


Machine-learning system flags remedies that might do more harm than good

December 9 2021, by Adam Zewe



A new machine-learning system could give physicians risk scores for different treatments. Credit: Christine Daniloff, MIT

Sepsis claims the lives of nearly 270,000 people in the U.S. each year. The unpredictable medical condition can progress rapidly, leading to a

swift drop in blood pressure, tissue damage, multiple organ failure, and death.

Prompt interventions by medical professionals save lives, but some sepsis treatments can also contribute to a patient's deterioration, so choosing the optimal therapy can be a difficult task. For instance, in the early hours of severe sepsis, administering too much fluid intravenously can increase a patient's risk of death.

To help clinicians avoid remedies that may potentially contribute to a patient's death, researchers at MIT and elsewhere have developed a [machine-learning model](#) that could be used to identify treatments that pose a higher risk than other options. Their model can also warn doctors when a septic patient is approaching a medical dead end—the point when the patient will most likely die no matter what [treatment](#) is used—so that they can intervene before it is too late.

When applied to a dataset of sepsis patients in a hospital intensive care unit, the researchers' model indicated that about 12 percent of treatments given to patients who died were detrimental. The study also reveals that about 3 percent of patients who did not survive entered a medical dead end up to 48 hours before they died.

"We see that our model is almost eight hours ahead of a doctor's recognition of a patient's deterioration. This is powerful because in these really sensitive situations, every minute counts, and being aware of how the patient is evolving, and the risk of administering certain treatment at any given time, is really important," says Taylor Killian, a graduate student in the Healthy ML group of the Computer Science and Artificial Intelligence Laboratory (CSAIL).

Joining Killian on the paper are his advisor, Assistant Professor Marzyeh Ghassemi, head of the Healthy ML group and senior author; lead author

Mehdi Fatemi, a senior researcher at Microsoft Research; and Jayakumar Subramanian, a senior research scientist at Adobe India. The research is being presented at this week's Conference on Neural Information Processing Systems.

A dearth of data

This research project was spurred by a 2019 paper Fatemi wrote that explored the use of reinforcement learning in situations where it is too dangerous to explore arbitrary actions, which makes it difficult to generate enough data to effectively train algorithms. These situations, where more data cannot be proactively collected, are known as "offline" settings.

In reinforcement learning, the algorithm is trained through trial and error and learns to take actions that maximize its accumulation of reward. But in a health care setting, it is nearly impossible to generate enough data for these models to learn the optimal treatment, since it isn't ethical to experiment with possible treatment strategies.

So, the researchers flipped reinforcement learning on its head. They used the limited data from a hospital ICU to train a reinforcement learning model to identify treatments to avoid, with the goal of keeping a patient from entering a medical dead end.

Learning what to avoid is a more statistically efficient approach that requires fewer data, Killian explains.

"When we think of dead ends in driving a car, we might think that is the end of the road, but you could probably classify every foot along that road toward the dead end as a dead end. As soon as you turn away from another route, you are in a dead end. So, that is the way we define a medical dead end: Once you've gone on a path where whatever decision

you make, the patient will progress toward death," Killian says.

"One core idea here is to decrease the probability of selecting each treatment in proportion to its chance of forcing the patient to enter a medical dead-end—a property that is called treatment security. This is a hard problem to solve as the data do not directly give us such an insight. Our theoretical results allowed us to recast this core idea as a reinforcement learning problem," Fatemi says.

To develop their approach, called Dead-end Discovery (DeD), they created two copies of a neural network. The first neural network focuses only on negative outcomes—when a patient died—and the second network only focuses on positive outcomes—when a patient survived. Using two neural networks separately enabled the researchers to detect a risky treatment in one and then confirm it using the other.

They fed each neural network patient health statistics and a proposed treatment. The networks output an estimated value of that treatment and also evaluate the probability the patient will enter a medical dead end. The researchers compared those estimates to set thresholds to see if the situation raises any flags.

A yellow flag means that a patient is entering an area of concern while a red flag identifies a situation where it is very likely the patient will not recover.

Treatment matters

The researchers tested their model using a dataset of patients presumed to be septic from the Beth Israel Deaconess Medical Center [intensive care unit](#). This dataset contains about 19,300 admissions with observations drawn from a 72-hour period centered around when the patients first manifest symptoms of sepsis. Their results confirmed that

some patients in the dataset encountered medical dead ends.

The researchers also found that 20 to 40 percent of patients who did not survive raised at least one yellow flag prior to their death, and many raised that flag at least 48 hours before they died. The results also showed that, when comparing the trends of patients who survived versus patients who died, once a patient raises their first flag, there is a very sharp deviation in the value of administered treatments. The window of time around the first flag is a critical point when making treatment decisions.

"This helped us confirm that treatment matters and the treatment deviates in terms of how patients survive and how patients do not. We found that upward of 11 percent of suboptimal treatments could have potentially been avoided because there were better alternatives available to doctors at those times. This is a pretty substantial number, when you consider the worldwide volume of patients who have been septic in the hospital at any given time," Killian says.

Ghassemi is also quick to point out that the model is intended to assist doctors, not replace them.

"Human clinicians are who we want making decisions about care, and advice about what treatment to avoid isn't going to change that," she says. "We can recognize risks and add relevant guardrails based on the outcomes of 19,000 patient treatments—that's equivalent to a single caregiver seeing more than 50 septic patient outcomes every day for an entire year."

Moving forward, the researchers also want to estimate causal relationships between treatment decisions and the evolution of patient health. They plan to continue enhancing the model so it can create uncertainty estimates around treatment values that would help doctors

make more informed decisions. Another way to provide further validation of the model would be to apply it to data from other hospitals, which they hope to do in the future.

More information: Medical Dead-ends and Learning to Identify High-risk States and Treatments. [papers.neurips.cc/paper/2021/f ...
74eb7c696c-Paper.pdf](https://papers.neurips.cc/paper/2021/f...74eb7c696c-Paper.pdf)

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

Citation: Machine-learning system flags remedies that might do more harm than good (2021, December 9) retrieved 27 April 2024 from <https://medicalxpress.com/news/2021-12-machine-learning-flags-remedies-good.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--