Article from

Predictive Analytics and Futurism

April 2018
Issue 17

Chairperson's Corner

By Anders Larson

I know that some of my fellow council members (and at least one former council member) will cringe when they see that I'm leading off the Chairperson's Corner with an anecdote about Microsoft Excel. But stick with me here. About five or six years ago, I realized that there was a formula that allowed you to do a conditional sum-product of two vectors. I was aware of the =sumproduct and =sumifs, but until then, I was unaware that the =sumproduct could be modified to add conditions. Mind = blown.

So what happened after that? I started noticing instances left and right where old workbooks could be improved with this "new" formula. A few years before that, I had a similar experience upon realizing the superiority of index-match functions to vlookups. I started cleaning up existing workbooks, but more importantly, I started thinking differently about setting up new workbooks. Of course, I didn't invent the conditional sum-product or the index-match. I just finally realized they existed, and all of a sudden I became a little bit better at my job.

I believe that actuaries can look at predictive analytics in much the same way. There are algorithms and techniques out there just waiting to be implemented into your existing work. Now, I realize that it's significantly more difficult to get comfortable with a support vector machine than a simple Excel formula, but the concept is the same. Once you start to see how a new approach can fit into one problem, it becomes that much easier to see how it can fit into countless others.

The obvious danger is that it is easy to start seeing everything as a nail once you have a cool new hammer to play with. In general, if a simpler model is just as effective as a more advanced approach, it's best to stick with the simpler approach. One of the key drawbacks I find with many machine learning algorithms is a lack of interpretability, particularly for those who don't work with them on a regular basis. In some cases, that's fine—I don't really care *how* my Amazon Alexa is able to understand speech, but a regulator may not be as willing to accept your estimates if they seem like they came from a black box.

But just because everything isn't a nail, that doesn't mean there aren't nails out there that you've been hitting with a spoon. Sure, the spoon will eventually drive the nail in there, but there's a better tool out there. In our July 2017 newsletter,[1] I wrote about a situation where we used a gradient boosting machine to predict primary care office visit utilization for individual patients. In the past, we might have attempted to predict primary care office visits using an existing risk score algorithm meant to predict health care costs. And while the existing risk score algorithm may have been useful, it was not really the best tool for this job. For instance, the sickest patients in a commercial population can have risk scores that are more than 100 times the population average, but very few patients will have even 10 times as many primary care visits as the population average.

Instead of thinking of each new algorithm as an all-purpose hammer, think of them as new tools to be added to your existing toolbox. Actuaries already have a wide array of traditional tools at their disposal, and those will continue to play an integral role in the future of actuarial science. But we can also improve our profession by incorporating new approaches into our work.

Here's another example from my own experience. I recently co-authored a paper[2] in which we identified the key drivers of gross savings for accountable care organizations (ACOs) participating in the Medicare Shared Savings Program (MSSP). We had more than 180 features about each ACO, many of which were highly correlated with each other. A few years ago, I likely would have approached this problem by limiting the data to a handful of reasonably independent features that I expected would be key drivers, and then running a simple linear regression. This would have still made for an interesting paper, but it likely would have been loaded with caveats that would have softened our conclusions. Instead, we used a random forest to estimate the relative importance of all 180+ features in predicting gross savings. This method allowed us to evaluate all the features together and let the machine identify which were most predictive. There were still caveats, of course—there is no silver bullet for a complex problem like this—but we felt the more rigorous statistical approach added credibility to our findings.

These predictive analytics tools are already out there. They've already been designed, built and tested for us. As actuaries, we just have to pay the small price of learning how to use them (and maybe some Amazon Web Services fees), and we can have them in our own toolbox. ■

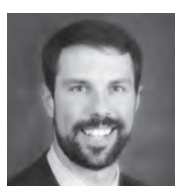

Anders Larson, FSA, MAAA, is an actuary at Milliman in Indianapolis. He can be reached at *anders.larson@milliman.com*.

ENDNOTES

1 *https://www.soa.org/Library/Newsletters/Predictive-Analytics-and-Futurism/2017/june/2017-predictive-analytics-newsletter-issue-15.pdf*

2 *http://www.milliman.com/insight/2017/What-predictive-analytics-can-tell-us-about-key-drivers-of-MSSP-results/*