Machines Learn to Read Hospital Records, Will Doctor's Handwriting Be Next?

By Chuck Dinerstein — June 20, 2018

Courtesy of Wellcome Trust [1]

Patient records are unruly; they consist of numbers, images, and text. Electronic health records (EHR) are more frequently digital "Xerox" versions of the patient chart's you'd associate with the kindly physicians depicted by Norman Rockwell. As a result of the jumble of data types and formats, data mining to identify predictive analytics initially require a careful selection of variables of interest which are abstracted from medical records and put in a machine-readable form. Data preparation often takes longer than the analysis itself, and because variables are determined before analysis, these datasets are one trick ponies unable to be used to investigate other issues. They are rate-limiting bottlenecks for the proponents of Big Data, in their quest to apply machine learning to medical problems. A paper in Nature Partner Journals/Digital Medicine provides a proof of concept that the bottleneck has been removed.

The researchers were given access to 216,221 de-identified hospital records from two major teaching hospitals, nearly 46 billion individual pieces of data. Algorithms for natural language processing were unleashed on the data automating the conversion of unruly hospital records into a machine-readable format. [1] The resulting data were then utilized to develop predictive models of patient mortality, readmissions, length of stay and discharge diagnosis. The study was retrospective, but because hospital records are time-stamped, the data could be ordered along a timeline to simulate real-time analysis

- Inpatient mortality was predicted correctly 95% of the time compared to 85% of the time using a commercially available algorithm, and the information was available 24 hours earlier.
- Hospital readmissions were predicted 77% of the time compared to 70% for the hospitals current systems

- Length of stay was predicted correctly 86% of the time compared with 76% for the hospitals current systems
- The entire set of discharge diagnosis was predicted correctly 87% of the time, by default the hospital codes at 100%

Perhaps the most significant achievement of the study was demonstrating that hospital records, which currently lack the ability for one software system to talk with another, can be reconfigured for use in machine learning by the machine's themselves. The bottleneck of data preparation has been broken with an automated and therefore scalable approach. While this is good news for Big Data advocates, it also serves as a hopeful sign to beleaguered physicians and nurses who spend increasing amounts of time entering data into records which cannot be used outside the hospital or their vendor's EHR platform.

An equally important benefit of the researchers model is that it could display that data points on a timeline, something our current hospital's systems do poorly and highlights those data points used by the algorithm in making its predictions – in other words, the software "explains" its decision. The lack of explanation by deep learning methods, so-called black box medicine, has been another rate-limiting step in the acceptance of these systems into clinical practice.

The study does have limitations. It is retrospective. While the researchers took pains to look at various domains of care, the clinical utility is limited. Mortality predictors are used to augment physician judgment, not replace it. Length of stay predictors and discharge diagnosis are more useful to administrators than in clinical care. Predictions of readmissions are clinically relevant because if we can understand the causes of readmission, we may be able to reduce the incidence (which typically runs at 20% of hospitalizations) improving patient care and saving the health system a great deal of money. But look at those percentages for the readmission predictor, 77%, was the lowest of the bunch. It may be that big data cannot supply all the answers. As the authors state, "…although it is widely believed that accurate predictions can be used to improve care, this is not a foregone conclusion, and prospective trials are needed to demonstrate this." That said, this was a big step forward in developing systems that can augment and assist physicians in their judgment and care.

[1] Fast Healthcare Interoperability Resources or FHIR – a set of standards for encoding health data. The current electronic health records foisted on the American public and its physicians, at great expense, use their proprietary formats. The federal government is now providing more funding to vendors to develop information exchanges to correct for this design flaw.

Conflict of Interest Disclosure: The author is a proud family member of an Alphabet employee.