INTERNATIONAL INSTITUTE FOR ANALYTICS™

RESEARCH BRIEF
RESEARCH & ADVISORY NETWORK

Operationalizing Analytics for Intelligent Fraud Detection and Case Management

MICHAEL AMES, Senior Director, Data Science and Emerging Technology, SAS

ROBERT MORISON, IIA Lead Faculty

JANUARY 2018

Introduction

Fraud is widespread and continues to grow, especially online. It's a major problem in a variety of industries and government agencies far beyond the familiar areas of financial and retail fraud, where credit card information is compromised and fraudsters use it for online purchases. The problem worsens as criminals get more organized and technologically sophisticated and operate at greater scale. Large data breaches expose data about millions of people. Fraudsters automate their activities to exploit thousands of accounts at a time. Account takeover fraud in banking is up nearly 300 percent year over year,[1] web application fraud is up 200 percent, and fraud of government services and payments is up 30 percent.[2]

Objectives of Fraud Solutions

As fraudsters innovate and scale up, fraud prevention and investigation become more challenging, and advanced analytics become a bigger part of the solution. An effective fraud solution has to do three things well:

- ***Surveillance and detection.*** The quest is to improve both accuracy and speed. On the accuracy side, we want to do a better job of identifying suspicious events and cut back on the "false positive" alerts that create unnecessary work for investigators. As fraudsters get more adept, detection has got to incorporate more tools and techniques, including machine learning as well as business rules. On the speed side, we've seen a big push towards real-time surveillance and detection. However, roughly 80 percent of fraud operations are still batch-processed. Part of the push for speed comes from consumers who expect ongoing monitoring of their accounts.

- ***Investigation productivity.*** The fraud solution should provide information and tools with which investigators can work efficiently and effectively. If we can automate their routine tasks and data manipulation, plus prioritize cases for investigation, investigators can be far more productive. They can handle more cases, spend more of their time actively investigating, and focus on the cases with higher likelihood of fraud and higher financial impact. Better workflow also improves cycle time.

- ***Integration and feedback.*** The solution has got to operate as a coherent system, not just a sequence of steps. The entire process must be "instrumented" to track what's happening, what decisions and actions are taken, and what the outcomes are. Those measures are then regularly fed back to improve the analytical models for surveillance and detection, the work of the investigators, and the overall process performance. We've got to understand how the whole system works in order to figure out what can be automated and how to enhance the work of fraud investigators.

[1] McKenna, Frank (2017). Top 10 Fraud Types for 2017 Based on Losses. *frankonfraud*. Retrieved from: http://frankonfraud.com/fraud-reporting/top-10-fraud-losses-for-2016-and-where-they-are-headed-now/

[2] Conroy, Julie (2017). Machine Learning: Fraud Is Now a Competitive Issue. *Aite Group*. Retrieved from: https://www.aitegroup.com/report/machine-learning-fraud-now-competitive-issue

Growing Role of Machine Learning

Machine learning is the automation of extraction of both known and unknown patterns from data. These results are captured in formulas or instruction sets so the patterns can be detected in new data, and they are fed back into the system so that it "learns" and adapts over time.

Traditional rules-based approaches to fraud detection are extremely useful but unequal to today's challenges. They capture experience – the patterns of conditions known to indicate possible fraud – and they're easy to understand. However, one can only encode a small amount of information into a business rule. If you get more than a half dozen or so conditions in a rule, it gets hard to understand. The other drawback is that fraudsters are constantly probing, and they can sometimes reverse-engineer the rules and avoid triggering alerts. Many rules have dollar amount thresholds, such as "If the amount is $9,000 or more, flag this transaction." The fraudster is deflected at $9,000, but then tries and succeeds at $8,500, and has learned how to increase his odds of success.

With machine learning, we can encode large numbers of conditions, variables, and events into models and detect patterns of interest that would slip by the business rules, and would never be noticed by human analysts. The combination of machine learning models and business rules has proven to be much more complete and accurate at finding fraud than business rules alone. When we use ensembles of machine learning methods – including some of the newer ones like gradient boosting, random forests, and even deep learning – the models become extraordinarily accurate.

Machine learning models are also getting more scalable, which means that they can be trained on very large datasets and incorporated into real-time systems. And the final bit of good news is that machine learning tools and techniques have become much more accessible. You don't need a Ph.D. to develop successful models, and there are plenty of educational opportunities in machine learning.

One big challenge remains, however. These models are black boxes, looking at so many things in so many ways that their outputs can be very hard to interpret. Yet the fraud investigator needs to understand the rationale behind a fraud alert and recommended an investigation path. So we recommend building a secondary or surrogate model to present and explain the results, a white box companion to the black box.

The white box may explain the behavior of its companion black box in the form of a scorecard or a set of visuals or an auto-generated narrative of the key conditions indicating fraud. For example, "This account has a very high fraud indicator score of 920. We expect to see accounts of this type to have average daily cash deposits of $1,500 and a maximum of 3x that amount. This account exceeded both thresholds by a factor of five." The objective is to provide the investigator with the data and insight to explore the case – and not be puzzling over how the model works.

Machine learning and business rules are always going to work together. Contrary to what many in the machine learning community seem to think, business rules aren't going to go anywhere. Instead, we see increasing analytic rigor being applied to the combination of machine learning models and business rule sets. We also are seeing more and more institutions using simulations to evaluate the package of machine learning and business rules against different conditions. While business rules have historically been used for workflow and prioritization, we are increasingly seeing machine learning and optimization models being incorporated into these

operational activities in an effort to maximize analyst efficiency in terms of fraud prevention potential.

There is a tradeoff when incorporating machine learning, however. When a business rule can capture the minimum necessary description of a situation or action, we'd rather use the simple rule rather than a complex formula. Keep things simple and understandable, and layer in complexity only as necessary

People and Technology Working Together

A basic objective of fraud solution technology is to enhance the work of investigators and enable them to be more productive. That starts with automation. Confronted with rising fraud rates, organizations have traditionally scaled up both by hiring more investigators and applying technology. But as fraudsters exploit online channels, and automate and scale up their methods, there's no way to keep pace by scaling up headcount. So we need to automate the investigator's workflow as comprehensively as possible:

- Automate data collection and integration so investigators don't spend the majority of their time pulling data from disparate systems. Simply fetching relevant transaction history can have an impact.
- Automate visual data presentation, including with charts and graphs that correspond to the type of fraud – for example, a network in some cases, a timeline in others.
- Automate the preparation and filing of suspicious activity reports and other standard outputs of investigators.
- Automate the workflow itself by prioritizing cases (based on variables including likelihood of fraud, dollar amount, and likelihood of recovery), recommending investigative steps, and fast-tracking straightforward cases.

An automated and informed workflow enhances the work of investigation, and the investigator's feedback makes the overall system smarter over time. Essentially, the system tells the investigator what it noticed and recommends, and the investigator tells the system whether she agrees and why, what decisions and actions are taken, and what the ultimate disposition of the case turns out to be. The most valuable feedback can take the form, "No, this isn't what's really going on."

This feedback loop enables the ongoing training and performance improvement of both models and investigators. By analyzing investigator activities, commentaries, and case dispositions, machine learning models can find ways to automate better. They can also identify the patterns that distinguish the highest performing investigators. The organization can share these as best practices and use them in coaching investigators. Most will embrace the opportunity to improve, especially when they are compensated on the amount of fraud found rather than simply cases handled.

For this closed-loop system to work, investigators and technologists need to work together, often in new ways. If the black box models are built in isolation (no matter how good the data they're trained on), fraud investigators won't trust them – especially if early versions do little more than point out the obvious. Investigators need to communicate with modelers, or modelers need to shadow the investigators, to uncover subtle or intuitive variables and steps in the investigative process. Investigators make the determinations. Models would like to understand how.

For their part, the analytical specialists building models have to be attuned not just to the data and algorithms, but to the user interface and workflow design. Especially when constructing the white box models, the key questions are, "How digestible can we make this? Are we presenting the right things to the investigators for the cases at hand?"

The ambition is to make case management more intelligent. Improve the processes of both detection and investigation. Guide the investigator's thinking, rather than just presenting data and instructions. And close the loop so both models and investigators can learn and improve their performance and value.

Avoiding the Pitfalls

What do organizations tend to get wrong when they're attempting to combat fraud?

One of the most common pitfalls is relying too heavily (or even exclusively) on a single approach, a single type of model, for detecting fraud. The variety of techniques that fraudsters are using today requires a combination of approaches to spotting them. For example, network and relationship analysis is a great tool for finding patterns in insurance claims fraud, but it doesn't detect all varieties of fraud and doesn't lend itself to real-time transactional fraud detection. At SAS, we take a hybrid approach, applying different analytic techniques from multiple disciplines including network and graph analytics and supervised and unsupervised machine learning, as well as techniques from text analytics and forecasting. When you combine multiple techniques, plus business rules, on both the surveillance and the investigation side, you have a more sophisticated approach that can then be tailored to different types of fraud detection.

Organizations may try to improve or accelerate just a piece of the process and lose sight of the end game. For example, a bank's core fraud detection process took an average of 19 hours. The staff was focused on accelerating a key piece of surveillance that ran in about an hour. But what happens if you could make that just a few seconds? You'd still have 18 hours of unnecessarily complex process, 18 hours of chances for something to fail. After simplifying and removing unneeded touchpoints, integration points, and informational clutter, then better automating the streamlined result, the whole process runs in an hour or less.

Many organizations also struggle with data capture and integration, especially when they take fraud detection from batch processing closer to real time surveillance. When you're integrating surveillance data and logic closer to the time of event, you really have to think carefully about what data are available and when, and what constitutes a complete enough picture to build and run your analytics on. It's a different way of thinking, more about latency and sufficiency, and implementation requires different data management technologies and techniques.

Organizations have a tendency to fall into the transactions-per-second or "TPS trap." TPS is simply the number of transactions executed or evaluated per second, a measure of *throughput*. It's an important metric but doesn't tell the whole story because it doesn't address the *latency* of the system. Latency is the calculation of how much time it takes to get a transaction from one point to the next, for example, how long it takes from the swipe of a credit card to receipt of the approve or decline decision.

Consider the latency of the bitcoin blockchain. At present, each new 1MB "block" of transactions is added to the blockchain, every ten minutes. Until this changes, it throttles bitcoin to somewhere in the

region of 7 TPS, and the latency of the blockchain is effectively 10 minutes. Compare that to credit card issuers who regularly handle an average of 2,000 TPS with less than 40 milliseconds latency response. It's not enough to do high throughput – you must make sure you can do it with low latency.

Challenging the System

Maintaining high performance in a fraud solution requires constant attention and adjustment. This starts with monitoring the inputs, outputs, and performance of analytical models to notice when underlying conditions change. Recognizing when models "drift" is basic – but not always practiced – model management. You also want to pay corresponding attention to the performance of investigators and how smoothly people and systems are working together.

An increasingly popular approach to keeping the best models in play is a challenger system. Periodically a new version or variation of a model is developed and tested by running a percentage of transactions through the challenger model and rule set and comparing results with the current champion model. If the challenger proves significantly better, it becomes the champion and the improvement cycle starts again.

The biggest problem in maintaining performance is, of course, that we don't know what we don't know, we don't know what our blind spots are. That makes it very difficult to recognize and adapt to new and emerging threats. In the course of ongoing surveillance, we'd really like to know, "Is this something that we've seen before or not?" We also need the ability to test new hypotheses as well as retest current ones, perhaps newly developed hypotheses about emerging forms of identity theft or account takeover.

Thus, we need to be constantly prospecting, looking for those new and emerging threats. It's like prospecting for gold. Effective gold-miners drill a lot of holes, and once they have identified areas that are most promising, they scale up and get to mining in earnest. And continue prospecting elsewhere.

The technological equivalent is an environment for experimentation, hypothesis testing, and learning what the data has to say. Call it a "sandbox," but the activity in this "play area" can be very sophisticated and rigorous. For example, specific types of unsupervised models can uncover anomalies of interest. There will be a high rate of false positives, but we can take samples of recent transactions and push them through the regular models to see if there is something there. This doesn't take a huge sample of data, but we've got to be able to trace it all the way through the process to determine whether we've hit upon an emerging pattern – whether our prospecting has found a new vein of fraud.

Follow the Leaders

Who is doing intelligent fraud detection and case management really well?

The large credit card processors have set the standard, starting with real-time detection and automatically issuing alerts, including to the customer, or placing holds on accounts. Their systems are typically very well-instrumented. They're using advanced analytics, including machine learning methods, alongside business rules both in surveillance and in case management. Their process may be a well-oiled machine, but they still have to be alert to change. Because chip technology in credit cards has dramatically reduced point-of-sale fraud, criminals are scaling up their online fraud activities.

We also see insurance companies using a variety of techniques, including network analytics, and doing a really good job of uncovering claims fraud and fraud rings where different parties are colluding together. Several government agencies are getting very good at up-front fraud detection, for example, under-reporting (to avoid taxes) or over-reporting (to launder money) in customs transactions. And financial services firms have made a lot of progress recently in combating money laundering. Business rules had tended to flag too many false positives, but with the help of machine learning, organizations can whittle down the number of cases under investigation and concentrate on the most suspicious.

Looking across industries, the leading indicator of high performance is how close the organization comes to operating in real time versus batch mode. The transition to real time involves more advanced analytics and data management, as well as streamlined and effective investigative processes, behind the scenes.

Actions to Take

By way of summary, here are four key actions we've discussed:

1. Incorporate machine learning models in conjunction with business rules in both fraud detection and case management. Neither approach is complete or sufficient on its own, but together they make the entire process faster, more accurate, and smarter.

2. Go that extra mile and explain the results of surveillance. The black box analytical model needs a white box explanatory and advisory models both to gain investigators' trust and to raise their productivity. Keep in mind that detection models, white box models, and supporting documentation all need to be "trained" and updated on the same schedule.

3. Make room for prospecting. Look for new patterns of activity, test new hypotheses alongside revalidating old ones, and listen to possibilities suggested by the data. This prospecting process shapes the ongoing performance of analytical models and the business processes using them.

4. Above all, operationalize fraud detection and investigation as a system, well instrumented with a continuous feedback loop. Make improvement a journey, a regular cycle of progress through evaluating and training the entire system. Along the way, coordinate the management of fraud technology and operations. Optimize the performance of models and investigators together. And maximize the ability to improve performance at every opportunity.

Additional Information

To learn more about this topic, please visit sas.com/fraud.

About the Authors

MICHAEL AMES

Within the SAS Fraud and Security Intelligence Practice, Mike Ames leads the Data Science group, which includes a global R&D team focused on bringing new and emerging technology from SAS to market. Here, he focuses on real-time surveillance and investigation tools for fraud detection and compliance monitoring. Ames holds a BBA in economics and did his graduate work in computer science at the University of Georgia. He also earned an MBA from the University of North Carolina at Chapel Hill.

ROBERT MORISON

Robert Morison serves as Lead Faculty for IIA's Enterprise Research Subscription. An accomplished business researcher, writer, discussion leader, and management consultant, he has been leading breakthrough research at the intersection of business, technology, and human asset management for more than 20 years. He is co-author of Analytics At Work: Smarter Decisions, Better Results (Harvard Business Press, 2010), Workforce Crisis: How to Beat the Coming Shortage of Skills And Talent (Harvard Business Press, 2006), and three Harvard Business Review articles, one of which received a McKinsey Award as best article of 2004. He holds an A.B. from Dartmouth College and an M.A. from Boston University.

IIANALYTICS.COM