodore P Cross, EUN Koh, Nancy Rolock, and Jennifer Eblen-Manning. 2013. Why do children experience multiple placement changes in foster care? Content analysis on reasons for instability. *Journal of Public Child Welfare* 7, 1 (2013), 39–58.
- [42] Amy D’andrade, Michael J Austin, and Amy Benton. 2008. Risk and safety assessment in child welfare: Instrument comparisons. *Journal of evidence-based social work* 5, 1-2 (2008), 31–56.
- [43] David Danks and Alex John London. 2017. Algorithmic Bias in Autonomous Systems.. In *IJCAI*. 4691–4697.
- [44] EW Dankert and Kristen Johnson. 2013. Risk assessment validation: A prospective study. *Los Angeles: California Department of Social Services, Children and Family Services Division* (2013).
- [45] Elizabeth Davoren. 1975. Foster placement of abused children. *Children today* 4, 3 (1975), 41.
- [46] Munmun De Choudhury and Emre Kiciman. 2018. Integrating Artificial and Human Intelligence in Complex, Sensitive Problem Domains: Experiences from Mental Health. *AI Magazine* 39, 3 (2018), 69–80.
- [47] Alan J Dettlaff, Stephanie L Rivaux, Donald J Baumann, John D Fluke, Joan R Rycraft, and Joyce James. 2011. Disentangling substantiation: The influence of race, income, and risk on the substantiation decision in child welfare. *Children and Youth Services Review* 33, 9 (2011), 1630–1637.
- [48] David Donoho. 2017. 50 years of data science. *Journal of Computational and Graphical Statistics* 26, 4 (2017), 745–766.
- [49] John D Fluke, Martin Chabot, Barbara Fallon, Bruce MacLaurin, and Cindy Blackstock. 2010. Placement decisions and disparities among aboriginal groups: An application of the decision making ecology through multi-level analysis. *Child Abuse & Neglect* 34, 1 (2010), 57–69.

- [50] Administration for Children and Families. *Comprehensive Child Welfare Information System* (106 ed.). Vol. 81. Federal Register: The Daily Journal of the United States.
- [51] Patrick J Fowler, Katherine E Marcal, Saras Chung, Derek S Brown, Melissa Jonson-Reid, and Peter S Hovmand. 2019. Scaling up housing services within the child welfare system: policy insights from simulation modeling. *Child maltreatment* (2019), 1077559519846431.
- [52] Sarah Fox, Jill Dimond, Lilly Irani, Tad Hirsch, Michael Muller, and Shaowen Bardzell. 2017. Social Justice and Design: Power and oppression in collaborative systems. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 117–122.
- [53] R Michael Furr. 2017. *Psychometrics: an introduction*. Sage Publications.
- [54] Eileen Gambrill and Aron Shlonsky. 2000. Risk assessment in context. (2000).
- [55] Eileen Gambrill and Aron Shlonsky. 2001. The need for comprehensive risk management systems in child welfare. *Children and Youth Services Review* 23, 1 (2001), 79–107.
- [56] Stuart Gray, Kirsten Cater, Chloe Meineck, Rachel Hahn, Debbie Watson, and Tom Metcalfe. 2019. trove: A digitally enhanced memory box for looked after and adopted children. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*. ACM, 458–463.
- [57] Robin Gregory, Lee Failing, Michael Harstone, Graham Long, Tim McDaniels, and Dan Ohlson. 2012. *Structured decision making: a practical guide to environmental management choices*. John Wiley & Sons.
- [58] Ellie Harmon, Matthias Korn, Ann Light, and Amy Volda. 2016. Designing against the status quo. In *Proceedings of the 2016 ACM Conference Companion Publication on Designing Interactive Systems*. ACM, 65–68.
- [59] 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*. ACM, 108.
- [60] Kori Inkpen, Stevie Chancellor, Munmun De Choudhury, Michael Veale, and Eric PS Baumer. 2019. Where is the Human?: Bridging the Gap Between AI and HCI. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, W09.
- [61] Luma Institute. 2012. *Innovating for people: Handbook of human-centered design methods*. LUMA Institute, LLC.
- [62] David Jackson and Gary Marx. 2017. Data mining program designed to predict child abuse proves unreliable, DCFS says. (dec 2017). <https://www.chicagotribune.com/investigations/ct-dcfs-eckerd-met-20171206-story.html>
- [63] Sigrid James. 2004. Why do foster care placements disrupt? An investigation of reasons for placement change in foster care. *Social service review* 78, 4 (2004), 601–627.
- [64] K Johnson and D Wagner. 2003. California Structured Decision Making. Risk Assessment Revalidation: A Prospective Study. Children’s Research Center. SAMSHA’s National Registry of Evidence-based Programs and Practices (NREPP). (2003).
- [65] Will Johnson. 2004. Effectiveness of California’s child welfare structured decision making (SDM) model: a prospective study of the validity of the California Family Risk Assessment. *Madison (Wisconsin, USA): Children’s Research Center* (2004).
- [66] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. 2015. Prediction policy problems. *American Economic Review* 105, 5 (2015), 491–95.
- [67] Eun Koh and Mark F Testa. 2008. Propensity score matching of children in kinship and nonkinship foster care: Do permanency outcomes differ? *Social Work Research* 32, 2 (2008), 105–116.
- [68] J Nathan Kutz. 2013. *Data-driven modeling & scientific computation: methods for complex systems & big data*. Oxford University Press.
- [69] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Management Science* (2019).
- [70] Mark D Lardner. 2015. Are restrictiveness of care decisions based on youth level of need? A multilevel model analysis of placement levels using the child and adolescent needs and strengths assessment. *Residential Treatment for Children & Youth* 32, 3 (2015), 195–207.
- [71] Min Kyung Lee and Su Baykal. 2017. Algorithmic mediation in group decisions: Fairness perceptions of algorithmically mediated vs. discussion-based social division. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 1035–1048.
- [72] 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*. ACM, 5.

- [73] 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).
- [74] Michael J MacKenzie, Jonathan B Kotch, and Li-Ching Lee. 2011. Toward a cumulative ecological risk model for the etiology of child maltreatment. *Children and youth services review* 33, 9 (2011), 1638–1647.
- [75] Diana MacLean, Sonal Gupta, Anna Lembke, Christopher Manning, and Jeffrey Heer. 2015. Forum77: An analysis of an online health forum dedicated to addiction recovery. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 1511–1526.
- [76] James G March. 1991. Exploration and exploitation in organizational learning. *Organization science* 2, 1 (1991), 71–87.
- [77] David B Marshall and Diana J English. 2000. Neural network modeling of risk assessment in child protective services. *Psychological Methods* 5, 1 (2000), 102.
- [78] Thomas P McDonald, John Poertner, and Gardenia Harris. 2002. Predicting placement in foster care: A comparison of logistic regression and neural network analysis. *Journal of social service research* 28, 2 (2002), 1–20.
- [79] Hannah Miller Hillberg, Zachary Levonian, Daniel Kluger, Loren Terveen, and Brent Hecht. 2018. What I See is What You Don't Get: The Effects of (Not) Seeing Emoji Rendering Differences across Platforms. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 124.
- [80] Terry D Moore, Thomas P McDonald, and Kari Cronbaugh-Auld. 2016. Assessing risk of placement instability to aid foster care placement decision making. *Journal of Public Child Welfare* 10, 2 (2016), 117–131.
- [81] Michael Muller, Shion Guha, Eric PS Baumer, David Mimno, and N Sadat Shami. 2016. Machine learning and grounded theory method: Convergence, divergence, and combination. In *Proceedings of the 19th International Conference on Supporting Group Work*. ACM, 3–8.
- [82] Michael J Muller. 2009. Participatory design: the third space in HCI. In *Human-computer interaction*. CRC press, 181–202.
- [83] Judge Michael Nash. 2017. Examination of using Structured Decision Making and Predictive Analytics in assessing Safety and Risk in Child Welfare. *County of Los Angeles Office of Child Protection* (May 2017).
- [84] Barbara Needell, M Alan Brookhart, and Seon Lee. 2003. Black children and foster care placement in California. *Children and Youth Services Review* 25, 5-6 (2003), 393–408.
- [85] Kathleen G Noonan, Charles F Sabel, and William H Simon. 2009. Legal accountability in the service-based welfare state: Lessons from child welfare reform. *Law & Social Inquiry* 34, 3 (2009), 523–568.
- [86] US Department of Health and Human Services. 2017. Child Maltreatment 2017. *Children's Bureau* (Ed.) (2017).
- [87] US Department of Health, Human Services, and others. 2017. The AFCARS report: Preliminary FY 2016 estimates as of Oct 2017. *Children's Bureau* (Ed.) 21 (2017), 6.
- [88] R Lyman Ott and Micheal T Longnecker. 2015. *An introduction to statistical methods and data analysis*. Nelson Education.
- [89] Sarah Pink, Kerstin Leder Mackley, Val Mitchell, Marcus Hanratty, Carolina Escobar-Tello, Tracy Bhamra, and Roxana Morosanu. 2013. Applying the lens of sensory ethnography to sustainable HCI. *ACM Transactions on Computer-Human Interaction (TOCHI)* 20, 4 (2013), 25.
- [90] Anthony T Pinter, Pamela J Wisniewski, Heng Xu, Mary Beth Rosson, and Jack M Carroll. 2017. Adolescent online safety: Moving beyond formative evaluations to designing solutions for the future. In *Proceedings of the 2017 Conference on Interaction Design and Children*. ACM, 352–357.
- [91] Emily Putnam-Hornstein, Barbara Needell, Bryn King, and Michelle Johnson-Motoyama. 2013. Racial and ethnic disparities: A population-based examination of risk factors for involvement with child protective services. *Child Abuse & Neglect* 37, 1 (2013), 33–46.
- [92] Richard E Redding, Carrie Fried, and Preston A Britner. 2000. Predictors of placement outcomes in treatment foster care: Implications for foster parent selection and service delivery. *Journal of child and family studies* 9, 4 (2000), 425–447.
- [93] 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).
- [94] Joseph P Ryan, Philip Garnier, Michael Zyphur, and Fuhua Zhai. 2006. Investigating the effects of caseworker characteristics in child welfare. *Children and Youth Services Review* 28, 9 (2006), 993–1006.
- [95] John R Schuerman and Lynn Harold Vogel. 1986. Computer support of placement planning: the use of expert systems in child welfare. *Child welfare* 65, 6 (1986), 531–543.
- [96] A James Schwab and Susan S Wilson. 1989. The continuum of care system: Decision support for practitioners. *Computers in Human Services* 4, 1-2 (1989), 123–140.

- [97] A James Schwab Jr, Michael E Bruce, and Ruth G McRoy. 1984. Matching children with placements. *Children and youth services review* 6, 2 (1984), 125–133.
- [98] Craig Schwalbe. 2004. Re-visioning risk assessment for human service decision making. *Children and Youth Services Review* 26, 6 (2004), 561–576.
- [99] Craig S Schwalbe. 2008. Strengthening the integration of actuarial risk assessment with clinical judgment in an evidence based practice framework. *Children and Youth Services Review* 30, 12 (2008), 1458–1464.
- [100] Aron Shlonsky and Dennis Wagner. 2005. The next step: Integrating actuarial risk assessment and clinical judgment into an evidence-based practice framework in CPS case management. *Children and youth services review* 27, 4 (2005), 409–427.
- [101] Ravi Shroff. 2017. Predictive Analytics for City Agencies: Lessons from Children’s Services. *Big data* 5, 3 (2017), 189–196.
- [102] Fiore Sicolu. 1989a. Computer-aided decisions in human services: Expert systems and multivariate models. *Computers in Human Behavior* 5, 1 (1989), 47–60.
- [103] Fiore Sicolu. 1989b. Prediction and decision making in child welfare. *Computers in Human Services* 5, 3-4 (1989), 43–56.
- [104] Douglas G Simpson, Peter B Imrey, Olga Geling, and Susan Butkus. 2000. Statistical estimation of child abuse rates from administrative databases. *Children and Youth Services Review* 22, 11-12 (2000), 951–971.
- [105] Susan Leigh Star and Anselm Strauss. 1999. Layers of silence, arenas of voice: The ecology of visible and invisible work. *Computer supported cooperative work (CSCW)* 8, 1-2 (1999), 9–30.
- [106] James P Stevens. 1984. Outliers and influential data points in regression analysis. *Psychological Bulletin* 95, 2 (1984), 334.
- [107] Norman M Stone and Susan F Stone. 1983. The prediction of successful foster placement. *Social Casework* 64, 1 (1983), 11–17.
- [108] Angelika Strohmayer, Rob Comber, and Madeline Balaam. 2015. Exploring learning ecologies among people experiencing homelessness. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2275–2284.
- [109] Rebecca Tushnet. 2018. The Difference Engine: Perpetuating Poverty through Algorithms. *Jotwell: J. Things We Like* (2018), 1.
- [110] Rhema Vaithianathan, Emily Putnam-Hornstein, Nan Jiang, Parma Nand, and Tim Maloney. 2017. Developing predictive models to support child maltreatment hotline screening decisions: Allegheny County methodology and implementation. *Center for Social data Analytics* (2017).
- [111] D Wagner, K Johnson, and W Johnson. 1998. Using actuarial risk assessment to target service interventions in pilot California Counties. *13th National Roundtable on CPS Risk Assessment, San Francisco, CA* (1998).
- [112] Hanna M Wallach. 2006. Topic modeling: beyond bag-of-words. In *Proceedings of the 23rd international conference on Machine learning*. ACM, 977–984.
- [113] James A Walsh and Roberta A Walsh. 1990. Studies of the maintenance of subsidized foster placements in the Casey Family Program. *Child Welfare* 69, 2 (1990), 99–114.
- [114] Daniel Webster, Richard P Barth, and Barbara Needell. 2000. Placement stability for children in out-of-home care: A longitudinal analysis. *CHILD WELFARE-NEW YORK-* 79, 5 (2000), 614–632.
- [115] Jane Webster and Richard T Watson. 2002. Analyzing the past to prepare for the future: Writing a literature review. *MIS quarterly* (2002), xiii–xxiii.
- [116] James K Whittaker. 2017. *The child welfare challenge: Policy, practice, and research*. Routledge.
- [117] Jacob O Wobbrock and Julie A Kientz. 2016. Research contributions in human-computer interaction. *interactions* 23, 3 (2016), 38–44.
- [118] Jill Palzkill Woelfer and David G Hendry. 2010. Homeless young people’s experiences with information systems: life and work in a community technology center. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1291–1300.
- [119] Susan P Wyche, Paul M Aoki, and Rebecca E Grinter. 2008. Re-placing faith: reconsidering the secular-religious use divide in the United States and Kenya. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 11–20.
- [120] Susan J Zuravin. 1999. Child neglect. *Neglected children: Research, practice, and policy* (1999), 24–46.