WHITE PAPER

Socure

Attack of the Fraud-Killer Robots

How Machine Learning and Artificial Intelligence Fight Digital Identity Crime

by Jeff Scheidel, Head of Training & Development for Socure

For decades, it's been said that the way to job security is to learn how to write computer code. Computers run the world, the thinking goes, and we'll always need smart people to make them run things better and faster. We've developed the ability to create code that can make decisions at a speed and scale that cannot be matched by humans – receiving , processing and pushing data from Point A to Point B. One of the perceived values of computer code has always been that it is based on infallible, hard-wired logic, unaffected by external forces. Complex decisioning based on code would therefore be based on equally complex, layered rules, following rigid paths.

When it comes to nuanced, decision-making constructs, this reliance on hardwired rules can become a house of cards. Tiers of rules-based code that flow from one bloc to the next, where each decision is informed by the previous one that feeds it, need to be reorganized with every proposed change. These strict dependencies produce accuracy and precision, but also have a high degree of internal weakness, as changes in any one tier can have a cascading effect.

Perhaps no part of financial services demonstrates this vulnerability better than identity verification, which relies on the use of decisioning software to determine who gets to participate in financial transactions and who does not. Identifying customers prior to allowing them to conduct financial transactions (as mandated in the U.S., for example, by the Patriot Act), fighting fraud, and combatting money laundering is a substantial burden for banks, credit issuers, lenders and other institutions. So how these decisions get made has come under ever greater scrutiny as it has evolved, not only by authorities but also by corporate resources tasked with enforcing regulatory compliance initiatives.

Rules versus Robotic Reasoning

The relatively brief history of identity verification illustrates just how important the move away from a rules-based approach has been. Just a few short years ago, the financial services industry's approach to identity verification was largely based on heuristic, concrete design, highly reliant on human analysis of available data sources. With machine learning still in its relative infancy, risk engines were dependent on human interpretation of an analysis of known outcomes. A small step in the right direction, and one that many organizations still employ, is a form of supervised machine learning, with predictive features still primarily based on human intuition. Therein lies its principal limitation: it requires actual eyeballs.

The rise of neural networks is a topic that gets a lot of column space in articles on artificial intelligence (AI). These constructs, which are meant to recognize words, sounds, images and other data the same way humans do, but faster and better – eyeballs on steroids. But are they actually performing better than humans? And what does better mean? Certainly this would include performing at a speed and volume that even a team of humans cannot match. But to fight fraud, simply doing what people do faster is insufficient. Identity verification and fraud prevention solutions (at account opening) need to learn how to view applicants in new ways, so as to recognize good ones versus bad ones. In this way, compromised identities don't just face faster scrutiny, but also smarter scrutiny.

Identity Fraud: Tailor-Made for AI

Perpetrators of identity fraud are often well-funded, whether by governments or larger criminal enterprises. This access to resources allows them to constantly perfect their methods, and even create entirely new forms of fraud, such as the use of synthetic identities from multiple individuals. The sophistication of such criminals requires an equally sophisticated response. In a sense, you could say that identity fraud is a problem that has been tailor-made for AI and machine learning.

We are already long past a point where credit checks alone should be responsible for keeping out criminals. Basic rules-based solutions have been available for years. They walk through the multiple characteristics of a person's profile, weigh them, and deduce if any number of them signal risk. But purveyors of identity theft and synthetic identities have also long foiled these rules. Synthetic identities are particularly virulent, in that patient and crafty criminals create and nurture these avatars for months or longer, bolstering them with financial history, before setting them in motion to obtain and abuse loans, credit cards, and other types of accounts.

Ben Franklin said, "Tell me and I forget. Teach me and I remember. Involve me and I learn." If you write code based on rigid sets of rules, meaning you simply tell systems explicitly how to process data in a way that's set in stone, you've created a platform that is quickly obsolete. This fragility will need to wait for humans to update the rules as criminals improve their own methods, since the platform cannot teach itself. Therefore the defensive mechanism is not involved, like a robot that marches through a risk landscape and applies what it learns, so much as it is wound up like a watch that always performs its task the same way every time, regardless of a changing environment of fraud.

A long structure of if-then-else statements is essentially a line of dominoes. If you knock over any one of them, the cascading effect can be devastating. Systems based on rigid rules can't learn by themselves. Take for example the automated vetting of a credit card applicant. An individual submits a set of personal attributes, such as his/her SSN, name, phone, address, email and other pertinent data, which is examined by code:

```
if SOCIAL-SECURITY-NUMBER is good,
        then check CREDIT-RATING.
else
        reject APPLICANT
otherwise hey, if CREDIT-RATING is good,
        then if ADDRESS checks out
                then check PHONE-NUMBER
                else etc.
```

There are inherent weaknesses in this kind of flow. To begin with, stacking up these discrete decision points, which don't act in concert so much as perform handoffs, increases the potential of false positives, in that any one of them saying "No" could negate an otherwise collective "Yes." This is especially true when the various decision points are executed by multiple, self-contained software packages or services.

Second, what happens when an organization's auditors decide to change the conditions for approvals, based on changes in compliance requirements – i.e. "Let's start checking out the profile first, before we run the more expensive credit check?" The rules engine must be put up on a rack and modified in a painful process. Franklin's prescient take on this situation might well have been, "Don't apply set-in-stone rules. Teach me what I need to know in order to make a decision, and help me keep learning."

Here we see the limitations of building large systems based on human-based rules, especially for something as complex and fast-evolving as identity fraud: the humans alone simply cannot possibly keep up with evolving risk profiles. When financial institutions think about the engine they are using to digest the massive amount of data coming into the organization as part of the account opening process, it's important to think not only of how they are digesting it, but what they are learning from the digestion process that will allow them to eat smarter.

The next logical step in the evolution of powerful mechanisms that make complex decisions, and replacing fragile layers of hard-coded rules, has been the development of intelligent, holistic, well-informed models. These models ingest the data they are meant to process and, selfishly enough, learn from that data how to do that processing. In this way, the data itself becomes the driver, rather than just the driven.

A process like this is known as Machine Learning (ML), and it is how enterprises are leveraging their own data assets as well as others' to perpetually evolving, automated decision engines.

Machine learning, beyond the buzzwords

There is a large collection of clichés about learning, tired and over-used:

- Experience is the best teacher.
- You learn as much from your failures as your successes.
- A little learning is a dangerous thing.
- And so on.

But these old adages can now be re-applied to a new paradigm. Experience in the form of the processing of relevant feedback is precisely how to create filters against repeating old mistakes. Failures in the form of realized fraud are as important as recognizing the profile of successful requests. And learning once but not continuing to evolve along with fraudsters who change their attack vectors can produce a false sense of security.

Humans can't possibly assimilate new knowledge as quickly as the technology that can learn on their behalf, which is why we build technology platforms that serve as super-powered extensions of ourselves. Systems can learn, at a vast scale, how to discern good prospects from bad prospects, so that when new prospects come along, they are appropriately classified and routed. Instead of performing analysis and then building rules based on what you think you've learned, let the platform analyze, then teach itself.

Fighting digital crime has long been an exercise in realizing the fraud, deducing the patterns so they can be recognized, and retroactively designing a defense. This is a constant game of catch-up. But more advanced institutions, in a supreme effort to keep pace with fraudsters, employ fully automated, robotic machine learning. It's faster, it scales, it sees all the angles, if you show it how and then let it run.

People often confuse machine learning with artificial intelligence. By definition, artificial intelligence is the ability for a computer to simulate intelligent behavior and even reasoning. But that simulation is only possible if the computer is taught how and what to simulate. This is where machine learning (ML) comes in. ML is a subset of AI that utilizes algorithms and statistical analysis to discover patterns and build models that drive AI's responses to inputs. Models are built on predictors, or mini-algorithms that react to the individual elements presented in a focused task. The more data a learning engine can train on, the more accurate its models and predictors. In an open-loop, continual learning mode, the engine keeps acquiring and learning from data, allowing it to keep digesting and learning.

Consider Tesla. It cannot be overstated, the intelligence required for self-driving. Huge amounts of data are fed into the platform (and that's what Tesla is: a platform), to teach it what a stop sign looks like, the differences between a dotted line and a solid one, how to recognize and alert on adverse conditions. In addition, every time a Tesla is on the road, it is absorbing and passing along additional data to the platform, to feed the collective intelligence. That platform is the machine learning piece, which informs the AI piece which actually drives the car. The execution is incredible, but without the data, there are no lessons to be leveraged.

The weather service gathers data via radar, and history, and sensors dropped into hurricanes from storm-chasing planes. This process feeds the artificially-intelligent predictive powers used in recommending preparations, evacuations, emergency response, supply chain, and other actions.

As pointed out by Zeynep Tufekci in the August 2019 Scientific American, Google Translate once consisted of a half-million lines of code, but is now down to a thousandth of that. This was possible because, instead of piles of logic, Google uses as its driver vast amounts of data fed into a machine-learning engine to build and tune models.

In an identity verification and fraud prevention funnel, the predictors would support a model that examines all the characteristics of a person requesting a financial interaction and determine 1) are they who they claim to be, and 2) are they a crook.

The goal is to flip the old approach, where the code is everything, in favor of a process where the traffic cop only needs a whistle, and knows well in advance what the traffic will look like., because his training data set has prepped him adequately. In the case of fraud, machine learning can process large training sets in order to model what a good citizen looks like, as well as how a bad apple behaves. A good citizen's name and address and so on are all connected to a specific individual, reflecting adequate history and quality, and aren't cobbled together from the results of breached databases, the way a faker's are.

More good citizens bypass the friction applied to iffy applicants (such as manual review or "tell me your mother's maiden name or the model of your first car"), and far fewer phony or compromised identities are able to get through the application funnel. Legit individuals may even provide insufficient data but still be recognized, while illegitimates providing perfectly real data can still be rejected, all because a well-trained model knows how to tell good from bad, and may even be able to suggest when to actually recommend opting for friction, but only when necessary.

This exponential leap in the development of automated learning is actually leading to the point where the good guys can keep ahead of the bad guys, instead of just a constant game of leapfrog. Rather than relying strictly on human research and analysis to recognize fraud patterns and then hand-coding the ramparts, Machine Learning does the analysis and powers the AI-based defenses.

Learning and Execution, Hand in Hand

Generalized Artificial Intelligence is meant to handle any task. This is where machine learning is meant to really shine, by processing large amounts of data and teach the intelligent agent how to recognize myriad obstacles and react appropriately. However, devoting that same level of intense learning to applied AI, meaning an intelligence geared toward a specific task, means focusing on specific goals, which further means learning specific lessons using targeted data sets.

Twitter employs AI to thwart hateful, illegal, and fake content; to edit images to make them more appealing; and to better recommend relevant comments. Not to trivialize this, but in the end, the goal is to (simply) shape and deliver content. But given that the gathering and regurgitation of content is its charter, Twitter is well served by this focused effort.

In this way, identity verification and fraud prevention are also well served by an applied focus. The point is not to be all things to all people, but rather 1) ascertain whether an individual is who he says he is, and 2) determine the likelihood that this individual may commit fraud.

In the construction of fraud-fighting models, the AI may participate in more than one place in the process. Naturally it comes in at the end, to produce the decisions that are informed by the copious machine learning and analysis. But intelligent models aren't born from whole cloth.

Not all data is clean and wholesome. It may need to be curated for precision. For example, many financial institutions label their data as first party fraud, third party fraud, synthetic fraud, and they don't always do it accurately. Often, perceived first party fraud, wherein an actual individual fails to pay his bills, is actually identity theft, meaning that the person who stole the identity is at fault, to the detriment of the first party. So the first step in using that data for model training is to make sure it's properly labeled. This allows for analysis against a data set that is relevant.

Multiple models can then be generated by a robotic process that creates its own set of challengers. These models each apply their own decision trees, and are judged by the AI process for which one processes the data the most quickly and accurately.

Another axiom goes like this: "Life is not a destination, it is a journey." Once an accurate model is created, it can be applied to its target function, such as the verification of identities and the prevention of fraud. But it does not have eternal life. There is always more data to ingest, new fraud vectors to recognize and deflect, more lines of business with new requirements, a journey of new commerce, the crooks who dog that commerce, and the corporate crusaders who dog those crooks. Cryptocurrencies with no borders, fintechs relying on sponsor banks, increased digital applications over in-branch interactions, these continually drive changing requirements in this journey.

The famous futurist Alvin Toffler pointed out, "The illiterate of the 21st century will not be those who cannot read and write, but those who cannot learn, unlearn, and relearn." It's not enough to acquire a large data set and pull out the nuggets that power your decision-making. As fraudsters continue to find new vectors to commit crimes and bypass defenses, the models must continue to learn new patterns and defenses, and in some cases unlearn patterns that no longer apply. There will always be the script-kiddie types of attacks that never go away, but the more sophisticated crooks, especially the ones who are funded by foreign powers, continue. To apply their own modes of learning how to pretend to be people they are not.

If a vandal sets off a stinkbomb in a crowded place, he will try to avoid getting caught by slipping out with everybody else stampeding to the exits. This is how third party and synthetic fraud happens. The perpetrators try to look like the rest of the crowd. Therefore the point of machine learning is to examine everybody in the crowd, the good ones and the bad ones, to build the appropriate profiles, so that the AI net lets the good ones through while snagging the bad ones.

One more nugget: if a learning process is limited to only an individual organization's data, the lessons learned are only limited. When possible, the acquisition of a wider samples naturally illustrates more use cases from which to discern patterns.

The Final Tangible, and Less Tangible Benefits

It goes without saying that machines function more quickly than humans. This allows them to operate at a scale that even an army of examiners, such as those who might perform manual review of iffy applicants, or perform credit or other checks, could never hope to approach. That's the easy part. The bias prevention is why someone invented model governance. This is the process by which compliance professionals ensure that models are built and maintained with strict controls for performance and accuracy, and do not inherit any form of bias that might skew the results or cause regulatory concerns down the road.

Deep learning models, such as for fraud, succeed in a way that a stack of if-thens, or a stack of individual product or components, cannot. A self-contained model, informed by ample data, can orchestrate all the salient points about an individual, view that person as a whole, and render a comprehensive decision. Congratulations on your loan (and by the way, start making payments). Or sorry, it's time for adverse action.

Bill Gates says, "Your most unhappy customers are your greatest source of learning." This is true from the standpoint of fraud (up to a point). We learn how to prevent new crime by investigating old crimes. When your fraud engine learns to detect various types of fraud, you know which applicants to reject, which ones to investigate, which ones aren't worth spending the money to investigate, and, most important, how to recognize trends and patterns as fraud not only occurs but evolves. In addition, because fraud losses are classified differently from credit losses, this recognition can also affect accounting practices.

Here's the part that Gates' unhappy people don't provide. While you certainly want to recognize fraud when it shows up asking for a loan, you also want to know what a good profile looks like, for the sake of automatically accepting the non-risky applicants. Friction, in the form of out-of-wallet questions of manual review, slows the process, costs more, and risks a good applicant abandoning the process. Auto-approval of good applicants means a faster start to that happy, profitable, and hopefully long-term relationship between bank and customer. So machine learning doesn't just prevent crime, it increases revenue and creates a better user experience.

Criminals largely have to teach themselves their devious methods, and for years corporate security professionals have had to learn how to defend against those methods. With well-funded intelligence efforts, trained on adequate data and competent learning algorithms, the good guys have the ability to not only stay even with fraud, but get ahead of it, by sending forth their artificially-intelligent, data-trained robots. And it only took us 250 years to learn this from Ben Franklin.

Socure