

Technique reveals whether models of patient risk are accurate

January 23 2020



Credit: CC0 Public Domain

After a patient has a heart attack or stroke, doctors often use risk models to help guide their treatment. These models can calculate a patient's risk of dying based on factors such as the patient's age, symptoms, and other

characteristics.

While these models are useful in most cases, they do not make [accurate predictions](#) for many patients, which can lead doctors to choose ineffective or unnecessarily risky treatments for some patients.

"Every [risk model](#) is evaluated on some dataset of patients, and even if it has [high accuracy](#), it is never 100 percent accurate in practice," says Collin Stultz, a professor of electrical engineering and [computer science](#) at MIT and a cardiologist at Massachusetts General Hospital. "There are going to be some patients for which the [model](#) will get the wrong answer, and that can be disastrous."

Stultz and his colleagues from MIT, the MIT-IBM AI Lab, and the University of Massachusetts Medical School have now developed a method that allows them to determine whether a particular model's results can be trusted for a given patient. This could help guide doctors to choose better treatments for those patients, the researchers say.

Stultz, who is also a professor of health sciences and technology, a member of MIT's Institute for Medical Engineering and Sciences and Research Laboratory of Electronics, and an associate member of the Computer Science and Artificial Intelligence Laboratory, is the senior author of the new study. MIT graduate student Paul Myers is the lead author of the paper, which appears today in *Digital Medicine*.

Modeling risk

Computer models that can predict a patient's risk of harmful events, including death, are used widely in medicine. These models are often created by training machine-learning algorithms to analyze patient datasets that include a variety of information about the patients, including their health outcomes.

While these models have high overall accuracy, "very little thought has gone into identifying when a model is likely to fail," Stultz says. "We are trying to create a shift in the way that people think about these machine-learning models. Thinking about when to apply a model is really important because the consequence of being wrong can be fatal."

For instance, a patient at high risk who is misclassified would not receive sufficiently aggressive treatment, while a low-risk patient inaccurately determined to be at high risk could receive unnecessary, potentially harmful interventions.

To illustrate how the method works, the researchers chose to focus on a widely used risk model called the GRACE risk score, but the technique can be applied to nearly any type of risk model. GRACE, which stands for Global Registry of Acute Coronary Events, is a large dataset that was used to develop a risk model that evaluates a patient's risk of death within six months after suffering an acute coronary syndrome (a condition caused by decreased blood flow to the heart). The resulting risk assessment is based on age, blood pressure, heart rate, and other readily available clinical features.

The researchers' new technique generates an "unreliability score" that ranges from 0 to 1. For a given risk-model prediction, the higher the score, the more unreliable that prediction. The unreliability score is based on a comparison of the risk prediction generated by a particular model, such as the GRACE risk-score, with the prediction produced by a different model that was trained on the same dataset. If the models produce different results, then it is likely that the risk-model prediction for that patient is not reliable, Stultz says.

"What we show in this paper is, if you look at patients who have the highest unreliability scores—in the top 1 percent—the risk prediction for that patient yields the same information as flipping a coin," Stultz

says. "For those patients, the GRACE score cannot discriminate between those who die and those who don't. It's completely useless for those patients."

The researchers' findings also suggested that the patients for whom the models don't work well tend to be older and to have a higher incidence of cardiac risk factors.

One significant advantage of the method is that the researchers derived a formula that tells how much two predictions would disagree, without having to build a completely new model based on the original dataset.

"You don't need access to the training dataset itself in order to compute this unreliability measurement, and that's important because there are privacy issues that prevent these clinical datasets from being widely accessible to different people," Stultz says.

Retraining the model

The researchers are now designing a user interface that doctors could use to evaluate whether a given patient's GRACE score is reliable. In the longer term, they also hope to improve the reliability of risk models by making it easier to retrain models on data that include more patients who are similar to the patient being diagnosed.

"If the model is simple enough, then retraining a model can be fast. You could imagine a whole suite of software integrated into the electronic health record that would automatically tell you whether a particular risk score is appropriate for a given patient, and then try to do things on the fly, like retrain new models that might be more appropriate," Stultz says.

Provided by Massachusetts Institute of Technology

Citation: Technique reveals whether models of patient risk are accurate (2020, January 23)
retrieved 23 April 2024 from
<https://medicalxpress.com/news/2020-01-technique-reveals-patient-accurate.html>

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.