

The clinical artificial intelligence department: a prerequisite for success

Christopher V. Cosgriff,¹ David J. Stone,^{2,3} Gary Weissman,^{4,5} Romain Pirracchio,⁶ Leo Anthony Celi ^{7,8}

To cite: Cosgriff CV, Stone DJ, Weissman G, *et al.* The clinical artificial intelligence department: a prerequisite for success. *BMJ Health Care Inform* 2020;**27**:e100183. doi:10.1136/bmjhci-2020-100183

Received 21 May 2020
Accepted 17 June 2020

Is artificial intelligence (AI) on track to usurp the electronic health record (EHR) as the most disappointing application of technology within medicine? The medical literature is increasingly populated with perspective pieces lauding the transformative nature of AI and forecasting an unforeseen disruption in the way we are practising.^{1,2} However, the reality of the available evidence increasingly leaves little room for optimism. As a result, there is a stark contrast between the lack of concrete penetration of AI in medical practice, and the expectations set by the presence of AI in our daily life.³ But medical AI need not follow the path of the EHR as a clinical tool that to many led to more workflow woes than it was intended to fix.⁴⁻⁶

As Atul Gawande so eloquently put, "... we've reached a point where people in the medical profession actively, viscerally, volubly hate their computers".⁷ If we are going to unavoidably add some disruption to workflow with AI, it should be as painless as possible to circumvent further, or perhaps even reduce, clinician burnout. We believe that this will require the combined and cross-disciplinary expertise of an organised and dedicated clinical AI department.

Historical precedents in radiology and laboratory medicine offer lessons for how to steward a new tool into the realm of safe and effective clinical use. Such accomplishments were due, in large part, to the gathering of relevant stakeholders under a single department. This approach ensured that the necessary clinical participants took the reins rather than ceding them to third-party developers. Thus, to secure AI's place in the annals of successful medical technologies, we propose the establishment of the first departments of clinical AI.

This proposal is deeply rooted in the history of American medicine. In 1890, the first X-ray image was generated at the University of Pennsylvania, although unbeknownst to

its creators, Goodspeed and Jennings.⁸ When the significance of this emerging technology was finally appreciated after the discovery of Roentgen rays, Goodspeed began informally collaborating with surgeons to deploy the technology clinically. This quickly led to the first division, and subsequently, department, of radiology. Under the auspices of this department, clinicians, researchers, engineers, managers and ethicists worked together on a shared mission to pioneer various technologies and methods that are intrinsic to the way medicine is practised today.

Within academic medicine, algorithms are currently developed in silos by researchers interested in the intersection of healthcare and machine learning. This has led to a panoply of published models trained on health data, yet only a handful have been prospectively evaluated on patients. In fact, when models have been prospectively evaluated on clinical outcomes, the results have frequently been unimpressive.⁹⁻¹² In contrast, the same multibillion-dollar technology companies that exploit patterns in our digital behaviour to sell advertising have now founded entire research programmes around health AI. We would argue that the lack of clinical results is the byproduct of a lack of coherence, leadership and vision. Hence, unless we change course, we should expect that AI deployment in healthcare will progress much the way the EHR revolution did before it, that is, mainly based on corporate and administrative benefits without requiring any demonstrable improvements in processes or outcomes for our patients or ourselves. As in the development of other areas that required full departmental support, the decision to establish a department of clinical AI has several logistical and policy implications.

First, leveraging the premises of AI to improve healthcare represents challenges in a number of ways such as implementation issues and applied policies. Therefore,



© Author(s) (or their employer(s)) 2020. Re-use permitted under CC BY-NC. No commercial re-use. See rights and permissions. Published by BMJ.

For numbered affiliations see end of article.

Correspondence to
Dr Leo Anthony Celi;
LCeli@mit.edu

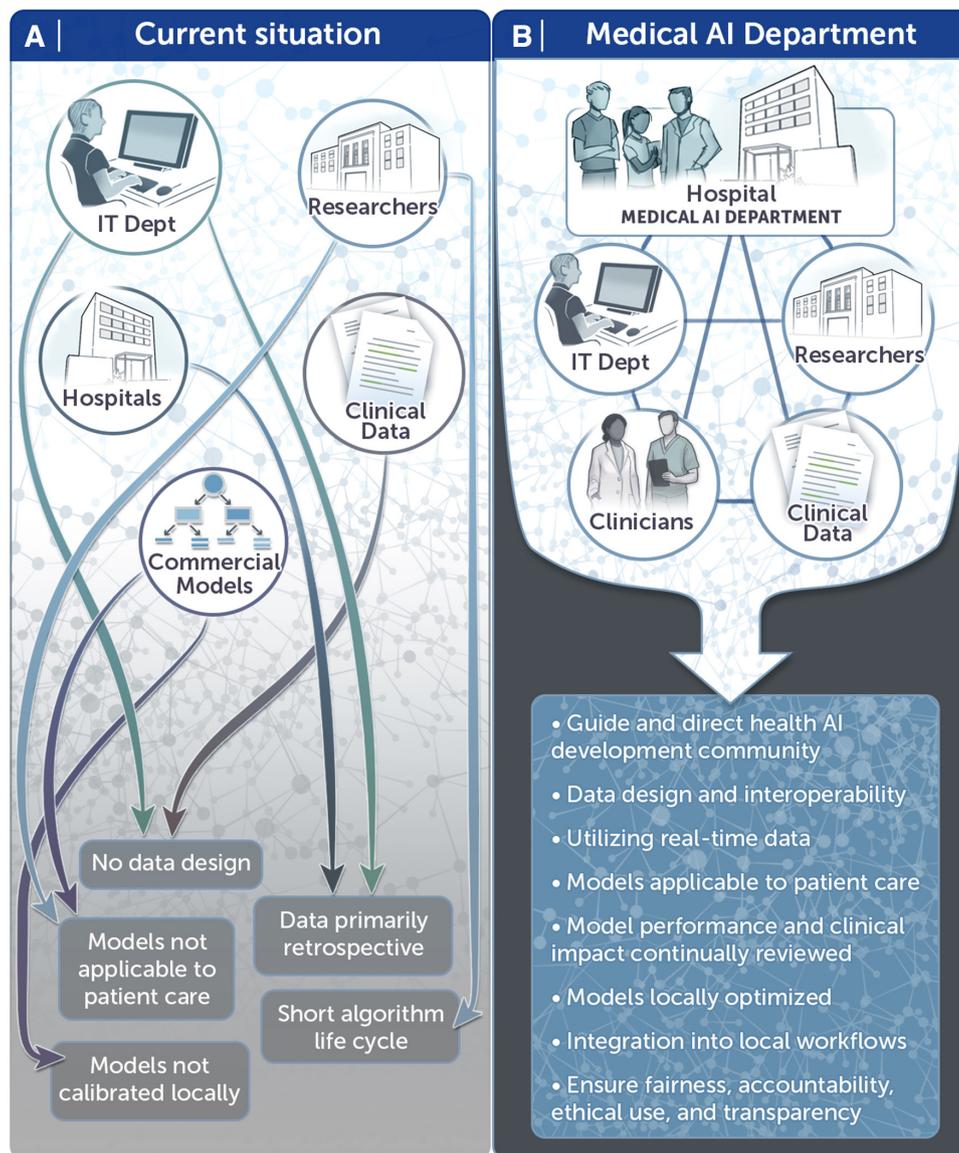


Figure 1 Medical artificial intelligence (AI) departments will provide the structure by which institutions can become AI Ready.

a chief mandate of department of clinical AI would be to make health centres *AI Ready*, a concept we illustrate in figure 1. These initiatives should lead to the development of models that will directly benefit the health of our patients, pioneer research that advances the field of clinical AI, focus on its integration into clinical workflows and foster educational programmes and fellowships to ensure we are training current practitioners as well as the next generation of leaders in this field. In addition to these traditional tripartite roles, AI departments should also play an essential role in the implementation, utilisation and enhancement of the infrastructures that underlie AI solutions. Central to this mission will be removing barriers to data access, and the proposed department would therefore assume partnered stewardship of the institution's data as part of its mandate. While the role of information technology specialists in maintaining a health system's computational infrastructure should not be subsumed, the department would be responsible for

integration, research and production databases that can support its broader mission. By centralising this role, we would finally overcome the chasms among ideas, development and effective deployment.

Second, these new departments will be instrumental as our country's financial and regulatory environments shift to acknowledge and incorporate AI's potential to improve care. The tasks and benefits involved may require a modified model of reimbursement such as that in place for laboratory tests. But as has been the case for corporate (eg, Amazon) AI, demonstrated improvements in clinical and financial outcomes could provide financial incentives to support the clinical use of AI and drive the increased deployment of predictive models. Market incentives will no doubt promote the proliferation of companies seeking to sell models to health systems. However, the need for model re-calibration precludes simply buying and deploying third-party models.¹³ Clinical AI departments will work to ensure that health systems are poised

for safe implementations that are tailored to their specific patient populations, and that the necessary data analytics will be readily available for negotiating with payers.

Third, the clinical utilisation of AI will require standardisation such as the establishment of best practice guidelines regarding workflow integration design, performance assessment and model fairness. Appropriate models should be tested on held-out current data to assess performance and safety, and only then prospectively evaluated first without, and then with, deployment in terms of accuracy and impact on clinical end points. From there, regular re-assessments of model calibration must occur to ensure the relationship between the inputs and the outputs has not changed, and to re-fit the model where it has. This requirement for re-assessment and recalibration in a specific clinical context has become evident when researchers have attempted to apply one site's data sets across institutional, system or geographic boundaries: AI applications can be sensitive to small input changes, and this potential fragility must be carefully and expertly monitored.¹⁴ While AI intrinsically manifests some degree of 'black box' characteristics, the functionality and reasons for its results should be as transparent and explicable as possible so that clinicians can incorporate these modalities into their workflows.¹⁵ As the introduction of information technology in medicine has heretofore demonstrated, successful technical solutions, from both software and hardware aspects, are different and much more difficult to accomplish when all decisions are not black and white, and lives are at stake.

Twenty years now into the 21st century, there is little question that AI will be a defining technology for the foreseeable future. We need visionary clinicians working with expert technical collaborators to establish the organisational structures requisite to translate technological progress into meaningful clinical outcomes. With the innumerable ways in which medicine could be improved, the hype around AI in healthcare will only be realised when the scattered champions of this movement emerge from their silos and begin formally working as a team under the same roof. Our patients are waiting for us to make use of these advances to improve their care, and every day wasted is a missed opportunity. Therefore, we ask—who will establish the first department of clinical AI?

Author affiliations

¹Department of Medicine, Hospital of the University of Pennsylvania, Philadelphia, Pennsylvania, USA

²Departments of Anesthesiology and Neurosurgery, University of Virginia, Charlottesville, Virginia, USA

³Center for Advanced Medical Analytics, University of Virginia School of Medicine, Charlottesville, Virginia, USA

⁴Division of Pulmonary and Critical Care Medicine, Hospital of the University of Pennsylvania, Philadelphia, Pennsylvania, USA

⁵Palliative and Advanced Illness Research Center, University of Pennsylvania, Philadelphia, Pennsylvania, USA

⁶Department of Anesthesia and Perioperative Care, University of California, San Francisco, San Francisco, California, USA

⁷Laboratory for Computational Physiology, Harvard-MIT Division of Health Sciences and Technology, Cambridge, Massachusetts, USA

⁸Division of Pulmonary Critical Care and Sleep Medicine, Beth Israel Deaconess Medical Center, Boston, Massachusetts, USA

Twitter Christopher V. Cosgriff @cosgriffc and Leo Anthony Celi @MITCriticalData

Contributors CVC produced the original draft under guidance from LAC. DS, GW and RP then gave input and provided edits. LAC oversaw the incorporation of these edits, led the discussions around the principle concepts and approved the final draft.

Funding LAC is funded by the National Institute of Health through NIBIB R01 EB017205.

Competing interests None declared.

Patient consent for publication Not required.

Provenance and peer review Not commissioned; externally peer reviewed.

Open access This is an open access article distributed in accordance with the Creative Commons Attribution Non Commercial (CC BY-NC 4.0) license, which permits others to distribute, remix, adapt, build upon this work non-commercially, and license their derivative works on different terms, provided the original work is properly cited, appropriate credit is given, any changes made indicated, and the use is non-commercial. See: <http://creativecommons.org/licenses/by-nc/4.0/>.

ORCID iD

Leo Anthony Celi <http://orcid.org/0000-0001-6712-6626>

REFERENCES

- Celi LA, Fine B, Stone DJ. An awakening in medicine: the partnership of humanity and intelligent machines. *Lancet Digit Health* 2019;1:e255–7.
- Topol EJ. High-Performance medicine: the convergence of human and artificial intelligence. *Nat Med* 2019;25:44–56.
- Panch T, Mattie H, Celi LA. The “inconvenient truth” about AI in healthcare. *NPJ Digit Med* 2019;2:77.
- Lau F, Price M, Boyd J, et al. Impact of electronic medical record on physician practice in office settings: a systematic review. *BMC Med Inform Decis Mak* 2012;12:10.
- Adane K, Gizachew M, Kendie S. The role of medical data in efficient patient care delivery: a review. *Risk Manag Healthc Policy* 2019;12:67–73.
- Sinsky C, Colligan L, Li L, et al. Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties. *Ann Intern Med* 2016;165:753–60.
- Gawande A. Why doctors hate their computers. the new Yorker 2018.
- Radiology department history, 2020. Available: <https://www.pennmedicine.org/departments-and-centers/department-of-radiology/about-penn-radiology/department-history> [Accessed 2/24, 2020].
- Baillie CA, VanZandbergen C, Tait G, et al. The readmission risk flag: using the electronic health record to automatically identify patients at risk for 30-day readmission. *J Hosp Med* 2013;8:689–95.
- Bedoya AD, Clement ME, Phelan M, et al. Minimal impact of implemented early warning score and best practice alert for patient deterioration. *Crit Care Med* 2019;47:49–55.
- Courtright KR, Chivers C, Becker M, et al. Electronic health record mortality prediction model for targeted palliative care among hospitalized medical patients: a pilot quasi-experimental study. *J Gen Intern Med* 2019;34:1841–7.
- Downing NL, Rolnick J, Poole SF, et al. Electronic health record-based clinical decision support alert for severe sepsis: a randomised evaluation. *BMJ Qual Saf* 2019;28:762–8.
- Davis SE, Lasko TA, Chen G, et al. Calibration drift among regression and machine learning models for hospital mortality. *AMIA Annu Symp Proc* 2017;2017:625–34.
- Shah ND, Steyerberg EW, Kent DM. Big data and predictive analytics: recalibrating expectations. *JAMA* 2018;320:27–28.
- Gennatas ED, Friedman JH, Ungar LH, et al. Expert-augmented machine learning. *Proc Natl Acad Sci U S A* 2020;117:4571–7.