

medConfidential's current understanding of the public timeline of events regarding the relationship between the Royal Free London NHS Trust and Google DeepMind.

1. On 29 September 2015, a document entitled "Information Sharing Agreement"¹ from the Royal Free London NHS Trust was signed by the "co-founder and head of applied AI" for "Google UK Ltd". This document was subsequently released under a Freedom of Information Act request by the New Scientist.
2. In February 2016, the Royal Free London NHS Trust and DeepMind announced the Streams app at what was described as a public event.²
3. We have found no evidence that historic bulk data was mentioned publicly at this time.
4. In April 2016, the Royal Free London NHS Trust published the minutes of its Trust Board meeting of 24 February 2016.³ The minutes state (emphasis added):
5. "The chief executive reported on a subject which had previously been comprehensively discussed by the board at its confidential meeting in January 2016. The board had agreed to enter into a memorandum of understanding with Google DeepMind to form a **strategic partnership to develop transformational analytics and artificial intelligence healthcare products** building on work currently underway on an acute kidney failure application. This was to be publicly launched at an event later that day."
6. The content of this memorandum of understanding remains secret. We understand the memorandum has been requested several times, and that Freedom of Information Act requests are pending.
7. On Friday 29 April 2016, the New Scientist published the Information Sharing Agreement alongside an article containing questions as to whether all applicable rules had been followed, and as to the justification for the scope of the data.⁴

¹ <https://drive.google.com/file/d/0BwQ4esYYFC04NFVTRW12TTFFRFE/view>

² <https://www.theguardian.com/technology/2016/feb/24/smartphone-apps-google-deepmind-nhs-hospitals>

³ http://s3-eu-west-1.amazonaws.com/files.royalfree.nhs.uk/Trust_board_papers/Trust_board_meeting_PUBLIC_6_April_2016_combined.pdf - '2016/35 Chair and Chief Executive's Report', page 6 of 79.

⁴ <https://www.newscientist.com/article/2086454-revealed-google-ai-has-access-to-huge-haul-of-nhs-patient-data/>

8. Page two of the Information Sharing Agreement contains section 4 “Principles”, and, under the line “Information sharing between organisations must always be consistent with the Caldicott Principles”, lists 6 of the 7 Caldicott Principles.⁵
9. The Agreement requires HL7 feeds of test results. A feed of data so that patients’ current test results may be displayed is uncontroversial, in that this would be for the patient’s direct care.
10. Questions are however raised by further details on page 3 of the Information Sharing Agreement, which lists other categories of information to be shared:
11. “CDS/ SUS submission - via VPN or SFTP
 - a. CPC - completed inpatient episodes
 - b. CC - critical care
 - c. AAE - accident and emergency”⁶
12. “Last 5 years data of all of the above, to aid service evaluation and audit of the new product.”
13. Following the New Scientist article, there was follow up media.
14. The first other article of which we are aware was published on the BBC website entitled, “Google given access to London patient records for research”.⁷ The piece included the line “the BBC understands that no artificial intelligence will be used.”
15. Following this, on 4 May 2016, the story appeared as a front page article in the Daily Mail,⁸ and in multiple other national and international media outlets.
16. At this stage, and given headlines such as on the BBC website, the question was why no research protocols had been followed.
17. Later on May 4th, following the press coverage that morning, BBC Radio 4’s PM programme ran a segment interviewing Prof Hugh Montgomery on the topic.⁹ He said (all emphasis added):

⁵ https://en.wikipedia.org/wiki/Caldicott_Report - the 7th principle was added in 2013.

⁶ Item (d) omitted, as it relates to specification documents, not patient data.

⁷ <http://www.bbc.co.uk/news/technology-36191546>

⁸ <http://www.dailymail.co.uk/news/article-3571433/Google-s-artificial-intelligence-access-private-medical-records-1-6million-NHS-patients-five-years-agreed-data-sharing-deal.html>

⁹ BBC iPlayer will delete the episode within a day or so of this letter being sent. We have placed a copy of the segment at <https://www.dropbox.com/s/cfimojgec24rlrj/20160504-deepmind-radio4-pm.mp3?dl=1>

18. “A lot of harm comes to patients in hospital, and that's partly because **there are increasing numbers of patients with increasingly complex diseases and increasing amounts of data, and the human brain, and the staffing levels, really aren't equipped to deal with that.** One of those is something called Acute Kidney Injury, and it's highly prevalent in hospitals. Of the 100,000 people who die in hospital every year, 50% of them will have had AKI. We think that, if we could address those, around 30% of those deaths could be saved. And so, the NHS put out a call, a safety alert, in 2014, to try and address that and say that it had to be hit.”
19. “What they demanded, was a little computer algorithm got written in to detect cases when blood results came in, to try and alert doctors to the presence of such a patient so we could act on it. But of course, that architecture, for information systems, means **that someone has to go and look for the result, recognise what it means, and think what they might do.**”
20. “So, a bunch of clinicians went to DeepMind to ask for their help, in creating a platform which would provide an immediate alerts to doctors, **the right doctors, in the right way, synthesise the data** so that the best care can be delivered”
21. It was later in this interview that Prof Montgomery confirmed for the first time that this was *not* a “research” project using the Secondary Uses Service dataset, but was rather “direct care”.
22. The Caldicott 2 Report¹⁰ defines direct care as follows: “For the purposes of direct care, relevant personal confidential data should be shared among the registered and regulated health and social care professionals who have a legitimate relationship with the individual”, i.e. direct care is the care of an identified patient by an identified clinical professional.
23. Around the time of the Montgomery interview, the Royal Free published a FAQ on the topic which stated: “The Streams app, which is being developed, uses data to provide diagnostic support and track patient outcomes. Therefore, a range of patient data must be analysed.”
24. “Historical data is used to analyse trends and detect historical tests and diagnoses that may affect patient care.”
25. After the PM programme, the technology outlet TechCrunch received further details.
¹¹ The TechCrunch piece included the line, “The Royal Free spokesman said it is not possible, under the current data-sharing agreement between the trust and DeepMind, for the company to apply AI technology to these data-sets and data streams.”

¹⁰ <https://www.gov.uk/government/publications/the-information-governance-review>

¹¹ <http://techcrunch.com/2016/05/04/concerns-raised-over-broad-scope-of-deepmind-nhs-health-data-sharing-deal/>

26. And further quoted the RFH spokesman as saying ““The only thing this data is for is direct patient care,” he added. “It is not being used for research, or anything like that.”
27. In light of this new information – which may explain why no research approval processes were followed – the question became one as to the purpose of all of the data being for direct care, and the justification of the volume of data. This issue might most easily be crystallised by asking, “What is the ‘direct care’ purpose for the release of data about patients who have never had a blood test when they visited A&E, and who will not return to the Royal Free?”
28. On Friday 6 May, the Guardian published an article entitled, “DeepMind has best privacy infrastructure for handling NHS data, says co-founder”.¹² This was somewhat strange, as at no point had there been a dispute as to whether DeepMind could keep the data it had securely - DeepMind is part of Google, with Google infrastructure. The article appeared more of a smokescreen for the real issue, i.e. what was the basis for the data to which DeepMind had access.
29. The May 6th Guardian piece also includes a quote from Mustafa Suleyman, stating: “When we developed our information governance toolkit and we submitted that for assessment to the health and social care information centre (HSCIC) ... we got 100% for our toolkit. There’s pretty much nobody else who’s been able to get a score as high as that.”
30. The ellipses in the quote above refer to a footnote: “This article was amended on 9 May 2016 to remove an incorrect statement in a quote that the HSCIC approves data-sharing agreements.”
31. The version of the IG toolkit in force at the time was v13, covering 2015 - 2016.¹³ As DeepMind personnel should have been aware when they filled in the toolkit,¹⁴ the “100%” is a self-reported score and requires an adequate understanding and careful consideration of issues of Information Governance, especially relating to the data being held.
32. An article published in Business Insider, also quoting Mustafa Suleyman, says: “The data is anonymised, meaning DeepMind staff shouldn’t be able to see whose data

¹² <https://www.theguardian.com/technology/2016/may/06/deepmind-best-privacy-infrastructure-handling-nhs-data-says-co-founder>

¹³ v14 is now in use

¹⁴ The submission is v13 visible here: <https://www.igt.hscic.gov.uk/AssessmentReportCriteria.aspx?tk=425159990792031&Inv=3&cb=840fe512-74ab-4643-89d9-2919133ff3eb&sViewOrgId=49979&sDesc=8JE14>

they are looking at.”¹⁵ The level of anonymisation is self-evidence from the image of streams on the Deepmind website, which includes name, NHS number, and hospital number, event dates.¹⁶ The full list of information considered identifying by the NHS is found in Appendix 7 to the 1997 Caldicott Review.¹⁷

33. A week later, on 13 May 2016, the New Scientist published a second article, entitled “Did Google’s NHS patient data deal need ethical approval?”,¹⁸ to which Google / DeepMind objected strongly. This piece was supplemented with information only provided to the New Scientist after publication.
34. This piece included the line, ““We and our partners at the Royal Free are in touch with MHRA regarding our development work,” Google said in a statement to *New Scientist* on 13 May. The app is currently offline.”
35. On 18 May 2016, TechCrunch posted a second long article,¹⁹ a more in-depth follow up of the issues, with more than an exchange of canned quotes. As such, the following paragraphs quote it at length:
36. MHRA are quoted as saying: “We have been in contact with Google since May 4 and are currently in discussions with them about whether or not their app needs to be registered as a device,”
37. “However both DeepMind and the Royal Free assert there was no requirement for them to gain prior approval to develop and pilot the app because they say they have not conducted any “clinical trials/investigations”.
38. “A Royal Free spokesman rather says they have carried out small “user tests” of the app.
39. “We’ve asked the MHRA at what point a pilot of a product would be considered to constitute a clinical trial or investigation in its view and will update this post with any response.
40. “The MHRA’s normal procedures mean it can issue a letter of ‘No Objection’ after reviewing an application to run a clinical investigation into a medical device –

¹⁵ <http://uk.businessinsider.com/deepminds-cofounder-has-defended-the-data-sharing-agreement-with-the-nhs-2016-5>

¹⁶ <https://deepmind.com/css/images/bg-health-streams.png>

¹⁷ Page 89 http://webarchive.nationalarchives.gov.uk/20130107105354/http://www.dh.gov.uk/prod_consum_dh/groups/dh_digitalassets/@dh/@en/documents/digitalasset/dh_4068404.pdf

¹⁸ <https://www.newscientist.com/article/2088056-did-google-need-ethical-approval-for-nhs-patient-data-deal/>

¹⁹ <http://techcrunch.com/2016/05/18/uk-healthcare-products-regulator-in-talks-with-googleddeepmind-over-its-streams-app/>

assuming it does not have any concerns about the proposal. The Streams app has not yet gone through this review process.”...

41. “DeepMind asserts that access to all patient data across the three hospitals is necessary for the app’s predictive function to work. It also claims it is not engaged in research, and says the Streams app is being used for direct patient care – an important distinction because additional regulatory and ethical approvals would likely be necessary if the Google-owned company was performing research on the data-set. Or applying any machine learning algorithms to the data, which it says it is not.” ...
42. “Our [RFH] agreement with DeepMind is our standard third-party data sharing agreement, with the trust being the data controller and DeepMind being the data processor.” ...
43. “All the identifiable data under this agreement can only ever be used to assist clinicians with direct patient care and can never be used for research. We and our partners at the Royal Free are in touch with MHRA regarding our development work.”
44. “So what then is meant by “development work”? The Royal Free spokesman told TechCrunch that in total three small “user tests” of Streams have been run so far, with each lasting between two and six days, and with a maximum of six clinicians using the app during each test.
45. “The spokesman declined to specify how many patients have been involved in the tests – although given that all three hospitals’ patients data is being fed into the algorithm powering the app then, in theory, all current and past patients (extending back five years) of the hospitals are in some sense ‘involved’ in these tests because their data is being used by the app. In all likelihood the vast majority of these people will be unaware their data is being used for this purpose.
46. “It is also not clear what criteria DeepMind/the Royal Free are using to evaluate their “user tests” of the Streams app. Nor which outside body – if any – is reviewing the tests.”
47. “Another interesting question here is what exactly is the role of DeepMind in the project? The design of the app was at least partially outsourced (described as ‘co-designed by’) to London based app design studio ustwo, while the [algorithm being used to process patients’ data](#) was, we are told, developed by the NHS. So why is a company famed for its artificial intelligence algorithms being engaged to act as, effectively, a project manager for a healthcare app?

48. Later on 18 May 2016,²⁰ medConfidential posted a second news update on the topic, mostly repeating earlier questions and placing other ad hoc comments in a single location.²¹ It included our work to infer the origins of an internally contradictory information sharing agreement, which suggested that it had been 'cut and pasted' together from multiple other documents.
49. This timeline document was written on 31 May 2016. Since the above articles were published, several quotes attributed to DeepMind have been put to us, which have not yet been published. These include:
50. "Thus far, we have conducted 3 small scale tests, each lasting from 2-6 days. Throughout these tests, Streams has always served as an additional information system for clinicians, tested alongside rather than replacing standard hospital safety systems."
51. "If we don't have the data then we can't predict who might get it."
52. "With any clinical data processing platform it is quite normal to have data lying in storage and it is nonsense to suggest that these platforms should only hold the data of those patients being treated at that very moment."
53. "We sought all the approvals we believed to be necessary."

Hark

54. All of the above points, and all public concern, has been in reference to the Streams app. The coverage of the February 2016 event also included a second app, developed by Imperial College London, called Hark. Writing about Hark, The Guardian said (emphasis added):²²
55. "**The app will pick up relevant information from the electronic patient record**, so that medical staff have a complete picture of the patient's condition and medical history."
56. In April 2016, it was announced Hark had been purchased by DeepMind.²³

²⁰ Some of this information was later merged into the piece dated 13th May from the New Scientist, as Google answered questions only after the New Scientist piece had been published.

²¹ <https://medconfidential.org/2016/update-on-google-deepminds-nhs-app-is-it-just-resting/>

²² <https://www.theguardian.com/technology/2016/feb/24/smartphone-apps-google-deepmind-nhs-hospitals>

²³ <http://www.imperialinnovations.co.uk/news-centre/news/imperial-spinout-hark-health-solutions-sold-google/>

–ends–

medConfidential

coordinator@medConfidential.org

31 May 2016