
DeepMind Health Independent Review Panel Annual Report

June 2018

About DeepMind Health

DeepMind was founded by Demis Hassabis, Shane Legg and Mustafa Suleyman in 2010. It was acquired by Google in 2014 and is a separate legal entity within the Alphabet Group.

DeepMind began working with the Royal Free London NHS Foundation Trust in July 2015 after being approached by clinicians from the Trust and DeepMind Health was founded in February 2016.

About the Independent Review Panel

The Independent Review Panel was announced at the same time as DeepMind Health.

They are currently:

Martin Bromiley OBE

Elisabeth Buggins CBE

Eileen Burbidge MBE

Richard Horton

Dr Julian Huppert

Professor Donal O'Donoghue

Matthew Taylor

Professor Sir John Tooke

Contents

Appendix links	02
Foreword	03
What healthcare needs are being addressed?	04
Report of the independent reviewers	06
Principles	07
DeepMind Health self-assessment on principles	08
Business model	12
Clinical outcomes	14
Clinical utility	15
Broader consequences and human factors/ergonomics	16
Public and patient engagement	19
Law, regulation and data governance	21
Security	23
Independent review panel governance	24

Appendix links

[Biographies Independent Review Panel](#)

[Independent Review Panel pledge with DeepMind Health](#)

[Royal Free Agreement Streams](#)

[Imperial College Agreement Streams](#)

[Taunton and Somerset Agreement Stream](#)

[Yeovil Agreement](#)

[Mammography timeline](#)

[NHS Innovation Accelerator Evaluation](#)

[Streams Preliminary data results](#)

[Human Factors review of Streams](#)

[Human Factors and Ergonomics review by Hu-tech](#)

[Streams screenshot](#)

[Patient and public engagement review](#)

Foreword

By Dr Julian Huppert, Chair, Independent Review Panel

When DeepMind Health was established in 2016, it created an Independent Review Panel. This is a very unusual and brave arrangement for a commercial company, and we are only aware of one other similar example, the Yoti Guardian Council. Nine Independent Reviewers, drawn from different walks of life but linked by their reputation, experience and integrity have been given broad access to what DeepMind Health is doing. Panel members are not required to sign non-disclosure agreements and are able to commission external analysis and opinion, with DeepMind Health meeting the costs. DeepMind Health make only one stipulation of the Panel; that it should publish an annual report. Our first report was issued in July 2017, and I am delighted to present this second report.

When we, the Independent Reviewers began our scrutiny of DeepMind Health, our brief was completely open. There were potentially many areas that we could have explored but for our first annual report, we chose to focus on a small number in depth. For instance, one was law, regulation and data governance. For this, we commissioned a legal opinion from a noted expert in law relating to technology and data use. As with all our work, the choice of expert was entirely our own.

In this second report, we have sought to do three things. Firstly, we wished to return to the areas of concern we identified in 2017, looking at DeepMind Health's progress over the last year. Secondly, we selected three topics which we did not address in depth last year – the business models for DeepMind Health and their relationship with Alphabet, some additional work on human factors and evidence of clinical utility.

We also wanted to agree a firm set of expectations for how a company like DeepMind Health – or any organisation in this area – should operate. This would enable us to move away from purely reactive work, although we would continue to do that where needed, and instead base our future analysis of DeepMind Health on a set of general principles. As DeepMind Health's business expands, in scale, scope, and location, these principles should act as a framework for them against which we can review their activities. In this report we present the 12 principles we developed, and DeepMind Health have agreed that we will use them as a framework for us to judge their activities. We intend to assess them against each one of these every year, as well as selecting some particular areas of focus.

In this report we highlight a number of areas where DeepMind Health need to build upon the progress made in responding to our first report. We have been clear from the outset that 'good enough' is not good enough for a company with such a close relationship to Google, a company which already reaches deep into all our lives. The issues of privacy in a digital age are if anything, of greater concern now, than they were a year ago and the public's view of the tech giants has shifted substantially. In particular, the acquisition of vast amounts of data from Facebook and its use by a third party, Cambridge Analytica to covertly subvert democratic elections has shocked many. All companies that wish to operate in the area of healthcare data ought to be held to high standards, but this onus is even greater for a company such as DeepMind Health.

As a result of our self-imposed term limit, this is my last report as Chair of the Panel, and I would like to take the time to thank all those who have helped so much in establishing this Panel and in our work over the last two years. In particular I would like to thank all of the eight other Reviewers I have had the pleasure of working with and learning from, to Vivienne Parry for turning our thoughts into words, and especially to Rebecca O'Leary, for making this whole project function.

Dr Julian Huppert

Chair

Please click [here](#) to view the biographies of the Independent Review Panel and [here](#) to view our agreement with DeepMind Health.

What healthcare needs are being addressed?

There are two principle threads of work for DeepMind Health. The first is Streams, which we described last year and is a direct clinical project, presenting clinicians with medical data. DeepMind Health is also involved in three medical research project collaborations within the NHS.

Acute kidney injury (AKI) at the Royal Free London NHS Foundation Trust

AKI is sudden damage to the kidneys that prevents them working properly. The damage ranges from minor loss of kidney function to complete failure requiring temporary or even lifelong dialysis. It can also cause death. It is common in hospitals, found in 13–18% of those admitted and up to 30% of patients in critical care. It is associated with a wide range of serious health problems including major surgery, trauma, burns and sepsis as well as with existing kidney or urinary tract conditions.

The major challenge of AKI is that it is symptomless in its early stage with few or no warning signs. Its detection depends largely on laboratory tests, particularly of a rise in blood levels of a chemical waste product of muscles called creatinine, which should normally quickly be excreted by the kidneys. Treatment requires rapid assessment of the underlying cause and appropriate treatment. Speed is of the essence.

The issue for clinicians is that they constantly have to check whether lab results are back before deciding what they might mean for the patient. With many hundreds of test results of all types coming back each day plus a full range of other clinical tasks and emergencies, delays in assessing patients for AKI occur.

DeepMind Health developed Streams, which is a secure mobile app designed to help clinicians better identify and manage patients at risk of deterioration. It was initially deployed at the Royal Free London NHS Foundation Trust in 2016. In its current state, the app helps a team of nurses and doctors at the Trust to respond to AKI.

Streams is a way of displaying alerts developed by the standardised NHS England AKI algorithm on portable hand-held devices (such as tablets or phones). It has result viewing functionality, and has tools for collaborating between, and within teams. Streams does not currently use artificial intelligence (AI) although it is envisaged that in the future, AI driven alerts could be delivered by it. The current generic functionality mentioned above could also be applied for many other areas of clinical practice including sepsis.

Streams will be deployed in a number of other NHS hospital Trusts over the course of 2018 including Imperial College Health NHS Trust, Yeovil District Hospital NHS Foundation Trust and Taunton and Somerset NHS Foundation Trust Hospitals. It is envisaged that Streams will provide a broad range of functionality at these hospital sites including test results viewing and vital signs monitoring.

Faster assessment of retinal imaging **Automated analysis of retinal imaging using machine learning**

A technique called Optical Coherence Tomography (OCT) is used to take cross sectional pictures of the retina, the light sensitive tissue at the back of the eye. The digital scans produced are used in the management of a range of common eye conditions including age related macular degeneration, glaucoma and diabetic retinopathy.

The challenge is that these scans are highly complex and require time consuming interpretation by an expert, despite a range of traditional analysis tools. The time taken for analysis has an impact on the number of patients that can be assessed and the time that experts can spend with patients.

The DeepMind Health approach involves the assessment of a number of automated analysis methods. As part of the project, historic de-personalised data has been ‘cleaned’ – for instance, removing any coding errors – to produce an enhanced, research-ready dataset.

This project has been undertaken in conjunction with Moorfields Eye Hospital.

Please click [here](#) to view the Royal Free London agreement, please click [here](#) to view the Imperial College agreement, please click [here](#) to view the Taunton and Somerset agreement, please click [here](#) to view the Yeovil agreement.

What healthcare needs are being addressed?

Better planning for head and neck tumours

Radiotherapy is a key treatment for head and neck tumours but their location, so close to the brain, eyes and other precious structures, means that great care must be taken to restrict radiation to the tumour alone. This means knowing where the 'edges' of the tumour are, so that no healthy tissue is irradiated and knowing its exact size, so that precisely the right amount of radiation is used to destroy it. Each tumour is unique in shape, so developing a 3D model of the tumour, derived from scans, a process known as segmentation, is the earliest and most critical step in the planning process. If this model is not correct it can lead to major problems for the patient.

There are several challenges. Segmentation is complex and very time consuming. It takes about 4 hours for each patient. This can delay the start of treatment and also has an impact on the number of patients that can be treated. Because of its complexity, experts may arrive at different interpretations of the tumour volume, even when given the same scanning information. The minimisation of this variability is an important aim as it should improve patient treatment and minimise side effects as well as release clinician time.

The DeepMind Health approach involves using applying machine learning techniques to perform automated segmentation of head and neck tumours volumes and organs on radiotherapy planning, CT scans.

This project has been undertaken with University College London Hospitals NHS Trust.

Improved breast cancer screening

DeepMind Health began a partnership with the Cancer Research UK Imperial Centre in November 2017, collaborating with a consortium of leading clinicians and academics led by the Centre, as well as with the AI health research team at Google, to explore the potential benefits that AI technology could have in identifying signs of breast cancer in mammograms (X-rays of the breasts).

Breast cancer is a significant global health problem, with 1.6 million people across the globe diagnosed with the disease every year. Early detection and treatment of breast cancer leads to a much higher chance of a full recovery, saving thousands of lives every year. However, accurately detecting and diagnosing breast cancer still remains challenging.

The latest machine learning technology is being used to carefully analyse historic mammograms; initially from around 7,500 women provided by the Cancer Research

UK-funded OPTIMAM mammography database, at the Royal Surrey County Hospital NHS Foundation Trust. OPTIMAM's historic digital images are depersonalised and have been stripped of any information which could be used to identify patients. These images have been available for use by research groups around the world for a number of years. DeepMind Health have also been supporting the expansion of the OPTIMAM programme to include a larger, more clinically representative database.

This project will investigate whether machine learning technology could more accurately identify signs of breast cancer in mammograms, potentially leading to earlier diagnosis and treatment for patients. Machine learning algorithms will be developed to perform detection (finding a malignant lesion) and classification (analysing the lesion type) of malignant lesions on mammography images. If it is successful, this research will improve the quality of reporting of screening mammograms, leading to fewer missed cancers, fewer false alarms and, ultimately save lives.

Report of the independent reviewers

Principles

Our initial work largely took a reactive approach, seeking to investigate whether DeepMind Health had in our opinion behaved correctly in various ways, and making recommendations about how they could improve. This was and will continue to be an important thread of work, and in particular where there are serious questions being raised in the press and public discourse about specific things DeepMind Health have done, as happened last year, it is important that we are able to look into them in detail.

However, we would like to take a more proactive role in establishing a framework for ethical behaviour, that goes well beyond simple regulatory and legal compliance. After extensive discussion and consideration amongst the Independent Reviewers, we have therefore produced a set

of 12 principles, that we believe would be appropriate for DeepMind Health, or indeed any company, large or small, working in the healthcare technology space. Our focus, is of course the activities of DeepMind Health, but we make no apology for wishing other companies to lift their standards.

In parallel, DeepMind Health has itself begun to develop a set of values with patients and the public. These values are currently being finalised. We have seen a draft version and they are entirely consistent and compatible with those we set down here. However, as might be expected, they focus primarily on the company's external impact, whilst the principles we set down here also include a range of more fundamental issues for the company's operation.

Principles for DeepMind Health and other companies in the healthcare technology space

1

Benefit to data providers

The company seeks to ensure that patients, service users, healthcare systems and organisations, who are the source of the data, benefit appropriately from the learning derived from it.

2

Public, patient and practitioner engagement

The company proactively engages with patients, carers, practitioners and members of the public, and is responsive to their inputs.

3

Design for safety and utility

The company always designs products and processes that make it easier for staff to do the right thing and minimise unintended consequences.

4

Evidence-driven

The company is committed to generating and sharing evidence of effectiveness for any interventions, including peer review as appropriate, and shall avoid over-claiming the effectiveness of any products and services.

5

Anti-monopoly

The company seeks to ensure that it promotes competition, and encourages other organisations, including SMEs, into the market; in particular, the company will ensure that their systems are interoperable, using open APIs.

6

A model employer

The company ensures that it is exemplary in employment practices, including promoting diversity in all dimensions, equal pay, flexible working, and paying the living wage.

7

Legal and ethical

The company obeys the letter and the spirit of all appropriate legislation and regulation, including taxation.

8

Protecting privacy

The company takes strong steps to protect patient's privacy by design and in implementation.

9

Secure

The company continuously ensures the highest level of security of all data it holds.

10

Transparency

The company promotes transparency in its own work and contracts, within the constraints of privacy.

11

Reasonable profit

The company will not use its assets or position to seek to extract excessive profits in its dealings with the public sector and will, as far as possible, operate contracts on an open book basis.

12

Openness

The company promotes a culture and maintains processes to encourage any member of staff to feel they can raise – without any fear of adverse personal consequences – concerns they have about risks or unethical behaviour.

DeepMind Health self-assessment on principles

We developed these 12 principles separately from DeepMind Health, and they have agreed that we will use them as a framework for us to judge their activities. They are specifically intended to go

beyond legally binding requirements on DeepMind Health. Every year we intend to ask DeepMind Health to self-assess against each of them, and their responses will help to feed back into our ongoing analysis and investigation.

Principle

DeepMind Health responses

1

Benefit to data providers

The company seeks to ensure that patients, service users, health care systems and organisations, who are the source of the data, benefit appropriately from the learning derived from it.

We recognise that the question around the value of data is an important one facing healthcare systems across the globe, especially as new technologies and suppliers enter the field. At DeepMind, we are committed to ensuring that value is delivered back to all the partners and patients involved in the process and are actively engaging with stakeholders to have these conversations.

We believe that AI will have a positive impact on healthcare over the coming years, and that the value of such technologies does not solely arise from the data itself. Value also stems from the complex combination of machine learning, clinical, and safe software expertise that are essential to building the technology – as well as the time and effort that goes into ensuring data is in a useable format. There are many ways to recognise and return benefit, including offering free or discounted use of the eventual product, providing direct payments or a share of revenue and returning an improved data set as a public asset for future research. This is a conversation we expect to continue for the months and years ahead, and we are committed to working with stakeholders to ensure that everyone benefits from our work.

2

Public, patient and practitioner engagement

The company proactively engages with patients, carers, practitioners and members of the public, and is responsive to their inputs.

Since August 2017, DeepMind Health have been working with Ipsos MORI, Crowe Associates, and The Social Kinetic, to proactively talk to a wide range of stakeholders – members of the public, patients and carers, and those working in NHS healthcare and technology – and hear their thoughts and concerns about how technology companies like DeepMind should work with healthcare systems. Through a series of 1:1 interviews, focus groups, workshops, and a day-long Summit in January, members of the DeepMind Health team worked with stakeholders to develop a set of public values which DeepMind Health could adopt into its everyday practices. These will be published shortly.

The programme builds on our existing work with patients and carers, started with Rosamund Snow in September 2016. We have hosted four events at the DeepMind offices and in London venues, working with patients to hear their ideas for advanced technology in the NHS. We now have an online User Group of over 50 patients and carers, as well as patient advisors who work with us on our research collaborations. In March 2018, we led our first design session with patients, mapping out their user journeys through health systems in order to inform our work. This session mirrored the existing design programme we have with clinicians and nurses to help ensure Streams meets the needs of those using it.

These engagement programmes are part of our long-standing commitment to engaging meaningfully with all of those affected by our work. We are also mindful of ensuring that we are accessible to those individuals, and that getting in touch with us is an easy and positive process. As such, we have two primary email aliases – sayhi@deepmindhealth.com and patients@deepmindhealth.com – that are regularly checked and responded to. We receive dozens of emails every week from patients and stakeholders, and we respond to every one.

Outside of these workstreams, we also regularly participate in medical, AI and technology conferences. We believe that this is an important part of contributing to the community, learning from others and sharing our work. In the last two years, we have spoken at over 50 conferences in the UK and abroad.

3

Design for safety and utility

The company always designs products and processes that make it easier for staff to do the right thing and minimise unintended consequences.

We have a very strong user-centred design focus at DeepMind Health with roughly 10% of DeepMind Health employees working on user-centred design. The team adopts the mantra ‘fall in love with the problem, not the solution’. This outlook informs our wider practice when designing health technology, working closely with staff on the ground to ensure that we truly understand the day-to-day problems they face. Through thousands of hours spent shadowing nurses and doctors on the wards at our partner Trusts, interviewing them about their daily routines, and hosting collaborative design workshops, the DeepMind Health team aims to ensure that clinicians have the tools they need to appropriately support them in saving lives and improving care.

As part of our wider programme for design, we host regular user testing sessions with clinicians. We have a group of more than 100 doctors and nurses who regularly come in to the office to test the latest iterations of the Streams clinician app. These sessions ensure that we are made aware of any issues and can respond prior to implementation. Seemingly minor changes, such as using the colour red in the interface, improved the way clinicians navigated the app and responded to patient alerts. We also make regular team visits to our Trusts to collaborate on our roadmap and define the features that will be most effective for clinicians and patients.

4

Evidence-driven

The company is committed to generating and sharing evidence of effectiveness for any interventions, including peer review as appropriate, and shall avoid over-claiming the effectiveness of any products and services.

DeepMind Health adopts an evidence-led approach for all of its projects, adhering to strict academic standards with regards to publication of research, and commissioning service evaluations of our direct-care partnerships.

With regards to Streams, whilst we have heard anecdotes from nurses at the Royal Free about the positive impact of the app, we are awaiting a peer-reviewed service evaluation to measure the overall effectiveness of Streams.

With our research projects, we are committed to following standard academic procedures through publication on peer-reviewed platforms. Where our partners agree, we also publish research protocols prior to any data analysis beginning.

We also have a group of clinical advisors who bring a range of expertise and varying perspectives in their respective fields. They meet on a weekly basis, and work alongside our team to ensure both the clinical efficacy of our projects and the best use of our technologies.

5

Anti-monopoly

The company seeks to ensure that it promotes competition, and encourages other organisations, including SMEs, into the market; in particular, the company will ensure that their systems are interoperable, using open APIs.

DeepMind Health believes a competitive market place is essential for a thriving health-tech ecosystem. We aim to be a valued contributor to the wider health-tech ecosystem working collaboratively in four main ways:

- i) commitment to interoperability standards (see below for more information)
- ii) working with partners to improve datasets which can then be shared with the wider research community
- iii) not requiring any of our partners to work exclusively with us
- iv) supporting the startup community

From its conception, Streams has been designed as an interoperable system, responding to the very real need for healthcare systems to talk to one another. Rather than simply creating a system that only our apps would understand, we have implemented state-of-the-art interoperability standards (such as FHIR) that will allow existing and future software and apps to work alongside our technologies.

DeepMind is also a member of INTEROPen, which includes members from NHS organisations, healthcare companies and software developers who work together to accelerate, and advocate for, the development of open standards for interoperability in the health and social care sector. In March 2017, we co-hosted a two-day event with INTEROPen for over 300 people at our London offices. The Interop Summit brought together local and national experts to discuss how we could further the interoperability agenda. We also took part in discussions run by NHS England on interoperability.

Our team regularly contribute to startup events at universities and in the commercial sector and a number of our clinical, design and engineering team provide mentorship to startup founders.

6

A model employer

The company ensures that it is exemplary in employment practices, including promoting diversity in all dimensions, equal pay, flexible working, and paying the living wage.

We want our team and their families to live happy and healthy lives, both in and out of work. Our employment practices and benefits are thoughtfully designed to create an environment that enables people to fulfil their potential. For example, we offer competitive pay, generous parental leave policies, retirement savings plans, relocation support, and access to excellent healthcare and wellbeing choices.

Our diversity and inclusion team work with our recruitment team to ensure that DeepMind is recruiting from the widest possible talent pool. For example, in 2016 and 2017 we have had a stand at the Grace Hopper Celebration, the world's largest gathering of women technologists. We also hosted over 300 children from disadvantaged and under-represented groups in the office last year, as part of a series of engagement activities and talks. Throughout the year, we run awareness campaigns and celebrations in the office for important moments like LGBT Pride, Mental Health Awareness Week and Black History Month. We seek to support under-represented groups into technology through multiple charity partners and outreach programmes, ensure our systems are equitable and checked for bias, and create an inclusive environment that celebrates diversity.

7

Legal and ethical

The company obeys the letter and the spirit of all appropriate legislation and regulation, including taxation.

We acknowledge the paramount importance of tight regulation of companies operating in healthcare, especially those that process sensitive health data. We have legal, clinical safety and information governance teams that work closely together to ensure all projects undergo appropriate due diligence prior to contract signature and ensure that all applicable legislation and regulation is complied with. In addition, these teams are engaged over the life of a project to manage our ongoing compliance.

As technology develops we appreciate the importance of businesses working with government to ensure legislative frameworks fit technology deployed in the market. To this end we have made ourselves available to regulators to discuss the work that we are doing, and plan to carry out in the future.

8

Protecting privacy

The company takes strong steps to protect patient's privacy by design and in implementation.

Across all areas of our work, privacy and security of data sits at the heart of what we do. As part of our Collaborative Listening and Values programme, this emerged as a priority for stakeholders and a prerequisite for any company working with NHS data.

In our research partnerships all the data that we use is approved through the relevant research and ethics bodies for example the Health Research Authority on ophthalmology and radiotherapy. All research data has been de-personalised by our partners before being transferred.

In our direct care partnerships, where we process patient information, we are legally and contractually bound to only using patient data to improve care, under the instructions of our partners. It would be against the law, and our ethics, to use the data for any other purpose.

Streams has been explicitly designed to uphold patient privacy. The app has passed all NHS audit and review processes, and the data is held within an accredited high security environment that meets all NHS standards with restricted access governed by the Trust. The data itself is encrypted both while it is being stored in the data centre, and when in transit. Only registered clinicians who have been verified by the Trusts can log into Streams, and they have to go through two levels of fingerprint access.

We work closely with our partners to ensure that patients and clinicians are informed about how patient data is being used, from helping design leaflets and information packs, to visiting the Trusts and participating in events.

9

Secure

The company continuously ensures the highest level of security of all data it holds.

Data privacy and security is paramount, and we hold ourselves to the highest security standards. Adhering to all NHS standards, our systems have been built by some of the world's leading security experts. Personal data that we process on behalf of our partners is encrypted at rest and in transit, and stored in a high-security facility. Only those who need to access this personal data are able to, after a robust verification process. We are developing new technology, called Verifiable Data Audit, that will ensure the logs of how data has been processed by our systems cannot be modified later, and eventually make them available for audit by third parties through open standards.

10

Transparency

The company promotes transparency in its own work and contracts, within the constraints of privacy.

Since its creation, DeepMind Health has taken proactive steps to becoming one of the most transparent companies working in healthcare – from appointing its own board of Independent Reviewers with unrestricted access to review its work, to publishing its contracts with NHS partners on its website, to developing the foundations for a Verifiable Data Audit system which will allow partners to track who has had access to health data, when and for what purpose.

In the Collaborative Listening and Values programme to shape the values by which DeepMind Health will operate, transparency came out as the principal value for all stakeholder groups. In those sessions, publication of contracts were less of a priority, but questions around company roadmap and priorities were of high interest. Moving forward, we will continue to place transparency as a top priority across all of our work.

11

Reasonable profit

The company will not use its assets or position to seek to extract excessive profits in its dealings with the public sector and will as far as possible operate contracts on an open book basis.

DeepMind Health works to address the IHI 'Triple Aim' laid out by Don Berwick: enhancing the experience of care; improving the health of populations; and reducing the per capita cost of healthcare. Over the last few months, we have been talking to a range of stakeholders about what values we should adopt when working in health-care. One of the values we have adopted is to 'operate sustainably' which commits us to being mindful of what projects we work on and why, as well as reinvesting profits back into our long-term mission to ensure that we have a positive impact in healthcare.

To operate sustainably, DeepMind Health must achieve a reasonable rate of return in order to make lasting positive impact and provide enduring value to our partners. We intend to show through robust and peer-reviewed evaluations that our technology delivers both outcome and experience improvements, as well as making financial sense. While we aim to work within an outcome-based framework, focused on fair remuneration based on results, health systems continue to rely on traditional fee for service models. So, in the near term, we will charge hospitals a fair fee for providing them with a useful service like other health technology providers do but with a commitment that, where possible, payments will be tied to measurable improvements in clinical outcomes and patient experience.

12

Openness

The company promotes a culture and maintains processes to encourage any member of staff to feel they can raise – without any fear of adverse personal consequences – concerns they have about risks or unethical behaviour.

DeepMind Health aims to foster an open culture in which any member of staff can raise concerns with ease and without fear of adverse personal consequences. We have put several processes in place to ensure that the voices of staff can be heard, and responded to, in an appropriate manner.

There is clear guidance for staff that any concerns can be raised to a manager or our Clinical Safety Officer (if applicable). We also have a weekly meeting specifically looking at risks and issues, where a member of staff can table a concern. Furthermore, we recognise that people may not always feel comfortable raising concerns in these ways, and if a member of staff wants to speak confidentially about an issue, they can speak to our Information Governance Manager in confidence. We are committed to ensuring that our actions are held to account both by external parties and by our own staff.

Business model

Principles covered:

1 Benefit to data providers

5 Anti-monopoly

11 Reasonable profit

The tide of public opinion has turned strongly against the tech giants. They are seen as monopolies that do not play their fair part in society, whether it is paying enough taxes or keeping harmful content off their platforms. Their motives are now regarded with increasing suspicion. The scandal involving Facebook and Cambridge Analytica, served to sharply underline the validity of that old maxim, ‘if a product is free, you’re the product’. Against this background, it is hardly surprising that the public should question the motivations of a company so closely linked to Google as DeepMind Health. Furthermore, DeepMind Health work with medical information, something that is regarded as deeply personal and which therefore attracts greater scrutiny. So the question, ‘Where are they making their money’ is a crucial one. It’s the question that many feel foolish for not previously asking of the tech giants.

It is important for the public to have reassurance about DeepMind Health’s business model. They want to know that DeepMind Health’s revenue is – or will be – coming from a source they consider appropriate. They also want to know that any promises made now by DeepMind Health will be kept in the future and that if this generation enjoys a service for little or no cost, that the next generation will not pay the price in excessive costs or lack of control further down the line.

There would be considerable sensitivity if DeepMind Health’s business model involved selling data, either in a depersonalised form or in a raw form. There would also be very significant concerns if any of the data were used to tune advertising. From what we have seen, neither of these are at all envisaged or considered as desirable revenue sources by DeepMind Health and clear restrictions on how DeepMind can use data are stated in each of the contracts we have reviewed. So, if it is intended that DeepMind Health should make a profit, where will it come from? It might be that DeepMind Health accepts it will never make a profit in the UK but instead wishes to use the experience gained here to develop substantial revenues

in other territories. But the principles we have developed are not UK exclusive, they should apply across the world. Alternatively, it is possible that DeepMind Health is not intended to make money. It could be thought of as a not for profit, whose purpose is building the brand for DeepMind or Alphabet, or driving the use of other Google services, for instance cloud storage, rather than making money. It could even be envisaged as a charitably-minded venture driven by an altruistic desire to improve health.

We have had detailed conversations about DeepMind Health’s evolving thoughts in this area, and are aware that some of these questions have not yet been finalised. However, we would urge DeepMind Health to set out publicly what they are proposing. If not, there is a real risk that the public will assume that the work is driven by a purely profit-making motivation and will have stronger grounds for cynicism about how data and machine learning will be used.

In addition, there is a crucial issue about the relationship between DeepMind and DeepMind Health and between both DeepMind and DeepMind Health with their parent company Alphabet, which is also the holding company of Google. To what extent can DeepMind Health insulate itself against Alphabet instructing them in the future to do something which it has promised not to do today? Or, if DeepMind Health’s current management were to leave DeepMind Health, how much could a new CEO alter what has been agreed today? We appreciate that DeepMind Health would continue to be bound by the legal and regulatory framework, but much of our attention is on the steps that DeepMind Health have taken to take a more ethical stance than the law requires; could this all be ended? We encourage DeepMind Health to look at ways of entrenching its separation from Alphabet and DeepMind more robustly, so that it can have enduring force to the commitments it makes.

Business model

The relationship with Google is a constant question that runs through many areas of DeepMind Health's business. The DeepMind Health website at <https://deepmind.com/applied/deepmind-health/deepmind-health-faqs> says 'data will never be connected to Google accounts or services'. Is this a firm assurance that can be made and relied upon? We accept that there is no proposal for data to be used for purposes such as advertising, but this statement might be taken to mean that, for example, DeepMind Health would only use non-Google cloud services. However, the Mammography project (detailed on P 5), clearly states that a Google cloud service will be used which might lead some to think that this promise was already being broken.

In reality, the high security provided by Google's cloud services is the reason why they are used by so many major companies. Nevertheless we urge DeepMind Health to clarify, in simple language, exactly what it means by 'data will never be connected' so that people can be sure that what is being done will not put their data at risk.

A separate issue relates to the extent to which DeepMind Health, even apart from its connections with Alphabet, could find itself in a position of being able to exert excessive monopoly power. There are many examples in the IT arena where companies lock their customers into systems that are difficult to change or replace. Such arrangements are not in the interests of the public. And we do not want to see DeepMind Health putting itself in a position where clients, such as hospitals, find themselves forced to stay with DeepMind Health even if it is no longer financially or clinically sensible to do so; we want DeepMind Health to compete on quality and price, not by entrenching legacy positions. We are therefore encouraged by DeepMind Health's stated commitment to interoperability of systems, and their adoption of the FHIR open API. This means that there is potential for many other SMEs to become involved, creating a diverse and innovative marketplace which works to the benefit of consumers, innovation and the economy. We also note DeepMind Health's intention to implement many of the features of Streams as modules which could be easily swapped, meaning that they will have to rely on being the best to stay in business.

DeepMind Health has given commitments to publishing details of the algorithms it develops from its machine learning studies, such as that developed for their mammography project in peer reviewed scholarly journals for scrutiny from the academic community. We commend them for that and hope that they will continue to do so. We are also

encouraged by the attitude they have taken with those who have provided data for learning, such as with Moorfield's Eye Hospital, who retain control of the cleaned up and enhanced anonymous data that DeepMind Health used, including the ability to make it available to other organisations and to cease making it available to DeepMind Health.

Given the current environment, and with no clarity about DeepMind Health's business model, people are likely to suspect that there must be an undisclosed profit motive or a hidden agenda. We do not believe this to be the case, but would urge DeepMind Health to be transparent about their business model, and their ability to stick to that without being overridden by Alphabet. For once an idea of hidden agendas is fixed in people's mind, it is hard to shift, no matter how much a company is motivated by the public good.

We recommend

That DeepMind Health should be transparent about its business model.

That DeepMind Health specify exactly how they will work with other elements of Alphabet, and what data could ever be transferred to them.

That DeepMind Health should investigate ways to ensure that public commitments can be relied on, even where they are not legally binding, in the event of a leadership change or decisions taken by their owners.

Clinical outcomes

Principles covered:

1 Benefit to data providers

4 Evidence-driven

In 2017 we raised concerns about:

The complexity and scale of the problems not being sufficiently understood and about the method of development which may have potentially adverse implications for successful adoption and diffusion across the wider health system.

Clinicians frequently mentioned their view that DeepMind Health did not understand the complexity of the problems they were addressing. However, we also recognised the converse; that untainted by experience, DeepMind Health have been able to tackle issues that would have daunted others or, been placed immediately in the 'too difficult' box. We are clear that as DeepMind Health have advanced their work, they have become all too familiar with the scale of the challenges they face.

The main issue here however, is that a lack of understanding can also lead to methods of development which could have potentially adverse implications for successful adoption and diffusion across the wider health system.

A recent evaluation of the NHS Innovation Accelerator highlighted the major barriers to adoption and scaling across a range of innovations. We feel that those which are of particular relevance to DeepMind Health are the level of service disruption, the short and long-term payback (both clinical and financial) of the investment in time and effort by Trusts and the number of people needed to support the innovation in a particular setting. It is vital to have clinical leadership in order to overcome NHS structure and process barriers, and that staff involved have a range of supportive personal characteristics.

We therefore recommended early engagement with the appropriate Royal Colleges and other clinical professional bodies for early identification of potential problems, as well as experts in implementation science and in quality improvement to maximise the potential for adoption and diffusion.

DeepMind Health responded

Over the last year, we have engaged with organisations including NHS Improvement, Royal Colleges of Physicians, Surgeons and Radiologists to increase transparency of our work and seek feedback. We have also met with the Royal College of Pediatrics and Child Health, the British Heart Foundation, Cancer Research UK, the Wellcome Trust, the NIHR, the MRC (OSCHR) and the Health Foundation. Professor Rosalind Raine, is leading our service evaluation at the Royal Free Hospital and is an expert in quality improvement and implementation science. A trial of our work at Moorfields will also be led by independent experts and we will seek oversight from relevant independent bodies (and patients) in the trial steering and monitoring committee.

This has resulted in the establishment of advisory and oversight groups including service users, system leaders, clinical and informatics professionals. This engagement has been supported by plain English summaries of the problems being addressed, research questions and project design including clear descriptions of the meaning of AI, machine and deep learning (these are produced on the back page). In addition, where data will be held, the purposes for which it will be used and access rights to the data have also been clearly identified.

Please go to the following link for the [NHS Innovation Accelerator evaluation report](#).

Clinical utility

Principles covered:

3 Design for safety and utility

4 Evidence-driven

Part of clinical utility is addressing problems that are important and relevant to patients and health systems, rather than cherry picking those that may be less relevant clinically but easier to address or commercially more attractive.

We are satisfied that DeepMind Health are addressing important and relevant problems in all the areas they deal with. A good example is their mammography project.

Two radiologists are used in the interpretation of each mammogram, in the UK breast screening programme, a practice known as double reading which increases accuracy. However, there is a shortage of radiologists which is particularly acute in breast imaging, with the Royal College of Radiologists repeatedly highlighting the crisis in workforce and its threat to the continuation of the national breast screening programme. This was before the recent identification of potentially 450,000 older women who will now require screening because they were not invited to their final routine screen, as they should have been. The scale of challenge here is formidable and DeepMind Health's work represents a potential solution to this extremely pressing problem.

Clinical evaluation

Formal evaluation of the three projects we considered last year that would inform a judgement on clinical utility are still awaited. DeepMind Health have provided some preliminary data, which looks promising, for Streams. However, rigorous robust evaluation of all the evidence is required before any formal comment can be made on Streams. We anticipate that this will be published within the next year.

Retinal imaging

There are several major challenges presented by the use of AI analysis of optical computed tomography (OCT) images. Firstly, it must be as good as, or better than experts. It must be capable of operating in real world clinical settings where patients have a wide range of eye conditions necessitating retinal scans and of fitting into standard clinical pathways. Finally, it should ideally be 'generalisable', that is, if other or newer imaging equipment is used, the AI analysis should remain accurate without having to develop a new dataset and training for every type of machine used.

The DeepMind approach was tested for patient triage in a typical ophthalmology clinical referral pathway, comprising more than 50 common diagnoses for which OCT provides the definitive imaging modality. The DeepMind Health team are excited by promising results, and are awaiting publication of the findings in a peer-reviewed journal

As part of this project, DeepMind Health collaborated with Moorfields Eye Hospital NHS Foundation Trust to 'clean' their dataset, for instance removing coding errors, thus making this enhanced dataset 'research ready'. Its value has quickly been appreciated and utilised by vision researchers. At a recent international ophthalmology conference, 8 posters were presented, demonstrating results from projects investigating important clinical problems based on use of this enhanced dataset.

Please go to the following link for the [Streams Preliminary Data Results](#).

Broader consequences and human factors/ergonomics

Principles covered:

- 3 Design for safety and utility
- 4 Evidence-driven
- 8 Protecting privacy
- 9 Secure

In 2017 we raised concerns about:

There may be problems with rolling out Streams to other hospitals where it may be seen as being parachuted in.

We noted that in deploying Streams at the Royal Free London NHS Foundation Trust, which is a large tertiary referral hospital, a dedicated renal response team had been factored into its development. This team are mobilised in response to alerts. The typical average sized district general hospitals are very unlikely to have a dedicated team and will rely on clinicians with many other roles and responsibilities.

DeepMind Health responded

DeepMind Health appreciate that rapid response renal teams may not be feasible across the wider health system.

The 'care pathway' component of the Streams intervention was determined by the Royal Free London NHS Foundation Trust and we accept that different clinical services will configure AKI response differently. Imperial College London's evaluation of Streams implementations at other hospital sites will capture unintended adverse implications regarding workflow, stress, communication and teamwork, independently overseen at the College. The RFH evaluation now considers AKI duration/severity/time to recovery, renal replacement therapy, escalation to ITU, readmission rate and survival. It reports to a monitoring committee including RFH/patient representatives, and extensive HSR experience.

Yeovil District Hospital NHS Foundation Trust is one of the roll out sites but again, this is not a typical Trust, but rather one that is already at the forefront of working with new technologies. It is important for DeepMind Health to engage with hospitals that are less research and innovation focused in order to develop a deeper appreciation of the adoption challenges that need to be addressed in the real world, which are discussed in more detail in the Clinical Outcomes section on P 14. It is also important that DeepMind Health work with as wide a range of users as possible, especially the "least capable user working in the most adverse conditions", such as a locum clinician on their first night's shift in A&E. DeepMind Health's implementation to date seems robust and user-centred and they are to be commended for the time they have put into building relationships and understanding ways of working in the Trust. There has been a significant degree of co-design in each new site, with a wide range of staff. Whilst it has the same functionality, Streams allows different hospitals to enable different features in different ways – Yeovil for instance does not use the 'alerts' system. This provides flexibility with fidelity. Tailoring the product in this way requires considerable input from DeepMind Health but lessens the likelihood of perceptions of it being 'parachuted in'.

Broader consequences and human factors/ergonomics

In 2017 we raised concerns about:

The broader implications of the use of Streams in relation to performance management, workforce and potential litigation have not yet been explored.

We recommended that DeepMind Health consider, with clinical and non-clinical professionals, the implications of their work for performance management, for litigation and for assessment of future workforce requirements.

DeepMind Health responded

As Streams is a patient-centric focused app, it has not been developed with the functionality for performance management nor for assessment of future workforce requirements in mind. With respect to litigation – all actions within Streams are fully audited and in the case of a medico-legal requirement these audit logs can be provided if required.

In 2017 we raised concerns about:

We recommended that DeepMind Health considers any infection risks and how they might be addressed.

DeepMind Health responded

As pointed out in the Hu-tech Human Factors report this is a Trust issue and not something within DeepMind's remit to either address or monitor. However, this has been highlighted to the Trusts' Infection Control Teams who are responsible for overseeing the infection control risk of all devices and instruments within the hospital. Staff already use their personal devices within the hospital environment and the current Infection Control policies and procedures are applied.

As part of our 2018 programme of work, two reports were commissioned to review DeepMind Health processes and the Steams app from a human factors perspective during the preparation for implementation of Streams at Taunton and Somerset NHS Foundation Trust. Dr Ian Randle of Hu-Tech, supported by Helen Gagg were asked to follow up on DeepMind Health's response to last year's recommendations, while Dr Jane Carthey was asked to review how the implementation of Streams in a second NHS organisation was being planned, and what lessons might be learned for the future roll out of such technology. However only preliminary insights into this second question can be drawn due to Streams not yet being live at Taunton and Somerset.

DeepMind Health have responded well to recommendations that were highlighted in last year's report and have continued to make Streams easier to use for frontline staff; for instance, the development of an alert to warn users of loss of wifi or mobile signal, continued development of Android capability (Streams is currently configured to work only on Apple/iOS devices) and considerations on accessibility standards (for instance larger fonts). An issue about surgical glove use and compatibility should continue to be a consideration in future roll outs as clinician behaviour may differ between Trusts. A suggestion that patient lists should be dynamic and should have filtering options (for instance the ability to filter by medical department) has been taken up, with some changes already made to the filtering options. Additional functionality in the form of the ability to triage alerts is regarded as a significant enhancement, particularly the removal of the yes/no response button 'will you see this person'. In the opinion of our experts, this response could be potentially ambiguous. The new functionality gives three options. The 'review recommended' that the alert remains in the inbox until it has been dealt with.

To view the Human Factors review of Streams click [here](#), To view the Human Factors and Ergonomics Report by Hu-tech click [here](#), and to view the Streams screenshot click [here](#).

Broader consequences and human factors/ergonomics

Two areas were considered in particular detail at the request of the Independent Review Panel. What the impact of 'alert fatigue' might be and where responsibilities for responding to alerts lie, especially given that one patient may potentially be under the care of several different clinicians.

The issue of responsibility is a complex one, and processes vary between individual Trust but the bottom line is that patient safety depends on there being a clear hierarchy, in order to avoid a patient 'falling between the cracks'. DeepMind Health have introduced a feature into Streams to support the National Early Warning Score (NEWS) which digitally support assessment of the patient, based on the six cardinal vital signs (respiratory rate, oxygen saturation, temperature, blood pressure, pulse/heart rate, AVPU response) and one other observation. Local Trust protocols around escalation are then sign-posted within the application. Further investigation of this functionality is recommended.

Alert fatigue may arise where high numbers of alerts are received and cause distraction or stress to such an extent that they may be ignored or switched off. Developing a smart system to manage high volumes of alerts was recommended in 2017. DeepMind Health are beginning to address this by looking at ways to deliver alerts to match specific roles or responsibilities, allowing users to 'subscribe' to particular alerts. Good progress has been made but it was recommended that it is kept under review.

Both reports emphasise that for each change made or required in Streams, a robust process of risk and hazard analysis should be carried out to mitigate the risk of unintended consequences.

However, there is a danger that in being very responsive to frontline needs the frontline themselves may succumb to the "wow" factor and may not always recognise why the design is developed a particular way or may even underestimate the learning needs of users less involved in the user-testing. This can lead to work-arounds which are sub-optimal and failure to respond to electronically highlighted patient issues. User organisations may need more assistance to ensure that training not only includes how to use Streams but also the underlying assumptions and risks.

While DeepMind Health are to be commended on their work to date it is also clear that continued focus on robust risk and hazard analysis, combined with continued user testing in the real and simulated world is essential. It's important to acknowledge that compared to the vast majority of providers of technologies and tools in healthcare DeepMind Health are setting a high standard in making Streams easy to use correctly and in the avoidance of unintended consequences.

We recommend

The expansion of user testing groups to include a wider range of staff and testing in more worst case scenarios.

That further development alongside human factors and safety experts of the risk and hazard management process, during proposed changes to Streams, occurs to ensure greater robustness and to focus on the development of more effective barriers to anticipate and avoid unintended consequences.

Public and patient engagement

Principles covered:

2 Public, patient and practitioner engagement

In 2017, we raised concerns about the lack of work on public engagement particularly in relation to links between DeepMind Health and Google and public perception that data processed by DeepMind Health could be shared with Google. We recommended that DeepMind Health developed principles for effective engagement and a strategy for public engagement in partnership with others, such as the Wellcome Trust. We also suggested that DeepMind Health should consider developing education programmes about AI and its uses in healthcare.

We commissioned an independent review of this work by Simon Denegri, the former Chairman of Involve and currently NIHR National Director for Patients and the Public in Research. This report is to be found in full in the Appendices. In particular, we wished to know not just if patients and clinicians had been engaged but how they were selected, reimbursed and if they were a representative demographic that would be able to articulate the differing points of view, particularly in relation to use of their data, that are known to exist in the wider public. We also wanted to know how lessons learned would be integrated and applied in the future.

Public engagement is the area of DeepMind Health's work that has seen the most radical transformation since our last report. We have been impressed by DeepMind Health's commitment to, and investment in, its developing partnerships with patients, carers and the public ('users'). Simon Denegri said DeepMind Health are 'showing an intent and commitment to public involvement, which clearly puts it in a leadership position in its own field of AI and data use'. We agree.

DeepMind Health can point to a number of key achievements including the establishment of a patient 'user group'. They held a number of patient involvement and engagement events during 2017, culminating in a 'Collaborative Listening Summit' which concluded with the development of a set of values, which will be published shortly. They have also established a set of five priorities for the coming year for

'involving patients and carers in the work we do'. These include growing and developing the DeepMind Health user group, working with partner Trusts, focusing the development of health technology and AI on the needs of patients and creating accessible education content for patients.

There is however, no cohesive plan for following up on these priorities and there is a need to develop a reportable programme of work by which they, the Independent Reviewers and others can measure progress. We also feel that the time scale is too short at a year. Three to five years would be more appropriate.

There seem to be diverse views within DeepMind Health about the purpose of their engagement work, with responses including 'social good', 'to increase the quality and utility of their products', 'to understand problems from the patient perspective'. All are valid reasons for engagement but we believe that being more explicit in their objectives would help them develop a plan of work more effectively against various objectives. Defining purpose would also help their user group, who whilst feeling valued by DeepMind Health, are not always sure what they are there to do. On the other hand, where there is clear purpose, for example in patient involvement in the breast screening project, it has resulted in an exemplary and important piece of co-production which is of benefit to all.

To view the Patient and public engagement review click [here](#).

Public and patient engagement

Realising their values

DeepMind Health have worked in partnership with the user group to develop a set of priorities and a 'ways of working' document. However, it is not clear what commitment DeepMind Health is making to users and it is recommended that they develop a leadership statement which articulates their commitment and why they are working this way. Their relationship with their user group and particularly their feedback to them needs to be addressed. They should think about developing a patient voice at the leadership team level, perhaps with, in time, a formal advisory board. Where they are working with NHS Trusts, they should include in their development agreements on involving and engaging users.

Sustainability

Their work to develop a toolkit for local projects is a good step towards more sustainable involvement and engagement. In addition, they should develop a mechanism by which local projects can share and learn from one another. It would be innovative if this were to include projects outside the UK. The User Group should be used in an advisory body on cross-project issues arising.

Communications

DeepMind Health have a very good story to tell about their involvement activity but are not yet telling it, at least not through their website. It is due for a complete overhaul but launch is planned in 2019. We believe this should be brought forward.

There is still a very notable gap in communication about their relationship to Google which we discussed earlier. (see P 12) We noted this last year and is still as conspicuous by its absence now, as then.

Internal engagement with their teams, educating them on user involvement and engagement, is almost as important as external engagement and again points to a single company narrative on purpose. DeepMind Health should be looking for opportunities to develop this further (induction programmes, learning and development opportunities, secondments).

Accessibility

In the same way that we have pointed to the need to engage Trusts that are behind the curve in their relationship with technology, DeepMind Health also need to ensure that they engage more patients from less advantaged communities, particularly those with disabilities and those from ethnic

minorities. This can be achieved through the user group but there is a need to review the way that participants are reimbursed (in M&S vouchers rather than cash). This should be considered as an accessibility issue, since it disadvantages the poorest most, and a solution sought urgently.

Broader public engagement work

We are impressed by the way that DeepMind more generally, has proceeded with its public engagement, trying out a number of approaches in a low-key way to establish what works. A good example is the company's involvement in the Cheltenham Science Festival, where in addition to some high-profile events in the main lecture programme, they also had a stand which helped explain AI to the public. This involved a number of more junior DeepMind staff, many of whom had received specific training in how to engage with the public. This has increased the number of DeepMind staff able to talk confidently to the public in non-technical language. There is good evidence that trust increases when 'ordinary' workers at a company, rather than the leadership, explain what their company is doing. In making a point of using as many female staff as possible, it shows that tech is not an exclusively male preserve. It also provides female role models, encourages young women to enter computer or data science careers which in turn provides a more diverse understanding of the impact of AI on society.

We recommend

That a company-wide clear purpose for their engagement work and a formal plan to realise their priorities in this area is developed.

That the lack of diversity in their user group, particularly in relation to ethnic and social background is addressed. Their current voucher payments system for their user group should be viewed as an accessibility issue and addressed.

Their website re-launch should be brought forward if possible.

Law, regulation and data governance

Principles covered:

7 Legal and ethical

10 Transparency

Our 2017 report was written in the wake of data issues raised by the agreement between the Royal Free Hospital and DeepMind Health and unsurprisingly, law, regulation and data governance had our greatest focus.

In our last report, we raised a concern about the lack of clarity in the original sharing agreement with the Royal Free London NHS Foundation Trust but by the time it was published, this had already been corrected.

In 2017 we recommended:

That DeepMind Health should respond positively to any recommendations that result from the ICO investigation.

The Information Commissioner concluded her investigation in July 2017, a day before our report was published. The ICO ruled that the Royal Free NHS Foundation Trust had failed to comply with the Data Protection Act when it provided patient details to DeepMind Health, for the purposes of clinical testing of the Streams application before it was fully deployed.

We note that her conclusion agreed with the legal opinion that we had commissioned for our report, that DeepMind had acted as a data processor for the Royal Free, which remained the data controller. DeepMind Health subsequently issued a statement in response to the ICO investigation.

ICO full letter to Sir David Sloman, Chief Executive of the Royal Free NHS Foundation Trust can be found [here](#).

DeepMind Health responded

Although today's findings are about the Royal Free, we need to reflect on our own actions too. In our determination to achieve quick impact when this work started in 2015, we underestimated the complexity of the NHS and of the rules around patient data as well as the potential fears about a well-known tech company working in health. We were almost exclusively focused on building tools that nurses and doctors wanted and thought of our work as technology for clinicians rather than something that needed to be shaped by and accountable to patients, the public and the NHS as a whole. We got that wrong and we need to do better.

In addition DeepMind Health made themselves available to the Royal Free as required to assist them with responding to the undertakings made to the ICO. Separately they were in contact with the ICO to proactively engage with them on the work they are doing with Streams outside of the Royal Free investigation. They have appointed a new Information Governance Lead with extensive experience of managing IG issues in the NHS.

Since this report went to press, the results of an ICO audit into the use of Streams by the Royal Free London Hospital have been published, conducted by Linklaters. We had not previously seen these results although were aware that the audit was being conducted. The audit concludes 'We consider the Royal Free's use of Streams is lawful. However, the audit has identified some areas in which further improvements should be made'. We note that this agrees with our own legal advice and urge DeepMind Health and the Royal Free London Hospital to follow the recommendations in the report, and to take steps to avoid concerns such as these being raised again.

Law, regulation and data governance

In 2017 we recommended:

That tech providers, the Department of Health and the Information Commissioner should discuss together a new system which protects patient data whilst allowing innovation and that collaborative discussions should take place in safe places, similar to Research Council 'sandpits' in order to create a new model for regulation.

This discussion has been ongoing throughout the last year and has involved a wide range of actors beyond DeepMind Health. The ICO is now engaging in early dialogue with developers (including DeepMind Health) and is also considering a 'sandpits' style approach.

We note that when developing workstreams outside of the United Kingdom, DeepMind Health have also dealt pre-emptively with privacy concerns, engaging with their regulators early in the process.

In 2017 we recommended:

That DeepMind Health should set, as a firm policy, that all future contracts with the public sector should also be published openly, with minimal or no redactions.

DeepMind Health responded

For all new agreements we have entered into for the provision of Streams (Taunton and the now executed Yeovil agreement), we have followed the form of the agreements signed with Royal Free London and ICHNT, with the data controller tailoring the provisions to their circumstances. These have more robust information processing provisions as compared to the original information sharing agreement. We have continued to suggest that partners join us in openly publishing agreements – and the Taunton and Yeovil agreements are on our website.

Although we believe DeepMind Health have encouraged their partners to publish agreements openly, this has not taken place in all cases. We believe that how much DeepMind Health insist on this is an important test of their commitment to transparency.

We recommend

That DeepMind Health continue to be very cautious about information governance, as a second wave of public concern, as happened with the Royal Free Hospital contract, could be very damaging to public trust.

We reiterate our previous recommendation that DeepMind Health should demonstrate its commitment to transparency by having a firm policy with current and prospective partners, that all contracts with the public sector should be published openly, with minimal or no redactions.

Security

Principles covered:

8 Protecting privacy

9 Secure

In 2017, we raised concerns about some minor security vulnerabilities that were found as a result of a security report commissioned by the IR Panel.

DeepMind Health has now dealt with these issues. Although we will continue to keep a watching brief on security, we see no need to repeat a further security audit this year.

DeepMind Health responded
A detailed internal review was undertaken. The three findings were mitigated by alternate design choices (for example, vulnerabilities that were specific to a web based product, while we have no web interface). We are working to address points anyway for the ease of future audits.
Specific work is in place to: <ul style="list-style-type: none">• place a limit on the number of requests from a given client within a given time, to prevent farming, is complete for login and planned for all other APIs• explicitly wipe sensitive data from memory etc is planned• remove the potential use of the 3DES cipher (which would not normally happen in practice and which is, anyway, secure but not preferred) is planned, though blocked on changes to 3rd party open source infrastructure – we will also remove use of SHA-1 at this time, for similar reasons.

Independent review panel governance

The Independent Reviewers raised a number of points about governance.

The Reviewers expressed concerns that, whilst members of the panel are currently unpaid as a marker of independence, this is not a sustainable model as the workload becomes more complex. They also commented on the ad hoc method of selection for Independent Reviewer members, suggesting that there will need to be thought given to how future members will be chosen, how long they serve for, and how Chairs are chosen. They recommended that DeepMind Health consider financial models for reimbursement of Independent Reviewers and that DeepMind Health should support Reviewers on membership decisions.

DeepMind Health responded

The Independent Reviewers proposed that there should be an honorarium – £6,157 per annum paid quarterly, with the payment being made to either the individual, to a limited company or charity. Each Independent Reviewer would be required to sign a pledge which can be found in the Appendices but there would otherwise be no contract requirement, other than to attend at least 3 out of 4 meetings during the year. They proposed an additional premium for the Chair to reflect the extended time commitment and workload.

It was suggested that Independent Reviewers will be appointed for 1 year with a suggested appointment of up to 3 years (potentially extendable) and that they would work with the Independent Reviewers on membership. Candidates for new Independent Reviewers will be nominated by an appointment panel consisting of the Chair of the Reviewers and a representative of DeepMind Health. However, the final decision whether to approve a new Independent Reviewer appointment or not will rest with the Head of Applied AI, DeepMind.

DeepMind Health requested that the Independent Reviewers give DeepMind reasonable notice (at least five working days) of anything that they intend to publish, separately from the yearly review statement, about what they've learned and concluded as part of their work, and give them opportunity to respond.

The Independent Review Panel agreed on all these points. Recently, three new Reviewers have been appointed: They are David Lock QC, Maxine Mackintosh from OneHealthTech, and Jeremy Taylor from National Voices. They will take over from Mike Bracken, Martin Bromiley, and Richard Horton.

The Independent Reviewers pledges can be found [here](#).

Machine learning and artificial intelligence are difficult concepts. We are glad to note that DeepMind Health are now producing easy to understand summaries, as the one below.

Machine learning Q&A

What do you mean by machine learning?

- In general, machine learning is a type of artificial intelligence that refers to the development of programs that are able to act without being explicitly programmed.
- A program will use a variety of techniques to build a mathematical description or model, based on the data shown to it and so 'learn from' this data.
- Machine learning has proven to be a very successful approach to solving a range of different problems from self-driving cars to medical diagnosis.

What do you mean by 'training models'?

- When we talk about training a model, we mean showing lots of information to the computer program so that it can learn how to interpret that information.
- For each example we show, we want the computer program to provide a decision – for example, to give a diagnosis for each mammogram scan we show it.
- If the decision is wrong – as confirmed by a labelled and correct dataset that we can compare the algorithm – then the model will learn from its mistake and amend its decisions in future.
- This process is known as training.

What about "Deep Learning"?

- When we build large computer programs for machine learning we often refer to this as deep learning.
- The word "deep" comes from the number of 'steps' information is passed between before the computer program provides a decision.

How does this differ from Artificial Intelligence?

- Artificial Intelligence, or AI, is a more general term than machine learning AI refers to a very wide branch of computer science that attempts to build machines capable of intelligent behaviour.
- Machine learning is a type of AI.

Why is machine learning better than other methods?

- There are several reasons you might want to solve a problem with machine learning.
- Machine learning programs will learn patterns that might otherwise not be picked up – this can be useful for problems involving complex data, such as those in the medical domain.
- Machine learning isn't always better than other methods, but for the type of imaging and healthcare related problems DeepMind Health hopes to address in breast screening and eye disease, for example, it is seen as the best solution.

