



on
heart and
circulatory
diseases

Putting patients at the heart of artificial intelligence

This is not an official publication of the House of Commons or the House of Lords. It has not been approved by either House or its committees. All-Party Parliamentary Groups are informal groups of Members of both Houses with a common

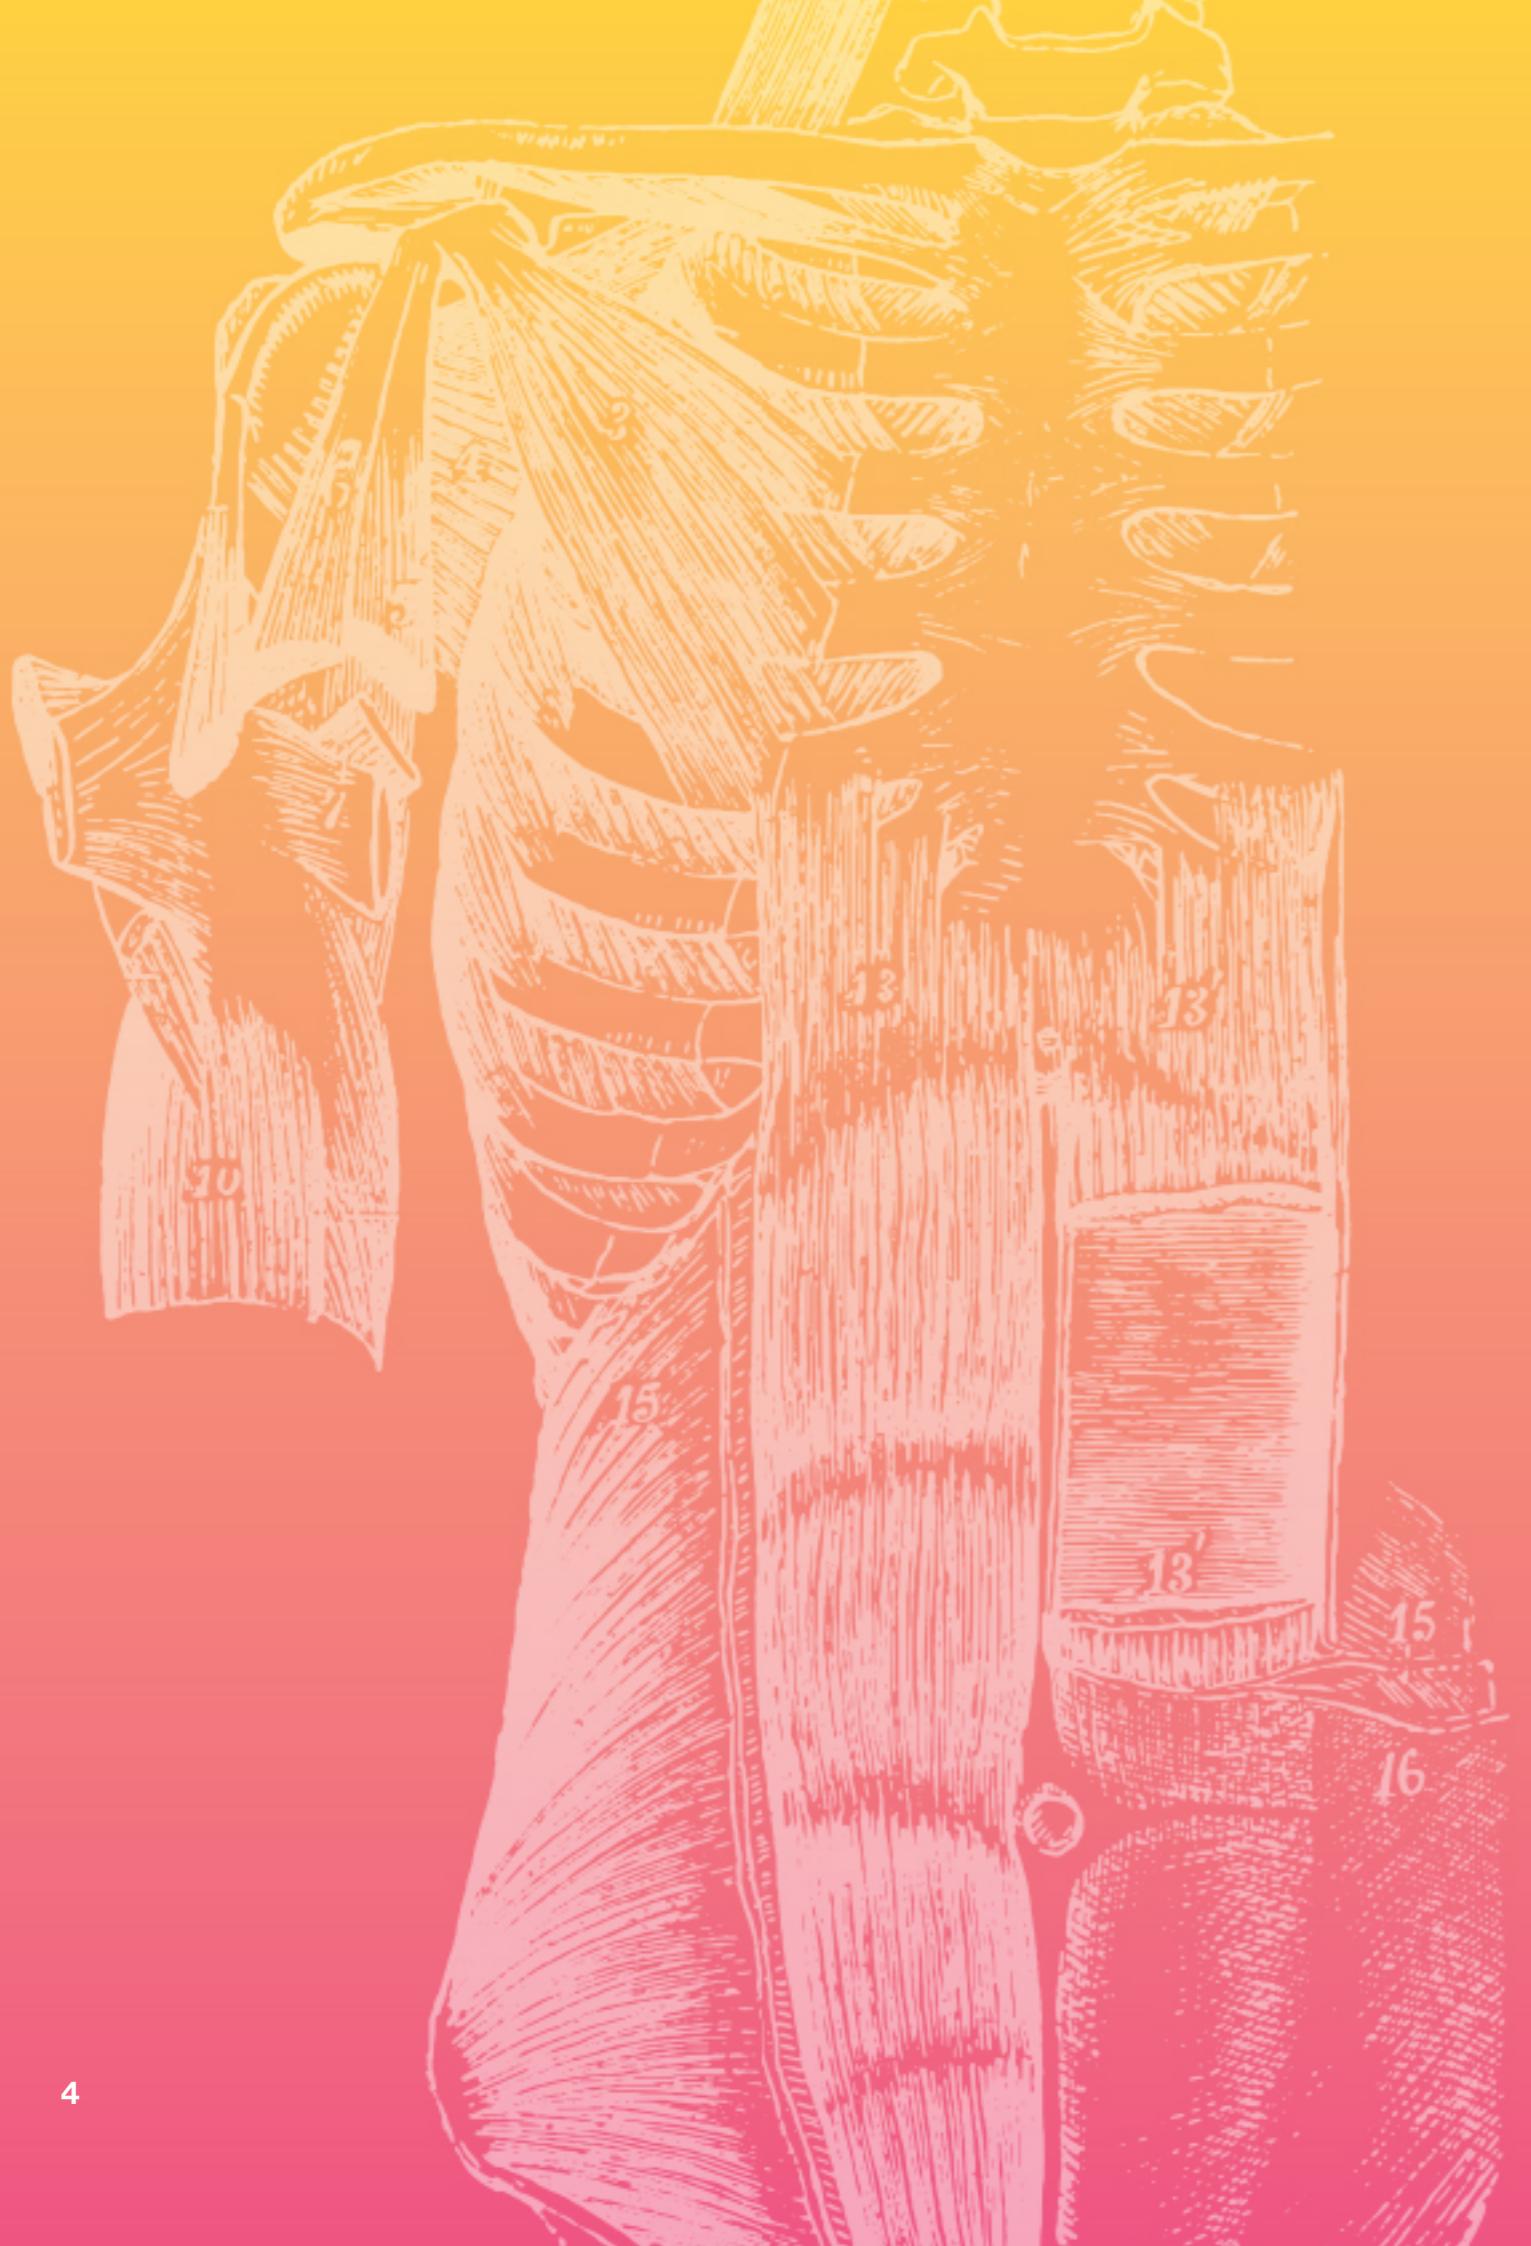
interest in particular issues. The views expressed in this paper are those of the group. This report has been compiled and funded by the British Heart Foundation, which provides the Secretariat for the APPG on Heart and Circulatory Diseases.



Contents

student content

<u>Advisory panel</u>	<u>5</u>
<u>Foreword</u>	<u>7</u>
<u>Summary of recommendations</u>	<u>8</u>
<u>Introduction</u>	<u>9</u>
<u>Policy context and inquiry process</u>	<u>11</u>
<u>Heart and circulatory disease: How can AI help?</u>	<u>12</u>
<u>Engaging with patients: Who, why, and when?</u>	<u>14</u>
<u>Learning from past mistakes</u>	<u>16</u>
<u>What do patients care about and expect?</u>	<u>18</u>
<u>Will artificial intelligence increase inequalities?</u>	<u>20</u>
<u>What do patients need to know?</u>	<u>22</u>
<u>Making clear who the public can turn to</u>	<u>24</u>
<u>Conclusion</u>	<u>27</u>
<u>Appendix</u>	<u>28</u>
<u>Endnotes</u>	<u>30</u>



Advisory panel

Expert advisors

Dr Natalie Banner
Understanding Patient Data Lead, Wellcome Trust

Dr Finn Catling
Critical Care Doctor & Machine Learning Researcher, UCLH

Anna King
Commercial Director, Health Innovation Network

Hilary Newiss
Chair, National Voices

Dr Declan O'Regan
MRC Senior Clinician Scientist and Consultant Radiologist and BHF-funded researcher, Imperial College London

Professor Simon Ray
President, British Cardiovascular Society

Jennifer Boon
Senior Policy Advisor, Centre for Data Ethics and Innovation

Dr Indra Joshi
Digital Health and AI Clinical Lead, NHS England

John Loder
Investment Director, Nesta

Jess Morley
Technology Adviser, DHSC, and Research Assistant, OII Digital Ethics Lab

Professor Nick Peters
Professor of Cardiology and Head of Cardiac Electrophysiology, Imperial College London

Dr Harpreet Sood
Associate Chief Clinical Information Officer, NHS England

Roundtable participants

Henry Smith MP
Chair of the APPG on Heart and Circulatory Diseases

Chris Green MP
Member of Parliament for Bolton West

Dr Vennela Boyalla
Research Fellow, Royal Brompton & Harefield NHS Foundation Trust

Sue Daley
Head of Programme for Cloud Data, Analytics and AI, techUK

Dr Indra Joshi
Digital Health and AI Clinical Lead, NHS England

George Howarth MP
Member of Parliament for Knowsley

Dr Adrian Baker
Policy Manager, British Heart Foundation

Dr William Bradlow
Consultant Cardiologist, University Hospitals Birmingham

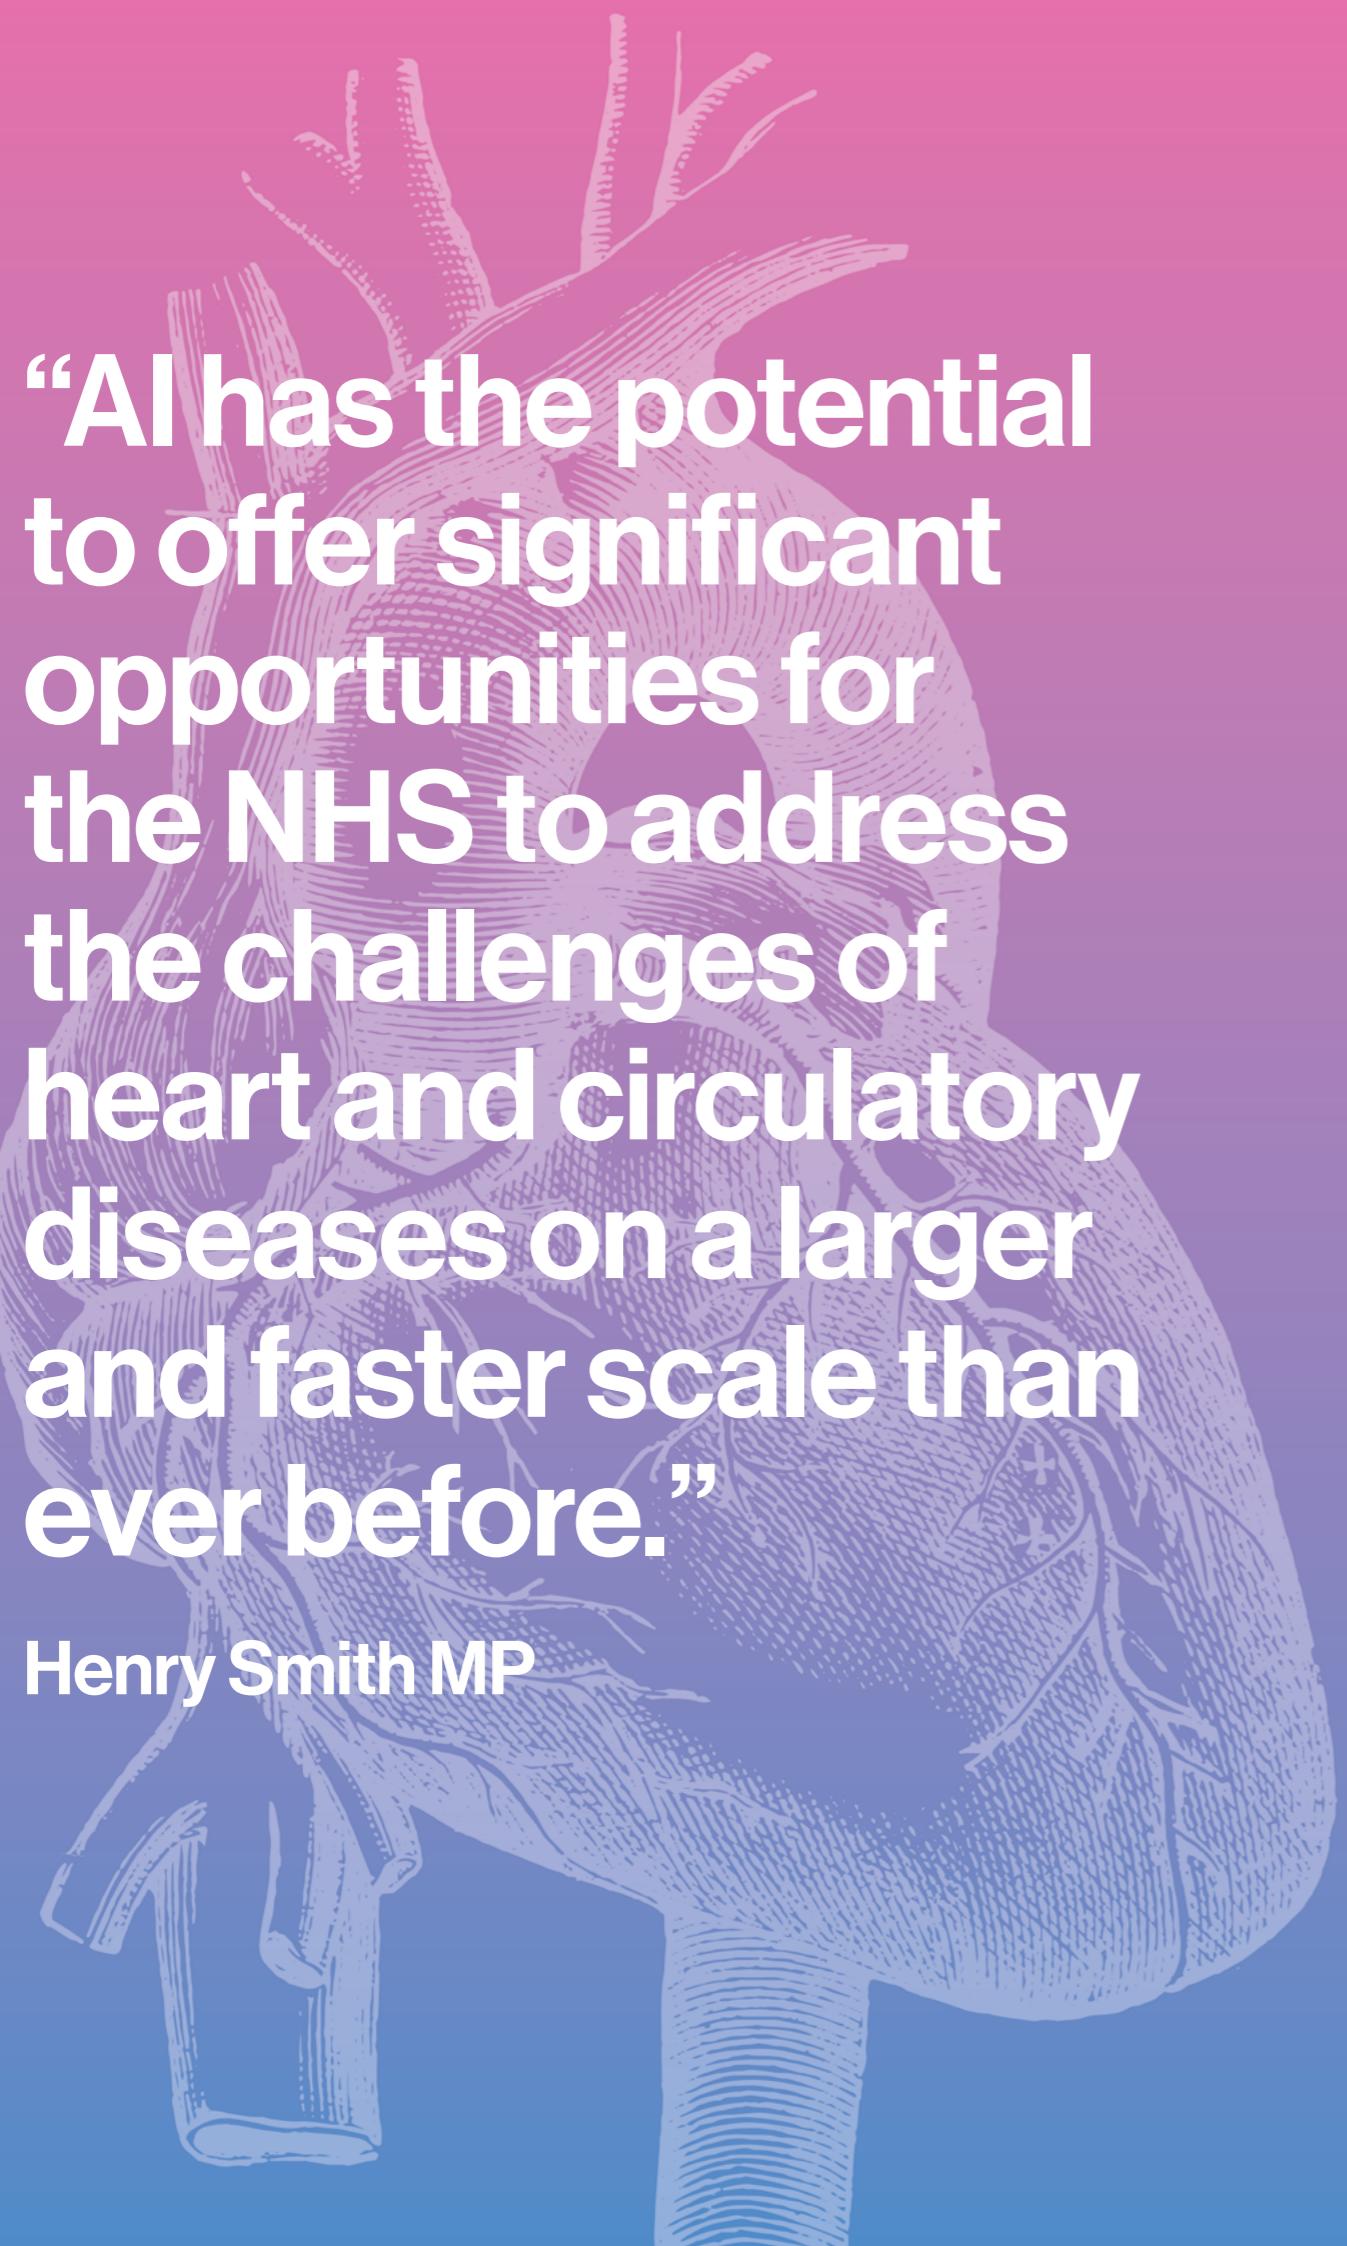
Simon Gillespie
Chief Executive, British Heart Foundation

Anna King
Commercial Director, Health Innovation Network

Patient representatives

Dr Phil Blakelock
Katy Chan
Eric Deeson

Robert Jones
Dr David Kelsall



“AI has the potential to offer significant opportunities for the NHS to address the challenges of heart and circulatory diseases on a larger and faster scale than ever before.”

Henry Smith MP

Foreword

Although major progress has been made in detecting, treating and managing heart and circulatory diseases, there are still seven million people living with these conditions in the UK and they cause a quarter of all deaths. It is crucial that we break new ground by developing better ways to prevent heart disease and detect it earlier as well as treat and manage conditions. Pioneering new research is improving our understanding of these devastating diseases - this includes the use of artificial intelligence (AI).

AI has the potential to offer significant opportunities for the NHS to address the challenges of heart and circulatory diseases on a larger and faster scale than ever before. However, the speed of new developments introduces a number of challenges to policy makers. These include ensuring a suitable regulatory framework, equitable access to technologies, strengthening ethical standards, and ensuring the NHS workforce is adequately trained to use new technologies. The APPG firmly believes that patients should have a fundamental say on the path these developments take.

This report brings together the views of patients, policy makers, healthcare professionals, parliamentarians and the health tech industry to highlight the need to engage with patients early and to better consider patient perspectives on AI. It presents a case for actively engaging patients and the public during the development and adoption of these data-driven technologies, to ensure they do not fall victim to the kinds of crises of public confidence that have hindered or derailed previous scientific advances. As Dr Eric Topol has highlighted in his recent report on preparing the workforce to deliver the digital future, *patient benefit must remain the driving criterion for AI design and use*¹.

I would like to thank all those who contributed to this inquiry, whether by participating in the roundtable discussions, responding to the survey, or reviewing drafts of this report. In particular, I would like to extend my gratitude to those who were kind enough to share with us their personal experiences of living with heart and circulatory diseases. Throughout the course of this inquiry we have heard many different examples of how AI could significantly transform the lives of people living with or at risk of heart and circulatory diseases. What is required now is a coordinated effort to involve patients in the design, use, and uptake of these new technologies.



Henry Smith MP
Chair
All Party Parliamentary Group
on Heart and Circulatory Diseases

Summary of recommendations

1

NHSX should set up a series of discussions with charities, the public, and others to:

- understand patients' views and concerns about the use of AI in healthcare
- understand the degree of granularity patients need in information sharing
- develop routes to ensure that information effectively passes between patients and policy makers
- explore the best timings and approaches for this engagement.

2

The APPG fully supports initiatives aimed at informing and engaging the public on how AI uses patient data, and the risks and benefits it brings. Similar mechanisms to those used by Understanding Patient Data (UPD) to engage with the public should be employed for the topic of AI. We therefore call on the UPD team to work with charities, patients, and stakeholders in the healthcare sector, to develop tools and resources for others to use in engaging the public about AI.

3

Academic Health Science Networks (AHSNs) should help facilitate exchanges of innovations through convening events and other communications, so that new developments in emerging technologies can be shared between patients living with heart and circulatory diseases, charities, and industry partners. This will allow innovators to better understand the challenges and needs of patients and the health system which supports them, thus allowing for more targeted and efficient development of new evidence-based solutions.

4

NHS England and NHS Digital should expand their good work on digital inclusion and digital literacy to explore the potential impact of wearables and AI on health inequalities, and begin working closely with charities and think tanks across different sectors in order to design strategies to mitigate potential inequalities.

5

NHSX should work with UPD, charities, and patient organisations to ensure that policy development in AI is designed with the explicit purpose of understanding, promoting, and protecting public values, and that this is clearly and openly communicated.

6

NHS England and NICE should encourage the development and use of reporting standards for AI research, in order to provide best practice for AI researchers. This can also provide a way for the media, policy makers, clinicians and the public to recognise quality in AI research.

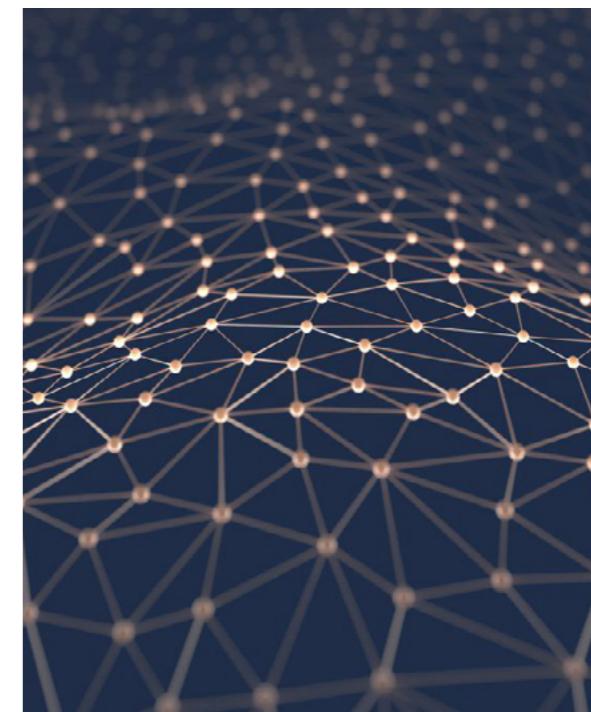
Introduction

Several reports have been written about the opportunities and benefits that AI can offer for healthcare^{2 3 4 5}. These reports have all contributed greatly to our understanding of the issue. Our inquiry does not seek to replicate the arguments and recommendations offered in these reports, but rather to begin a discussion about the right timing and approaches to engaging patients and the public with regards to AI in healthcare.

This inquiry has found that the speed at which AI has been developing over the last decade has made it challenging for policy makers and healthcare professionals to ensure that patients are meaningfully engaged with and informed about AI. Data-driven health and care technologies using AI depend on patient data. In order to make use of the full benefits of AI, it is vital that policy makers, the NHS, charities, healthcare professionals, and the health tech industry engage and empower patients in the design, development and diffusion of AI in the NHS; otherwise we may risk losing the great strides we have made to misinformation and apathy.

Looking at other areas of our lives gives an indication of just how transformative AI could be to healthcare. Every day, we interact with something that uses AI. From entertainment to online shopping, wearable devices to virtual assistants, chat bots to advertising; the use of AI is ubiquitous. Whether it is through faster or more accurate diagnosis, more personalised treatment, better triaging of demand, improvements in service planning and delivery, or better predictions, AI has the potential to touch upon all aspects of healthcare delivery and management.

Heart and circulatory diseases, including coronary heart disease, stroke, and vascular dementia are causing significant impact for millions of families across the UK. We must therefore utilise the enormous potential of AI across all of these areas, to transform the way we prevent, diagnose, treat and support those living with or at risk of heart and circulatory diseases. Whilst the focus of this inquiry is on heart and circulatory diseases, the recommendations and discussion are applicable more broadly.



What is artificial intelligence?

AI is the term given to a set of computer actions that mimic human intelligence. What separates modern AI, such as machine learning, from other types of computer programs is that they can 'learn' and improve at tasks. AI is particularly strong at finding patterns and trends in data that are not obvious through human analysis. Machine learning is a type of AI that uses algorithms (a set of rules that a computer uses to make a calculation) to look for patterns in data and then makes decisions based on these patterns. It looks for patterns in many different types of data – from scrutinising images to analysing genomic data.



Policy context

The APPG on Heart and Circulatory Diseases welcomes the great strides that have been made in recent years to speed up the development and diffusion of AI in the NHS. In outlining the benefits of AI to healthcare, a report by the House of Lords Select Committee on Artificial Intelligence in the UK called on the NHS to do more to become digital and take advantage of AI⁶. This has been met with a concerted effort to ensure that the NHS is a receptive environment for digital innovations.

In addition, the Government's AI Mission calls for the use of AI in prevention, diagnosis, and treatment, as a way to improve care and to reduce the incidence of late stage treatments. Of particular note is the emphasis on the need for cross-sector collaboration between the NHS, charities, and industry to accelerate the adoption of AI. The Office for Artificial Intelligence⁷ and the AI Council⁸ have huge potential to bolster the UK's position as a world leader in AI as part of the government's AI sector deal.

The Centre for Data Ethics and Innovation can also cement the UK's leadership in ethical AI and make sure society can shape the direction of travel and reap the benefits of AI. It will draw on evidence from academia, industry, and the public, and use those insights to advise the government about gaps in regulation. It will also work across sectors to highlight best practice in the use of data and AI⁹. In addition, NHS England's Long Term Plan has signalled a move away from focusing on the necessary but short-term target of going paperless by 2020, to one of longer-term digital transformation¹⁰.

The second iteration of the Code of Conduct for Data Driven Health and Care Technologies¹¹ rightly recognises that, while data-driven health technologies will undoubtedly deliver huge benefits to patients, clinicians, and carers, it is the duty of NHS England and the government to capitalise on these opportunities responsibly¹². These developments offer significant opportunities for AI and other digital innovations to radically transform the delivery of healthcare to patients with heart and circulatory diseases.



Inquiry process

The first step in the inquiry was to examine the UK policy and regulatory landscape and the existing literature on AI in UK healthcare. This gave us a good understanding of how the policy landscape is changing to respond to the pace of technological innovations. Our review encouraged us to address the gap in the literature on patient perspectives, as well as to look at the broader political and regulatory challenges that have been identified by others.

On 15 and 22 October 2018, the APPG held two roundtable discussions with patients, policymakers, cardiologists, researchers, and industry representatives to hear a range of perspectives on the opportunities and challenges of using AI in healthcare.

The APPG also conducted a survey to find out what people living with heart and circulatory diseases felt about the use of AI in healthcare. The survey was open for two months between November and December 2018 and attracted a total of 128 responses. The literature review, roundtable discussions, survey results, and our individual meetings with experts in the fields of AI and healthcare all contributed to the development of this report and its recommendations.

How can AI help?

As health care becomes increasingly digitised, the pool of patient data available is growing exponentially – this allows new AI technologies to continually develop and existing AI technologies to learn and improve. Medical imaging, predictive analytics¹³, and improved models that speed up drug discovery are some of the areas that could benefit from the use of large data sets. It is important to note that increases in computing power have enabled substantial changes in the potential for AI in healthcare. Within the last five years, AI has gone from struggling to identify images of cats, to being able to identify skin cancer in histological sections of biopsies just as well as a team of specialised doctors with decades of combined experience^{14 15}.

Pioneering research using AI is being undertaken in the UK, highlighting the potential to improve outcomes for people with or at risk of heart and circulatory diseases. For example, Google Deep Mind is using AI to predict risk factors for heart and circulatory diseases. By training a deep-learning algorithm on the retina scans of 284,355 patients, a team of scientists were able to predict risk factors such as age, gender, smoking status, systolic blood pressure, and major cardiac events¹⁶. Whilst the research is at the early stages, prediction via retinal scans opens up new opportunities to prevent heart attacks and strokes.

Medical research charities such as the British Heart Foundation (BHF) have also been funding pioneering research using AI. Since 2017, the BHF has partnered with the Alan Turing Institute to fund research in this area—bringing together data scientists and cardiovascular researchers to find solutions to important health problems¹⁷.

The case studies highlighted are examples of AI being used to predict risk in patients. This is important because by predicting risk more accurately (e.g. the risk of an adverse event or the risk that a patient's condition deteriorates), doctors are able to better target treatments. Better prediction therefore helps personalise treatment, ensuring that patients get the right type of treatment when it is needed.

The inquiry also heard possible scenarios of how AI could influence healthcare five to ten years from now. For example, complex AI systems such as chatbots would likely be used by a greater percentage of the population. By employing

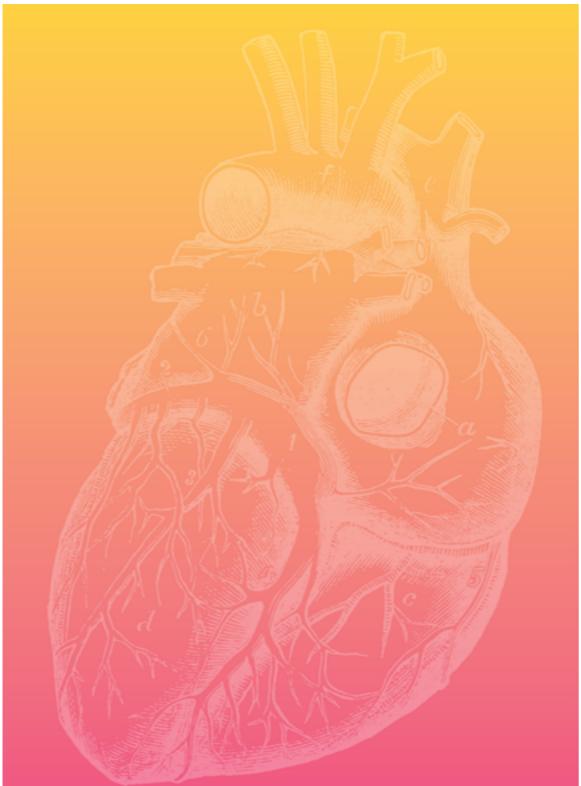
machine learning and having access to all guidelines, standards, and the latest evidence, chatbots could act as a triaging system that reduces demand on primary care and gives people immediate access to personalised health information. If patients do need to go to a GP, AI could be used to automatically create GP notes, freeing up the GP from the computer screen and improving the interaction between doctor and patient. Machine learning could also be used on GP referral systems, much like a more sophisticated version of the chatbots that people would use. This could reduce inappropriate referrals to secondary care, in-turn improving the resource allocation in secondary care.

AI could also be used to a greater extent in the self-management of care. Patients could use wearable devices and sensors to manage their condition at home and in the community instead of in hospital. AI systems could then monitor for unusual patient-specific patterns (such as deterioration in a heart failure patient), and relay that information to a clinical team for further intervention. Indeed, this last scenario represents an opportunity to put patients in much better control of their care.

Using machine-learning for personalised risk prediction of heart and circulatory disease

Angela Wood & Emanuele Di Angelantonio, University of Cambridge. Mihaela van der Schaar, University of Oxford

These researchers believe that the NHS health check could be better at distinguishing the risk of different types of heart condition, so that a person can receive the most suitable treatment. They are adapting cutting-edge statistical and machine learning methods, which use computers to learn from data and apply these techniques to existing sources of medical data to learn how to identify who is most at risk. If this new way of using health data can improve the accuracy and efficiency of NHS health checks, it could save both money and lives.



Using machine learning to predict clinical outcomes in heart failure

Dr Declan O'Regan, Imperial College London

Heart failure occurs when the heart is not pumping blood around our body as effectively as it should – and affects over 900,000 people in the UK. Changes in certain genes increase the risk of developing heart disease, but genetic testing alone isn't enough to tell if or when someone will develop heart problems. Artificial intelligence can now do many complex tasks faster and more accurately than humans – and it could be used to make better decisions about diagnosis and treatment by combining automated analysis of cardiac imaging with genetic information. Dr O'Regan and his team are using the latest advances in AI – across the domains of image analysis, genomics and prediction – and applying them to these challenging problems in healthcare. If successful, training of the new machine learning approaches with massive imaging datasets could be followed by their introduction into decision making and risk prediction in cardiology.



Engaging with patients Who, why and when?

Few would argue against engaging the public on AI. Where differences occur is in the timing of awareness raising, the resources needed to raise awareness of the potential of AI in healthcare, and the mechanisms for ensuring that patients and the public are effectively engaged. It is important to highlight just what can happen if we don't engage, as this provides useful examples of the imperative to get this right. This inquiry has found that it is vital that efforts to engage the public on AI in healthcare begin right away. If the knowledge gap with regards to AI in healthcare is not filled by the correct information, it will be filled by misinformation.

A report by the Academy of Medical Sciences on data-driven technologies highlighted the importance of public engagement and partnership. The report noted that early patient involvement in the development, evaluation and implementation of technologies is essential to maximise impact¹⁸. Likewise, the House of Lords Artificial Intelligence Select Committee's report on AI in the UK stated that public acceptance and understanding of the use of patient data in AI, and maintaining public trust with regards to data and AI, are of the utmost importance to ensure that patients and the health system benefit from this technology¹⁹. In her evidence to the Select Committee, the National Data Guardian, Dame Fiona Caldicott, argued that not enough was being done to engage the public on the use of patient data and AI²⁰.

This inquiry has found that it is vital that patients are involved in the development, implementation, and evaluation of AI in healthcare. Evidence shows that involving patients in the development of healthcare innovations, for example through co-production, can improve health outcomes and lower costs^{21,22}.

A consensus statement calling for the need for public involvement and engagement with health research data was signed by universities across the globe, alongside the Wellcome Trust and Farr Institute. The statement emphasised the need for public engagement and to approach the public as part of the solution to ensuring that data-intensive health research benefits society²³.

Our survey of 128 patients found that a majority of respondents (64%) had at least some awareness about the potential future uses of AI to diagnose and treat heart and circulatory diseases. However, only 17% of respondents were aware of any current uses. This represents a huge opportunity to inform patients as to the opportunities of AI. What follows is a selection of the questions asked and the responses we received.

"I'm not sure the general public know that much about AI, personally. My fear is that the perception is that it is all to do with robots, which I've heard many times, I really don't think the message has got anywhere near out there yet."

Patient representative

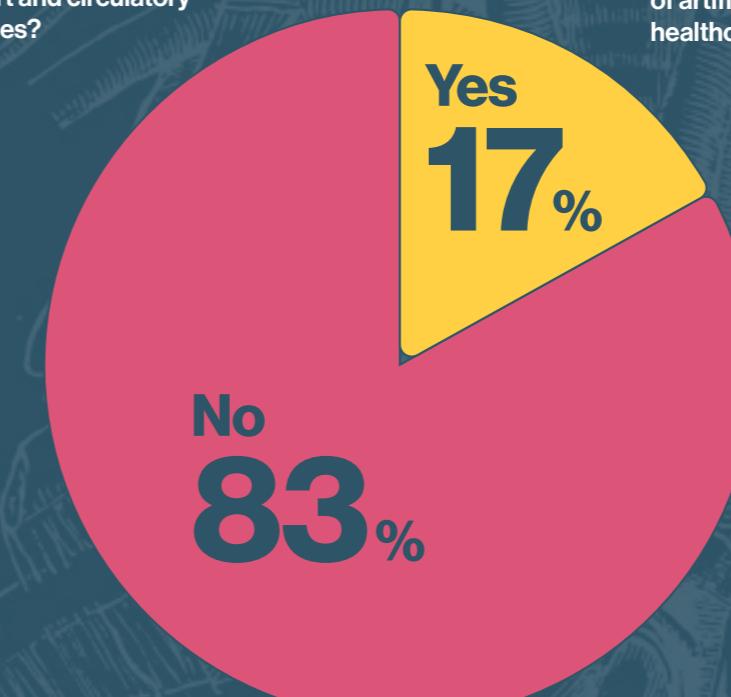
"There are a lot of scare tactics about AI, especially on the robotics side of it."

Patient representative

Recommendation

- NHSX should set up a series of discussions with charities, the public, and others to:
 - understand patients' views and concerns about the use of AI in healthcare
 - understand the degree of granularity patients need in information sharing
 - develop routes to ensure that information effectively passes between patients and policy makers
 - explore the best timings and approaches for this engagement.

Are you aware of any current cases of artificial intelligence being used in the diagnosis and treatment of heart and circulatory diseases?



Who do you think has a responsibility to inform the public about current and potential future uses of artificial intelligence in healthcare?



How aware are you about the current and potential future uses of artificial intelligence in the diagnosis and treatment of heart and circulatory diseases?



To what extent do you agree that the public should be well-informed about the use of artificial intelligence in healthcare?



Learning from past mistakes

It may be possible to diminish fears about AI if knowledge and understanding is widespread. One frequently cited example is what happened to genetically modified (GM) food. Despite the potential benefits of the technology, public mistrust led to a narrative that moved away from a measured discussion of the costs and benefits of GM produce²⁴.

In 1998 the biochemist and nutritionist Árpád Pusztai claimed on television, and later in a press conference in the House of Commons on 12 February 1999, that feeding rats with genetically modified potatoes caused damage to their stomach lining and immune system. Despite Pusztai's paper being published in a peer-reviewed journal²⁵, the Royal Society (1999) stated that "we found no convincing evidence of adverse effects from GM potatoes"²⁶. Although Pusztai's claim was not supported by evidence²⁷, the headlines of many newspapers from that period created an environment of mistrust^{28 29}.

Care.data³⁰ is a recent and well known example where initial public doubts over a programme were made worse by lack of transparency and poor communication^{31 32 33}. The programme aimed to bring together health data from across the NHS into a single database which could then be used by researchers to develop new treatments and assess performance of NHS services. The lessons of care.data are being learned and successfully applied to the National Data Opt-Out Programme. This includes utilising a range of communications approaches, a longer timetable for implementation and sustained dialogue with the public post-implementation³⁴.

When should we engage the public?

Public engagement and awareness-raising should be proactive rather than reactive. Attempts to reverse negative public perceptions can be difficult³⁵. The cases described above are clear examples of what happens when misinformation and mistrust take hold due to a lack of early and proactive engagement. Lack of information can lead to uncertainty, and uncertainty can lead to fear. However, providing information about AI early on is not enough to meaningfully engage with patients. Patient engagement is needed not only at the policy level, but also at service design and implementation. Meaningful, early, and proactive engagement on how AI is used in healthcare is essential for effective implementation and sustainability. This means that

patients and the public need to be involved with the design of the overarching system and the paths that AI will take in healthcare³⁶. Indeed, interviews conducted by the inquiry found that policy makers should not approach AI as if it is on a predetermined path and that patients should just be told what is being done. In reality, there are multiple paths AI can take, but this will depend on the quality of public discussion and engagement.

Despite the support from this inquiry for the use of AI in heart and circulatory conditions, policy makers should be fully prepared to accept unfavourable public attitudes, and not presume that they are merely a result of lack of engagement or misinformation. This means that some developments in AI or some paths that AI could take would need to be changed if, after meaningful engagement, there is lack of support from patients and the public. At the same time, an environment needs to be created where people have a right to disengage with these issues should they wish. Policy makers should also note the importance of differentiating between patients who do not want to engage with this topic, and patients who do but do not have an outlet as both can look the same from the outside, particularly if no effort is made to identify the latter group. Under-represented groups by definition may not have the capacity or resources to be involved in public discussions, which is why the role of charities and other intermediaries is important.

What does good patient engagement look like?

The Understanding Patient Data initiative (UPD), funded by the Wellcome Trust and several Government bodies, provides a useful template for stakeholders to collaborate in engaging and informing the public. UPD supports better conversations about the uses of health information. Its aim is to explain how and why data can be used for care and research, and how personal information is kept safe. It was developed in part to

"Public engagement is the real key. The massive success – or not – of AI will rest upon it ... Everybody has to understand and be positively part of this. If not, AI is likely to be the latest GM in our society."

Lord Holmes of Richmond

Parliamentary debate, 19 November 2018

support the implementation of the National Data Opt-out Programme (NDOP) and has, since 2016, worked with patients, charities and healthcare professionals to champion responsible uses of data. An evaluation of UPD in 2018 found that it had brokered information sharing between parties and enabled feedback across the landscape where previous tension existed as well as defused some of the sting in the patient data debate³⁷.

UPD is not resourced to develop and deliver direct public engagement campaigns. They instead develop tools and resources for other stakeholders to utilise. Some of the creative resources UPD has developed to help others engage with patients includes the use of animations. These explained complicated topics, such as clarifying to patients how their data flows through the healthcare system and how it is used. Some of the animations focused on specific conditions and medical events, and had large input from condition-specific charities. These animations enabled organisations to better tailor messages and information to their locality or target audience. Such collaborative working helped pool resources and expertise, resulting in patient engagement that had both credibility and legitimacy.

Importantly, UPD sought to delve deep into understanding patient attitudes and beliefs with regards to health data; exploring what people thought and why they thought it. This depth of understanding is needed when developing policies on AI because it helps meet or mitigate the underlying cause of public beliefs. UPD provides a valuable model for future patient engagement, and should build on its work by engaging the public with regards to AI on its use potential in healthcare.

Other stakeholders have an important role to play in engaging the public, or providing a forum where stakeholders can convene. This is particularly true of Academic Health Science Networks (AHSNs) and programmes such as DigitalHealth.London, which have a vital role to play in bringing together industry, patients and charities so that information can be shared. This has a number of benefits, as it will allow innovators to better understand the needs of patients and the health system, so the development of new technologies and services better meet those needs.

The efficient signposting of innovators to the support infrastructure should allow them to develop innovations of a high quality quicker, particular

through linkages with research, but also real-world validation of solutions. High impact solutions should then be supported by the AHSNs and NHS England to be adopted by the NHS, utilising programmes such as the Accelerated Access Collaborative. One of the key strengths of the 15 AHSNs across England is the focus on promoting regional improvements in health and economic growth. This regional aspect, coordinated by the AHSN Network, allows for national discourse to be contextualised to local needs and will be an essential part of effective engagement, as local knowledge is used to drive change³⁸.

Recommendations

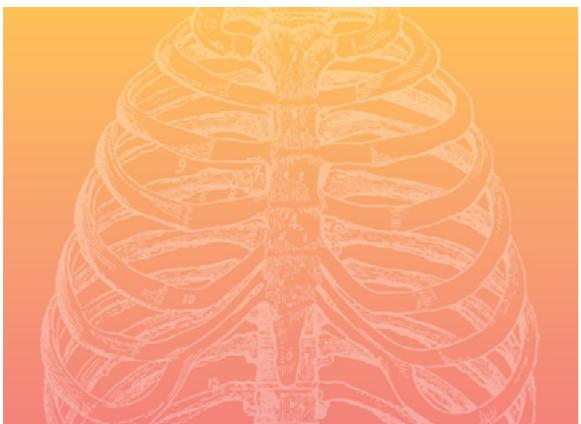
2 The APPG fully supports initiatives aimed at informing and engaging the public on how AI uses patient data, and the risks and benefits it brings. Similar mechanisms to those used by Understanding Patient Data (UPD) to engage with the public should be employed for the topic of AI. We therefore call on the UPD team to work with charities, patients, and stakeholders in the healthcare sector, to develop tools and resources for others to use in engaging the public about AI.

3 Academic Health Science Networks (AHSNs) should help facilitate exchanges of innovations through convening events and other communications, so that new developments in emerging technologies can be shared between patients living with heart and circulatory diseases, charities, and industry partners. This will allow innovators to better understand the challenges and needs of patients and the health system which supports them, thus allowing for more targeted and efficient development of new evidence-based solutions.

What do patients care about and expect?

This report does not seek to prescribe what should be included in patient engagement strategies, as this would require more detailed discussion amongst stakeholders. However, as a starting point, it is important to highlight some patient expectations with regards to AI in healthcare to frame and guide future discussions.

AI is a complex topic, with nuances on how it can be defined and applied³⁹, made more difficult by emerging developments. It is also clear that there are varying degrees of awareness amongst the general population with regards to AI. A survey by the Royal Society for the Encouragement of Arts, Manufactures and Commerce (RSA) reported that only 19% of respondents were aware that 'automated decision systems' could be used to aid decisions about healthcare, and only 20% of respondents supported the use of 'automated decision systems' to assist with decisions about healthcare (with 23% neither supporting nor opposing). Whilst the report does not break the data down into detail, it was suggested that the more unfamiliar people are with AI and what it does, the less likely they are to show support⁴⁰.



"The human interface is crucial, there must always be people keeping an eye on what's going on and be prepared to step in but I'm concerned about the replacement of all medics by machines, robots, whatever you call them."

Patient representative

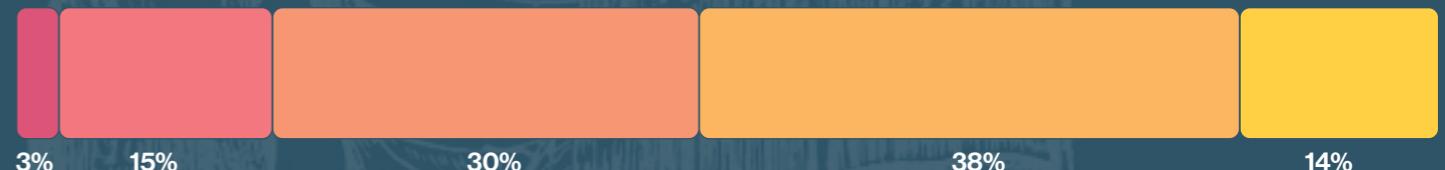
AI in healthcare The patients' view

To what extent do you support doctors using artificial intelligence technologies to assist them in diagnosing and treating heart and circulatory diseases?



Respondents were presented with a hypothetical scenario designed to assess how they may feel about AI technologies once they are more widespread in the health system, and how AI may impact the relationship between patient and doctor: "In ten years' time, artificial intelligence technologies are commonly used across UK hospitals to help doctors diagnose and treat heart and circulatory diseases. In practice, these technologies have been shown to be able to independently diagnose a wide range of heart and circulatory conditions better than most human doctors working without access to artificial intelligence technologies."

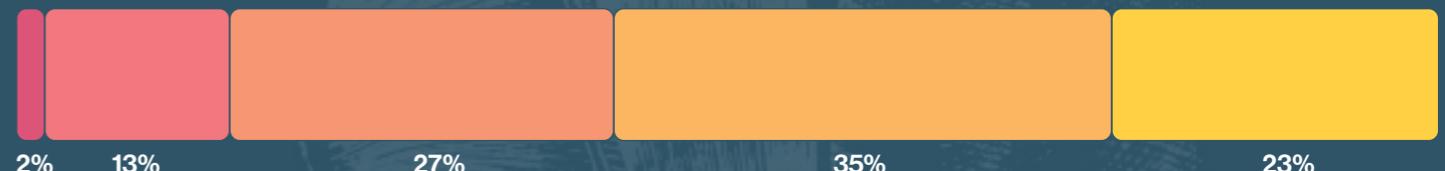
In this scenario, how comfortable would you be if your doctor chose not to consult artificial intelligence technologies when diagnosing your medical condition?



In this scenario, how comfortable would you be being diagnosed by a human doctor who used their own judgement, along with recommendations of an artificial intelligence assistant to help inform their decision?



In this scenario, how comfortable would you be being diagnosed solely by an artificial intelligence technology, without the input of a human doctor?



Legend for support:

- Strongly support
- Support
- Unsure
- Do not support
- Strongly do not support

Legend for comfort:

- Very comfortable
- Comfortable
- Unsure
- Uncomfortable
- Very uncomfortable

Legend for comfort:

- Very comfortable
- Comfortable
- Unsure
- Uncomfortable
- Very uncomfortable

Legend for comfort:

- Very comfortable
- Comfortable
- Unsure
- Uncomfortable
- Very uncomfortable

Will artificial intelligence increase inequalities?

This inquiry found that one of the major challenges of introducing AI into every day practice in healthcare is its potential to exacerbate health inequalities. Wearable devices provide a good example. As the Chief Medical Officer (CMO) has highlighted, these devices can be empowering - they can allow users to see their own data and take control of their health⁴⁴. Wearable devices can use AI on the data they generate, and offer real opportunities for self-management of a number of heart and circulatory conditions. Whilst the accuracy of wearable devices varies depending on a range of factors such as what is being measured or what device does the measuring⁴⁵, a YouGov poll highlighted that 57% of respondents were interested in wearables that monitor blood pressure⁴⁶.

This inquiry found that demand for technologies from patients and patient organisations can often be the catalyst for innovations to be adopted or funded. Whilst initially new technologies are taken up by 'early adopters' who are more interested than the general population in that particular innovation, the active role patients and consumers can play in spreading adoption of AI-driven technologies should be noted.

Age, ethnicity, and socio-economic demographic factors can influence access to the best technologies. Although there is currently a lack of evidence that explores inequalities of access to AI technologies by different demographics, it is possible to extrapolate lessons from similar work in other areas. For example, the Pew Research Centre has highlighted variations in smartphone usage between different demographic groups in the US⁴⁸. In 2018, 94% of people aged 18-29 owned a smartphone compared to 46% of people aged 65+. However, research by Ofcom has shown a rising trend in smartphone and digital usage among the over-65s. Indeed, almost 90% of UK adults have used the internet in the last three months, and almost half of adults over the age of 65 shop online^{49 50}.

Nevertheless, variations in access exist, and are relevant because AI is currently being implemented in consumer-facing technologies such as smartphones that help manage adherence to blood pressure medication, smart watches that track and analyse heart rates and voice-activated assistants that can act as useful reminders to take medications. As AI becomes increasingly used

91%

Percentage of patient respondents agreed that the public should be well-informed about how AI is used in the healthcare sector

for health-related advice, variations in access may result in variations in using personalised advice and in self-management of conditions.

Due to its data-driven nature, AI can also reflect inequalities and biases that already exist in the health system. For example, lack of representation of people from Black, Asian, and minority ethnic (BAME) communities in clinical trials may mean that AI could work less well for the BAME community, and this may similarly be the case for other under-represented groups. As discussed later in this report variations in terms of language skills and digital literacy can also lead to inequalities, particularly for patients to get the most out of consumer-facing products.

The work being done by NHS England and NHS Digital to widen digital participation amongst those less likely to use digital technologies is therefore a vital part of ensuring that the public are empowered in using AI⁵¹.

Recommendation

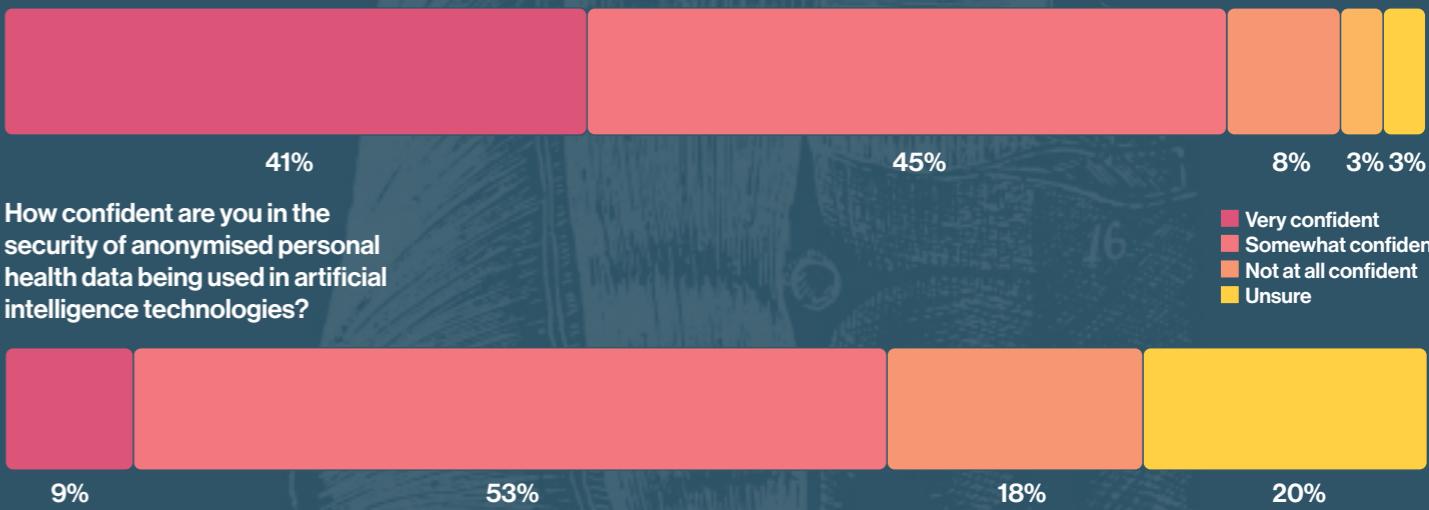
4 NHS England and NHS Digital should expand their good work on digital inclusion and digital literacy to explore the potential impact of wearables and AI on health inequalities, and begin working closely with charities and think tanks across different sectors in order to design strategies to mitigate potential inequalities.

Patient data

Policy makers should feel confident that patients support the use of AI in healthcare if it is done to improve health outcomes. 86% of respondents to our survey were comfortable with their personal health data being used to help better diagnose medical conditions. This is in contrast to a Future Advocacy survey conducted by YouGov which found 40% of respondents were comfortable with sharing their data⁵². Whilst these findings should not be compared like-for-like due to differences in the way questions were framed, differences in opinions between patients and the general public with regards to data should nevertheless be further explored.

This broad level of support should be encouraging for policy makers, providing an impetus to be brave, push forward timescales, and expand initial pilots. However, there are caveats. To maintain patient support, we need to meet the conditions of their support. Our survey shows that whilst the majority of respondents are happy for their data to be used in AI technologies, they care who uses their data, and are only somewhat confident in its security.

Artificial intelligence technologies can use your anonymised personal health data to help better diagnose medical conditions in other people. How would you feel about your anonymised data being used for this purpose?



synthetic data can produce novel data for simple or highly defined environments (e.g. a chess board), they cannot yet produce novel data for highly complex situations and real-world environments. This means that whilst synthetic data has its advantages and uses, it cannot be a substitute for using real data for certain types of research^{57 58}.

Ultimately, maintaining public trust with regards to data and AI is of utmost importance to ensure that patients and the health system benefit from this technology. That is why initiatives aimed at engaging the public on how AI uses data are so important. No use of technology is without risk, and honest discussions about the potential risks of the use of AI is necessary in order to ensure that progress on AI in healthcare is sustainable.

Public trust cannot be engendered if only benefits are highlighted. Nevertheless, it is equally important not to shy away from highlighting the benefits of AI in an attempt to be balanced. Organisations that work with patients, such as charities, can play a significant role in engaging patients and the public on these issues.

■ Very comfortable
■ Comfortable
■ Unsure
■ Uncomfortable
■ Very uncomfortable

■ Very confident
■ Somewhat confident
■ Not at all confident
■ Unsure

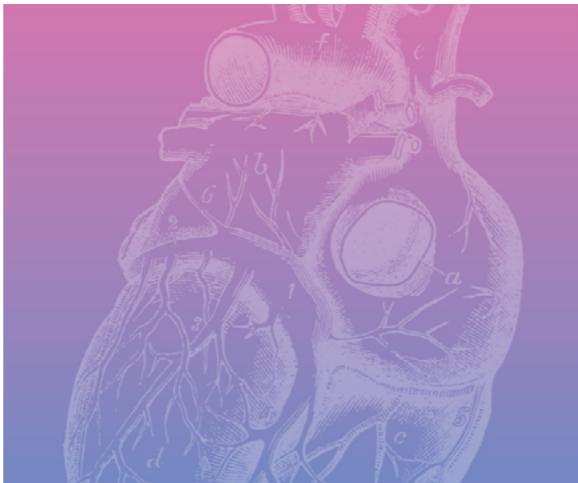
What do patients need to know?

One of the issues identified during this inquiry was that of the 'black box' of AI. This refers to the difficulty in understanding how AI models reach their decisions⁵⁹. This can be problematic because regulators may want to scrutinise how models reach their conclusions and whether, for example, they include any biases. This can be exacerbated with the increasing autonomy AI is given (i.e. the less human supervision is involved). This inquiry found that research into explainable AI is rapidly growing, and began in the 1960s. Different techniques can be used to help understand why an AI has come to an outcome^{60 61 62}. This can be useful in healthcare because such techniques can help AI to both suggest a diagnosis and highlight the reasons why such a diagnosis was made.

Another issue with 'black box' AI is understanding not only *why* an AI has given a certain outcome ('explainability') but also *how* it has arrived at such an outcome⁶³. Some experts have argued that the 'black box' issue is not really a problem, because humans are equally opaque in how they arrive at decisions⁶⁴. However, the ability to scrutinise, conduct quality assurance, and undertake due diligence are important parts of regulating the health system and ensuring patient safety is paramount.

Using non-'black box' models could also come at a cost. Currently, one of the greatest costs is accuracy. That is because AI models that are explainable tend to be simpler than AI models that are less explainable. More complicated 'black box' models can produce more accurate results in certain (typically more complex) tasks. Only using AI models that are explainable therefore reduces the types of models that can be used, some of which may produce better results in a given task.

There are also tensions between being able to understand how an AI model works and comes to a decision and privacy. One tension lies in the training of models themselves. Traditionally, training AI models requires data to be sent to a centralised location. This is because the processing power to train models on large datasets tends to be found on large servers. This form of training risks making personal data discoverable by locating it in one place. However, new forms of training, such as 'federated learning' are emerging, which do not require data to be centralised; for example, users get to keep their data on their smartphones⁶⁵. Another tension is that the outcomes of an AI model are rarely understood by looking at only the



"We are currently on the journey of working collaboratively with all parts of the system to ensure what is being developed is safe, effective and ethical, but we also need to ensure this is an embedded behaviour when we deploy this technology. Together we can create an ecosystem of distributed responsibility that encourages innovation whilst being mindful of its impact and outcomes."

Dr Indra Joshi
NHS England

algorithm and source code. Instead, understanding the outcomes requires looking at the input data, the output data, and their interaction with the algorithm⁶⁶. Federated learning techniques such as the one mentioned above can help in this case by protecting the input data and making just the algorithm and outputs transparent.

This inquiry found that when engaging the public on the transparency of AI, it is important to specify what type of transparency, the intended outcomes of that transparency, and for whom the transparency is intended. For example, if the intended outcome

of transparency is to ensure accountability if something goes wrong, that may require different levels of transparency and different things being made transparent compared to transparency aimed at ensuring trust in the developers or users of the AI. Transparency can include outlining why an algorithm was developed, what types of data were used, and how development was funded. Other forms of transparency, such as developers publishing their source code, may therefore not be needed, and indeed can jeopardise a vibrant ecosystem and result in disincentives for the development of innovations that could ultimately lead to patient benefits⁶⁷. More work is needed to explore patient attitudes to different types of transparency as it is conceivable that patients will place greater importance on transparency that leads to improved trust and safety than transparency for other intended outcomes.

However, maintaining public trust over the use of AI in healthcare requires confidence and faith in regulators being able to unpick an AI algorithm should there be a need. Transparency is also required by the research community, where AI-driven recommendations could have profound impact and where understanding the reasons behind the recommendations could lead to new insights and research. Transparency is also useful for research because it allows for reproducibility, which is an important part of assuring quality.

Finding ways to resolve these issues is important, and is one of the range of issues that will be explored by the Centre for Data Ethics and Innovation (CDEI). However, in discussions on what the right level of transparency should be, and on the details of how to communicate this transparency, it is important to note the views of patients. This is not something the Centre for Data Ethics and Innovation can do by itself, and Roger Taylor, the Chair of the CDEI has stated that it is incumbent on the organisations using AI to engage with the people who will be affected by its use⁶⁸.

Digital literacy

As we have mentioned earlier, AI is already being widely used in consumer-facing health technology products. Companies such as Apple and FitBit are releasing ever more sophisticated wearables that can detect conditions such as atrial fibrillation, and companies such as Babylon and Ada are providing AI-powered health information through smartphones. As the CMO report noted, AI-

Recommendation

5 NHSX should work with UPD, charities, and patient organisations to ensure that policy development in AI is designed with the explicit purpose of understanding, promoting, and protecting public values, and that this is clearly and openly communicated.

powered wearables could empower patients to be equal partners in their care. However, these products may also have unintended consequences as patients may misinterpret information from this technology and request treatment from clinicians based on misperceptions. Awareness and understanding of AI in healthcare must be complemented by digital literacy⁶⁹.

The latest generation of wearable technologies can provide potentially useful tracking for patients with intermittent symptoms, but engaging with the public now about how to make the best use of their wearables might help stem a spike in primary care demand. Wearable technologies can pick up transient abnormalities that are not symptomatic of disease. When a patient normally sees a doctor, they tend to already be symptomatic. This means that their 'risk score' is higher, and therefore the benefits of any treatment or investigation outweigh the risks. However, in cases of sub-clinical symptoms⁷⁰, the risks of treatments and investigations may outweigh the benefits. Given the potential number of people who may use wearables in the future, if only 1 per cent generate a false positive, it could represent significant costs both to the individual patient and wider health system⁷¹.

In the long term, wearables have good self-management and prevention potential. Machine learning may result in fewer false positives, provide more sophisticated advice and triaging recommendations, and might offer people who use wearables better information on sub-clinical or transient abnormalities, all of which could help reduce unnecessary demand. However, there are examples where demand can occur even with little evidence of efficacy, as we have seen with antibiotics for the common cold.

Making clear who the public can turn to

The rapid development of technological innovations means that health systems are becoming even more complex environments to regulate, and the increasing involvement of technology companies introduces an extra layer of complexity. There is currently a great deal of uncertainty and lack of clarity regarding who the public can turn to if they encounter an issue with a data driven technology. Several excellent reports have been published exploring regulation in AI^{72 73 74} and the Centre for Data Ethics and Innovation will also be looking at regulation. However, there is still work to be done to highlight what good looks like in terms of the regulatory landscape.

This need for clarity is becoming ever more pertinent with the speed of technological developments, and consumer healthcare apps using AI. Chatbots, for example, have epitomised the difficulty in regulation. Following complaints about Babylon's symptom checker the Care Quality Commission (CQC) referred them to the Medicine and Healthcare products Regulatory Agency (MHRA)⁷⁵. However, five public bodies were responsible for assessing the impact and safety of Babylon's products (CQC, MHRA, NHS Digital, NHS England, and Hammersmith and Fulham CCG), demonstrating the lack of clarity that currently exists in the oversight of technology products such as chatbots^{76 77}. Whilst the MHRA currently regulates medicines and medical devices, apps that use AI are not subject to assessment, certification, or approval by the MHRA, unless they fall under the device category. This situation does not provide the general public the clarity that is required. The problem with a patchwork quilt of regulators is that it does not provide the public with a single point of contact to raise their concerns or provide feedback.

Data protection regulations move more slowly than the speed of AI and technological developments, and should therefore cover the most important elements. That is why more agile frameworks such as the Code of Conduct for data-driven health and social care technologies⁷⁸ and the NICE Evidence Standards for digital health technologies⁷⁹ are steps in the right direction. Such flexibility is important and something rarely afforded by regulation.

The Code of Conduct for data-driven technology, developed by the Department for Health and Social Care and NHS England, represents a meaningful step for AI to be routinely used in healthcare,



and with further iterations can be a bastion for patient empowerment. The Code aims to provide a foundation to "deepen the trust between citizen, clinician and the wider health and care system [and] we should not wait to take this opportunity"⁸⁰. It aims to do this by outlining a set of principles that can help commissioners, innovators, healthcare professionals, patients, and the general public better work together and ensure outcomes are mutually beneficial. Some principles of the Code are pertinent to direct patient and public engagement, such as the need to understand users and being transparent about the limitations of the data and algorithms used. Indeed, the development of the Code is rooted in the importance of ensuring that transparency, accountability, liability, explicability, and fairness are realised⁸¹. This creates considerable opportunity to provide clarity for patients, and will benefit from patient-facing communication as to what the Code means for them.

At the same time, it is important not to add to the current regulatory burden that new innovations have to go through before they are used in clinical practice. Unnecessary regulation can stifle the development, spread, and adoption of innovations. Existing regulatory and advisory bodies such as the MHRA, NICE and NHS Digital are well placed to ensure that the clinical and health applications of AI are safe and effective for public use, albeit with additional resources and powers, as recommended by the House of Lords Artificial Intelligence Select Committee⁸². However, a single known point of contact for patients is needed to ensure that information and concerns can effectively be passed to and from patients. This has been successfully achieved in several sectors, for example with Ofwat.

Managing expectations

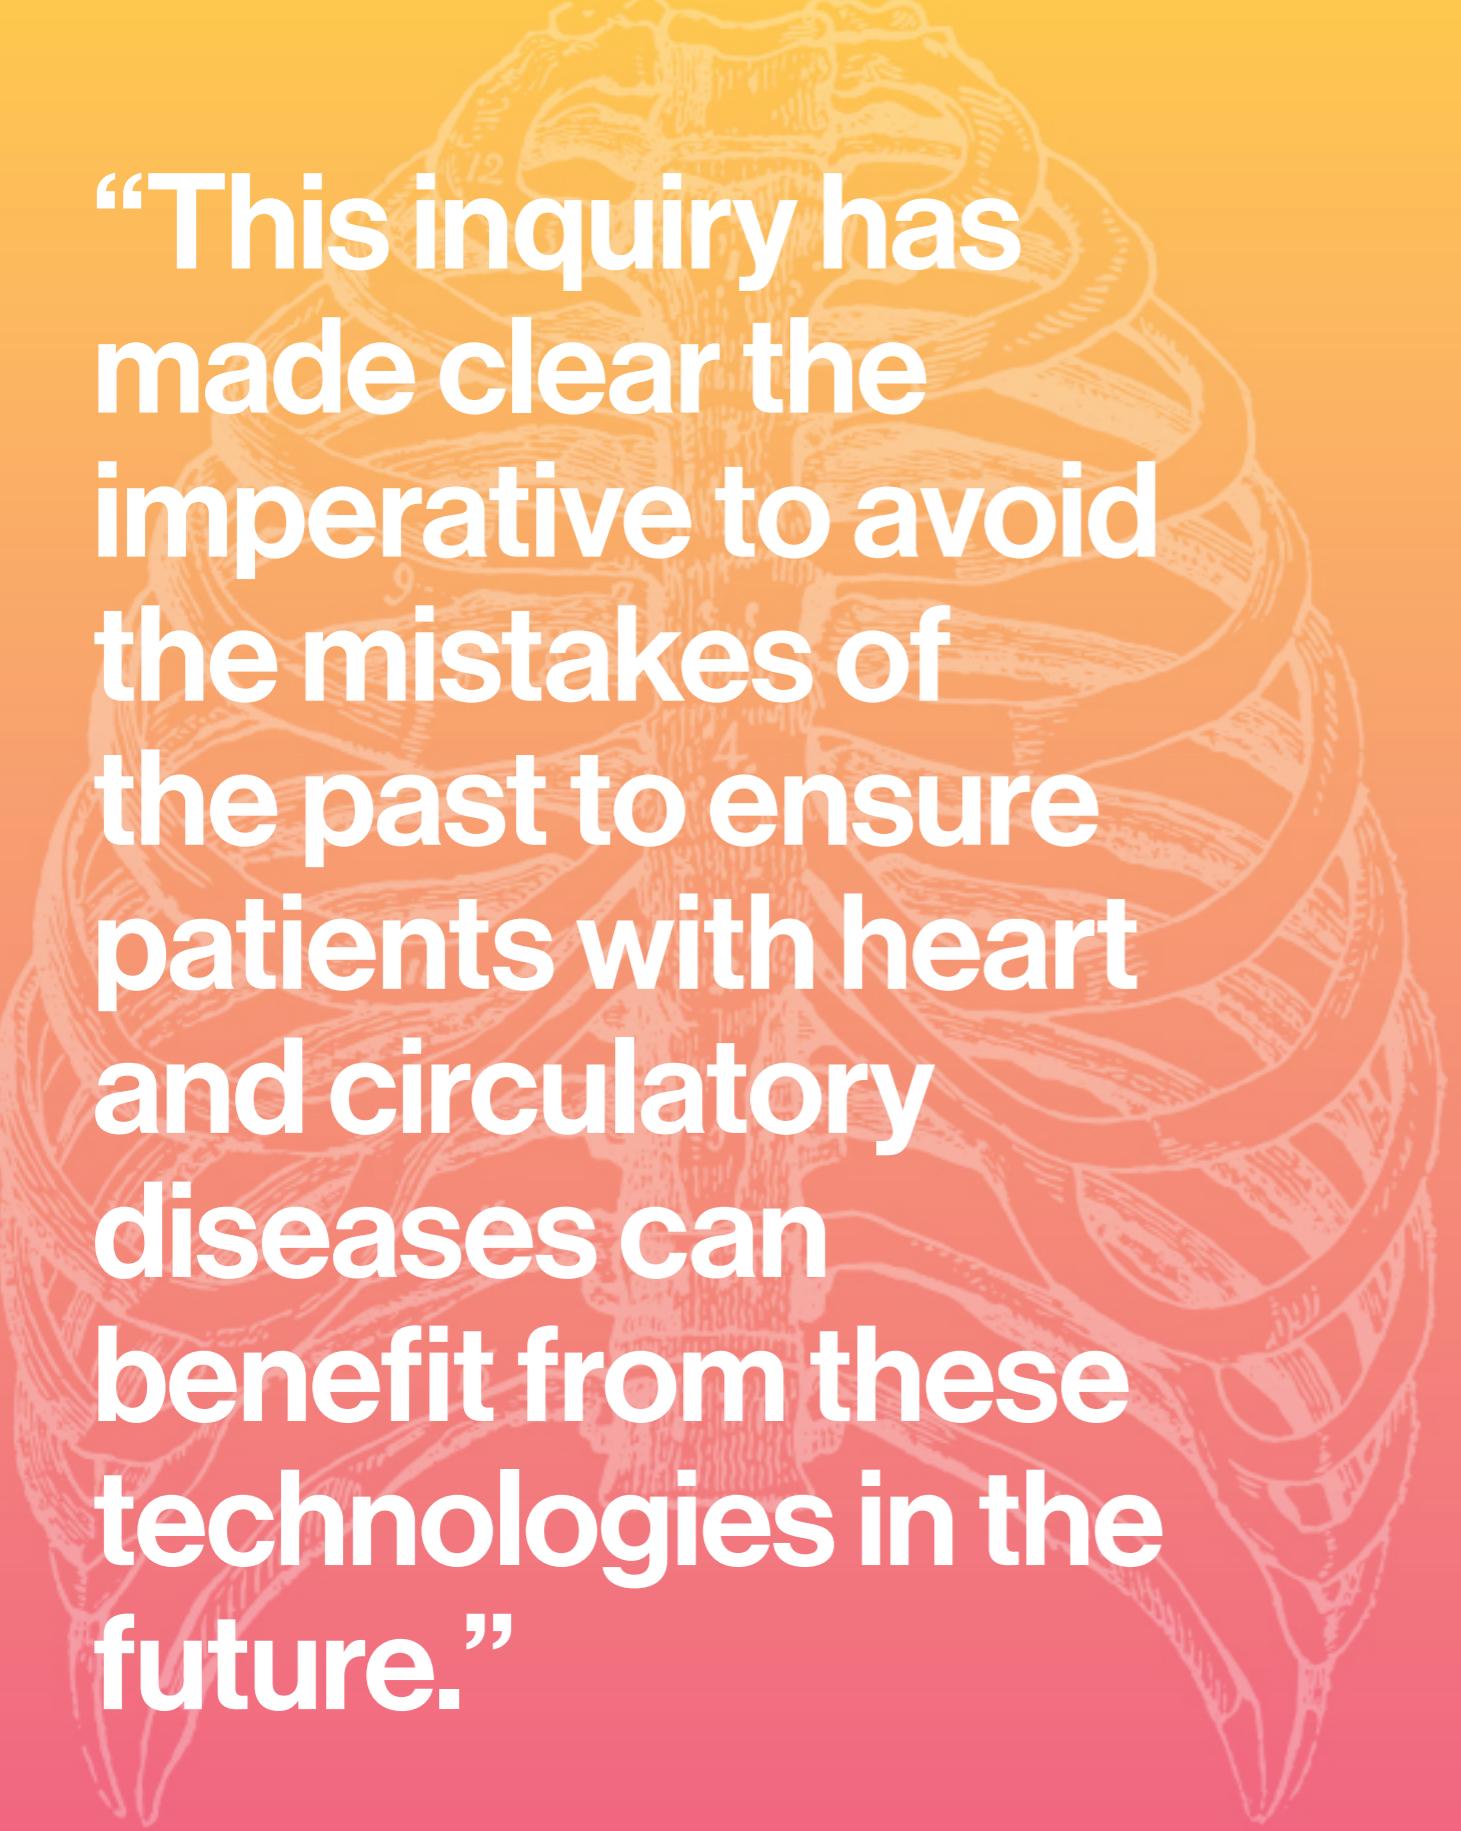
The public may be less willing to embrace AI if they think their data will be misused. At the same

time, the public may lose faith in AI if they feel that innovations are not being implemented quickly enough or are not bringing about the improvements that the media or policy makers might talk about. The public may also lose faith in AI in healthcare if they feel that it is being developed without their needs and outcomes in mind. A balancing act between managing expectations and encouraging hope and enthusiasm is always challenging but nevertheless important. When we say patients should be informed and clear on what AI can do for the NHS, it is not a tick-box measure. It is to provide the clarity that is needed for better diffusion of AI. The inquiry found that publication standards are needed for AI research. This is important because it helps the media, policy makers, clinicians, and the public identify hyperbole or even misinformation. The importance of better publication in AI research was also highlighted by Dr Eric Topol⁸³ and a major paper that was co-authored by representatives of Health Data Research UK (HDRUK), the Alan Turing Institute, and NICE⁸⁴. The latter paper set out 20 questions to begin a discussion of what good looks like in terms of AI research.

The need for standards in how research is designed, conducted, analysed, and communicated cuts across all research disciplines. Additionally, the thematic framework suggested by the authors of the paper (transparency, replicability, ethics, and evaluation) are relevant to other areas. However, within these themes, the paper sets out questions specific to AI research; for example, whether the computer and software resources are available in real-world settings⁸⁵. Such guiding questions are an important foundation to ensure that public information and debate is based on high quality research.

Recommendation

- 6** NHS England and NICE should encourage the development and use of reporting standards for AI research, in order to provide best practice for AI researchers. This can also provide a way for the media, policy makers, clinicians and the public to recognise quality in AI research.



“This inquiry has made clear the imperative to avoid the mistakes of the past to ensure patients with heart and circulatory diseases can benefit from these technologies in the future.”

Conclusion

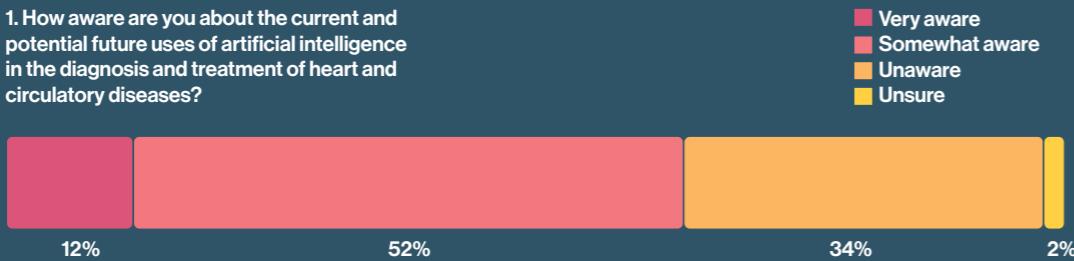
AI provides an opportunity to give patients more control over their health, through the use of wearables and personalised health recommendations. This in turn can change the way patients interact with the health system and how they understand and manage their condition. Despite this, patients have been largely left out of the discussion. Although the widespread use of AI in the NHS is some years away, it is already being used in consumer products such as chatbots. This report argues that policy makers, in conjunction with other stakeholders such as charities, industry, clinicians, and the research community, should begin engaging patients now about AI in healthcare. This engagement should not only take place at the policy level, but also at service design and technology implementation.

This inquiry has not been prescriptive as to the form and content of engagement with patients, as it is important for a wider range of stakeholders to be involved in this discussion. However, we have highlighted several issues that should at least be considered when planning how to engage patients with regards to AI in healthcare. Importantly, this inquiry has made clear the imperative to avoid the mistakes of the past to ensure patients with heart and circulatory diseases can benefit from these technologies in the future.

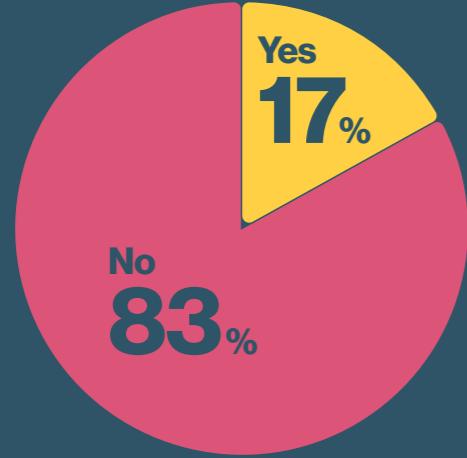
Appendix

Summary of patient responses

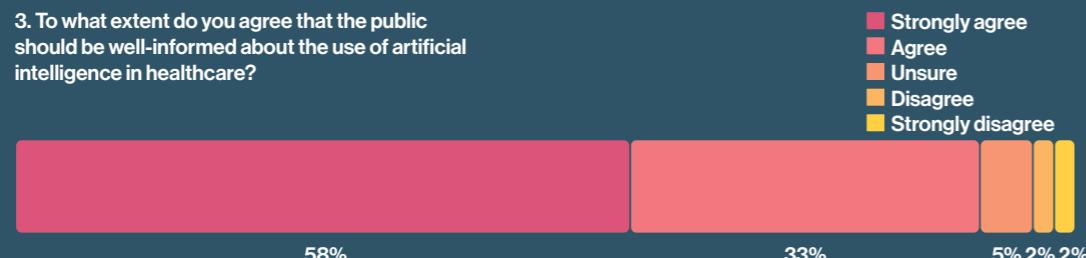
1. How aware are you about the current and potential future uses of artificial intelligence in the diagnosis and treatment of heart and circulatory diseases?



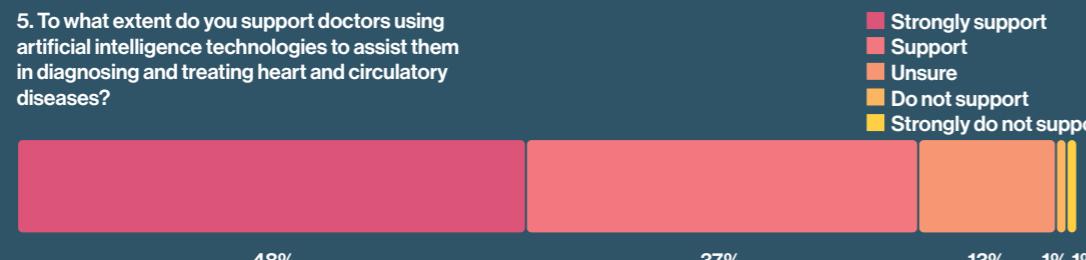
2. Are you aware of any current cases of artificial intelligence being used in the diagnosis and treatment of heart and circulatory diseases?



3. To what extent do you agree that the public should be well-informed about the use of artificial intelligence in healthcare?



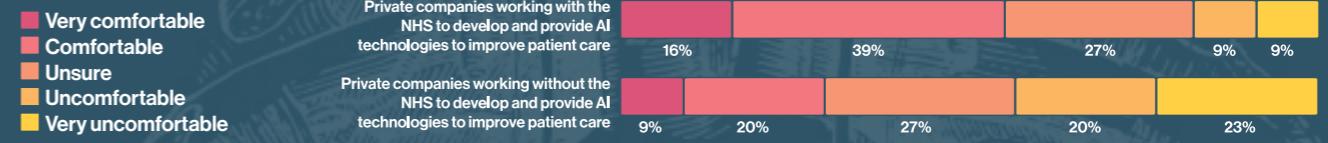
5. To what extent do you support doctors using artificial intelligence technologies to assist them in diagnosing and treating heart and circulatory diseases?



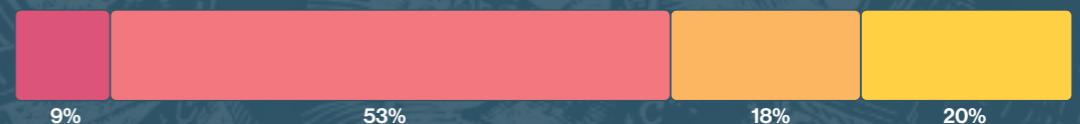
6. Artificial intelligence technologies can use your anonymised personal health data to help better diagnose medical conditions in other people. How would you feel about your anonymised data being used for this purpose?



7. How comfortable would you be with the following institutions using your anonymised data in AI technologies?



8. How confident are you in the security of anonymised personal health data being used in artificial intelligence technologies?



Scenario-based support questions:

In the final section of the survey, respondents were presented with a hypothetical scenario designed to assess how patients may feel about AI technologies once they are more widespread in the health system, and how AI may impact the relationship between patient and doctor. In short, our findings show that patients anticipate that—even when AI is well-proven and widespread in healthcare—their preference will be for AI to support doctors in making their decisions, and that they will feel uncomfortable about decisions about their healthcare being taken without human input. However, advancements in AI will also create expectations that doctors will consult such technologies if they are available, and feel uncomfortable if doctors do not.

"In ten years' time, artificial intelligence technologies are commonly used across UK hospitals to help doctors diagnose and treat heart and circulatory diseases. In practice, these technologies have been shown to be able to independently diagnose a wide range of heart and circulatory conditions better than most human doctors working without access to artificial intelligence technologies."

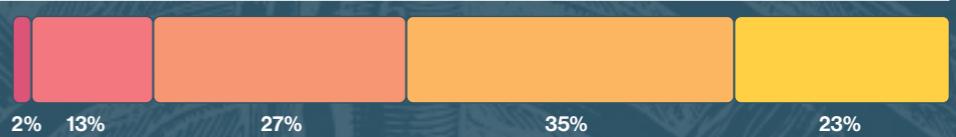
9. In this scenario, how comfortable would you be if your doctor chose not to consult artificial intelligence technologies when diagnosing your medical condition?



10. In this scenario, how comfortable would you be being diagnosed by a human doctor who used their own judgement, along with recommendations of an artificial intelligence assistant to help inform their decision?



11. In this scenario, how comfortable would you be being diagnosed solely by an artificial intelligence technology, without the input of a human doctor?





- 1 Topol E (2019) 'Preparing the workforce to deliver a digital future: An independent report on behalf of the Secretary of State for Health and Social Care', P.55. Health Education England, February 2019**
- 2 The AHSN Network (2018) 'Accelerating artificial intelligence in health and care: results from a state of the nation survey'.**
- 3 Wellcome Trust (2018) 'Ethical, Social, and Political Challenges of Artificial Intelligence in Health' London: Wellcome Trust, 2018.**
- 4 House of Lords Artificial Intelligence Select Committee (2018) 'AI in the UK: ready, willing, and able?', HL Paper 100**
- 5 Loder J and Nicholas L (2018) 'Confronting Dr Robot: Creating a people-powered future for AI in health' London: Nesta, 2018**
- 6 House of Lords Artificial Intelligence Select Committee (2018) 'AI in the UK: ready, willing, and able?', HL Paper 100**
- 7 The Office for AI is based in the Department for Digital, Culture, Media and Sport and the Department for Business, Energy and Industrial Strategy. It is tasked with overseeing the implementation of the UK Government's AI strategy.**
- 8 The industry body set up as a response to the Government's Industrial Strategy and tasked with improving the adoption of AI across sectors.**
- 9 Centre for Data Ethics and Innovation**
- 10 NHS England (2019) 'The NHS Long Term Plan', p.91**
- 11 The Code of Conduct is a set of ethical and behavioural principles for those developing or using data-driven technologies.**
- 12 Department of Health and Social Care (2018) 'Code of conduct for data-driven health and care technology'**
- 13 Predictive analytics is the use of data, statistical algorithms and machine learning techniques to identify the likelihood of future outcomes based on historical data.**
- 14 Markoff J (2012) 'How many computers to identify a cat? 16,000', New York Times, June 25, 2012**
- 15 Esteva A et al (2017) 'Dermatologist-level classification of skin cancer with deep neural networks', Nature, 542, 115-118**
- 16 Poplin R et al (2018) 'Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning', Nature Biomedical Engineering, 2, 158-164**
- 17 BHF-Turing Cardiovascular Data Science Awards**
- 18 The Academy of Medical Sciences (2018) 'Our data-driven future in healthcare: people and partnerships at the heart of health related technologies', London: The Academy of Medical Sciences, 2018**
- 19 House of Lords Artificial Intelligence Select Committee (2018) 'AI in the UK: ready, willing, and able?', HL Paper 100**
- 20 House of Lords Artificial Intelligence Select Committee (2018) 'AI in the UK: ready, willing, and able?', HL Paper 100**
- 21 Batalden M et al (2016) 'Coproduction of healthcare service', BMJ Qual Safe, 2016; 25:509-517**
- 22 Foot C et al (2014) 'People in control of their own health and care: the state of involvement', London: The King's Fund, 2014**
- 23 Aitken M, Tully MP et al (2019) 'Consensus Statement on Public Involvement and Engagement with Data', International Journal of Population Data Science, 2019; 4:1:06**
- 24 The Royal Society (2018) 'Portrayals and perceptions of AI and why they matter', London: The Royal Society, 2018**
- 25 Ewen SW, Puszta A (1999) 'Effects of diets containing genetically modified potatoes expressing Galanthus nivalis lectin on rat small intestine', Lancet 354: 353-354**
- 26 Burke D (2004) 'GM food and crops: what went wrong in the UK?', EMBO Rep, 2004 5(5): 432-436**
- 27 Chen ZL, Gu H et al (2003) 'Safety assessment for genetically modified sweet pepper and tomato', Toxicology 188: 297-337**
- 28 The Royal Society (2018) 'Portrayals and perceptions of AI and why they matter', London: The Royal Society, 2018**
- 29 Burke D (2004) 'GM food and crops: what went wrong in the UK?', EMBO Rep, 2004 5(5): 432-436**
- 30 Care.data was a programme aimed at uploading patient data from GPs into a single central database.**
- 31 Triggle N (2014) 'Care.data: How did it go so wrong?'**
- 32 Understanding Patient Data 'Is the new national data opt-out just care.data all over again?'**
- 33 Godlee F (2016) 'What can we salvage from care.data?', BMJ 2016;354:i3907**
- 34 Understanding Patient Data 'Is the new national data opt-out just care.data all over again?'**
- 35 Van Staa T-P (2016) 'Big health data: the need to earn public trust', BMJ 2016;354:i3636**
- 36 The Academy of Medical Sciences (2018) 'Our data-driven future in healthcare: people and partnerships at the heart of health related technologies', London: The Academy of Medical Sciences, 2018**
- 37 Understanding Patient Data (2018) 'Evaluating Understanding Patient Data'**
- 38 The AHSN Network (2017) 'About Academic Health Science Networks'**
- 39 The AHSN Network (2018) 'Accelerating artificial intelligence in health and care: results from a state of the nation survey'**
- 40 Royal Society of Arts, Manufactures and Commerces (2018) 'Artificial intelligence: real public engagement', London: RSA, 2018**
- 41 APPG on Heart and Circulatory Diseases – 2019 Inquiry, Roundtable 1 [patient representatives]**
- 42 APPG on Heart and Circulatory Diseases – 2019 Inquiry, Roundtable 1 [patient representatives]**
- 43 Topol E (2019) 'Preparing the workforce to deliver a digital future: An independent report on behalf of the Secretary of State for Health and Social Care', Health Education England, February 2019**
- 44 Davies S (2018) 'Annual Report of the Chief Medical Officer: Health 2040 – Better health within reach', Crown Copyright: Department of Health and Social Care, 2018**
- 45 Shcherbina A et al (2017) 'Accuracy in Wrist-Worn Sensor-Based Measurements of Heart Rate and Energy Expenditure in a Diverse Cohort', Journal of Personalized Medicine, 2017**
- 46 Feldman R (2017) 'What does the future hold for wearables?'**
- 47 Rogers E (2003) 'Diffusions of Innovations, 5th Edition', Free Press, 2003**
- 48 Pew (2018) 'Mobile Fact Sheet'**
- 49 OFCOM (2018) 'Adults' Media Use and Attitudes Report'**
- 50 Office for National Statistics (2018) 'Internet access – households and individuals, Great Britain: 2018'**
- 51 NHS Digital (2018) 'Digital inclusion guide for health and social care'**
- 52 Future Advocacy (2019) 'Artificial Intelligence: maximising opportunities, minimising risks'**
- 53 Smart W (2018) 'Lessons learned review of the WannaCry ransomware cyber attack', Crown Copyright**
- 54 Information Commissioner's Office (2018) 'Bupa fined £175,000 for systemic data protection failures'**
- 55 APPG on Heart and Circulatory Diseases – 2019 Inquiry, Roundtable 2 [expert panel]**
- 56 Harwich E and Lasko-Skinner R (2018) 'Making NHS data work for everyone', London: Reform**
- 57 O'Regan, D (2019, January 4). Personal interview**
- 58 Murray, J (2018) 'Training AI with Fake Data: A Flawed Solution?', Binary District**
- 59 Mukherjee S (2017) 'AI versus MD: what happens when diagnosis is automated?', Annals of Medicine April 3**
- 60 Park D et al (2018) 'Multimodal explanations: justifying decisions and pointing to the evidence'**
- 61 Lundberg S and Lee S (2017) 'A Unified Approach to Interpreting Model Predictions', NIPS Proceedings, 2017**
- 62 Ribeiro M et al (2016) 'Local Interpretable Model-Agnostic Explanations (LIME): An Introduction'**
- 63 PwC (2018) 'Explainable AI: Driving business value through greater understanding'**
- 64 Pande V (2018) 'Artificial intelligence's 'black box' is nothing to fear', New York Times, January 25, 2018**
- 65 Google AI Blog (2017) 'Federated Learning: Collaborative Machine Learning without Centralized Training Data'**
- 66 Hosanagar K and Jair V, (2018) 'We Need Transparency in Algorithms, But Too Much Can Backfire', Harvard Business Review**
- 67 House of Commons Science and Technology Committee (2018) 'Algorithms in decision making', HC Paper 351**
- 68 Science and Technology Committee (2019) 'Oral evidence: Work of the Centre for Data Ethics and Innovation', HC Paper 1927**
- 69 Davies S (2018) 'Annual Report of the Chief Medical Officer: Health 2040 – Better health within reach', Crown Copyright: Department of Health and Social Care, 2018**
- 70 A subclinical disease has no recognisable clinical findings. It is distinct from a clinical disease, which has signs and symptoms that can be recognised.**
- 71 Med Life Crisis (2018) 'Cardiologist's scientific analysis of the Apple Watch'**
- 72 Hall W and Pesenti J (2017) 'Growing the artificial intelligence industry in the UK', Department for Digital, Culture, Media & Sport and Department for Business, Energy & Industrial Strategy**
- 73 Harwich E and Laycock K (2018) 'Thinking on its own: AI in the NHS', London: Reform**
- 74 House of Lords Artificial Intelligence Select Committee (2018) 'AI in the UK: ready, willing, and able?', HL Paper 100**
- 75 Heather B (2018) 'Safety regulators reviewing concerns about Babylon's 'chatbot', Health Service Journal**
- 76 McCartney M (2018) 'AI in health must be strictly tested', BMJ, 2018; 361:k1752**
- 77 Evans H (2017) 'Regulators against the machine', British Journal of Healthcare Computing**
- 78 Department of Health and Social Care (2018) 'Initial code of conduct for data-driven health and care technology'**
- 79 NICE (2018) 'Evidence standards framework for digital health technologies'**
- 80 Department of Health and Social Care (2018) 'Initial code of conduct for data-driven health and care technology'**
- 81 Department of Health and Social Care (2018) 'Initial code of conduct for data-driven health and care technology'**
- 82 House of Lords Artificial Intelligence Select Committee (2018) 'AI in the UK: ready, willing, and able?', HL Paper 100**
- 83 Eric J. Topol, (2019) 'High-performance medicine: the convergence of human and artificial intelligence', Nature Medicine volume 25, pages 44–56**
- 84 Vollmer S et al (2018) 'Machine learning and AI research for Patient Benefit: 20 Critical Questions on Transparency, Replicability, Ethics and Effectiveness', arXiv:1812.10404**
- 85 Vollmer S et al (2018) 'Machine learning and AI research for Patient Benefit: 20 Critical Questions on Transparency, Replicability, Ethics and Effectiveness', arXiv:1812.10404**



on
heart and
circulatory
diseases

