What Works for Children's Social Care

Machine Learning in Children's' Social Care

Does it work and how should we be doing it?

30 September 2020

Dr Michael Sanders

Chief Executive

What Works for Children's Social Care

Housekeeping

- We’re recording the webinar
- We have an hour and a half
- We’ll have time for questions at the end
- Please type your questions into the chat box
- We’ll do our best to get to as many as possible, and may group similar questions
- If your question is for a particular panellist, please include this!

Speakers

Michael Sanders, Chief Executive, What Works for Children's Social Care

Dr Lisa Holmes, Director of the Rees Centre, Department of Education, University of Oxford

Vicky Clayton, Data Science Manager, What Works for Children's Social Care

What Works *for* Children's Social Care

About WWCSC

IMPROVING EVIDENCE FOR BETTER OUTCOMES

How?

- Pulling together what we already know
- Supporting the good work that is already happening
- Commissioning new research
- Giving practitioners, young people and families a platform to share their experience
- Improving the accessibility and relevance of the evidence

Why?

To ensure the best possible outcomes for children, young people and families

What Works *for* Children's Social Care

Background

What is machine learning?

- We used a subset of machine learning called predictive analytics
- We look at historical cases to understand what information we have about a case is associated with a particular outcome e.g. if the referral is from the school the child's case may be more likely (on average) to escalate
- We then apply those associations to new cases to predict the outcome 6-12 months down the line

For example:
Does the child / young person's case progress to the child being subject to a Child Protection Plan (CPP) or being looked after (CLA) within 6-12 months of a contact?

Why were we doing this research?

- There are a small but growing number of local authorities in England (and other jurisdictions) investing in predictive models to assist social workers assess risk

- There is disagreement amongst the sector over in which circumstances it is appropriate and desirable to apply predictive techniques in children's social care (CSC)

- WWCSC identified the use of predictive models in CSC as an area which would benefit from more evidence on the technical aspects as well as the ethical perspective.

Dr Lisa Holmes

Director of the Rees Centre

Department of Education

University of Oxford

What Works for Children's Social Care

The Alan Turing Institute

Rees Centre

Department of Education

University of Oxford

Ethics review of machine learning in children's social care

@whatworksCSC

Lisa Holmes

Project background

“Is it ethical to use machine learning approaches in children’s social care systems and if so, how and under what circumstances?”

- Should we be doing this?
- Can we do this right?
- What is to be done?

The Alan Turing Institute

Rees Centre

Methodology

Study comprised three components:

- Desk based review of the literature
- Roundtable discussion
- Family engagement workshop

The Alan Turing Institute

Rees Centre

The bigger picture

Local authorities already using predictive analytics and ML

- Internal analyses
- Commissioned external companies

Transparency of approaches, underpinning principles and approach to ethics

The Alan Turing Institute

Rees Centre

Should we be doing this?

The Alan Turing Institute

Rees Centre

What does ethical ML in children's social care require?

We focused on 19 sets of ethical standards for social work and AI, and we found that they had a lot in common. We identified:

- *Ethical values* to set the direction of responsible ML innovation in CSC
- *Practical principles* to establish the justifiability of design and deployment
- *Professional virtues* for professional integrity

We used these to propose a **Commitment to care, collaboration and understanding**

The Alan Turing Institute

Rees Centre

Challenges to ethical ML innovation in CSC

Public management in the context of austerity

- Concept of value
- Relationship-based practice

System, organisation, and participant readiness

- Infrastructures
- Resources
- Cultural receptivity

Social inequality and cycles of poverty and discrimination

- Patterns of inequality could be perpetuated

The Alan Turing Institute

Rees Centre

Can we do this right?

The Alan Turing Institute

Rees Centre

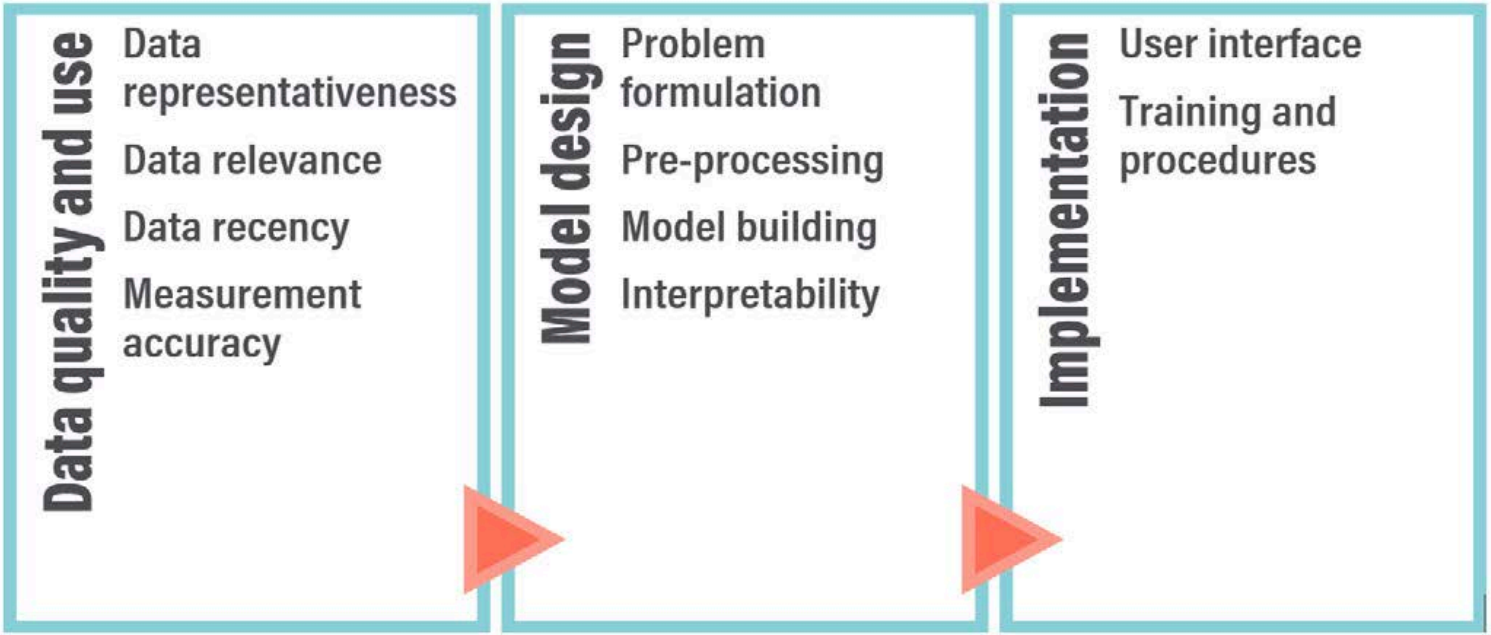

The Alan Turing Institute

Rees Centre

A few key questions

- Do we currently have data of sufficiently high quality to provide a holistic picture of families?
- How should we reliably measure success of children's social care services?
- Can we ensure the justifiability of designer and user decisions at every step of the design, development, and use of ML?
- Do we have the resource necessary for responsible development, use, and maintenance of AI now and in the future?

The Alan Turing Institute

Rees Centre

What is to be done?

The Alan Turing Institute

Rees Centre

Mandate the responsible design and use of ML models in CSC at the national level

Families stressed the need for nationally mandated public standards to guide the ethical design and deployment of ML in CSC.

The Alan Turing Institute

Rees Centre

Connect practitioners and data scientists across local authorities to improve ML innovation and to advance shared insights in applied data science through openness and communication

Siloing of research and innovation practices stifle progress, notion of ‘black box’ analysis

The Alan Turing Institute

Rees Centre

Inclusive and consent-based practices for designing, procuring, and implementing ML models

Families emphasised the importance of consent and deliberative involvement at all points across the ML lifecycle.

Local authorities should listen.

The Alan Turing Institute

Rees Centre

Fund, initiate, and undertake active research programmes in system, organisation, and participant readiness

There is a need for empirical research to identify and better understand the barriers and enablers to effective integration of responsible ML innovation in CSC.

The Alan Turing Institute

Rees Centre

Understand the use of data in CSC better so that recognition of its potential benefits and limitations can more effectively guide ML innovation practices

We must better understand how data has and is being used in the field and to what degrees of success or failure. This will help manage wider expectations about what is possible in the use of ML in CSC.

The Alan Turing Institute

Rees Centre

Use data insights to describe, diagnose and analyse the root causes of the need for CSC, experiment to address them:

Direct data science to understand, diagnose, and address the root causes behind deeper social-structural problems and dynamics that generate the expanding needs for CSC services.

Use knowledge to better design policies and interventions and rigorously evaluate them.

The Alan Turing Institute

Rees Centre

Focus on individual- and family-advancing outcomes, strengths-based approaches, and community-guided prospect modelling

Research is needed to explore how positive outcomes can be integrated into data analytics in CSC.

This will require inclusive, multi-stakeholder, and interdisciplinary approach to objective setting, model planning, and implementation.

The Alan Turing Institute

Rees Centre

Improve data quality and understanding through professional development and training

Data collection, analysis, and use of ML models should be built into social care and social work training.

Focus should be on accurate and impartial data collection and use of ML in line with ethical values.

The Alan Turing Institute

Rees Centre

Further information

Project report and summary: https://whatworks-csc.org.uk/research-report/ethics-review-of-machine-learning-in-childrens-social-care/

lisa.holmes@education.ox.ac.uk

@LisaHolmes_Rees

The Alan Turing Institute

Rees Centre

Vicky Clayton

Data Science Manager

What Works for Children's Social Care

What did we want to find out in the technical research?

- How well do the models predict outcomes in children's social care? In how many cases is the concern overstated ("a false alarm")? In how many cases is the concern understated ("a child missed")?
- Does including text improve how well the models predict outcomes?
- Are the models more error prone for children / young people with particular characteristics?

What did we do?

We worked with 4 local authorities across North West, South West, West Midlands and South East. They had Good or Outstanding Ofsted ratings.

We worked with the intelligence teams at each local authority to agree which outcomes to predict.

We spent 12 weeks with each local authority preparing the data and conducting the analysis.

What data did we use?

- The local authorities provided us with 3-7 years of data extracted from their case management systems.
  - Structured data: usually in the form of the Annex A report which they prepare for Ofsted.
  - Text data: accompanying notes from early help contacts, referrals, assessments, initial and review child protection conference reports and strategy discussion (depending on the outcome being predicted)
- We only used data which would be available to the social worker at the time of the decision.
- This gave us datasets of c.700-c.24,000 cases.

How did we test the models?

The models weren't tested in practice or used to make any decisions about a child or young person.

To test each outcome, we split the historical data into 'training' and 'test' data, training the model on the training data and then testing whether the patterns learned generalise to the test data. This simulates how well the models would perform if used in practice to predict the outcome for new cases.

What did we find overall?

WE FIND THAT

On average, if the model identifies a child is at risk, it is wrong six out of ten times. The model misses four out of every five children at risk.

None of the models' performances exceeded our pre specified threshold for 'success' .

Adding information extracted from reports and assessments does not improve model performance.

Our analysis of whether the models were biased was unfortunately inconclusive.

There is a low level of acceptance of the use of these techniques in children's social care amongst social workers.

How well do the models predict outcomes in children's social care?

0 = worst possible model
1 = best possible model
Random = 0.02-0.05

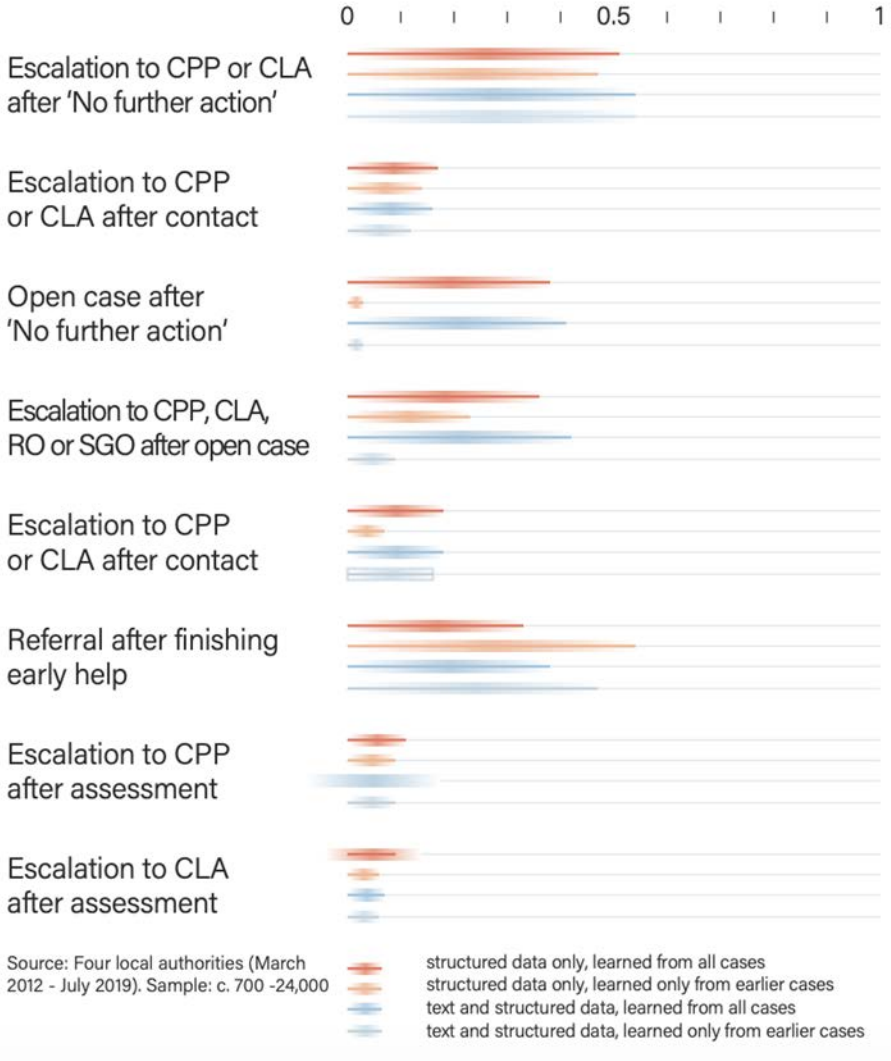

After successfully finishing early help, is the child / young person referred to statutory children's services within 12 months?

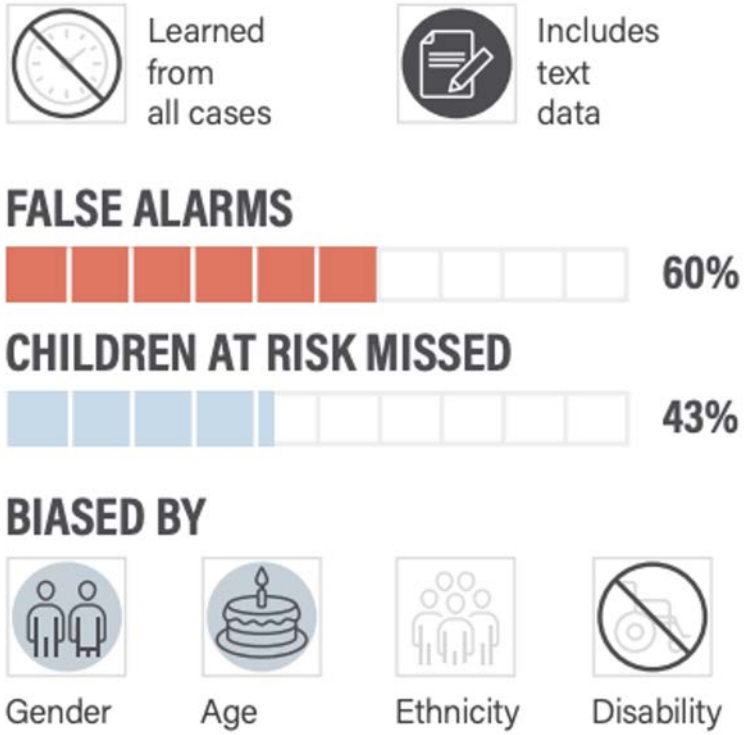

Does including text improve how well the models predict outcomes?

0 = worst possible model
1 = best possible model
Random = 0.02-0.05

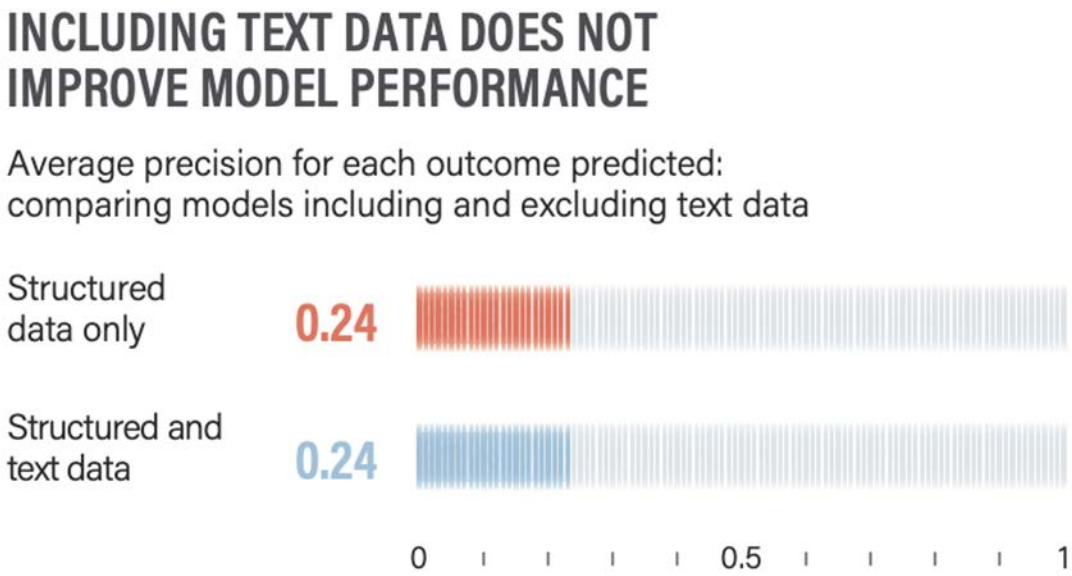

Are the models more error prone for children / young people with particular characteristics?

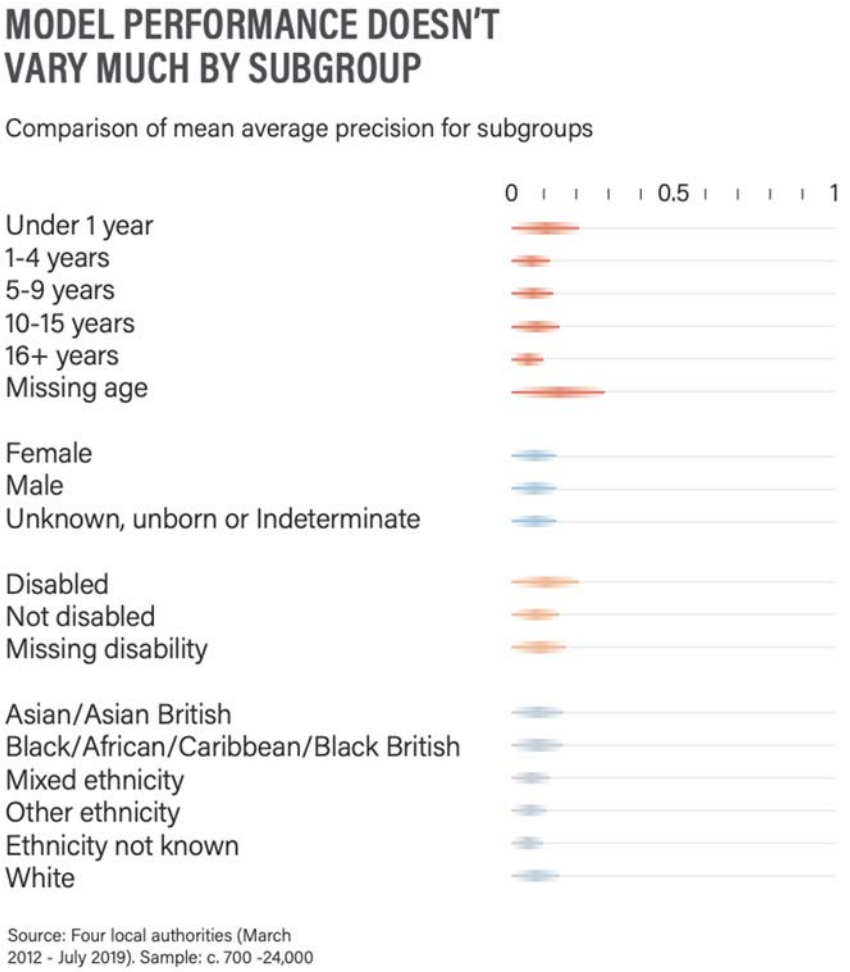

0 = worst possible model
1 = best possible model
Random = 0.02-0.05

What did social workers think?

Source: WWCSC social worker poll, March 2020. Sample: 129

What does this mean for machine learning in CSC?

- Other outcomes may be easier to predict
- It seems that more observations would help but these are not easily available
- Adding richer data may not help as the model is already trying to over-generalise the nuances of particular cases
- We would like to see more transparency on how well machine learning models correctly identify children / young people at risk and children / young people not at risk


QUESTIONS
AND
DISCUSSION

What Works for Children's Social Care

Thank you

@whatworksCSC