Data Ethics Committee minutes: 8th September 2025

Notes of the Data Ethics Committee (DEC) meeting 8th September 2025 via Teams

Attendees

Presenters
Emily Moss (*Intelligence Manager, Essex County Council)*
Amos Turin *(Data Science Fellow, Essex County Council)*

ecda members
Clare McLean (*Operational Lead, ecda*)
Wyshnavi Kanesalingam *(Data Partnerships Officer, ecda)*

DEC members
Katerina D Hadjimatheou (Vice Chair)
Yani Tyskerud
Anton Beer
Karthik Durgaprasad

Project overview
During this committee meeting Essex County Council (ECC) presented two Artificial Intelligence (AI) Proof of Concepts (POC's) for review - *Summarisation and Thematic Retrieval*, and *Front door triage* – as part of it's exploration into AI solutions to help address different business challenges. The projects brought forward in the meeting were for test and trial purposes only, to understand what works and what doesn't, there is no intention to deploy either project.

Commented [KH1]: I think this wording could be cut as it invites the assumption that it will be deployed in the future, but it may never be deployed.

The Front Door Triage project

The first project reviewed was ECC's Front Door Triage Project, which explored how AI can support the increasing demand for referrals. With requests rising annually, ECC partnered with Faculty to enhance an existing AI model that helps triage incoming cases - sorting them into urgent action or no further action (NFA) - to support staff at the Front Door.

The goal was not to replace human decision-making, but to assist practitioners in making faster, more informed decisions, allowing them to focus on high-priority cases. Faculty also worked with ECC Data Analysts to build internal understanding and confidence in maintaining the model.

User interviews identified three key requirements for the tool to be effective: accuracy, reliability, and explainability. Faculty's improvements led to an AUROC score of 0.9 (this metric tells us how well the model can tell the difference between urgent and NFA cases), with the model correctly identifying urgent cases 88% of the time and achieving 82% overall accuracy. Notably, 25% of predictions were made with very high confidence. To ensure transparency, ECC used Shapley values to explain model decisions - revealing that predictions were largely driven by the free-text content in referral forms.

The Summarisation and Thematic Retrieval project

This project explores the potential of large language models (LLMs) to support automating the summarisation of case file audits.

Currently, ECC teams review 20 cases per month, with each case reviewed by two independent auditors. The team then compares both submissions and produces a summary document - a process that takes approximately four working days per month. These summaries are critical for evaluating service standards.

ECC deployed an LLM built in-house and worked with Faculty AI for expert feedback. To test its effectiveness, ECC used previously audited cases with existing summaries and compared them to AI-generated versions. The aim was to assess whether the model could produce accurate, concise summaries without omitting key information or introducing inaccuracies. Initial estimates suggest that using this model could significantly reduce manual workload, saving 2.5 months of staff time annually.

Conclusion from Committee discussion

The Front Door Triage project

While the model showed promising results, ECC identified challenges around integrating it into existing systems used by front door staff. The current setup would require social workers to switch between platforms, offsetting the time savings gained through the AI-model. It was acknowledged that further work is needed to design a technical framework that aligns with frontline workflows ensuring the tool is not only effective but also practical and user-friendly for staff.

The Summarisation and Thematic Retrieval project

ECC acknowledged that more time was needed to review and evaluate the model to improve the outputs. It outlined the need for further iterations to work on improving AI performance and accuracy of outputs.

Committee observations raised and responses provided:

The Front Door Triage project

- ECC highlighted a key limitation in the current care management system: the lack of real-time data reporting. This prevents timely feedback on triage requests, particularly those that fall outside of the expected four-hour response window. The committee acknowledged that delays between data upload and algorithm processing undermine the model's effectiveness within this timeframe. It was agreed that this issue could be explored further with ECC's system provider to enable more immediate data flow and better align the tool with operational needs.
- The committee sought clarification on the model's reported 88% accuracy rate, specifically how it compares to human decision-making and whether a 12% error rate is acceptable. ECC advised that the model's performance is broadly in line with human accuracy. Further analysis of the 12% of cases where the model's predictions differed revealed that some were re-referred or aligned with social worker decisions.
- The committee raised concerns about how the model ensures fairness in its predictions, particularly regarding protected characteristics. ECC confirmed that such characteristics were not used as input variables in the model to avoid bias. However, they were analysed separately to assess whether prediction accuracy varied across different groups, e.g. by ethnicity.
- The committee asked about the scale of the data used, ECC confirmed that the model was built using three years of data, comprising approximately 45,000 requests for

Commented [KH2]: What was the result of the analysis?

support. The model draws on both demographic information and free-text content from referral forms to determine whether a case requires urgent action or no further action (NFA).

- ECC clarified that decisions are made through a multi-step process: a frontline worker at the front door initially reviews the case to see if the case is NFA or urgent (this is the stage at which the AI model would assist), which is then passed to a senior practitioner for sign-off, followed by contact with the family before allocation to a social worker. This layered approach ensures additional oversight.
- The committee asked how frontline staff feel about integrating the model into their workflow. ECC shared that Faculty's conversations with potential users revealed initial hesitation and confusion, with some apprehension around the use of AI. To address this, Faculty worked to demystify the technology and clarify its role in supporting, not replacing, professional judgement. They also maintained fortnightly engagement with service managers, who were consistently supportive and provided valuable feedback throughout the development process.

Commented [KH3]: And how would the AI step fit into this? And how would oversight be maintained?

Summarisation and Thematic Retrieval

- ECC advised that although just the auditors' reports on the case file were used and not the actual case files themselves, these reports still contained sensitive information regarding care recipients that could not be anonymised or removed without negatively impacting the quality of outputs. ECC therefore liaised with Microsoft to ensure that all information input into the tool was encrypted both at rest and in transit. Additionally, the model was deployed on a local UK-based server, further safeguarding sensitive information.
- ECC acknowledged that due to delays in activating the generative model, there was only time for two rounds of prompt engineering, limiting opportunities to refine outputs. Despite this, initial results were encouraging: 41.54% of AI audit summaries were rated by the audit team as accurate and reliable. ECC advised further iteration was needed to improve model performance and output quality.

Project recommendations

The Front Door Triage project

- Committee advised on looking into co-developing the model along with a system integrator or the CRM provider to explore real time data flow.
- Committee requested further clarity on what kinds of errors were identified and were these errors different to the errors made by humans as part of the project evaluation. ECC were advised to look into this closely to understand the model.
- The committee recommended looking more extensively at how the model interpreted protected characteristics within its algorithm as part of the evaluation and advised to include experts in recognising bias in data or AI.
- ECC to investigate further the return on investment when it comes to time saved in the triaging process by comparing how long the current manual process takes for a social worker with how long it takes with the algorithm being used.

Summarisation and Thematic Retrieval

- ECC to do ethics assessments earlier within their project testing phases. If all questions cannot be answered at that early stage keep the assessment as a living document to refer back to.
- Committee recommended to include questions within the ethics review that ask the auditors what they think the risks might be, to allow the auditor to provide their own input.

General recommendations

Ensure senior team members and tool developers consult and collaborate with end users on all projects to ensure AI tools/models work to the users benefit and any inefficiencies or errors identified by the user can be rectified early on.

Engage with focus/consultation groups made up of lay people and those with lived experience of the social care system to gather feedback of models and tools prior to launching to get an understanding of community concerns.

Have a clear success criteria at the beginning of any testing phase for projects to guide evaluation and ensure transparency.

Consider small-scale rollouts (e.g. 5–10%) of tested tools within ECC systems to demonstrate value and build organisational confidence in AI technologies.

When reviewing impact and feedback, consider how users adapt to new tools. Encourage critical thinking and assess risks across short, medium, and long term horizons when developing ethical frameworks.

ECC requested general feedback on their ethics framework.

Committee requested a process to be put in place for ad-hoc meetings or an easier/quicker way to get assistance from the committee for reviews on future proof of concepts.