Greenbank

AI and health:
Risks, opportunities and considerations

On Thursday 27 June 2024, Greenbank hosted a roundtable for charities on the topic of AI and health. This document is an overview of the topic and a summary of the discussion.

Contents

Speakers

Dariel Burdass
CEO
The Physiological Society

Dariel Burdass has been CEO of The Physiological Society since 2016.
The Physiological Society brings together around 3,500 scientists from over 60 countries and its membership has included more than 20 Nobel Prize winners. Prior to this Dariel spent 25 years at the Microbiology Society, progressing to Deputy Chief Executive and Director of Strategy and Communications. Dariel has a degree in microbiology and has published several textbooks on microbiology. She is a Trustee of the Science Council and FSRH (Faculty of Sexual and Reproductive Healthcare) and sits on the Finance, Audit and Risk Committee.

Professor Colin Gavaghan
Professor of Digital Futures
Bristol Digital Futures Institute, University of Bristol Law School

Colin Gavaghan is Professor of Digital Futures at the University of Bristol, based between the Bristol Digital Futures Institute and the Law School. Prior to that, he was the inaugural holder of the Law Foundation Chair of Law & Emerging Technologies at the University of Otago in New Zealand. His research interests are in health law and the regulation of emerging technologies. His most recent book was the co-authored Citizen's Guide to Artificial intelligence (MIT Press, 2021). In addition to academic activities, Colin has advised the New Zealand government on digital technologies, served as the first chair of NZ Police's Expert Panel on Emerging Technologies, and spent several years on New Zealand's regulatory body for assisted reproduction.

Dr Bilal Mateen, MPH, MBBS
Executive Director
Digital Square

Dr. Bilal Mateen is the Executive Director for Digital Square, a global initiative that supports digital health transformation in low- and middle-income countries. In this role, he supports ministries of health across Africa and Asia with effective planning, procurement, and implementation of technology to help address the global health equity gap. Bilal is a physician by training with an academic background in health-related applications of data science and machine learning (PI on almost a quarter-of-a-billion in grants and contracts); in recognition of the contributions Bilal has made to the field of AI4health, he holds an honorary professional appointment at University College London and a fellowship at the Alan Turing Institute (the UK's national institute for data science and AI). Previously, he served as the Clinical Technology Lead and Senior Manager for Digital Technology at the Wellcome Trust, where his team provided funding (>£100M per annum) for digital public goods to address mental health, the impacts of climate change on health and escalating infectious diseases. Bilal is passionate about the role of research in advancing ethical and inclusive applications of AI to improve human lives and serves as an investment adviser to several philanthropic foundation and impact investing organisations, Newcastle University and is a Chartered Ecologist.

Kai Johns
Senior Ethical, Sustainable and Impact Researcher
Greenbank

Kai conducts analysis of investments against a range of environmental, social and governance criteria for both new ideas and companies in Greenbank's investment universe. He uses a range of data sources and integrates them into the Greenbank research process. Kai's focus is on climate impact and risk assessment across the investment portfolios Greenbank manages and his areas of interest include climate solutions, net zero and smart cities. Kai joined the ethical, sustainable and impact research team in March 2019 after graduating from the University of Cambridge with a BA in Law.

Key takeaways

1. AI hype may be justified – but tread warily.

2. There is a tension between getting tools out quickly and adequate testing, bias removal, and regulation.

3. The expectation of better outcomes from AI may be too high. Are we expecting too much for AI? No drug works 100% of the time and human diagnosis is not 100% accurate. What is the comparator? The perfect human doesn't exist, however is 'just a bit better than a human' enough?

4. AI can sometimes fail in sensitive situations. Does AI deprive us of the human touch in the healthcare setting? Are chatbots smart enough to escalate when necessary?

5. We need to be more intentional and avoid biases by including diversity in the room at creation, otherwise we won't see issues until a product hits the market. Only with assessment and evaluation through randomised clinical trials can that be determined.

6. Cost effectiveness will be a factor in the development of AI products and diagnostic tools.

An introduction to the ethical risks and implications of AI in healthcare

Colin Gavaghan
Professor of Digital Futures
Bristol Digital Futures Institute
University of Bristol Law School

Dr Colin Gavaghan introduced the roundtable session by providing an overview of developments in Artificial Intelligence (AI) and healthcare. Although the use of AI in healthcare has the potential to be highly transformative, successfully and ethically integrating it into healthcare practice will require detailed consideration, agreement on best practices, and regulatory oversight. For example, one ethical implication of using AI is understanding and removing the cognitive bias already found in its existing and future outputs. AI is susceptible to cognitive bias as a result of human biases being translated into machine learning algorithms[1].

Many parts of society are embracing or indeed 'jumping on the bandwagon' of AI, but this increased attention and rapid proliferation of its use serve to highlight a requirement for suitable regulation and governance. Despite these concerns, the AI hype is not unjustified, particularly given its potential benefits when applied to healthcare. For example:

— In the area of stroke and cancer assessment, AI may be able to outperform humans in its ability to spot cancer indications early, preventing unnecessary interventions where indications are not malignant.

— AI has the potential to act as a triage facility, enabling targeted support systems and resource allocation which could improve patient outcomes, particularly for older patients and those with complex histories.

— AI could be very helpful as a support system for responding to chronic worldwide shortages in the provision of healthcare. If applied effectively, it could address gaps, support humans, and possibly take over roles previously done by humans.

[1] A machine learning algorithm is a set of rules or processes used by an AI system to conduct tasks, most often to discover new data insights and patterns, or to predict output values from a given set of input variables. Algorithms enable machine learning to learn. Source: IBM, **ibm.com/topics/machine-learning-algorithms**

While these applications of AI in healthcare systems are positive in theory, in practice there are ethical and governance concerns which raise questions and require complex consideration. For example:

— **Transparency**
How can we ensure transparency in AI's decision-making? A core healthcare principle is informed consent[2] but this is brought into question if we cannot explain the decision-making process of AI and the intervention or procedure recommended as a result.

— **Accuracy and bias**
What if the AI system is trained on poor datasets? Healthcare organisations may have a high quantity of data, but the quality of that data may not be fit for purpose. For example, a study by the Office for National Statistics (ONS) revealed that there are significant errors within the ethnicity data of patients[3]. Furthermore, medical data is sensitive information that requires permission from the patient before it can be used. If whole cohorts do not grant this permission, the AI dataset will be non-representative, which could likely lead to different health outcomes as a result of using it.

— **Human touch**
What are the consequences of compromising on the human touch? AI can sometimes fail in sensitive situations. For example, failures of AI chatbots to escalate serious abuse and other issues underscore the limitations of the technology as it stands.

These issues apply across a range of sectors that are looking to harness AI, which is why regulators have looked to intervene. Regulatory approaches vary, with the EU proposing specific AI regulations whilst the UK rely on existing frameworks. The EU AI Act proposes bespoke rules for AI, categorising them by risk and imposing stringent requirements for high-risk applications. This includes mandates for transparency, accuracy, and human control. In the UK context, existing regulatory bodies such as Ofsted is responsible for governing the use of AI in education, and the Medicines and Healthcare products Regulatory Agency (MHRA) has set out a strategic approach to the use of AI in medicine and science[4]. It remains to be seen which method will prove to be more effective, but regulation lag is a key risk. When the pace of innovation is so high, regulation will need to keep up if it is to safeguard us effectively.

When the pace of innovation is so high, regulation will need to keep up if it is to safeguard us effectively.

[2] Informed consent is the process in which a health care provider educates a patient about the risks, benefits, and alternatives of a given procedure or intervention. The patient must be competent to make a voluntary decision about whether to undergo the procedure or intervention. Source: National Library of Medicine, **ncbi.nlm.nih.gov/books/NBK430827**

[3] Source: Office for National Statistics, **ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthinequalities/articles/understandingconsistencyofethnicitydatarecordedinhealthrelatedadministrativedatasetsinengland2011to2021/may2024**

[4] Source: Medicines and Healthcare products Regulatory Agency, **gov.uk/government/news/mhras-ai-regulatory-strategy-ensures-patient-safety-and-industry-innovation-into-2030**

The growing importance of AI in healthcare

Dariel Burdass
CEO
The Physiological Society

Dariel Burdass provided an overview of what AI is, how AI learns and how it could be beneficial in physiology.

AI is an umbrella concept to describe a range of digital technologies. Although AI has been around for some time, it is ChatGPT that has caught wider attention. AI systems such as ChatGPT fall into the category of generative AI, which enables users to generate new content based on a variety of inputs or instructions. The content is created by a process of identifying patterns and structures within existing content, which it then applies to generate the new content.

AI is bigger than just generative AI. Predictive AI, as its name suggests, can be used to make predictions by taking in and assessing historical data. Both generative and predictive AI use machine learning – a method for training a system to identify data without human intervention – but by different processes for distinct outputs.

As a key risk of AI comes from poor datasets resulting in poor predictions, the quality and availability of data is essential. Considering this, the European Commission's strategy for data[5] focuses on a human-centric approach to the development of technology that upholds and promotes 'European values' in digital rights.

> "The availability of data is essential for training artificial intelligence systems, with products and services rapidly moving from pattern recognition and insight generation to more sophisticated forecasting techniques and, thus, better decisions."
>
> "Moreover, making more data available and improving the way in which data is used is essential for tackling societal, climate and environment-related challenges, contributing to healthier, more prosperous and more sustainable societies."
>
> A European Strategy for Data, **European Commission**

[5] A strategy for Data, European Commission, 2020, 2-3. Source: **digital-strategy.ec.europa.eu/en/policies/strategy-data**

The role of physiology and physiologists in unlocking the potential of artificial intelligence for health

Physiology is the branch of biology that aims to understand the mechanisms of livings things from the basis of cell function at the ionic and molecular level to the integrated behaviour of the whole body and the influence of the external environment.[6] Physiological research helps us to understand and determine what goes wrong in disease and facilitates the development of new treatments.

Physiology is a discovery science with physiologists working in basic, translational and clinical research. Physiology data has been harnessed by artificial intelligence to advance health care, such as assisting in the early and accurate diagnosis of diseases and providing personalised medicine.[7]

Physiologists play a crucial role in building better AI tools in healthcare by contributing their expertise and knowledge of the human body's complex systems and functions in two critical areas:

— By providing insight into the underlying biological processes and mechanisms that drive various health conditions. This domain of knowledge can help inform the development of AI algorithms and ensure that they accurately model the plausible physiological processes and reduce the risk of identifying confounding factors.

— By helping assess, interpret and contextualise the data used to train AI models, ensuring that they contain plausible measurements and are representative against known standards for the target end users.[8]

AI has the potential to transform healthcare, enhance health outcomes and improve population health, but it must be implemented thoughtfully to manage the risks.

With a deep understanding of using big datasets, physiologists can provide oversight to ensure that the quality of data being used in AI systems is accurate and does not inadvertently exclude data cohorts. Integrating physiological knowledge into relevant AI models and systems can enhance the understanding and interpretation of complex health data, ultimately leading to better informed decision-making.[9]

With a deep understanding of using big datasets, physiologists can provide oversight to ensure that the quality of data being used in AI systems is accurate and does not inadvertently exclude data cohorts.

[6] Source: The Physiological Society, **physoc.org/explore-physiology/what-is-physiology/**

[7] Zhang, A. et al. Leveraging physiology and artificial intelligence to deliver advancements in health care, July 2023. Source: **journals.physiology.org/doi/abs/10.1152/physrev.00033.2022?journalCode=physrev**

[8] Source: The Physiological Society, **physoc.org/policy/artificial-intelligence/artificial-intelligence-and-health/**

[9] Source: The Physiological Society, **physoc.org/magazine-articles/developing-trusted-healthcare-tools-how-physiology-can-unlock-the-potential-of-artificial-intelligence-for-health/**

Transforming AI in low- and middle-income countries

Bilal Mateen MPH, MBBS
Executive Director
Digital Square

Bilal Mateen outlined a broad history of AI evolution from the 1990s, when the first AI tool was authorised by the Food and Drug Administration (FDA), to the early 2000s, when computer-aided detection (CAD) image checkers became prolific in breast cancer screening, to today's potential. The principal question discussed was – how transforming could AI be for healthcare?

Magpie approach
Society has swung from one shiny technology to another, without investing adequately in generating robust evidence of effectiveness. Bilal outlined examples in breast cancer detection where AI is just one cog in the machine, and there is limited (if any) evidence that increased detection rates improve patient health outcomes. More evaluation through randomised clinical trials, and novel study designs are needed. If the answer is that AI tools are non-inferior (in terms of patient outcomes) to expert physicians, then what is the value proposition? In short, this becomes an issue of cost-effectiveness – do the efficiency gains from an AI tool justify the investment needed to evolve the clinical workflow and the day-to-day operating costs?

Who pays?
Clinical effectiveness may not be enough, AI based tools need to be cost-effective too. Health systems do not have strong mechanisms to reimburse the private sector for the development of these tools – in fact, this is a source of substantial innovation in terms of policy such as Digital Health Applications (DiGA) from Germany, and coding infrastructure such as The American Medical Association's AI CPT[10] codes. This then leads to prioritised development and commercialisation in countries such as the US where privatisation and the willingness to pay is different to many other places in the world. The result is that although developing AI medical devices can be done at pace (as demonstrated by the over 800 FDA-cleared devices), it also means that they are based on the demographics of the US-insured consumer.

[10] Current Procedural Terminology (CPT) codes offer doctors and healthcare professionals a uniform language for coding medical services and procedures to streamline reporting, increase accuracy and efficiency. Source: **ama-assn.org/practice-management/cpt/cpt-overview-and-code-approval/what-is-a-cpt-code**

In the UK we have (in principle) a single payer in our health system via the NHS which should provide a strong bargaining position with the private sector. However, the procurement process for digital technology has been devolved to NHS trusts – via the development of tools like the digital technology assessment criteria and NHSX[11] buyer's guide. The question is if this process should follow the same approach as that of medicines, wherein procurement is done centrally to secure volume-based discounts (as one example of the benefits).

Study biases when the end user is a US insurance company
As Americans are the biggest market and have one of the most responsive regulatory ecosystems, they can drive innovations in their own market. However, there is a huge problem with a lack of diversity, with most tools (more than 70%) having data from the same three states in the US.[12]

The UK has a high quality of routinely collected data, however, there are ethnicity flaws due to the way people report their ethnicity. For example, in the case of the 'Arab' ethnic group – ⅓ have their ethnicity reported incorrectly in their health record. In essence, biases and incomplete data are not new. Most drug development historically used data from white middle-aged men and this issue needs to be addressed as AI use increases.

Notably though, more needs to be done to illustrate the return-on-investment from 'cleaning up' and addressing data bias, as often this is currently reliant on 'philanthropic capital' that is not bound by any incentive mechanism and is not a sustainable longer term solution.

The UK has a high quality of routinely collected data however there are ethnicity flaws due to the way people report their ethnicity.

[11] Now integrated into the Transformation Directorate. See: **transform.england.nhs.uk/**

[12] Kaushal, A. et al. Geographic Distribution of US Cohorts Used to Train Deep Learning Algorithms, September 2020. Source: **jamanetwork.com/journals/jama/fullarticle/2770833**

The investor perspective

Kai Johns
Senior Ethical, Sustainable and Impact Researcher
Greenbank

Broadly, the implications of AI and thus the investment opportunity set can be broken into three categories.

— **Picks and shovels**
Who are the companies providing the essential products and services needed for AI? The semiconductors value chain; physical infrastructure; data centres; cooling equipment; and data specialists as a lot of companies have data, but it is not structured or digitalised.

Companies within the 'picks and shovels' category require careful oversight of additional environmental and social characteristics: the human rights risk associated with the precious metals needed, water use within cooling and the high levels of renewable energy required to power a data centre cleanly.

— **Entrench or disrupt**
The more customers a company has, the more data points it has access to. This enables it to create better tools, making it harder for competitors with fewer datapoints to catch up. Legacy, complicated, and underperforming technology leaves itself vulnerable to a competitor harnessing AI.

Within the 'entrench or disrupt' category, opportunities include an assessment of the management of risks such as governance controls, privacy and permission.

— **Innovation pipeline**
If we look at the rate of innovation in healthcare historically, it has become increasingly expensive. This is known as Eroom's Law, the concept that the cost of developing healthcare solutions has increased exponentially in the last several decades despite improvements in technology.[13] AI has the potential to alter this trend, either by reducing the cost of future innovation or speeding it up – ideally both.

The central questions framing this category are: What new products and services will AI enable, and what sustainability challenges could it solve?

[13] Zhavoronkov, A. When Will AI Beat The Eroom's Law In The Pharmaceutical Industry? Forbes, August 2022. Source: **forbes.com/sites/alexzhavoronkov/2022/08/22/when-will-ai-beat-the-erooms-law-in-the-pharmaceutical-industry/**

Additional reading

The Physiological Society, **From 'Black Box' to Trusted Healthcare Tools: Physiology's role in unlocking the potential of AI for health**, June 2023. — static.physoc.org/app/uploads/2023/06/23142343/WEB-From-Black-Box-to-Trusted-Healthcare-Tools.pdf

BMJ, **Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness**, March 2020. — bmj.com/content/368/bmj.l6927.full

Han, R. et al. **Randomised controlled trials evaluating artificial intelligence in clinical practice: a scoping review,** May 2024. — thelancet.com/journals/landig/article/PIIS2589-7500(24)00047-5/fulltext

IBM Technology, **How large Language Models Work**, July 2023. — youtube.com/watch/how-large-language-models-work

IBM Technology, **Risks of Large Language Models (LLM)**, April 2023. — youtube.com/watch/risks-of-large-language-models-llm

Omiye, J. et al. **Large language models in medicine: the potentials and pitfalls**, August 2023. — arxiv.org/pdf/2309.00087

Chiat Ling Ong, J. et al. **Ethical and regulatory challenges of large language models in medicine**, April 2024. — thelancet.com/journals/landig/article/PIIS2589-7500(24)00061-X/fulltext

Perrigo, B. **Exclusive: OpenAI Used Kenyan Workers on Less Than $2 Per Hour to Make ChatGPT Less Toxic**, Time, January 2023. — time.com/6247678/openai-chatgpt-kenya-workers/

Perrigo, B. **The Indian Startup Making AI Fairer—While Helping the Poor**, Time, July 2023. — time.com/6297403/the-workers-behind-ai-rarely-see-its-rewards-this-indian-startup-wants-to-fix-that/

Greenbank

For further information on the services Greenbank provide, or to arrange a meeting, please contact us.

Call
0117 930 3000

Email
enquiries@greenbankinvestments.com

For more information, please visit
greenbankinvestments.com

@GreenbankInvest

Greenbank Investments

Our UK offices
London
Bristol
Edinburgh
Glasgow
Liverpool

Information valid at date of publication, August 2024.

Rathbones, Greenbank and Greenbank Investments are trading names of Rathbones Investment Management Limited, which is authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority. Registered office: Port of Liverpool Building, Pier Head, Liverpool, L3 1NW. Registered in England No. 01448919.