Machine Learning and Financial Security: Selecting and Using Data

This brief is the 3rd part of our series on emerging technology and financial security. Read our other briefs, Blockchain Technology and Financial Security & Machine Learning and Financial Security: Defining the Problem, to learn more.

Commonwealth previously explored the importance of defining the desired outputs when developing machine learning technology. For these technologies to be effective in helping build financial security, the problem they are trying to solve must be clearly defined. Even when the use case is clear, financial institutions may still face problems in having the right data to build recommendations on, especially for underrepresented populations. Commonwealth connected with Daniel Shenfeld, a Boston-based AI consultant with experience working in the financial services industry, to explore these pressing data challenges and how financial institutions must account for them.

Machine Learning and its Use of Data

Machine learning is a branch of artificial intelligence that involves machines learning from examples as opposed to working by a set of rules.

Traditional software applications rely on rules to make decisions: "If the applicant's credit score is higher than 650, then approve the loan." But that's not how the human brain works. You can distinguish between an image of a dog and one of a cat in a fraction of a second, but it is very hard to come up with rules to do this task. You can do that because your brain has seen many dogs and cats and trained itself to recognize the difference. In other words, humans learn from examples. Machine learning tries to do the same thing with computers: teach them how to interpret images, text, sounds, and any other source of information using a large number of examples.

Accurate machine learning algorithms require a lot of data, which is just another word for examples. They also require significant computing power to process them. As the amount of data and access to computing power have both increased exponentially in recent years, machine learning and AI research is regularly breaking new ground, and applications to business have become ubiquitous.

To learn from examples, machines need a lot of them. Without enough examples or a broad enough range of scenarios, these technologies can draw false conclusions. These issues caught public attention this year. Companies have shut down their AI-powered bot after it learned offensive responses and stopped use of AI recruiting tools when the system unintentionally taught itself that male candidates were preferable to women. In each case, the algorithms and machine learning tools themselves were not biased. Problems in the underlying data set, such as the historical underrepresentation of women in Amazon's hiring trends, were repeated by the technology because the algorithms and machine learning tools were trained on bad data.

Data and Financial Services

These issues around bias and underrepresentation are just as likely, or even more likely, to occur in financial services. Because those who are the most financially vulnerable also tend to engage the least with their financial institution, there is less data available on the population that these machine algorithms are designed for.

Due to this lack of data, machine learning tools that are applied to make financial recommendations or help build financial security are likely to make mistakes. A tool that is only using data of people who are already banked and financially secure to make recommendations on building savings may not present realistic recommendations to Americans that struggle to make ends meet, have a foot in both traditional and nontraditional financial services channels, etc.. An algorithm that recommends certain financial products may learn from historical inequities and not suggest products for populations that have been historically excluded from using them.

What should a financial institution consider?

Know your data set. Is the input data credible? Is there enough data to have a training set? Even financial institutions that have large amounts of data may not have the right data in the right form

Know the factors that are impacting the machine learning model. A financial institution should look into how a model is making predictions and what factors may be biasing the outputs before rolling it out for broader use. Researchers are continuing to develop best practices on testing algorithm outputs for bias, and a financial institution using machine learning models will need to commit time and resources to staying up to date on the practices and tools available as technology evolves.

Understand consumer trust issues that may arise from using machine learning. Some machine learning algorithms produce outcomes that cannot be easily explained in ways people understand. This is machine learning's "black box" problem. Models in traditional computer programming are, by definition, explainable: a human wrote the if/then statements that a computer is acting on. In machine learning, the rules the model is going to follow are not built out in advance. The machine learning program's inputs and outputs can be seen, but not the process to arrive at the output. This raises issues around consumer trust. If a financial institution cannot explain how a model reached its conclusion, how will consumers trust its recommendation?

How can the broader fintech ecosystem address this problem?

Lower the barrier to entry for data. To build a successful model, a financial institution may need access to data outside of their own institution. The ecosystem can create more partnerships that facilitate anonymous data sharing, and give more players access to the data they need.

Curate relevant data sets. The ecosystem as a whole must make a deliberate effort to include a broader representation of financial resources in the data sets they are using. If a financial institution's own data sets are excluding low income or underbanked users, other sources may need to be pulled in to give a fuller picture. More research needs to be done to understand what types of data are the most fair and the best predictors of financial well-being. Financial institutions will need to continually re-evaluate what types of data are available, as well as stay up to date on techniques to test for unintentional bias.

Build new partnerships to collect data. Efforts have been made by government agencies and academics to collect valuable data on household finances and financial securities. Still, the way these data sets are structured make them useful for understanding large trends, but unhelpful in designing specific tools. New collaborations between academics and fintechs could make the data researchers are already collecting more actionable for building products.

What is training data?

Training data is the first set of data used to help an algorithm begin to make connections and find patterns.The training data must contain examples of the goal the model is trying to achieve.

For example, if a machine is being programmed to correctly identify between images of dogs and cats, the training data would consist of images that are already classified as being of either a dog or a cat. The learning algorithm then takes these examples and tries to find common patterns or themes that would let it correctly identify future images.

It is essential that the learning algorithm has enough training data- with too few examples, or examples that don't cover a wide enough range of scenarios, the algorithm will produce poor predictions. For instance, if the training data only includes examples of poodles, the machine learning algorythm may miscategorize a chihuahua as a cat.

Summary

As the financial services industry embraces data science, it can work to avoid some of the mistakes made by other industries. A focus on quality data inputs and preparation can ensure that machine learning tools are tested for bias. This may involve testing machine learning algorithms for adverse impact on financially vulnerable consumers, or making sure that, just because a consumer type isn't part of a financial institution's current customer base, they aren't left out of data sets that are supposed to be representative.

While this access to unprecedented levels of data holds clear risks, it also opens up new opportunities to correct for practices that have historically blocked people out of the financial system. New data could lead to new ways of assessing credit, offering access to financial products, or understanding the underlying causes of financial insecurity.

To realize this potential, players across the industry - from fintechs to data scientists to academics - have a responsibility to make sure data used in decision making is not only effective, but just.

What steps can we take to ensure that machine learning is trained with adequeate and accurate data to improve the financial security and opportunity of all Americans, especially those who are financially vulnerable? Contact us at info@buildcommonwealth.org to start a conversation.