Department for Education

Responsible Technology Adoption Unit

Guidance for developing data analytics tools: Explainers

For local authorities developing data analytics tools in children's social care

April 2024

Contents

Summary

This document is a compendium of “explainers”: short chapters which provide more detailed advice and good practice on common challenges that local authorities can face in developing data analytics tools in children’s social care. Each chapter covers a different aspect of running data analytics projects in this sector:

- Chapter one covers how to **communicate** transparently and take a **co-design** approach to developing data analytics tools
- Chapter two covers **sharing data** legally and ethically
- Chapter three covers making sure your **data quality** is fit for purpose
- Chapter four covers **bias mitigation**
- Chapter five covers how to **test and evaluate** your projects’ impact
- Chapter six covers establishing effective **governance and oversight** on your data analytics projects
- Chapter seven covers the main issues and challenges with **procuring** third party services and tools

Who this publication is for

- **senior leaders in children’s services** will find the chapters on Co-design & communication, Data Quality, Governance & Oversight useful.
- **data practitioners or developers** will find the chapters on Co-design & communication, Data sharing, Data Quality, Mitigating Bias, Testing & Evaluation, Governance & Oversight, and Procurement useful.
- **procurement officers or commissioners of children’s services** will find the chapters on Mitigating Bias, Testing & Evaluation and Procurement useful.
- **information governance teams** will find the chapters on Data sharing, Mitigating Bias, Governance & Oversight useful.

How this publication should be used

This document should be used in tandem with the guidance for developing data analytics tools for using data analytics in children’s social care. You should use this document as an index of individual issues which are likely to arise at various stages of the development lifecycle. When an issue in the guidance for developing data analytics tools points towards an explainer, you should consult this document appropriately.

Chapter One: Co-design & Communication

Main points

- Clear communication is vital to build trust when developing data analytics tools for children's social care.
- Co-design is an approach that involves directly involving stakeholders (such as frontline practitioners, managers, children and families), in the design process throughout the development of a tool.
- Co-design can help to identify and address problems in a tool early on, and can save time and resources in the long term.
- This chapter will help you:
    - communicate with a range of audiences about data analytics tools
    - identify relevant stakeholders to involve in communication and co-design
    - engage with stakeholders across the development lifecycle

Communication on data analytics projects

Clear and transparent communication is a key element of responsible innovation. Good communication can encourage public participation and involvement, allow experts to ask questions to make the tool more robust, and support a wider and more informed public conversation about the use of data and data analytics in society. Failing to communicate, or communicating poorly about data analytics can leave the public wary and distrustful, and can result in the public rejecting a tool or approach.

How to communicate on data analytics projects

- project managers are advised to complete and publish information about a tool ahead of a tool being deployed. This can help to inform experts, civil society groups and the general public about the tool that you have developed.
- senior leaders can accompany any communication about data analytics alongside broader communication around how the council uses data. This can help the public to be more informed and able to be involved in conversations about how data is used.
- procurement officers or commissioners should ensure that any procurement contract allows information about the tool and how it was developed to be published.

- developers may wish to publish technical details about the tool and its performance, for example on GitHub, to encourage validity of the tool from external technical experts.

Complete the Algorithmic Transparency Recording Standard

Project managers are advised to complete and publish information about a tool ahead of it being deployed. This can help to inform experts, civil society groups and the general public about the tool that you have developed. The Cabinet Office, supported by the Responsible Technology Adoption Unit, has developed the algorithmic transparency recording standard, which you can use to publish information about your tool and what it will be used for.

Resources for communication

- the Ada Lovelace Institute published Participatory data stewardship which explores how citizens can gain increasing control and agency over their data through different communication and engagement.
- the ICO and Alan Turing Institute have developed a guide, Explaining AI decisions. It provides a framework for explaining processes, services and decisions delivered by AI, to improve transparency and accountability.
- the Ada Lovelace Institute, AI Now and Open Government Partnerships have explored Algorithmic accountability for the public sector, comparing different models and standards of transparency.
- Your website team should ensure that information about a tool meets the international WCAG 2.1 AA accessibility standard.
- You may wish to use Government guidance to make your website accessible.

What constitutes co-design

Co-design is an approach that involves directly consulting those who are specifically impacted by a tool, project or service as you develop it. The methods for doing this can vary, but all of them involve stakeholders taking an active role in co-creating how a tool is developed.

Co-design goes beyond merely informing or consulting individuals about how their data is used. It is about collaborating with and empowering those who are likely to be impacted by a tool. This can help those impacted feel that they have a degree of ownership over the process.

Why co-design is important

It's important that those who may be affected by data analytics are given the opportunity to shape its design and development.

In addition, social care is a highly intuitive and context-specific profession. The expertise and professional judgement of social workers and the lived experience of children and families in the social care system should be central to the development of any data analytics tool.

Co-design may also result in a better-designed tool. The process is likely to allow users, affected stakeholders and other subject-matter experts to point out problems early during the design and development, potentially saving time and resources.

Finally, co-design can contribute to the success and sustainability of a tool where users and affected stakeholders can see that the tool has considered their expertise, needs and lived experience.

Conversely, lack of engagement with frontline practitioners and managers can mean that data analytics tools aren't used, and money and resources may be wasted in development.

Implementing a co-design approach

A co-design approach consists of three steps:

- identify affected stakeholder groups
- establish effective communication
- run co-design sessions

Identify affected stakeholder groups

We recommend that you identify the stakeholders who are likely to be affected by your data analytics tool. Consider how you will best engage with each group throughout the development of a data analytics tool.

Affected stakeholders will likely include the following:

- Children and families
    - care leavers, via local authorities' children in care council or care leaver board
    - community organisations or voluntary services who work directly with children and families

    - relevant expert organisations where your tool addresses a specific issue. For example: sexual exploitation, you may engage with Barnardos or the Independent Inquiry into Child Sexual Exploitation (IICSA)

- Social workers or other frontline practitioners
- Service leads, principal social workers, directors or heads of service
- Safeguarding leads in schools and hospitals, police, extended family, community organisations providing support

Establish effective communication and engagement

The following approach can support proactive engagement with users and affected stakeholders.

Engage with affected stakeholders from the very beginning of the project

Engaging early about the challenges that stakeholders face, to diagnose a problem that can be addressed by data is crucial.

Bring senior managers and principal social workers onboard to promote co-design and champion the importance of data

Creating a shared narrative around the value of information and co-design can ensure that impacted stakeholders are brought into the process.

Assign adequate resources to engage with stakeholders throughout the development process

You could create a stakeholder engagement role. They would be responsible for planning and delivering co-design. A closer relationship between data and performance teams and frontline teams can help identify where better access or use of information could improve service delivery.

Communicate regularly with stakeholders

Regular updates with stakeholders throughout the design and development process can be an effective means of maintaining engagement and getting input.

They should be kept informed about:

- the tool's progress
- future decisions and opportunities to feed-in
- how their feedback has been incorporated.

Build relationships between child and data practitioners

Those communicating about data and data analytics should:

- understand data science
- be able to communicate effectively with child and family practitioners

Better relationships between child and data practitioners can enable identification of how and where data tools could support the provision of care.

Outline how and when stakeholders can get involved in co-design

Be transparent from the outset about:

- the level of influence participants in co-design can expect to have
- where there may be limited scope for change.

Run Co-design sessions

Co-design sessions can take many forms, but all of them involve directly consulting those affected by a tool on its development.

Select a method for co-design engagement

Structure co-design with practitioners in a way that requires the minimum amount of organisation or time commitment.

- surveys may be appropriate for more simple yes or no answers, and can be an effective way of gathering views quickly and anonymously.
- drop-in lunchtime sessions can mean that practitioners, managers, and other staff can engage without creating additional pressures on their working day.
- monthly social worker forums which engage senior leaders, principal social workers and practitioners can be used to provide updates, get feedback, and flag ways for practitioners to get involved in the next phase of work.
- a focus group run by experienced facilitators may be most appropriate if you want to discuss more complex issues or a use of data analytics which is high risk
- children's or social worker scrutiny panel can allow for stakeholders to provide constructive advice and challenge, and help to ensure that a tool is grounded in what they want and need

- A one-to-one chat may be more appropriate to capture the views of those who are less confident in group settings, or for children with language barriers or disabilities.

Communicate the rationale for exploring, and then deploying, data analytics in children's social care

Explain how data analytics tools can help to support decision making and improve existing processes and service delivery.

Some local authorities have found it helpful to organise a workshop with practitioners to demonstrate how data analytics can be useful. In the workshop, they have taken real life case review scenarios, and explored where better use and sharing of information could have made a difference.

Before a tool is deployed, you should share:

- the rationale for the tool and why it is being introduced
- expected benefits it will deliver
- protections that are in place to ensure data and the tool are used responsibly.

Topics for co-design

Taking a co-design approach means asking different questions throughout the lifecycle of a project so that stakeholders are co-creating the tool. **Annex A** contains a full list of questions that you can use to shape a co-design workshop or consultation for each stage of the development cycle.

Encourage practitioners and families to raise concerns about data analytics

Develop channels for affected stakeholders to:

- challenge decisions
- voice concerns
- seek advice or further information

You should be honest about the opportunities and risks around using data, and be clear about how you are addressing those risks. You may point out that data can be used to address bias, and that human judgements on their own may be biased.

Incorporate co-design into procurement

It is crucial that any external third party working with the local authority engages with domain experts and those with lived experience.

When engaging with third party consultants, ensure that they incorporate the co-design process into their way of working.

When a software developer is developing your tool, you should ensure that they participate in the co-design approach, for example, by observing discussion sessions and speaking to key stakeholders.

We would recommend that you include this as part of your service agreement.

Chapter Two: Data Sharing

This chapter explores some of the challenges to sharing data effectively, and describes how to mitigate some of the technical, ethical and legal risks brought about when sharing data. It sets out actions and decisions you may need to take to:

- identify what data you need to access to be able to develop your chosen data analytics solution, and what barriers there may be to accessing this.
- adopt good practices for sharing and using data, such as implementing role-based access controls, a robust logging and auditing infrastructure, or developing APIs for making certain datasets accessible.
- work with information governance teams to identify a legal basis for data sharing, and develop data sharing agreements with partners where appropriate.

Main points

- data sharing involves an organisation providing access to information it holds to one or more other organisations.
- poor data sharing can mean social workers don't have a complete picture of a child's safety or wellbeing, leading to ineffective interventions and poor outcomes.
- effective data sharing can help deliver high quality services, reduce the burden on social workers, and drive more efficient central planning and resource allocation.
- ensuring data is stored, shared and used securely and effectively requires a holistic approach that includes appropriate technical, legal and organisational mechanisms.

What data sharing encompasses

Data sharing involves an organisation providing access to information it holds to one or more other organisations.

Sharing data can ensure social workers, analysts, and other individuals working in a local authority have timely access to relevant information. This information enables them to make informed decisions and carry out their role effectively.

Moreover, integrated datasets - created by sharing and linking together previously siloed datasets - can increase the utility of data, enabling deeper insights to be derived from the data.

Why data sharing is important

Poor information sharing can mean social workers don't have a complete picture of a child's safety or wellbeing. Reviews and inquiries into children's social care have often pointed to the lack of data sharing between agencies as a reason for failures.

Effective sharing of data between agencies can help to:

- build a more holistic understanding of a child and their circumstances. This can enable a more informed assessment to be made of a child's vulnerability.
- make sure that vulnerable children do not fall out of the system when, for example, they move to a different local authority. The establishment of multi-agency safeguarding hubs (MASH) have helped ensure information is shared between agencies for this purpose.
- ensure that a child does not need to retell their story to multiple professionals

Better sharing of data can also reduce the amount of time and effort social workers spend inputting data to case management systems. This can allow them to spend more time directly engaging with children and families.

Finally, effective sharing of information by local authorities with central government can drive more efficient central planning and resource allocation.

Challenges to data sharing

However, data sharing can be challenging and is not without risk, particularly when handling datasets containing sensitive or personal data. There are also further barriers to data sharing:

Technical barriers, such as a lack of common data standards or variable data quality can make it difficult to link datasets held by different authorities and agencies.

Legal barriers, such as difficulty in establishing the lawful basis for sharing, or the need to create new data sharing agreements for each partner you are sharing data with.

Cultural barriers, such as risk aversion resulting from historic data breaches, or uncertainty about compliance with data protection law.

Technical and legal safeguards, implemented alongside a robust information governance function, can help to overcome these barriers.

Security safeguards for data sharing

One of the greatest challenges in data sharing is that it raises the chance of data loss or theft.

Successful cyberattacks or data breaches can significantly damage trust in a local authority's ability to safely handle data about citizens. An attack may also compromise the functionality of an organisation's systems, impacting day-to-day operations relying on such systems.

Cybersecurity standards and certification

There is a mature cybersecurity ecosystem providing standards and certifications that organisations can implement to minimise their vulnerability to a cyberattack.

The following resources may be useful standards for your organisation or project partners:

- ISO27001 standard which can help keep information assets secure
- The National Cyber Security Centre's (NCSC's) Cyber Essentials certification. NCSC Cyber Essentials is a government backed scheme to help protect organisations against a range of the most common cyber attacks. Cyber Essentials is mandatory for central government contracts involving personal or sensitive data. Local authorities should strongly consider adopting Cyber Essentials for their own IT systems, and contractually obliging partners to do the same when procuring services.
- For analysing your approach to protecting a specific bulk data set, the NCSC provides an excellent good practice guide to protecting bulk personal data.

Data security best practice

This includes:

- implementing role-based access controls, where only authorised personnel are able to access higher sensitivity datasets. This can minimise the risk of sensitive data about children and families being used inappropriately.
- a robust logging and auditing infrastructure. This can create an audit trail of how data is accessed, used and modified. This is invaluable in troubleshooting errors or problems caused by a data-driven tool.
- developing APIs for making certain datasets accessible. Making data accessible via APIs can enable sharing in a more efficient, automated, and secure way. APIs can limit sharing of a dataset to the specific records that are required.

The LGA has developed Tools and services supporting the use of local data to support data publication and sharing by local authorities. The resources include:

- open data standards and schemas
- catalogues of published data
- tools to support data linking, validation and aggregation

The above safeguards provide a strong foundation for safely and securely collecting, sharing and processing data.

Technical safeguards for data sharing

In addition to the security techniques described above, technical mechanisms described below can help unlock greater access to sensitive datasets.

Apply de-identification techniques, such as anonymisation

Traditionally, de-identification techniques (such as redaction, tokenization, or k-anonymity) have been applied to sensitive datasets to conceal all but the minimum amount of information necessary from the person accessing the data.

Local authorities should apply such de-identification techniques to sensitive datasets in line with the ICO's Anonymisation Code of Practice.

Explore privacy enhancing technologies

Privacy Enhancing Technologies (PETs) are a set of techniques to enable data to be accessed with strong guarantees of privacy and security, whilst maintaining the data's full utility.

When effectively implemented, PETs can be transformative in enabling safe and secure access to previously siloed datasets. Data mature local authorities in particular should consider how these technologies could be applied to unlock valuable use of sensitive data.

The Responsible Technology Adoption Unit has published a PETs Adoption Guide, which sets out information and use cases on how to develop and deploy PETs in practice.

Explore data intermediaries

Data intermediaries provide another mechanism for reducing the complexity and regulatory burden of sharing data between local authorities and agencies.

Data intermediary is a broad term referring to organisations providing technical infrastructure, expertise and/or governance processes for enabling data sharing and processing between organisations.

The Responsible Technology Adoption Unit has published a research paper, Unlocking the value of data: Exploring the role of data intermediaries, describing the types of data intermediary and how they could lead to more effective use of data.

Legal safeguards for data sharing

Data protection law is relevant where personal data is being shared or used to develop data analytics tools. However, many local authorities are uncertain about what they can legally share and whether consent is needed. This can mean that data isn't shared about children and families when it should be.

Sharing data to safeguard children

The Government has developed specific Guidance to safeguard children.

The Government has also developed information sharing advice which is specifically designed for child and family practitioners working in a safeguarding context.

Legal bases other than consent to support safeguarding

It is not necessary to seek consent to share information for the purposes of safeguarding and promoting the welfare of children. Other appropriate legal bases to use and share information can include: 'legal obligation' or 'public task'.

The ICO hosts a data sharing information hub. The hub:

- provides guidance and tools for organisations to share data lawfully
- can help local authorities to standardise their approach to data sharing

Legal gateways to share data under Part 5 of the Digital Economy Act 2017

Part 5 of the Digital Economy Act 2017 introduces new powers to share information in certain situations. This includes sharing data to improve the delivery of public services.

The Government has published a code of practice for organisations sharing information under Part 5 of the Digital Economy Act 2017. The guidance includes details at section 2.5 about the process that public organisations should follow when using the power.

Frameworks and templates for data sharing agreements

A data sharing agreement is a document setting out the purpose of data sharing between parties. It covers:

- what happens to the data at each stage
- standards for how data is used.

A data sharing agreement can help the parties involved better understand each of their roles and responsibilities.

Whilst it is not a legal requirement to create a data sharing agreement, it can help organisations demonstrate that they are meeting their accountability obligations under UK GDPR. Being transparent about the data sharing agreements that are in place can also promote greater trust.

Chapter Three: Data Quality

This chapter sets out the importance of data quality, and the actions you can take to make sure your data is fit for purpose.

Main points

- data quality is best understood as the ‘fitness for purpose’ of a dataset. Is the data good enough for what you want to use it for?
- there will be a number of dimensions of data quality that need to be considered. This includes: the representativeness, relevance, recency, and accuracy of the datasets.
- investing in tooling and mechanisms that facilitate accurate and complete recording of data can help improve data quality in the long-term, making it easier to develop quality data analytics tools.
- senior leaders should develop a culture that prioritises data quality
- developers should identify the dimensions of data quality that are relevant to their use case and take steps to mitigate any relevant data quality issues
- data practitioners should establish mechanisms for regularly assessing and auditing data quality

What constitutes data quality

Data quality is best understood as the ‘fitness for purpose’ of a dataset - is the data good enough for what you want to use it for?

For any given use case, there will be a number of dimensions of data quality that need to be considered. Which dimensions are relevant will depend on the specific context.

Why data quality is important

Data quality underpins all data analytics. The performance of a data-driven tool is dependent on the quality of data used to develop it: if you put poor data into a tool, you’ll likely get poor outputs.

Effective use of good quality data can:

- Drive efficiency
- Free up time for frontline staff to engage with children and families
- Help identify at-risk children who may previously have fallen through the cracks of the care system.

Such changes could be transformative for local authorities and the communities they serve, but good data quality is a prerequisite for this potential to be realised.

Local authorities that are seeking to develop their data analytics capabilities should prioritise data quality.

Ensuring good quality data is a continual process that requires senior leaders to:

- develop a data quality 'culture', underpinned by effective technical and organisational measures.
- invest resources into appropriate interventions and mechanisms for ongoing maintenance and improvement of data quality
- empower staff to dedicate time and effort to this.
- The Government Data Quality Framework provides guidance to help public sector organisations develop a 'culture' of data quality. The framework sets out principles, practices and tools for improving data quality.
- The framework is accompanied by an introductory course, and also provides a data quality action plan to help organisations understand the steps required to reach a more mature level of data quality.

Dimensions of data quality

Research by the What Works Centre for Children's Social Care (now Foundations - What Works Centre for Children and Families) has highlighted four key dimensions that are likely to be important when considering data quality in children's social care:

- representativeness
- relevance
- recency
- accuracy

Whilst no single tool or intervention will lead to perfect data quality, the options below describe practical steps organisations can make to improve the quality of their data in these four dimensions.

Representativeness

Data used for developing a tool should accurately reflect the underlying population the data will represent. It is important to consider the representativeness of all data sources that contribute to the operation of the tool.

Where data is about individuals or families, there is a risk that over- or under-represented groups could be disproportionately affected by the outputs of the tool. Such imbalance in a dataset can be a source of algorithmic bias.

In children's social care, it may be particularly challenging to ensure representativeness. Families from poorer socioeconomic groups are disproportionately represented in social care, as well as in other datasets, such as criminal justice records.

This can lead to:

- over-intervention against poorer families
- missing at-risk children from more affluent families.

Improving the representativeness of a dataset may necessarily mean collecting data about a larger number of children and families, including those not in the care system.

There is an important tradeoff between:

- The number of children and families about which information is collected
- The risk of perpetuating bias through the use of more limited datasets.

Relevance

It is important to consider the relevance of the datasets, and specific fields in those datasets, to the task at hand.

Domain knowledge is crucial here, and so it is important to engage with relevant stakeholders to understand which data it is reasonable to include (e.g. frontline workers, academics).

Positive factors that can influence outcomes should also be considered by domain experts. For example, a child having a particularly supportive and engaged wider family network may reduce their vulnerability.

Capturing this in the data used to train a tool can therefore improve the effectiveness of its outputs.

Recency

Data should be recent enough to capture and represent the phenomena being modelled.

Insights or predictions that are based on outdated data or inferences learned from a context that no longer exists will be inaccurate.

A data-driven tool should be updated whenever policies, procedures, and practices are changed.

Information about the recency of data should be recorded and provided to users of a tool to help them understand its limitations. More generally, statistical models should be accompanied by information about the data the model was trained on.

It may be appropriate to deploy a new model with a 'sell by date' after which its performance should be re-assessed, and the model potentially retrained on more recent data or retired completely.

Accuracy

Inaccurate or incomplete data means factors are not fully represented, which can undermine the performance of a tool.

Inaccuracies in data recording can stem from:

- human biases
- errors
- linked records from separate datasets which erroneously connect separate individuals, or fail to connect data about the same individual.

Any mechanisms put in place to address accuracy must still provide social workers the autonomy to record their subjective judgements.

Whilst no single tool or intervention will lead to perfect data quality, there are a number of practical steps organisations can take to maximise the accuracy and completeness of data at the point it is recorded.

Professional development and training

Frontline workers should be trained on how to effectively use the case management systems and other data entry tools used by their local authority.

This can be beneficial for:

- practitioner's professional development
- fostering greater trust and confidence amongst staff in the technology being used
- lead to more comprehensive and accurate data entry that improves the quality of data used to develop data-driven tools.

Invest in infrastructure for recording data

The quality of data will only ever be as good as the system used to record it. Senior leaders should provide frontline workers with modern data entry software that:

- is easy to use
- has an intuitive user interface
- minimises the chance of incorrect data being recorded

Introduce mandatory data standards

Enforcing standards for data can improve consistency and accuracy, and ensure important data fields are not omitted.

This can improve the ability to link information from different datasets within a local authority.

Sharing and linking information across different authorities and agencies can be improved by adopting a shared set of standards.

Continual assessment and audit

Processes for continually assessing aspects of data quality relevant to a tool's performance should be established and documented.

Depending on the outcome of this assessment, a tool may need to be refreshed or retired if the data used to develop it is sufficiently outdated.

When procuring a tool, you should have appropriate contractual agreements in place with the provider to ensure the long-term maintenance of the tool.

Resources

- The Central Digital and Data Office's guidance, Assessing if artificial intelligence is the right solution, provides a longer list of dimensions of data quality that may be particularly pertinent when developing and deploying a machine learning-based tool.
- The Local Government Association's Data maturity tool helps local authorities to self-assess different components of data maturity. Whilst not explicitly aimed at assessing data quality, the tool poses a number of relevant questions in the "Use" and "Management" sections.

Chapter Four: Bias Mitigation

This chapter describes how to mitigate bias in data analytics tools, and support fair outcomes in children's social care. It defines algorithmic bias, outlines the risks it can cause when using data analytics, and sets out steps to mitigate these biases occurring.

Main points

- decisions made by humans can often be affected by individual and societal biases, which are often unconscious.
- algorithmic bias is when the results or insights from a data analytics tool produces an outcome that is deemed to be unfair. Algorithmic bias can enter into a decision making process at many points including when: data about children and families is collected and recorded, a tool is designed and built, and insights are interpreted by a human decision-maker.
- understanding how to address and mitigate bias risks is crucial to ensure that tools are accurate and have a positive impact. In addition, unfair outcomes as a result of protected characteristics are discriminatory under the Equality Act 2010.
- many organisations are aware of the risks of bias, but can be unsure how to address it in practice. Risks can be managed with appropriate transparency, oversight and testing. Data can help to identify where bias is occurring, help to investigate why, and measure whether efforts to address it are effective.
- It is also important to engage with impacted stakeholders to think about what they think is a 'fair' outcome for algorithmic tools.

Algorithmic bias explained

Algorithmic bias is when biased decisions about the design of a data analytics tool means that the results of a tool produce unfair or inaccurate outcomes.

Where bias can enter into a data-driven process

The following are examples of how bias might be encoded within a data analytics tool in children's social care:

- practitioners and local authorities collect and record data about children and families. Recording data can hold conscious and unconscious biases. It can also reflect wider historic and societal inequalities.
- data teams collect and select data to build a tool. This data may over or under represent certain groups, and mean that a tool is more or less accurate for certain people. It can also give a skewed picture of a group.

- professionals use their judgement to interpret the insights from a tool, and decide what action to take. There is a risk that bias enters the process as the professional applies their own conscious and unconscious biases.

Why mitigating bias is important

Understanding how to address and mitigate bias risks is crucial to ensure that tools have a positive impact and are deemed to be fair.

Unfair outcomes as a result of protected characteristics are discriminatory under the Equality Act 2010. The Public Sector Equality Duty (PSED) is an ongoing duty that requires public bodies to have due regard to the need to eliminate discrimination, advance equality of opportunity and foster good relations between different people when carrying out their activities.

Failing to consider whether software could have a discriminatory impact may be a breach of the Public Sector Equality Duty. This Responsible Technology Adoption Unit blog explains more about the Public Sector Equality Duty and bias in algorithms.

Technical measures to mitigate bias

Make sure your datasets are as representative as possible

You should consider how to make your dataset more representative. You should also record and communicate the ways that the dataset is not representative.

Make sure that the tool's results can be interpreted by practitioners

Practitioners will be able to provide relevant context and judgement, and may be able to challenge bias.

Professionals should understand:

- the capabilities and limits of the tool
- why human judgement is important
- where human judgement is needed.

Using protected characteristic data can help to identify bias

Protected characteristic data can help you to understand how different groups and communities are affected by a policy or project. This data can also enable you to monitor and identify potential bias.

- The Information Commissioner's Office has developed Guidance on AI and data protection which advises organisations to take care when processing special category data to assess and address discrimination in data-driven tools.

Regularly monitor outcomes for bias

Analysis can help to measure and compare disparities between different characteristics (e.g. by race, socio-economic status).

If and when bias has been identified, you should take steps to address it. If it isn't possible to address, you should consider whether the tool is safe to deploy.

- The Responsible Technology Adoption Unit published a report to help organisations understand algorithmic bias and fairness, including Bias identification and mitigation in algorithmic decision-making. It includes an implementation handbook to guide technical teams through the process of building, monitoring and reviewing algorithms. The accompanying website explores technical ways to mitigate bias.

- The Child Welfare Inequality Project app can help local authorities to investigate trends, patterns and links between inequality, ethnicity, and child protection interventions.

Organisational measures to support fairness

Build a diverse team

Organisations building and deploying algorithmic decision-making tools should make increasing the diversity of their workforce a priority.

Collecting data about the diversity of the workforce in recruitment and progression can help you to monitor and improve.

Work with key stakeholders to identify a 'fair' outcome

You should identify the definition of fairness that is most appropriate to the decision making process the tool supports.

'Fair' decision making can be categorised as procedural and outcome fairness:

- procedural fairness is about treating all people fairly within how a decision is made. Will individuals be treated fairly in the decision that the data analytics tool is supporting?

- outcome fairness is concerned about what decisions are made, i.e. measuring average outcomes of a decision-making possess and assessing how they compare to a baseline.

What a fair outcome means is highly subjective. You need to discuss and decide on what 'fairness' means in your context before a tool is built and designed.

Establish clear responsibilities for addressing bias

Senior leaders in children's services should:

- understand the capabilities and limits of tools, in order to make informed decisions about whether a tool should be developed or used.

Oversight and scrutiny groups should:

- carefully consider how individuals will be treated by the decision making process that the tool supports. It's important that an overall decision-making process is fair, and not just that algorithms are unbiased.
- make a conscious decision about the appropriate level of human involvement in the decision making process.
- understand legal obligations, and carry out relevant impact assessments.

Data practitioners and developers should:

- put structures in place to gather protected characteristic data and monitor outcomes.

Legal measures

Complete an Equality Impact Assessment

An Equality Impact Assessment (EIA) is designed to help organisations make sure that their decision-making processes are fair.

An EIA can help you to:

- understand potential effects of the tool by assessing the impact on different groups
- remove and mitigate any adverse impacts
- ensure decisions are transparent and clear.

An oversight or scrutiny group may consider your EIA, and advise on the most appropriate way to move ahead.

- no barriers or negative impacts might mean that the project can proceed carefully;
- evidence of bias might mean that the project should be paused, and work done to explore how bias can be addressed or eliminated.

Resources

- The Equality and Human Rights Commission has developed Guidance for meeting the equality duty when using AI and data-driven technology. The Commission has also developed a checklist for public bodies.
- The Local Government Association has developed an Equality framework which helps authorities to address the public sector equality duty.

Incorporate relevant terms into your procurement contract

When procuring third party expertise, ensure that your contract includes terms that require any consultancy or organisation to:

- share information on the representativeness of the dataset that is collected to build the tool
- test the accuracy of the tool with a representative sample of your local population
- test for fair outcomes

Any technical testing completed by a third party must consider the accuracy rates for different groups, as well as the whole population.

Chapter Five: Testing & Evaluation

The guidance for developing data analytics tools includes an explicit testing stage, which provides guidance for carrying out testing in a controlled environment. The results from the testing stage are used to consider whether the tool is ready to be used in a live setting.

This chapter provides more detailed information on how you can effectively perform such an evaluation, and methods for ongoing testing and evaluation of a tool once it has been deployed.

Main points

- testing and evaluation assesses and assures the performance and effectiveness of a data analytics tool. There are different types of testing and evaluation that can be performed when developing and deploying a tool.
- system testing assesses whether a tool performs to the required level of accuracy (or other relevant metrics), primarily through quantitative measures. System testing normally occurs before a tool is deployed operationally.
- process evaluation assesses the success of the tool's implementation and integration into existing processes, through qualitative measures such as interviewing users of the tool.
- impact evaluation assesses the overall effect of introducing the tool, and the extent to which intended outcomes have been achieved. This can only occur after a tool has been deployed operationally for a period of time.

Testing and evaluation explained

Testing and evaluation assesses and assures the performance and effectiveness of a data analytics tool.

Data analytics tools that are deployed should demonstrably improve outcomes in children's social care.

We recommend that any tool deployed in a local authority be rigorously evaluated to ensure it:

- improves existing processes
- leads to better outcomes for children and families
- has provided value for money.

There are different types of testing and evaluation that can be performed when developing and deploying a tool. Here, we cover three types:

System testing

This assesses whether a tool performs to the required level of accuracy (or other relevant metrics), primarily through quantitative measures. System testing should occur before a tool is deployed operationally. It is also known as statistical validation.

System testing can be split into alpha and beta testing.

- alpha testing involves retrospective validation on historic data, to assess the statistical validity of the tool in a 'lab environment'.
- beta testing means piloting the tool in a live context for the purposes of evaluating its accuracy 'in the field', when presented with new, unfamiliar data.

Process evaluation

This assesses the success of the tool's implementation and integration into existing processes, through qualitative measures such as interviewing users of the tool.

Impact evaluation

This assesses the overall effect of introducing the tool, and the extent to which intended outcomes have been achieved. This can only occur after a tool has been deployed operationally for a period of time. It is also known as outcome evaluation.

Why testing and evaluation is important

Successful alpha and beta testing can provide a good level of confidence that the tool will be effective when fully deployed.

Effective system testing ensures a tool is:

- statistically valid
- achieves the desired level of accuracy
- free from unacceptable levels of bias before it is deployed operationally.

Establishing success criteria

The first step of any testing and evaluation is to establish criteria that define what the successful introduction of a data analytics tool looks like. Criteria should be established before the evaluation takes place.

You should consider what data will need to be collected to be able to evaluate the tool against these criteria.

Collecting data for benchmarking

You may need to collect data to benchmark the effectiveness of existing processes. For example, if introducing a tool to predict whether a child is vulnerable you need to understand how accurate your current process is in order to understand whether a tool is improving on this.

Data practitioners should collect record qualitative and quantitative data to help to evidence the status quo. This will depend on your specific use case, but is likely to include:

- how the decision making process currently works
- current user behaviour in the decision making process
- data to evidence outcomes among children and families
    - children's outcomes
    - support they receive
    - services they have interacted with
    - demographic data, for the purposes of monitoring for biased outcomes.

You should carefully consider and document where such missing data could affect your measurements, and consider how you could reduce this impact. Where this data collection involves personal data, this can only be done if there is a lawful basis for collecting or processing it.

System testing

System testing is used to determine whether the tool achieves the required levels of accuracy and robustness, (or other relevant metrics) prior to it being widely deployed into a live environment.

These metrics should be compared against the pre-defined success criteria. If those criteria are not met, we would advise that the tool should undergo further development before being more widely deployed.

Which metrics are relevant will depend on the specific use case.

Process and impact evaluation

Process evaluation focuses on:

- Identifying the strengths and weaknesses of the processes surrounding the use of a tool
- Making recommendations for adjusting the structure and/ or implementation of the tool.

Outcome evaluation aims to provide an overall assessment of the outcomes of a tool in terms of its benefits, merits and worth. Outcome evaluation may be a longitudinal process taking place over several months or even years.

Interviews and observations of users

Users of the tool can be interviewed and observed as they use the tool. This can enable an understanding of:

- Its impact on their work
- Their perception of whether it is providing benefits, for example by improving their decision-making, or helping them use their time more efficiently.

Examples of things that could be assessed through such a methodology include:

- How regularly the intended users are actually using the tool
- Whether they are using the tool as intended
- Which components of the tool they are using most
- What they think the tool's limitations are and how it could be improved

Controlled trials

Impact can also be assessed through controlled trials that account for factors other than the tool that may influence outcomes.

Controlled trials are a form of impact evaluation in which one cohort of the population are subject to the data analytics tool (and will likely receive one intervention), while the second cohort acts as a control group (i.e. are not subject to the data analytics tool, and may receive a different intervention).

Well-designed controlled trials can provide increased confidence in the efficacy of a data analytics tool. They can help you know what is being achieved as a result of the data analytics tool. This can ultimately lead to better outcomes for children.

We recommend that local authorities communicate transparently about any controlled trial, to minimise the perception that a trial may be “experimenting” on children.

Having an independent expert evaluator conduct the trial can build greater trust in the process.

Independent evaluation

Contracting an independent evaluator can help ensure that the evaluation is robust and rigorous. It should be considered in higher risk use cases.

Independent evaluation can ensure that the assessment of the tool’s performance is objective, impartial and of high quality.

Independent evaluation may be particularly important for high risk use cases or tools that have been controversial in the past.

Local authorities could partner with academics or researchers to perform the independent evaluation, or with charities or companies offering evaluation services who have experience of the children’s social care sector.

Communicating results of the evaluation

Results of the evaluation may need to be communicated to different stakeholders, including:

- technical developers, data scientists and auditors
- senior leaders and managers
- end users
- those that may be affected by the tool

It is important that they are presented in an intuitive format that provides information at a level of detail appropriate for the stakeholder.

Chapter Six: Governance & Oversight

This chapter sets out the importance of setting up effective governance and oversight structures on your data analytics project, and how to make sure that decision making is accountable.

Main points

- governance is about ensuring that there are appropriate structures in place to support decision-making.
- teams should work with information governance to establish robust mechanisms for governance before making significant steps to invest in or develop new data analytics
- decisions about how technology is designed and used to support children's social care should be subject to oversight. Discussion and deliberation should occur inside and outside of the project team, with appropriate levels of independence where necessary
- where particularly high risk or sensitive projects are being considered, establish specialised independent bodies. This includes making sure that they are properly supported and there are clear feedback loops to project teams and decision-makers

Why governance and oversight is important

Good governance enables the trustworthy use of data analytics in children's social care. It provides:

- clear responsibility and accountability for decision making across development and implementation
- clear mechanisms to facilitate input, challenge and advice by experts to decision makers
- consistent and standardised processes for audit and improvement

Establishing internal governance and oversight

Robust mechanisms for governance should be established before a local authority makes significant steps to invest in new data analytics.

Assign clear responsibility and accountability at the outset

You should:

- identify who is responsible for tasks and milestones across the development and implementation of a tool
- identify key decision-makers' and decision points along the project lifecycle
- develop a risk process which identifies, monitors and raises risks during development to appropriate levels of seniority

Develop consistent and standardised processes

A key part of good governance is consistency and standardisation. Repeatable processes should be put in place to identify, address and test ethical considerations. This can help to make sure that there is a consistent approach and facilitate auditability.

Similarly, standardised processes should be established to:

- evaluate and capture the learning from each data analytics project
- disseminate lessons learned with the wider local government community

Local authorities should draw on existing internal governance structures and scrutiny committees to ensure the appropriate governance of data analytics.

These could include:

- children in care councils or care leaver boards
- data or technology boards
- ethics committees
- safeguarding partnership board
- scrutiny committees (e.g. children and families services)
- elected members

Establishing a project-level advisory group

Internal governance could involve setting up a project-level, advisory group. They would be involved throughout the lifecycle of a data analytics project.

The group would:

- provide input, advice and challenge on proposed data analytics projects

- review and sign off on initial business justification documents and project proposals.

The group may comprise:

- users of analytics tools
- in-house data practitioners or business intelligence developers
- data protection and legal expertise, including the designated data protection officer (DPO) and senior information risk owner.

The group should provide a centralised function for:

- assessing strategic priorities for data analytics
- identifying emerging business requirements
- procuring, testing and evaluating new data analytics tools to build an evidence base of what works
- discussing and addressing the legal, data protection and equality considerations of new data analytics tools.

External governance

Where particularly high risk or sensitive projects are being considered, local authorities may wish to establish a specialised independent body or group. They are also known as external advisory groups or data ethics committees.

This group would provide additional expertise, guidance and critical review of projects.

Example

- Essex Data Ethics Committee. An independent committee has been established to provide independent advice to the Essex Centre for Data Analytics. The Committee is made up of data science specialists, human rights experts, senior policy representatives, residents of Essex, and information governance specialists. The Committee considers project and policy proposals which relate to innovative and novel uses of data, and citizens' views are sought around data use and public benefits.

Developing terms of reference

The following are some suggestions about what any external body (such as an independent data ethics committee) should do:

- be an advisory, not a decision-making body.
- provide advice before decisions are made by the project team
- operate independently from the project team
- have a clear purpose and way of providing advice. For example: To advise leadership on whether - and with what amendments - data analytics projects should proceed, before significant decisions are made
- have rights to information related to the projects. This will make sure that they are able to give informed advice and maintain confidentiality of the projects discussed.

Whilst these bodies should bolster expertise and scrutiny, their role is advisory rather than decision-making. Project teams will remain ultimately responsible and accountable for their use of data analytics.

Chapter Seven: Procurement

Many data analytics projects in children's social care will involve procuring tools or data from the private sector. This introduces a set of particular issues not faced by building in-house systems.

This chapter sets out the key considerations when procuring data analytics tools from third parties, and steps you can take to help these tools be developed responsibly.

Main points

- Third party providers can bring valuable expertise and tools to the children's social care sector.
- It's essential to assess any systems or tools procured to ensure they are effective, accurate, relevant to the sector you are deploying it in, and are explainable and accountable.
- Local authorities should set out clear legal responsibilities for third party providers to comply with, and agree where some elements of responsibility for the outcomes of a tool sit with third parties.

What procurement looks like in this sector

In children's social care, third party providers include:

- software providers who build case management systems
- public sector consultancies who develop data governance capabilities
- academic teams who provide technical supporting in developing and building a tool
- academic teams who provide support and advice on public engagement, including co-design
- data science contractors who build analytical tools.

These systems or tools can be used for a range of services, for example:

- tracking the number of referrals received and the average caseload
- identifying children or families eligible for a support programme.

Ultimately, local authorities should be able to benefit in a lasting, sustainable way from the expertise and innovation in the private sector.

Supporting effective procurement

Procurement officers may find the following toolkits and guidance helpful in supporting robust and effective procurement:

- The London Office for Technology and Innovation (LOTI) and PUBLIC has developed an interactive toolkit to support procurement of technology products, services and support in the local government sector.
- The Office for AI has developed Guidelines for AI procurement in the public which provides advice around the procurement process.

Assessing effectiveness

Third party providers can come to local authorities with off-the-shelf tools which have been developed in different contexts and on different datasets from your own.

This can be beneficial as it means that:

- a local authority is not starting from scratch in developing a tool
- the off-the-shelf tool may factor in improvements that have been made when being tested elsewhere.

Training and testing

To make sure that the tool is effective in the context where it will be used, we advise that local authorities should - or ask third party providers to:

- train any model or tool on representative data
- test the tool for effectiveness

Ensuring maximum explainability and transparency

Some third party providers may want to protect the methodology (i.e. source code) underpinning an analytical tool. However, this should never inhibit a local authority in meeting its explainability and transparency requirements.

Local authorities should be able to fully understand how the tool works. It is their responsibility to:

- monitor the tool in an ongoing way

- provide explanations for how it works to users and affected stakeholders.

There should also be maximum transparency around how it was developed and the data used. This should not be limited by Intellectual Property considerations.

Share information publicly about procurement

We advise that you should include details about any procurement of tools or services in information that you share publicly about the tool and its development. There are numerous ways to share information publicly, for example, on your website and in public engagement with citizens.

Setting out clear responsibility and accountability

Legal responsibility and ultimate accountability for the tool sits with the local authority, although some elements of legal responsibility may sit with providers as part of service-level agreements.

This means that there needs to be widespread knowledge, understanding and visibility of the tools among senior leaders and elected members who are accountable for decisions made by the council.

Clarify legal responsibilities

Any third party is responsible for:

- confirming that their tool complies with the relevant legal frameworks, including the Equality Act and Data Protection Act
- providing necessary information to support project teams to complete Impact Assessments.

Request ongoing support

It can be valuable for local authorities to request that third parties provide ongoing support beyond the initial deployment of the tool. This can include:

- monitoring how it is performing
- identifying any tweaks that need to be made.

This will need to be discussed with the third party, and incorporated into your commercial agreement.

Annex A: Suggested questions for co-design

This section identifies some questions that you may discuss with end users and affected stakeholders. It corresponds to the lifecycle stages in the **Guidance for developing data analytics tools.**

Exploratory Stage

At this stage, you should engage with children's services teams to understand more about the department's data maturity.

Data maturity

- How do you use data to inform decisions and run services?
- What are your attitudes towards data?
- How is data stored, accessed and shared?
- How is data governed, and kept private?

Problem identification stage

At this stage, you should engage with children's service teams to understand their experience using data and their reactions to data analytics. You should also explore key demands and pressures that they face in their roles and their views on whether data can play a role in addressing the demands and pressures.

Experience using data

- Do you always have access to the information that you need?
- Could additional information help you to do your job?
- What are the main challenges in using data in your role?
- What are your reactions to data analytics?
- Do you have any questions or concerns about data analytics?

Reactions to data analytics

- Do you have any questions or concerns about exploring data analytics?

- How would you like to be involved in the design and development of a tool (e.g. lunchtime sessions, workshops)?

Identifying problems in children's social care

- What are the key demands, pressures and challenges that you face in delivering children's social care?
- What key outcomes are you driving towards?
- How could better information and analysis address key demands and/ or support desired outcomes?
- What other (non-data) solutions could address key demands and/ or support desired outcomes?

Exploring data-related solutions

- Who could benefit and how?
- What data would be needed?
- What action could be taken as a result of knowing the information?
- How would this change (and improve) how decisions are currently made?
- What, in addition to data, might be needed?

Choosing a solution stage

At this stage, you should engage with your children's services team to consider their views on the tool that you want to develop.

Reactions to data analytics development

- What are the potential impacts - positive or negative - for you? What are the potential impacts - positive or negative - for children and families?
- Do you have any concerns about the different options? What are they? Can you think of ways that your concerns could be addressed?
- How appropriate is the tool for your local context? To what extent do you think it will support and improve the provision of children's social care?
- Do you support the development of the tool?
- Are there other challenges that you feel need to be addressed?

Design stage

At this stage, you should map out the 'end users' and 'affected stakeholders', who should be involved in co-design. You should focus your engagement with future users of the tool, and those who will likely be affected by it. You should explore affected stakeholders' views on acceptable uses of data and co-design the tool's interface with end users.

Co-design outcomes and impact statements

- What outcomes and impacts do you want to see from the development and use of the tool?
- How could these be achieved?
- What would redlines or undesirable outcomes look like?
- How could these be avoided?

Identifying acceptable use of data

- What are your views/ reactions to the datasets that we want to use to develop and test the tool? Do you think they are relevant?
- What data do you think users should be able to access and use?
- Conversely, what information do you think is inappropriate for users to access and use?
- Do you have any concerns about how data could be shared, used or accessed?
- How should access to data be restricted and monitored?
- How can we improve children's safety and privacy?

Co-designing the interface

- How do you normally use data?
- Do you feel comfortable using data and technology? To what extent do you trust or rely on it?
- Which tools have you used in the past that have worked well?
- What has frustrated you about tools you have previously used?
- Which aspects have been good or bad?

Development stage

At this stage, you should share a prototype of the tool with stakeholders, and seek their views on it. You should also explore how the tool will be explained and communicated to end users and affected

Responding to the prototype

- What do you like and dislike about it?
- How could it be improved?
- How easily do you think this tool would fit into the existing decision making process?
- Is it easy to use?
- Would you use it as part of your work?
- What kind of training would you like to have to support you to use this tool?

Testing stage

At this stage, you should provide training for users on how to use the tool, and have users test the tool in a controlled environment.

User feedback

- What do you like and dislike about the tool?
- How do you think it could be improved?
- To what extent is the tool intuitive and easy to use?
- To what extent do you think that it has a positive or negative effect in supporting your decision-making?
- Would additional training be helpful to support you to use the tool?
- To what extent do you think that the tools fit well into existing decision making processes?
- To what extent do you trust the tool and its insights?
- Do you feel confident to flag an error in the tool? To what extent do you feel confident that it would be addressed?

- To what extent do you feel able to reach your own conclusions?

Reviewing testing and results

- What are your responses to the results?
- Are you concerned about any aspects of the tool or its test results?
- How confident are you that the tool would be accurate and effective where deployed live?
- Could you envision any adverse impacts?

Deployment and monitoring stage

At this stage, users and affected stakeholders should be involved in review and evaluation processes to understand whether the tool is meeting its objectives.

User review

In addition to the questions described above in ‘user feedback’ in the **testing stage**, you may wish to ask:

- Have you used the tool? How often?
- What have the benefits been?
- Do you have any concerns about the tool? What are they?
- Do you think the tool is meeting/has met the outcomes and objectives you set? Do you think that these are still relevant?
- Do you understand how to - and feel comfortable to - flag concerns about the tool?
- Have you provided feedback? Are you happy with how it was addressed? (Why?) How could the feedback process be improved?

Department for Education

Responsible Technology Adoption Unit

About this publication:

enquiries https://www.gov.uk/contact-dfe
download www.gov.uk/government/publications

Follow us on Twitter: @educationgovuk
Connect with us on Facebook: facebook.com/educationgovuk