Adaptive Stress Testing for Adversarial Learning in a Financial Environment

Khalid El-Awady*

July 9, 2021

Abstract

We demonstrate the use of Adaptive Stress Testing to detect and address potential vulnerabilities in a financial environment. We develop a simplified model for credit card fraud detection that utilizes a linear regression classifier based on historical payment transaction data coupled with business rules. We then apply the reinforcement learning model known as Adaptive Stress Testing to train an agent, that can be thought of as a potential fraudster, to find the most likely path to system failure – successfully defrauding the system. We show the connection between this most likely failure path and the limits of the classifier and discuss how the fraud detection system's business rules can be further augmented to mitigate these failure modes.

1 Introduction

Transactional volume facilitated by branded and domestic payment cards (credit, debit, and prepaid) amounted to \$42.274 trillion in 2019, up 4.2% over 2018 [1]. Gross fraud losses on these cards reached \$28.65 billion, up 2.9% from \$27.85 billion in 2018. This amounts to 6.5 cents on every \$100 of payment volume (or 6.5 bps). Further, LexisNexis reports that the payment card industry collectively incurs an additional \$2 for every \$1 of fraud in detection, prevention, and remediation efforts [2]. The secular trend in the growth of share of online payments, where fraud rates average 3-4$\times$ those of in-person transactions, is further accelerating these costs.

Approaches to detecting credit card fraud run the gamut of traditional machine learning techniques, including some of the earliest commercial implementations of neural networks [4]. Yet the field remains active in both research and industry. This can be attributed to the adaptability of fraudsters who have the luxury of testing and learning with an abundance of cheap stolen user data [5] and in an online environment that is ideal for programmatic and anonymous trials. This requires the financial system to constantly update their defenses and employ novel techniques.

Adversarial learning, in which the actions of the fraudsters are explicitly modeled is a relatively new approach to payment card fraud detection. This approach incorporates an adversary's potential strategies when attempting to improve defenses. In this report we present an adversarial fraud detection and prevention approach using reinforcement learning. We formulate a model for fraud classification from publicly available credit card fraud data and model the fraudster's actions as

---

a Markov Decision Process (MDP). To devise the fraudster's optimal policy we will explore the application of Adaptive Stress Testing to this environment to see what it might offer.

2   Related Work

Most work in the area of adversarial learning in payment card fraud credits Liu and Chawla [6] as one of the earliest inspirations for this approach. In that work, the authors present a game theoretic approach to email spam detection. Spam detection shares important similarities with payment card fraud including large class imbalances, a similar amount of high-volume observations that can be generated at low cost, and the presence of malicious actors that want to fool the classifier. This work and its predecessors showed that a spam classifier that takes the adversary's best actions into account greatly increases their model's predictive power.

In [7] and [8] Beling's teams present a reinforcement learning approach to credit card fraud detection that includes an MDP model of fraudster decisions and an active defense by the card issuer. The authors utilize the learned adversary policy to disrupt the adversary's learning process, altering classifier probability classification thresholds at regular intervals to effectively limit the amount of successful fraud. One drawback of their work is the artificial action and reward spaces they constructed for the fraudster. In [7] the entire data set is split into 3 bulk clusters using Gaussian Mixture Models and the fraudster is allowed to choose as its action a transaction from the cluster that is most profitable. In [8] the action space is limited to a binary choice of a low $-amount transaction of a high $-amount transaction. There is also no accounting for the element of time in fraud detection, which is a critical feature of any payment card fraud detection system (e.g., through the imposition of rate limits on transactions).

Recently, Lee et al. [9] have a presented a new approach termed Adaptive Stress Testing (AST) that offers a framework for finding the most likely path to a failure event in a system modeled by an MDP. Discovering failures that are most likely to occur is an intriguing enhancement and offers a natural priority for adding incremental mitigation measures. While there is literature describing the application of AST in areas such as collision avoidance systems for air traffic control [10] and autonomous driving [11], we have not found any significant literature on its use in a financial setting.

3   Problem Statement

Adversarial learning in fraud prevention has been shown to increase effectiveness over static models that do not account for changing fraudster behavior. Further, reinforcement learning has been demonstrated to be an effective way to learn a fraudster's attack strategy with an eye towards improving defenses. We show how the performance of a fraud detection system can be enhanced using Adaptive Stress Testing to model a fraudster's behavior and mitigating likely strategies they might employ.

4   Approach

Figure 1 shows the reinforcement learning model of our problem. The system represents an approver of an attempted payment card transaction (typically a card issuer). The system receives a request to authorize a payment from a merchant and classifies it as either fraud or not fraud based on

learned historical patterns. The system also applies certain business rules such as limits on the number of accepted daily transactions.
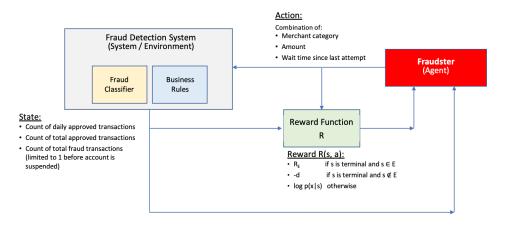

Figure 1: Reinforcement learning model for adversarial learning of the payment card fraud detection system.

In the AST approach, the learning problem is setup to determine the path of most likely failure (fraud). The agent (fraudster) interacts with the system to learn the 'best' way to commit fraud. In fact this mimics practice closely where it is common for fraudsters to purchase batches of stolen cards test and learn where vulnerabilities are (often with low dollar transaction), and then commit high dollar fraud.

4.1 The System: A Model for Payment Fraud Detection

Modern payment card fraud detection systems utilize at their core a statistical classifier that is trained on historical transaction data, some of which have been reported as fraud. This will use the combined attributes of the available payment data elements to determine whether to approve or decline a payment authorization request. In addition, these systems will also incorporate business rules such as imposing hard limits on amounts of individual transactions and frequency of transacting (amongst other rules). The systems also lean toward conservatism – at the first strong indication of potential fraud the account is generally suspended and prevented from further transacting.

In our simplified system model we define the **state** to include 3 variables:

1. Cumulative approved transaction count on the card: this tracks the usage of the card and correlates with the age of the card.

2. Cumulative daily transaction count on the card: this is a key fraud control parameter and is generally limited to a certain number (say 10) as a stop gap against an undetected fraud run.

3. Indicator if fraud has been detected: this represents the output of the classifier. It is assumed that the first detected instance of fraud results in immediate suspension of the card.

To train a fraud classifier we located a appropriate public dataset on Kaggle with sufficient richness and transparency for our exercise: https://www.kaggle.com/kartik2112/fraud-detection. We reduced this to a set of key payment transaction data fields which are either necessary for

3

modeling purposes or are known (from experience) to have a strong correlation to fraud and are subject to fraudster influence: the payment account number, a merchant category index (one of 14 possible values – see table 1), the transaction amount, and the elapsed time interval since last transaction.

Finally we augment our system with a set of simple **business rules** that reflect common practice:

1. The card is suspended at the first detection of a fraudulent transaction.

2. The card is suspended if the number of daily transactions exceeds 10 transactions.
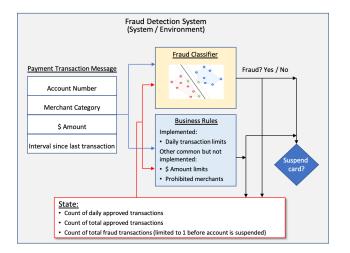

Figure 2: The fraud detection system.

The system is depicted graphically in Figure 2. Note that with these assumptions a fraud episode (using a single payment card) ends whenever the system detects fraud or if the fraudster performs 10 transactions in a single day.

4.2 The AST Adversarial Model

Now we detail the actions and rewards of the adversarial agent (fraudster) of figure 1. Table 1 shows merchant categories available in our dataset and the associated statistics on overall and fraud transaction amounts.

The table shows that fraudsters prefer merchants such as online shopping. This is due to their convenience and the ready availability of high value fencible goods. Fraudsters generally also prefer to accumulate as much fraudulent activity in as short a time period as possible as they know detection is just a matter of time. Fraudsters also generally know the approximate age of a stolen card and therefore can choose the age of a card used in a fraud attempt. These behaviors motivate the following choice of **actions**:

[card age]×[merchant category]×[amount]×[interval between transaction].

The form of the **reward function** distinguishes AST from other reinforcement learning approaches. It is designed to find failure events as the primary objective, and maximize the path likelihood as a secondary objective [9]. We define the following:

| | Overall | | | Fraud | | |
Merchant Category	Portion	$ Mean	$ Stdev	Portion	$ Mean	$ Stdev
entertainment	7%	64	64	3%	510	74
food / dining	7%	51	48	3%	122	14
gas / transport	10%	64	16	7%	12	5
grocery online	3%	54	23	2%	12	3
grocery in person	9%	116	52	23%	313	27
health / fitness	7%	54	48	2%	20	2
home	9%	58	48	3%	258	47
kids / pets	9%	58	49	3%	20	3
misc online	5%	79	164	13%	804	87
misc in person	6%	62	134	3%	193	316
personal care	7%	48	49	3%	26	12
shopping online	8%	83	237	24%	994	95
shopping in person	9%	77	232	10%	887	131
travel	3%	112	596	0%	9	2

Table 1: Comparison of overall vs fraud payment transaction statistics.

- **Terminal state**: if fraud is detected or if 10 transactions are completed in a single day.

- E the **event space** of interest: the fraudster is <u>not</u> caught.

- R_E the **reward if the fraudster is successful**: given by the total accumulated fraud amount: $R_E = \sum_{\text{fraud trx}} v_k$ where v_k is the amount of fraud transaction k.

- d the **miss distance**: this represents a penalty for being caught. Through experimentation we find that setting this to a large negative constant works best. https://www.overleaf.com/project/60b6f393891￼

- $p(a|s)$ the **likelihood of an action in a given state**: this is modeled from the data set by assuming independence of the action components and examining the histograms of each in the dataset where we find:

$$
\begin{aligned}
p(n_{tot} = \text{total card trx}) &\sim \exp(\lambda = 1/(390 \text{ trx})) \\
p(m = \text{merch categ}) &\sim \text{as in column "Fraud - Portion" of table 1} \\
p(v = \text{amount}|m) &\sim \mathcal{N}(\mu_i, \sigma_i^2) \quad \text{per table 1,} \qquad\qquad i \in \texttt{enum}(\text{merch categ}) \\
p(\delta_i = \text{interval}|m) &\sim \mathcal{N}(\mu_i', \sigma_i'^2) \quad \text{similar to table 1 but not shown,} \quad i \in \texttt{enum}(\text{merch categ}) \\
\Rightarrow p(a|s) &= p(n_{tot}) \cdot p(m) \cdot p(v_i|m) \cdot p(\delta_i|m)
\end{aligned}
\tag{1}
$$

With these the reward function is given by

$$
r(s,a) = \begin{cases} R_E & \text{if } s \text{ is terminal and } s \in E \\ -d & \text{if } s \text{ is terminal and } s \notin E \\ \log(p(a|s)) & \text{otherwise} \end{cases}
\tag{2}
$$

To reflect the strong preference of fraudsters for quick profit we employ a discount factor that strongly discounts fraud that takes more than a day to collect. Thus we assume a daily $\gamma = 0.2$.

Algorithm 1 AST Q-learning

Require: Action space $\mathcal{A} =$[card age] $\times$ [merch. categ.] $\times$ [amount]$\times$[interval bet transactions],
   action selection rule $\pi(\cdot;\cdot) \sim U[\mathcal{A}]$, step size $\alpha \in (0,1]$, max daily transactions $T = 10$
 1: Initialize $Q(s,a)$ for all $s \in \mathcal{S}, a \in \mathcal{A}$
 2: **for** each episode, $n = 0, \ldots, N$ **do**
 3:    Initialize the state, s_0: agent randomly selects a new or old payment card account to use,
   daily transactions are set to 0, and the fraud indicator set to 0.
 4:    **for** $t = 0, \ldots, T$ **do**
 5:       $a_t \leftarrow \pi(s_t; Q)$
 6:       Determine $\log p(s_t|a)$ according to equation (1)
 7:       Classify the resultant payment transaction as fraud / no fraud
 8:       Determine reward r_t according to equation (2)
 9:       Observe next state x_{t+1}
10:       $Q(x_t, a_t) \leftarrow Q(x_t, a_t) + \alpha(r_t + \gamma \min_{a' \in \mathcal{A}} Q(x_{t+1}, a') - Q(x_t, a_t))$
11:    **end for**
12: **end for**

---

5   Experiments

5.1   Fraud Classifier

The dataset comprises approximately 1.2 million training transactions across several thousand accounts, and an approximately 500,000 transaction test set, both with fraud rates under 0.4%. This represents a highly imbalanced dataset. We utilize the synthetic minority oversampling technique (SMOTE) to mitigate this [12], [13]. We first oversample the fraud transactions in our dataset to where they form 10% of the samples, then create additional synthetic data such that our final training set is 1/3 fraud and 2/3 non-fraud. This approach helps to improve the classifier's predictive capabilities [14].

For simplicity and transparency, we chose to use a logistic regression model with the limited features detailed in section 4.1 as our classifier. With this classifier we get the following confusion matrix on our validation dataset:

- Accuracy: 96.8% – these are transactions correctly identified as either fraud or not fraud.

- Decline rate: 3.85% – these are good transactions mis-identified as potential fraud and erroneously declined and represent missed revenue for merchants.

- Uncaught fraud rate: 0.1% – these are fraudulent transactions not caught by the classifier. Of the attempted fraud this represents about 26% (meaning the classifier catches about 3/4 of attempted fraud).

Note that the above values are in line with commercial fraud detection performance.

5.2   AST Experiments

We employ algorithm (1) to estimate the Q function associated with this reinforcement learning environment.

For our first experiment we discretize the action space as

- Initial card total transactions: {10 (new card), 500 (slightly over average age card)}.

- Merchant categories: {entertainment, shopping online, shopping in person}. These are chosen due to their affinity to fraud activity (fencible products, potential for high $ ticket items)

- Amounts: {$10, $100, $1,000}

- Intervals between transactions: {1 min, 50 min, 150 min} (all same day transactions)

This represents an action space of size $|\mathcal{A}| = 54$. Figure 3 shows the convergence of the Q function as measured by the Frobenius norm of the change over iterations. The resultant most likely fraud path from the learned policy is as shown in Table 2.
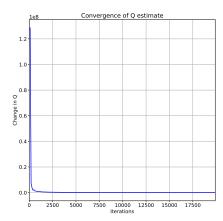

Figure 3: Convergence of the Q-function for experiment 1.

	New Card			Middle Age Card		
Daily tran	Category	$ Amount	Int. (min)	Category	$ Amount	Int. (min)
1	Entertainment	100	150	Shopping online	100	50
2	Shopping online	100	150	Shopping online	100	150
3	Shopping online	100	1	Shopping online	100	50
4	Shopping online	100	1	Shopping online	100	150
5	Shopping online	100	1	Shopping online	100	1
6	Shopping online	100	1	Shopping online	100	50
7	Shopping in person	100	50	Shopping online	100	1
8	Shopping in person	100	1	Shopping online	100	50
9	Shopping in person	100	50	Shopping online	100	150
10	Entertainment	100	150	Entertainment	100	150

Table 2: AST most likely fraud path detected for the first experiment (no new business rules).

Inspection of the results reveals that the fraudster will always revert to a moderate size transaction amount ($100) and mostly use the online shopping channel. There is a diverse range of intervals utilized. We might intuit from this that the fraud classifier does a good job of detecting "typical" fraud such as large amount transactions. But if a fraudster chose to pursue a repetitive moderate amount transaction strategy the system may not catch it.

It is difficult to modify or shape the logistic regression classifier behavior to attempt to mitigate this fraud strategy. But the business rules are well equipped to handle this. A rule can be added to monitor for long successions of similar transaction sizes. In fact we do this and re-run the experiment and obtain the most likely path strategy of Table 3. Most notably here we see the fraudster must modify the transaction amounts in this case.

Daily tran	New Card			Middle Age Card		
	Category	$ Amount	Int. (min)	Category	$ Amount	Int. (min)
1	Shopping in person	100	150	Shopping online	100	50
2	Shopping online	1,000	150	Shopping in person	100	150
3	Shopping online	1,000	1	Shopping online	1,000	50
4	Entertainment	1,000	1	Shopping online	1,000	150
5	Shopping online	100	1	Shopping online	100	1
6	Shopping online	1,000	50	Shopping online	1,000	50
7	Shopping in person	100	1	Shopping in person	100	1
8	Entertainment	1,000	1	Shopping online	100	50
9	Entertainment	100	150	Entertainment	100	150
10	Shopping in person	100	150	Entertainment	1,000	150

Table 3: AST most likely fraud path detected for the second experiment with rule preventing lengthy successions of similar transaction amounts.

This second experiment also plants the seed for subsequent study, namely utilizing this framework in an interactive game setting where the fraudster develops a strategy, the system mitigates it, and iterations continue until a more refined fraud detection system is developed.

6 Conclusions and Future Work

In this report our contributions include:

- Casting the Adaptive Stress Testing paradigm as an adversarial learning problem for payment card fraud detection.

- Creating a simple but realistic payment card fraud detection system that incorporates a statistical classifier and a set of business rules to mitigate fraud and that resembles systems used in practice by card issuers.

- Showing the most likely fraud path as a sequence of actions that can be taken by a fraudster to maximize their profit and providing an intuitive explanation why it works.

- Modifying the fraud detection system through the addition of a business rule to prevent the identified strategy and showing that it results in a change in fraudster strategy.

- Motivating continued study of the use of AST in an interactive game-like setting between the system and fraudster to fine-tune the system's performance.

Future work would mainly revolve around exploring the interactive game-setting implementation of this adversarial learning approach. Some areas of specific investigation might include: what is a process for automating a business rule in response to a fraud strategy, and how to we protect against unintended consequences of new business rules. Also, the fraud classifier is ripe for additional sophistication utilizing more data fields and features and more powerful models such as a neural network.

References

[1] Nilson Report. Issue 1187, Dec 2020.

[2] LexisNexis Risk Solutions. 2020 True Cost of Fraud Stud yE-commerce/Retail Report. July, 2020.

[3] Smart Insights. Forecast e-commerce growth in percentage of online retail / e-commerce sales 2017 to 2023. April 6, 2021. https://www.smartinsights.com/digital-marketing-strategy/online-retail-sales-growth/

[4] Samaneh Sorournejad, Zahra Zojaji, Reza Ebrahimi Atani, Amir Hassan Monadjemi, "A Survey of Credit Card Fraud Detection Techniques: Data and Technique Oriented Perspective," The Arxiv, November, 2016.

[5] The Identity Theft Resource Center. 2019 End of Year Breach Report.

[6] W. Liu and S. Chawla. "A game theoretical model for adversarial learning." In Data Mining Workshops, 2009. ICDMW'09. IEEE International Conference on, pp. 25-30, IEEE, 2009.

[7] M. Zeager, A. Sridhar, N. Fogal, S. Adams, D. Brown, and P. Beling, "Adversarial Learning in Credit Card Fraud Detection", 2017 Systems and Information Engineering Design Symposium (SIEDS), April 2017.

[8] A. Mead, T. Lewris, S. Prasanth, S. Adams, P. Alonzi, and P. Beling, "Detecting Fraud in Adversarial Environments: A Reinforcement Learning Approach," 2018 Systems and Information Engineering Design Symposium (SIEDS).

[9] R. Lee, O. Mengshoel, A. Saksena, R. Gardner, D. Genin, J. Silbermann, M. Owen, M. Kochenderfer, "Adaptive Stress Testing: Finding Likely Failure Events with Reinforcement Learning," The Arxiv, Nov. 2018.

[10] R. Moss, R. Lee, N. Visser, J. Hochwarth, J. G. Lopez, M. Kochenderfer, "Adaptive Stress Testing of Trajectory Predictions in Flight Management Systems, " The Arxiv, Nov. 2020.

[11] M. Koren; S. Alsaif; R. Lee; M. Kochenderfer, "Adaptive Stress Testing for Autonomous Vehicles", 2018 IEEE Intelligent Vehicles Symposium (IV), June 2018.

[12] N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," The Arxiv, June 2011

[13] A. Dal Pozzolo. "Adaptive Machine Learning for Credit Card Fraud Detection", Thesis, Université Libre De Bruxelles, 2015, Web, ¡http://difusion.ulb.ac.be/vufind/Record/ULB-DIPOT:oai:dipot.ulb.ac.be:2013/221654/Holdings¿.

[14] SMOTE for Imbalanced Classification with Python. https://machinelearningmastery.com/smote-oversampling-for-imbalanced-classification/