Pay Attention to the Man Behind the Curtain

ACM TiiS · Follow
Published in ACM TiiS
4 min read · Aug 13, 2018

By Henry Lieberman
MIT Computer Science and Artificial Intelligence Lab (CSAIL)

Machine learning often seems like magic — and it is.

Machine learning is often applied to prediction problems. In the past, those problems would have been approached by careful prior analysis of the problem domain, trying to analyze the underlying processes, and making rules and representations to perform inference.

With modern machine learning techniques, once an algorithm, parameters, and data sets have established, for a wide range of problems, you can just push a button and get an answer. That answer turns out to be right, a certain percentage of time. It feels like pulling a rabbit out of a hat.

But like the Sorcerer's Apprentice, you have to be a bit careful what you wish for. As with stage magic, the trick may have some limits.

The magician won't explain to you exactly how the trick works. Often, even machine learning researchers can't supply plausible explanations for why their techniques succeed.

The trick may be dependent on certain aspects of the situation which you may not have noticed, but are crucial for the trick to succeed. You may not have been paying attention to the alignment between the bottom of the hat and the table it rests on, but if you move the hat, the trick might not work. It may not work for problems that seem similar but really aren't. The magician may be able to cut a lady in half, but that doesn't mean they could cut a tiger in half.

If you want to learn how to become a good magician, you can't just sit back and be wowed by the trick. You have to dig into the details of how it happens, and what it can and can't do. Successful machine learning takes some knowledge, and a lot of experimentation.

And while magicians do mind-reading tricks, machine learning can't actually read your mind. So there needs to be some communication between the machine, and people who serve as its trainers, implementors, and users.

Recently, a movement has arisen to focus on the human aspects of the machine learning process: Human-Centered Machine Learning.

This TiiS special issue provides an introduction to this topic, ten articles that envision a new partnership between people and machine learning algorithms. They show how judicious use of human interaction can lead to successul applications of machine learning in a wide variety of applications.

The first article is a survey of interaction design for machine learning, by Dudley and Kristensson. This is the place to start, along with the issue introduction by Gillies and Fiebrink. If you only have time for one article, this is it. They deliver six principles that serve as guidelines for UI design for applications. If you're a UI designer tasked with machine learning interfaces, this gives you something to post on your wall.

The rest of the articles cover a wide variety of applications, from social science to medicine to visual and audio presentation. There are nine of them, so there won't be space in this post to even give a summary of each one. See the issue introduction for that.

The breadth of topics covered give you a feel for how each approached the design challenges of human interaction. They might be an eye-opener, if your whole experience with machine learning is reading about mathematical algorithms, and playing the one-upmanship game of comparing accuracy rates on standardized datasets.

Many of them center around the machine learning practice of "labeling" data. This is taken for granted in many machine learning discussions. But the process of asking a person to label something brings up a whole host of issues: labels can be ambiguous or misleading; people may disagree about labels; new concepts may arise. Some introduce innovative interfaces for doing the labeling interactively, presenting the task in user-friendly ways, and providing ways to assure that the labels meet the users' intent. See particularly Chen, Dumitrache, and Kim.

Also necessary is to present the output of machine learning in user-friendly ways. Visualization uses the interpretive power of the visual system to integrate information in the blink of an eye. Presenting results in a manner appropriate to the application context helps give users an understanding of the confidence they should have in results, and how far they can be generalized. See Morrison, Smith, and Francoise.

Not to be neglected is what happens in between the labeling (input) stage and the prediction (output) stage. Many machine learning techniques come with a 747-cockpit of parameters and modeling choices. Most of the articles do treat the question of how to manage this complexity, from the perspective of the machine learning approach they have chosen, and for the application area they are working in. We have to get past the temptation to just treat machine learning as a black box.

Arthur C. Clarke said, “Any suffficiently advanced technology is indistinguishable from magic.” Magic is fun, and it’s undeniably impressive. But it’s Human-Centered Machine Learning that wil get us to understanding machine learning as the sufficiently advanced technology that it is.

Machine Learning

Follow

Written by ACM TiiS

16 Followers · Editor for ACM TiiS

ACM Transactions on Interactive Intelligent Systems (TiiS), a journal dedicated to publishing original research combining AI and human-computer interaction.

More from ACM TiiS and ACM TiiS