MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 2 UNIT 2
Casebook Video 2 Transcript

DELIVERED IN COLLABORATION WITH getsmarter

Module 2 Unit 2 Casebook Video 2 Transcript

THOMAS MALONE: Two of the most important ways machine learning is used in the financial world are to manage consumer credit risks and to prevent fraud. Here's Professor Andrew Lo to tell you about these AI applications.

ANDREW LO: Hello. I'm Andrew Lo from MIT, and I'm here to tell you about some of our research on machine learning models of consumer credit risk. Consumer credit is a really big deal – $3.8 trillion of consumer credit outstanding as of July 2017, and 995 billion of that is revolving credit, meaning credit cards or home equity lines of credit, and as of 2017 the average household held about $6,662 of credit card debt. That's a lot of debt. And one of the things that we've been working on is to try to understand some of the exposures that financial institutions have to consumer credit card default and delinquencies. Now, if you take a look at this graph here, which is a graph of standard credit scores across various different years starting from 2005 all the way to 2008, one of the things you'll observe is that, while the credit scores are pretty good at ranking consumers in terms of their likelihood of default – so, for example, at the left side of the graph you've got consumers that have very low credit scores, which means they tend to have a high default rate. On the right side of the graph you have lower default rates, higher credit scores – the credit scoring algorithms that are used in the industry generally work pretty well, but what they don't do very well is to try to see whether or not these credits deteriorate across years.

So, the different lines that you're looking at in the graph are virtually indistinguishable from 2005 all the way to 2007, and that's because from year to year these kinds of credit scores don't change a whole lot, despite the fact that consumer credit conditions can actually deteriorate pretty significantly. So, one of the things that we've been working on is to try to understand whether or not we can use big data for consumer credit. A few years ago, we were approached by a major US commercial bank offering to allow us to look at a very large anonymized data set of their credit card customers. We got data for their credit card transactions, we got data from the credit bureau, but we also got data from their checking accounts and other aspects of their banking information. And from the combined set of data, what we were able to do is to identify certain features that allowed us to tell whether or not a consumer was likely to be delinquent or defaulting on his or her consumer credit card debt. So, it's a very large data set. About a 1% sample was about 10 terabytes, and so it took a while for us to go through this, but the results that I'll show you in a few minutes told us that we were actually on to something.

So, here's an example. On this graph we show in the blue line the likelihood of a consumer becoming 90 days or more delinquent on paying their credit card debt over a six-month period. And what we see is that over this period from 2005 to 2008 there's not a lot of variation in that rate of delinquency, but now if we separate out those consumers who've been getting direct deposits on their checking accounts, and those direct deposits stopped sometime over the last 30 days, that subset of consumers, their delinquency rates are graphed in red, and you can see that there's a big difference. Using this one simple feature, we can actually identify a certain subgroup of consumers that are going to end up having difficulty in making their payments, and while we have lots of different features that we can put together, there's a challenge in trying to figure out how to do that in the best way. That's where machine learning models come in. Using powerful techniques like random forests and support vector machines, what we're able to do is to choose lots and lots of features and combine them in interesting ways to be able to come up with the very best forecast of delinquency and default.

So, in a paper that my co-authors, Amir Khandani and Adlar Kim, and I published in 2010, we described the outcome of this experiment. We had 600,000 credit card customers per month, and within that set what we would do is to apply our particular machine learning algorithm to come up with a likelihood of default, and then we compared that to the usual credit score probability, and that's what you're looking at now in this graph of blue dots. Every single one of these blue dots corresponds to two numbers. On the horizontal axis, we're plotting the typical credit score. On the vertical axis, we're plotting the probability of delinquency based upon our machine learning algorithms, and one of the things that you notice on this graph is that it doesn't look like there's any kind of a relationship between the horizontal and vertical axes. In particular, the higher credit scores, which is on the right side of the graph, don't necessarily have lower probability of default according to our machine learning forecasts. So that suggests that what we're picking up is actually different from the typical credit scores. Now, the question is, is it any good? We know it's different, but is it more accurate?

Now, if I superimpose on this graph some color coding to show you what happened over the next six months, we can begin to see the power of these machine learning models. Color coded in green are those consumers that were not delinquent over the following six months. In blue are the consumers that were delinquent 30 days, in yellow are 60 days, and red are 90 days or more. And so, here's an interesting observation. If we use a vertical line to divide up the consumers into good or bad credits, in other words, if we use the traditional credit scoring models, what we're seeing is that there are some very good customers that actually would be eliminated from the kind of credit provision that banks provide. On the other hand, if we use a horizontal line, in other words if we use the machine learning scores to actually group consumers into good or bad credits, that's much more likely to take the red and separate them from the green and the blue. So, this is a clear demonstration that machine learning adds value well beyond the traditional credit scoring models, and here's another demonstration. If we look at out of sample credit forecasts over time, you can actually see that our forecasts are quite a bit closer to the actual than you would have expected with the traditional credit scoring models.

In our current research we're actually applying this to a number of banks and we're working with co-authors from the Office of the Comptroller of the Currency and allowing them to use our algorithms to see whether or not they can tell various different banks across the entire spectrum and how their credits are evolving. We see the same kind of patterns. If you use all of the data across the six banks, then the credit scores do not distinguish between good or bad credits as well as machine learning scores, but more interestingly, if you apply these models bank by bank, you see that the factors that are driving the particular credit scoring models by machine learning algorithms can be quite different. In other words, these tools allow us to customize these models very easily.

So that's a little bit about what we're doing on applying machine learning techniques to consumer credit risk. If you want to learn more, please visit our website. Our papers are online. Thank you.

THOMAS MALONE: Did you understand all the concepts covered in this video? If you'd like to go over any of the sections again, please click on the relevant button.