

A Vision for Health decision support AIs in the NHS and beyond

The US Government chooses to spend the biggest slice of its discretionary budget on the military. By contrast, the biggest amount of the UK's discretionary budget goes on the NHS; as a nation, we choose to commit resources to saving lives and helping people. Both choices have resonated around the globe. Technology allows us to move fast, the question is do we choose to help people, or harm them?

The aim of NHS efforts around data and AI are to support people's health; the UK's ARPA should accelerate that support, as the US's DARPA has done for US priorities.

If you are admitted to a hospital requiring a blood test in the UK, you will likely be screened for kidney injury. The NHS has a standard flow chart for assessing those results, and deciding what to do about them. That flow chart – the NHS AKI algorithm – is published openly online¹ so it's easy to find, and easy to check.

As a result of publishing this algorithm – a decision taken by the NHS, in the interests of the NHS and NHS patients, for NHS-only reasons² – if you are admitted to hospital anywhere else in the English-speaking world, you will quite likely be screened for kidney injury using the same algorithm, precisely because it is freely available and supports high quality decision making for all. There are other positive side-effects as well, both in terms of NHS recruitment and the UK's 'soft power' more generally.

Replicating this culture and practice for AI models and processes would offer significant benefits to health across the NHS that could be reused around the world. Rather than being grounded in exploitation, the UK could lead a system where a diverse range of stakeholders – each working in their own best interests but taken together – are greater than and not dependent on any individual actor, either public or private.

Reference datasets as the first building block...

One example of a reference dataset would be the COVID-19 chest imaging dataset (NCCID), which aims to “enable the validation and development of automated analysis technologies, and to promote research projects”.³

As this statement suggests, NCCID should become a dataset which covers all types of chest image, for all those the NHS serves. Supporting high quality clinical assessment, potentially with retrospective re-diagnosis and known outcomes; covering all relevant equipment, i.e. manufacturer, output style, etc. and human factors, e.g. gender, age, ethnicity, body type,⁴

¹ <https://www.england.nhs.uk/akiprogramme/aki-algorithm/>

² 'Make things open, it makes things better' – NHSX

³ <https://medphys.royalsurrey.nhs.uk/nccid/guidance.php>

⁴ Not all of these need to be available to dataset users, however they are necessary for the validation datasets to test bias and accuracy

that the NHS experiences – as well as common mis-diagnoses and missed diagnoses. The dataset should certainly cover all outputs from all equipment that is in use in NHS hospitals.

The same should be replicated in other clinical domains, with copies of reference datasets available on application to those who seek to build tools to help clinicians. (Several separate validation datasets should also be retained, for testing purposes.) The lead applicant for every project should, however, be a registered clinician – making it clear this is a clinically-led process to aid clinicians, and thereby patients. Just as every pharmacy has a responsible pharmacist whose name and GPhC registration number must be on a sign on the wall of the establishment for which they are responsible, every step of interaction with the NHS or NHS data must have someone clinically responsible.

All projects using such data or datasets should be recorded in a register that is accessible *by those whose data is being used*. This full transparency and accountability prevents shortcuts by charlatans and AI fraudsters who use less ethically-sourced and rigorously maintained datasets – who claim compliance with the rules, without the burden of following any of them.

For those with interests beyond the UK, equipment used elsewhere should also be included in the core (or an equivalent) dataset – whether that be EU-wide, or beyond.

...to stimulate Demand... (procurement for competitive markets + commodity pricing)

Many X-ray machines in use throughout the world would benefit greatly from the same diagnosis support tools that the NHS would seek to use, and this would be entirely affordable with commodity pricing.

For services built using the reference datasets, the NHS should commit to enabling a competitive market using both its procurement power and guidelines. The most straightforward way to achieve this is to insist that, where a care provider wishes to use AI as a diagnosis support mechanism, they must procure three *independently* built, maintained and sold systems,⁵ to ensure both patient safety and a diversity of systems in use.

Procurement rules should enforce high standards; the equivalent of sterile scalpels, rather than rusty knives. Procurement for any AI system must therefore also require transparency and accountability around the testing datasets used to build and validate that tool.

Those who take shortcuts around data protection, ethical sourcing of data and information governance should find their wares harder to sell, not easier. It should be impossible to ‘informally’ acquire a dataset and then expect to sell the resulting tool to any public body.⁶ Such rules would make a level playing field for those who are reputable and who follow the rules, and discourage those who like cutting corners.

⁵ We note a large AI company’s team in London has taken this approach internally, claiming their own three AIs satisfy this prerequisite; by definition they cannot.

⁶ https://www.whatdotheyknow.com/request/ai_agreements_with_orthai

...Supply... (from the UK economy)

Managed procurement can encourage a competitive market, e.g. providing incentives for new entrants where there is an insufficient diversity of suppliers.

The NHS can ensure a functional market that encourages investment which will meet NHS needs. Useful tools can be built by junior doctors with clinical and technical skills, who are currently falling out of the NHS⁷ – give them somewhere to go that contributes to the NHS, and provide those companies that hire them mechanisms to supply the NHS, in a market where no-one has to be worried they'll get flattened by one large provider. Win. Win. Win.

There are plenty of clinical areas where a modest amount of technical skill combined with meaningful clinical experience will provide approaches that would not be obvious to technical generalists driven mostly by large-scale profit motives.

Many large companies see the £110 billion NHS budget and think of the slice they would like to acquire – “such a small slice”, in their heads – but, as with the pharma budget, what goes to one new place has to come from elsewhere. Small, diverse, clinically-led suppliers are far more likely to create real value while also reducing overall costs to the NHS, than transferring as much value as possible to the tech giants.⁸

This is another reason why all projects should be clinically led by a named, responsible clinician. It also ensures there is a UK base for all of the AI that the NHS purchases.

...more Demand... (worldwide)

By insisting on a scheme of commodity pricing per-assessment for usage, plus additional charges for support and ‘value added services’, the NHS can ensure the maximal worldwide market is available. Indeed, as AI diagnosis support tools become able to be run locally⁹ on mobile devices just like any other local app on a smartphone, the marginal per-assessment cost will drop to zero – as marginal costs are borne by the hospital or clinic that chose to use the tool.

This has the potential to open up services worldwide, where a mobile phone, an X-ray machine and the skills to use both may be present, but where the clinical skills to diagnose marginal results are unreasonably far away. One sustainability model for Google Health is to offer free and ongoing disease monitoring to everyone with a Gmail account – as a service which monitors your blood sugar is one you'll continue to use... while continuing to look at the Gmail ads. (In the same way, if you are reassured that Apple watch ‘fall detection’ is monitoring your elderly relatives, you'll keep buying them new Apple watches and iPads to pair them to.)

⁷ <https://medconfidential.org/wp-content/uploads/2019/08/NHS-AI-detail.pdf>

⁸ <https://medconfidential.org/wp-content/uploads/2019/10/business-models.pdf>

⁹ i.e. without any reliance on processing in any cloud (akin to the Apple data privacy model)

In the expectation that NHS tools will receive worldwide use, there will need to be 'booster samples' for some ethnic minorities, so as to provide reference samples sufficient for populations around the world. We note that while this may be extremely difficult for some groups, there are likely to be few others on the planet other than the NHS who could do so – or who would choose to do so.

...more Supply will come, from somewhere... (VC, better products, faster)

The Venture Capital world is pouring money into AI all over the place, but is disproportionately cautious about investing in health services because it expects to be flattened by Google DeepMind. In the broadest terms, ensuring a competitive market for health AI will give others confidence that they will be able to get a reasonable return for a good investment.

Of course, if this doesn't happen in the UK, we will get to watch it happen in the EU¹⁰ or elsewhere – and the NHS could simply buy their products at the prices others set instead, while the UK misses out on all the skilled jobs that could have come with clearer vision.

medConfidential,
September 2020
coordinator@medConfidential.org

¹⁰ <https://www.politico.eu/pro/how-finland-plans-to-kickstart-an-eu-health-data-revolution/>