Paper 2525-2018

A Fraud Management Solution for Middle Market Banks and Ways to Reduce False Positives

Gourish Hosangady, Aithent Inc.

ABSTRACT

Financial institutions generate enormous amounts of transaction data each day. The pressure on compliance and the need for quick detection of fraud continues to increase. As a consequence, the same institutions need to reduce losses from penalties and fraud. The challenge lies in how best to use a select set of rules coupled with modeling—using data science and machine learning techniques to address this challenge. Suspicious transactions should be flagged with minimal false positives. The process also should maximize productivity and create a degree of seamlessness in both alert creation and investigations. Once compliance and fraud are both addressed, further analysis of customer and transaction data might be performed to gain insights into customer behavior. Such an approach can achieve the following goals:

a) Reduce false positives, achieve cost benefits. This outcome also maintains customer satisfaction as excessive false alerts cause customer attrition for banks, in addition to reputational damage.

b) An ability to create new rules and thus be ahead of the game with respect to fraudsters. Rules can get outdated quickly, so tweaking thresholds and modifying rules is much needed.

c) Create an end-to-end process from alert generation to case management to reporting.

d) Create a closed loop system so that data about true fraud can be fed back into the source data for corrective modeling.

INTRODUCTION

A recent study showed that nearly 20% of consumers who experienced a fraud-related authorization decline had no future spend on that card 6 months after the declining event. Rules pick up certain aspects and capture a theme of a transaction, but don't see the whole story, hence at times, a good customer looks the same to them as a fraudster. With rapid advancement in technology, there are alternatives that can directly address this issue. For example, the use of machine learning and its ability to leverage the power of big data, combined with optimal levels of human intervention and modern technology can provide a faster, more flexible solution that can not only block fraud, but also filter out false positives more intelligently. Therefore, it becomes imperative for companies to adopt this approach quicker. The costs associated with false positives and its financial impact on business is highly underestimated. A recent research report indicated that millions in fraud can be stopped today with just machine learning. Banks are currently facing significant rates of false positives, up to and beyond 85%. The business implications with false positives can be damaging, for example, blocking a transaction and requesting a customer to perform due diligence can lead to lower customer satisfaction and potential attrition. How can a bank optimize its efforts so the level of false positives is neither too high (which increases work load for investigators and increases cost) or too low (when real fraud slips through the cracks causing losses and facing noncompliance?

TYPICAL SCENARIOS

Data shows that half of the staff in banks indicated they capture less than 75 percent of fraudulent activity. About 20 percent of their peers responded that they identified and prevented 90 percent or more cases of attempted fraud. These figures are critical to mid-market banks (by definition, between $10B and $50B in asset size) due to profitability and resource pressures and typically experience the following:

- About 40 percent may be reporting a false-positive rate higher than 25 to 1. (In other words, 95% of transactions that are legitimate are predicted to be potential fraud).
- Rules need to be delivered in a timely manner and ahead of fraudsters changing strategies. Building and modifying rules can be an open-ended task, one which needs to be accomplished within a matter of hours.
- Existing predictive models may not be current and using older data. Alternate, or revised models need to be sufficiently explored and assessed for improved results.
- There is no optimally designed closed loop system. This implies having to capture the fraud status of each alert and feeding the information back to the source data, whether the core system or a data warehouse.
- Not enough insights are generated around patterns of fraud. A thorough analysis of past fraud is needed to gain smarter insights.

RULES

A big part of fraud detection is the deployment of rules that govern alerts. In general, rules define the criteria (using Boolean logic) in flagging potentially fraudulent transactions. The rules use thresholds, or cutoffs, which define the limits within which transactions are considered legitimate. There are no perfect thresholds for each rule and require several iterations before falling to a “steady state”. Managing the rules is essential with some executed in batch mode and others (such as for credit cards) run in real time. Applying excess rules can suppress alerts and cause fraud to slip through the cracks. On the other hand, deploying too few causes too many alerts. The key is to use business intelligence in deciding the optimal set of rules to be deployed in production. Every alert generated should have a link to the rule that triggered it. Over time, the quality and efficacy of rules may be assessed quicker. There are some rules that may be applied universally across products, and others that are specific to each product.

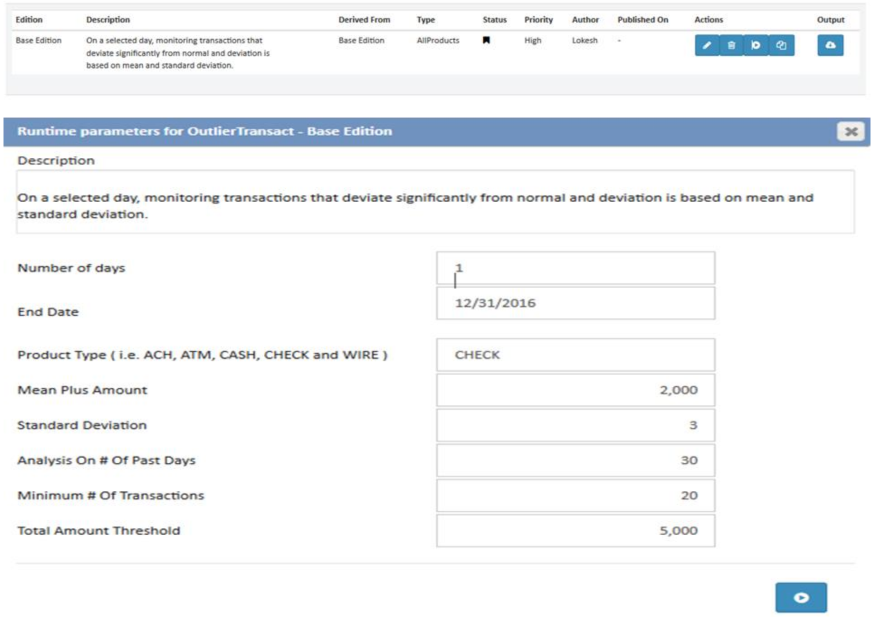

Edition	Description	Derived From	Type	Status	Priority	Author	Published On	Actions	Output
Base Edition	On a selected day, monitoring transactions that deviate significantly from normal and deviation is based on mean and standard deviation.	Base Edition	AllProducts	■	High	Lokesh	-		

Runtime parameters for OutlierTransact - Base Edition

Description

On a selected day, monitoring transactions that deviate significantly from normal and deviation is based on mean and standard deviation.

Number of days	1
End Date	12/31/2016
Product Type (i.e. ACH, ATM, CASH, CHECK and WIRE)	CHECK
Mean Plus Amount	2,000
Standard Deviation	3
Analysis On # Of Past Days	30
Minimum # Of Transactions	20
Total Amount Threshold	5,000

Figure 1: A screen shot of a sample rule with parameters

MACHINE LEARNING

The use of rules to flag fraudulent transactions can have their limitations due to changing behavior by customers. The use of machine learning can only improve detection accuracy. Much of the accuracy depends on the availability of data, both primary and secondary. Today, the abundance of third party data makes for a robust approach to developing machine learning based models to predict fraud. Once the data is defined, the modeling part is less time consuming and involves for example, the SEMMA methodology from SAS for data mining. Multiple models are generated in selecting the best one: regression, decision trees, neural networks and others. Partitioning the data sets is key to delivering quality results. The three data sets are: training, validation and testing. The scores from the model will reflect the likelihood of fraud expressed as a probability. The diagram below indicates part of the data mining process for generating the model. This machine learning model in addition to generating likelihood scores will also help provide root causes of fraudulent transactions through the display of key statistics of the predictor variables. Some of these variables may be statistically significant and others not.

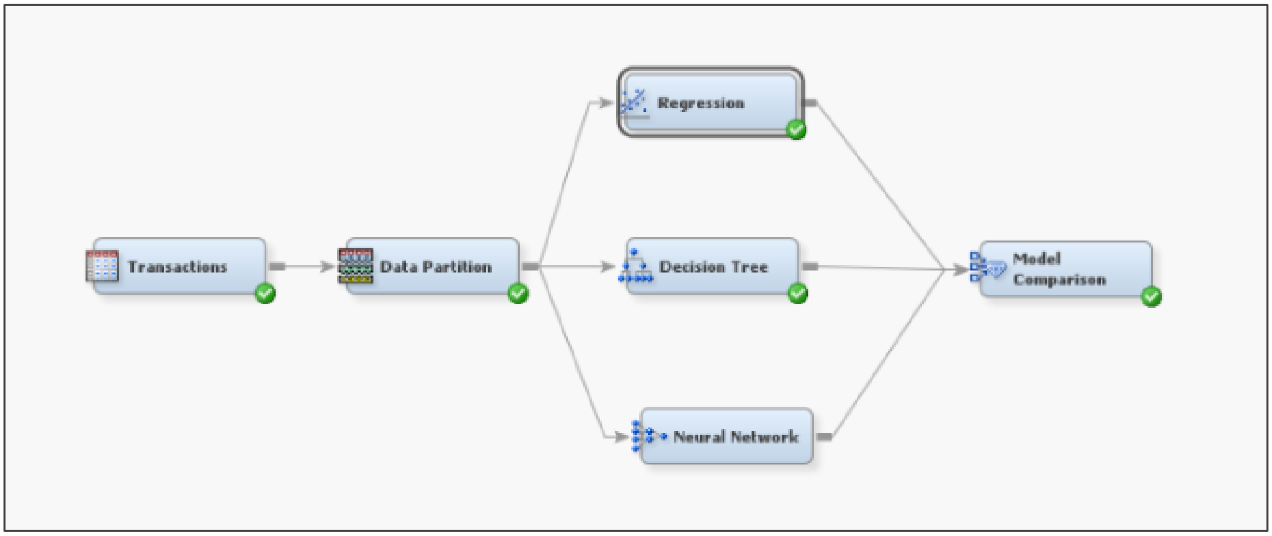

Figure 2: Business process flow for modeling

THE OPTIMAL SOLUTION

Most mid-sized banks are grappling with resource issues as well as with speed of deployment of solutions that are flexible and easy for business analyst to adopt with minimal training. The detection system should necessarily integrate seamlessly with the case management module so that high quality alerts are immediately available to be investigated. The quality of alerts will improve through the use of both rules and machine learning. This is achieved by using the probability score as one of the parameters for rule setting. A sample list of parameters most commonly used are listed below:

- Amount: the transaction amount
- Cumulative Total: a running total of transactions over time
- Time: definition of the time period (in days)
- Product/Channel: product used in the transaction (check, ACH etc)
- Z: a multiplier used for defining distance from the average transaction amount
- Mean Plus: the amount by which the transaction deviates from the average

- Number: number of transactions
- Risk: the risk score for a customer
- Probability: the probability of fraud (from the machine learning step)
- Time Limit: a time period prior or subsequent to which occurrences are monitored (for example a withdrawal within 2 days of a check deposit)

Each rule will have its own set of parameters and thresholds. Choosing the correct thresholds, including for the probability of fraud, will both reduce false positives and improve true positive rates. The more rules deployed for detecting fraud, the tighter it gets since a given transaction will have to meet the criteria on EVERY rule to qualify as an alert. Hence, good judgement is required in selecting the right set of rules. Too many rules will cause fraud to escape, whereas too few rules will generate too many alerts. Setting a threshold for the probability of fraud further controls the degree of false positives.

THE CASE FOR A CLOSED LOOP

A full implementation of a fraud solution comprises the detection and case management components. There are several advantages to implementing one system, one of which is the ability to creating an automated closed loop (see Figure below). The case system is usually the part of the solution that can access the customer information file (CIF). This allows an investigator to capture additional information against alerted customers. All alerts that are triaged in the case system and registered/stored as fraud/non-fraud become a source of information for future rules and modeling. Building a process that automatically and dynamically updates transaction files with such information will create an efficient closed system for keeping up with the goals of false positive reduction.

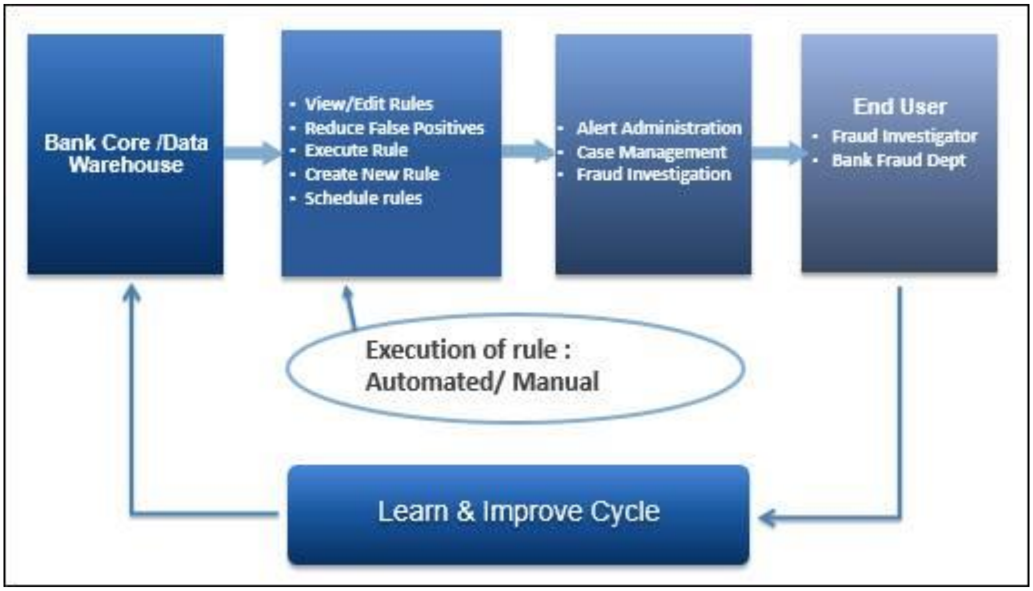

Figure 3: The Closed Loop

RESULTS

Adopting best practices on rules development, machine learning techniques for predictive modeling, deploying an enterprise solution and implementing a closed loop system yields multiple benefits, some of which are summarized below:

- Lower false positive rate
- Managing fraud upstream by focusing on root causes
- Higher flexibility with rules and models
- Reduction in unrecoverable financial losses
- Smarter prioritizing of alert investigations

The confusion matrix below for a sample of dummy alerts illustrates the misclassification rates. The top right-hand corner shows the false positives and represents the percentage of legitimate transactions that were alerted falsely (=10/60). By minimizing this rate, the precision rate (a highly desirable metric) automatically rises leading to a high percentage of predicted fraud that were actual fraud, thus lending testimony to the efficacy of the solution.

n=165	**Predicted: NO**	**Predicted: YES**	
Actual: NO	TN = 50	**FP = 10**	60
Actual: YES	FN = 5	**TP = 100**	105
	55	110	

Table 1: Confusion Matrix

CONCLUSION

The optimal approach to reducing false positives should involve a combination of select rules and modeling, machine learning tools that come together in smart “decisioning” of transactions. The timing is critical too. A big part of the effort is to organize and manage data, both internal and third party. The more data the better the quality of alerts. Establishing the rules with the correct set of parameters coupled with use of the likelihood scores for fraud makes for a robust solution that is both flexible, accurate and easy to use. Both the rule parameters and the predictive model need to be updated to account for changed behavior of fraudsters. Changing parameters for rules is best done with deeper insights from historical data. The closed loop system enables input data to stay current, a requirement for ongoing tweaking of rules and models into the future.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Gourish Hosangady
Aithent Inc.
+1.212.725.7646 ext 1007
ghosangady@aithent.com