

## PERSPECTIVE OPEN

## The “inconvenient truth” about AI in healthcare

Trishan Panch<sup>1,2</sup>, Heather Mattie<sup>2,3</sup> and Leo Anthony Celi<sup>4,5</sup>

npj Digital Medicine (2019)2:77; https://doi.org/10.1038/s41746-019-0155-4

As the UK sits in painful deadlock over Brexit, it is important to remember that governments are regularly faced with crises, and their responses can create enduring benefit for future generations. Back in 1858, for example, the UK parliament was dealing with another messy crisis: “the great stink.” In a world before sanitation, the river Thames had become an open latrine, and as summer blossomed parliament was engulfed in a pestilential stench. £2.5 million (about £300 million in today’s money) was hastily approved to build a network of sewers throughout the capital.<sup>1</sup> This particular model of sanitation, developed by Bazalgette, was adopted by other cities around the world and the rest, as they say, is history. It is now unthinkable that a developed nation would not have sanitation infrastructure. However, back in 1858 the debate was whether sanitation infrastructure was worthy of investment and whether it was a public or private good. A similar debate has been simmering for some time regarding health data infrastructure, defined as the hardware and software to securely aggregate, store, process and transmit healthcare data. Is data infrastructure necessary for healthcare organizations and if so, is it the responsibility of individual healthcare organizations, of local health systems, or is it a public good?

In the 21st Century, the age of big data and artificial intelligence (AI), each healthcare organization has built its own data infrastructure to support its own needs, typically involving on-premises computing and storage.<sup>2,3</sup> Data is balkanized along organizational boundaries, severely constraining the ability to provide services to patients across a care continuum within one organization or across organizations. This situation evolved as individual organizations had to buy and maintain the costly hardware and software required for healthcare, and has been reinforced by vendor lock-in, most notably in electronic medical records (EMRs). With increasing cost pressure and policy imperatives to manage patients across and between care episodes, the need to aggregate data across and between departments within a healthcare organization and across disparate organizations has become apparent not only to realize the promise of AI but also to improve the efficiency of existing data intensive tasks such as any population level segmentation<sup>4</sup> and patient safety monitoring.<sup>5</sup>

The rapid explosion in AI has introduced the possibility of using aggregated healthcare data to produce powerful models that can automate diagnosis<sup>6</sup> and also enable an increasingly precision approach to medicine by tailoring treatments and targeting resources with maximum effectiveness in a timely and dynamic manner.<sup>7,8</sup>

However, “the inconvenient truth” is that at present the algorithms that feature prominently in research literature are in

fact not, for the most part, executable at the frontlines of clinical practice. This is for two reasons: first, these AI innovations by themselves do not re-engineer the incentives that support existing ways of working.<sup>2</sup> A complex web of ingrained political and economic factors as well as the proximal influence of medical practice norms and commercial interests determine the way healthcare is delivered. Simply adding AI applications to a fragmented system will not create sustainable change. Second, most healthcare organizations lack the data infrastructure required to collect the data needed to optimally train algorithms to (a) “fit” the local population and/or the local practice patterns, a requirement prior to deployment that is rarely highlighted by current AI publications, and (b) interrogate them for bias to guarantee that the algorithms perform consistently across patient cohorts, especially those who may not have been adequately represented in the training cohort.<sup>9</sup> For example, an algorithm trained on mostly Caucasian patients is not expected to have the same accuracy when applied to minorities.<sup>10</sup> In addition, such rigorous evaluation and re-calibration must continue after implementation to track and capture those patient demographics and practice patterns which inevitably change over time.<sup>11</sup> Some of these issues can be addressed through external validation, the importance of which is not unique to AI, and it is timely that existing standards for prediction model reporting are being updated specifically to incorporate standards applicable to this end.<sup>12</sup> In the United States, there are islands of aggregated healthcare data in the ICU,<sup>13</sup> and in the Veterans Administration.<sup>14</sup> These aggregated data sets have predictably catalyzed an acceleration in AI development; but without broader development of data infrastructure outside these islands it will not be possible to generalize these innovations.

Elsewhere in the economy, the development of cloud computing, secure high-performance general use data infrastructure and services available via the Internet (the “cloud”), has been a significant enabler for large and small technology companies alike, providing significantly lower fixed costs and higher performance as well as supporting the aforementioned opportunities for AI. Healthcare, with its abundance of data, is in theory well-poised to benefit from growth in cloud computing. The largest and arguably most valuable store of data in healthcare rests in EMRs. However, clinician satisfaction with EMRs remains low, resulting in variable completeness and quality of data entry, and interoperability between different providers remains elusive.<sup>11</sup> The typical lament of a harried clinician is still “why does my EMR still suck and why don’t all these systems just talk to each other?” Policy imperatives have attempted to address these dilemmas, however progress has been minimal. In spite of the widely touted benefits of “data

<sup>1</sup>Division of Health Policy and Management, Harvard T.H. Chan School of Public Health, Boston, MA, USA; <sup>2</sup>Wellframe Inc., Boston, MA, USA; <sup>3</sup>Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, USA; <sup>4</sup>Institute for Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, MA, USA and <sup>5</sup>Division of Pulmonary, Critical Care and Sleep Medicine, Beth Israel Deaconess Medical Center, Boston, MA, USA  
Correspondence: Leo Anthony Celi (lceli@mit.edu)

Received: 28 April 2019 Accepted: 26 July 2019

Published online: 16 August 2019

liberation”,<sup>15</sup> a sufficiently compelling use case has not been presented to overcome the vested interests maintaining the status quo and justify the significant upfront investment necessary to build data infrastructure. Furthermore, it is reasonable to suggest that such high-performance computing work has been and continues to be beyond the core competencies of either healthcare organizations or governments<sup>16</sup> and as such, policies have been formulated, but rarely, if ever, successfully implemented. It is now time to revisit these policy imperatives in light of the availability of secure, scalable data infrastructure available through cloud computing that makes the vision of interoperability realizable, at least in theory.

To realize this vision and to realize the potential of AI across health systems, more fundamental issues have to be addressed: who owns health data, who is responsible for it, and who can use it? Cloud computing alone will not answer these questions—public discourse and policy intervention will be needed. The specific path forward will depend on the degree of a social compact around healthcare itself as a public good, the tolerance to public private partnership, and crucially, the public’s trust in both governments and the private sector to treat their healthcare data with due care and attention in the face of both commercial and political perverse incentives.

In terms of the private sector these concerns are amplified as cloud computing is provided by a small number of large technology companies who have both significant market power and strong commercial interests outside of healthcare for which healthcare data might potentially be beneficial. Specific contracting instruments are needed to ensure that data sharing involves both necessary protection as well as, where relevant, fair material returns to healthcare organizations and the patients they serve.<sup>17</sup> In the absence of a general approach to contracting, high profile cases in this area have been corrosive to public trust.<sup>18,19</sup> Data privacy regulations like the European Union’s General Data Protection Regulation<sup>20</sup> (GDPR) or California’s Consumer Privacy Act<sup>21</sup> are necessary and well intentioned, though incur the risk of favoring well-resourced incumbents who are more able to meet the cost of regulatory compliance thereby possibly limiting the growth of smaller healthcare provider and technology organizations. Initiatives to give patients access to their healthcare data, including new proposals from the Center for Medicare and Medicaid Services<sup>22</sup> are welcome, and in fact it has long been argued that patients themselves should be the owners and guardians of their health data and subsequently consent to their data being used to develop AI solutions.<sup>16</sup> In this scenario, as in the current scenario where healthcare organizations are the de-facto owners and guardians of patient data generated in the health system alongside fledgling initiatives from prominent technology companies to share patient generated data back into the health system,<sup>23</sup> there exists the need for secure, high-performance data infrastructure to make use of this data for AI applications.

If the aforementioned issues are addressed, there are two possible routes to building the necessary data infrastructure to enable today’s clinical care and population health management and tomorrow’s AI enabled workflows. The first is an evolutionary path to creating generalized data infrastructure by building on existing impactful successes in the research domain such as the recent Science and Technology Research Infrastructure for Discovery, Experimentation and Sustainability (STRIDES) initiative from the National Institutes of Health<sup>24</sup> or MIMIC from the MIT Laboratory for Computational Physiology<sup>13</sup> to generate the momentum for change. Another, more revolutionary path would be for governments to mandate that all healthcare organizations store their clinical data in commercially available clouds. In either scenario, existing initiatives such as the Observational Medical Outcomes Partnership (OMOP<sup>25</sup>) and Fast Healthcare Interoperability Resources (FHIR) standard<sup>26</sup>

that create a common data schema for storage and transfer of healthcare data as well as AI enabled technology innovations to accelerate the migration of existing data<sup>27</sup> will accelerate progress and ensure that legacy data are included. There are several complex problems still to be solved including how to enable informed consent for data sharing, and how to protect confidentiality yet maintain data fidelity. However, the prevalent scenario for data infrastructure development will depend more on the socio-economic context of the health system in question rather than on technology.

A notable by-product of a move of clinical as well as research data to the cloud would be the erosion of market power of EMR providers. The status quo with proprietary data formats and local hosting of EMR databases favors incumbents who have strong financial incentives to maintain the status quo. Creation of health data infrastructure opens the door for innovation and competition within the private sector to fulfill the public aim of interoperable health data.

The potential of AI is well described, however in reality health systems are faced with a choice: to significantly downgrade the enthusiasm regarding the potential of AI in everyday clinical practice, or to resolve issues of data ownership and trust and invest in the data infrastructure to realize it. Now that the growth of cloud computing in the broader economy has bridged the computing gap, the opportunity exists to both transform population health and realize the potential of AI, if governments are willing to foster a productive resolution to issues of ownership of healthcare data through a process that necessarily transcends election cycles and overcomes or co-opts the vested interests that maintain the status quo—a tall order. Without this however, opportunities for AI in healthcare will remain just that—opportunities.

## ACKNOWLEDGEMENTS

L.A.C. is funded by the National Institute of Health through the NIBIB grant R01 EV017205.

## AUTHOR CONTRIBUTIONS

T.P. wrote the first draft. All authors contributed to both the subsequent drafting and critical revision of the manuscript.

## ADDITIONAL INFORMATION

**Competing interests:** The authors declare no competing interests.

**Publisher’s note:** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## REFERENCES

1. OpenLearn, The Open University. How London got its Victorian sewers. <https://www.open.edu/openlearn/science-maths-technology/engineering-technology/how-london-got-its-victorian-sewers> (2018).
2. Rajkumar, A., Dean, J. & Kohane, I. Machine learning in medicine. *N. Eng. J. Med.* **380**, 1347–1358 (2019).
3. Panch, T., Szolovits, P. & Atun, R. Artificial intelligence, machine learning and health systems. *J. Glob. Health* **8**, 020303 (2018).
4. Yan, S. et al. A systematic review of the clinical application of data-driven population segmentation analysis. *BMC Med. Res. Methodol.* **18**, 121 (2018).
5. Pronovost, P. J. et al. Paying the Piper: investing in infrastructure for patient safety. *Jt Comm. J. Qual. Patient Saf.* **34**, 342–348 (2008).
6. Keane, P. A. & Topol, E. J. With an eye to AI and autonomous diagnosis. *NPJ Digit Med* **1**, 40 (2018).
7. Shaban-Nejad, A., Michalowski, M. & Buckeridge, D. Health intelligence: how artificial intelligence transforms population and personalized health. *NPJ Digit Med.* **1**, 53 (2018).
8. Fogel, A. L. & Kvedar, J. C. Artificial intelligence powers digital medicine. *NPJ Digit Med.* **1**, 5 (2018).

9. Gijsberts, C. M. et al. Race/ethnic differences in the associations of the framingham risk factors with carotid IMT and cardiovascular events. *PLoS ONE* **10**, e0132321 (2015).
10. Hermansson, J. & Kahan, T. Systematic review of validity assessments of Framingham risk scoreresults in health economic modelling of lipid-modifying therapies in Europe. *Pharmacoeconomics* **36**, 205–213 (2017).
11. Fry, E., Schulte, F. Death by a thousand clicks: where electronic health records went wrong. <https://www.fortune.com/longform/medical-records> (2019).
12. Collins, G. S. & Moons, K. G. M. Reporting of artificial intelligence prediction models. *Lancet* **393**, 1577–1579 (2019).
13. Johnson, A. E. W. et al. MIMIC-III, a freely accessible critical care database. *Sci. Data* **3**, 160035 (2016).
14. U.S. Department of Veterans Affairs. VA Informatics and Computing Infrastructure. [https://www.hsrd.research.va.gov/for\\_researchers/vinci/cdw.cfm](https://www.hsrd.research.va.gov/for_researchers/vinci/cdw.cfm) (2014).
15. Park, T., Sivak, B. Health Datapalooza IV Tops Off a Huge Year in Health Data Liberation & Innovation. <https://obamawhitehouse.archives.gov/blog/2013/06/07/health-datapalooza-iv-tops-huge-year-health-data-liberation-innovation> (2013).
16. Mandl, K. D., Szolovits, P. & Kohane, I. Public standards and patients' control: how to keep electronic medical records accessible but private. *BMJ* **322**, 283–287 (2001).
17. Panch, T. et al. Artificial intelligence: opportunities and risks for public health. *Lancet Digit Health* **1**, e13–e14 (2019).
18. Ornstein, C., Thomas, K. Sloan Kettering's cozy deal with start-up ignites a new uproar. <https://www.nytimes.com/2018/09/20/health/memorial-sloan-kettering-cancer-paige-ai.html> (2018).
19. Revell, T. Google DeepMind's NHS data deal 'failed to comply' with law. <https://www.newscientist.com/article/2139395-google-deepminds-nhs-data-deal-failed-to-comply-with-law/> (2017).
20. European Parliament, Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32016R0679> (2016).
21. California Legislative Information. AB-375 Privacy: personal information: businesses. [https://leginfo.ca.gov/faces/billTextClient.xhtml?bill\\_id=201720180AB375](https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375) (2018).
22. Centers for Medicare & Medicaid Services. CMS Advances Interoperability & Patient Access to Health Data through New Proposals. <https://www.cms.gov/newsroom/fact-sheets/cms-advances-interoperability-patient-access-health-data-through-new-proposals> (2019).
23. Apple Newsroom. Apple announces effortless solution bringing health records to iPhone. <https://www.apple.com/newsroom/2018/01/apple-announces-effortless-solution-bringing-health-records-to-iphone/> (2018).
24. National Institutes of Health. STRIDES. <https://datascience.nih.gov/strides> (2019).
25. Observational Health Data Sciences and Informatics. OMOP Common Data Model. <https://www.ohdsi.org/data-standardization/the-common-data-model/> (2019).
26. HL7. Introducing HL7 FHIR. <https://www.hl7.org/fhir/summary.html> (2018).
27. Rajkomar, A. et al. Scalable and accurate deep learning with electronic health records. *NPJ Digit Med.* **1**, 18 (2018).



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019