Machine-learning system could aid critical decisions in sepsis care

November 7 2018, by Rob Matheson

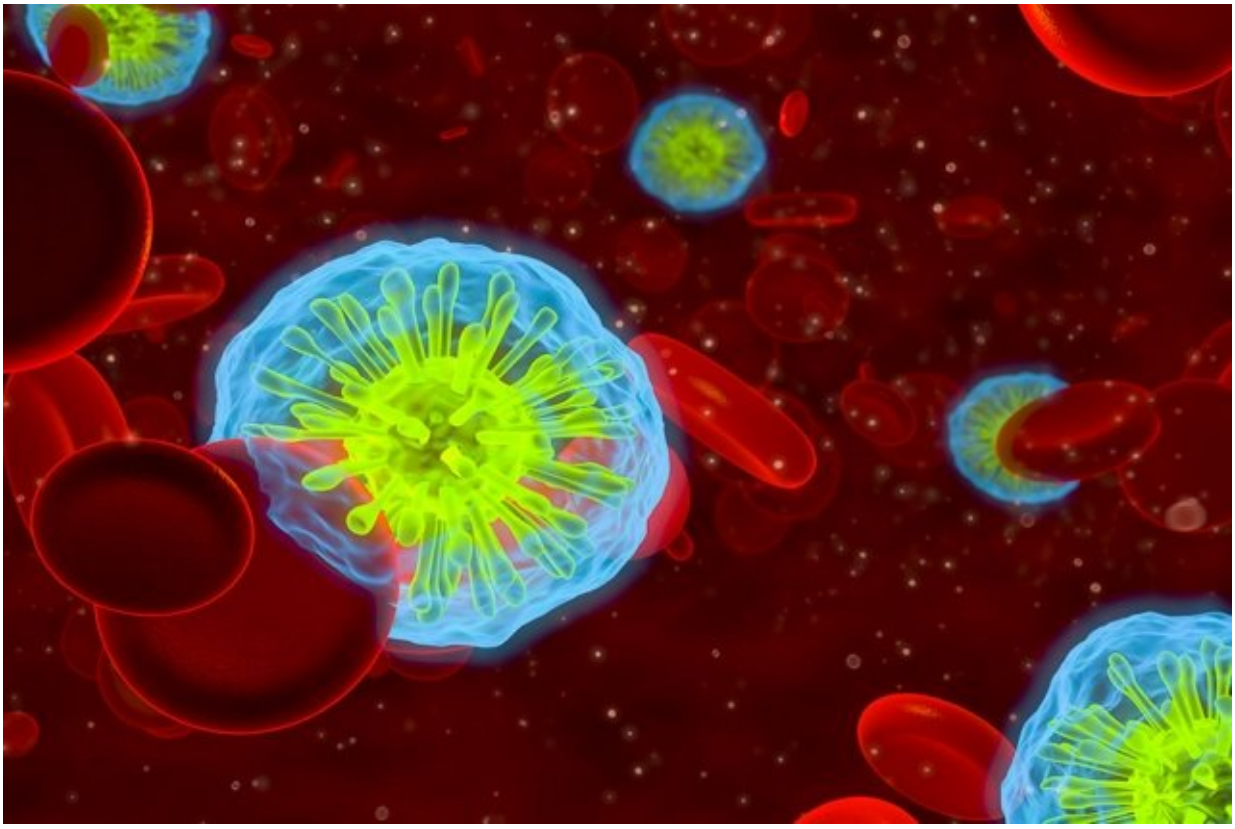

A new machine-learning model that predicts whether ER patients suffering from sepsis may need to be switched to certain medications could help guide clinicians in sepsis care. Credit: Massachusetts Institute of Technology

Researchers from MIT and Massachusetts General Hospital (MGH) have

developed a predictive model that could guide clinicians in deciding when to give potentially life-saving drugs to patients being treated for sepsis in the emergency room.

Sepsis is one of the most frequent causes of admission, and one of the most common causes of death, in the intensive care unit. But the vast majority of these patients first come in through the ER. Treatment usually begins with antibiotics and intravenous fluids, a couple liters at a time. If patients don't respond well, they may go into septic shock, where their blood pressure drops dangerously low and organs fail. Then it's often off to the ICU, where clinicians may reduce or stop the fluids and begin vasopressor medications such as norepinephrine and dopamine, to raise and maintain the patient's blood pressure.

That's where things can get tricky. Administering fluids for too long may not be useful and could even cause organ damage, so early vasopressor intervention may be beneficial. In fact, early vasopressor administration has been linked to improved mortality in septic shock. On the other hand, administering vasopressors too early, or when not needed, carries its own negative health consequences, such as heart arrhythmias and cell damage. But there's no clear-cut answer on when to make this transition; clinicians typically must closely monitor the patient's blood pressure and other symptoms, and then make a judgment call.

In a paper being presented this week at the American Medical Informatics Association's Annual Symposium, the MIT and MGH researchers describe a model that "learns" from health data on emergency-care sepsis patients and predicts whether a patient will need vasopressors within the next few hours. For the study, the researchers compiled the first-ever dataset of its kind for ER sepsis patients. In testing, the model could predict a need for a vasopressor more than 80 percent of the time.

Early prediction could, among other things, prevent an unnecessary ICU stay for a patient that doesn't need vasopressors, or start early preparation for the ICU for a patient that does, the researchers say.

"It's important to have good discriminating ability between who needs vasopressors and who doesn't [in the ER]," says first author Varesh Prasad, a Ph.D. student in the Harvard-MIT Program in Health Sciences and Technology. "We can predict within a couple of hours if a patient needs vasopressors. If, in that time, patients got three liters of IV fluid, that might be excessive. If we knew in advance those liters weren't going to help anyway, they could have started on vasopressors earlier."

In a clinical setting, the model could be implemented in a bedside monitor, for example, that tracks patients and sends alerts to clinicians in the often-hectic ER about when to start vasopressors and reduce fluids. "This model would be a vigilance or surveillance system working in the background," says co-author Thomas Heldt, the W. M. Keck Career Development Professor in the MIT Institute of Medical Engineering and Science. "There are many cases of sepsis that [clinicians] clearly understand, or don't need any support with. The patients might be so sick at initial presentation that the physicians know exactly what to do. But there's also a 'gray zone,' where these kinds of tools become very important."

Co-authors on the paper are James C. Lynch, an MIT graduate student; and Trent D. Gillingham, Saurav Nepal, Michael R. Filbin, and Andrew T. Reisner, all of MGH. Heldt is also an assistant professor of electrical and biomedical engineering in MIT's Department of Electrical Engineering and Computer Science and a principal investigator in the Research Laboratory of Electronics.

Other models have been built to predict which patients are at risk for sepsis, or when to administer vasopressors, in ICUs. But this is the first

model trained on the task for the ER, Heldt says. "[The ICU] is a later stage for most sepsis patients. The ER is the first point of patient contact, where you can make important decisions that can make a difference in outcome," Heldt says.

The primary challenge has been a lack of an ER database. The researchers worked with MGH clinicians over several years to compile medical records of nearly 186,000 patients who were treated in the MGH emergency room from 2014 to 2016. Some patients in the dataset had received vasopressors within the first 48 hours of their hospital visit, while others hadn't. Two researchers manually reviewed all records of patients with likely septic shock to include the exact time of vasopressor administration, and other annotations. (The average time from presentation of sepsis symptoms to vasopressor initiation was around six hours.)

The records were randomly split, with 70 percent used for training the model and 30 percent for testing it. In training, the model extracted up to 28 of 58 possible features from patients who needed or didn't need vasopressors. Features included blood pressure, elapsed time from initial ER admission, total fluid volume administered, respiratory rate, mental status, oxygen saturation, and changes in cardiac stroke volume—how much blood the heart pumps in each beat.

In testing, the model analyzes many or all of those features in a new patient at set time intervals and looks for patterns indicative of a patient that ultimately needed vasopressors or didn't. Based on that information, it makes a prediction, at each interval, about whether the patient will need a vasopressor. In predicting whether patients needed vasopressors in the next two or more hours, the model was correct 80 to 90 percent of the time, which could prevent an excessive half a liter or more of administered fluids, on average.

"The model basically takes a set of current vital signs, and a little bit of what the trajectory looks like, and determines that this current observation suggests this patient might need vasopressors, or this set of variables suggests this patient would not need them," Prasad says.

Next, the researchers aim to expand the work to produce more tools that predict, in real-time, if ER patients may initially be at risk for sepsis or septic shock. "The idea is to integrate all these tools into one pipeline that will help manage care from when they first come into the ER," Prasad says.

The idea is to help clinicians at emergency departments in major hospitals such as MGH, which sees about 110,000 patients annually, focus on the most at-risk populations for sepsis. "The problem with sepsis is the presentation of the patient often belies the seriousness of the underlying disease process," Heldt says. "If someone comes in with weakness and doesn't feel right, a little bit of fluids may often do the trick. But, in some cases, they have underlying sepsis and can deteriorate very quickly. We want to be able to tell which patients have become better and which are on a critical path if left untreated."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology