

A recent IEEE Pulse on Stage forum offers some perspective.

©ISTOCKPHOTO.COM/MUSTAFA BOZDEMIR

Can AI Truly Transform Health Care?

By Ahmed Morsy

As scholars have predicted and researchers have now shown, we are entering an age of global artificial intelligence (AI) convergence. Health care is just one area in which AI is gaining a foothold, as evidenced by two parallel conferences held in March 2018 in Las Vegas, Nevada. The first of these was the annual meeting of the Hospital Information and Management

Systems Society, attended by nearly 45,000 participants; the second was a more focused, engineering-oriented conference on biomedical health informatics and biosensor networks, attended by about 450. Straddling those two meetings, *IEEE Pulse* held its second IEEE Pulse on Stage event, focusing (not surprisingly) on the use of AI in health care (Figure 1).

2018 IEEE Pulse on Stage Forum

IEEE Pulse on Stage is meant to be a mixing-pot forum for engineers, medical doctors, clinical researchers, informatics

Digital Object Identifier 10.1109/MPUL.2018.2834998
Date of publication: 16 July 2018

researchers, computer scientists, health information technology professionals, and industrialists. The goal of such a forum is to ask and try to provide answers for fundamental questions of interest to a broad spectrum of the health-care delivery community. This IEEE Pulse on Stage event asked a particularly loaded question: “Can AI truly transform health care?” Implicitly, we posed three fundamental issues related to employing AI in healthcare delivery: 1) how much progress has already been achieved on the ground, 2) why is the hype so overwhelming, and 3) what is the true potential of AI in transforming health care?

To help answer these critical questions, the event showcased a lineup of distinguished speakers representing a variety of organizations (Figure 2):

- ▼ IBM (engineer Salwa Rafee, representing the chief technology officer of the Watson and Cognitive Computing team)
- ▼ Microsoft Research (Dr. Rich Caruana, principal researcher)
- ▼ Humetrix, a Silicon Valley-based start-up (Dr. Bettina Experton, founder and chief executive officer)
- ▼ Georgia Tech Health Informatics Center (Dr. Jon Duke, director).

The speakers addressed the topic from various enlightening angles, and the audience reciprocated with many probing questions. The take-home messages were as numerous as they were vital.

The forum emphasized the importance of building highly interpretable predictive models from healthcare data. Highly interpretable models can help not only in identifying possible problems within the employed model but also in identifying issues that may be present within the data—an all too common problem in itself. It is worth mentioning that both the U.S. Food and Drug Administration regulations in the United States and the European Union General Data Protection Regulation advocate strongly for “explainability.” While the debate over how limiting this requirement is likely to continue, it can be argued that model interpretability can help bridge a few gaps in this debate. But model interpretability comes at the expense of model accuracy, as is evident in much of the work that appears in research literature. One would certainly hope to find a highly interpretable model that is also very accurate, but, in the presence of such a fundamental tradeoff, it is important to wisely sacrifice some of the interpretability to maintain decent accuracy levels.

Even so, the importance of model accuracy comes under more intense

scrutiny when we remember that medical errors have now risen to the frightening position of third leading cause of death, at least in developing countries. This is an area in which it is reasonable to expect AI to play a major role: preventing many of these errors by augmenting human knowledge and preventing lapses caused by fatigue or hastiness. Technologies to help reduce medical errors and wrong diagnoses can be enabled at a large scale via initiatives like Blue Button (<https://www.healthit.gov/buzz-blog/consumer/announcing-blue-button-toolkit/>)—particularly after it is made available on mobile platforms—along with technology frameworks such as IBM Watson.

Clinical Decision Support Systems

Mending certain gaps in medical professionals’ knowledge and prevention of lapses caused by fatigue or hastiness both lend themselves well to the domain formally known as *clinical decision support systems*. Such systems are also built on predictive models designed using machine-learning techniques. While clinical decision support systems hold tremendous promise, they also face many challenges.

For example, trust is a major factor when it comes to deciding to acknowledge a machine-generated alert (and not override it). Inaccuracies of predictive models in this case may generate a large number of alerts, mostly of the false positive or “uncalled-for” type, causing what is commonly known as *alert fatigue*.

In reality, alert override occurs at a rate that cannot be accounted for solely based on alert fatigue or excessive machine false positive errors. Rather, it is a culture ingrained among doctors to ignore (or not trust) many machine-generated alerts. This culture-induced lack of trust can be addressed by providing analytics on how many peers or superiors have accepted a particular alert in cases that prove “close enough” to the case at hand. An explanation of the decision-making process is another factor that can help address this trust issue.

FIGURE 1 An announcement for the 2018 IEEE Pulse on Stage event.



FIGURE 2 Ahmed Morsy, editor-in-chief of *IEEE Pulse* (far right), poses with three of the four keynote speakers from the recent IEEE Pulse on Stage event. From left: Rich Caruana, Ph.D., Microsoft Research; Bettina Experton, M.D., Humetrix; and Jon Duke, M.D., Georgia Tech. (Photo courtesy of Cynthia Weber.)

At a more fundamental level, designers of clinical decision support systems need to understand what level of innate confidence is associated with which alert (or clinical recommendation) and what level of seriousness (or criticalness) is associated with the clinical problem being addressed. To reduce alert fatigue, inspire trust in the system, and improve clinical outcomes, alerts must be prioritized based on the two important parameters of confidence and criticality. In fact, improved results can, in themselves, provide another incentive. Clinical decision support systems can gain stronger traction in the medical community if they are proven to help doctors secure financial incentives related to increased quality of care and attaining better health outcomes.

AI in the Future: The 80/20 Rule

So where should our focus be with regard to decision support over the shorter range? It appears that a variant of the 80/20 rule may be needed here. Over the next few years, 80% of mundane tasks and decisions now performed by less experienced healthcare professionals may lend themselves well to being replaced or heavily augmented by AI tools in a fairly non-conservative approach. This can also provide a great learning environment for the young and less experienced by continually preparing analytics on the performance of such AI systems with ample analyses for what went right or wrong and why. On the other hand, the more critical and sensitive 20% can use

There was also a clear consensus that most of the progress in this area is still ahead of us.

enthusiasm about what AI can do and the conservatism induced by the fact that machines still, and very often, make mistakes that are not only embarrassing but also capable of putting patients at real risks.

There was also a clear consensus that most of the progress in this area is still ahead of us. For that, similar dialogs are extremely important to cover more, and in more depth, fundamental aspects of the problem. These include the security and vulnerability of deployed models, along with traceability in terms of which models were used where and when to provide what type and quality of outcome. Interoperability of models invented and tested by different groups and entities will also be a major success determinant.

We hope to address most of these topics in another IEEE Pulse on Stage forum in Honolulu, Hawaii, during the 40th meeting of the IEEE Engineering in Medicine and Biology Society (17–20 July 2018). This will be an opportunity to further advance discussion of the critical role of AI in health care, and *IEEE Pulse* invites you to join the conversation at this special event (<https://pulse.embs.org/onstage>).

Ahmed Morsy (amorsy@ieee.org) is the editor-in-chief of IEEE Pulse and a professor of biomedical engineering at Cairo University, Egypt.

some AI guidance in a strongly conservative manner. This balance can fit both the needs and the realities on the ground. As AI models become more mature and our experiences with deploying them become more positive, the 80/20 mix can be tilted toward a more liberal application of AI in various clinical settings.

Progress So Far

At the closing of the question and answer session, all four speakers agreed that appreciable progress has already been made on more than one front. On one important front, we now certainly have better models and machine-learning tools. On another front, we have the ability to ingest huge amounts of healthcare data, do lots of number crunching, and, finally, give an output based on the intrinsic value of those tons of data. Open-data initiatives can be considered another type of progress driving further future progress. Finally, an important sign of progress is an increased balance between en-