



Editorial DeepMind

Luis Kun¹

Published online: 26 April 2018

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

This Editorial exists as a direct result of the formal complaint from Google DeepMind that was presented before the Editorial Board, and the Editor-In-Chief of The Journal of Health & Technology (IUPESM – WHO – SPRINGER NATURE), who cited certain objections to the publication of the peer-reviewed article titled: “Google DeepMind and healthcare in an age of algorithms,” authored by Julia Powles and Hal Hodson. On the 15th of May 2017, Dr. Dominic King, Clinical Lead at DeepMind, dispatched a communication, titled: “*Letter from DeepMind in response to Google DeepMind and healthcare in an age of algorithms,*” to this Editor-In-Chief (E-I-C).

That first email read as follows:

*“Dear Professor Kun,
Please find attached our response to the following ‘original article’ published online in the Journal of Health and Technology on the 16th of March 2017.*

Google DeepMind and healthcare in the age of algorithms

Julia Powles and Hal Hodson

This article provides an inaccurate representation of DeepMind’s activities and intent and we look forward to your editorial response to the issues raised. Please get in touch if you would like me to submit this response through an online system.”

The objections, as stated by King and colleagues on behalf of DeepMind specified that there were purportedly, “*factual inaccuracies and unvidenced statements alongside multiple failures to follow Springer journal author guidelines for an academic article,*” as it was being presented by Powles and Hodson.

✉ Luis Kun
hcai@aol.com

¹ W. Perry, Center for Hemispheric Defense Studies at National Defense University, Washington, DC, USA

As Editor-In-Chief (E-I-C) of the Journal of Health and Technology, the submitted complaint was discussed with the Journal’s Editorial Board. A consensus was reached among all concerned that DeepMind, should be permitted to present their complaint in full, and for the authors to also be afforded an equal opportunity to present an opportunity to rebuttal to the complaint.

Subsequent to the Editorial Board’s decisions, and E-I-C’s further considerations, additional communications between the E-I-C, DeepMind (through Dr. King), and the authors (through Dr. Powles) transpired during the months of May, June, July, August, October and December 2017.

The decided-upon action by the Journal’s Editorial Board and the E-I-C constitutes the following three parts. 1) To publish the entire complaint from DeepMind; 2) To allow the authors to address the complaint in a direct, transparent and unobstructed manner, 3) To publish the totality of the complaint, and the authors’ reply to the complaint - in a future issue of the Journal, and, lastly, 4) to address DeepMind’s complaint regarding “*failure to follow Springer Journal guidelines,*” by the E-I-C.

Action 1 - To this end, the entire content of the DeepMind letter is included in Appendix 1, of this Editorial.

Actions 2 and 3 - Following this Editorial from the E-I-C, two other related editorials appear in this Journal Issue:

1. The first, written by Dr. King’s team at DeepMind, entitled *Letter to the Editor* (#2), received on July 28, 2017, and a subsequent, entitled: “Letter in response to Google DeepMind and healthcare in an age of algorithms”; and
2. The second, written by Dr. Powles’ and her co-author, responding to the concerns of DeepMind, which is labeled: *Letter to the Editor* #3

Submissions from parties are published herein, in their entirety.

Action 4 (“*failure to follow Springer Journal guidelines,*”) In Appendix 1, Journal readers will find items 38 through 47 under the subtitle: “*Lack of adherence to Springer journal*

guidelines”. DeepMind’s claim is that, the article (Powles/Hodson) does not adhere to the Journal’s “Instructions for Authors” (http://www.springer.com/engineering/biomedical+engineering/journal/12553?detailsPage=pltc_1325108) in numerous areas which include:

1. 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.
2. Personal communications and unpublished works should only be mentioned in the text’.
3. Journalist opinion are included as references,
4. Some References are not published work but open-access referee comments,
5. Footnotes and endnotes are repeatedly used as substitutes for references lists,
6. Personal blogs are included as references against Springer guidelines.

The Journal of Health & Technology, showcases two main types of publications. Unsolicited and Solicited manuscripts are the two main types of submissions that the Journal considers at the present. Unsolicited manuscripts, refers to papers written by authors routinely interested in publishing their work, in the Journal, and independently, of any Special Issues of the Journal. Solicited manuscripts are generally submitted in response to a Special Issue announcement, on a very specific topic, or subject. In both cases manuscripts undergo peer review. In addition, Special Issues have the involvement of at least one (sometimes more than one) Guest/Associate Editor (s). The paper in question is part of a Special Issue (on Privacy and Security of Medical Information). Ultimately, peer-reviewed manuscripts are finally cleared by the E-I-C, prior to its acceptance and publication.

The Journal’s Inaugural Editorial (*Just another journal? No, a different one!*) explained just how different this Journal was to be (read Bos, L. & Kun, L. Health Technol. (2011) 1: 1. <https://doi.org/10.1007/s12553-010-0001-9>). As a result, with regard to Special Issues, Guest / Associate Editors are afforded greater latitude so that they are able to address a much wider and internationally diverse audience including: physicians, nurses, economists, communications specialists, biomedical engineers, medical physicists, computer scientists, lawyers, policy makers, etc. In spite of the diverse aims of the Journal, Springer Nature Publication standards are not compromised.

For example under, the Journal’s: “Instructions to authors”, which can be found at: http://www.springer.com/engineering/biomedical+engineering/journal/12553?detailsPage=pltc_1325108, the use of footnotes is encouraged in certain cases. Within the instructions, one can find, an *Editorial Procedure* section, which reads:

“As indicated in the Aims and Scopes, the audience for this journal is a very broad one. The authors are therefore requested to ensure that their writing will also be accessible for a less, or non-scientific audience. Authors are kindly requested to give a basic explanation (if not in text than by footnote (not endnote)) of formulas or equations used”.

The Journal’ guidelines, are, by definition, not a strict and inflexible set of rules, rather, “*a line by which one is guided*” (ref. Merriam Webster). The Journal has not deviated from ACCEPTABLE / ACCEPTED standards. Guidelines to AUTHORS detail “guide-posts”; albeit that published peer-reviewed sources are most-preferred, they do NOT represent the SOLE fountainheads for relevant, meaningful and credible information. It was the synthesis and decision of the Editorial Board and the E-I-C, that the manuscript, after peer-review, was deemed to be value to our readership, and therefore, entirely appropriate for publication in the Journal.

Compliance with ethical standards

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

Appendix 1

15th of May 2017.

Dear Editor,

Google DeepMind and healthcare in an age of algorithms. Julia Powles and Hal Hodson.

We would like to respond to this “original article” published in the Journal of Health and

Technology ([doi:https://doi.org/10.1007/s12553-017-0179-1](https://doi.org/10.1007/s12553-017-0179-1)). The article concludes that DeepMind was “*given unfettered, unexamined access to population-wide health datasets*” by the Royal Free London NHS Foundation Trust and involved “*secretive deals and specious claims*”. This is a wholly inaccurate representation of the company’s activities and intent. We list below 37 factual inaccuracies and unevicenced statements alongside multiple failures to follow Springer journal author guidelines for an academic article.

We are raising these concerns with you so that you can determine the most appropriate editorial action.

We look forward to your response.

Sincerely yours.

Dr. Dominic King BSc MBChB MEd PhD MRCS.
Senior Clinician Scientist and Clinical Lead, DeepMind Health.

Dr. Alan Karthikesalingam MA MBBChir MSc PhD MRCS.
Senior Clinician Scientist, DeepMind Health.

Dr. Cian Hughes MB ChB MSc MRCS.
Senior Clinician Scientist, DeepMind Health.

Professor Geraint Rees MA MB BCh PhD MRCP FMedSci.
Senior Scientific Advisor, DeepMind Health.

Professor Hugh Montgomery, MBBS BSc FRCP MD.
Clinical Advisor, Deepmind Health.

Professor Rosalind Raine MBBS, BSc, MSc, PhD, FFPH.
Clinical Advisor, Deepmind Health.

On behalf of The DeepMind Health Team.

Factual inaccuracies and un evidenced statements.

The article contains numerous errors of fact identified 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.
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.¹ 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*.
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.

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).
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

¹ The video of the launch of DeepMind Health at Royal Society of Medicine can be viewed at www.youtube.com/watch?v=CEvMXuAw5DY&t=1467s

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.

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.’).
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.² 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).
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.
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).
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***

² AKI is well known to be associated with acute orthopaedic issues not limited to trauma and bone and joint infections. 25% of elderly patients with a hip fracture develop an AKI www.bmcnephrol.biomedcentral.com/articles/10.1186/s12882-017-0437-5

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.

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.
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).
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).³ 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.
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.

³ YouTube video of DeepMind Patient and Public Engagement Event. September 20th. 2016. <https://www.youtube.com/watch?v=wBfBiD38x34>

- Accordingly, it is incorrect to suggest DeepMind has any ability to repurpose the data.
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.
 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.
 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.
 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.
 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.
 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.
 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 mistakenly 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.

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 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 guidance⁴ 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.
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.⁵ 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 processor.⁶ 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

⁴ <https://ico.org.uk/media/1546/data-controllers-and-data-processors-dp-guidance.pdf>

⁵ <https://ico.org.uk/media/for-organisations/documents/1546/data-controllers-and-data-processors-dp-guidance.pdf> (p7)

⁶ http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2010/wp169_en.pdf

- 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.
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 website⁷ 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.
 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).
 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.
 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’).
 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.
 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.⁸ The press have been actively informed about our work, leading to extensive coverage (e.g. in the press including the BBC and Guardian^{9,10})
 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

⁷ <https://deepmind.com/applied/deepmind-health/transparency-independent-reviewers/>

⁸ <https://deepmind.com/applied/deepmind-health/>

⁹ <http://www.bbc.co.uk/news/technology-37230806>

¹⁰ <https://www.theguardian.com/technology/2016/jul/05/google-deepmind-nhs-machine-learning-blindness>

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.

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

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.
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.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.

Lack of adherence to Springer journal guidelines.

The article does not adhere to the Journal’s Instructions for Authors in numerous areas identified 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.
39. Personal communications and letters are included as references, against Springer guidelines (2, 16, 31, 50, 53, 56).
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).
41. References 81–83 are not published work but open-access referee comments.
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).
43. Personal blogs are included as references against Springer guidelines (71).
44. Reference 10 does not support the claim made in paragraph 2.1, as it does not refer to the factual matters being discussed.
45. Reference 60 is an example of a references that does not guide the reader to the salient discussion point.
46. Reference 83 points to a published research protocol of no relevance to the point being made.
47. Published links to references 37, 51 do not work.

¹¹ <https://deepmind.com/applied/deepmind-health/our-health-team/>