DeepMind Health Independent Review Panel Annual Report

July 2017

About DeepMind Health

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

DeepMind began working with the Royal Free in July 2015 after being approached by clinicians from the Trust and DeepMind Health (referred to from hereon as DMH) 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:

Mike Bracken
Martin Bromiley OBE
Elisabeth Buggins
Eileen Burbidge
Richard Horton
Julian Huppert
Professor Donal O'Donoghue
Matthew Taylor
Professor Sir John Tooke

Contents

Appendix links

Independent Review Panel biographies

Independent Review Panel pledge

Streams protocol

Streams timeline

Streams metric

Streams service agreement

Retinal imaging protocol

Retinal imaging at Moorfields timeline

Radiotherapy for head and neck tumours protocol

Radiotherapy timeline

Verifiable data audit blog

First Base datacentre and development review

RFL Information Processing agreement

First Base datacentre server based review

First Base mobile iOS application review

First Base web API penetration review

Clinical utility assessment AKI Streams

Clinical utility assessment retinal imaging

Clinical utility assessment radiotherapy

Patient and public involvement & engagement (PPIE)

Rosamund Snow recommendations

Hu-Tech ergonomics report

IR declaration of interest

Foreword

By Dr Julian Huppert, Chair, Independent Review Panel

When a patient is treated, their care is informed not only by their own health data – their medical history, test results, imaging and so on – but also by the health data of thousands of other people. In fact, it is health data from others that will probably play the biggest role in making a patient well again, for it gives medical teams the essential insights and knowledge on which to base their care. Current medical treatments are almost entirely dependent on previous studies.

Access to the health data of many brings undoubted benefits for patients and for health systems. It can show what works and what does not. It can bring better care which may also be more cost effective. It is a platform for innovation in prevention, care, surgery, diagnostics and medicines.

But many of the data systems in UK hospitals are still paper-based. They are complex, unwieldy and insecure, and the data they contain is difficult to manage, which means that care improvements that could be made, as a result of research using the data, are not discovered or acted upon. This is a particular concern when there is huge and unacceptable variation in care across the country. DeepMind Health (DMH), through their systems of managing data, envisage a world where data is used securely for the benefit of patients and the public and where unwarranted variability becomes a thing of the past. This is an immensely exciting prospect.

However, there is also a concern about private companies having access to health data. This information is often intensely personal. Whilst many studies have shown that in general people are happy for their personal health data to be used for research, there are also very real and understandable concerns about its privacy and confidentiality. These are magnified if a commercial organisation is involved.[1] In the case of Google, which is part of the same corporate group as DeepMind, under Alphabet Inc.

There may be additional worries given how much Google is already perceived to know about us. At what stage does an organisation like Alphabet become simply too powerful?

These concerns should be applied to any organisation dealing with health data. We are well aware that many organisations, both private and within the NHS, do not always act according to the appropriate standards for data security and to maintain privacy. Indeed, we have heard some alarming messages about how medical data is currently dealt with in many cases. However, that does not excuse DeepMind Health from any failings. 'Good enough' is not good enough for a company linked so closely to Google, a company that already reaches into every corner of our lives. Even if untrue, the perception that Google might acquire highly sensitive health records, in addition to data already held by them about an individual, makes some feel uncomfortable no matter what the benefits might be.

We believe that it is right that DeepMind Health should be held to higher standards, even if that means they are singled out as a lightning rod for public concerns.

We believe that it is in DMH's own interest to live up to a higher standard, and to be seen to do so. All their efforts will be in vain if the general public have a catastrophic lack of faith that they will act correctly.

It is refreshing that when DeepMind Health was established in 2016, it unusually – and possibly uniquely – for a commercial company, also established an Independent Review Panel, with almost unfettered access to what they are doing. Our panel consists of nine reviewers, each working in different but related areas, all encouraged to take part because of their reputation, experience, and a belief in their integrity.

We are entirely independent, and are not subject to any form of non-disclosure agreement. Indeed, our agreement with DeepMind Health, is explicitly clear that we are not subject to binding secrecy rules, and are free to speak to the press however and whenever we wish.

We are self-governing, with our own secretariat and are free to set our own agenda and timetable – our only constraint is to produce an annual report.

1. Wellcome Trust 2016 'Public attitudes to commercial access to health data' www.wellcome.ac.uk/publicattitudes

Foreword

We have an annual budget of £50,000 from DMH, with full budgetary control, to fund the secretariat and hire outside expertise where necessary. We are free to hire whomever we wish, with no approval required from DeepMind Health.

We know of no other commercial organisation that has set up an independent panel in this way. We hope that other organisations will consider taking such a bold step towards transparency, with its risks and benefits.

This is our first annual report. In the year since DeepMind Health was founded, a lot has happened. They launched, expanded, and started development projects, research schemes, and worked on clinical applications with the NHS. There is a huge amount to cover and we cannot scrutinise every single aspect. We have therefore made the decision to focus our attention on a number of key issues; governance and their legal consequences, data security, clinical outcomes, patient and public involvement and engagement, and the broader consequences of DeepMind Health's work.

Our work has been informed by our various experiences and discussions with DMH, who have been very open with us. We have also commissioned a wide range of independent external people and organisations to investigate specific aspects. This work was all funded by DMH, but the decision as to who to ask, and what the brief should be, was ours. We very much wanted this report to be a readable length, to encourage people to read it in its entirety, and so the reports themselves are largely in the appendices.

In our second report, we will report back on the issues raised in this first report. We will also cover other important areas that we have not yet been able to cover in depth. We anticipate that this will include studying the DMH business model and its consequences on other companies, and how their systems would roll out around the NHS, including the wider consequences of that. We are open to suggestions for further work or topics that any members of the public, patients or other experts feel should be considered.

Please click **here** to view our full biographies and **here** to view our agreement with DeepMind Health.

The Independent Review Panel can be contact via email at **independentreviewers@deepmindhealth.com**

I would like to add my own deep thanks to all the members of the Review Panel. It has been a pleasure to work with you all and to learn from you. I would also like to thank Rebecca O'Leary and Vivienne Parry for their support with our organisation and turning our ideas and conclusions into text.

We believe that DeepMind Health was set up with a desire to act ethically, and to be seen to act ethically. We describe here how we feel that they have performed so far. We have highlighted areas where we believe they deserve commendations, areas where we have raised concerns, and recommendations for the future. We hope that they will respond constructively.

Julian Huppert
Former MP, Director of the Intellectual Forum at Jesus College, Cambridge, and Audit Chair, Cambridgeshire and Peterborough Clinical Commissioning Group (CCG)

Why an Independent Review Panel?

By Mustafa Suleyman, Co-Founder, DeepMind Health

When DeepMind was founded in 2010, there were two core parts of our mission; to develop AI and to use it to make the world a better place. Initially, we concentrated on the first but soon began to ask ourselves what ethical applications of AI might look like. We investigated many fields – food distribution, nano materials, transport, drones and drug discovery but it was healthcare where we felt we could contribute most, especially in the processing and managing of data, in real life, high stakes health environments. We were convinced that this could not only lead to better care for patients but that the data generated, would provide a crucial asset for research in many different fields.

We chose to begin with Acute Kidney Injury. It is a significant problem in hospitals, its prevention is challenging, it has important consequences, it affects people of all ages and it is very generalizable to other conditions. We understand that what we build must be intuitive to use and have fidelity with flexibility, allowing it to be customised and personalised to suit local context.

There is no road map for what we are doing and we started with a blank slate. This is an opportunity because we can build the diverse community from many different sectors and stakeholders that is an essential foundation for genuine co-production and with it, trust. But it also means that we must be humble and admit that we don't have all the answers.

We know we will make mistakes but we should not be afraid of them. Setting up a panel of independent reviewers builds a stable set of partnerships and creates a safe space for critical comment, where we can share our mistakes and then start to correct them. Companies that open themselves to scrutiny are companies whose practices will change.

I did not know any of the panel members before their appointment. We knew which fields needed representation – for instance human factors or clinical outcomes – and then sought exemplars in those categories. We chose people who had specific expertise but also reputations for integrity, who did not hold back, who could be angry and critical. We need people like that, to yell at us. That's good for us and makes us better and more reflective.

I strongly believe that good people will do the right thing. If we expect them to trust us, we should trust them. We should not need them to sign the extensive non-disclosure agreements typical of companies, although I admit that our lawyers took a little more convincing on this point. Instead, we agreed a simple pledge.

We have also provided funding to enable them to commission additional expertise where they feel it is required and to support a secretariat. We will increase the funding if required.

My background is in charities and in public policy. From the outset public and patient involvement and engagement was very important to me. I saw it as an opportunity to start afresh and begin something that was genuinely co-produced. I met up with Rosamund Snow, then Patient Editor of the BMJ and asked her advice. She was fearless, straight talking and sceptical. We had three meetings before she agreed to become involved. She was radical and wanted to change the world but equally was pragmatic, considered and deliberate. She drew up the principles for public and patient involvement and engagement. Her sudden death was a shock and an enormous loss but we are strongly committed to building on the principles she established with us, patients and public.

Rosamund was a truly remarkable person. I hope that what we started together will continue to reflect her vision as well as ours and that what she started, we will finish in her spirit.

What healthcare needs are being addressed?

DeepMind Health (DMH) are engaging in a wide range of activities in healthcare. Some of these are essentially more technologically focused, such as ensuring that data can be transferred and stored securely, and presented to the right clinician at the right time. Others involve the use of machine learning to achieve tasks that are currently difficult for clinicians to perform accurately and rapidly.

Whilst some, rightly question the security of electronic data, currently highly insecure paper-based systems predominate. Paper-based data can easily be lost, copied, or retrospectively changed, and it is hard to securely transfer it to where it is needed. The advantages of having a secure and functional electronic data system are clear, if it is done correctly.

DMH has selected a few specific healthcare needs to address so far, which we outline below. These are the first tentative steps in showing how the digital revolution that has swept through the rest of society, but which has been late coming to health services, might benefit patients in the future. We are very excited at the potential this might have to improve patient experiences and in particular, to eliminate the unacceptable levels of variability in healthcare seen across the country.

Acute kidney injury at the Royal Free London NHS Foundation Trust – Clinical use

The Streams Project: a prospective, observational study of a new digitally-enabled care pathways for the recognition and management of acute kidney injury on a single hospital site

Acute Kidney Injury (AKI) is sudden damage to the kidneys that prevents them working properly. This ranges from minor loss of kidney function to complete kidney failure requiring temporary or even lifelong dialysis. It can also result in death. It is common in hospitals, found in 13–18% of those admitted and is a major issue in critical care where up to 30% may be affected. It is associated with many different causes including a wide range of serious conditions such as major surgery, burns, trauma, sepsis, heart and other organ failure and reactions to certain medicines as well as existing kidney and urinary tract problems.

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

Lab results are normally viewed by hospital clinicians on office-based computers at best, on scribbled bits of paper or written on their hands at worst. Finding the time to check whether results are back and what they might mean for the patient can cause delays of hours or even days. By the time the patient has developed symptoms of AKI, substantial damage is already likely to have occurred.

NHS England has recently introduced a standardised algorithm, which helps to identify patients that have already begun to develop AKI, although this depends on patients having had a baseline measurement of their creatinine levels.

Streams is a way of displaying the alerts developed by the standardised algorithm to clinicians on portable hand-held devices. It does not currently use machine learning although it is envisaged that in the future AI driven alerts could be delivered by it.

Please click **here** to view the Streams protocol and **here** to view the full project timeline.

What healthcare needs are being addressed?

Streams use at Imperial College Healthcare NHS Trust – Clinical use

Following final approval of the Trust Executive Board, a project is planned at Imperial College Healthcare NHS Trust where the use of Streams for clinical task management will be evaluated within a hospital using the Streams interface as a viewing platform for task management only.

Retinal imaging at Moorfields Eye Hospital – Research

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 to diagnose a range of common, sight threatening 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. There is also an impact on the time taken to diagnose the condition.

The DMH approach involves the assessment of a number of automated analysis methods.

Please click **here** to view the retinal imaging protocol and **here** to view the full project timeline.

Radiotherapy planning for head and neck tumours – Research

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 incorrect 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 delays the start of treatment and it 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 DMH approach involves applying machine learning to perform automated segmentation of head and neck tumour volumes and organs on radiotherapy planning CT and MRI scans.

Please click **here** to view the full radiotherapy protocol and **here** for the project timeline.

Law, regulation and data governance

The most important question faced by the Independent Reviewers related to the use of patient data supplied by the Royal Free Hospital, in connection to the Streams project. However, a number of other issues were also considered in our analysis.

The wider landscape

Our analysis needs to be located within a wider landscape relating to the use of healthcare data. First, the failed introduction in 2014 of care.data, an NHS scheme intended to widen the collection of NHS hospital data to include GP data. One of its most controversial aspects was the sharing of this data with commercial companies. Following heavy condemnation, the scheme was dropped but widespread concern about commercial access to health data has persisted, even when it is legal and of clear patient benefit. During this fiasco, it became clear that there was little public understanding of the purposes or benefits of the use of data in healthcare, and little understanding in the NHS and government of the public's concerns.

Secondly the large number of cases of data loss or breach reported in the media, further heightened concerns about the confidentiality of sensitive personal health data. Many people are, quite reasonably, concerned about third parties seeing their medical data.

Finally, the corporate relationship between DeepMind and Google which in the public's mind at least, is widely presumed to know everything about everyone through the services it supplies. This means that DMH has to assume a much greater level of corporate responsibility than is the case with most start-ups. As we have said before, it is right that DMH should be held to a higher standard.

The regulatory and technical landscape

Following the concerns about care.data, a National Data Guardian (NDG) (Dame Fiona Caldicott) was announced in 2014 by the Department of Health. The role of the NDG is to strengthen the security of health and care information and help people make choices about its uses. Data protection issues are overseen by the Information Commissioner (ICO), the UK's independent authority set up to uphold information rights in the public interest.

Of note is that regulation, both of data and of medical devices, relies on definitions which are rapidly being superseded by new developments in technology and which have simply not caught up with modern day usage. To some extent, regulation will always be struggling to catch up with advances in technology. However, the current position means that patients are≈sometimes denied the benefits from new technologies inappropriately, and that their data is sometimes over-shared in ways they are not comfortable with. This creates a perhaps inevitable tension between public goods – for instance between that of privacy and patient benefit.

There is a compelling need for a new mechanism which protects the public interest whilst not delaying the introduction of new technologies. It is often suggested that there is a necessary trade-off between the benefits of privacy and the benefits for patient care. We do not accept that. We believe regulation can and must drive improvements in both. We believe that the current system both overprotects some information and under protects other information, and can be improved.

We think it would be helpful if there was a space, similar to the 'sandpits' established by the Research Councils, which would allow regulators, the Department of Health and tech providers to discuss these issues at an early stage of product development. The protection of data during testing is an issue that should be discussed in a similar collaborative forum. We believe that there must be a mechanism that allows effective testing without compromising confidential patient information.

Background to the relationship between DeepMind and the Royal Free Hospital

DMH began working with the Royal Free in July 2015 following an approach made to it by clinicians. Later that same year, DMH began receiving data from the Royal Free to enable it to test its application Streams. Between October and November 2015, confidential, identifiable information of 1.6 million patients covering a 5-year period was transferred from the Royal Free to DeepMind for the purposes of testing of the Streams product.

The original Information Services Agreement (ISA) was between Google UK Ltd and the Royal Free. Following criticism that appeared in a feature in New Scientist in April 2016 and subsequently in a number of media outlets, a new ISA between DeepMind Technologies and the Royal Free was agreed in November 2016 and the original ISA terminated. This is much improved and more typical of other NHS customer/IT supplier relationships.

As a matter of policy, to promote transparency, DMH have published their contracts with the Royal Free (and others) with minimal redactions. We have seen unredacted versions.

In March 2017, DMH announced that it was building an auditing system for healthcare data. DMH's system will record and later verify every event related to hospital health data, thus allowing it to track data. This so called 'immutable logging' has been widely compared in the media to the way blockchain works to track every event related to bitcoins, although it does not in fact use this technology.

This approach has huge potential to prevent misuse of data, and provide much increased reassurance to patients and regulators about how data has actually been used. It is not a substitute for open discussion about how and why data should be used, but should allow many fears to be eliminated – and appropriate remedies to be applied if things have gone wrong.

A further recent and very important development by DMH is the establishment of large-scale simulated hospital data. This is data that is manufactured to look like genuine patients on a typical hospital ward but which has not in fact come from real patients. It is carefully clinically curated to simulate genuine patient data, but does not carry the risks of inadvertent disclosures for developers nor generate concerns from patients that might arise were developers to be working with anonymised or pseudonymised test data. This will mean that the demonstration and testing of systems can take place without fear of exposing any patient identifiable data. DMH have indicated that it has begun to make this simulated data stream available to a few other companies to develop their own products, and have indicated that they will make this more widely available.

Legal opinion

In order to assist our scrutiny, we commissioned legal opinion from Peter Wainman, an expert in law relating to technology and data use at legal firm **Mills & Reeve**. We asked for analysis of DMH's data sharing arrangements and relevant privacy measures for compliance with applicable law, regulation and best practice. We specifically asked:

- Whether DMH was acting as a data controller or a data processor in respect of the patient data that it received from the Royal Free?
- Whether DMH had breached the Data Protection Act and what implications there might be for DMH if this were to be the case?
- Whether the redactions shown in contracts were appropriate?
- Whether DMH has breached confidence?

We made it clear that the purpose of the analysis was not to provide a defence of DMH but instead to lay out an entirely independent legal opinion, which would assist the Independent Reviewers in their scrutiny. DMH made full access to unredacted documents available and offered Mills & Reeve the opportunity to ask any questions or request any additional documentation necessary.

During the course of our work, there have been a number of exchanges of letters between the Royal Free, DMH and the NDG and Information Commissioner. We have seen these and one has been leaked to the media. The key issue they are looking at is the data transfer between the Royal Free Hospital and DMH. We had expected regulators to conclude their investigations before we released our report, and do not want to prejudice their detailed investigation. In addition, we have no role in reviewing the actions of the Royal Free, and have not discussed any of these issues directly with them.

A blog by DMH on verifiable data audit can be viewed **here**.

However, given the importance of this issue, we did want to report publicly on the legal advice we received about the position specifically of DeepMind Health. In summary, our legal advice found that DMH had acted only as a data processor on behalf of the Royal Free, which has remained the data controller.

It found no evidence that DMH had violated the data sharing agreement or any other contractual arrangements with the Royal Free. It found no evidence to suggest that DMH has breached confidence.

Consequentially, any issues about data protection obligations or confidentiality obligations arising from the use of patient data during testing are in law, matters for the Royal Free as data controller, and we will not comment further on them. We look forward to the final outcome of the ICO's investigation, and will return to this matter when that is published.

In addition, the legal opinion notes that, they believe DMH did not do enough to allay concerns expressed in the media, that patient data would be combined with other data held by Google and used for other purposes (even though this would have breached the contracts between DMH and the Royal Free). There is no evidence that DMH had or have any intention of doing this.

We commend

DMH for openly publishing their contracts with the Royal Free Hospital with only minimal redactions.

DMH for their developing work on immutable logging, so that people can be assured of how their data is used.

DMH for generating detailed simulated hospital data for testing purposes, and allowing others to use it.

We have raised concerns about

The lack of clarity in the original information sharing agreement with the Royal Free Hospital, although this has since been corrected.

We recommend

DMH should respond positively to any recommendations that result from the ICO investigation.

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

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.

Technology and security

Patient data is highly sensitive and its security and privacy has always been a central concern for the Independent Reviewers. We wanted to assure ourselves that the data was being handled securely, and that there were no serious vulnerabilities at any stage of data being stored. In broad terms, there are four potential risk areas given the current scope of the DMH projects.

The front door – The Streams app

The Streams app runs on iPhones and iPads only, not Android-enabled devices. It is used by clinical staff to both input and retrieve information and is the "first line" or "front door" for all the DMH systems that are behind it. We wanted to be assured what vulnerabilities there were to this interface. For example, if someone deliberately stole or accidentally found a clinician's device, could they access patient details or change them? Or could data that was input through the app be accessed or compromised through other access mechanisms?

Border security – APIs

Application programming interfaces (APIs), are the "borders" or integration points for DMH applications and any other services and applications, either in the NHS now or, in the future, perhaps outside the NHS (for instance in social care). The APIs should allow data from known, trusted and secure parties to pass to other trusted and secure parties, but not any others.

Data storage – Datacentre

The datacentre is where all the data referenced and incorporated with the DMH applications reside. For that reason, security and access controls here are of paramount importance. This data must be protected from both remote hacking and physical entry.

How the app and other projects are developed in the future – Development processes

This area is critically important because while we can make best efforts to assess what is currently available and in use with the NHS, looking at the DMH general development environment should give reassurances or alternatively expose risk with respect to maintenance of existing systems and any future developments. Security by design is an important principle.

The review process

We do not have the expertise to assess the actual technical implementation and security of the DMH applications ourselves. We therefore sought specialist advice, inviting three different security analysis companies to provide an initial scope and quotation and ultimately, we commissioned First Base Technologies Limited. Over a two-week period, security analysts from First Base liaised with the Independent Reviewers and DMH to arrange the necessary access and resources to properly conduct their work, which included entry to the DMH datacentre. The cost of this review was met by DMH as part of our annual budget but they were not involved in the procurement.

Only those systems that were entirely within DMH's control were reviewed by First Base. They were not asked to consider, for instance, the security of data systems within the Royal Free Hospital or evaluate the rigor and reliability of NHS standards or requirements which DMH was required to adhere to.

The Independent Review team were pleased with the comprehensive nature of First Base's work and reports relating to each of the four areas above are in the appendices. It is very reassuring to note that there were no critical or high-level vulnerabilities detected. In total 11 vulnerabilities were identified, all of which were low in severity with a single exception which was medium. The Independent Reviewers endorse the recommendations made by First Base, including and most especially the recommendation for on-going penetration and other testing to continue on a regular basis.

In order to review the full report supplied by First Base, please go to:

- **First Base datacentre and development review**
- **First Base datacentre server based review**
- **First Base mobile iOS application review**
- **First Base web API penetration**

Mobile iOS app (Streams app)

"...Overall security of the application is excellent, with no critical or high-risk vulnerabilities identified" The report also highlighted the points that need to be considered when implementing a mobile/BYOD solution, given the increased risks associated. Some form of mobile device management is necessary, and DMH

Detailed report commentary & highlights **Aggregate vulnerability count**

	Streams iOS app	Web API	Datacentre server build	Datacentre & development
Low	2	8	0	0
Medium	0	0	1	0
High	0	0	0	0
Critical	0	0	0	0
Total vulnerabilities	2	8	1	0

should ensure that this is always considered properly when schemes are being implemented.

Border security – APIs

The API is not intended to be openly accessible, but only via a dedicated VPN. This significantly reduces the risks, although the report found a number of issues, which they rated as low vulnerability. We have chosen to highlight one of these, related to authentication. The system is not configured to require regular logins, but preserves an access token on the device, including the user name. Although it requires high levels of technical capability to extract a user name from the APIkey, it is still a legitimate vulnerability in the context of a sophisticated attack. Therefore, it would be technically straightforward to require password credentials at every use. However, we recognise that not requiring this may have been a trade-off between usability, especially in high pressure situations, and security. DMH may wish to consider other ongoing forms of validation, such as fingerprint or voice recognition.

Datacentre – Server build

The study of the server design revealed one medium level vulnerability. The report says that this vulnerability "should be addressed but are not thought to present an immediate threat to the environment or data handled by it".

In addition, as DMH employs a Linux-based environment, it was relatively protected from the very recent WannaCry global ransomware attack and other similar types of attacks which target Windows servers/environments. Even so, staying abreast of the latest versions and patching would go a long way in mitigating such threats.

Data storage – Datacentre, physical security & development review

This report was quite clear about the robust nature of DMH's development and datacentre processes and identified zero vulnerabilities. The report stated that "The review of both the datacentre and the software development processes revealed a very high standard in both areas... with no vulnerabilities of any risk level being identified",

"The overall security of the datacentre is excellent" and "The software development practices follow industry best practices and an extremely rigorous development process is in place".

We commend

DMH for their overall high level of data security, reaching industry best practice in many areas.

We have raised concerns about

The relatively minor vulnerabilities that were found.

We recommend

DMH take appropriate steps to deal with all the vulnerabilities identified and mitigate any risk revealed.

Clinical outcomes and clinical utility

For innovative technology to be widely adopted in medicine, it must have genuine clinical utility. By this we mean that it not only addresses an important clinical problem, but also that its deployment would result in significantly improved outcomes for patients, it is not too complex or resource intensive to implement and use, it does not solve one problem only to create others and finally that it has the potential for widespread adoption and diffusion.

Innovative technology always evokes excitement and inevitably claims are always made that it is going to 'change the face of healthcare'. We genuinely believe this has the potential to be the case for the technologies being introduced by DMH. What they offer is not IT, although it may be perceived as such, but a step change in the way problems are managed, which creates huge potential for improvement in clinical practice in a health service which is incredibly complex. These improvements will also help reduce the unacceptable variations in healthcare that are seen currently. It also has implications for workforce, for training and management.

The digital revolution has largely bypassed the NHS, which, in 2017, still retains the dubious title of being the world's largest purchaser of fax machines. Many records are insecure paper based systems which are unwieldy and difficult to use. Seeing the difference that technology makes in their own lives, clinicians are already manufacturing their own technical fixes. They may use SnapChat to send scans from one clinician to another or camera apps to record particular details of patient information in a convenient format. It is difficult to criticise these individuals, given that this makes their job possible. However, this is clearly an insecure, risky, and non-auditable way of operating, and cannot continue.

'The average NHS trust has 160 different computer systems in operation'

Part of the reason has been that digital solutions have been laid on top of the hundreds of different ageing IT systems employed, even within a single NHS Trust or Authority. The technologies employed by DMH use FHIR (Fast Healthcare Interoperability Resources) a very secure interoperable data system that allows both the interrogation and receipt of data from other systems, using a common language. In lay terms it is the babel fish of data sets allowing them to talk to each other. Thus, its operation does not require a complete re-engineering of existing datasets or computer systems.

However, at the heart of DMH's work is an interest in the clinical benefit that can be achieved, not just the pure technological innovation. The clinical areas that they are looking at are important, but none of them has yet had its operation in a real clinical situation fully assessed. This is being done for the Streams project, and although early results are very promising, we await the end of the project for fuller analysis.

However, we did want to look into the clinical potential, as it is so core to the purpose of DMH. We therefore, sought experts who could look into the work so far and produce a review as to the clinical utility of each of the projects. In each case we contacted the relevant Royal College for their advice on the most appropriately qualified reviewer.

For the Streams project, the reviewer was Dr Andrew Lewington, Consultant Renal Physician at Leeds Teaching Hospitals Trust.

For the retinal imaging research project, the reviewers were Professor Peter Scanlon, Consultant Ophthalmologist, Gloucestershire and Oxford Eye Units and Dr Irene Stratton, Senior Statistician in the Gloucestershire Retinal Research Group.

For the radiotherapy segmentation research project, the reviewer was Dr Petra Jankowska, Consultant Clinical Oncologist and Head of Radiotherapy at Musgrove Park Hospital Taunton.

We very much appreciated the enthusiasm of the Royal Colleges and the time and commitment of the reviewers.

Please go to the following links to review each clinical utility report, **Steams AKI report**, **retinopathy report** and **radiotherapy report**.

Clinical outcomes and clinical utility

A number of common themes have emerged:

- All the reviewers were positive about the potential opportunities raised by this technology.
- There were challenges in relation to definitions and benchmarking in each of the projects.
- The complexity and scale of challenge of introducing a new approach had not always been fully appreciated.
- There should be greater focus on implementation challenges from the outset.

A number of specific technical issues were raised by reviewers which are detailed in the reviews. With regards to Streams, it is worth noting that early detection and alerts are just one part of the AKI pathway which needs to be examined in totality, so that fixing one bottleneck does not simply lead to other problems being overlooked or assuming greater importance.

More generally it is clear that if initiatives are to achieve scale and reach, they need to have broad engagement as part of the initial planning and scoping process. This engagement needs to go beyond clinicians and patients to include clinical colleagues in other departments, managers and experts in quality improvement and implementation science.

If a technology is to be genuinely game-changing, there is great value in thinking through potential implementation problems up front. We think that the medical royal colleges and other professional bodies of clinicians, including nurses, could play an important role in this process, particularly in supporting a peer reviewed development of protocols and also the dissemination of learning.

We commend

The efforts DMH have made to seek to improve the quality of medical care in the NHS.

The potential opportunities presented by the use of these technologies.

The commitment from DMH to using the open and interoperable FIHR standard, which will help other providers to innovate.

We have expressed concerns about

The complexity and scale of the problems not yet being sufficiently understood.

The method of development which may have potentially adverse implications for successful adoption and diffusion across the wider health system.

We recommend

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.

Patient and public involvement and engagement (PPIE)

Patient and Public Involvement and Engagement (PPIE) has been an important element of the DMH approach from the start. Rosamund Snow was commissioned by DMH to develop recommendations for PPIE. Following Ros' untimely death, the Independent Panel commissioned Simon Denegri, National Director of Patients and the Public in Research at NIHR and Chair of INVOLVE to provide a further review.

Principles for effective involvement projects

1. All contributors know what their roles and scopes are.
2. The patients are asked to advise because of specific perspectives, experience, connections or expertise relevant to the project.
3. The patient's expertise informs the project from the earliest design stage.
4. The patients have the support they need to discuss ideas equally alongside other contributors.
5. The patients are valued equally alongside other contributors including remuneration.
6. All contributors have a chance to make informal or social connections with each other as well as formal ones.

The Independent Reviewers have been impressed with the commitment that DMH has shown for PPIE. Rosamund Snow laid out six key principles that are core to effective involvement projects in her excellent initial report and DMH will need to continue to build on these.

Simon Denegri has suggested that DMH might wish to work with citizens to turn these definitions of 'effective involvement projects' into a set of business principles or values for how they work with patients and the public. We agree with Simon that DMH should be more ambitious with regards to the second principle; creating a space which would enable idea creation with citizens on a continuous basis with an equal commitment to develop the capabilities of a number of citizens every year.

Simon identified some practical options for how DMH might respond to Principles 3, 4 and 5 (see report below). No matter which route is followed by DMH, we would like the active involvement and engagement of patients, their carer's and representatives to continue.

Some work is needed on the positioning of DMH's work with patients and the overall tone and style of its narrative which can sound old-fashioned. For example, on the website, DMH has a tab which says 'For patients' rather than the preferable 'With patients.'

Patients are members of the public and members of the public will become patients. Nevertheless, they are quite distinct audiences with different concerns, attitudes and priorities. For example, patients tend to have fewer concerns about sharing their health data than the well public but it is the well public's attitude on the use of health data that tends to prevail in the media. What both audiences share are concerns about commercial access to health data which were the main driver for the furore over care.data. It should come as no surprise, especially in light of the toxic debate on care.data, that DMH's proximity to Google would attract particular anxieties. As far as we can ascertain, DMH does not share its data with Google, yet the public perception that this might be the case, now or in the future, will be difficult to overcome and has the potential to delay or undermine work that could be of great potential benefit to patients. It is of particular concern to the Independent Reviewers that DMH's public engagement work is less coherent or visible than its work with patients.

We do not wish to suggest that DMH has not been active in the media. In fact, it has repeatedly made the point that data collected from the app is not shared with Google. However, media outlets frame messages in their own way, often with headlines that may infer an opposite intent and this message is not gaining traction. This speaks to the need for a wider engagement with the public. The question of trust and health data is not one that DMH faces alone, so a collaborative and strategic

work with other organisations with interests in the use of data for medical benefit, such as the Health Data Research UK (previously the Farr Institute), Wellcome, MRC, Association of Medical Research Charities and Genomics England might be helpful in building trust. There are a number of other potential avenues for public engagement including deliberative work with the British Science Association for instance or debates at any of the very well attended UK science festivals. There is also the opportunity to model a new type of engagement strategy using some of the models developed by the World Economic Forum Center for the Fourth Industrial Revolution in San Francisco. Verified data audit could be the game-changer here, but again, it will require wide collaboration with others to bring the concept to the public.

Where DMH has a special role is in explaining the potential of and problems with Artificial Intelligence (AI) in healthcare as this is little understood by medical professionals, let alone the public. Education projects in this area would be valued.

We commend

DMH's commitment to patient involvement and engagement and its open and transparent method of working. We are impressed by the way that they have listened to advice about patient involvement and acted on it.

We have expressed concerns about

The lack of work on public engagement particularly in relation to links between DMH and Google and public perception that data processed by DMH could be shared with Google.

We recommend

That DMH develop the principles for effective involvement laid out by Rosamund Snow, in co-production with citizens, into a set of values for how they work with patients and the public. DMH Talk to Wellcome Trust about their work with patient engagement.

That DMH urgently develop a strategy for public engagement, in partnership with others, as part of a wider conversation about health data and trust.

That DMH consider developing education programmes about AI and its uses in healthcare.

Please go to the following links to view each review **Rosamund Snow PPIE recommendations'** and **Simon Denegri PPIE Review**.

Broader consequences and human factors/ergonomics

A major consideration for the Independent Review Panel has been the potential for broader consequences arising from the use of both Streams in relation to healthcare delivery but also, more broadly, in the longer term, of the use of AI across healthcare. We wanted to know if DMH had explored or considered such consequences and if the Independent Review Panel felt there were additional areas or questions which should be considered by DMH either now or in relation to future work.

Streams

Streams is essentially a means of viewing data, including intelligent alerts, clinical noting and task management. It does not currently involve Artificial Intelligence (AI) but in the future AI driven alerts could be delivered. It has the potential to change the working habits of clinicians and the management of many health conditions, not just Acute Kidney Injury (AKI), It also has the potential for impact on the overall habits and systems of frontline healthcare delivery. We were particularly interested to know how DMH envisaged the development of Streams beyond the Royal Free setting.

DMH is clear that the primary motivation for Streams is improving patient safety. The Royal Free team has developed a set of protocols and a care pathway under the "Think Kidneys" initiative and Streams is modelled on that care pathway. We explored how the team had developed the app and also discussed the implementation process which involved an implementation team of clinicians and developers, behavioural work and the setting up of a dedicated clinical response team to AKI alerts. A great deal of work, from both clinicians and DMH has clearly gone into this.

Response to alerts

We were reassured that the development had included recognition of the potential for 'alert fatigue' and also extensive consideration of what might be expected of clinicians in response to alerts to ensure that the pathway and protocols were followed properly. We asked how locums might be expected to use the application and protocols, given that they would have no prior knowledge of the system. At this stage, it is not clear how a familiarisation process might take place on the frontline given current locum rates. However, this is an existing problem for locums and potentially Streams should make it easier for them to act correctly.

Adoption and diffusion

A key feature of the Royal Free development model is the enthusiasm of those involved. We wondered how Streams might be implemented if it were to be applied to a different location where it would be "parachuted in". In this situation, the co-design element would no longer exist and there would a danger of Streams simply being viewed as yet another "IT intervention". In this scenario, clinicians might not use it in the way intended. The team felt enthusiastically, and probably with justification, that Streams has been co-designed to tackle a specific, important problem that almost all clinicians agree needs to be resolved. It is therefore less likely to suffer the same fate as IT approaches for more general clinical use. The team also felt that being designed to be "beautiful and easy to use" made Streams a far more compelling tool. In other meetings DMH have said that they recognize that all doctors are 'tinkerers', in other words that they will naturally want to tinker with a system to reflect their own priorities in a local setting. The Streams system makes this possible, providing fidelity with flexibility. Nevertheless, we felt this was an area that requires more thought before Streams is rolled out in other Trusts.

The wealth of data provided by Streams could be used to monitor/measure clinician and organisational performance, such as looking at individual response times to alerts. Its data could also be used in the event of an incident, such as patient harm occurring because of a delayed response. We believe it is important to recognize the value of accurate data for investigation and this should be welcome especially where paper records are known to go missing. However, in other industries in the UK such as aviation, factual data is used for investigation and is also admissible in Court. If Streams data were to be admissible it could be a potential barrier for implementation, with Trusts and clinicians being wary of increasing their risk exposure. However, this is not a new situation for the medical profession which extensively documents the responses to changes in patient status and indeed many existing health IT systems time stamp the order and review of tests and investigations.

Broader consequences and human factors/ergonomics

We understand that extensive user research has been undertaken, including discussions with the BMA and with the RCN.

Use for other conditions

The DMH team has been inundated with requests from clinical staff to apply the Streams data viewing approach to other conditions. It was clear that some conditions would be suitable but the team stressed that each and every condition would require the same careful approach as undertaken with the Streams AKI model, i.e. a long process of identifying the clinical evidence based care pathway, potential false triggers and reliability.

More generally there is clearly a large scale of requests for DMH to use their expertise in other areas. We are unclear how DMH prioritises such requests but picking the right ones, may in itself speed implementation. This must include those with the greatest unmet need, those which are tractable, and in particular we would urge DMH to consider those which will better knit healthcare together, with both primary care and mental health care, rather than only looking inside hospitals.

Human factors/ergonomic review

As part of its work, the Independent Reviewers commissioned an ergonomics and human factors review of the DMH Streams AKI application from Hu-Tech, a specialist human factors and ergonomics consultancy. This technical discipline focuses on understanding how we make it easy for humans to do the right things consistently and better. Their review is reproduced in full and contains a number of detailed observations. We note here their principal recommendations that DMH should:

- Broaden the testing protocol to include all end-user demographics, including the least able working in the worst environment.
- Identify how Android users will be accommodated in the rollout of the AKI app.
- Consider increasing the accessibility options of the app interface.
- Consider developing voice input, with reference to protocols for safety critical verbal communications.
- Develop filter and sort functions to aid task management.
- Further develop the functionality of the 'Will you see this person' button.
- Provide a clear hierarchy of responsibilities and sharing of responses in the app.
- Provide alert to warn of loss of signal or reduced functionality.
- Develop a 'smart' system to manage high volumes of alerts and reduce the risk of alert fatigue, with reference to alarm management protocols used in other high-hazard sectors.
- Investigate whether good practice in human error analysis and safety critical system integrity has been included in a systematic manner in the risk analysis.
- Widen stakeholders to include IT front-line staff.
- Broaden the qualitative study on identifying barriers to implementation and usability to include a wider group of stakeholders.

Please go to the following link for the **Hu-Tech ergonomics review** which contains a number of detailed observations.

Of particular note, was that clinicians could not use their screens when wearing surgical gloves. This is not typical use but requires attention. Infection control is a matter for hospitals, not DMH, however rapid uptake is more likely to involve clinician's own devices and there may be an issue here if these cannot be appropriately cleansed.

Unintended consequences of the use of artificial intelligence in healthcare

There are potential ethical and financial dilemmas raised by the application of AI to health conditions. Most are not new and apply to other screening methodologies, for instance, patients receiving unanticipated findings whilst being investigated for a different condition. Another might be the early detection of a dread condition which was untreatable, or detection of conditions at such an early stage that preventive or treatment options were not yet available. A major issue would be the early detection of conditions which created overwhelming demand, which the NHS was not resourced to address or which

Broader consequences and human factors/ergonomics

encroached on funding for other important services. It is clear that the AI itself does not yet create new "unintended consequences", but it could create them quicker or at a scale which would potentially require greater resources to resolve.

These issues do not pertain equally to all three projects but are particularly relevant to the retinal scan project at Moorfields.

DMH have already recognised a number of potential issues:

- Use of AI could provide highly accurate predictors that may drive medical intervention. If diagnosis rates were substantially increased, resources to treat newly diagnosed patients might not be available. This becomes a particular problem if the false positive rate increases, either because more borderline conditions are being detected, or because the parameters being measured clinically, are not totally related to the underlying condition.
- AI 'learning' might develop information on conditions which were not the intended aim of the programme. A retinal scan might for example, also pick up a neurological condition. The World Health Organisation's screening tool identifies principles for screening. These are intended for use while developing a screening tool and might be difficult to apply retrospectively. However, in some senses this is no different to current situations in the NHS, where not everything that is found beyond the initial indicator for screening, is shared with patients but is held anonymously to inform research and future policy.

The DMH team has a clear vision, which they articulate frequently, that they want DMH to be different from other organisations which provide services to healthcare. A key driver was their desire to ensure that anything their expertise is applied to, meets the highest ethical and social purposes. They will turn down development in areas where this is not the case. This is an admirable aspiration. It was recognised by the independent reviewers that the process by which DMH judge requests to develop tools or apply AI might benefit from the development of some clearer principles.

It is evident that there are many potential implications of AI that society is currently unaware of and that even developers have not yet fully realised – the so called 'unknown unknowns'. We would encourage scenario work with clinicians and the wider public to try and tease out what these might be and to generate other perspectives and ideas.

We commend

The way DMH has co-developed Streams with energy and commitment showing a clear awareness of the importance of user interfaces and an appreciation of the challenges being addressed.

We have expressed concerns that

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

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

We recommend that

DMH considers any infection risks and how they might be addressed.

DMH consider with clinical and non-clinical professionals the implications of their work for performance management, for litigation and for assessment of future workforce requirements.

Independent Review Panel governance

The Independent Review Panel was deliberately set up without rules, a model or a modus operandi. These were for the panel to establish for itself. One year on, we need to look back, reflect and consider what has worked and what needs improvement.

The Panel was established at speed and has no comparable models. Without a blueprint, we found ourselves establishing rules very quickly. In addition, we had much to come to grips with, not least familiarising ourselves with an area of technology which is so complex and which is changing so rapidly.

As Mustafa Suleyman has laid out in his introduction, integrity was one of the values sought when seeking independent reviewers for the panel. Trust is clearly essential – DMH would not have given the extensive access that we have been offered, if they do not trust us to behave responsibly. Equally, if we are viewed as uncritical cheerleaders for DMH, we will have failed to keep our independence and will not deserve any trust from the public. We hope we have found the correct balance, and fulfilled our mandate.

Membership

We were appointed personally by Mustafa Suleyman, based on his assessment of our expertise and integrity. None of us had worked with him before, and many had not previously met him. However, this is an ad hoc method of selection and there will need to be thought given to how future members will be chosen. We are open to suggestions about how to attract a wide range of diverse experts, with clear trustworthiness and integrity.

Initially, each panel member had a one year term of appointment but it is already clear that this is too short a timeframe to come to grips with this complex area. We suggest three-year appointment terms, staggered so that current members serve a further one, two or three years. No more than a third should leave in any given year and extensions should be granted if needed to achieve this. No person should have more than one extension.

The current Chair was elected by the group and we suggest that this should continue, with annual election and a maximum term of two consecutive years.

We have insisted that a clear declaration of interests is made by each reviewer.

Payment

Members of the panel are currently unpaid, as a marker of independence. Members do not gain financially from their roles, in fact for some there is a meaningful net cost. There is however a significant and complex workload, which makes it hard for some reviewers, particularly those who do not have an employer prepared to donate their time, to give as much time as required. This is not sustainable.

We explored four potential options:

- DMH pays an honorarium to each member, in exchange for a time commitment. This needs to tie in to fixed terms of office and to more independent appointment. A similar model exists for members of governing boards for Arms-Length Bodies.
- Another body provides funding for panel members. This might be a Foundation interested in transparency in business and technology, who wish to see this model succeed. This is more independent, but some might question why a third party would want to save DMH money.
- DMH makes an agreed contribution to a charity for each member's work. There is a less direct financial incentive for members, but still has a sense of reward for time commitment. However, this might be at the expense of diversity since only those who could afford to give their time could accept an offer to join the panel.
- As now, with consequent effects on motivation, recruitment and time commitment available from reviewers.

We favour the first option, with the option for Reviewers to receive the payment directly or nominate an organisation or charity to receive it on their behalf.

Travel expenses are remunerated and secretarial and commissioned roles are paid from a notional operating budget of £50,000 made available by DMH. This year we spent £59,315 and we appreciate DMH agreeing to increase the budget somewhat, at our request.

Independent Review Panel governance

We are grateful for assurances that more funding would be provided if needed and that this will be an ongoing commitment. We are clear that we should not waste money unnecessarily but equally we need the ability to service our work and to commission and investigate where required.

Please go to the following links to view **Independent Review Panel Declaration of interest**.

Designed by Soapbox, soapbox.co.uk