September 2017

Response submitted by:

Sobia Raza
sobia.raza@phgfoundation.org

Alison Hall
alison.hall@phgfoundation.org

Tanya Brigden
tanya.brigden@phgfoundation.org

Inquiry on artificial intelligence

Whilst the field of A.I. has existed for decades, in recent years significant advances in this area have expanded the range of A.I. based applications. Our organisation is specifically interested in the use of artificial intelligence for health and healthcare and the associated opportunities, risks, ethical and social implications, and wider policy considerations. Our responses to this inquiry are therefore in the context of A.I. for health.

The PHG Foundation is a not-for-profit think tank that aims to make science work for health, with twenty years' experience in issues surrounding the responsible and effective use of genomics within health services.

The pace of technological change

What is the current state of artificial intelligence and what factors have contributed to this? How is it likely to develop over the next 5, 10 and 20 years? What factors, technical or societal, will accelerate or hinder this development?

Accepting there is not a universally accepted definition of A.I., in this consultation response we use the term to denote the development and use of computing systems concerned with making machines work in an intelligent way, including those that iteratively learn from data to improve their performance with experience.

A.I. already underpins a plethora of mainstream technologies across many life domains e.g. web search engines, fraud detection, marketing systems. Rapid developments in associated technologies and sub-areas of A.I. such as machine learning, computer vision and natural language processing, combined with the increasing availability of 'big data' are expanding the prospective applications of A.I. and transforming previously hypothetical uses into more tangible prospects; autonomous vehicles are one case in point.

In health and medicine there is great scope for A.I. based applications to be developed using patient health records and other health related datasets.

In health and medicine there is great scope for A.I. based applications to be developed using patient health records and other health related datasets.

When applied to these datasets, the view is that A.I. approaches may:

- Help to identify new disease biomarkers and refine understanding of disease
- Be used to make predictions about health and disease risk and potentially to stratify populations according to these predictions to better target appropriate interventions
- Inform and underpin new medical diagnostics, helping to develop more targeted treatments, and treat and manage patients and individuals on a more 'personalised' basis.

A.I. approaches are already beginning to demonstrate potential utility for very specific medical applications, examples including:

- Automation of medical image analysis e.g. in radiology
- Risk management support tools e.g. to identify patients at high risk of hospital readmission, or acute kidney injury

Beyond these highly targeted applications, the wider and large-scale use of A.I. in health is further from realisation due to practical, technical, and societal factors.

In our view, the key factors that are most likely to impact on the pace of A.I. based developments in health include:

- Data availability for 'training' i.e. developing A.I. based algorithms
- Cross-sector collaboration – particularly between the computing (A.I.) and the healthcare and medical research domains
- The ability to collate enriched health datasets and share data within and between those sectors collaborating to develop health related A.I. applications
- The challenge in securing public trust in sharing health data, particularly with private sector developers
- Difficulties in predetermining user perception and preference concerning A.I. based health devices – especially where the tools interface directly with patients and the publics

Is the current level of excitement which surrounds artificial intelligence warranted?

There is a great degree of excitement and discourse surrounding the potential impact of A.I. in health and medicine, which stems from the potential to derive new insights from health datasets. Whilst A.I. does hold great promise to benefit patients and health systems, we believe the current levels of excitement should be tempered by the immediate practical challenges and wider considerations to developing health related and medical A.I. applications, these include:

- Technical obstacles to obtaining health data sets: not least due to the slow pace of health record digitisation, but also the lack of data standards

and interoperability

- Technical challenges to collating citizen generated health-relevant data (e.g. from wearables and monitors) and integrating this with health records
- The need for greater collaboration between A.I. experts and medical professionals in order to better define and prioritise the areas to which A.I. could be applied
- Uncertainly surrounding the impact of upcoming regulatory changes (such as the EU General Data Protection Regulation (GDPR) coming into force) on the legitimacy of data processing and data profiling
- Uncertainty regarding the implementation of the proposals set out by the National Data Guardian (for Health and Social Care) on 'data security and consent and 'opt-outs' recently accepted by the Government , and specifically the impact of an 'opt-out' on the availability and completeness of datasets , as this will influence the ability to develop and use A.I. tools which can serve a diverse U.K. population.

Impact on society

How can the general public best be prepared for more widespread use of artificial intelligence?

Across a range of sectors the more widespread use of A.I. is expected to impact upon the current job market, including in healthcare. Whilst some forms of health employment may be displaced by A.I. technologies, there is also the potential for new types of employment to be created, and the need for collaboration between health professionals and the A.I. sector will be increasingly important.

Whilst A.I. does hold great promise to benefit patients and health systems, we believe the current levels of excitement should be tempered by the immediate practical challenges and wider considerations to developing health related and medical A.I.

The public must therefore be prepared for an A.I. integrated healthcare workspace. This will require education and training to place greater focus on skillsets that arguably cannot easily be displaced by A.I. such as creativity, effective social interaction, manual dexterity and intelligence.

As the future job market may be much more fluid, support and incentives for life-long learning will be important to enable healthcare workers to acquire new skills and retrain for new work.

To prepare the public for the future widespread use of A.I. it is crucial to provide accessible information and ongoing engagement that highlights the existing use of A.I. in many domains of life, including its emerging use for health and healthcare. Encouraging awareness surrounding current uses of the technology may help dissolve misconceptions that fuel opposition. Early engagement, and raising awareness around the potential of A.I. to support, inform and improve healthcare, will prepare the public and health professionals for more extensive interactions with A.I. in the future.

Public perception

Should efforts be made to improve the public's understanding of, and

Since the development of health based A.I. applications will require collaboration between different sectors, it will be important to embed appropriate frameworks that can both support cross sector data sharing.

engagement with, artificial intelligence? If so, how?

Since the development of A.I. applications for healthcare is contingent on data availability, public discourse around A.I. should be accompanied by greater engagement around the benefits and risks of sharing health datasets, and a concerted effort to build public trust for sharing health data.

Within the UK, Understanding Patient Data, set up following the National Data Guardian's Review of consent and opt-outs is one important initiative to support conversations with the public, patients, and healthcare professionals about the uses of health information for care. To continue to improve awareness and engagement it is crucial that such efforts are an ongoing rather than a transient programme of work.

The Global Alliance for Genomics and Health (GA4GH) is an organisation committed to improved genomic data sharing on a global basis particularly for medical research. It is developing a number of demonstration projects that explore the use of A.I. to facilitate effective data sharing.

Since the development of health based A.I. applications will require collaboration between different sectors, it will be important to embed appropriate frameworks that can both support cross sector data sharing and also build and reinforce public trust through transparency and engagement about how health data are used.

In the context of healthcare, as A.I. applications develop they have great potential in the future to underpin, inform or support medical enquiries, diagnoses, health monitoring and tailored care. If integrated effectively, there is the opportunity for A.I. to not only enable greater healthcare personalisation, but also alleviate some of the current pressures on the health system. The success of these transformative technologies will in part rely upon the publics' and health professionals' willingness to use them. To realise the benefits of A.I. in health and medicine it will be crucial to encourage public and health professional engagement and provide a factual and transparent view of how developments in A.I. technologies facilitate better health.

What are the ethical implications of the development and use of artificial intelligence? How can any negative implications be resolved?

If the datasets used for developing (training) A.I. algorithms which underpin health applications are not sufficiently representative of the populations they are intended to serve, then it is possible the A.I. predictions may not function correctly for sections of the population underrepresented in the 'training' sets.

1. To mitigate against potential disparities, it will be crucial for policy makers and those developing A.I. based tools for healthcare to carefully consider population diversity when collating datasets for developing A.I. algorithms

2. The objective of equity should therefore be taken seriously by the sector. For example, questions about securing equitable access to research are

already included as part of the NHS research ethics review process, and could be replicated within this sector.

It is possible that issues of liability may arise if incorrect health / medical predictions are made based on tools underpinned by A.I.

There is currently a lack of clarity in the literature surrounding who will be liable for errors made through use of A.I. tools. Such errors will be inevitable, especially at early stages of development. Mechanisms will need to be developed which address this problem. For example, extension to existing NHS Indemnity could be employed (where the NHS adopts liability for negligent acts of professionals employed or owing a duty of care), or something similar adopted whereby the risks to users is shared between the manufacturer and the health service.

Since the A.I. approaches can reveal novel insights within datasets, mechanisms for dealing with incidental health findings may be necessary.

1. There is considerable debate about the extent to which use of novel technologies such as whole genome sequencing creates an ethical obligation to actively search for additional clinically actionable findings and/or to validate and treat any unsolicited incidental health findings that may arise through use of these technologies.

2. Similar challenges are likely to arise in the context of A.I. Thresholds for reporting potentially actionable findings will need to be identified; validation and reporting obligations evaluated; pathways clarified; and funding secured.

3. If these technologies are used by health care professionals, there will also be a need to assess how these technologies impact upon existing professional duties and responsibilities (both ethical and legal). If technologies are used for self-testing, then routes for further advice/ action need to be clearly articulated.

The Information Commissioner's Office has powers to prepare appropriate codes of good practice and we suggest that guidance could be developed to clarify what constitutes best practice when processing data, including specific guidance for transparency/ accountability.

In what situations is a relative lack of transparency in artificial intelligence systems (so called 'black boxing') acceptable? When should it not be permissible?

In the health sector, recent regulatory changes will necessitate increased transparency, particularly where algorithms are used for diagnosis or risk prediction. We welcome these changes to the extent that they ensure that such algorithms are used in ways that are safe and effective for patients and consumers.

Some A.I. algorithms are already regulated under the EU In Vitro Diagnostic Devices Directive 1998, but the scope of regulation will increase pursuant to the EU In Vitro Diagnostic Devices Regulation (2017) which comes into force in May 2022.

Under this Regulation, standalone medical software used for certain purposes will be regulated as IVD devices, and in order for them to be

Empirical work on public attitudes and commercial access to data has suggested that understanding the broad uses of data, and who will be involved are seen as being even more important than ensuring that effective regulation and safeguards are in place.

placed on the market within the EU, will have to satisfy requirements for clinical performance, performance evaluation, labelling and information provision. This will require developers and manufacturers to clearly articulate the uses for which A.I. algorithms will be put, and for the algorithms to have demonstrable clinical utility within a designated clinical population. Compliance with this Regulation is likely to be challenging for the sector.

The role of government

What role should the Government take in the development and use of artificial intelligence in the United Kingdom? Should artificial intelligence be regulated? If so, how?

As mentioned in answer to the previous question above, the UK Government has confirmed that it will be implementing the EU IVDR and the EU GDPR since these Regulations come into force before the UK exits the EU: the EU IVDR directly regulates algorithms that are used for certain health related purposes since such algorithms are classified as in vitro medical diagnostic devices.

The EU GDPR (to be implemented in May 2022) specifically regulates profiling which is defined as automated processing of personal data for certain applications including health (GDPR Article 4(4)) . Article 13(2)(f) of this Regulation requires data controllers using profiling to disclose 'the existence of automated decision-making' and 'meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject' and Article 22 clarifies the legal bases under which such processing may be lawfully undertaken.

The scope of the GDPR regulates personal data (including some pseudonymised data). More detailed guidance is currently being prepared by the Information Commissioner's Office: depending on what this concludes, there may be a need for additional regulation of areas that fall outside of the EU IVDR and EU GDPR.

Facilitating effective governance and regulation as one of the means in which public trust and confidence can be facilitated, however empirical work on public attitudes and commercial access to data has suggested that understanding the broad uses of data, and who will be involved are seen as being even more important than ensuring that effective regulation and safeguards are in place.