Affected Community Perspectives on Algorithmic Decision-Making in Child Welfare Services

Alexandra Chouldechova
Estella Loomis McCandless Assistant Professor of Statistics and Public Policy

IOF on Algorithmic Fairness and Social Justice
September 25, 2020

Carnegie Mellon University
Heinz College
INFORMATION SYSTEMS • PUBLIC POLICY • MANAGEMENT

It is estimated that **37.4%** of children experience a CPS investigation by age 18

Source: Kim et al. (2017)
Lifetime Prevalence of Investigating Child Maltreatment Among US Children

2018 US Child Maltreatment Statistics

Source: *Child Maltreatment 2018* based on 2018 NCANDS data

4.3 million referrals

7.8 million children

3.5 million children received investigation or alternative response

The New York Times Magazine

GET UPDATES

FEATURE

Can an Algorithm Tell When Kids Are in Danger?

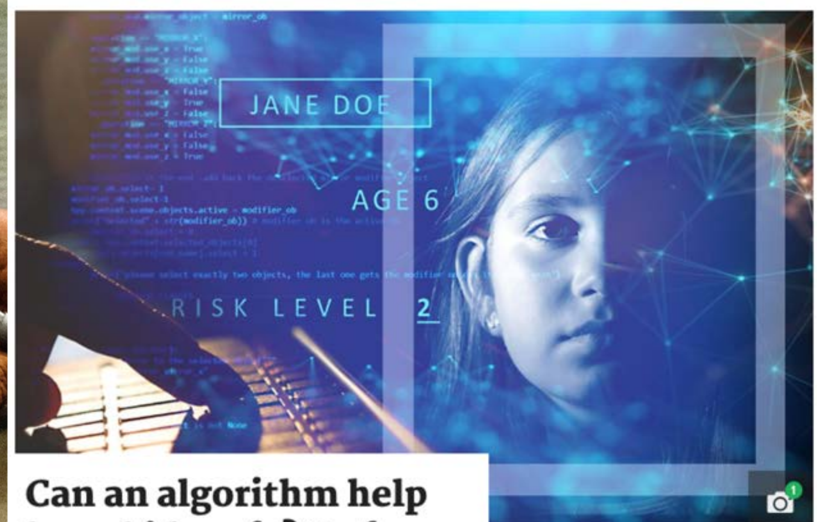

Daniel Marsula/Post-Gazette illustration

Can an algorithm help keep kids safe? So far, Allegheny County's screening tool is improving accuracy

WIRED SUBSCRIBE

VIRGINIA EUBANKS BUSINESS 01.15.2018 08:00 AM

A Child Abuse Prediction Model Fails Poor Families

Why Pittsburgh's predictive analytics misdiagnoses child maltreatment and prescribes the wrong solutions

"We definitely oversample the poor," says Erin Dalton, Director of Allegheny County's Office of Data Analysis, Research and Evaluation. "All of the data systems we have are biased. We still think this data can be helpful in protecting kids." NKBIMAGES/GETTY IMAGES

Pittsburgh's child welfare agency goes full Orwell

Starting in 2020, Allegheny County, Pa. will attempt to, in effect, stamp EVERY child born in the county with a "scarlet number" risk score that could haunt the child and her or his family for life.

Need systems that are *reliable* and *trusted*

- Many widely used approaches to building and evaluating risk assessment models fail to be reliable, often for underappreciated reasons
- Even if a tool is reliable, there many be significant obstacles in getting the public to trust it and to trust how it is used
- Issues go deeper than compliance with laws and regulation

  "There is enormous opportunity for positive social impact from the rise of algorithms and machine learning. But this requires a licence to operate from the public, based on trustworthiness. […] We have seen before in the case of genetic modification what can happen when science is pushing forward but loses public trust—this set the take-up of the science back significantly."

Shah, *Algorithmic Accountability* (2018)


— Schuler & Namioka, 1993
‘In Participatory Design the people destined to use the system play a critical role in designing it.’

THREE CENTRAL QUESTIONS

i

How do people who are most likely to be subject to or affected by algorithmic decision-making feel about the deployment of such systems?

ii

What are the primary sources of community discomfort surrounding the development and deployment of such tools?

iii

What can researchers and designers do in the development and deployment stages to raise comfort levels among affected communities?

Organizational Justice

Procedural justice

- Perceived fairness of the process that produces the decisions/outcomes

Distributive justice

- Perceived fairness of the decisions/outcomes

Informational justice

- Sufficiency and completeness of information provided to explain and justify decisions/outcomes

Interpersonal justice

- Extent to which people are treated with dignity and respect by those making and communicating decisions

Colquitt (2001) *On the dimensionality of organizational justice: A construct validation of a measure.*

WORKSHOP PARTICIPANTS

USA study

5 x workshops
Single USA county
83 participants

Group make-up

FAMILIES

FRONTLINE PROVIDERS

SPECIALISTS

A FICTIONAL SCENARIO

Nicole's Story

1a You are doing well, looking after your three-year old daughter Nicole on your own with support from your parents and extended family. One night, after a long day at work you come home, feed and play with Nicole and put her to bed. You then take a hot bath, put some headphones on and relax for a half hour – but when you get out of the bath and check on Nicole you find she is not in her bed.

Nicole is found by a neighbour, barefoot, cold and lost, trying to find her Nana's house (where you have walked with her many times). The worried neighbour settles Nicole, who is very upset, and returns her to your care. On returning home the neighbour calls Child Welfare Services and **the call centre worker decides to recommend an investigation**.

The next day you pick up Nicole from daycare and are very embarrassed to hear that a **social worker has visited the daycare centre to make inquiries about you and your family situation. The social worker also visits your home and informs you that you are under investigation, which involves checking on any previous welfare or criminal records, and that the outcome could take some weeks.**

Nicole's Story

1b Imagine the same scenario, but where Child Welfare Services have introduced a computer tool to help workers make more informed decisions about whether to recommend an investigation or not.

When you neighbour calls Child Welfare Services, a computer tool runs a statistical analysis of historical data which includes family case history, public health data, and criminal justice records, and also the history of other calls and how they turned out. The result is a risk score which predicts the level of child risk, from 1 (low risk) to 10 (severe risk).

Since you have no history with Child Welfare Services and no criminal record, **Nicole's risk score is assessed as being low and does not result in an investigation.**

Nicole's Story

1c It is now nine years later and Nicole is 12. She is going through some difficult times – she has fallen in with a bad crowd and you ground her because of bad behaviour.

Late that evening she jumps out her bedroom window and runs away to a friend's place. Police are concerned to see her walking alone late at night and bring her home.

You explain what has happened, but they are required to inform Child Welfare Services.

The next day you receive a visit from a social worker who says that the computer tool they are using scored Nicole as being at high risk. This is because your new partner has a previous domestic violence investigation on his record and together with your data record in the Child Welfare Services system there is possibility of a pattern.

Nicole's Story

1d **i).** Four years later, Nicole is now 16 and doing well. She is in a stable relationship with her boyfriend and becomes pregnant. She has a lovely boy, Anthony. You are very supportive, and Anthony is a happy, healthy baby. While still at the hospital Nicole receives a visit from a nurse explaining that she has been **identified by a statistical tool as needing support**, and offering her home visits and access to other services over Anthony's first year.

ii). During the course of the conversation the nurse tells Nicole that **statistics show that 1 in 5 mothers identified as needing support — like she has been — end up having their child removed by Child Welfare Services.** Home visits tend to reduce the risk of placement.

iii). Would your response be any different if she tells Nicole that **statistics show that 4 in 5 mothers identified as needing support — like she has been — end up having their child placed by Child Welfare Services?**

SCENARIO STEPS

‘Nicole’s Story

Scenario 1A
Reactive + human decision-making
Presents a child welfare decision being made by an intake worker in reaction to a call from a member of the public. There is no mention of personal data or a computer tool being involved in making the decision.

Scenario 1B
Reactive + algorithm-assisted decision-making
Presents a child welfare action that is reactive, and a decision that is made by a human assisted by an algorithm.

Scenario 1C
Reactive + algorithmic decision-making + using family’s child welfare data
Presents a child welfare action that is reactive, and a decision that is made by an algorithm including associative data (child welfare investigation).

Scenario 1D (i), (ii), (iii)
Proactive + algorithmic decision-making + using family’s child welfare data
Presents a child welfare action that is proactive, and a decision to offer services that is made by an algorithm — with three different ways of communicating this decision to the family.

Scenario 1E
Reactive + algorithmic decision-making + using administrative data beyond child welfare
Presents a child welfare action that is reactive, and a decision that is made by an algorithm including associative data (criminal justice data).

Scenario 1F
Proactive + algorithmic decision-making + using administrative & community data
Presents a child welfare action that is proactive, and a decision that is made by an algorithm including nonassociative data (criminal records, neighborhood, age).

SYSTEM LEVEL CONCERNS

1.

System-level concerns were the most common reasons given for low comfort in algorithm-assisted and algorithmic decision-making.

SYSTEM-LEVEL CONCERNS

Low trust and low benefit

FAMILIES

"It's been me versus the system."

"They would look at me more because I had previous experience than because they wanted to help me with my daughter."

FRONTLINE PROVIDERS

"It seems like a deficit model — let's weigh up all the dirty things in your life, nothing good though."

SPECIALISTS

" 'Investigation' says that you've been judged already."

DATA CONTEXT CONCERNS

2.

All groups raised concerns about potential bias on the part of case workers involved in the decision process, as well as bias present in the data or the algorithm.

SYSTEM-LEVEL CONCERNS

Concerns about bias

FAMILIES

"The system here in America just lets us down, especially if you are Black."

SPECIALISTS

"How honest are we allowed to be? Most of our systems were not made for people of colour, or by people of colour, or have People of Color in them."

FRONTLINE PROVIDERS

"My neighbour might be shooting up heroin and their six year old is out in the street, but they have private insurance so their records aren't part of this system. The computer tool is only capturing people who have to use public health so there's a bias to poorer people in the system."

EXPLAINABILIY AND TRANSPARENCY

3.

Participants questioned whether a statistical model could adequately account for all relevant decision elements, and emphasized the need for a human in the loop approach.

SCENARIO SPECIFIC CONCERNS

Data context and interpretation

FAMILIES

"A computer cannot understand context. My son has autism — how does the data account for this?"

FRONTLINE PROVIDERS

"The score should be a flag rather than a definitive 'go'. It needs to be approached with curiosity: Where are you at? What are you facing? What are your needs? Would you benefit from home visits, more community? Help put it back together. If Child Welfare was just a score we wouldn't be sitting here."

SPECIALISTS

"Use data without removing human decision-making."

ACCOUNTABILITY

4.

Participants wanted more information on how the algorithm weighs different factors, and the ability to dispute the score.

EXPLAINABILIY AND TRANSPARENCY

5.

Even potentially beneficial decisions resulted in discomfort due to concerns about how and whether risk information was communicated to families and case workers.

1d

i). Four years later, Nicole is now 16 and doing well. She is in a stable relationship with her boyfriend and becomes pregnant. She has a lovely boy, Anthony. You are very supportive, and Anthony is a happy, healthy baby. While still at the hospital Nicole receives a visit from a nurse explaining that she has been **identified by a statistical tool as needing support**, and offering her home visits and access to other services over Anthony's first year.

ii). During the course of the conversation the nurse tells Nicole that **statistics show that 1 in 5 mothers identified as needing support — like she has been — end up having their child placed by Child Welfare.** Home visits tend to reduce the risk of placement.

iii). Would your response be any different if the nurse tells Nicole that **statistics show that 4 out of 5 mothers identified as needing support — like she has been — end up having their child placed by Child Welfare?**

✓ Participants approved of young mothers being offered supportive services such as home visits.

Saying that mothers like her have a 1 in 5 chance of having their child placed (removed) was perceived as a threat.

Saying there's a 4 in 5 chance was perceived as a bigger threat

Participants were wholly opposed to any mention of a statistical tool in this context.

RESEARCH TEAM

Predictive Risk Models

Rhema Vaithianathan
Professor of Health Economics
Centre for Social Data Analytics
Auckland University of Technology, NZ

Alexandra Chouldechova
Assistant Professor, Statistics & Public Policy
Carnegie-Mellon University, USA

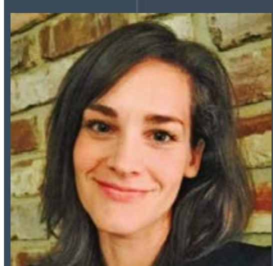

Emily Putnam-Hornstein
Associate Professor, Social Work
University of Southern California, USA

Participatory Design

Anna Brown
Associate Professor
Toi Āria: Design for Public Good
Massey University, NZ

Andrew Tobin
Associate
Toi Āria: Design for Public Good
Massey University, NZ

We are grateful to the SIGKDD Community Impact Program for their financial support