In: E.H. Shortliffe, J.J. Cimino, et al., *Biomedical Informatics: Computer Applications in Health Care and Biomedicine*, Springer, 2020 (forthcoming).

The Future of Biomedical Informatics: Bottlenecks and Opportunities

Eric Horvitz, MD, PhD

I see the rich, interdisciplinary field of biomedical informatics as the gateway to the future of healthcare. The concepts, methods, rich history of contributions, and the aspirations of biomedical informatics define key opportunities ahead in biomedicine—and shine light on the path to achieving true evidence-based healthcare.

Progress with influences of biomedical informatics on healthcare over the last three decades has been slower than I had hoped. However, I remain optimistic about a forthcoming biomedical informatics revolution, made possible by a confluence of advances across industry and academia. Such a revolution will accelerate discovery in biomedicine, enhance the quality of healthcare, and reduce the costs of healthcare delivery.

From my perch as an investigator and director of a worldwide system of computer science research labs, I view key opportunities ahead as hinging on the dual of (1) addressing the often underappreciated bottleneck of *translation*—moving biomedical informatics principles and prototypes into real-world practice, and (2) making progress on persisting challenges in principles and applications of artificial intelligence (AI). I am optimistic that we will make progress on both fronts and that there will be synergies among these advances.

On challenges of translation, I believe that the difficulties of transitioning ideas and implementations from academic and industry research centers into the open world of medical practice have been widely underappreciated. Numerous factors are at play, including poor understanding of how computing solutions can assist with the tasks and day-to-day needs of healthcare practitioners and patients, inadequate appreciation of the needs and difficulties of developing site-specific solutions, poor compute infrastructure, and a constellation of challenges with human factors, including entrenched patterns of practice and difficulties of integrating new capabilities and services into existing clinical workflows.

Multiple advances coming with the march of computer science will help to address challenges of translating ideas and methods that have been nurtured by biomedical informaticians for decades. At the base level, such advances include ongoing leaps in computing power and in storage, but also key innovations with computing principles and methods in such subdisciplines as databases, programming languages, security and privacy, human-computer interaction, visualization, and sensing and ubiquitous computing.

Faster and more effective translation of ideas and methods from biomedical informatics will also be enabled by jumps in the quality of available computing tools and infrastructure. Increases in the power and ease-of-use of cloud computing platforms are being fueled by unprecedented investments in research and development by information technology companies—companies that are competing intensely with one another for contracts with enterprises that are hungry for digital transformation and the latest in modern computing tools. Cloud computing companies are packaging in their offerings sets of development tools and constellations of specialized services. Many of these offerings are relevant to biomedical informatics efforts, including machine learning toolkits, suites for analysis and visualization of data, and computer vision, speech recognition, and natural language analysis services made available via programmatic interfaces.

Beyond developing generic platform capabilities, cloud service providers are motivated to gain understandings in key vertical markets, such as healthcare, finance, and defense, and have been working to custom-tailor their general platforms with tools, designs, and services for use in specific sectors. For example, there is incentive to support rising standards on schemata (e.g., Fast Healthcare Interoperability Resources (FHIR)) for storing and transferring electronic health records and on methods to ensure the privacy of patient data. There is also pressure to develop special versions of computing services for medicine, such as language models and, more generally, natural language capabilities specialized for medical terminology, enabling more accurate understanding and analysis of medical text and speech. Competitor cloud computing companies have also worked to identify and provide efficient methods and tools for important vertical needs, such as the rising importance of determining DNA sequences and interpreting protein expression data. Such special needs of researchers and clinicians have led to the availability of efficient and inexpensive cloud-computing services for genomic and proteomic analyses.

Moving on to the second realm of opportunities, around harnessing advances in the constellation of technologies that we call AI, I believe that our community can do more to leverage existing methods and also to closely follow, push, and contribute to advances in AI subfields. Beyond methods available today, key developments will be required in principles and applications to realize the long-term goals of biomedical informatics. I am seeing good progress and am optimistic that the advances coming over the next decade will be deeply enabling.

On existing technologies, and focusing on the example of developing effective decision support systems, we have been very slow to leverage the visionary ideas proposed by Robert Ledley and Lee

Lusted in 1959.[1] Ledley and Lusted published a blueprint for constructing differential diagnoses and to use decision-theoretic analyses to generate recommendations for action. Biomedical informatics investigators have been top leaders with exploring prototypes for decision support systems, and systems constructed over 60 years of research have been shown to perform at expert levels. However, real-world impact has been limited to date. A key bottleneck has been the scarcity and cost of expertise and data. I believe that harnessing advances in machine learning will be particularly critical for delivering on the vision of evidence-based clinical decision making. Machine learning techniques available today can and should be playing a more central role in healthcare for assisting with pattern recognition, diagnosis, and prediction of outcomes. There are multiple opportunities to build and to integrate pipelines where data flow via machine learning to predictions and via automated decision analyses to recommendations about testing and treatment. Making key investments to build and refine effective *data-to-prediction-to-decision* pipelines will provide great value in multiple areas of medicine.

Opportunities ahead for biomedical informatics includes leveraging recent advances in deep learning in medical applications, especially for image recognition and natural language tasks. These multilayered neural network architectures are celebrated for providing surprising boosts in classification accuracy in multiple application areas and for easing engineering overhead, as they do not require special feature engineering. The methods have been shown to perform well for recognition in the image-centric areas of pathology and radiology. Different variants of deep learning are also being explored for building predictive models from clinical data drawn from electronic health records. Beyond direct applications, deep learning methods have led to enhanced capabilities in multiple areas of AI with relevance to goals in biomedical informatics, including key advances in computer vision, speech recognition, text summarization, and language translation.

With the all of recent fanfare about deep learning, it is easy to overlook the applicability of other machine learning methods, including probabilistic graphical models, generalized additive models, and even logistic regression for serving as the heart of predictions in recommendation engines. While excitement about deep learning is appropriate, it is important to note that the methods typically require large amounts of data of the right form and that such datasets may not be available for medical applications of interest. Other approaches have proven to be as accurate for clinical applications and come with other benefits such as providing more intelligible, explainable inferences. Also, when sufficiently large corpora of data labeled with ground truth are not available, knowledge acquisition

[1] R.S. Ledley and L.B. Lusted (1959). *Reasoning Foundations of Medical Diagnosis*, Science. Vol. 130, No. 3366 (Jul. 3, 1959), pp. 9-21.

techniques, referred to broadly as *machine teaching*, can provide value. While work is moving forward on machine teaching, existing methods and tools can be valuable in building models for prediction and classification.

I believe that it is important to note that having access to powerful machine learning procedures may be insufficient for addressing goals in biomedical informatics. Key challenges for moving ahead with developing and deploying effective decision support systems include identifying where and when such systems would provide value, collecting sufficient amounts of the right kind of data for applications, developing and integrating automated decision analyses to move from predictions to recommendations for action[2], maintaining systems over time, developing means to build and apply learned models at multiple sites, and addressing human-factors, including formulating means for achieving smooth integration of inferences and recommendations into clinical workflows, and providing explanations of inferences to clinicians.[3] Providing explanations of predictions generated by machine-learned models is a topic of rising interest.[4] I hope to see revitalized interest and similar enthusiasm extended to addressing challenges identified in bioinformatics with the intelligibility and explanation of the advice provided by other forms of reasoning employed in decision support systems, including logical, probabilistic, and decision-theoretic inference.[5]

Key opportunities in AI research for progress with developing and fielding effective decision support systems include efforts in principles and applications of *transfer learning*, *unsupervised learning*, and *causal inference*. Transfer learning refers to methods that allow for data or task competencies learned in one area to be applied to another. Unsupervised and semi-supervised learning refers to methods that can be used to build models and perform tasks without having a complete set of labeled data, such as labels about the final diagnoses of patients when working with electronic health records data. Causal inference refers to methods that can be used to identify causal knowledge, versus statistical associations

[2] M. Bayati, M. Braverman, M. Gillam, K.M. Mack, G. Ruiz, M.S. Smith, E. Horvitz (2014). *Data-Driven Decisions for Reducing Readmissions for Heart Failure: General Methodology and Case Study*. PLOS One Medicine. October 2014.

[3] R.L. Teach and E.H. Shortliffe (1981). *An analysis of physician attitudes regarding computer-based clinical consultation systems*. Computers and Biomedical Research, Volume 14, Issue 6, December 1981, pp. 542-558. https://doi.org/10.1016/0010-4809(81)90012-4.

[4] R. Caruana, P. Koch, Y. Lou, M. Sturm , J. Gehrke, N. Elhadad (2015). *Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission.* KDD, August 10-13, 2015, Sydney, NSW, Australia, August 2015.

[5] E. Horvitz, D. Heckerman, B. Nathwani, L.M. Fagan (1986), *The use of a heuristic problem-solving hierarchy to facilitate the explanation of hypothesis-directed reasoning*, October 1986, Medinfo, Washington, DC, North Holland: New York, pp. 27-31.

that are commonly inferred from data. Advances in these areas promise to provide new sources of biomedical knowledge, and to address the challenge of data scarcity and related difficulties with the generalizability of data resources for healthcare applications.

On data scarcity and generalizability, an important, often underappreciated challenge in biomedical informatics is that the accuracy of diagnosis and decision support may not transfer well across institutions. In our work at Microsoft Research, we found that accuracies of a system trained on data obtained from a site can plummet when used at another location. The poor generality of datasets is based on multiple factors, including differences in patient populations—with site-specific incidence rates, covariates, and presentations of illness, site-specific capture of evidence in the electronic health record, and site-specific definitions of signs, symptoms, and lab results. As an example, we found site-specificity when my team studied the task of building models to predict the likelihood that patients being discharged from a hospital would be readmitted within 30 days. The accuracy of prediction for a model learned from a massive dataset drawn from single large urban hospital dropped when the model was applied at other hospitals. This observation of poor generalizability was behind our decision to develop a capability for performing automated, recurrent machine learning separately at each site that would rely on local data for predictions. This local train-and-test capability served as the core engine of an advisory system for readmissions management, named Readmission Manager, that was commercialized by Microsoft.

Moving forward, research on a set of methods jointly referred to as *transfer learning* may help to address challenges of data scarcity and generalizability. Transfer learning algorithms for mapping the learnings from one hospital to another show promise in medicine.[6] Such methods include *multitask learning*. Also, obtaining spanning datasets, composed of large amounts of data drawn from multiple sites, may provide a path to effective generalization. In support of this approach, methods called multiparty computation have been developed that can enable learning from multiple, privately held databases, where there is no violation of privacy among the contributing organizations.

Beyond the daily practice of healthcare, and uses in such applications as diagnosis and treatment, methods for learning and reasoning from data can provide the foundations for new directions in the clinical sciences via tools and analyses that identify subtle but important signals in the fusing of clinical, behavioral, environmental, genetic, and epigenetic data. I see many directions springing from applications

[6] J. Wiens, J. Guttag, and E. Horvitz (2014). *A Study in Transfer Learning: Leveraging Data from Multiple Hospitals to Enhance Hospital-Specific Predictions*, Journal of the American Medical Informatics Association: 21(4):699-706. doi.org/10.1136/amiajnl-2013-002162

of machine learning, reasoning, planning, and causal inference for healthcare delivery as well as in supporting efforts in healthcare policy and in the discovery of new biomedical understandings.

I remain excited about advances in biomedical informatics and see a biomedical informatics revolution on the horizon. Such a revolution will build on the glowing embers of decades of contributions and the flames of late-breaking activities that address long-term challenges and bottlenecks.