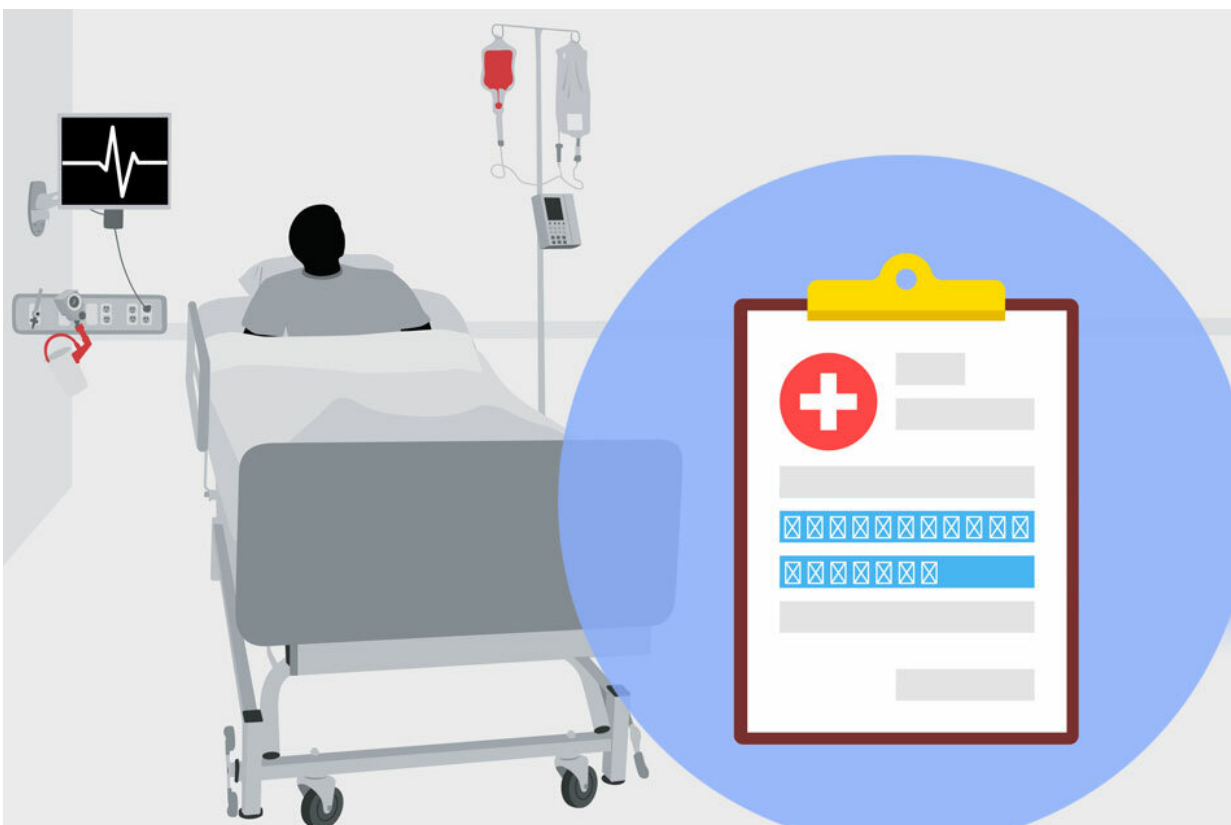# Neural network assimilates multiple types of health data to help doctors make decisions with incomplete information

January 28 2019, by Rob Matheson

New research describes a single neural network that can collect missing data from a patient and add it to their chart. Credit: MIT News

MIT researchers have developed a model that can assimilate multiple

types of a patient's health data to help doctors make decisions with incomplete information.

The field of "predictive analytics" holds promise for many health care applications. Machine learning models can be trained to look for patterns in patient data to predict a patient's risk for disease or dying in the ICU, to aid in sepsis care, or to design safer chemotherapy regimens.

The process involves predicting variables of interest, such as disease risk, from known variables, such as symptoms, biometric data, lab tests, and body scans. However, that patient data can come from several different sources and is often incomplete. For example, it might include partial information from health surveys about physical and mental well-being, mixed with highly complex data comprising measurements of heart or brain function.

Using machine learning to analyze all available data could help doctors better diagnose and treat patients. But most models can't handle the highly complex data. Others fail to capture the full scope of the relationships between different health variables, such as how breathing patterns help predict sleeping hours or pain levels.

In a paper being presented at the AAAI Conference on Artificial Intelligence next week, MIT researchers describe a single neural network that takes as input both simple and highly complex data. Using the known variables, the network can then fill in all the missing variables. Given data from, say, a patient's electrocardiography (ECG) signal, which measures heart function, and self-reported fatigue level, the model can predict a patient's pain level, which the patient might not remember or report correctly.

Tested on a real sleep study dataset—which contained health surveys, and ECG and other complex signals—the network achieved 70 to 80

percent accuracy in predicting any one of eight missing variables, based on the seven other known variables.

The network works by stitching together various submodels, each tailored to describe a specific relation among variables. The submodels share data as they make predictions, and ultimately output a predicted target [variable](#). "We have a network of models that communicate with each other to predict what we don't know, using the information we do know from these different types of data," says lead author Hao Wang, a postdoc at the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). "If you have, say, eight different types of data, and I have full information on a patient from seven, the communication between the models will help us fill in the missing gaps in the eighth type of data from the other seven types."

Joining Wang on the paper are Chengzhi Mao, an undergraduate student at Tsinghua University; CSAIL Ph.D. students Hao He and Mingmin Zhao; Dina Katabi, the Andrew and Erna Viterbi Professor of Electrical Engineering and Computer Science at MIT and director of the MIT Center for Wireless Networks and Mobile Computing; and Tommi S. Jaakkola, the Thomas Siebel Professor in the Department of Electrical Engineering and Computer Science and the Institute for Data, Systems, and Society.

Bi-directional predictions

Using traditional machine-learning models to analyze the number of variables the researchers' network can handle is practically infeasible, because the number of models scales exponentially with the number of variables.

"We asked, 'Is it possible to design a single model that can use all these groups of data, despite the fact that in each group we have different

information?'" Wang says.

The key innovation was breaking the network into individual submodels each tailored to fit a different type of input data. A neural network is an interconnected network of nodes that work together to process complex data. One node does relatively simple computations before sending the output to the next node. In networks with submodels, however, each node can function as a separate network that can handle more complex computations. Submodels can be much more efficient, depending on the application.

In their work, the researchers created one probabilistic submodel for each variable output. They also developed a technique to let the submodels communicate with one another while making predictions, called Bi-directional Inference Networks (BIN). This technique leverages a neural network training technique known as backpropagation. In training, backpropagation sends computing errors back through the nodes to update the network's parameter values. But this technique is never used in testing, especially when there are complex conditional dependencies involved. Instead, in traditional testing, inputted data get processed from node to node in one direction, until a final node at the end of the sequence outputs a prediction.

The researchers programmed their network to use both the traditional method and backpropagation during testing. In this context, backpropagation is basically taking a variable output, then predicting an input from that output, and sending the input value backward to a previous node. This creates a network where all submodels are working together and co-dependent on one another, to output a target probability.

Filling in the blanks

The researchers trained their network on the real-world Sleep Heart

Health Study 2 (SHHS2) dataset. The data include electroencephalography (EEG) readings, which measure brain function; ECG; and breathing pattern signals. It also includes information from a health survey to measure eight health variables—including emotional well-being, social functioning, and energy/fatigue—ranked on a scale of 0 to 100.

In training, the network learns patterns for how each variable may affect another. For instance, if someone holds their breath for long durations, they may be tense, which can indicate physical pain. In testing, the network is able to analyze the relationships to predict any of the eight variables, based on any of the other information, with 70 to 80 percent accuracy.

The network could help quantify sometimes-ambiguous health variables for patients and doctors, such as pain and fatigue levels. When patients sleep after surgery, for instance, they may wake up in the middle of the night in pain, but may not remember an appropriate pain level the next day.

Next, the researchers hope to implement the network as a software component for a device they built, called the EQ Radio, which can track someone's breathing and heart rate using only wireless signals. Currently, the device analyzes that information to infer if someone is happy, angry, or sad. With the [network](link), the device could potentially make continuously updated predictions about a patient's health, passively, given only partial information. "This could be so helpful in assisted-living facilities, where doctors can monitor both emotional and physical dimensions of a patient's health all day, every day," Wang says.

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