Tom J. Pollard
Research Scientist
Massachusetts Institute of Technology
Cambridge, USA

tpollard@mit.edu

@tompollard

Leo Anthony Celi
Intensivist and Principal Research
Scientist
Massachusetts Institute of Technology
Cambridge, USA

Beth Israel Deaconess Medical Center
Boston, USA

lceli@mit.edu

Enabling Machine Learning in Critical Care

Critical care units are home to some of the most sophisticated patient technology within hospitals. In parallel, the field of machine learning is advancing rapidly and increasingly touching our lives. To facilitate the adoption of machine learning approaches in critical care, we must become better at sharing and integrating data. Greater emphasis on collaboration—outside the traditional "multidisciplinary" realm and into the engineering, mathematical, and computer sciences—will help us to achieve this. Meanwhile, those at the forefront of the health data revolution must earn and maintain society's trust and demonstrate that data sharing and reuse is a necessary step to improve patient care.

Critical care units are home to some of the most sophisticated patient technology within hospitals. Devices such as vital sign monitors, mechanical ventilators and dialysis machines, to name a few, are used to support patients whose bodies need time to recover and repair. Data, a by-product of technology, contains information with the potential to improve our understanding of health and disease. Outside critical care units, we are seeing increasing adoption of digital health systems in place of the paper-based systems of the past.

In parallel, the field of machine learning is advancing rapidly and increasingly touching our lives. Algorithms built upon large volumes of data have beaten world champions in the complex game of Go, driven cars on the open roads, and matched doctors in diagnosing skin cancers (Esteva et al. 2017) and diabetic eye conditions (Gulshan et al. 2016).

With these advances, are we now on the cusp of transformed, algorithm-driven care? Not yet, it would seem. In critical care, and medicine as a whole, the massive troves of data needed to pour into machine learning algorithms are difficult to find and access. For the development of machine learning-based approaches to care, data must first be properly

archived, integrated across data sources and shared for reuse. Most hospitals have a long way to go in this regard. Data is still treated as a currency for clinical researchers to build careers, harming efforts to combine data to an extent that can fuel progress. The absence of incentives for data integration and data sharing has hindered us from understanding health and disease in new ways from analysing 'real-world' data collected in the process of care. Too many of our study models still rely on either absurdly small datasets or on large-scale but coarse registry data devoid of the rich details that are required to unleash the value of machine learning.

With funding from the National Institutes of Health, the MIT Laboratory for Computational Physiology (MIT-LCP) develops and maintains the publicly available Medical Information Mart in Intensive Care (MIMIC), a database of patients admitted to the intensive care units of a large teaching hospital of the Harvard Medical School. The current version, MIMIC-III, contains data associated with 53,423 distinct hospital admissions for adult patients admitted to critical care units between 2001 and 2012 (Johnson et al. 2016). Data include vital signs, medications, laboratory measurements, charted notes,

billing codes and out-of-hospital survival data. With over 4,000 users in academia and industry from over 30 countries. MIMIC-III has been used for clinical research studies, exploratory and validation analyses performed by pharmaceutical and medical technology companies, as well as university, conference and online courses, tutorials and workshops. At least 24 courses in the United States alone use the database to teach concepts in machine learning, medical informatics and biostatistics.

Spurred by the success of MIMIC, MIT-LCP recently released the eICU Collaborative Research Database in collaboration with Philips Healthcare, comprising de-identified health data associated with over 200,000 critical care admissions from patients admitted to >200 hospitals throughout the United States between 2014-2015. The dataset is itself a subset drawn from a pool of nearly 3 million ICU admissions and provides a unique and invaluable resource for health research and education. Like MIMIC-III, the database includes detailed clinical data such as vital signs, pharmacy medication orders, laboratory results and severity of illness scores, giving researchers comprehensive insights into patient care. The database presents an opportunity to assess

heterogeneity in treatments, patient populations and settings, which was not possible with large single-site research databases such as MIMIC-III.

To encourage research transparency and collaboration around the databases, MIT-LCP creates and supports collaboratively maintained, open code repositories. For example, the MIMIC Code Repository is a centralised code base for generating reproducible studies on the MIMIC-III dataset (MIMIC in press). All code is made open source under an MIT License and is freely available online (**github.com/MIT-LCP/mimic-code**). Executable documents reproduce published studies end-to-end, providing a template for future researchers to replicate. The repository's issue tracker enables community discussion about the data and concepts, allowing users to collaboratively improve the resource. Consistent application of the same code for underlying concepts is a key step in ensuring research studies in critical care are comparable and reproducible.

But it is not enough to create high-resolution databases to propel the application of machine learning in critical care medicine. The most daunting challenge, as with most complex problems of our time, is the lack of collaboration across the key players who represent the disciplines and who continue to work in their own silos (Celi et al. 2016a). To this end, MIT-LCP organises critical care datathons, a portmanteau of data + hackathon, focusing the application of the hackathon model on data analytics (Aboab et al. 2016; Celi et al. 2016b). The goal of these datathons is to unify clinical experts, data scientists, statisticians and those with domain-specific knowledge to brainstorm ideas and contribute clinically relevant research.

Datathons

During the past year, MIT-LCP has helped to organise and host several international datathons to gain new insights from routinely collected patient data. The events were held in Beijing in October 2016 (funded by the People's Liberation Army General Hospital), London in December 2016 (funded by the UK Intensive Care Society and the MIT International Science and Technology Initiatives Global Seed Fund), Melbourne in March 2017 (funded by the Australian and New Zealand Intensive Care Society, Alfred Hospital and Philips Healthcare), Sao Paulo in May 2017 (funded by the Hospital Israelita Albert Einstein and the MIT Brazil Seed Fund) and Singapore in July 2017 (funded by Merck Sharpe & Dohme and the National University of Singapore).

> ❝ **power of freely accessible data repositories for crowdsourcing knowledge creation and validation** ❞

Bringing together clinicians (including nurses, pharmacists and therapists) and data scientists at datathons serves to demonstrate the value of each other's expertise. We have found these events to be an important tool in demonstrating the power of freely accessible data repositories for crowdsourcing knowledge creation and validation. Perhaps most importantly, they have generated interest amongst participants to contribute new, high-resolution critical care databases to the research community, supplementing existing resources such as MIMIC-III and the eICU Collaborative Research Database.

For a more in-depth tutorial in secondary analysis of health records, MIT-LCP teaches a fall course at MIT on Collaborative Data Science in Medicine. Students learn the basics of research using routinely collected health data, including data extraction, processing and analysis, and acquire skills from a diverse set of fields including epidemiology, databases, statistics and machine learning. In addition, students team up with Boston-area clinicians for a course project using either MIMIC-III or the eICU Collaborative Research Database to produce novel research, often leading to publication in a clinical journal. An open access textbook accompanies the course and has been downloaded more than 90,000 times since its publication in October 2016 (Celi et al. 2016c).

Change is on the horizon with growing interest in digital health, the application of machine learning on health data and the dawn of artificial intelligence to assist healthcare providers and patients. In the United States alone, venture capital investments in digital health grew at an annual rate of 30% from 2011 to 2016 and last year totalled US$4.2 billion (Tecco 2017). More large companies, from Apple, Microsoft, IBM, Alphabet (Google's parent), Merck, Aetna, to UnitedHealth Group (through its Optum subsidiary), are investing in digital health products (Swanson 2016; CB Insights 2016; Bergen 2015). But for a true health data revolution to occur in healthcare, the environment—the technology, the policies, and the people, both providers and patients—needs to be supportive of change.

Within hospitals we will need to begin adapting culture and education to prepare for the changes to come. Greater emphasis on collaboration – outside the traditional "multidisciplinary" realm and into the engineering, mathematical, and computer sciences – will help us to create the right environment for a move towards algorithm-driven care. Meanwhile, those at the forefront of the health data revolution must earn and maintain society's trust and demonstrate that data sharing and reuse is a necessary step to improve patient care.

Conflict of Interest
The authors have received funding from Philips and Merck. ∎

References

Aboab J, Celi LA, Charlton P et al. (2016) A datathon model to support cross-disciplinary collaboration. Sci Transl Med, 8(333): 333ps8.

Celi LA, Davidzon G, Johnson AE et al. (2016) Bridging the health data divide. J Med Internet Res, 18(12): e325.

Esteva A, Kuprel B, Novoa RA et al. (2017) Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542: 115-8.

Gulshan V, Peng L, Coram M et al. (2016) Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA, 316(22): 2402-10.

Johnson AEW, Pollard TJ, Shen L et al. (2016) MIMIC-III, a freely accessible critical care database. Scientific Data, 3: 160035.

Johnson AEW, Stone DJ, Celi LA et al. (2017) The MIMIC code repository: enabling reproducibility in critical care research. JAMIA: Journal of the American Medical Informatics Association. In Press.

For full references, please email editorial@icu management.org or visit https://iii.