DAC BEACHCROFT

AI IN HEALTHCARE

Transforming the UK's health system

CONTENTS

Healthadviser

DAC Beachcroft's Health Adviser publications seek to provide insight, foresight and thought-provoking features and articles that provide practical solutions for the issues of the day, for health and social care professionals.

FOREWORD

We are at the start of a medical revolution. The NHS alone already uses more than 130 artificial intelligence (AI) based products to diagnose or treat 70 different conditions. There is no doubt that AI is going to play an increasing role in the future of medicine.

The potential reservoir of countrywide data available through the NHS means we have the opportunity to deliver effective, safe AI-assisted healthcare of the highest quality.

However, the nature of AI brings with it novel challenges to the provision of healthcare. It raises vital and fundamental issues where law, medicine, technology, business and ethics come to the debating chamber to argue how best to save and improve lives. It forces us lawyers to think afresh about the application and interpretation of contractual rights and duties, data management, confidentiality, regulation, medical liability and public law issues.

Through the use of algorithms, AI deploys machine and deep learning to bring intelligence to applications. Based on the analysis of large bodies of data, AI will provide fast and reliable solutions to healthcare problems: that is the goal. If developed and used in the right way, it will transform the way healthcare is delivered.

We are already at a tipping point. The speed with which new AI-based technology is coming to the health market is testing the UK's traditional regulatory and legal frameworks. The tension is there.

Meanwhile, patient buy-in is essential – the public needs to trust that people's rights and data are properly protected, but equally that their health is also safeguarded. But many patients mistrust the technology. They fear the unknown and worry that AI will drive compassion and humanity out of their medical care.

Yet change is inevitable. The Covid-19 pandemic has forced us to see the world in a different way and adapt. During the pandemic and driven by the need to avoid hospital attendance, cystic fibrosis sufferers who previously had to attend regular hospital appointments for assessment are now successfully and better treated at home through AI devices, which constantly monitor their condition and provide early diagnosis of any deterioration. It is a small but important example of how patients have seen first-hand how AI can benefit their care.

Meanwhile, the EU has recently published proposals for regulations for AI including its use in healthcare. The proposed regulations are based around the evaluation of risk – generally AI healthcare products will fall into the 'high-risk' category and be subject to high-risk management and quality management systems.

Will the UK mirror the EU proposals or provide a more bespoke approach to healthcare products?

At DAC Beachcroft, we believe that the time is ripe for a fresh look at the law and how it is applied to AI across a range of interconnected issues.

Our philosophy is simple. We believe that AI in healthcare calls for a holistic, coordinated approach, in which the various interested parties work together for the common cause.

The term 'AI' is often used loosely to cover a variety of technology-based products. There is no accepted definition of what AI is. What we want to do with this thought leadership report is to show you some of the legal issues that arise from software systems that are self-running, self-enforcing, that can operate independently of their creators and operators and that are equipped with adaptive, learning abilities.

DAC Beachcroft has sought out the views of key voices in the future of AI healthcare. We have spoken to clinicians, regulators, developers and investors. Even other lawyers. On their own, each of our contributors would provide valuable insight into AI but bring them together and they give a kaleidoscopic overview of the potential that AI brings to the delivery of healthcare. Through their insights, we examine the ethical and legal challenges that the development and use of this radical new technology presents and offer practical advice and some potential solutions.

So we would like to say a big thank you to our contributors: Dr Junaid Bajwa (Chief Medical Scientist at Microsoft Research); Dr Nicola Byrne (National Data Guardian for health and adult social care in England); Dr Peter Feldschreiber (Barrister at 4 New Square chambers); Sarah-Jane Green (Head of Artificial Intelligence Regulations and Policy at NHSX); Johan Ordish (Head of Software and AI - Innovation Devices Division, at the Medicines and Healthcare products Regulatory Agency); and Seb Wallace (Investment Director at Triple Point).

Darryn Hale
Partner at DAC Beachcroft

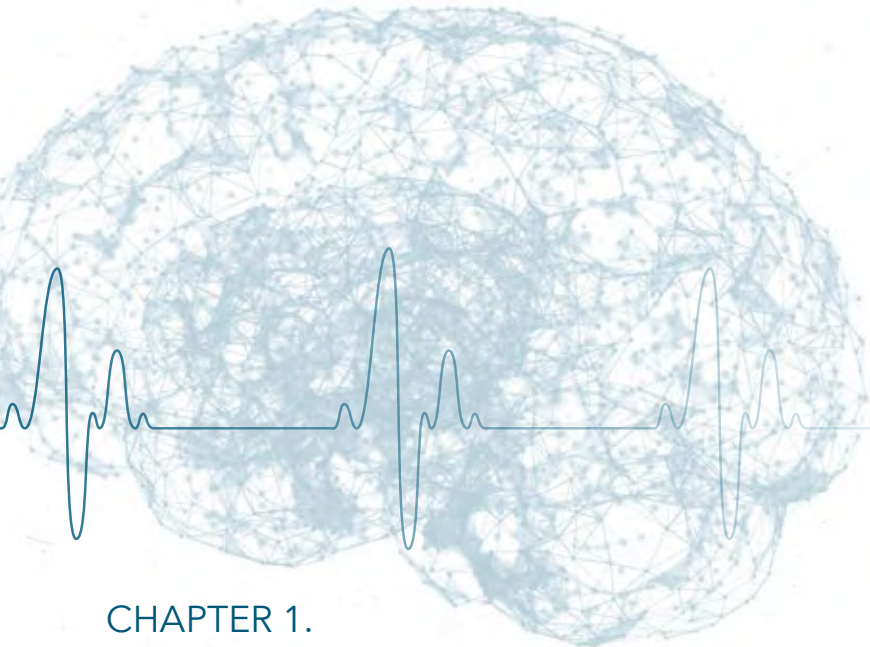

CHAPTER 1.

THE INNOVATOR AND THE INVESTOR

AI-related healthcare is one of the fastest growing areas for investment and has the potential to innovate and transform the way healthcare is delivered. Different roles need to work together towards achieving the same goal – a healthcare system that will meet demands for decades to come.

As Chief Medical Scientist at Microsoft Research, Dr Junaid Bajwa acts at the intersection of medicine and technology, as both clinician and technology innovator. Amongst his other duties, he works with a multidisciplinary team tasked to consider how they might transform the practice of medicine with trusted, reliable human-centred AI.

By way of example, he cites the NHSX award-winning 'Project Inner Eye' that uses state-of-the-art machine learning technology to build innovative tools for the automatic, quantitative analysis of three-dimensional medical images.

A typical (human) radio-oncologist, he explains, would design a treatment plan for a patient based on mapping out tumour segment images, and this process could take anywhere between 30 minutes and potentially three hours to complete.

Researchers at Microsoft and Cambridge University Hospitals, leveraging Project Inner Eye Technology, have found that AI can augment and accelerate clinicians' ability to perform radiotherapy planning 13 times faster, meaning that waiting times for starting potentially life-saving radiotherapy treatment can be dramatically reduced.

To build trusted, reliable AI tools, at least three ingredients need to exist: access to data, access to domain expertise and compute power. Regarding access to data, Dr Bajwa likens access to raw health data to access to crude oil: by itself, crude oil holds little utility, it needs to be refined and go through a series of processes to create something more useful or valuable (for example, plastic, petrol, etc). Likewise, raw health data by itself holds little utility, it needs to be cleansed, refined and contextualised to transform it into something of insight (use), to ultimately lead to better health outcomes (value).

"With great data, comes great responsibility."

Dr Junaid Bajwa, Chief Medical Scientist at Microsoft Research

As such Dr Bajwa is clear about the obligations that surround data collection and use. He suggests that "with great data, comes great responsibility", and the importance of all parties as custodians of health-related data, engaging with and following appropriate regulations, such as GDPR, and having their own appropriate information governance procedures in place.

In the UK, there continues to be a need to have a wider debate across society on the importance of these data-driven issues. The Wellcome Trust has done some great work in this space under the banner of 'understanding patient data', where it worked with patient groups, charities, the NHS and policymakers to bring transparency, accountability and public involvement to the way patient data is used – but more needs to be done to ensure trust, consistency and the adoption of any standards across the wider UK health and social care ecosystem.

In healthcare, the value of AI is strongest in the context of the human activity it supports. We need to think of how AI tools can act as enablers, built into a workflow that makes sense. Medicine is both an art and a science, and AI can augment the superpowers that clinical teams, and indeed non-clinical teams, have to help drive precision, accuracy and potentially productivity (amongst other benefits) – but the tool will not substitute the care and compassion needed to ultimately have impact in someone's life.

We are on a journey towards precision medicine, which will be enabled by connected care, precision diagnostics and precision therapeutics – all leveraging AI to different degrees. This will take time and the collective effort of payers, providers, regulators, policy makers, alongside appropriate public engagement, to make this a reality.

Lessons from the private sector

Seb Wallace of Triple Point echoes many of the themes mentioned by Dr Bajwa, but from a different vantage point – that of the investor looking at investing in AI products. Triple Point is an active healthcare investor through its early stage B2B technology venture fund. Healthcare is a core sector focus for it, with the fund having backed medical data AI diagnostics businesses, electronic health record (EHR) software and credential and passporting medical HR software.

Wallace comments that what he often finds in his dealings with the NHS is a drive to innovate, but a patchy aversion to private sector innovation, summed up in some cases as 'public sector good, private sector bad'.

"Obviously, it's good for patients that they don't pay for healthcare," he says. However, he argues that "underneath the patient, it should be irrelevant to the end user whether or not part of the supply chain is sourced from the private sector".

Wallace sees some NHS bias against the private sector as a reflection of the politicised 'privatisation of the NHS' debate which, at its most basic level, conflates private healthcare providers operating within the NHS with patients directly paying for healthcare, like a traditional private health system requires.

"Clearly these two outcomes are different," Wallace adds. Instead, he believes price and efficiency should be the key determinants in NHS procurement. "What the NHS should be exploiting is the system's scale and ability to drive economies of scale. When it comes to software, and NHS Digital's remit," he asks rhetorically, "the NHS doesn't make drugs or hospital beds, so why should it make software?"

Asked what he looks for in an AI start-up, top of Wallace's requirements is a team with technical expertise, for example, people with a data science background who also understand how to commercialise in different health systems, such as the NHS. "It's one thing building the product but quite another selling it, particularly in the UK."

"To get new technology approved in the UK can take months or years. In a start-up world, that is bad for business. It can demotivate the team or impact the return for investors."

Seb Wallace, Investment Director, Triple Point

Wallace has worked with several UK health tech companies, which started selling their products to the NHS but eventually pivoted to the US, where sales were easier and procurement less politically charged. The companies then moved their HQs stateside. If you can sell into the NHS, you will grow your business in the NHS first and stay in the UK. "As investors, we like to see homegrown talent flourishing; but the NHS clearly has a problem with procurement. If it didn't, you wouldn't see healthcare start-ups leaving for other countries," he explains. "The NHS is one of the harder systems to sell into for innovative new technology providers. That's the case even when tech founders are NHS-trained doctors."

In terms of due diligence, Wallace and Triple Point will start by looking at the start-up founders' backgrounds and talking to customers, which may include the NHS. They will perform selective audits of the technology, but this is not always straightforward. "Often AI models are 'black box', so it's difficult to diligence the AI models themselves. You can't see what is inside them when they synthesise the data in as little as quarter of an hour.

Instead, what helps is to have articles in respectable journals, alongside randomised controlled trials (RCTs). Essentially the team is looking for evidence that shows a product's application in healthcare, rather than just a new idea which could be used in health.

Innovations come at the cost of mistakes, but you can't have mistakes in healthcare. Given this, Wallace feels that there is a need to create sandbox environments where people can access data and make mistakes in controlled ways; innovating while ensuring patient safety. Wallace points to London's Chelsea and Westminster NHS Trust for what can be achieved. "Chelsea and Westminster has been particularly innovative in engaging with the private sector and helping to foster new technology adoption." He encourages other trusts to be as innovative.

When looking at medical data, the UK has a head start. "The exciting thing here," Wallace explains, "is that the UK has by far some of the best, centrally maintained medical data in the world. It's a single player with hundreds of bodies. If you get access to that data, you can train AI to do what pretty much no other country's datasets can."

But there is resistance to data sharing among the public. This unwillingness is to some extent based on a poor understanding of data sharing and what it means in practice. "If we make it hard for tech in the UK to access some of the most comprehensive health data in the world, we hand the keys to the next generation of companies in the US, where they will find ways to make it work. The US is much more forgiving about sharing data than Europe is. We need to support data sharing with new, cutting-edge healthcare businesses. Doing this could be the key to building the next generation of global health tech unicorns – and improving patient outcomes in the process," Wallace explains.

When people say 'personal data' they think of it as details attached to them. "People imagine their name and address being shared next to their medical condition," he outlines. "In fact, shared data is anonymised or pseudonymised. The media and the NHS could do better at explaining the distinction. It's just a matter of education. The public needs to be provided with more information and context about the benefits of data sharing and AI generally.

"For example, there are some who see AI as seeking to replace clinicians. This is not the case. The idea is to help overworked doctors and nurses, so they can focus on the things that matter. AI is really about tech doing the grunt work."

"Industry, and that includes its funders like me, need to improve our explanation and communication when it comes to why these things [AI products] are improving healthcare. If there were better communication, many people would be comfortable with what we want to achieve," he adds.

Wallace finishes by saying that there needs to be an improvement to the process used for regulating new healthcare technology. "To get new technology approved in the UK can take months or years. In a start-up world, that is bad for business. It can demotivate the team or impact the return for investors."

Wallace's message is simple: reduce regulator processing time, not rigour, and you reduce the cost of the product. "If the MHRA (Medicines and Healthcare products Regulatory Agency) aims to be a globally recognised stamp of approval like the FDA (Food and Drug Administration) in the US, it will help UK tech. If you get FDA approval, other nations' regulators look favourably upon you – even if they may not publicly admit it. But to get to that point, we need an MHRA approval process that is pragmatic, commercial and robust. Over time it can become a gold standard."

Perhaps Wallace's position is best summed up as the NHS being close to unleashing the potential of UK health tech, but not quite being there yet. It has the data, but it needs to be more fluid and flexible. Without that, the UK is unlikely to create the cutting-edge technology of the next decade.

"Rather than having to buy off the US private innovators, why not support the private innovators here?" he wonders.

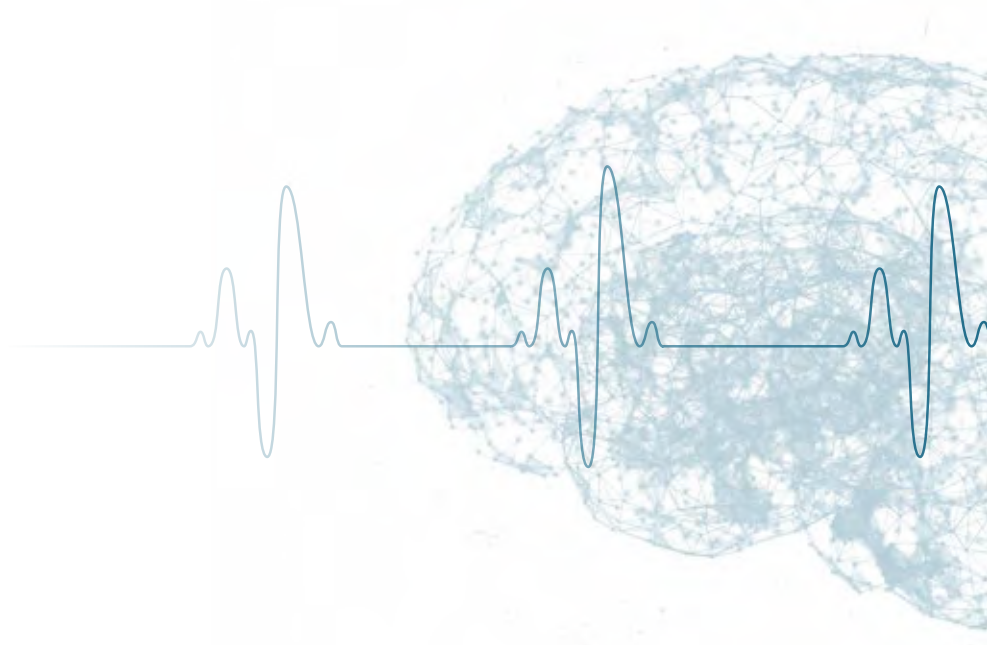

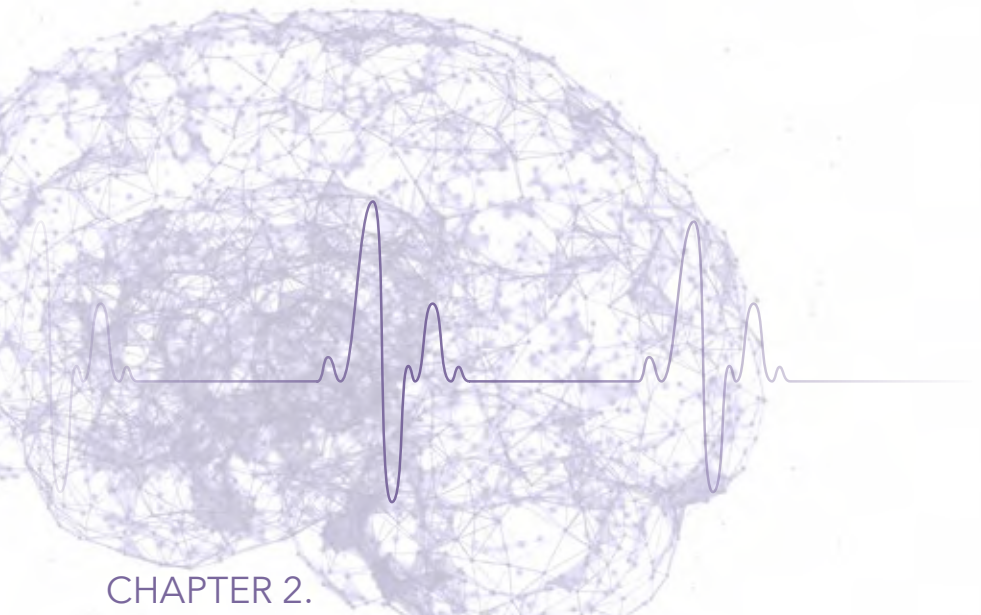

CHAPTER 2.

HANDLING DATA WITH CARE

There are many challenges when it comes to the use of data in AI such as anonymisation and transparency. Ensuring society is comfortable with an AI-assisted future in healthcare is paramount.

The variety and breadth of potential deployments of AI in healthcare mean that the data protection challenges, and solutions, are not uniform – they are fact and context-specific depending on the nature and volume of data being used, and the purpose(s) for which that data is to be used. For example, is identifiable data even needed at all, or could the same purpose be achieved with fully anonymised data (whether dummy or synthetic)? Does the AI aim to drive individual or population-level decision-making, such as risk stratification for individual health interventions versus planning healthcare resources across a particular geographical footprint? Relatedly, does the AI act as the sole decision-maker or is it designed to support decision-making by clinicians?

Is the AI genuinely self-learning and, if so, how easily and transparently can it be explained to patients? This potential dissonance between the practical operation of AI and individual data protection rights is a key area of concern.

Statutory data protection laws, most notably the UK GDPR and Data Protection Act 2018, are necessarily generic in their application and so it is not always straightforward to apply them to the unique factors which relate to the health sector. In addition, there are particular legal challenges relevant to health that do not necessarily arise elsewhere – most notably the common law duty of confidentiality. This imposes additional restrictions on the use of health data, especially in respect of any such uses unrelated to an individual's direct care. In those circumstances the patient's consent is typically required before their data is used unless an alternative overriding justification can be found (such as an overriding public interest).

The potential dissonance between the practical operation of AI and individual data protection rights is a key area of concern.

It is clear there are a vast array of legal issues and obligations arising under data protection law, and we intend to focus on key challenges, as we see them, and to offer our assessment on how they can be addressed. In order to do so it is important to have the following fundamental legal concepts in mind:

1. Data protection law only governs 'personal data', meaning information from which a living individual can be directly or indirectly identified (including through a combination of information). It does not, therefore, apply to deceased individuals although broader duties of confidentiality still place restrictions on the use of such information;

2. Legal responsibility is attributed to controllers, those directly responsible for determining why and how to use personal data, and processors, those who are simply instructed to use personal data in a specific way by controllers. This is likely to be a particularly live issue in AI-based collaborations across public and private sector organisations;

3. All uses of personal data must satisfy a number of data protection 'principles', including that it is used lawfully, fairly and transparently; and

4. Simply complying with data protection law is not enough – it must also be possible to demonstrate how compliance has been achieved.

It is against that backdrop we highlight particular challenges below. Clearly, enforcement for non-compliance is a hugely significant concern but the focus is on how to get it right and not what happens if it goes wrong. That said, however, it would be remiss to ignore the regulatory and enforcement angle entirely; unlike other areas explored in this piece data protection liability is relatively predictable in the sense that controllers and/or processors who fail to discharge their legal obligations could face fines from the ICO (Information Commissioner's Office), claims from affected individuals, or a combination of the two.

Anyone operating in the AI space in healthcare should think beyond the minimum steps required in order to comply with data protection law and instead ask, 'how can I make this understandable to patients in a way that will enable them to place their confidence in it?'

Transparency

Given the specificity of the challenges applicable to the health sector, the perspective of the National Data Guardian (NDG) on the issues facing AI in healthcare is vital to our understanding of the correct approach. The NDG is a statutorily appointed post, enabled to issue guidance to which all persons (whether acting in a public or private sector capacity) must have regard whenever they use NHS data. The current NDG is Dr Nicola Byrne and she emphasises that developers need to proceed with extreme caution when handling any identifiable health and care data.

One of the most critical aspects to get right when using health data is transparency. There are specific obligations under the UK GDPR to inform individuals about how their personal data will be used. However, anyone operating in the AI space in healthcare should think beyond the minimum steps required in order to comply with data protection law and instead ask, 'how can I make this understandable to patients in a way that will enable them to place their confidence in it?'. Negative publicity can often be the death knell of any initiative and ensuring that patients are brought along with you as part of the journey is critical.

"Boundaries between different public sector organisations and government departments will also need to be consciously maintained; access for one purpose within health and care might lead to unanticipated findings or uses outside, which potentially could undermine trust in a confidential health and care system, if data is used in ways the public does not expect or necessarily support."

Dr Nicola Byrne, National Data Guardian for health and adult social care in England

This is a point emphasised by the World Health Organization's (WHO) guidance on Ethics & Governance of Artificial Intelligence for Health[1], noting that information about a particular deployment of AI 'should facilitate meaningful public consultation and debate on how the AI technology is designed and how it should be used'. It is not just about publishing information to demonstrate that transparency has been ticked off, but instead actually explaining in easily digestible terms what the technology is, how it works and what it means for individuals.

The work of the NDG directly supports the fundamental importance of public transparency, following some public dialogue work last year in collaboration with Understanding Patient Data into views on public benefit assessment[2]. The key messages from that work were:

1. The public want health and care organisations to be ambitious for data use for better planning, research and innovation, but conditions of transparency and authentic public engagement in decision-making about data access must be met;

2. Confidential identifiable data must be handled with utmost care; and

3. Benefits must be fairly distributed, with public benefit outweighing any profit.

The public and healthcare professionals alike are significantly concerned about the role of commercial companies – who will 'get hold of' their data, how secure will it be (who will potentially be able to see it?) and who will be making a profit – but equally the development of AI, and the potential for significant improvements in patient outcomes which go with it, is critically dependent on them. As Dr Byrne notes: "Those issues need to be transparently and actively addressed on a case-by-case basis for people to be able to challenge and potentially be satisfied that commercial involvement is necessary, the safeguards around use are sufficiently strong and that any commercial profit will not be disproportionate to the benefits the data use will bring to the public."

Dr Byrne adds: "Boundaries between different public sector organisations and government departments will also need to be consciously maintained; access for one purpose within health and care might lead to unanticipated findings or uses outside, which potentially could undermine trust in a confidential health and care system, if data is used in ways the public does not expect or necessarily support."

Even a perception that data is being shared outside the confines of a specific hospital or technology provider could be enough for the public to lose confidence in it and, in turn, undermine its credibility.

This issue is also firmly on the radar of NHSX, the government unit with responsibility for setting national policy and developing best practice for NHS technology, digital and data. Sarah-Jane Green, Head of Artificial Intelligence Regulations and Policy, advocates a similar message.

[1] Ethics and governance of artificial intelligence for health www.who.int/publications/i/item/9789240029200

[2] Putting Good into Practice: A public dialogue on making public benefit assessments when using health and care data www.gov.uk/government/publications/putting-good-into-practice-a-public-dialogue-on-making-public-benefit-assess-ments-when-using-health-and-care-data

NHSX supports innovation, including those built around the use of data, "but as an industry and as the health sector, we need to be super clear what data we are being asked to share and why because I am not sure that I would want my data to be used to train an algorithm when it has confidential medical information about me if it is not anonymised".

She adds that the NHS "needs to make it clear why we are collecting that data and that it will be anonymised. We need to help more people share their data so it will be easier to train these algorithms".

This public discourse is critical and we very much welcome the NDG's intention to publish draft public benefit guidance for consultation, which aims to be a useful tool for organisations setting out to build and maintain public trust in their use of data for purposes beyond direct care. This is a very important first step towards building public confidence and consensus around the use of health data in an AI context.

Automated decision-making

The UK GDPR imposes specific restrictions and safeguards in respect of decision-making based solely on automated processing of personal data. In particular, individuals have the right not to be subject to decisions based on automated processing, which produce 'significant effects' on them. It is easy to see how this could be the case in the context of AI deployed in healthcare, for instance if an algorithm were to be used in order to assess whether a particular individual requires a specific healthcare intervention. Further, even if the relevant decision-making can take place permissibly in accordance with the UK GDPR then specific safeguards are required, including the right of an individual to obtain human intervention. This runs the risk of undermining the utility of the AI, if its purpose is to alleviate the pressures on clinicians by automating otherwise lengthy or time-consuming tasks.

Green notes that there are different types of AI: "There is the AI that has sat in the computer basement for years, used for diagnostic purposes as an assistant to the physicians and there is AI that has the ability to make decisions. What we need to be mindful of here is whether or not it considers the element of 'meaningful human control'. That terminology comes from the weapons sphere but it's also applicable in the healthcare space because it's all about clinician and patient trust and regulation.

"If your AI is acting fully autonomously and decisions are not being reviewed by a clinician, there are questions of liability and accountability and what about risk? Are there multiple risks or is there less of a risk because you are removing human error? We have research programmes looking into these areas."

This emphasises the importance of the functionality of AI and developers understanding this by reference to the relevant legal framework. It is much more difficult to retrospectively mould technology to avoid transgressing legal and regulatory hurdles. There are also ethical considerations to bear in mind, as per the WHO ethics guidance, which focuses on protecting autonomy by ensuring that we do not design humans out of critical decision-making in a healthcare context. Clearly, anyone developing AI needs to be aware of these issues and we are therefore encouraged that work is on-going to bring together key organisations, including NICE (the National Institute for Health and Care Excellence), the MHRA and the ICO, into the Multi Agency Advice Service (MAAS)[3], to ensure there is clear guidance for organisations both developing and implementing AI systems.

"If your AI is acting fully autonomously and decisions are not being reviewed by a clinician, there are questions of liability and accountability and what about risk? Are there multiple risks or is there less of a risk because you are removing human error? We have research programmes looking into these areas."

Sarah-Jane Green, Head of Artificial Intelligence Regulations and Policy, NHSX

[3] The multi-agency advice service (MAAS) www.nhsx.nhs.uk/ai-lab/ai-lab-programmes/regulating-the-ai-ecosystem/the-multi-agency-advice-service-maas/

Anonymisation

One way to avoid the myriad obligations of data protection and confidentiality law is to use fully anonymised data. A new proposed duty in the Health and Care Bill would impose an obligation on health and care organisations to share anonymised data for the benefit of the system as a whole and, in turn, the widespread availability of such data could potentially accelerate the development of AI.

However, it is fair to say that anonymising data to legal standards is easier said than done. Pseudonymised data, which is information from which direct identifiers have been removed but could be re-applied to render individuals identifiable again, is still considered personal data and so within the scope of data protection law.

Dr Byrne agrees that this is a challenge: "I'm keen to support the sharing of anonymous data to improve health and care but I don't underestimate the complexity of organisations navigating what can be considered a truly anonymous dataset. Challenges include balancing the theoretical possibility of reidentification through linkage with other datasets, against the likelihood of that happening within the context of both technical safeguards and legal sanctions."

The ICO has now published the first two (of several planned) chapters of guidance on anonymisation, pseudonymisation and privacy enhancing technologies[4] for consultation, and the NDG is continuing to work with the ICO to input into the further development of this guidance. Completion of that work by the ICO is of central importance for clarification in this area and it is hoped that progress on this will be seen in the near future.

[4] ICO call for views: Anonymisation, pseudonymisation and privacy enhancing technologies guidance. https://ico.org. uk/about-the-ico/ico-and-stakeholder-consultations/ico-call-for-views-anonymisation-pseudonymisation-and-priva-cy-enhancing-technologies-guidance/

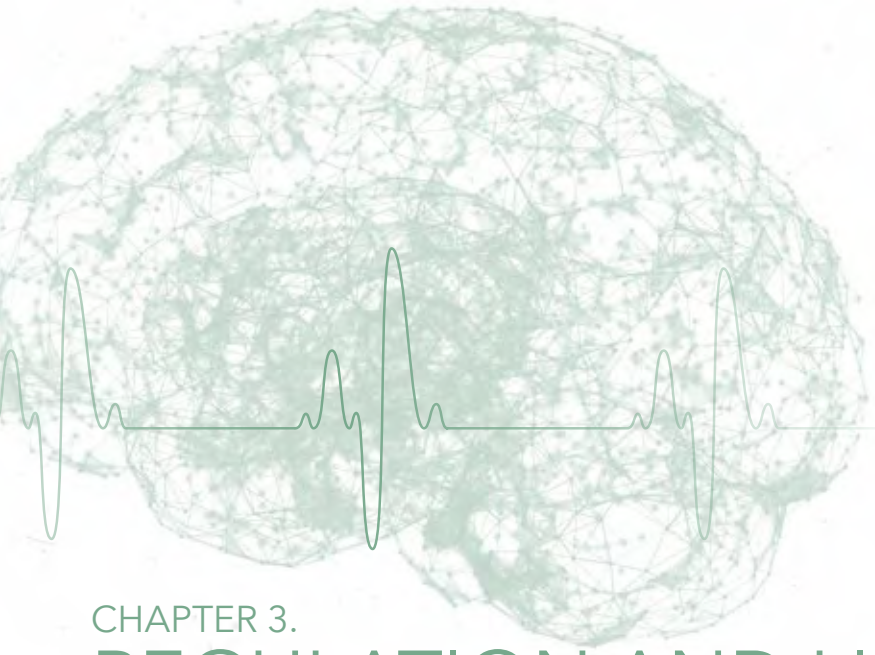

CHAPTER 3.

REGULATION AND LIABILITY

The medical device landscape and in particular the use of AI has changed enormously in recent years and is set to continue evolving, not least due to the acceleration of change created by the Covid-19 pandemic. Building trust and ensuring patients' safety in an AI-assisted healthcare system are key.

If AI is to play an increasing role in health systems, it is vital that patients, public, and healthcare professionals know that the medical technology employed is safe and fit for purpose. The use of AI presents unique regulatory challenges but building trust is crucial if the NHS is to reap the benefits of AI at a time of immense pressure. Patients must be confident that risks are minimised and any clinical errors that do occur will be compensated appropriately.

"The results of the product are only as good as the data that went into it. So the human aspect is absolutely vital and the evaluation of results is a human thing not an AI thing."

Dr Peter Feldschreiber, Barrister, 4 New Square chambers

Science not magic

Dr Peter Feldschreiber, dually qualified Physician and Barrister and Editor of the Law and Regulation of Medicines and Medical Devices (Oxford University Press 2021), specialises in product liability and clinical negligence law. He emphasises that "it is desperately important that the public understand that AI is just another application of science and not some magic trick".

Big datasets, with all their potential for unlocking valuable findings, need to be handled carefully. Possible bias must be identified and avoided and outcomes must not be manipulated or misleadingly represented. Dr Feldschreiber explains: "The results of the product are only as good as the data that went into it. So the human aspect is absolutely vital and the evaluation of results is a human thing not an AI thing. All that AI can do, however it is defined, is produce hypotheses and signals of what is happening with biological data and it's up to human intelligence to evaluate that. This is particularly important in the assessment of causation as opposed to evidence relationship."

"Where will responsibility lie if AI starts to take over more clinical decision-making and if it is opaque to healthcare professionals trying to interpret its outputs?"

Johan Ordish, Head of Software and AI (Innovation Devices Division) at the Medicines and Healthcare products Regulatory Agency (MHRA)

Of course, AI comes in many different forms and products. Johan Ordish, Head of Software and AI (Innovative Devices Division) at the MHRA, says: "We don't want to promote AI exceptionalism – sometimes it presents a modest performance improvement over standard methods.

"People make the standard assumption that AI takes a lot of data to train. That's not always true – some are not data hungry at all. People also often assume that AI is opaque, you don't know how it reaches the conclusions that it does. That is not true of some models either."

Protecting patients and enabling clinicians

Safeguards are already in place. Medical devices in the UK are currently regulated under the Medical Devices Regulations 2002 (MDR), with the MHRA responsible for the sector including the increasingly prominent role of software and AI medical devices within health systems.

Manufacturers wishing to place an AI medical device on the market in Great Britain need to register with the MHRA. In order to demonstrate that a medical device meets the requirements of the MDR, a conformity assessment process is required. Manufacturers of some devices may be able to self-declare their conformity against the MDR while others may need to apply to an Approved Body to approve and certify their products. Once this process is completed, a UKCA (UK Conformity Assessed) mark is placed on the medical device to show it conforms to the MDR requirements and that it is fit for its intended purpose and meets legislation relating to safety and performance. Manufacturers can use either the UKCA marking or the CE marking on devices they place on the GB market until 30 June 2023.

AI medical devices are being developed with so many potential applications that Ordish considers they must be regulated each according to the risk presented in a proportionate way.

While the MHRA and regulations are there to ensure AI medical devices function as intended and have a favourable risk/benefit ratio, clinicians play a pivotal role in the safe and effective use of AI solutions. The device's UKCA marking extends to the instructions for use, which must enable healthcare professionals to use it safely as intended by the manufacturer, giving them all the information and warnings they need.

Medicine is not a riskless endeavour

Dr Feldschreiber notes that "if the AI algorithms don't work or they work and produce equivocal results, which are subsequently used to make diagnoses and monitor disease, that could cause problems". He explains that the demographics and population clinical histories in one place where AI datasets might have been developed might be totally different to where they are potentially offered for use. "We must be comparing like with like as much as possible," he says.

"Medicine is not a riskless endeavour," Ordish points out and whether AI merely assists the healthcare professional to make decisions or replaces them is a question that presents both practical and legal challenges. Where will responsibility lie "if AI starts to take over more clinical decision-making and if it is opaque to healthcare professionals trying to interpret its outputs"?

Ordish explains: "In my view, the more of the critical workflow that a product takes up then the more appropriate it would be to sue in product liability rather than clinical negligence and that is also heightened by the fact that product liability is strict – you don't have to prove fault so it is a more attractive route versus clinical negligence." Dr Feldschreiber says: "When there is damage, you are looking at cause and effect, and the analysis of causation in medical terms is very similar to what would be happening in engineering terms."

While an injured patient may well consider that AI is the problem and decide to bring a product liability claim against the manufacturer, in practice questions may still arise about the actions of particular clinicians and the standard of care that should apply when they use and rely upon AI-driven decision-making tools.

All involved will want to ensure that AI products are deployed as safely as possible reducing current levels of risk to patients. However, where harm occurs mechanisms to ensure efficient and fair apportionment of liability between the manufacturer and providers of care will need careful consideration and development.

An added challenge is that not all problems will come to light quickly. Ordish notes that "approximately 80% of AI devices on the UK market are diagnostic in nature. The interesting thing about harm in the diagnostics world is that harm is not proximate to when the device malfunctioned".

The route to compensation for patients who suffer harm under clinician-led care is relatively clear. However, this model may well not work so smoothly, when AI products are introduced into the equation. Some would argue that although contained in a product, given its self-learning abilities, the AI itself should be judged as a human clinician, but this may not sit well with the current litigation procedure. It is also not easy to see how liability should be apportioned between manufacturers and clinicians in any given case. These are issues which will need to be actively addressed by AI stakeholders at some point especially if it becomes necessary to unlock and assess the validity of the 'black box' data at the heart of many AI devices. Ordish sums up the situation when he comments that: "There is some work to be done in ensuring liability is in the right place at the right time as we go forward."

If complex questions of attribution arise between the AI device and how it has been used by a clinician then perhaps the answer may lie in dedicated dispute resolution procedures so that proper safeguards are in place. Such measures along with better regulation, improved evaluation and more transparency to enable product validation will all help healthcare professionals and the public benefit from AI in the future as it plays an increasingly important role in our healthcare.

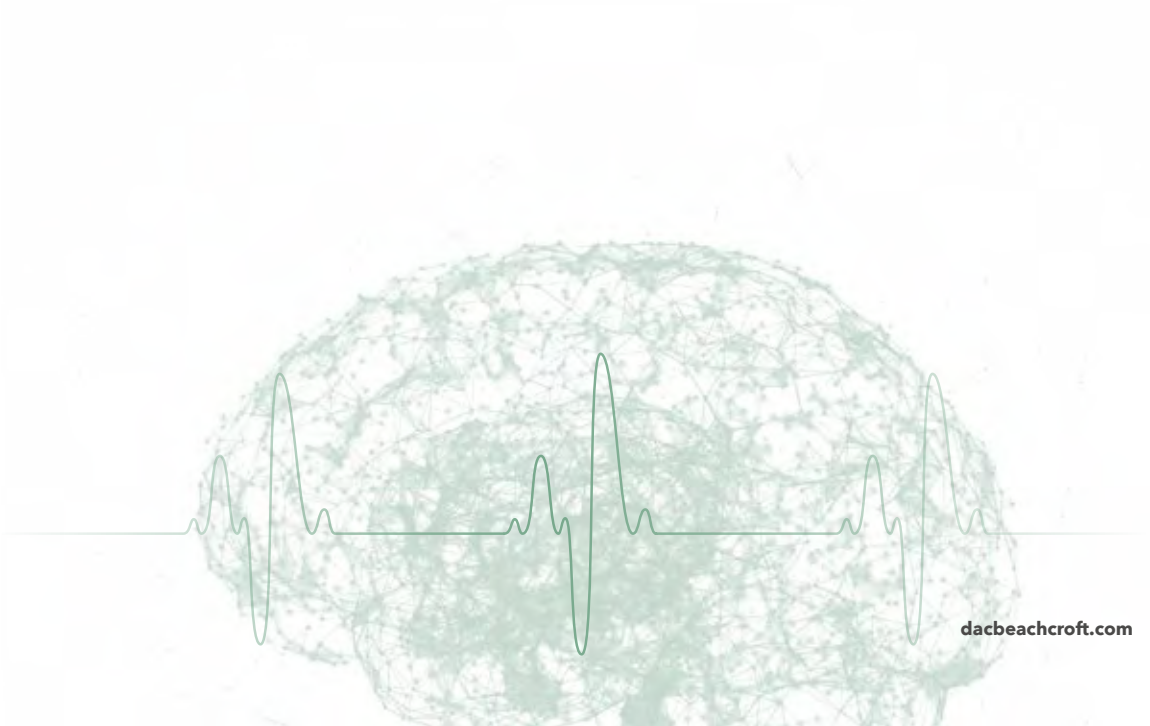

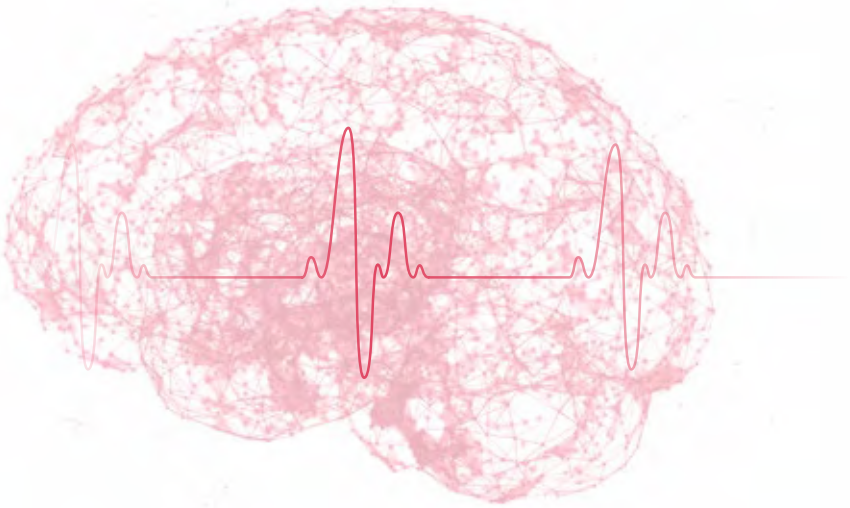

IN SUMMARY

This thought leadership piece is designed to give an overview of some of the many issues that we at DAC Beachcroft are encountering as AI products become more prevalent in the UK healthcare market.

AI-assisted healthcare is a rapidly developing and fast changing area. Here, we give some pointers of what is and may be happening now and in the near term, where we would like the market to go and how we can get there:

Take a patient-first approach to the protection of data
It may be tempting to look at the complexity of data protection law, and either view it as a purely technical exercise ticking off specific obligations and/or too complex to even contemplate navigating. In reality, it should be considered neither of those things.

Undoubtedly ensuring all individual instances and deployments of AI comply with data protection law is crucial but achieving this in a manner which prioritises the perspective of the patient, particularly by reference to transparency, is critical.

Consider compliance alongside and throughout the development of the technology
This will ensure that fundamental considerations such as the form and necessity of data, the roles of healthcare and technology providers, and the way AI fits in with clinical services enable the adoption of technology. The ongoing work of the ICO, the MHRA, the NDG and others with a role in the regulation of AI will enable further consensus about the ways legal obligations are met. A robust and well-planned start-to-end development process will also ensure no critical steps and considerations are missed.

Undoubtedly ensuring all individual instances and deployments of AI comply with data protection law is crucial but achieving this in a manner which prioritises the perspective of the patient, particularly by reference to transparency, is critical.

Increased regulation will improve trust amongst patients, public and healthcare professionals
Work is underway to improve the use and regulation of AI and a range of measures are being introduced. For example, the MHRA's autumn 2021 'Consultation on the future regulation of medical devices in the United Kingdom' set out proposals for changing the current regulatory regime, including for Software as a Medical Device (SaMD) and AI as a Medical Device (AIaMD). Following the government's response and the publication of draft regulations, the new regime is expected to come into effect from 1 July 2023.

The MHRA has published guidance on Software and AIaMD Change Programme which contains proposals for reform across the SaMD lifecycle.

In August 2021, the MAAS was also established. It aims to provide a one-stop shop for support, information and guidance on the regulation and evaluation of AI technologies in healthcare.

The issues surrounding healthcare and AI may be complex and interrelated but understanding how the various elements interlink is key to a successful outcome.

Alone we can do little, together we can do so much
Listening to our contributors it is clear that all parties should be working in partnership. The investors and developers need an innovation-friendly environment for their AI products to come to market. The NHS needs to use safe AI products based on shared reliable data. The regulators need to ensure that the public is safeguarded in ways that protect patient rights and promote the population's wellbeing.

The issues surrounding healthcare and AI may be complex and interrelated but understanding how the various elements interlink is key to a successful outcome.

If you want to find out how we see AI in healthcare and how we can help you navigate your way to an appropriate AI solution, we would be happy to talk.

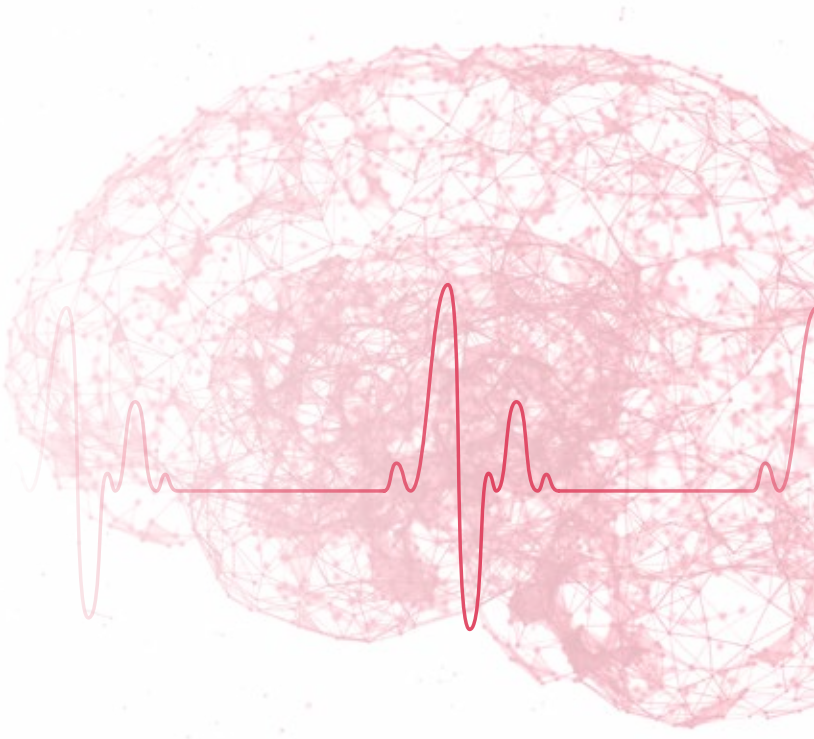

CONTRIBUTORS

Dr Junaid Bajwa
Chief Medical Scientist at Microsoft Research

Dr Nicola Byrne
National Data Guardian for health and adult social care in England

Dr Peter Feldschreiber
Barrister at 4 New Square chambers

Sarah-Jane Green
Head of Artificial Intelligence Regulations and Policy at NHSX

Johan Ordish
Head of Software and AI (Innovation Devices Division) at the Medicines and Healthcare products Regulatory Agency (MHRA)

Seb Wallace
Investment Director at Triple Point

CONTACTS

For any questions relating to this report, please contact:

Hamza Drabu
Partner
T: +44 (0) 207 894 6411
M: +44 (0) 786 059 2042
hdrabu@dacbeachcroft.com

Darryn Hale
Partner
T: +44 (0) 207 894 6125
M: +44 (0) 786 059 2042
dahale@dacbeachcroft.com

Alison McAdams
Head of Life Sciences
T: +44 (0) 207 894 6588
M: +44 (0) 777 172 5558
amcadams@dacbeachcroft.com

Jonathan Bonser
Legal Director
T: +44 (0) 161 934 3107
M: +44 (0) 781 801 3817
jbonser@dacbeachcroft.com

dacbeachcroft.com

Follow us: @healthlawuk

Connect with us: DAC Beachcroft LLP