

AI in healthcare: A perspective on the practicalities



AI in healthcare: a perspective on the practicalities

Max Dillon, Data Scientist
James Thatcher, Data Scientist

At a glance

Recent advances in the technologies underpinning artificial intelligence have led to a rising tide of attention, both positive and negative. Healthcare is commonly cited as an industry that stands to gain from emerging technology, attracting interest from leading researchers around the globe. AI can be the source of a host of benefits to healthcare, resulting in lower costs and improved patient outcomes as it is used to enhance the ability of providers of care¹. However, there is some caution from stakeholders which needs to be managed: a recent survey conducted by YouGov and PwC found that a quarter of patients see no benefit to healthcare from the introduction of AI².

Whilst the media has tended to focus on the theoretical implications of introducing this technology, to date, little attention has been given to the practical steps that need to be taken to ensure that the full benefits are realised. These include: appropriate handling of patient data; robust and reliable infrastructure; and building a culture of trust in emerging technology. These are not trivial tasks. With inspiration drawn from other industries, this report lays out possible steps that those working in healthcare can take to aid a symbiotic integration of this advanced technology into the industry, bringing the most benefit to patients in the UK.

The big picture

The NHS Long Term Plan³, published in January 2019, identifies the need to 'accelerate the redesign of patient care to future-proof the NHS for the decade ahead.' Not only is this necessary to meet evolving health challenges, such as increasing longevity, but also to increase the efficacy of care in the face of tightening budgets⁴.

Technology, specifically artificial intelligence (AI), is often seen as the key to transforming organisations, however this comes with practical risks. In October 2018, the Department of Health and Social Care published their vision for the future of healthcare, which lists privacy & security, interoperability and skills & culture as areas of focus, highlighting areas that have been the undoing of previous initiatives⁵. More recently, the NHS announced the creation of NHSX, a joint unit which will work with the NHS to leverage cutting edge technologies in the delivery of health and care services, whilst also upgrading much of the NHS's existing pre-internet era systems⁶.

This, however, will not be an easy task. Recent research conducted by YouGov and PwC has found that 1 in 4 members of the public see no benefit to using AI in healthcare, highlighting the magnitude of the challenge⁷.

This comes from both misinformation about the role technology will play in healthcare delivery and a subsequent lack of understanding surrounding its merits.

This wariness is not without cause. Embedding emerging technology within healthcare is a complex task, an example being the National Programme for IT in the NHS: announced in 2002 and cancelled over a decade later⁸. This project highlighted the fact that a single solution, one-size-fits-all approach to technological transformation is less effective in a multi-entity organisation like the NHS with differing legal and governance structures. Taking valuable lessons from experiences such as these, the NHS is seeking to leverage initiatives like NHSX and a new principles based code of conduct for the use of AI, setting standards for the development and use of secure and trustworthy AI in healthcare⁹.

AI offers many potential benefits to the NHS, resulting in both better patient care and lower operating costs, but these benefits require careful and coordinated introduction of the technology. In this report we explore the following six areas that can guide practitioners in managing the integration of AI systems into the NHS:

¹ <https://www.wiltonpark.org.uk/wp-content/uploads/Thinking-on-its-own-AI-in-the-NHS.pdf>

² https://pwc.blogs.com/health_matters/2018/10/is-the-public-ready-for-ai-in-health.html

³ <https://www.longtermplan.nhs.uk/wp-content/uploads/2019/01/nhs-long-term-plan.pdf>

⁴ <https://www.kingsfund.org.uk/projects/nhs-in-a-nutshell/nhs-budget>

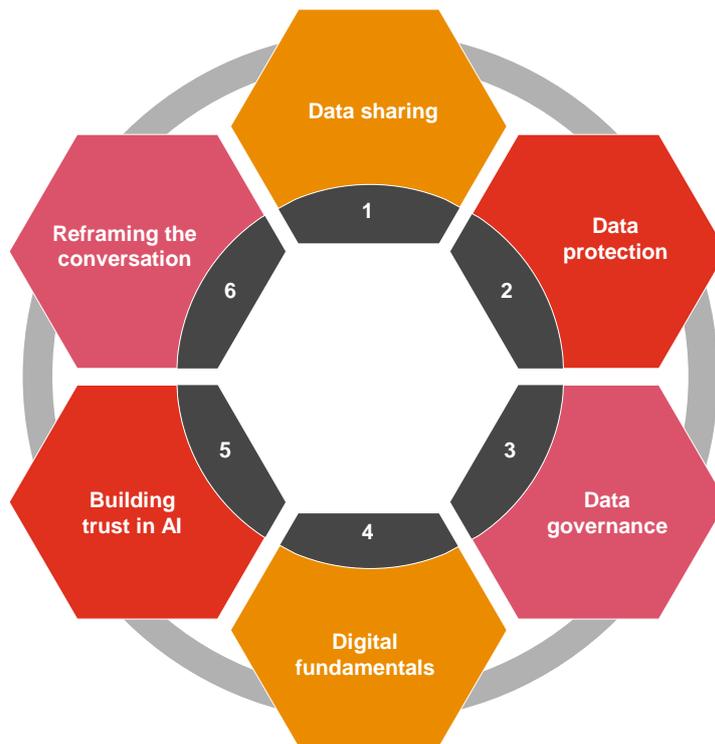
⁵ <https://www.gov.uk/government/publications/the-future-of-healthcare-our-vision-for-digital-data-and-technology-in-health-and-care/the-future-of-healthcare-our-vision-for-digital-data-and-technology-in-health-and-care#guiding-principles>

⁶ <https://www.gov.uk/government/news/nhsx-new-joint-organisation-for-digital-data-and-technology>

⁷ https://pwc.blogs.com/health_matters/2018/10/is-the-public-ready-for-ai-in-health.html

⁸ <https://publications.parliament.uk/pa/cm201314/cmselect/cmpubacc/294/294.pdf>

⁹ <https://www.gov.uk/government/news/new-code-of-conduct-for-artificial-intelligence-ai-systems-used-by-the-nhs>



How can professionals practically manage data?

A key component in capitalising on emerging technologies such as AI and robotics is the data itself – however more technologically mature industries have often overlooked the practical steps to safely capturing, storing and transferring this data. This is due to a combination of factors including: the ‘move fast and break things’ mantra of technology development from the late 2000s; industry regulators (where they exist) playing perpetual catch-

up with accelerating tech growth; and a lack of understanding or appreciation of the risks inherent with statistically driven insight. This attitude cannot be adopted by healthcare as it improves its technological capabilities. Previous lapses in diligence involving patient data have warned the public of the need for strong oversight and communication around who is using their data and the purposes it is serving¹⁰.

‘Technical debt’ is a term used in software development to refer to the cost of future work that will be necessary due to making

quick fixes today that are not robust in the medium-term. we propose that there is a related concept that applies to the data itself: ‘data debt’. Opting for quick fixes at the expense of quantity or quality of data captured accrues data debt that will need to be disentangled later. The following three sections give practical examples of how to avoid a build up of data debt, helping the long term performance of the individual organisation and the wider healthcare system it sits within.

1. Innovative methods for sharing data

Advances in AI have come about in the last decade thanks to major advances in computing hardware and the availability of digitally represented data. As an organisation, the NHS stores vast amounts of data about both patients and procedures. However, this is housed in disparate locations on different systems, contributing to low interoperability and long wait times. If the profound benefits of AI and related technologies are to be extended to healthcare in the UK, there needs to be integration and consistency in data and storage systems. Here, we offer two ways in which trusts can do this:

Integrated datasets

Over the past five years, Kent County Council and the seven Kent Clinical Commissioning Groups (CCGs) have collaboratively created the Kent Integrated Dataset (KID). This is one of the largest integrated healthcare databases in the country, containing de-personalised data from GPs, mental health, social care and public health services. Impressively, it has been used to inform both individual and public health strategies and decisions, including forecasting future service pressures and permitting controlled trials on the impact of certain programmes. KID has been built from existing systems distributed across orphaned groups and can act as a blueprint for other districts to follow, or indeed as a proof of concept for integrated datasets on a larger scale.

¹⁰ <https://www.bbc.co.uk/news/technology-40483202>

1. Innovative methods for sharing data

Data sharing agreements

An alternative arrangement to an integrated dataset is a multilateral data sharing agreement like the Care Information Exchange (CIE)¹¹: a platform that puts patients in control of who sees their data. Currently being trialled in Northwest London with 20,000 patients, CIE uses Application Programming Interfaces (APIs) to share data between care providers across the territory¹².

One groundbreaking feature of the CIE is the web interface patients can access to view their patient records, their appointment information, their test results and to control who this information is being shared with. Learning from the previous NHS care.data scheme which shared patient data with commercial companies without the consent of the patient, this is a key offering of the CIE¹³. It has the impact of increasing transparency, patients' control over their data: generating public trust in the system.

2. Data Protection Protocol

GDPR and the UK's Data Protection Act, like preceding regulation, categorise health data as 'special category' personal data. As data controllers processing special category health data, NHS trusts must meet these requirements, applying necessary technical and organisational measures to protect the sensitive data they process on an ongoing basis. These requirements also apply to any private organisations that are processing patient data, whether or not it is being used for a commercial purpose.

At a practical level, the NHS' Data Security Centre is available to help health and care organisations manage cybersecurity risk by offering services like training, monitoring, assessment and advice on how to manage risk and recover swiftly in the event of a breach¹⁴. On top of this, the NHS Information Governance Alliance Code provides a guide for practitioners managing patient records, but a standard protocol for navigating the data processing pipeline is yet to be introduced¹⁵.

There are examples of organisations developing components of such a protocol. The Verifiable Data Audit (VDA)¹⁶ is a method developed by DeepMind Health that acts as a digital ledger in a similar fashion to block chain, making a log of every interaction with the data, detailing what operation was performed on the data¹⁷. It is hoped that this system will increase transparency around what data is used for, eventually being opened up to the public so they can see exactly how their personal records are utilised.

3. Effective governance of patient data

It will not be possible to democratise data throughout the NHS without careful, coordinated governance. To assert this, the south east London NHS Sustainability and Transformation Partnership has appointed a Chief Clinical Information Officer (CCIO)¹⁸. This is a key appointment, as many attempts at integrating new technology into the NHS fall short through a lack of clinical engagement. This appointment signals the desire for coordination and consistency in incorporating new technology.

This is a key learning from previous technological transformations and a potential benefit of the more decentralised structure of the NHS we see today. It will be key in leveraging the data available to health and care organisations.

There is a trade off that needs to be made between making it easy for practitioners and researchers to transfer data between organisations, but making sure there are sufficient processes in place to protect the subjects of the data.

There is cause for optimism: NHS Digital recently announced the UK Health Data Research Alliance, which will bring together health data specialists and leading academics to provide better quality health data for more open research¹⁹. The sooner an equilibrium between openness and security is reached, the greater the long term benefits to patients and the healthcare system will be.

¹¹ <https://www.patientsknowbest.com/cie.html>

¹² <https://www.healthylondon.org/our-work/digital/london-health-care-information-exchange/>

¹³ <https://www.wired.co.uk/article/care-data-nhs-england-closed>

¹⁴ <https://digital.nhs.uk/services/data-security-centre#summary>

¹⁵ <https://digital.nhs.uk/data-and-information/looking-after-information/data-security-and-information-governance/information-governance-alliance-iga>

¹⁶ <https://deepmind.com/blog/trust-confidence-verifiable-data-audit/>

¹⁷ <https://pwc.blogs.com/data/2018/09/blockchain-the-safest-place-for-our-data.html>

¹⁸ <https://www.greenwichccg.nhs.uk/News-Publications/news/Pages/Better-sharing-of-care-records-offers-people-faster,-safer-treatment.aspx>

¹⁹ <https://digital.nhs.uk/news-and-events/latest-news/new-uk-health-data-research-alliance-to-boost-medical-research-and-improve-future-health-and-care>

Getting the fundamentals in place

It may appear obvious, but a key component in any successful rollout of technology is having the ecosystem in place to support it. In the rush to make scheduled release dates, corners can be cut and aspects overlooked. This section highlights two components pertinent to the health sector.

Security

With data sharing between organisations within the NHS becoming more commonplace, comparisons can be drawn between health and financial data processing: a high volume of data; a high volume of transactions; and highly sensitive information. Whilst the finance industry has been developing its information security environment for decades, this has not been a deterrent to attackers: cyberattacks are estimated to cost each US financial services firm \$18m annually²⁰. To date, financial service firms have received the vast majority of cyberattacks, but these have been getting more numerous and more sophisticated each year. For example, the Financial Conduct Authority were notified of 69 attacks in 2017, compared to 24 in 2015²¹. It is prudent to assume that the sophistication of the attacks will be transferred across industries, so healthcare will have to adopt the latest defense strategies swiftly. Taking best practices from security experts will help shield the NHS, a large target for nefarious agents.

The importance of taking simple preventative measures against malicious attacks was highlighted by reports stating that recommended software updates to IT systems could have prevented the WannaCry attack of 2017, saving the NHS an estimated £92m^{22,23}.

Infrastructure

A common problem when releasing new technological solutions is not preparing for, or underestimating the demand for the new service. As part of the Affordable Care Act released in 2013, the US government launched www.healthcare.gov, the centralised exchange for comparing and selecting health insurance policies²⁴. 34 states elected to use this service, but on the first day of going live the system was overwhelmed and had to be shut down, with President Obama stating the demand '... exceeds anything that we had expected'.

It took around 8 weeks to remedy the issues, with reports that the site had received 50 times more users than it could handle on the opening day²⁵. It is critical that new systems are stable upon release as this builds trust with users, but also anchors their wider expectations about the use of technology in healthcare. It is arguably better in the long run to delay the release until sufficient tests have been passed, especially in the face of last minute design changes as were reported with www.healthcare.gov.



²⁰ <https://www.forbes.com/sites/bhaktimirchandani/2018/08/28/laughing-all-the-way-to-the-bank-cybercriminals-targeting-us-financial-institutions/#657bb48c6e90>

²¹ <https://www.mortgagesolutions.co.uk/news/2018/02/21/fca-stats-reveal-rise-data-hacking-financial-services/>

²² <https://www.england.nhs.uk/wp-content/uploads/2018/02/lessons-learned-review-wannacry-ransomware-cyber-attack-cio-review.pdf>

²³ <https://www.telegraph.co.uk/technology/2018/10/11/wannacry-cyber-attack-cost-nhs-92m-19000-appointments-cancelled/>

²⁴ <http://www.healthcare.gov>

²⁵ <https://www.nbcnews.com/politics/politics-news/stress-tests-show-healthcare-gov-was-overloaded-flna8C11548230>

Building Trust in AI

Artificial Neural Networks (ANNs), the driving force behind most recent advances in AI, have proved effective in a range of tasks such as cancer detection, hospital readmission prediction and machine translation. Increasingly complex architectures are driving improved performance, however as algorithms become more complex, so does the task of understanding how specific outputs were reached.

Governance regarding the explanation of machine decision making is critically important from both a clinical and an ethical perspective. It doesn't feel acceptable for a healthcare professional to take the output of an automated system and use that as the sole justification for a decision regarding a patient. Indeed, this is not permitted under GDPR.

This is where the concept of having a 'human in the loop' is gaining ground. This is seen as a paradigm shift away from that of full autonomy: a machine learning system is put in place, but the output is reviewed by a human expert before being acted upon. This has been in use for decades in aviation with autopilot systems, but is now being explored in healthcare in areas such as ICU care where recommendations made from interpreting the data of multiple monitoring machines, previously assessed in disparate fashion, will be presented to clinicians in real-time²⁶.

Furthermore, a key part of this strategy that must be clear to all stakeholders at the outset is the intended use of the technology: which functions it is, and is not suitable for. This is important in gaining the trust of not only patients, but critically, medical professionals increasingly required to use these technologies, and making sure that expectations about the technology match the reality. It is crucial that those using the technology are given the correct level of training and understand that the technology is there to aid and not replace.

The NHS has indicated an awareness of this with an April 2018 announcement launching a review into how workers can be upskilled to capitalise on the new technology²⁷.

Reframing the conversation

The relationship between humans and technology is commonly framed as adversarial, with one group losing out. However, rather than make humans redundant, emerging technology within healthcare is being used to either: perform the easier, time consuming tasks so that more time can be spent on the multi-faceted, abstract tasks like patient care or medical research²⁸; or to aid humans by supplementing their work, for example by aiding radiographers with cancer detection²⁹. Both of these lead to better patient outcomes whilst avoiding a detrimental effect on the medical labour market.

Moreover, research carried out by PwC as part of the 'Will Robots Really Steal Our Jobs?' study shows that the roles within healthcare that are most at risk of automation in the UK are those making up the smallest percentage of the workforce³⁰. In particular, clerks, cleaners and helpers were found to be the jobs containing the largest number of automatable tasks with automation risks of 60% and 40% respectively, but cumulatively make up less than 15% of the workforce. At the other end of the spectrum, personal care workers make up about 40% of the workforce, but have an automation risk of less than 20%. This ties in with the point above regarding humans in the loop: the technology can be used to enhance interpersonal caregiving, not replace it. It also highlights the importance that the distribution of jobs across the sector has in the conversation surrounding automation: it is crucial in understanding the total impact emerging technologies are expected to have.

Automation risk across occupations in UK Healthcare



²⁶ https://spectrum.ieee.org/biomedical/devices/ai-could-provide-momentbymoment-nursing-for-a-hospitals-sickest-patients?utm_campaign=Data_Elixir&utm_medium=email&utm_source=Data_Elixir_202

²⁷ <https://www.gov.uk/government/news/review-announced-into-training-nhs-staff-to-use-ai-and-robotics>

²⁸ <https://deepmind.com/blog/ai-uclh-radiotherapy-planning/>

²⁹ <https://deepmind.com/blog/moorfields-major-milestone/>

³⁰ <https://www.pwc.co.uk/economic-services/assets/international-impact-of-automation-feb-2018.pdf>

Of equal importance to this conversation is the discussion around job creation: to estimate the net effect that automation will have on the labour market, we must attempt to quantify the roles that will be required but do not exist today. Research conducted by PwC in 2018 found that the sector containing health, public services, education and recreation would see the largest uplift in global GDP by 2030, with a 21% rise assuming all other variables follow their projected trajectory³¹. Whilst it is impossible to say for sure what jobs will be created, it is likely that the increased demand associated with a rise in GDP will lead to an increased demand for labour in some areas of healthcare, whether that be logistical, operational, or managerial. This angle is often missing from media reports surrounding artificial intelligence, however incorporating this into the narrative will help balance public portrayal of the AI's use in healthcare.

A perspective on the practicalities

The benefits of integrating new technologies into healthcare are numerous, profound and well documented, ultimately leading to improved patient care. However the optimism surrounding this potential often forgets the reality of embedding the technology into existing systems, both technical and societal. Whilst the challenges are numerous, this report presents examples of measures that medical practitioners can take to support the responsible stewardship of AI into healthcare, helping to release the potential sooner, to a wider audience.

Acknowledgements

We would like to thank Olivia McDonald and Arnav Joshi for their help in producing this report.



³¹ <https://www.pwc.co.uk/economic-services/assets/macroeconomic-impact-of-ai-technical-report-feb-18.pdf>

This document has been prepared only for techUK and solely for the purpose and on the terms agreed with techUK. We accept no liability (including for negligence) to anyone else in connection with this document, and it may not be provided to anyone else.

© 2019 PricewaterhouseCoopers LLP. All rights reserved. PwC refers to the UK member firm, and may sometimes refer to the PwC network. Each member firm is a separate legal entity. Please see www.pwc.com/structure for further details..

190529-195122-OM-OS