

Active learning model for computer predictions

December 3 2013, by Mark Riechers



Rob Nowak and Kevin Jamieson. Credit: Nick Berard.

Computers serve as powerful tools for categorizing, displaying and searching data. They're tools we continually upgrade to improve processing power, both locally and via the cloud. But computers are simply the medium for big data. "We really need people to interact with the machines to make them work well," says McFarland-Bascom Professor of Electrical and Computer Engineering Rob Nowak. "You can't turn over a ton of raw data and just let the machine figure it out."

Unlike computers, people cannot be upgraded. They work at a finite speed and at rising costs, so Nowak is improving interactive systems that

can optimize the performance of both humans and machines tackling big data problems together.

Typically, human experts—people who categorize data—will receive a large, random dataset to label. The computer then looks at those labels to build a basis of comparison for labeling new data in the future. However, Nowak suggests the model should be flipped. "Rather than asking a person to label a random set of examples, the machine gets the set of examples, then asks a human for further classification of a specific set of data that it might find confusing," says Nowak.

With support from the National Science Foundation and Air Force Office of Scientific Research, Nowak has been exploring an active learning model, in which the machine receives all the data up front. Initially, with no labels, the machine makes very poor predictions, improving as a human expert supplies labels for some of the data. For example, if a new data point is similar to one that a human has labeled, the machine can predict that this point should probably have the same label. The machine can also use the similarities and labels to quantify its confidence in the predictions it makes. And when the confidence for a certain prediction is low, it asks the human expert for advice.

To explore these sorts of human-machine interactions, Nowak and his student Kevin Jamieson have applied the idea to a technology that's a natural fit in Wisconsin—an iOS app that can predict which craft beers a user will prefer.

In this case, the similarities between data points—beers—are based on flavor, color, taste and other characteristics defined by the spectrum of terms used to describe beers in reviews on Ratebeer.com. Using that existing data, the researchers' algorithm can find the closest match for beers the user might enjoy, in much the same way that a bartender might: presenting the user with two beer choices, then using the user's

preference between the two to hone in on a specific point in the "beer space."

"Basically, if I already know that you prefer Spotted Cow to Guinness, then I'm probably not going to ask you to compare Spotted Cow to some other stout," says Nowak. "Because there are relationships between every beer, I don't have to ask you for every comparison."

These sorts of "this-or-that" determinations tend to be more stable than categorizations based on ranking scales or other more subjective measures, which are more vulnerable to psychological priming effects and can change over time. Finer point comparisons offer the machine more reliable data to improve its categorization and prediction over time.

And most importantly, it allows machines to process data much, much faster, since they require less human help to categorize the data. For example, pulling from thousands of possible beers, Nowak says the app can make a personalized beer recommendation based on only 10 to 20 comparisons.

That sort of efficiency becomes important as data sets get bigger and human labor can't keep up. In a collaboration with UW-Madison psychology colleagues, Nowak has applied his model to the relative emotionality of words; without the active machine learning model, learning the similarities between 400 words could require as many as 30 million total comparisons. "Even if you could recruit a cohort of 1,000 undergraduates, that would still be 30,000 trials apiece," he says.

Understanding human judgments about the similarity of word meanings is a fundamental challenge in cognitive science and absolutely crucial in order to make machines capable of understanding the subtleties of human language. Optimizing ways to apply machines and people toward problems like that could be key to making big data analysis economical

and effective in many more situations. "There's no research to be done on the infrastructure side," he says. "We have big data infrastructure. What we don't understand is how to optimally yoke humans and machines together in big [data](#) analyses."

Provided by University of Wisconsin-Madison

Citation: Active learning model for computer predictions (2013, December 3) retrieved 12 May 2024 from <https://phys.org/news/2013-12-active-learning-model-for-computer.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.