Catching Fraudsters In Real Time

Aaron Tietz

aaron.tietz@tufts.edu

Mentor: Ming Chow

Abstract

Unlike physical store retailers, e-retailers are responsible to repay customers for money lost due to fraudulent credit card transactions [5]. The cost of these chargebacks as well as shipping goods for these transactions hurts profits, so detecting fraudulent transactions and stopping them before they're authorized is an important task for any e-retailer. A common technique used by most retailers is to conduct a manual review of all orders that pass certain thresholds, such as reviewing all orders over $300. These techniques work, but miss a lot of fraudulent transactions and have trouble keeping up with changing fraud trends. Third party businesses have begun offering the service of using machine learning algorithms to analyze patterns in real time and predict which transactions are likely to be fraudulent ones. Using large numbers of transaction, customer, and cross-merchant data points, these techniques have been shown to do a better job at detecting fraudulent transactions and adapting to the changing techniques of fraudsters (the industry term for those who commit such transactions).

Introduction

E-commerce retailers protect against fraud by using manual reviews of pending transactions, with the transactions to be reviewed being flagged by hand-crafted rules systems (e.g., orders with three or more of the same item might be flagged). Around a quarter of all e-commerce orders are manually reviewed [3], with 75% of reviewed orders being accepted [3]. This results in a lot of wasted time and money for retailers, as these reviews take around 5 minutes per order [3]. Machine learning techniques allow for ruling out false positives (transactions that seem like they might be fraudulent but are not) faster and more effectively than humans, so that employee time can be spent looking at the transactions that are truly suspicious. Manual reviews also do a poor job of stopping fraud, as those orders reviewed and approved are five times more likely to be fraudulent than industry averages [3]. Machine learning techniques can help with this as well as they are also more effective in determining false negatives (orders that appear honest but are actually fraudulent).

The machine learning techniques used are supervised learning techniques, which means that they require human reviewers to teach algorithms how to be more effective. In order to learn what data points and patterns most suggest fraudulent transactions, humans are needed to label some of the transactions as good or fraudulent. The more examples of good and bad transactions provided by humans, the better the algorithms are at finding fraudulent transactions in real time. Machine learning analysis of customer and transaction data is therefore a very effective aid to, as opposed to a replacement for, the manual reviews e-retailers already conduct.

To The Community

This is an important topic because e-commerce represents a huge and fast growing global industry. In 2013 it stood at $580 billion and was growing at at rate of 17% annually [1,2]. Unlike physical retail stores, online (or, card-not-present) merchants are responsible for confirming shopper's identities [5]. As a result, merchants are responsible to foot the bill for fraudulent transactions, and so lose far more than banks and customers due to online fraud each year (10x and 20x, respectively)[5]. In 2012, it was estimated that $3.5 billion dollars were lost to e-commerce fraud in North American markets alone [3]. Also in 2012, the fraud rate by order was 0.8% for domestic orders, and 1.6% for international orders [3]. This may still seem like a small number, but if retailers have chargeback rates above 1% for a month, they can be put on probation and eventually lose the rights to accept credit card transactions [5]. Therefore, there is a great deal of motivation for e-retailers to catch fraudulent transactions before approving them.

Additionally, this topic is important because stories of stolen credit card data are sadly the norm. Storing sensitive customer data more securely is vitally important, but a world in which credit card information is never stolen is unrealistic. Therefore, developing and using effective strategies for guarding against the use of stolen credit card data is equally important. While businesses have developed fairly effective strategies to detect fraud, machine learning techniques have proven to be more effective than manual review of orders using rules-based systems. Preserving trust between retailers and the transactions their customers make, and reducing profit lost to fraudulent transactions, is important to sustain the health and growth of this large and growing industry.

Applications

E-retailers have historically used hard-coded rules-based systems to have their computers flag orders for human review. These rules are developed over time by looking at internal data and using the experience of the employees who review orders. An example rule might be flagging all order that have a different billing and shipping address and are over $200. These rules work because there are general trends in fraudulent orders. For example, in 2012 the average fraudulent order was $200, compared to $149 for non-fraudulent orders [3]. However, orders approved by manual review are five times more likely to be fraudulent than orders as a whole [3]. If businesses are able to have better information around which orders are worth reviewing, and on the true fraud likelihood of those orders, employee time and money lost to fraud could be used instead to grow the business.

Machine learning techniques have been shown over the last few years to be a more effective fraud detection aid than rules-based systems. By analyzing millions of data points taken from the lifecycle of customers and transaction attempts, these algorithms can determine which pieces of information (or, signals) best determine the likelihood of fraud. The main types of information used by machine learning approaches are customer and order data, customer-website interaction data, velocity checks, and device fingerprinting. By using a combination of all of these, a even more accurate assessment of an individual transaction can be made.

The more customer and order data (e.g., billing and shipping addresses, email address, etc.) stored for a given customer and transaction attempt, the more opportunities to match with patterns of previous fraudsters. For instance, a shipping address that is extremely far from a billing address can signal potential fraud. As a counter example, historical data points may inform the algorithm that this customer has used the address before, and so the order might be

deemed unlikely to be fraudulent. Another example is that a high number of users linked to the same billing address or device can signal a fraudulent transaction. Going beyond basic customer account data, any interaction a customer has with a website can be stored and used for analysis. This includes login attempts, account information changes, pages viewed, and so on. This type of data allows for the flagging of suspicious behaviors such as logging in and directly visiting the checkout page multiple times without browsing a site [6]. By combining customer, order, and customer-site interaction data, a well-defined assessment of the transaction can be made.

Velocity checks are another effective measure for detecting when fraudsters make transaction attempts. One type of velocity check relates to the size of an order. While ordering one or two of something is commonplace, ordering 20 watches can indicate fraud. A second type of velocity check relates to the frequency of use of any parameter involved in the purchase: credit card number, IP address, e-mail address, shipping address, and so on. If any of these items are used a high number of times within the same day, for example, then any new transactions involving them might be flagged as suspicious. Velocity data for a transaction parameter can be compared to averages for an industry, a retailer, and a specific customer profile to make the most accurate assessments possible.

Another type of useful information comes from identifying the device being used to make the transaction attempt. Device fingerprinting, the name for gathering this information, is a relatively new and effective technique for linking devices to fraudulent purchase histories. By detecting things such as operating system, browser version, screen size, browser and system settings, and more, a fairly unique picture of a device can be made that's more powerful than just using IP and customer account info alone. This can allow algorithms to effectively find fraudsters who attempt to use multiple ip addresses or customer accounts from the same device [5]. While fraudsters may adapt in time to make it more difficult for device fingerprinting to be

used effectively, the strength of supervised machine learning is that it can grow and adapt to recognize and thus combat those new trends.

Machine learning techniques benefit from having a lot of data. This means that when a business starts using a machine-learning-based fraud detection system, it can take a few weeks for the system to accurately flag transaction attempts as suspicious or not. However, when e-retailers use third party fraud detection services, they can also benefit immediately from the fact that those services pool information across multiple sites [5]. When a piece of customer information like an e-mail address gets flagged as suspicious on one site, it can then be more likely to be flagged on another site that shares that database, as the likelihood of being linked to a fraudster has increased [6]. Over time, the algorithms used will find a good balance between retailer-specific and industry-wide data points creating unique, quickly creating the most relevant rules for each retailer.

Flagging transactions in real time (as well as adapting to changing fraud trends) requires constant recalculation of customer and transaction scores. Every time a new piece of data is sent to a third party provider using machine learning techniques, the relevant database is updated and a new score is determined. When a transaction attempt is finalized (i.e. the order is submitted), a score is ready and waiting so that the retailer can make an instant decision to either put the order on the queue of those to be manually reviewed or to go ahead and approve the order. As manual review takes an average of five minutes per order, and customer satisfaction is important, being able to act as quickly as possible is a requirement.

The type of machine learning used by these systems is supervised machine learning. This means that the algorithms need feedback from humans on what are good and fraudulent orders in order to improve. When algorithms suggest that a transaction might be fraudulent, reviewers can indicate whether the algorithm was correct or not. This might involve using data the business has about a trusted customer, making a phone call to verify the purchase, or

another method. The most effective human feedback is identifying false positives and false negatives. False positives are transactions flagged as likely fraudulent that are not fraudulent, while false negatives are the opposite. As these are instances where the algorithms makes mistakes, information provided to correct them is extremely helpful. Manual reviewers and machine learning algorithms, then, work in harmony to improve each other's effectiveness. Reviewers teach algorithms how to be more accurate, and those algorithms in turn flag fewer false positives, freeing up reviewer time to focus on those transactions that actually matter.

Conclusion

While retailers have developed methods for reviewing pending transactions for potential fraud, human reviewers using rule-based systems do a relatively poor job of catching fraud. In the last few years, machine learning techniques have been applied to large amounts of customer data points in an attempt to quickly detect patterns of behavior that suggest fraud. These techniques have proven to be more effective than rules-based systems, both at determining if a transaction is fraudulent and at adapting to changing trends of fraudsters. Preventing fraud in online transactions will always be a cat and mouse game between retailers and fraudsters, but the techniques described above have show great promise. By analyzing transaction data in real time and comparing it to large databases of fraudulent and non-fraudulent transactions, machine learning methods have shown that they can catch fraudsters in the act more effectively than previous methods.

References

1. ATKearney. The 2013 Global Retail E-Commerce Index. 2013. http://www.atkearney.com/consumer-products-retail/ideas-insights/featured-article/-/asset_publisher/KQNW4F0xInID/content/online-retail-is-front-and-center-in-the-quest-for-growth/10192

2. Cushman & Wakefield. Global Perspective On Retail: Online Retailing. July 2013. http://www.cushmanwakefield.com/~/media/global-reports/Global%20Perspective%20on%20Retail%201st%20July%202013.pdf

3. CyberSource. 2013 Online Fraud Report: Online Payment Fraud Trends, Merchant Practices, and Benchmarks, 14th Annual Edition. 2013. http://www.cybersource.com/resources/collateral/Resource_Center/whitepapers_and_reports/CyberSource_2013_Online_Fraud_Report.pdf

4. CyberSource. Mobile Payment Management Trends 2012-2013. 2013. https://www.cybersource.com/resources/collateral/Resource_Center/service_briefs/Mobile_Trends_DS.pdf

5. FirstData. Strategies for Reducing the Risk of eCommerce Fraud. October 2010. http://www.firstdata.com/downloads/thought-leadership/ecommfraudwp.pdf

6. Roush, Wade. Sift Science Uses Machine Learning to Weed Out Credit Card Fraud. Xconomy. August 22, 2013. http://www.xconomy.com/san-francisco/2013/08/22/sift-science-uses-machine-learning-to-weed-out-credit-card-fraud/2/