SOCIAL FINANCE
Social Finance Limited Tintagel House, 92 Albert Embankment, London SE1 7TY
T +44 (0)20 3963 5151 W socialfinance.org.uk

Social Finance response to CDEI Call for Evidence - Review into bias in algorithmic decision-making

This document outlines our response to the Centre for Data Ethics and Innovation's call for evidence review on bias in algorithmic decision making.

Background

Over the last five years, Social Finance – a non-profit – has supported 35+ local authorities in the use of data and technology to improve social services and outcomes. Our extensive experience of working with local government puts us in a unique position to comment on the use of algorithms in social care.

Many of our local government partners are currently experimenting with algorithmic decision making or are heading in that direction. To date, we have been engaged to explore the use of algorithms in issue areas such as homelessness, healthcare, 'Not in Education, Employment or Training,' (NEET), children's social care, school readiness and school exclusion.

While we recognize the potential for algorithms to play a positive role in improving public services, we believe that this needs to be balanced with a series of ethical and practical considerations. We have six key recommendations:

- **Recommendation 1: Local authorities and their partners need guidance and tools to help them assess data quality**
- **Recommendation 2: Local authorities and their partners need data that captures people, not just processes.**
- **Recommendation 3: Local authorities and their partners need to align algorithmic decision making to user needs**
- **Recommendation 4: Local authorities and their partners need to co-design algorithms with end users to build trust and credibility**
- **Recommendation 5: Local authorities and their partners need additional resources to build capacity for the use of algorithms and avoid introducing bias**
- **Recommendation 6: Local authorities and their partners need an ethics framework to explore the issues and opportunities of algorithmic decision making**

We expand on these recommendations below, drawing on examples of our work.

Recommendation 1: Local authorities and their partners need guidance and tools to help them assess data quality. Data captured in local authorities and partner organisations are often low quality and incomplete due to factors such as data entry errors, data modelling inconsistences and changes to practice models. These issues need to be accounted for to prevent erroneous conclusions, but they often go unnoticed. The data visualisation below highlights this point.

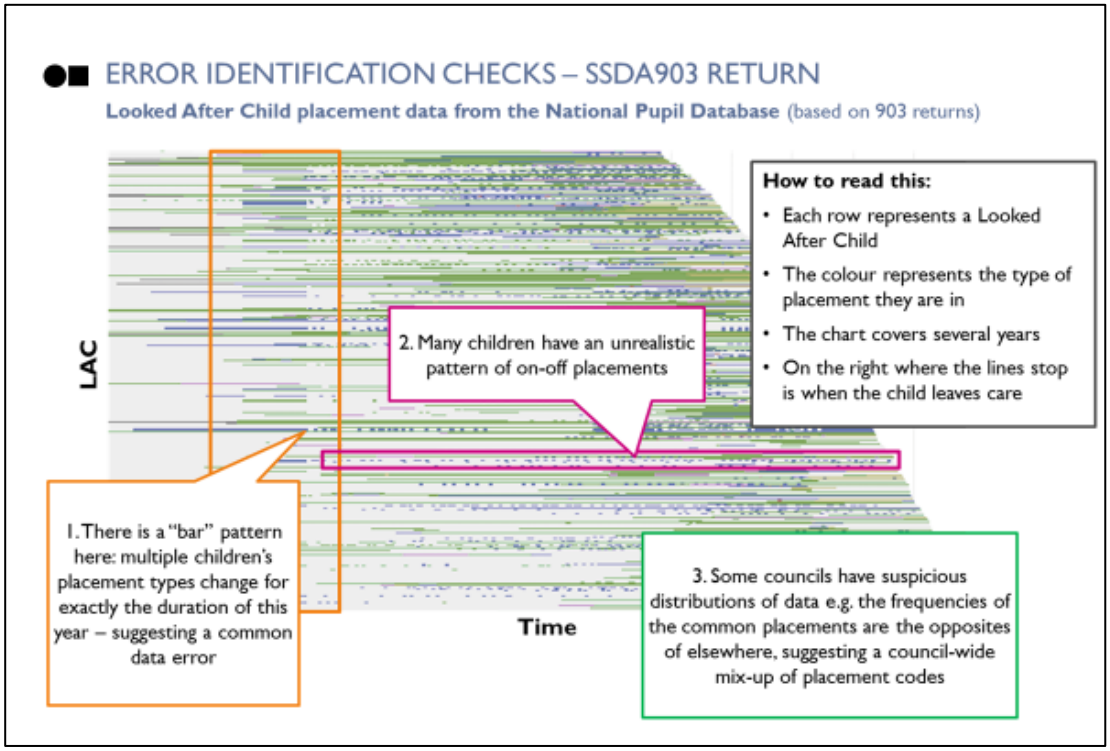

The data shows a series of non-random data entry errors previously unidentified by the council despite already being cleaned as part of the statutory return process. These errors include placement duration corresponding directly to year end dates and wrong placement codes. These errors could produce spurious insight – especially if the algorithmic approach focuses on finding patterns and correlations in the data. Highlighting and assessing potential errors requires a combination of domain/sector experience and technical expertise.

Identifying and accounting for data quality issues: Social Finance co-developed a commissioning tool with 15 local authorities to help them build the case for investment in intensive edge of care services to prevent children from unnecessarily going into care.

To build trust in the data and analysis, we used a series of algorithms and tools and user testing to identify bias and errors in the data – including quantitative error detection, record linkage and de-duplication, rule-based detection and pattern enforcement and transformation. We then conducted dip-sampling of errors to assess with practitioners whether they were the result of recording issues or accurate representations of families' experiences.

The approach gave the Greater London Authority the confidence to launch a new edge of care service to help 450 families access the support they need.

Recommendation 2: Local authorities and their partners need data that captures people, not just processes. Case management systems emerged from a need to monitor processes and comply with statutory requirements and reporting. Most data captured is on inputs and outputs and limited to the specific area where that service is being delivered. This creates risks that: (1) little to no information about a child or a family is collected unless there has been some interaction with the local authority and (2) information that is available is only about their experience of a process.

This is a key concern in the context of algorithmic decision making because data collected in this manner may be incomplete in a non-random way if some groups are more likely to access public services than others. Furthermore, this data can only be used to optimise processes rather than improve outcomes.

Designing data around people: Social Finance is currently helping Leeds City Council create a common dataset for their preventative "Early Help" services, which is focused on people, not process. We have conducted 60+ user interviews to date with front line practitioners to build a picture of what they really need to understand on the families they support. We're also engaging with families to ensure they're comfortable with how the information is being captured and used.

These insights are informing a case management system redesign around collected data focused on families' needs, strengths and the outcomes they achieve and less on the process they go through. This builds a more holistic view of people, better enabling both front line support and decision making.

Recommendation 3: Local authorities and their partners need to align algorithmic decision making to user needs. Algorithms must solve a clear user need if they're going to improve decision making and services. However, vendors and local authorities often develop algorithms without clearly understanding the user or the service. We've seen products fail time and time again because they require major changes to service delivery to add value or provide insight in a usable format.

User research is key. Local authorities and their partners need to spend more time with users to understand the needs, pain points and opportunities around algorithmic decision making and clearly outline how the insights developed from algorithms can be incorporated into services. This is the only way to unlock the potential for data to improve decision making, services and, most importantly, outcomes for children and families.

Define a clear user need: Social Finance worked with a North West of England local authority who wanted an algorithmic risk stratification system to identify children with poor school readiness, so they could support them earlier. We began with a discovery project to understand the user needs around school readiness. The team conducted user research with 30+ practitioners and families – creating detailed personas and an end-to-end service blueprint of Early Years and Early Help Support. We found that the most pressing user need came from the Health Visitor service. They wanted better access to information that is

already collected by the council about the families they visit to better assess risk and make appropriate onward referrals where necessary.

The discovery recommended that the council focus on addressing this user need through better information sharing rather than invest in an algorithmic risk stratification system, which would be more complex and present more ethical issues but not necessarily add more value.

Recommendation 4: Local authorities and their partners need to co-design algorithms with end users to build trust and credibility. We believe that algorithms should be co-designed with local authority decision makers so that everyone understands and trusts the results and outputs. Algorithms cannot be built away from users. Users must be part of the process to provide constant feedback in creating a product that empowers them and improves their decision making. This helps to ensure that algorithms and their outputs align with practice models and support, rather than replace, professional judgement.

Working collaboratively to build trust and understanding: To help two large metropolitan councils rethink their approach to tackling NEETness, we developed a common longitudinal data model for understanding young people's experiences with universal and statutory services.

To enable everyone, from frontline practitioners to service directors, to feed into the analysis, we developed an interactive decision-tree tool. Despite the team using a supervised machine learning algorithm, the visualisations presented the findings in a clear, intuitive format.

Together, we identified the five main drivers of young people becoming NEET, such as experience in care. Manchester used the insight to recommission their NEET reduction service around these drivers, providing a service to support the most vulnerable.

The inclusive process made our analytical approach and results transparent, trusted and easy to understand for both Newcastle and Manchester and the broader public – so much so that the Department of Education subsequently replicated our analysis across the country.

Recommendation 5: Local authorities and their partners need additional resources to build capacity for the use of algorithms and to avoid introducing bias. Many local authorities lack the resources and capacity needed to build, support and use algorithms in social care decision making. There is a wide range of aspects that local authorities need to manage carefully, in addition to making sure that the model is fair and free from biased outputs. This includes complying with privacy and data processing legislation, setting out the accountability structure, constantly training and monitoring the algorithm, and managing the public narrative.

This means that local authorities need to be thoughtful and intentional around the resources needed to manage projects that use algorithms. Who is responsible for changing and modifying the algorithm? Does the technical expertise currently exist within the local council? How does the algorithm change the process of delivering social care? These are some key questions that local authorities need to consider when deciding if it is sensible to proceed.

Capacity Building and Tech decisions: Social Finance worked closely with Newcastle to build a data model and matching algorithm to join data from universal and statutory services to understand the impact of preventative support.

We conducted a skills audit at the start of the project to understand how to develop the analysis and matching algorithm using technologies the council could use, maintain and update. We also provided training to build confidence in the approach and ensure its sustainability.

This was vital to ensure that the council could both continue to run the analysis, and to update it based on new learnings and changes to the practice model.

Recommendation 6: Local authorities and their partners need an ethics framework to explore the issues and opportunities of algorithmic decision making. There are 343 authorities in England. Each take their own approach to determining what the ethical and appropriate use of data and algorithms is. They have different understandings of risk, privacy, information governance, and impact. They also work within different political environments, budget constraints, and practice models. For example, we've seen neighbouring councils take opposite opinions on the use of data to target support at children who are unlikely to be school ready.

This is a problem. Despite living in the same country, the use of data and algorithmic decision making across social care varies by postcode, with children and families living a street apart experiencing drastically different services and having their data used in drastically different ways. Local authorities and their partners need guidance on what ethical issues to consider and how best to explore the ethical implications of how they're using data.

Data ethics and user research: Social Finance is testing and trialling different approaches to bringing data ethics into the forefront of algorithmic design and use – from a data ethics canvas to user research to algorithmic auditing. User research with those most likely to be affected by the use of data and algorithmic decision making – including front-line practitioners and residents – has been core to our approach.

The School Readiness and Family Context project teams, for example, undertook user research with residents to understand their perceptions, comfort and expectations around how data would be used to inform and target services. The user research synthesis highlighted the need to provide guidance to social workers and other key workers on how they should use this information during assessments and conversations with families.