Editor: **Daniel B. Neill,** H.J. Heinz III College, Carnegie Mellon University, neill@cs.cmu.edu

Using Artificial Intelligence to Improve Hospital Inpatient Care

Daniel B. Neill, *Carnegie Mellon University*

The US healthcare system faces many challenges, including skyrocketing costs, high rates of drug-resistant and hospital-acquired infections, and failures of care delivery leading to preventable adverse health events. Overtreatment, poor execution of care, and failure to adopt best practices for preventive care and patient safety have huge and directly measurable impacts on both healthcare costs and patient outcomes.[1,2] On the other hand, both the increasing availability of electronic health data and the ongoing development of methodological approaches to analyze these data suggest the potential for the use of artificial intelligence and machine learning methods to improve the quality and lower the cost of patient care.

Electronic health records (EHR) have become more available due to the guidelines of the Health Information Technology for Economic and Clinical Health (HITECH) Act, which offers incentives to healthcare providers to adopt EHR to advance clinical processes and improve outcomes. Meanwhile, health insurance providers and nonprofits such as the Health Care Cost Institute have committed to providing health insurance claims data with the goal of reducing costs while improving the quality and availability of coverage. Such sources provide detailed, time-stamped, and highly multivariate data for a large patient population, enabling the use of AI techniques to connect care practices and outcomes. However, the data's size and complexity—as well as the variability in content and format between different providers, data types, and care settings—create huge challenges. Additionally, the potential danger of the violation of patients' privacy has significant moral and legal ramifications, requiring extreme care in the use of health data.

Recent Advances in AI for Patient Care

Clinical decision support systems (CDSS) were one of the first successful applications of AI, focusing primarily on the diagnosis of a patient's condition given his symptoms and demographic information. Work on CDSS for medical diagnosis began in the early 1970s with Mycin[3]—a rule-based expert system for identifying bacteria that cause severe infections and recommending antibiotics to treat these infections. David Heckerman and his colleagues[4] developed Pathfinder, which used Bayesian networks (a graphical model that encodes probabilistic relationships among variables of interest) to help pathologists more accurately diagnose lymph-node diseases. AI has also been useful for computer-aided detection of conspicuous structures (such as tumors or polyps) in medical images. Such approaches assist in the screening of mammography images,[5] as well as the diagnosis of various forms of cancer, coronary artery disease,[6] and congenital heart defects.

More recent advances in machine learning and AI build predictive models and make real-time inferences from a large patient population for purposes including alerts,[7] stratifying risk,[8,9] and predicting the length of stay.[10] Several of these approaches focus on critical care, using physiological data that are routinely recorded in intensive care units. For example, Ying Zhang and Peter Szolovits[7] developed an intensive care monitoring system to model individual patients' vital signs and produce patient-specific models and alarm thresholds. Decision trees and neural networks were used to generate binary classifiers of the patient state and determine when to issue an alarm. Suchi Saria and her colleagues[8] developed a physiological assessment score for preterm newborns, using time-series data captured from the newborn's

1541-1672/13/$31.00 © 2013 IEEE
Published by the IEEE Computer Society

first three hours of life, and a hierarchical Bayesian model, the *time-series topic model*. This approach lets healthcare providers accurately estimate the probability of an infant's risk of such severe problems as infections and cardiopulmonary complications. Interestingly, Saria and her colleagues noted that physiological parameters such as short-term variability in respiratory and heart rates had greater predictive power than invasive laboratory studies, suggesting the potential for new and less-invasive neonatal care practices.[8] Scott Levin and his colleagues[10] also focused on intensive care patients, but extracted provider orders (laboratory tests, procedures, and medications) from the hospital's computerized order entry system. They then used a logistic regression model to predict the length of stay and demonstrated significant improvements in prediction accuracy. Jenna Wiens and her colleagues[9] estimated patients' risk of hospital-acquired illness (specifically, the risk of infection by *Clostridium difficile*) by extracting more than 10,000 variables for each day of each hospital admission, using a support vector machine (SVM) to produce a time series of daily risk scores and applying various approaches (including Hidden Markov Models and SVMs) for time-series classification. They demonstrated that the use of temporal information representing the evolution of each patient's health state (and therefore, risk of infection) over time leads to improved classification accuracy, as compared to classifiers that only consider the patient's current state.

Answering General Questions for Patient Care

Although these advances demonstrate the potential of AI and machine learning to improve patient care, nearly all of the aforementioned techniques focus on the prediction problem (either classification for predicting a discrete-valued attribute or regression for predicting a real-valued attribute). As such, they are typically limited in scope to specific diseases or diagnoses or only applicable to a small subset of the patient population. Perhaps the next great challenge for AI in healthcare is to develop approaches that can be applied to the entire population of patients, monitoring huge quantities of data to automatically detect problems and threats to patient safety (including patterns of suboptimal care, as well as outbreaks of hospital-acquired illness), and to discover new best practices of patient care.

Two very different AI approaches, each having great potential for addressing these challenges, are respectively based on question-answering (QA) and on large-scale anomalous pattern detection. Continued advances in general QA led to the design of the DeepQA architecture by IBM Research,[11] in collaboration with Carnegie Mellon University, and the well-publicized victory of IBM's Watson system over human champions on the well-known TV quiz show, *Jeopardy*. IBM is currently partnering with the Memorial Sloan-Kettering Cancer Center to enable patient-specific diagnostic test and treatment recommendations for various types of cancer.[12] Many of Watson's features that led to its Jeopardy Challenge victory are also relevant to the healthcare domain, including its ability to incorporate huge volumes of unstructured text data (patients' electronic health records, medical literature, and so on), respond to natural language queries, provide probabilistic reasoning to assist clinicians in making evidence-based decisions, and improve its performance through learning from user interaction.[12] Other QA systems such as the Semantic Research Assistant (SRA)[13] focus specifically on the medical domain. SRA extends the large-scale knowledge base Cyc to answer ad hoc queries by physicians, justifying each answer with general medical facts, expert-articulated rules, and specific patient records. SRA is currently used by the Cleveland Clinic to answer clinical research queries involving cardiothoracic surgery, cardiac catheterization, and percutaneous coronary intervention, and has reduced the typical time to produce a satisfactory answer to such queries from weeks to minutes.[13]

Anomalous Pattern Detection for Patient Care

Another new approach that might improve patient care focuses on statistical machine learning methods for detecting anomalous patterns in massive quantities of healthcare data. We recently developed a variety of machine learning methods based on *fast subset scanning*[14,15] to detect patterns in massive datasets, efficiently identifying subsets of data records and attributes that are collectively anomalous or that maximize some measure of interest, such as a likelihood ratio statistic. In the patient care setting, our primary focus is to detect anomalous patterns of care that influence patient outcomes. Consider the natural variation in care practices between different groups and different clinicians. For example, when presented with a patient with severe breathing difficulties, clinicians might choose to administer different types and dosages of medications, use different criteria to decide whether to place the patient on a ventilator, and so on. Similarly, hospital staff members have different care practices (such as hand-washing and isolation precautions) and adherence to physician orders. This variation in type and quality of care can have huge impacts on patient outcomes, such as mortality and morbidity rates, hospital re-admissions, and hospital-acquired infections.

We are currently developing a system that will automatically detect substantial variations in care between groups that have significant impacts on patient outcomes. These impacts can either be negative (systematic errors, for example), in which case we can detect and correct these suboptimal patterns of care, or positive. In the latter case, our system will have discovered a new potential best practice, which can then be investigated further, and if appropriate, shared with other groups. As a simple concrete example, we may discover that, in our data, certain patients with high blood pressure experience fewer complications if given drugs X and Y one hour instead of 30 minutes before surgery. By integrating health insurance claims with patient data, and treating cost of care as an additional outcome to be optimized, we hope to identify care practices that are cost-effective and improve outcomes.

We view this role of the system as focusing on hypothesis generation. The identified patterns represent alternative care practices that can be rigorously evaluated for potential use by the medical community. Such a system would ideally integrate huge amounts of data of multiple types, from multiple healthcare providers, in multiple care settings. Even within a single hospital, there might still be sufficient variation in care to discover new best practices. At a minimum, real-time detection of anomalous patterns should enable early warning systems for outbreaks of hospital-acquired illness, systematic errors in care (for example, poor hand-washing practices), or patterns of adverse events.

We believe that recent advances in fast and scalable detection methods are an important first step towards identifying and optimizing patterns of patient care. For example, the recently proposed *fast generalized subset scan* (FGSS)[15] can identify self-similar subsets of data records for which some subset of attributes is anomalous; the *multidimensional subset scan* (MD-Scan)[16] and *disjunctive anomaly detector* (DAD)[17] identify combinations of attribute values for which the corresponding number of data records is significantly higher or lower than expected. All of these methods incorporate the *linear-time subset scanning* property,[14] which enables rapid identification of the most anomalous subset, into an iterative algorithm. FGSS iterates between optimizing over subsets of records (for the given subset of attributes) and optimizing over subsets of attributes (for the given subset of records), while MD-Scan and DAD iterate over each attribute, optimizing over subsets of values for that attribute conditioned on the current subsets of values for all other attributes.

Although these techniques enable accurate and efficient anomalous pattern detection in general datasets, several important challenges remain for their application to identifying anomalous patient care patterns. First, even though any patterns identified by the system would undergo rigorous evaluation by the medical community before being directly applied to patient care, a practical and usable system must assist this process by focusing attention on those patterns that are most likely to be medically relevant. We wish to identify patient care patterns that are not just correlated with outcomes, but are likely to be causal factors influencing those outcomes.

For example, if we observe that patients in a given hospital bed have higher rates of hospital-acquired infection, we would like to distinguish the hypothesis that the given bed causes illness from the alternative explanation that more severely ill patients are placed in that bed (because it is right next to the nurse's station, perhaps) and such patients are also more susceptible to hospital-acquired infection. One possible solution is to integrate anomalous pattern-detection with econometric techniques such as propensity-score matching[18] or with machine learning approaches to causal structure discovery.[19]

A second set of challenges is posed by the use of massive quantities of streaming data for real-time monitoring of patient health and safety. Current techniques might be insufficient to analyze such massive quantities of data, and thus techniques for dimensionality reduction, clustering, aggregation, and data summarization might be useful.

Third, the extension of anomalous pattern detection beyond the hospital setting to incorporate data from outpatient settings, such as preventive care and management of chronic disease, creates further challenges, including patient non-compliance to prescribed treatments and preventive care. Additionally, the huge increase in variability between patients' behaviors and environment in the outpatient setting—as well as the much longer time scale—present challenges in attributing differences in outcomes to the greater variety of potential causal factors.

Although the primary roles of AI in patient care to date have mainly been in patient diagnosis and image analysis, the future holds great potential for applying AI to improve many aspects of the patient care process. Some example include personalizing treatments to maximize efficacy while minimizing side effects, recommending appropriate sequences of diagnostic tests, monitoring the patient population's health and safety, and discovering new medical knowledge that can directly impact the quality of care. Great challenges remain due to the health data's size and complexity, but the AI community is well on its way to meeting these challenges by developing new pattern detection techniques, scalable algorithms, and novel approaches that use massive quantities of health data to answer general questions.

Acknowledgments

This work was partially supported by the US National Science Foundation, grants IIS-0916345, IIS-0911032, and IIS-0953330, and a University of Pittsburgh Medical Center (UPMC) Healthcare Technology Innovation grant. All views are solely those of the author, and have not been endorsed by NSF or UPMC.

References

1. L.T. Kohn, J.M. Corrigan, and M.S. Donaldson, eds., *To Err Is Human: Building a Safer Health System*, Inst. of Medicine/Nat'l Academy Press, 1999.
2. N.C. Lallemand, "Health Policy Brief: Reducing Waste in Health Care," *Health Affairs*, 13 Dec. 2012.
3. B.G. Buchanan and E.H. Shortliffe, *Rule-Based Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, 1984.
4. D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani, "Toward Normative Expert Systems: The Pathfinder Project," *Methods of Information in Medicine*, vol. 31, no. 2, July 1992, pp. 90–105.
5. J.J. Fenton et al., "Influence of Computer-Aided Detection on Performance of Screening Mammography," *New England J. Medicine*, vol. 356, no. 14, 2007, pp. 1399–1409.
6. K.W. Kang et al., "Feasibility of an Automatic Computer-Assisted Algorithm for the Detection of Significant Coronary Artery Disease in Patients Presenting with Acute Chest Pain," *European J. Radiology*, vol. 81, no. 4, 2012, pp. 640–646.
7. Y. Zhang and P. Szolovits, "Patient-Specific Learning in Real Time for Adaptive Monitoring in Critical Care," *J. Biomedical Informatics*, vol. 41, no. 3, 2008, pp. 452–460.
8. S. Saria et al., "Integration of Early Physiological Responses Predicts Later Illness Severity in Preterm Infants," *Science Translational Medicine*, vol. 2, no. 48, 2010, pp. 48–65.
9. J. Wiens, J.V. Guttag, and E. Horvitz, "Patient Risk Stratification for Hospital-Associated C. diff as a Time-Series Classification Task," *Proc. Advances in Neural Information Systems*, Neural Information Processing Systems (NIPS) Foundation, vol. 25, 2012, pp. 247–255.
10. S.R. Levin et al., "Real-Time Forecasting of Pediatric Intensive Care Unit Length of Stay Using Computerized Provider Orders," *Critical Care Medicine*, vol. 40, no. 11, 2012, pp. 3058–3064.
11. D. Ferrucci et al., "Building Watson: An Overview of the DeepQA Project," *AI Magazine*, vol. 31, no. 3, 2010, pp. 59–79.
12. M.S. Kohn and P.C. Skarulis, "IBM Watson Delivers New Insights for Treatment and Diagnosis," Digital Health Conf., presentation, 2012.
13. D. Lenat et al., "Cyc to Answer Clinical Researchers' Ad Hoc Queries," *AI Magazine*, vol. 31, no. 3, 2010, pp. 13–32.
14. D.B. Neill, "Fast Subset Scan for Spatial Pattern Detection," *J. Royal Statistical Soc.* (Series B: Statistical Methodology), vol. 74, no. 2, 2012, pp. 337–360.
15. E. McFowland III, S. Speakman, and D.B. Neill, "Fast Generalized Subset Scan for Anomalous Pattern Detection," *J. Machine Learning Research*, 2013, forthcoming.
16. T. Kumar and D.B. Neill, "Fast Multidimensional Subset Scan for Outbreak Detection and Characterization," *Online J. Public Health Informatics*, vol. 5, no. 1, 2013.
17. M. Sabhnani, A. Dubrawski, and J. Schneider, "Searching for Complex Patterns Using Disjunctive Anomaly Detection," *Online J. Public Health Informatics*, vol. 5, no. 1, 2013.
18. P.R. Rosenbaum and D.B. Rubin, "Reducing Bias in Observational Studies Using Subclassification on the Propensity Score," *J. Am. Statistical Assoc.*, vol. 79, 1984, pp. 516–524.
19. P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search,* Lecture Notes in Statistics 81, Springer, 1993.

Daniel B. Neill is the Dean's Career Development Professor and associate professor of information systems at Carnegie Mellon University's Heinz College, where he directs the Event and Pattern Detection Laboratory. His research interests include machine learning, data mining, and event detection in massive datasets. Neill has a PhD in computer science from Carnegie Mellon University. Contact him at neill@cs.cmu.edu.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*