



Data ethics case study - NHS AI Lab

Summary

The National Health Service (NHS) will be trialling an Algorithmic Impact Assessment (AIA) to ensure that those who seek to use its data to train their artificial intelligence applications are open about the design of their algorithms and what impact they will have on research outcomes.

To do this they will:

- Create a model AIA that can be applied when researchers and developers request NHS data
- Work with those who design algorithms to ensure early considerations of possible impact, such as algorithmic bias
- Collaborate with an external organisation (The Ada Lovelace Institute) to learn from best practice and develop the model
- Use citizen panels to promote engagement with patients and the public in the use of medical data

Origins

The NHS established the Artificial Intelligence (AI) Lab to accelerate the safe, ethical and effective adoption of AI technologies in health and social care. It has £140 million for the AI in Health and Care Award to support the testing, evaluation and 'scale of promising AI-driven technologies'.

The AI Ethics Initiative is one programme of work under the NHS AI Lab and they work on the basis that data ethics is a practice and not an end point. They want to embed data ethics and move away from the idea of it as a burden.

One of the ways they are doing this is through creating a model AIA, which would complement Equalities Impact Assessments and Data Protection Impact Assessments. The AIA is a tool to support innovators to engage patients and the public early on in the development process for AI, when there is greater flexibility to make adjustments when potential risks are identified.

Description of the approach

The NHS AI Lab will be trialling an AIA on the proposed National Medical Imaging Platform (NMIP). The NMIP is a proposed centralised database for medical images. It has the potential to be a valuable national data asset that will help train and validate AI models.



The AIA concept is not new or unique to the NHS AI Lab; they are working with the Ada Lovelace Institute on a research project to determine how an AIA can be applied in the NHS in relation to the NMIP. The Ada Lovelace Institute published a [description of the project](#) with the NHS AI Lab, explaining the nature of the project as such in March 2021:

“It [the project] will develop actionable steps on AIAs for the NHS AI Lab, as well as to help inform wider AIA research and practice at the intersection of health and data science, and in the public and private sectors.”

In practical terms, what is proposed is that the AIA would be completed by researchers and developers when they apply to access data from the NMIP. The AIA would prompt them to reflect on ethical considerations, such as fairness and data privacy and security, and they would then explore potential risks and benefits with patients and the public as part of a participatory process.

The participatory process would be facilitated by the NHS AI Lab to ensure independence and would help developers take into account different perspectives on the possible impact of their proposed solutions prior to further development.

Review

If access has been granted on the basis of a satisfactory AIA, the applicants would be expected to review it at intervals (according to recommendations from the Data Access Committee for the NMIP).

Citizen considerations

The participatory process would involve members of the patients and the public. There would be a static panel of 25 - 30 citizens recruited to a panel that can be drawn on for each participatory session about the proposed AI solution in question. Similar to the model for citizen juries, participants would be recruited to be broadly representative of the population.

Further reading:

[Ethics as a service: a pragmatic operationalisation of AI ethics](#)