



Response to DeepMind

Julia Powles^{1,2,3} · Hal Hodson⁴

Published online: 26 April 2018

© IUPESM and Springer-Verlag GmbH Germany, part of Springer Nature 2018

We thank DeepMind’s staff and advisors for responding to our article, *Google DeepMind and Healthcare in an Age of Algorithms*,¹ on the public record. The response criticizes our work, and defends DeepMind’s. We are grateful to the journal’s editorial board for inviting us to reply.

A point at the outset. Our investigation and criticism of DeepMind is not motivated by any desire to see public health institutions bereft of advanced digital services. Quite the opposite. We are technology optimists, and truly believe that new technology can help humans solve real problems. We share the collective goal of saving lives and ensuring the right care reaches patients. At the same time, we are concerned that if long held principles are ignored or discarded in the process, the technology’s promise will turn sour.

There is a pressing need for open, fact-based, public discussion of technology companies entering into the provision of public services. The case study at the heart of our article is a particularly illuminating example. Google DeepMind’s first deal with the British National Health Service, involving a gift of at least 1.6 million detailed, non-anonymized health records from every patient in London’s Royal Free hospital trust in order for DeepMind to offer a smartphone app for kidney injury alerts, has been and remains troubling. We think this assessment extends to DeepMind’s response.

We must give a number of clarifications. DeepMind’s letter addresses arguments that we never make, and its continual assertions of “factual inaccuracies” and “unevidenced statements” are nothing of the kind. For example, readers will

find no claims in our paper that DeepMind engaged in “illicit and illegal data mining,” “secret” use of patient data, “exploitation for other purposes,” or “misuses of data.” There is no need for us to imply such wrongdoing when there are more immediate and demonstrable concerns. DeepMind’s response diverts attention from the real challenges we raise: (1) the continuing absence of a valid legal basis for processing every Royal Free patient’s data from November 2015 to at least January 2017, when deployment of clinical app Streams commenced, offering direct care to some proportion of patients being monitored for kidney injury; (2) the total inadequacy of contractual and institutional protections against the *possibility* of misuse; and (3) broader issues about value and power.

Similarly, we have made no assertions that DeepMind’s parent company Google/Alphabet will access, or is accessing, British patient data. Rather, we exposed the absence of contractual precautions against that *possibility*, as well as the transparency paradox that precludes independent voices such as ours from scrutinizing corporate data arrangements, particularly those in which Google/Alphabet has ultimate control.² A complete explanation of the relationship between DeepMind and its parent—including the nature and limits of flows of intellectual property, data, algorithms, and finances—as well as sister companies, must be forthcoming if DeepMind is to continue providing services to the public sector, in health, energy, or any other domain.

In its letter, DeepMind seeks to diminish the value of our work for its focus on a “now-superseded 2015 agreement,” and by suggesting that a July 2017 ruling by Britain’s top regulator, the Information Commissioner’s Office (ICO), in part “came to completely different conclusions.” We expressly cabined our research to focus only on the period July 2015 to October 2016, precisely to combat the asynchronicity between privatized technological progress and public mechanisms of redress. The premature revocation of the 2015 agreement halfway through its term can only be read in our favor, given it involved a realization by the parties that they were collaborating on an untenable, possibly unlawful, foundation. As for

¹ Powles J. & Hodson H. Health Technol. (2017) 7:351. <https://doi.org/10.1007/s12553-017-0179-1>

✉ Julia Powles
julia.powles@nyu.edu

¹ New York University School of Law, New York, NY, United States

² Cornell Tech, New York, NY, USA

³ University of Cambridge, Cambridge, UK

⁴ The Economist Newspaper, London, UK

² DeepMind states that “data access can be verified through audit,” but it has rebuffed numerous attempts by us to do precisely this and provide us with results, as has Royal Free.

the ICO ruling, there is only one point of difference and it is inconclusive. We questioned the self-designation of DeepMind as a mere “data processor” and said it was arguably a joint “data controller,” offering evidence and reasons in support. Similar arguments have been made by many data protection professionals, including the Chair of the National Association of Data Protection and Freedom of Information Officers.³ The National Data Guardian has also shed doubt on the arrangement, stating “the contract appears to contain elements of a data sharing agreement (i.e. Data Controller to Data Controller), and therefore does not provide all of the necessary controls for the sharing of this data.”⁴ The ICO, by contrast, simply took as an accepted fact that DeepMind’s self-designation was true. It made no positive finding and provided no evidence to support this classification, which remains in tension with the definition of processors/controllers under UK law.

We appreciate that neither of us are clinicians. Nor are the DeepMind letter’s authors legal or policy experts. As affirmed in countless reports on the subject, it is vital that a diverse range of voices are heard and heeded as new digital technologies are introduced to the NHS, by DeepMind and others. Indeed, since publishing our article, numerous clinical professionals have endorsed our work and expressed their shared concerns. They believe, like us, that it is perfectly possible to both respect individual rights and to enable promising technologies. Despite DeepMind’s continued assertions that it wishes to be one of the most open digital companies, it has continuously prioritized internal views and those of self-selected advisors, while seeking to discredit external opinions such as our own. We hope this will change.

Readers will find that DeepMind’s response continually returns to the touchstone of clinical will. We accept that Royal Free clinicians are well placed to understand what services will benefit their patients. But there is a regulatory regime in play, and it does not justify any scale of data transfer, to any third party, merely because such transfer is desired by clinicians—the vaguest possible term for hospital staff. Patients, and patient rights, *must* be respected. This is the overriding message of our article and of the ICO ruling. An activity does not become direct care simply because clinicians say so. Similarly, one is not a data processor by intent alone. Such questions are crucial matters of substance and circumstance. We explore them as such in our paper, and are disappointed that they have received so little examination in return.

A final clarification is that, despite the length and tone of [DeepMind’s letter and appendix](#), with its formidable looking 47 points of disagreement, the arguments in fact address only a very small percentage of our article. DeepMind’s points involve

considerable repetition, internal inconsistencies, and consistently ignore context in favor of oblique semantic arguments. We disentangle these issues, point-by-point, in the [Appendix 1](#), seeking where possible to elevate the conversation constructively.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Appendix 1—Authors’ point-by-point response

A preliminary note from Powles & Hodson:

DeepMind’s complaint is a very long document. For readers not steeped in the issues, its sheer volume and highly-credentialed backing from DeepMind’s employees and retained advisors may give the suggestion of a weighty critique. However, as we seek to demonstrate in our point-by-point responses, the vast majority of what is asserted collapses on examination, or otherwise lacks substance.

The journal’s decision to publish this detailed appendix has value for one overriding reason: it is public and transparent. DeepMind’s response is an exemplar of a particular mode of engagement to which we have both become accustomed in investigating technology companies, and which predictably follows whenever a critical independent article receives any traction. No new facts are presented—instead, accusations of inaccuracy are used to reassert and recast differences of *interpretation*, which are then escalated in the strongest terms to the highest editorial tier. This practice, though familiar, is usually invisible. We appreciate the forbearance of readers in assessing it for themselves.

DeepMind’s assertions of “factual inaccuracies and un evidenced statements”:

DeepMind claim that “the article contains numerous errors of fact identified [in enumerated points 1-37] below”:

1. Section 2, para 1: ‘**about developing software using patient data from the Trust**’. This is factually incorrect as patient data were not used to develop the application. As per the journal’s ‘Instructions for Authors’ and guidance from Springer’s ‘Author Academy’ every statement of fact or description requires a supporting reference, and this is not provided.

Powles & Hodson: We are not making a claim about how patient data “were used.” The context of the lifted snippet is a description of Royal Free’s initial approach to DeepMind. We simply record a salient fact that DeepMind has repeatedly tried to obscure: that DeepMind’s project to develop software following a July 2015 approach by Royal Free (support:

³ Lomas N. Patient data API pivotal to DeepMind’s push into UK’s NHS. TechCrunch. 22 Nov 2016. <http://tcm.ch/2fD5v4v>.

⁴ Letter from National Data Guardian to Royal Free, 12 Oct 2016, <https://www.gov.uk/government/publications/request-for-information-about-royal-frees-work-with-deepmind> Accessed 1 Aug 2017.

original ref. [1]) has been underpinned by a large-scale data transfer of every Royal Free patient’s data (support: original ref. [15]), in a way that deserves sustained investigation.

The sentence does not claim that DeepMind used Royal Free’s data to undertake technical development of the Streams application. In fact, in footnote 5, we cite DeepMind’s statement that it relies on “synthetic data” for this particular purpose. In our ongoing work, we are seeking to clarify the source, structure, and nature of this synthetic data through freedom of information requests.

Spoiler alert: This first criticism is emblematic of DeepMind’s response throughout the pages that follow. Our quote is divorced from its context and used to critique the paper for something that it is not saying, while diverting attention from what it actually does say. The response only turns the question back to DeepMind: if (a) patient data was not used to develop the software; (b) only a tiny sample of data was used for testing; and (c) the software has not been deployed for the Trust population at large—and certainly wasn’t between November 2015 and January 2017—what was the data transferred for? And how is its possession retrospectively justified? These questions remain unanswered at the date of this response.

2. Section 2.1 para 1: **‘DeepMind’s publicly announced purposes for holding sensitive data on Royal Free’s patients, i.e. the management and direct care of AKI, were narrower than the purposes that contractually constrained its use of the data.’** In our initial launch of DeepMind Health (Author’s Reference: 5) we stated a broad interest in patient deterioration, the prevention of avoidable harm and the benefits of technology in patient care (fn1). DeepMind has always been clear that the Streams application would be capable of providing benefits beyond acute kidney injury alerting. Indeed coverage of the launch cited by the the author in Reference: 9 states that: *“Suleyman said the company hoped to work on alerts for other life-threatening conditions too, such as sepsis – or blood poisoning.”* - Guardian Feb 2016.

Powles & Hodson: “Purpose” is a specific data protection term, and is not the same as “interest” or “hope.” When asked about the purposes for processing patient data, DeepMind has consistently stated that the detection and prevention of acute kidney injury (AKI) was the purpose justifying the holding of patient data (see article Section 2.2, para 3, particularly original refs [15] and [16]). This is also an official finding of the UK Information Commissioner’s Office (ruling of 3 July 2017, Section 2.0, para 3).

Both this and DeepMind’s previous point attempt to deflect attention from the central research question motivating the paper: how and why the patient data of every single Royal Free patient have been sitting on Google/DeepMind servers since November 2015.

3. Section 2.1, para 2: **‘The ISA was superseded, prematurely, by a new set of agreements signed on 10 November 2016. Those agreements are beyond the scope of the present article and will be considered in future work’** The article relies on an agreement terminated 18 weeks prior to the publication of this article. The Information Sharing Agreement entered into in September 2015 (the “2015 ISA”) was not terminated prematurely, but, rather, was replaced (on 10th November 2016) by the 2016 Services Agreement and Information Processing Agreement with both parties taking the opportunity to introduce additional compliance controls into these documents, including many that exceeded what was legally required for such an agreement at the time.

Powles & Hodson: The original Information Sharing Agreement (ISA) would have run for two years. It was replaced after one. This is patently a premature termination.

The article was submitted on 14 November 2016, four days after the ISA’s termination, precisely to respond to a revisionist narrative emerging about the legality of the initial data transfer. The journal’s publication timeline post-submission was outside the authors’ control, though was relatively efficient from a general publishing perspective.

While the post-November 2016 agreements are interesting and seek to respond to some of the compliance inadequacies of the ISA, they do not negate the importance of assessing the first ISA and the data transfer behind it, as we have carefully done. We do not rely on the ISA—we examine it.

4. Section 2.1, para 3: **‘The reality is that the exact nature and extent of Google’s interests in NHS patient data remain ambiguous.’** This is factually incorrect. DeepMind have unambiguously and publicly explained the utilisation of data (and its justification) for direct patient care with the application Streams. Examples cited by the article itself include public presentations by DeepMind representatives (References: 40, 41, 42), the DeepMind website (43), a blog on Medium (12) and the Royal Free Website (17).

Powles & Hodson: Our quote is clearly talking about Google, not DeepMind.

Section 2.1, para 3 specifically commenced with questions about Google/Alphabet’s interest in NHS data, not DeepMind’s. We asked how, legally, Google’s interests were represented. Our paragraph acknowledges DeepMind’s public relations statements: “DeepMind has made regular public assurances that Royal Free data will never be linked or associated with Google accounts, products or services.” It is the lack

of legal clarity around Google/Alphabet's control and interests that is ambiguous, not DeepMind's public communications.

It is striking that in its letter, DeepMind terms the prospect of Google having access to Royal Free patient data "frightening for the public." DeepMind could best respond to and mitigate this fear by including the following express legal guarantee in all current and future contracts, instead of retaining it simply as a placating public relations statement: "No patient data were, are or will ever be connected to Google accounts or services." DeepMind should also give a full and auditable account of the current and future arrangements for flow of intellectual property, data, algorithms, and finances between DeepMind and its parent company, as well as to "other Bets" within the Google/Alphabet corporate group. Through freedom of information requests, we have discovered that DeepMind's former strategy lead Will Cavendish (an ex-government official under a prohibition against lobbying the government) wrote to the Medicines and Healthcare products Regulatory Agency (MHRA) a number of times in October and November 2016 concerning another Google/Alphabet subsidiary, Verily, thus bringing clearly into question the fluid boundaries between different parts of the corporate group.

5. Section 2.2, para 2: **'For patients who had the necessary precursor renal blood test and were then progressed to being monitored by clinicians for AKI, the appropriate direct care relationship would exist to justify this data processing, through the vehicle of implied consent. However, the dataset transferred to DeepMind extended much more broadly than this. In fact, it included every patient admission, discharge and transfer within constituent hospitals of Royal Free over a more than five-year period (dating back to 2010). For all the people in the dataset who are never monitored for AKI, or who have visited the hospital in the past, ended their episode of care and not returned, consent (explicit or implied) and notice were lacking.'** This is factually incorrect. The dataset processed by DeepMind under the direction of the Royal Free London NHS Foundation Trust (RFL) extended only to the data deemed clinically necessary by RFL, the data controllers. The authors assume that if the patient has had no prior blood test or had no guarantee that they would require future treatment in the hospital, there is no value in holding their data for future use by clinicians. However, healthcare organisations cannot anticipate which patients will be readmitted or need monitoring, nor is this existing practice for clinical software applications. Nor can they determine, in advance, what clinical data may be relevant to any new presentation.

Streams holds, on behalf of RFL, securely and under strict clinician-only access controls (as per all other patient management systems), historical results and data on the patient's past medical history for future use should the need

arise. This is standard in all electronic patient management systems; that is, such systems store data on past encounters for potential future use by clinicians. If the clinical need does not arise the data are not used. The alternative - to delete all data of past healthcare encounters (lest future access be unnecessary) - would necessitate hospitals deleting all historical records of their patients and would have adverse consequences for patient management. If a patient is admitted and has kidney function tests (one of the commonest routine tests) it is necessary for the system to access historical data and, if the blood test is abnormal, notify clinicians. This can happen equally to patients with and without prior kidney tests results as AKI can be detected from sequential blood tests following such an admission. Here the context of background historical information is just as relevant. In essence, the position is no different to a GP surgery (using a commercial software application hosted by a third party) retaining medical records of all of its registered patients, not just those seeking treatment at a particular point in time.

Powles & Hodson: There is no factual inaccuracy in the quotation. DeepMind's response is an interpretation that we consider in the paper, though do not ultimately agree with. It relies on an assertion to the effect "clinicians said have all the data, therefore it is right" and an inappropriate analogy of Streams, an AKI alert app, to a general electronic patient management system.

Stimulating discussion, elaboration, and argument along the lines of DeepMind's response to this point is one of the reasons we thought it important to write the original research paper. However, the logical conclusion of DeepMind's reiterated argument, which we identify in Section 4.3, para 1 of our paper, is that any third party developing an app for healthcare can have access to all of the patient data in the trust they are working with, or even the whole NHS. We do not consider this to be a sustainable position.

There is a notable inaccuracy in DeepMind's rebuttal: the suggestion that our concerns about third party data access imply that hospitals should partially delete their own health records. This has never been our argument. To the contrary, the logic of medical information governance, which we fully accept and endorse, is that a hospital is the appropriate steward of highly sensitive medical data, except and until a clinical need arises or there is another approved data use.

6. Section 2, para 3: **'DeepMind said it was building a smartphone app [which] DeepMind claimed... would act as a mere interface to patient medical data controlled by the Royal Free'** [9]. This is not true. Reference 9 is used to support the statement made. It is to a newspaper article that quotes Mustafa Suleyman who makes clear that the app may well integrate AI or machine learning approaches in the future ('Despite DeepMind's

expertise in artificial intelligence (AI) and machine learning, the smartphone app being piloted does not use either technology. Mustafa Suleyman, co-founder and head of applied artificial intelligence at DeepMind, said “that may change in the future”. It also makes clear that the lack of AI use applies only to ‘just early pilots.’).

Powles & Hodson: The article deals exclusively with DeepMind’s activity between November 2015 and October 2016. What DeepMind might do in future is not relevant.

As with point 2, this point confuses the legal basis and purposes for which DeepMind was processing patient data between November 2015 and October 2016 with the “interest” that DeepMind had in future work.

7. Section 2.2, para 4: **‘The data package described in the ISA and destined for DeepMind is patient identifiable, and includes the results of every blood test done at Royal Free in the five years prior to transfer [18]. It also includes demographic details and all electronic patient records of admissions and discharges from critical care and accident and emergency. It includes diagnoses for conditions and procedures that have a contributory significance to AKI, such as diabetes, kidney stones, appendectomies or renal transplants, but also those that do not, such as setting broken bones.’** This is factually incorrect. Standard clinical practice requires that the clinician treating a patient with acute kidney injury has access to the full past medical history. Streams provides clinicians with the ability to access all historical pathology results to obviate the need for them to access such results from other desktop electronic systems or paper records. Such data are routinely accessed from existing electronic and written patient records or through direct communication with patients during clinical encounters. The data presented on mobile are a subset of those available on existing desktop platforms. The utility of these data for such clinical teams is not a hypothetical one – the data being processed is currently being utilised, live, on the Royal Free London NHS Foundation Trust (RFL) patients with early feedback from clinicians that Streams is speeding up their response times to AKI patients.

As for **‘setting broken bones’**, AKI is a recognised complication of both trauma and elective orthopedic procedures (fn2). Prior to a patient presenting to the RFL (or any acute NHS hospital), it is not possible to predict who will develop acute kidney injury, which is why the NHS national algorithm is applied across all patients. For example, fit 25-year-olds can develop AKI as a consequence of an appendicitis as much as an 80-year-old with a hip fracture, and it is always essential to pick all these cases up early. It is beyond the scope of this response to discuss clinical practice further

- but the past existence of broken bones (the exemplar used) may well have relevance to a diagnosis of renal impairment (through medication use, disease states which can predispose to both fracture and renal dysfunction, and more).

Powles & Hodson: Again, this response offers a different interpretation, and casts it improperly as a factual correction.

The response to the first half of this point is the same as to point 5. Streams, originally, was an AKI alert app, not a patient record presentation app. DeepMind’s cover letter to this appendix confirms this. Our paper expressly deals only with the original agreement, not DeepMind’s November 2016 pivot into wider patient record management. We have never contested the utility of Streams for a subset of patients with identified clinical need.

Broken bones was offered as an example, and we don’t dispute that such a treatment can contribute to AKI for patients receiving ongoing care. Our point was a different one. We were challenging DeepMind’s justification for holding historical data and its breadth of ongoing data retention. A patient might come to Royal Free for an accident involving a broken bone and never return—should their data continue to be processed by DeepMind? What about patients who have since passed away? The rights of patients must be considered in addition to the will of clinicians. It is specious to claim that any and all historical data is relevant to future AKI detection.

8. Section 2.3, para 1: **‘Both DeepMind and Royal Free claim that Streams relies solely on a ‘national algorithm’ for AKI published by the NHS; a process designed to assist in the rapid diagnosis of AKI from the starting point of a renal blood test for creatinine levels. The implication is that all that Streams does is host this algorithm, and pump the Royal Free data (as stored, structured, formatted and delivered by DeepMind) through it to generate alerts.’** This is incorrect: The AKI detection algorithm has been mandated by NHS England. Streams, as per national requirements, utilises this to detect possible AKI. Neither DeepMind nor the Royal Free London NHS Foundation Trust have ever claimed that all Streams does is run this detection algorithm. Streams provides (from existing systems) data to help clinicians to treat AKI once it has been detected as well as the detection function. This includes data on past medical history and historic pathology results, all of which are critical to helping clinicians to care for their patients.

Powles & Hodson: This objection appears to restate our point in different words, while willfully misinterpreting the context. The paper examines DeepMind’s processing of data for Streams as an AKI alert app, which was its clearly stated purpose during the period under examination. Casting Streams as a patient record

management app, rather than an AKI alert app, is a consistent theme of DeepMind's new position, even though Streams only sought to adopt this position under new agreements in November 2016, which falls outside the paper's stated scope.

DeepMind insisted repeatedly that the national algorithm was the only decision-making component of Streams. That insistence was important, as any other algorithm would consist of development/research and certainly require approval from the Health Research Authority or Confidentiality Advisory Group. Our point here is clearly about the decision-making software implemented in Streams, not a claim that Streams is nothing more than the national algorithm.

9. Section 2.3, para 1: **'Adding any new functions to the app, or fulfilling any of the broader contractual purposes described in the ISA, would comprise research.'** This is factually incorrect. The addition of further software to deliver clinical functionality (e.g. displaying allergy status) for direct clinical care does not constitute research as defined by the NHS Health Research Authority (as detailed in point 34).

Powles & Hodson: This, again, is willful misinterpretation of context. Of course we do not claim that minor changes such as changing the appearance of the app, or displaying some data field such as allergy status, would constitute research. But providing any new alerting functionality, or using any algorithms other than the national algorithm, would.

10. Section 3, para 1: **'Between late April 2016, when the scale of the data transfer from Royal Free to DeepMind and the relative lack of constraints on its use became publicly known, and until at least October 2016, DeepMind and Royal Free maintained the narrative that the entire purpose of transferring millions of patient records was to assist with AKI diagnosis and alerts, under a relationship of direct patient care. This position, however, fails to justify both the initial breadth of the data transfer and the continued data retention.'** This is factually incorrect. The authors are stating an unsubstantiated opinion. As previously stated, Streams provides clinicians with necessary clinical data to treat AKI patients, data that are already available from other sources. The clinical utility of this dataset (the size and content of which was defined by clinicians at the Royal Free London NHS Foundation Trust (RFL) and directed and approved by the RFL acting as Data Controller) is manifest from its current live clinical use in treating AKI patients through the Streams application. This is entirely justified and a widespread everyday practise across NHS information technology systems.

Powles & Hodson: Our entire paper documents the evidence and identifies the stakes for the second sentence which DeepMind now claims to be "an unsubstantiated opinion." Reinforcing our concern, the ICO has asked for an Undertaking justifying continued data holding—it wouldn't have had any cause to do so if the historical and current data use was "entirely justified."

11. Section 3.1, para 1: **'Royal Free states that AKI affects "more than one in six in-patients". If, as DeepMind claims, it only uses patient data in the service of monitoring and treating AKI, then it follows that as many as five sixths of patients (though this quantity is very unclear on the current state of the evidence) are not in a direct care relationship with the company. The distinction between being monitored or treated for AKI and not being monitored matters, because under British medical information governance guidelines, a direct care relationship between an identified patient and an identified clinical professional or member of a clinical care team obviates the need for explicit consent. Without such a direct care relationship, however, and without another basis such as consent, a formal research authorization from the HRA CAG, or otherwise satisfying necessity requirements and introducing appropriate safeguards, it is unlawful to continue to process patient data under the UK Data Protection Act 1998 (DPA).'** This is factually incorrect. Direct care relationships in the context of medical confidentiality exist between a clinician or healthcare worker and a patient: they do not exist between a data processor (whom that health care worker's organisation may choose to contract) and a patient. The direct care use of Streams by the Royal Free London NHS Foundation Trust is predicated upon the secure storage of data for potential future use by their clinicians, which applies to standard patient data storage systems including laboratory, radiology and electronic patient records. Similarly, monitoring systems (such as those that monitor physiological observations or blood sugar level in diabetics) must accrue, store and interpret identifiable patient data, regardless of whether clinical action is required through the monitoring of such data. The processing of patient data in Streams operates under the same governance framework as other existing clinical software offered by commercial third parties acting as data processors. The utilisation of data to treat patients within Streams is not research under UK Health Research Authority definitions and so the reference to requiring research authorisations is also factually incorrect. Point 17, below, is also relevant to understanding the nature of direct clinical care relationships.

Powles & Hodson: This response completely ignores context, as do many others. DeepMind’s response would be true and valid if DeepMind was only processing creatinine blood tests. But DeepMind is processing deeply personal and intimate data, Trust-wide, so it is not the same as other narrow laboratory systems. The comment also seems to ignore the fact that to merely “hold” data is, as a legal matter, to “process” it under the DPA.

Our reference to research authorizations is because that could have been, if granted, a legal avenue for wider access to health data.

12. Section 3.2 Para 1: **‘Despite the public narrative’s exclusive focus on AKI, it is clear that DeepMind and Royal Free have always had designs on much grander targets.....These are vast ambitions, considerably out of step with DeepMind and Royal Free’s narrow public relations orientation towards their collaboration being entirely founded on direct care for AKI.’** This is factually incorrect. DeepMind and the Royal Free London NHS Foundation Trust have been public about their long term ambitions to broaden the range of their collaboration to benefit patients (see point 4 and Guardian article).

Powles & Hodson: Again, this paper only addresses the work that is legally underpinned by the September 2015 ISA. Aspirations outside that ISA are not relevant. The ISA itself states that DeepMind may process data for broader purposes than just AKI detection, but the PR response has always been that between November 2015 and October 2016, data was being processed only for AKI detection.

To reiterate once more, we are talking about ambitions for this dataset, not for DeepMind more widely.

13. Section 3.2, para 3: **‘Nascent indications of DeepMind’s plans for datasets that not only span a large healthcare trust such as Royal Free, but the entire NHS, have not yet received critical discussion.’** This is misleading. DeepMind has opened itself to repeated questioning through conference presentations, journalists’ interviews (referenced by the authors) and engagement with respected critical authorities including the Kings’ Fund, the Wellcome Trust, the British Medical Journal, the OSCHR (Office for Strategic Coordination of Health Research) sub-board and the National Institute for Health Research (NIHR). DeepMind held an engagement event in September 2016 attended by 150 members of the public with wide questioning of Mustafa Suleyman and other members of the team (subsequently viewed on YouTube >6800 times) (fn3). At the DeepMind Health launch in February 2016, Mustafa Suleyman announced that a

respected group of public figures (the Independent Review Panel) would act in the public interest as unpaid independent reviewers of DeepMind Health.

Powles & Hodson: Our point here turns to the future. It is about DeepMind’s general plans for datasets reaching across the NHS; not just its public communications around Streams and Royal Free.

Materials and discussion about DeepMind’s work with the NHS have very rarely been critical in the mentioned fora, and this has continued with the July 2017 report of DeepMind’s self-selected Review Panel. The case for genuinely independent, critical discussion in an academic context is clear from the degree of interest that has followed our article’s publication.

14. Section 3.2, para 5: **‘repurposing of Trust-wide Royal Free data.’** This is factually incorrect. There has not been and will not be any “repurposing” of the data by DeepMind as claimed here, and no evidence is presented by the authors to support their allegation that these data were repurposed. Both the 2015 and the 2016 agreements make it clear that the Royal Free London NHS Foundation Trust continues to determine the purpose of the processing exclusively, and that DeepMind can only act on its instructions for the purposes set out in therein. Accordingly, it is incorrect to suggest DeepMind has any ability to repurpose the data.

Powles & Hodson: When read in context, this sentence is clearly referring to Suleyman’s plans. There is no claim that data “were” repurposed. The point, however, is that it would be hard to carry out Suleyman’s plans without repurposing the data.

As in other points, it is clear that the instructions in the 2015 ISA were inadequate in substantively constraining DeepMind, to the point where they bring into question DeepMind’s classification as a mere data processor.

15. Section 4, para 1: **‘The most striking feature of the DeepMind-Royal Free arrangement is the conviction with which the parties have pursued a narrative that it is not actually about artificial intelligence at all, and that it is all about direct care for kidney injury—but that they still need to process data on all the Trust’s patients over a multi-year period.’** This is misleading. Data are being processed under the terms explained and made public and covered amongst others in point 4.

Powles & Hodson: We differ in the conclusion we draw from examining the public record.

16. Section 4.1, para 1: **‘five months after the data had been transferred into DeepMind’s control and during which product development and testing had commenced’**. This is factually incorrect. At all points the Royal Free London NHS Foundation Trust have remained the data controller, and DeepMind a data processor. Furthermore, product development was carried out on simulated data, not on real patient data.

Powles & Hodson: This is spurious. The quotation contains two incontrovertible facts, concerning (1) data being under DeepMind’s physical control; and (2) development and testing having commenced. We do not claim that DeepMind was a data controller (although we discuss this likelihood later in the paper, stating that it is certainly an arguable position). Nor do we claim that product development was carried out using patient data (although it remains the case that, beyond assurances from DeepMind staff, we have no way of truly knowing).

The word “control” is not the same as the legal phrase “data controller.” We critically examine the issue of DeepMind’s self-characterization as a data controller/data processor in a different part of the article. At the time of this correspondence, it remains the subject of regulatory investigation.

17. Section 4.1, para 3: **‘We do not know—and have no power to find out—what Google and DeepMind are really doing with NHS patient data, nor the extent of Royal Free’s meaningful control over what Google and DeepMind are doing’**. This is factually incorrect: the basis and justification for data processing has been fully explained and is publicly available in many documents the authors reference. The full limitations on the purposes for which DeepMind may make use of the data are set out in the publically available Information Processing Agreement entered into by the parties at the same time as the Services Agreement and DeepMind cannot independently determine to move outside of the scope of this permitted usage. Regular project governance meetings are held to monitor processing. The 2015 ISA contains the right for the Royal Free London NHS Foundation Trust (RFL) to monitor compliance with the ISA, including an express right for an on-premises audit. This addresses the concerns expressed in the article (e.g. in Section 4.1), that there is no ability to find out what DeepMind are doing with NHS patient data. Moreover, the RFL and DeepMind are subject to oversight from competent regulators, including the Information Commissioner’s Office (ICO). The ICO has a range of statutory investigatory powers exercisable against the RFL as a data controller, including the right to request information and to conduct a compulsory audit into the RFL’s data processing practices, including its

arrangements with its processors (contractual and otherwise). See also point 23 below.

Powles & Hodson: Again, DeepMind refuses to constrain its comments to the period covered by the paper. The November 2016 Services Agreement and Information Processing Agreement (IPA) are simply not relevant. DeepMind also refuses to acknowledge Google/Alphabet’s ultimate, total control over the company.

The problem is that the public has no way of knowing or finding out about DeepMind’s processing, and it is not clear that Royal Free has the capacity to constrain it. Royal Free’s consistent obstruction of FOI requests over our investigation has only highlighted its refusal to publicly examine data processing done by a Google/Alphabet company. Patients remain in the dark.

18. Section 4.1 para 3 **‘The data transfer was done without consulting relevant regulatory bodies, with only one superficial assessment of server security, combined with a post-hoc and inadequate privacy impact assessment’**. This is factually incorrect. It is not necessary to consult regulatory bodies to approve data processing agreements for direct clinical care. The authors view that the privacy impact assessment was inadequate is unsubstantiated opinion. The authors criticise the timing of the Royal Free London NHS Foundation Trust’s Privacy Impact Assessment (“PIA”), because it was conducted after the 2015 ISA was signed (see Section 3.3 of the Article). However, in many controller-processor relationships this will be entirely appropriate, as it is only after the parameters have been set that the parties are able to assess the impact of the project. Before this point, it may not be possible (or at least as effective) for the parties to assess the potential impact on individuals and determine the necessary safeguards which can be put in place. Note that the General Data Protection Regulation (GDPR) specifically requires a data processing agreement to oblige processors to assist with PIAs (Article 28(3)f) of the EU General data Protection Regulation (2016/679)), suggesting that it would be usual practice for PIAs to be conducted after a contract has been signed.

Powles & Hodson: We make no claims about the mandatory nature or timing of the privacy impact assessment or consultation. The quoted sentence is a statement of fact, and it is not incorrect. We do imply that such consultations/assessments might have been useful in avoiding numerous pitfalls.

19. Section 4.1, para 3: **‘The amount of data transferred is far in excess of the requirements of those publicly stated needs, but not in excess of the information**

sharing agreement and broader memorandum of understanding governing the deal, both of which were kept private for many months;'. This is factually incorrect. The data processed in the application have been defined by and are currently being used by clinicians for the direct monitoring and care of AKI patients. The authors cite no justification for the assertion that such data processing exceeds that which is clinically required than their own (non-clinical) opinions and their unsubstantiated and incorrect view that the data are being accessed by DeepMind for other purposes.

Powles & Hodson: Again, this is a wilful misinterpretation. We compare the needs to verifiable public statements, not to what clinicians said to DeepMind in private. It is a statement of fact that the data transferred is broader than the requirements of AKI. How, for instance, do dead patients need AKI diagnosis? We do not need to make any insinuation about any use for other purposes.

According to correspondence with the Trust, Royal Free only has a record of its Caldicott Guardian reviewing the deal months after it was signed, apparently in response to public criticism.⁵

20. Section 4.1, para 3: '**None of the millions of identified individuals in the dataset were either informed of the impending transfer to DeepMind, nor asked for their consent**'. This is misleading: It is not necessary or required to obtain consent for routine clinical software applications, nor is there any way that such consent could be practically obtained for such large scale hospital software deployments.

Powles & Hodson: This is not misleading, it is true. There was no consent. DeepMind argues in response to this that consent was not necessary. Our assessment of the facts suggests that it may have been, and that there are other large scale projects where it has been practical, as we discuss.

Further, DeepMind does not address the fact that patients were not informed about the processing—and *still* have not been informed at scale, despite all that has passed since the deal was first made. This highlights the company's conspicuous inaction on the core recommendation of DeepMind's advisor on patient and public involvement and engagement, the late Rosamund Snow, given to the company at some point in 2016. Criticizing arrangements where "clinicians expect to decide what information to elicit from patients rather than patients or carers contributing to those conversations," Ms. Snow made a number of recommendations, driven by what

she termed one overarching principle: "at every level where clinicians have influence, ensure patients do too."⁶

21. Section 4.2, para 2: '**Under the UK Data Protection Act, DeepMind needs to comply with a set of data protection principles**'. The authors have incorrectly referred to DeepMind having to comply with the data protection principles, including having a lawful basis for the processing. Data processors are not directly subject to the Data Protection Act. In fact, it is the Royal Free London NHS Foundation Trust (RFL) who must ensure that any processing of data for which it is the controller is carried out in compliance with the Data Protection Act and ensure it has a lawful basis for the processing. DeepMind, as RFL's data processor, is then reliant on the same lawful basis and subject to certain contractual restrictions and obligations placed on it by RFL to ensure RFL's continued compliance. Accordingly, the assessment in Section 4.2 of the Article regarding explicit consent or "necessary for medical purposes" should be applied to RFL, and not DeepMind. Similarly, the authors are mistaken of the view that DeepMind has positioned itself as having a direct care relationship with patients. Such a relationship arises only between RFL and the patients, with DeepMind (as a processor) acting only on behalf of RFL and not for patients directly.

Powles & Hodson: Compliance via the control of Royal Free is still compliance. We have made it clear in the paper that Royal Free is unquestionably a data controller. However, DeepMind is the subject of study here. The paper questions DeepMind's self-designation as a data processor. The question remains open. Like "direct care," DeepMind's desire to be a data processor does not become fact no matter how many times DeepMind asserts it.

22. Section 4.2, para 4: '**Data protection law relies on a key distinction between 'data controllers' and 'data processors'. A data controller is defined as "a person who (either alone or jointly or in common with other persons) determines the purposes for which and the manner in which any personal data are, or are to be, processed", while a data processor is "any person (other than an employee of the data controller) who processes the data on behalf of the data controller". It is crucial to define controller and processor status in any information sharing arrangement because legal**

⁵ Email from Royal Free Press Office to Hodson, 14 Jul 2016.

⁶ Snow R. Recommendations on patient and public involvement and engagement (PPIE) at DeepMind Health. 2016. In: For Patients. <https://deepmind.com/applied/deepmind-health/patients/>. Accessed 1 Aug 2017.

obligations and liabilities flow from it, with significant real-world consequences.’ The authors suggest that DeepMind may in fact be acting as a joint controller, and rely on the ICO guidance (fn4) in support of this. The ICO guidance to which the authors refer identifies the following seven items which should be determined by the controller (see para 16 of the guidance): (1) to collect the personal data in the first place and the legal basis for doing so; (2) which items of personal data to collect, i.e. the content of the data; (3) the purpose or purposes the data are to be used for; (4) which individuals to collect data about; (5) whether to disclose the data, and if so, who to; (6) whether subject access and other individuals’ rights apply i.e. the application of exemptions (7) how long to retain the data or whether to make non-routine amendments to the data. All of the above continue to be exercised by the Royal Free London NHS Foundation Trust exclusively. None of these decisions are determined by DeepMind.

Powles & Hodson: DeepMind’s recitation of the ICO guidance on data controllers fails to address our arguments based on that guidance and, in particular, the broadness of purpose in the ISA, which allows DeepMind scope to create tools based upon its analyses.

23. Section 4.2, para 6: **‘DeepMind seems to have considerable discretion, in addition to Royal Free, to determine the purposes and manner in which any personal data is processed. The company is storing, structuring and formatting the Trust-wide dataset, testing it, preparing to deliver data and visualizations to clinician’s devices and, most recently, discussing technical infrastructure that could enable it to be repurposed. These factors all point very strongly to DeepMind assuming the role of a joint data controller.’** The authors base their arguments around joint controllership on the discretion DeepMind has in relation to the processing. However, the ICO accepts that certain decisions may be delegated to the data processor. These include what IT systems to use, how to store the data, the security details and how to destroy the data. This accords with the Article 29 Working Party’s (“WP29”) view that the data controller must determine the purpose of the processing, but may delegate the determination of the means to the processor⁶. The WP29’s opinion is subject to the requirement that substantial questions as to the means of the processing (e.g. what data will be processed, length of storage) continue to be determined by the controller. This is precisely the case in respect of the arrangement with the Royal Free London NHS Foundation Trust (RFL): DeepMind may have some discretion as to

aspects of the how (i.e. the means), but it has no discretion as to the why (i.e. the purpose).

The 2015 ISA also contains explicit restrictions on how DeepMind can use the data, which contradict the authors’ assertions (for example on p.2 and 3, and the list on p.14–15) regarding the lack of assurances on how DeepMind will use the data. The 2015 ISA states that only DeepMind and Google UK appointed staff can access the data dealing with authors’ concerns regarding access by Google, Inc. It also prohibits data leaving the European Economic Area (again, preventing access by US-based Google, Inc). Most importantly, the 2015 ISA states that “The processor will act in accordance with the Data Controller’s instructions and will only use the personal data to provide the services under this Agreement” (see p.6 of the 2015 ISA). Accordingly, regardless of any comments made outside the 2015 ISA or lack of specific restrictions (e.g. constraining the use of artificial intelligence or further use of the data), all use of the data by DeepMind can only be as instructed by RFL. Therefore there can be no suggestion of DeepMind linking the data with Google accounts, for example (see the authors’ concerns at Section 2.1 of the Article that there is “no legal foundation” for the assurance that the data will never be linked with Google data”). This would not be permitted under the 2015 ISA unless DeepMind was specifically instructed to do so by the RFL (and, even in that circumstance, only for the purposes identified in the 2015 ISA). Established practice is that data processing agreements should be drafted to contain a restriction on processing the data in any manner other than as instructed. This is a far more effective means for controllers to exercise control over their data processors, rather than attempting to specify every restriction. The wording in the 2015 ISA is deliberately narrowly restrictive, rather than generally permissive with express restrictions.

Powles & Hodson: We present the legal characterization of DeepMind as a question. DeepMind’s status as an artificial intelligence company means that its “means of processing” are particularly of interest, and very hard for Royal Free, a cash-strapped NHS trust, to even understand. The fact that Royal Free is completely unable to answer questions about the size and details of the dataset being processed by DeepMind is indicative of a lack of exclusive control.

In response to post-publication FOI enquiries, which required consultation with DeepMind, Royal Free was unable to provide any “instructions” sent from the Trust to DeepMind,⁷ providing

⁷ Royal Free response to Hodson freedom of information request 17/ 206, 4 Aug 2017.

instead only the original, loosely-constrained ISA. Absent other binding instructions, DeepMind’s argument rather evaporates.

24. The Article focuses its analysis on the 2015 ISA. The 2015 ISA is based on the Royal Free London NHS Foundation Trust (RFL)’s standard form, and there is no suggestion by the authors that it does not comply with the requirements of the Data Protection Act 1998. However, this agreement has now been terminated and replaced with two longer documents, dated 10 November 2016, between DeepMind Technologies Limited and the RFL. The 2016 Services Agreement and Information Processing Agreements are both available on DeepMind’s website (fn7) but, according to the authors, review of these documents is “beyond the scope” of the Article and “will be considered in future work”. The 2016 documentation currently governs the Streams project with RFL, and has been drafted to meet the requirements of both the Data Protection Act 1998 and in consideration of the incoming EU General Data Protection Regulation (2016/679). The 2016 documentation contains express restrictions on combining and/or disclosing the data, sets out an information governance process, and contains a clear explanation of the data flows.

Powles & Hodson: This point repeats the complaint stated in point 3 and implicit in others, which is DeepMind’s desire to focus on its revised agreements, not the prematurely superseded 2015 ISA. We think the foundations of this deal cannot be ignored, and that it is essential to examine the conduct and legality of one year of data processing on millions of patients.

25. Section 4.2, para 6: **‘It seems clear that Royal Free (sic) have contracted with DeepMind to analyse complex data and come up with solutions by applying DeepMind’s own expertise in analysis to an extent that Royal Free (sic) cannot begin to do’.** This is factually incorrect. DeepMind are not processing the Royal Free London NHS Foundation Trust’s data using machine learning or any advanced analytics. Such work would not be covered by the current agreements and would need to be covered by separate contractual agreements and governance approvals (including research ethics).

Powles & Hodson: Our point is primarily based on the asymmetry in technical expertise between the two entities, which has been consistently affirmed in all of the communications about the arrangement, as well as by the subsequent evolution of DeepMind’s NHS work. DeepMind’s rebuttal is directly refuted by the 2015 ISA, which states that the information is being shared to do analytics. It is also worth noting that DeepMind

had already applied for HRA ethics approval to apply machine learning to the data in October 2015—and had that application granted in November 2015. See further points 28 and 34 below.

The discussion of machine learning is, nevertheless, separate to the point we actually make in Section 4.2, para 6. There is no need to prove that DeepMind engaged in machine learning in order to question its self-classification as a “data processor.” The evidence we cite instead is what DeepMind has done with the data—“storing, structuring and formatting the Trust-wide dataset, testing it, preparing to deliver data and visualizations to clinician’s devices and, most recently, discussing technical infrastructure that could enable it to be repurposed”—and the weak arrangements under which this was done.

26. Section 4.3, para 3: **‘For individuals who are escalated to clinical intervention based on the results of applying the AKI algorithm after a preliminary blood test, clearly this direct care scenario applies. However, for the remainder of patients whose data has been transferred to DeepMind, no plausible necessity for DeepMind’s processing of their data arises.’** This is factually incorrect: the authors assert this again and the response to points 5, 6, 7 and 11 applies.

Powles & Hodson: Again, our analysis is developed in consideration of the UK Data Protection Act and Caldicott Guidelines, not the opinions of clinicians at Royal Free. A difference of opinion is not a factual inaccuracy. DeepMind’s position also utterly and inappropriately negates the rights of patients.

27. Section 4.3, para 3: **‘It is, instead, a classic situation of health services management, preventative medicine, or medical research that applies to the overall provision of services to a population as a whole, or a group of patients with a particular condition. This is the very definition of indirect care’** This is factually incorrect: the rationale for the data processing arrangements for Streams cannot, by any reasonable interpretation, be described as “health services management, preventative medicine or medical research”. The article describes indirect care applications as “research on identifiable individuals or risk prediction and stratification” - not consistent with the functionality of Streams at the Royal Free London NHS Foundation Trust (a clinical application to monitor patients and access data in the course of clinical care, which can only be described as a tool for ‘direct patient care’).

Powles & Hodson: Yet again, DeepMind’s assertion that what it is doing is “direct care” across the totality of the dataset does not make it so as a question of fact.

28. Section 4.4, para 1: **‘At the heart of this deal is a core transparency paradox. Google knows a lot about all of us. For millions of patients in the Royal Free’s North London catchment, it now has the potential to know even more.’** This is misleading - Google does not have the ability to know more about patients in North London through linkage of accounts. The assertion that ‘Google’ can access data processed for the Streams application has no basis in fact and such access would be illegal.

Powles & Hodson: DeepMind is part of Google. DeepMind is controlled by Google. The ISA under scrutiny in the article was signed by Google, and it poorly constrained DeepMind’s behavior with the data in question. Had DeepMind trained models on this data, as it applied for permission to do in October 2015, those models would have been the property of Google, and Google would indeed know more about those patients, even if not definitively in a personally identifiable manner. It is exactly this kind of behavior, hypothetical but not unexpected from a data mining company, that is a key warning signal in our paper.

29. Section 4.5, para 1: **‘The deal-making between DeepMind and public institutions continues to be secretive’.** This is factually incorrect. Relevant details of our partnerships are publicly available on the DeepMind website (fn 8). The press have been actively informed about our work, leading to extensive coverage (e.g. in the press including the BBC and Guardian (fn 9, 10).)

Powles & Hodson: DeepMind’s deals have been publicly announced only after signing. There has been no public consultation prior to this point. Doing public relations does not necessarily mean lack of secrecy.

30. Section 4.3, para 6: **‘Given Streams is characterized as a clinical app, there are more elegant—and less legally and ethically dubious—solutions available than simply running a mirror copy of the Royal Free’s repository of patient data on third-party servers controlled by DeepMind, for every single hospital patient, entirely independently of AKI susceptibility and diagnosis. One solution is for DeepMind to pull in historical data only on patients who have had the gateway blood test that is prerequisite for AKI diagnosis. If Royal Free’s systems cannot currently handle real time data requests in this manner, they ought to.’** This is misleading: the proposed “elegant” solution only “pulling in” data on patients who have had “the gateway blood test” (presumably creatinine) - would not be a clinically safe proposition and we are unsure if Powles and Hodson make these recommendations based

on external expertise given that neither are clinicians. In order to safely initialise the digitally-enabled care pathway for improved AKI management, consultant clinical nephrologists at the Royal Free Hospital instructed DeepMind to process five years of historical data for all patients, providing valuable clinical context for all patients inclusive of those who have not had a creatinine measured and enabling alerts to be generated to display this clinically-relevant information to clinicians.

Powles & Hodson: This objection is another version of “the clinicians said we should, so we’re right.” Our article examines the November 2015 – October 2016 deal against the law, not clinical opinion.

31. Section 4.7, para 1: **‘Offering DeepMind a lead advantage in developing new algorithmic tools on otherwise privately-held, but publicly-generated datasets....’.** This is factually incorrect. There is no basis for the assertion the DeepMind have been ‘offered’ a ‘lead advantage’ in the deployment of algorithmic tools in healthcare or in fact more widely. Such efforts are global and have multiple actors including but not limited to IBM and Microsoft, and have been ongoing many years before DeepMind was founded in 2010.

Powles & Hodson: Our point is specific, not general. The evidence supporting this assertion is a complete absence of similar research in the UK. No-one has ever been given access to the Moorfields dataset before, for instance. The fact that such datasets are generated and curated by publicly-funded institutions is crucial, and unaddressed in DeepMind’s response. In the wake of the ICO ruling, Royal Free and DeepMind have termed their partnership “unprecedented.” Which is it, a standard deal that did not deserve all this attention, or a fundamental arrangement to the future of healthcare?

32. Section 4.7, para 4: **‘The value embodied in these NHS datasets does not belong exclusively to the clinicians and specialists who have made deals with DeepMind.’** We agree. To clarify in the Streams application (and separately for DeepMind research projects) no data ownership is transferred to DeepMind. Partnership agreements were approved by the Royal Free London NHS Foundation Trust’s executive board and not by individual clinicians or specialists supporting the projects.

Powles & Hodson: Our point here is that the value embodied in these datasets for patients and the public, as those that generated the data and without whom it would not exist, has been

entirely omitted from consideration. DeepMind have zero response to this.

33. Section 4.8, para 2: **‘First mover advantage’ exists as it does whenever private companies exploit public resources’**. This is factually incorrect. DeepMind is not a first mover in either health analytics or computer aided diagnostics. Many large multinational technology companies including IBM, Microsoft, GE and Siemens have been delivering these services to the NHS and international health systems for decades. In addition there are many small and medium sized enterprises currently working in this sector with the NHS.

Powles & Hodson: The context of our discussion of public resources is the Royal Free case study, where DeepMind has clearly been given a first mover advantage. We are not making any broader assertions about competition in health analytics and diagnostics in general.

34. Section 2.3 of the Article states that: **‘Both DeepMind and Royal Free claim that Streams relies solely on a ‘national algorithm’ for AKI published by the NHS, a process designed to assist in the rapid diagnosis of AKI from the starting point of a renal blood test for creatine levels. The implication is that all Streams does is host this algorithm and pump the Royal Free data (as stored, structured, formatted and delivered by DeepMind) through it to generate alerts... Adding any new functions to the app, or fulfilling any of the broader contractual purposes described in the ISA, would comprise research. DeepMind did not have the requisite approvals from the Health Research Authority (HRA)...’** The above comment appears to oversimplify the definition of research requiring approval by the HRA, and it is worth noting that whether an activity falls within the HRA’s definition of research requiring HRA approval should be considered on a case by case basis. The definition of research used by the HRA is where: (a) participants in the study are randomised to different groups, and/or (b) the study protocol demands a change in treatment/patient care from accepted standards for any of the patients involved, and/or (c) the findings will be generalisable (i.e. will be used to derive generalisable new knowledge). In this case, approval from the HRA was not required since the activities did not constitute research as defined above.

Powles & Hodson: The question is more whether data processors should apply to HRA CAG. The CAG approves processing of identifiable patient data for more than just research.

Furthermore, the fact that DeepMind applied for an REC opinion on machine learning for AKI in October 2015 suggests that this was the plan (original ref. [22] and [30]). See also our response to point 8 and reporting by TechCrunch.⁸

35. Section 3.3 of the Article states that **‘... The Medicines and Healthcare products Regulatory Agency (MHRA) regulates medical devices. None of these bodies were approached about the November 2015 data transfer; ...not to go through an official and required device registration process with the MHRA before starting live tests of Streams at Royal Free in December 2015... (DeepMind has subsequently been in discussion with all of these parties in reference to its Royal Free collaboration and, for several months from July 2016, stopped using Streams until the MHRA-required self-registration process was completed).’** Contrary to the statement of the authors, it was not mandatory for DeepMind to consult the MHRA in advance of any data processing in 2015 nor was it a requirement for DeepMind to go through an official device registration with the MHRA since the project at this stage involved neither the use of a non-CE marked medical device, or the use of a CE marked medical device outside of its intended purpose. Article 1(2)(e) and Section 2.1 of Annex X (implemented into national law by the Medical Device Regulations 2002, Regulations 2(1) and 16.) of the Medical Devices Directive (93/42/EEC) explain that a clinical investigation is an investigation designed to establish that the performance of a device claimed by the manufacturer can be adequately demonstrated and that a device is judged to be safe to use on patients taking into account any risks associated with its use weighed against the intended performance. MHRA guidance acknowledges that clinical investigations are not needed for all types of novel software (MHRA Guidance: Medical device stand-alone software including apps (including IVDMDs). The project did not fall within the definition of a clinical investigation for which a notification to the MHRA was required. Similarly, the project did not involve the use of a medical device which was required to be CE marked at that point in time, given the definition of a medical device in Article 1(2)(a) of Directive 93/42/EEC (Regulation 2(1) of the Medical Device Regulations 2002.), since it did not involve the use of a device for a medical purpose as set out in that Directive (or even a finished device). Therefore DeepMind was not required to have gone through an “official and required” device registration process at this time. As

⁸ Lomas N. Documents detail DeepMind’s plan to apply AI to NHS data in 2015. TechCrunch. 31 Aug 2017. <http://tcrn.ch/2x9IIsn>.

the authors note, as the project has progressed DeepMind have actively engaged with the MHRA and obtained all necessary approvals for Streams.

Powles & Hodson: At the time of writing our original article, we did not have clarity on the MHRA's determination of the Streams app's status from December 2015. It is only through MHRA responses to FOI requests received on 26 May 2017, well beyond the date of submission of the article, that the MHRA stated that DeepMind was not required to go through any processes whatsoever, even a simple registration, in order to test Streams at Royal Free. We find this surprising, and question the sense of being able to test medical apps in live settings without going through any regulatory processes.

The FOI response of 26 May 2017 included correspondence from a DeepMind clinician arguing strongly that Streams is not a medical device at all, and did not need to be registered as such. The email thread showed further an MHRA official disagreeing with DeepMind, and DeepMind promptly taking steps to register the device. This demonstrates that our earlier deductions, written without knowledge of these communications, were sound at the time.

36. Section 2, Para 3: **'Why DeepMind, an artificial intelligence company wholly owned by data mining and advertising giant Google, was a good choice to build an app that functions primarily as a data-integrating user interface, has never been adequately explained by either DeepMind or Royal Free.'** This is factually incorrect. DeepMind has gone to great lengths to explain its aims in healthcare technology and why it feels it is beneficial to incorporate efforts to not only deploy AI but clinically useful clinical software and data processing applications. These explanations have been provided through press releases, website content, public presentations at several events and a globally streamed patient and public event as outlined in points 4 and 12.

Powles & Hodson: We do not claim that DeepMind has not explained its aims. We claim that it has not explained why it is a good choice to build a medical data presentation app.

37. Section 5, Para 1 (and Section 2, para 1): **'It has done so without any health-specific domain expertise'**. The statement that DeepMind has no health specific expertise is incorrect. Publicly available information about the DeepMind team makes clear that it draws on decades of experience across clinical medicine, health policy, research and health IT (fn 11). DeepMind's work in AKI

has been led throughout by experts in kidney disease, intensive care and informatics, and supported by broader clinical teams in the Royal Free London NHS Foundation Trust. This is publicly described on the DeepMind website.

Powles & Hodson: The word "expertise" could have been more wisely used. DeepMind has, of course, hired and worked with health experts on these projects. The point remains that DeepMind has no background delivering healthcare software (as we correctly state in Section 2, para 1), but it is owned by a company with a deep background in data processing and advertising. Given that, DeepMind's first foray into health service provisioning should have been more cautious.

DeepMind's assertions of "lack of adherence to Springer journal guidelines":

DeepMind claims "the article does not adhere to the Journal's Instructions for Authors in numerous areas identified [in points 38-47] below":

38. Only 28 out of 102 references could be said to come from peer-reviewed journals and academic publications. Journal guidance states that 'the list of references should only include works that are cited in the text and that have been published or accepted for publication. Personal communications and unpublished works should only be mentioned in the text'. The following examples are references not from published research (1, 2, 3, 4, 5, 9, 11, 12, 13, 17, 18, 21, 22, 23, 30, 21, 32, 37, 40, 41, 42, 43, 45, 47, 48, 49, 50, 51, 53, 56, 61, 84, 100, 101).

Powles & Hodson: This point and those that remain seem to be crafted to question the academic integrity and value of the paper, but from a laughable premise. The response fundamentally misunderstands the nature of journal guidelines (they are standards, not rules). Academic research on contemporary issues must, of course, rely on blogs, newspaper articles, personal communications, letters, and open-access reviewer comments. Quite simply, other verifiable sources do not exist.

There is a particular irony in this point coming after so many that effectively distil to "this article should have relied on DeepMind's blog, speeches, public communications, and placed pieces of access journalism"—all unpublished themselves, of course (e.g., point 36 above). It also demonstrates how unreflective DeepMind has been about the transparency paradox that we address in Section 4.4 of our paper.

39. Personal communications and letters are included as references, against Springer guidelines (2, 16, 31, 50, 53, 56).

Powles & Hodson: See response to point 38.

40. Journalist opinion are included as references, against Springer guidelines (15, 16, 33, 36, 38, 39, 46, 49, 54, 55, 57, 62, 72, 73, 85, 86, 89).

Powles & Hodson: See response to point 38.

41. References 81–83 are not published work but open-access referee comments.

Powles & Hodson: See response to point 38.

42. Footnotes and endnotes are repeatedly used as substitutes for references lists against Springer guidance (3, 4, 10, 13, 18, 21, 22, 23, 26, 45, 61, 68, 73, 74).

Powles & Hodson: There were extensive rounds of review during the production process. The final separation of references and footnotes was done by the Springer production team.

43. Personal blogs are included as references against Springer guidelines (71).

Powles & Hodson: See response to point 38.

44. Reference 10 does not support the claim made in paragraph 2.1, as it does not refer to the factual matters being discussed.

Powles & Hodson: This is an obscure target, and it misses.

45. Reference 60 is an example of a references that does not guide the reader to the salient discussion point.

Powles & Hodson: Having enumerated extensively and comprehensively in points 38–42 the footnotes in issue, it seems odd to just select “an example” for this point, which again misses its target.

Reference 60 is a particularly unfortunate example, given that it is a citation to a report that eviscerated projects that do not engage patients—a criticism that applies equally to DeepMind’s work with Royal Free.

46. Reference 83 points to a published research protocol of no relevance to the point being made.

Powles & Hodson: This repeats point 41. We differ in our assessment of the relevance of the protocol.

47. Published links to references 37, 51 do not work.

Powles & Hodson: As irony would have it, these are DeepMind’s own links. We invite the company to make available more stable links and to be transparent in signaling modifications to their content.