



[coordinator@medconfidential.org](mailto:coordinator@medconfidential.org)

10<sup>th</sup> June 2017

## Re: Google DeepMind and the Royal Free

Dear National Data Guardian,

Following the publication of your letter to the Royal Free and Google DeepMind by Sky News on 15<sup>th</sup> May, a number of statements put to journalists and others by both of these organisations appear to cast previous assurances into doubt.

Specifically, and with regard to medConfidential's complaint of 1 June 2016 about Google DeepMind, the Royal Free London NHS Trust, and the copying of data on 5 years' of patients' hospital events,<sup>1</sup> public statements at that time asserted this copying was being done for direct care purposes only.

Despite telling the press in 2016, "We have, and will always, hold ourselves to the highest possible standards of patient data protection",<sup>2</sup> statements by DeepMind in recent weeks make it clear the company either misled the public in 2016, is misleading the public now – or both.

Such duplicity suggests the usual reduction of Monetary Penalties by the ICO – often used as an encouragement for prompt rectification – cannot be available to DeepMind, because the company has refused to acknowledge the problem, as evidenced by its officers' public statements in the timeline below.

To further quote the press debate from last summer: "Explaining why they did not speak to patients first the Royal Free said: 'Health professionals may rely on implied consent when sharing personal data in the interests of direct care.'"<sup>3</sup>

medConfidential's complaint, however, was around the legal basis for the copying of the data of patients who *had never had* a blood test at the hospital, and/or who *will never return* to it, whose data Google DeepMind copied – and, according to the company's own statements, continues to copy.

---

<sup>1</sup> HES / SUS.

<sup>2</sup>

<http://www.dailymail.co.uk/news/article-3573286/NHS-trust-handed-private-patient-details-Google-says-implicit-permission-emerges-hospital-talks-internet-giant.html>

<sup>3</sup> Ibid.

Some patients whose data is displayed in the Streams app may quite possibly be receiving direct care, but this is not about them.

Given the Streams application is now in active use, have the ICO and NDG confirmed that their understanding of the Streams app's operation is identical to that which MHRA understands it has approved – including all of the data the app shows, and consequently, how that data is sourced?

Public statements by the Royal Free and briefings by Google DeepMind suggest there are discrepancies.

## **An alternate legal basis?**

Following the Sky News report, Google DeepMind has been briefing journalists and suggesting in private that it has an “alternate legal basis” for data processing, directly contradicting its public statements.

While an “alternate legal basis” has not (yet) been published, DeepMind's statements present a limited set of options:

1. Is there is a *third* set of paperwork<sup>4</sup>? If there is, DeepMind's claim to have published all of its contracts with hospitals – an assertion currently made on its website<sup>5</sup> – would be false and misleading.
2. If the November 2016 contract remains in force, then data is being copied by Google DeepMind from the Royal Free outside of the scope of the contract, i.e. there is no valid data sharing contract in place for everything that is being done.
3. The “alternate legal basis” is a PR fiction, and DeepMind continues to receive a regular data file of all patients from the hospital, irrespective of whether Streams is involved in their direct care – i.e. the National Data Guardian's assessment is as applicable to the second (Nov 2016) contract as it was to the first.

If there is a fourth option – perhaps deriving from DeepMind's belief that it is somehow ‘made of magic’<sup>6</sup> and that the law does not apply to them – then we may well have missed it. However, medConfidential's work with the NHS in past years has uncovered a whole series of egos believing they are above the rules; we expect this to work out no better here than elsewhere.

---

<sup>4</sup> The first set being the original two documents exposed by the New Scientist in April 2016; the second being the November 2016 information processing and services agreements; and then this possible third ‘secret deal’, alluded to by DeepMind...

<sup>5</sup> See ‘Our agreements with NHS partners’, <https://deepmind.com/applied/deepmind-health/transparency-independent-reviewers/>

<sup>6</sup> <https://www.theguardian.com/technology/2017/may/23/alphago-google-ai-beats-ke-jie-china-go>

The most obvious cultural and regulatory analogue to this would be Uber's recent attempt to 'legalise' running red lights in San Francisco.<sup>7</sup> While Uber's approach would have put people in hospital, Google DeepMind wants their data once they get there.

Irrespective of which turns out to be the case in this instance, is it appropriate for any part of the NHS to surprise patients with unclear data usage, whether lawful or otherwise?

### **Transparency vs secrecy?**

Are DeepMind misleading the public, are they misleading the regulators, and/or are they misleading themselves?

An ethics approval for research is not a lawful basis for data reuse – though it does undermine statements DeepMind has made to others about *not* using machine learning to improve the prediction of acute kidney injury and general patient deterioration.<sup>8</sup> medConfidential has not seen the undertakings that were given as part of that ethics approval process, but the public summary suggests the strong possibility of material falsehoods.

DeepMind claims, "We've been taking steps to become the most transparent company working in NHS IT."<sup>9</sup> As part of this purported 'commitment to transparency', DeepMind provides links to copies of its Information Processing Agreements and Service Agreements with a number of hospitals, including the Royal Free London NHS Foundation Trust.

While on the face of it, such publication would appear to meet DeepMind's claimed commitment to transparency, the precise content of these documents bears further scrutiny.

Given the controversy over what data is copied from the Royal Free, the single key page in the published Information Processing Agreement<sup>10</sup> – Schedule 1, on page 12, which lists the data to be copied under the new contract – is missing. Page 12 has simply been removed from the published document, not redacted, as has been done for, e.g. Schedule 3.

DeepMind asserts it is not copying data in bulk under this contract. However, item (d) of Schedule 1 shows this public statement to be false; DeepMind's so-called "transparency" is insufficient for the public to assess the integrity of its statements.

As part of your concluded investigation, can you cast any light on which of Google Deepmind's claims are true, and which are contradicted by stated facts?

---

<sup>7</sup> <https://www.wired.com/2016/12/ubers-self-driving-car-ran-red-light-san-francisco/>

<sup>8</sup> <http://www.hra.nhs.uk/news/research-summaries/using-machine-learning-to-improve-prediction-of-aki-deterioration/>

<sup>9</sup>

<https://deepmind.com/applied/deepmind-health/transparency-independent-reviewers/what-weve-learned-so-far/>

<sup>10</sup> <https://deepmind.com/documents/23/REDACTED%20-%20FULLY%20EXECUTED%20DeepMind%20RFL%20Information%20Processing%20Agreement.pdf>

## Updating our time line

Assuming the dates in the letter from the National Data Guardian are correct, in the period following our complaint, the time line seems to have been:

1. 2016-07-01<sup>11</sup>: Google DeepMind provides more information to the NDG, and explicitly confirms that data was copied for the purposes of direct care; **seemingly also confirming their only legal basis was implied consent.**
2. 2016-11-22: DeepMind and RFH rewrite their contract.<sup>12</sup> **The press release announcing this explicitly states that the data copied is unchanged.**<sup>13</sup>
3. 2016-12-07: Parliament is told that SUS is “maintained for purposes other than direct care”, which engages questions of dissent, not implied consent.<sup>14</sup>
4. 2016-12-16: **NDG writes to Google DeepMind to notify them that their use “cannot be regarded as direct care”**
5. 2017-01-20: DeepMind contradict their public statements in writing to the NDG panel.
6. 2017-01-24: DeepMind seemingly contradict their public statements in person to the NDG panel.
7. 2017-02-15: According to the NDG letter, on this date, DeepMind seemingly reconfirm that direct care was the sole legal basis for copying the data.
8. 2017-02-20: **NDG reconfirms the finding from 16 December 2016: that the development use “cannot be regarded as direct care”.**
9. 2017-02-27: **DeepMind and RFH issue a press release lauding Streams’ use in direct care**, citing “more than 26” patients helped.<sup>15</sup> The press release does not mention the NDG finding of one week earlier.<sup>16</sup>
10. 2017-05-15: RFH tell Sky News that the app is in routine use in the hospital, displaying patient information for direct care.<sup>17</sup>

---

<sup>11</sup> We refer to dates in a YYYY-MM-DD format, to assist clarity.

<sup>12</sup> <https://www.newscientist.com/article/2113701-googles-deepmind-agrees-new-deal-to-share-nhs-patient-data/>

<sup>13</sup> medConfidential press release, 16/5/17, Notes to editors, 5(d):<https://medconfidential.org/2017/press-release-google-deepmind-unlawfully-copied-the-medical-records-of-1-6-million-nhs-patients/>

<sup>14</sup>

<http://www.parliament.uk/business/publications/written-questions-answers-statements/written-question/Lords/2016-12-07/HL3943>

<sup>15</sup> <https://www.royalfree.nhs.uk/news-media/news/new-app-helping-to-improve-patient-care/>

<sup>16</sup> <https://medconfidential.org/2017/technology-company-deepmind/>

<sup>17</sup>

<http://news.sky.com/story/google-received-16-million-nhs-patients-data-on-an-inappropriate-legal-basis-10879142>

## Intent

One question on which we have repeatedly been asked to give a view by observing parties, is whether we felt DeepMind committed active malfeasance or whether its actions were simply ‘uninformed ineptitude’.<sup>18</sup>

In light of events during and following November 2016, we would now modify our answer. Initially, their actions were ineptitude borne from ignorance – DeepMind institutionally didn’t understand what it had been given, and also didn’t understand what it didn’t understand. Since the press coverage in November 2016, the academic paper of March 2017, and the media coverage in May 2017, that naïvety has been replaced with something more deliberate.

For example, following the Sky News piece, DeepMind’s press team told journalists that DeepMind had not seen the two page letter from the National Data Guardian, while those journalists were already looking at the third page which proved that statement entirely false.

There is now a strategic choice for DeepMind to do what happens in every other commercial or political field: spin, and issue ‘lines to take’ that are defended with the level of press office resources that a subsidiary of the second largest company in the world<sup>19</sup> will inevitably choose to deploy to protect its reputation, irrespective of truth – or, tell the truth.

DeepMind says it is committed to transparency. Yet, in the most controversial data project it has, with the highest sensitivities, in the middle of a 14-page PDF – a PDF which has other information redacted for entirely appropriate reasons – the one page (p12) that has been removed is the Schedule that covers the actual data that is copied. DeepMind has been told this page was missing; it promised the page would be restored. It wasn’t.

medConfidential does not at this point believe that DeepMind is hiding intentional malfeasance; we have seen no evidence of that. What we do believe is that this is a large organisation – like any other – hiding its own embarrassment. But that it is hiding anything at all is a demonstrable breach of the standards it tells the public that it has honoured, and will always honour. It is unfortunately the case that the checks and balances of medical ethics do not, and never will, constrain the commercial or public relations worlds.

Item (d) of Schedule 1 from the November 2016 Information Processing Agreement<sup>20</sup> shows DeepMind’s statements to journalists and to the public, and its officers’ assertions of difference between the two contracts, to be materially false. But, until your inquiry concludes, the public cannot know that because – despite its protestations to the contrary – Google DeepMind is hiding the evidence.

---

<sup>18</sup> We paraphrase. The question as asked in one conversation was, “Are we talking malice or muppetry?”

<sup>19</sup> [https://en.wikipedia.org/wiki/List\\_of\\_public\\_corporations\\_by\\_market\\_capitalization#2017](https://en.wikipedia.org/wiki/List_of_public_corporations_by_market_capitalization#2017)

<sup>20</sup> <https://deepmind.com/documents/23/REDACTED%20-%20FULLY%20EXECUTED%20DeepMind%20RFL%20Information%20Processing%20Agreement.pdf>

## Development and Innovation: AI across the NHS and beyond

Any argument that determining the data transfer in this project to be unlawful will undermine lawful app development are self-evident nonsense. Whether this turns out to be ignorant nonsense or wilful dishonesty will depend on the substance of that claim.

Data-driven development and iteration is entirely possible, and necessary – especially in direct care.

For example, looking at the same company, medConfidential currently has no major concerns with the approach taken between DeepMind and Imperial, DeepMind and UCLH, or DeepMind and Moorfields.<sup>21</sup> We also doubt that the only innovation at the Royal Free is with DeepMind.

One failed project, appropriately accounted for, should also not affect the vast majority of high quality research and innovation in the NHS. Following proper process is not the problem here.

Using an application in direct care with real patient data – and real patients – is something that should take place towards the end of a process, not the beginning. The approach taken by DeepMind in this instance is akin to a drug company synthesising a new drug, then trialling it directly on patients who are supposed to consider themselves fortunate to be in the hospital at a time convenient to the manufacturer.

In the clinical trial world, leaping straight from simulation to phase IV trials is evidently abuse of process – possibly profitable, certainly reckless.

And while apps are not mind-altering chemicals, there should be a process for prototyping applications (using synthetic data), testing them with clinicians outside a care setting (with approvals, and a process that moves towards real data), and then inside a care setting.

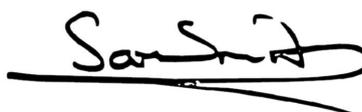
The good news is: that process already exists.

The question this whole sorry episode exemplifies is what happens when those processes are ignored.

Yours sincerely,



Phil Booth, medConfidential



Sam Smith, medConfidential

Cc: ICO, DH.

---

<sup>21</sup> This also covers our understanding of the still-secret Manchester project which is apparently targeting sepsis.