Value Detection Rate: A Performance Metric for Payments Fraud Detection

Danial Dervovic[1], **Saeid Amiri**[2], **Michael Cashmore**[1]

[1]JP Morgan AI Research. Edinburgh, UK
[2]JP Morgan AI Research. New York City, NY, USA
danial.dervovic@jpmchase.com, saeid.amiri@jpmchase.com, michael.cashmore@jpmorgan.com

Abstract

Any plan for preventing fraud in financial transactions networks requires a well specified metric for success. In the industry, a well-known metric used to benchmark payments fraud detection algorithms is *Value Detection Rate (VDR)*. This metric explicitly considers monetary savings due to prevented fraud in contrast to metrics in the academic literature. In this short note we motivate and introduce this metric via a formal description and show it constitutes a consistent asymptotically unbiased statistical estimator to a relevant probabilistic query about test data drawn iid from the transaction population. We show via experiments on two publically available fraud datasets that fraud prevention policies guided by VDR have smaller financial losses than those guided by the popular Precision-at-k ($P@k$) metric. Moreover, we empirically show the VDR estimate converges given a practical number of samples, recapitulating the limiting case proven formally.

Introduction

Payments Fraud is a damaging problem across industry, eroding trust, impacting customer loyalty and employee morale, alongside financial losses incurred due to the fraud itself and attendant regulatory action. Indeed the problem is widespread, with 71% of surveyed organisations being targeted by fraudsters in 2021 according to the study by the Association for Financial Professionals (2022). In the 2022 survey by PwC (2022), responding organisations reported total losses of US$42B.

Fraudsters typically belong to one of several archetypes (Saporta and Maraney 2022; Association for Financial Professionals 2022) where often there is a sequential component to their actions. For example, a fraudster may initially attempt several smaller fraudulent transactions involving a compromised account before attempting a larger transaction. On the defender's side, constraints on how many resources can be used to fight fraud are balanced against how much fraud can be prevented (Hassanzadeh et al. 2021).

These characteristics of the payments fraud prevention problem suggest this problem is well-suited to AI Planning methods and sequential decision-making algorithms, yet the literature is scant. The works by Shen and Kurshan (2021); Vimal et al. (2021) formulate fraud prevention as an MDP and use Deep Reinforcement Learning to solve this problem. Rigter et al. (2022) formulate an online task allocation problem that applies to fraud prevention as a hybrid MDP and solve it using a dynamic programming approach with state abstraction. The paper by Dervovic et al. (2021) considers fraud prevention as a constrained sequential decision-making problem, solved via a dynamic threshold based policy. Planning-based techniques are used by Borrajo, Veloso, and Shah (2021) for the – related, but different – task of Anti-Money Laundering.

Indeed, the vast majority of literature on applying AI to fraud detection uses Machine Learning or Data Mining approaches. This literature is too extensive to present here, so we reference a number of surveys on these approaches (Bolton and Hand 2002; Phua et al. 2010; Lucas and Jurgovsky 2020; Ali et al. 2022; Narayan, Madhu Kumar, and Chacko 2023).

In this work we describe the fraud detection problem and present an important metric, VDR, for evaluating fraud detection policies. Moreover, we show that VDR has several desirable theoretical and experimental characteristics.

Formulating Fraud Detection

Fraud detection is usually presented as a supervised learning problem with a large class imbalance. Approaches to imbalanced learning are extensively surveyed (Abd Elrahman and Abraham 2013; Wagle and Manoj Kumar 2023). Concretely, individual transactions are modelled as samples drawn from a distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$, where $\mathcal{X}$ corresponds to transaction features available to the defender or decision-maker responsible for detecting fraud. The data labels $\mathcal{Y} = \{0, 1\}$ are binary, with 1 corresponding to a true fraud and 0 corresponding to a non-fraudulent transaction. Each sample $(X, Y) \sim \mathcal{D}$ corresponds to one transaction. Notably, the fraudulent ($Y = 1$) population is orders of magnitude smaller than the non-fraudulent population ($Y = 0$). The defender has a corpus of training data drawn from $\mathcal{D}$. Their goal is then to prevent as many fraudulent transactions as possible at test time – with transactions drawn again from $\mathcal{D}$. In practise, the training data will be taken before a certain time period and test from the time period immediately following (Saporta and Maraney 2022; Le Borgne et al. 2022).

This testing procedure is carried out to allow for violations of the iid assumption on the data distribution, that is, the test data being drawn from some alternative distribution $\mathcal{D}'$. In this case the decision-maker wishes to be confident that any policy trained on data from $\mathcal{D}$ still performs well on $\mathcal{D}'$. The performance in fraud prevention is measured along a number of different axes in the literature.

Existing Performance Metrics

The fraud detection problem is usually stated in terms of unbalanced supervised classification, so most metrics are based on the test-set confusion matrix based on a fixed threshold, or threshold free methods such as Receiver Operator Characteristic (ROC) and Precision-Recall (PR) curves, with the latter being more popular (Boyd, Eng, and Page 2013; Saito and Rehmsmeier 2015). The work by Davis and Goadrich (2006) discusses the relationship between PR and ROC curves with respect to imbalanced classification. Hassanzadeh et al. (2021) explicitly consider tradeoffs in these metrics when there are constraints on the number of positive label assignments the defender can give – as is often the case in practise. The framework of *Cost-sensitive learning* (Elkan 2001) seeks to minimise the misclassification cost between positive and negative examples, even on an example-by-example basis (Bahnsen, Aouada, and Ottersten 2014), but often the methods are tuned to a specific classification algorithm. The Precision-@-k metric (Fan and Zhu 2011), or $P@k$, selects the transactions with top-k highest predicted fraud probability by a model and evaluates the precision on this set of data. Currently there is no consensus on the best metric to use for evaluating fraud detection policies (Le Borgne et al. 2022), but instead one must decide on the tradeoffs one is willing to make.

That being said, what is missing from these evaluation metrics is considerations of monetary costs incurred by undetected fraudulent transactions. The Value Detection Rate (VDR) is a metric used in industry that explicitly considers this. Namely, maximising VDR corresponds to minimising monetary costs. In this brief note we formally introduce VDR. We proceed to show some theoretical and experimental properties of the VDR metric, namely its estimator is asymptotically unnbiased and consistent. We empirically validate these theoretical bounds on public payments fraud datasets. Furthermore, we show that an inspection policy that is based on VDR minimises financial costs due to missed fraudulent transactions in contrast to existing fraud detection metrics.

Value Detection Rate

Value Detection Rate (VDR) is a quantity taking values in $[0, 1]$ designed to capture the financial damage prevented by a plan or policy. Normalising to $[0, 1]$ allows a fair performance comparison of the same policy across different data slices.

We assume that there is a trained classifier $\hat{y} : \mathcal{X} \to [0, 1]$, where the output $\hat{y}(x)$ is interpreted as the classifier's subjective probability that the transaction $x \in \mathcal{X}$ is fraudulent. For brevity we define the function $r : \mathcal{X} \to \mathbb{R}_{++}$ that denotes the monetary value of a transaction.

The classifier subjective probability and transaction value allow us to define a *scoring function*, $\Psi : \mathcal{X} \to \mathbb{R}_{++}$,

$$\Psi(x) = \hat{y}(x) \cdot r(x), \tag{1}$$

that is, the scoring function is the monetary value of the transaction weighted by the classifier's subjective probability that the transaction is fraudulent. Suppose we have a validation set, $D_n = \{(x_j, y_j)\}_{j=1}^n$, drawn iid from $\mathcal{D}$. We fix a proportion of transactions that are to be inspected $\epsilon \in (0, 1)$ and define the *inspection function*, $\psi_\epsilon : \mathcal{X} \to \{0, 1\}$,

$$\psi_\epsilon(x) = \mathbb{1}\{\Psi(x) \geq \Psi_\epsilon\}, \tag{2}$$

where

$$\Psi_\epsilon := \min \left\{ \Psi(x_j) \,\middle|\, j \in \underset{S \subset \{1,\ldots,n\}}{\arg\max} \left\{ \sum_{i \in S} \Psi(x_i) \,\middle|\, |S| \leq \epsilon n \right\} \right\}. \tag{3}$$

The score threshold Ψ_ϵ represents the score needed to be included in the top ϵ-scoring fraction of the validation set D_n. The inspection function is the classifier that marks transactions for inspection ($\psi_\epsilon = 1$) that exceed the score threshold Ψ_ϵ. We note that while the mathematical formula is complicated, this is easy to implement – solving the inner optimisation problem by sorting $\Psi(x_i)$ in descending order and choosing the top $\lfloor \epsilon n \rfloor$ transactions is sufficient.

Given an inspection function ψ_ϵ we want to know how effective this inspection function is at surfacing fraud, prioritising the most valuable transactions. Let $(X, Y) \sim \mathcal{D}$ be the random variable corresponding to a transaction, along with its fraud label. The monetary value of a transaction, assuming it is fraudulent, is given by $\mathbb{E}[r(X) \mid Y = 1]$. The expected fraudulent value captured by our inspection function ψ_ϵ is given by $\mathbb{E}[\psi_\epsilon(X) \cdot r(X) \mid Y = 1]$. We construct a dimensionless quantity that summarises the effectiveness of an inspection function ψ_ϵ, the Value Detection Rate (VDR), VDR_ϵ.

Definition 1 (Value Detection Rate)**.** Let $\psi : \mathcal{X} \to \{0, 1\}$ be a fraud classifier. Then the Value Detection Rate (VDR) of ψ over $\mathcal{D}$ is defined as

$$\mathrm{VDR}(\psi) = \frac{\mathbb{E}_{(X,Y)\sim\mathcal{D}}[\psi(X) \cdot r(X) \mid Y = 1]}{\mathbb{E}_{(X,Y)\sim\mathcal{D}}[r(X) \mid Y = 1]},$$

where for brevity we denote the VDR of an inspection function ψ_ϵ as defined in Eq. (2) by $\mathrm{VDR}_\epsilon \equiv \mathrm{VDR}(\psi_\epsilon)$.

Intuitively, VDR encompasses the notion that we want any fraud prevention scheme to capture as much of the fraudulent monetary value as possible given a randomly sampled transaction. Note that a perfect classifier ψ will achieve $\mathrm{VDR}(\psi) = 1$. Moreover, the trivial classifier $\psi : x \mapsto 1$ also has a perfect VDR score – it is the restriction on inspecting only an ϵ-fraction of transactions while also measuring fraudulent value captured that renders VDR_ϵ an effective metric for measuring performance.

We note that there is an additional metric used in the industry, *Total Detection Rate (TDR)*, corresponding to the special case where we impose $r : x \mapsto 1$. TDR is very similar to Precision@k, where the sole difference is k is a constant number of transactions to inspect vs a fraction ϵ, indeed the metrics are identical when $k = \lceil \epsilon n \rceil$.

Estimating VDR

Let us now assume we have test data $D_m = \{(x_j, y_j)\}_{j=1}^m$ drawn iid from $\mathcal{D}$. How can we estimate VDR_ϵ using D_m. An estimator that immediately springs to mind is the following one.

$$\widehat{\mathrm{VDR}}_\epsilon := \frac{\sum_{i=1}^m \psi_\epsilon(x_i) r(x_i) y_i}{\sum_{i=1}^m r(x_i) y_i} \tag{4}$$

We shall see that this estimator is the appropriate one to use, as it is asymptotically unbiased and consistent. We first state a lemma that will greatly help in proving these propositions.

Lemma 2. *(Cochran 1977, Theorem 6.4) Suppose we have random variables U, V with finite variances S_u^2, S_v^2, correlation coefficient $\rho_{u,v}$, and finite means μ_U, μ_V such that $\mu_V \neq 0$. The limiting distribution of $\hat{\mu}_u / \hat{\mu}_v := \sum_{i=1}^n u_i / \sum_{i=1}^n v_i$, from random samples of size n from an infinite population, is normal $\mathcal{N}(\mu, \sigma^2)$ with*

$$\mu = \frac{\mu_U}{\mu_V}, \quad \sigma^2 = \frac{1}{n} \frac{\mu_U^2}{\mu_V^2} \left[\frac{S_u^2}{\mu_U^2} - \frac{2\rho_{u,v} S_u S_v}{\mu_U \mu_V} + \frac{S_v^2}{\mu_V^2} \right].$$

We can now show asymptotic unbiasedness and consistency. Note in the following we assume that the distribution $r(X)$ has a finite mean and variance.

Proposition 3 (VDR Asymptotic Unbiasedness)**.** *The estimator $\widehat{\mathrm{VDR}}_\epsilon$ is asymptotically unbiased, that is*

$$\mathbb{E}\left[\widehat{\mathrm{VDR}}_\epsilon - \mathrm{VDR}_\epsilon\right] \to 0 \quad as \quad m \to 0.$$

Proof. The estimator $\widehat{\mathrm{VDR}}_\epsilon$ is a ratio. Consider first the expectation of the numerator of (4),

$$\mathbb{E}[\textstyle\sum_{i=1}^m \psi_\epsilon(x_i) r(x_i) y_i] = \sum_{i=1}^m \mathbb{E}[\psi_\epsilon(x_i) r(x_i) y_i], \tag{5}$$

by linearity of expectations

$$= \mathbb{E}[\psi_\epsilon(x_i) r(x_i) \mathbb{1}\{y_i = 1\}] \tag{6}$$

collecting terms in the sum and rephrasing the y_i variables

$$= \mathbb{E}[\psi_\epsilon(X) r(X) \mid Y = 1] \mathbb{P}(Y = 1). \tag{7}$$

Similarly, for the denominator of (4) we have

$$\mathbb{E}[\textstyle\sum_{i=1}^m r(x_i) y_i] = m\mathbb{E}[r(X) \mid Y = 1] \mathbb{P}(Y = 1). \tag{8}$$

Thus the ratio of expectations of the numerator and denominator of (4) is $\mathbb{E}[\psi_\epsilon(X) r(X) \mid Y = 1] / \mathbb{E}[r(X) \mid Y = 1] = \mathrm{VDR}_\epsilon$. Indeed, from Lemma 2 we have for large m that the expectation of the ratio is the ratio of the expectations and so $\widehat{\mathrm{VDR}}_\epsilon$ constitutes an unbiased estimator of VDR_ϵ for a sufficiently large test set D_m. □

Proposition 4 (VDR Consistency)**.** *The estimator $\widehat{\mathrm{VDR}}_\epsilon$ is consistent, that is for all $\varepsilon > 0$*

$$\lim_{m \to \infty} \mathbb{P}\left[\left|\widehat{\mathrm{VDR}}_\epsilon - \mathrm{VDR}_\epsilon\right| > \varepsilon\right] = 0.$$

Proof. From the proof of Proposition 3 and Lemma 2 we have that $\widehat{\mathrm{VDR}}_\epsilon$ follows a normal distribution for large m, with mean VDR_ϵ. Notice that the variance is $O(1/m)$, from which consistency immediately follows. □

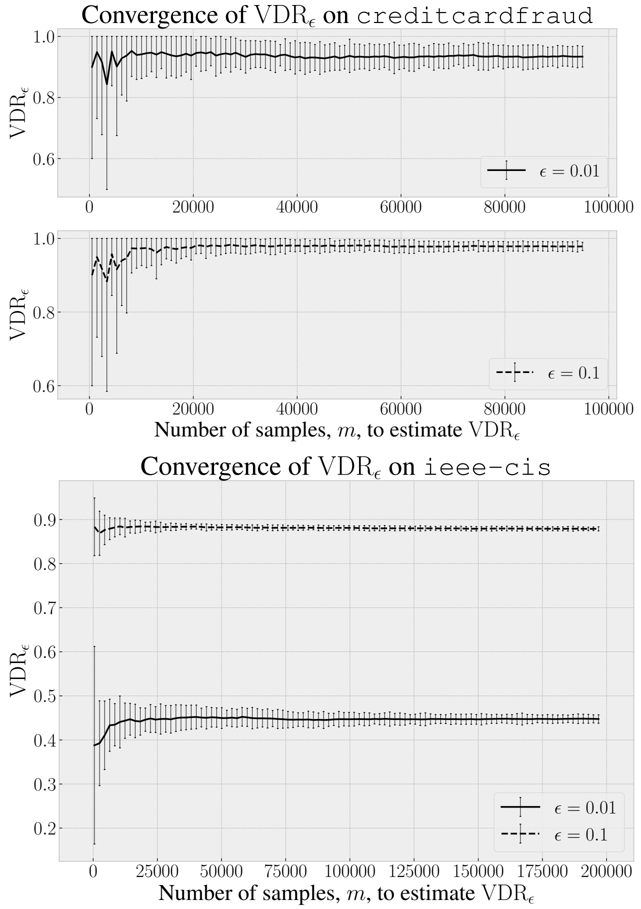

Figure 1: Plot of size of test set D_m against VDR_ϵ showing convergence for large number of samples on datasets `creditcardfraud` (top two panels) and `ieee-cis` (bottom panel). These give a practical estimate for the rates of convergence of Propositions 3 and 4. Reported error bars are standard deviation over 20 samples.

Experiments

In this section we conduct two experiments: *Experiment 1. VDR Convergence*; and *Experiment 2. VDR Effectiveness*. Our domain is fraud in financial transactions – there are few public datasets with realistic exemplars of this data (Le Borgne et al. 2022). We choose the `creditcardfraud` (Dal Pozzolo et al. 2015) and `ieee-cis` (IEEE-CIS 2019) datasets as they are based on real data, provide realistic features and explicit monetary values for transactions. Dataset preparation was identical to that in (Dervovic et al. 2021) and the model used was XGBoost (Chen and Guestrin 2016) with default hyperparameters.

In Experiment 1 (Figure 1) we vary the size of the test set D_m used to estimate VDR_ϵ for $\epsilon \in \{0.01, 0.1\}$, with m varying from 5000 to the size of the respective validation sets. There are 20 replicates sampled without replacement for each value of m, of which there were 100 linearly spaced values. From Figure 1 we see for both

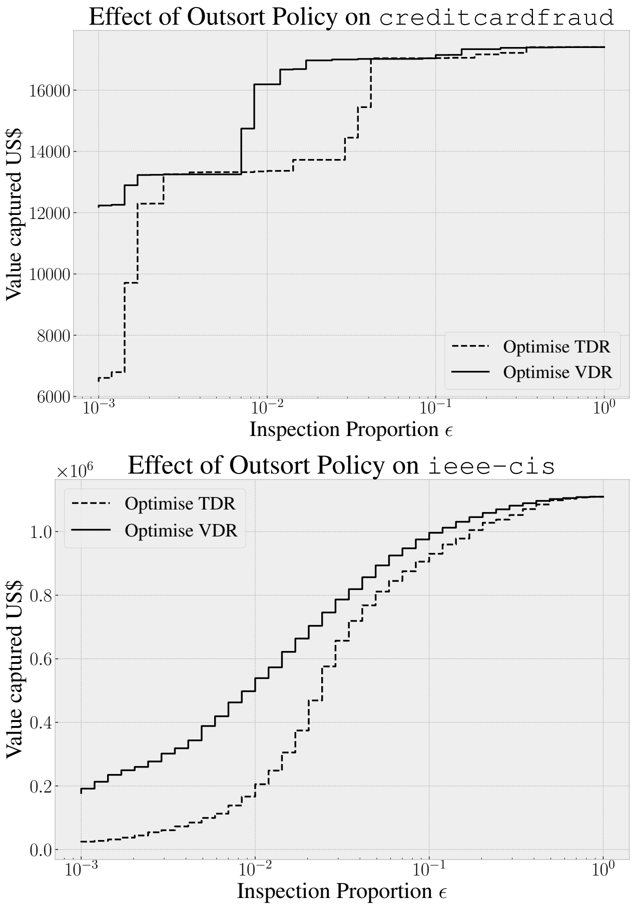

Figure 2: Curves of Fraudulent US$ captured against outsort rate $\epsilon \in (0,1)$ using two different policies: Total Detection Rate (equiv. to $P@k$) and VDR. Top panel is `creditcardfraud` dataset and bottom panel is `ieee-cis` . Notice that using an outsorting policy informed by VDR captures more fraudulent value over a wide range of outsort rates. The x-axis is log-scaled showing this effect persists over multiple orders of magnitude.

`creditcardfraud` and `ieee-cis` that the VDR_ϵ estimates are converging to approximately normal after a reasonable number of samples m, as evidenced by the stabilisation in the size of the error bars. This empirically validates the theoretical asymptotic bounds established in Propositions 3 and 4.

In Experiment 2 we consider fraud prevention via static threshold policies, namely for a given threshold τ, there are two policies:

$$\begin{aligned} \text{VDR policy} &:= \text{Inspect } x \text{ iff } \Psi(x) \geq \tau, \\ \text{TDR policy} &:= \text{Inspect } x \text{ iff } \hat{y}(x) \geq \tau. \end{aligned} \tag{9}$$

If a truly fraudulent transaction is inspected its monetary value is said to be captured and if it is not inspected the money is assumed to be lost by the decision maker. The VDR policy is implictly optimising VDR as we are thresholding based on the value-weighted model score, whereas the TDR policy is thresholding solely on model score. The TDR policy stands in for optimising existing fraud detection metrics such as $P@k$.

For the VDR policy we take $\tau_\epsilon^{(\text{VDR})} = \Psi_\epsilon$ as defined in Eq. (3) and for the TDR policy we assume $\tau_\epsilon^{(\text{TDR})}$ takes the form

$$\tau_\epsilon^{(\text{TDR})} := \min\left\{ \hat{y}(x_j) \,\middle|\, j \in \operatorname*{arg\,max}_{S \subset \{1,\ldots,n\}} \left\{ \sum_{i \in S} \hat{y}(x_i) \,\middle|\, |S| \leq \epsilon n \right\} \right\}. \tag{10}$$

Both thresholds $\tau_\epsilon^{(\text{VDR})}$ and $\tau_\epsilon^{(\text{TDR})}$ are computed with respect to the model training data.

In Figure 2 we plot the fraudulent value captured in the test set using the VDR and TDR informed policies, over a range of ϵ values in $(0,1)$ covering several orders of magnitude. For the `ieee-cis` dataset we see that the VDR policy strictly dominates the TDR policy for all ϵ and dominates for the `creditcardfraud` dataset. This confirms that using VDR to inform a fraud prevention policy gives superior monetary savings due to avoided fraud than existing metrics.

Conclusion

In this short paper we formally introduce a metric used for fraud detection in industry, VDR, and motivate its use in prevention plans and policies. VDR is shown to have several desirable statistical properties. We encourage the use of VDR in works by the academic community on planning for fraud detection and prevention.

Disclaimer. This paper was prepared for informational purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co. and its affiliates ("JP Morgan"), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

References

Abd Elrahman, S. M.; and Abraham, A. 2013. A review of class imbalance problem. *Journal of Network and Innovative Computing*, 1(2013): 332–340.

Ali, A.; Abd Razak, S.; Othman, S. H.; Eisa, T. A. E.; Al-Dhaqm, A.; Nasser, M.; Elhassan, T.; Elshafie, H.; and Saif, A. 2022. Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review. *Applied Sciences*, 12(19).

Association for Financial Professionals. 2022. 2022 AFP® Payments Fraud and Control Survey. Underwritten by JP Morgan. https://www.afponline.org/publications-data-tools/reports/survey-research-economic-data/Details/payments-fraud.

Bahnsen, A. C.; Aouada, D.; and Ottersten, B. 2014. Example-Dependent Cost-Sensitive Logistic Regression for Credit Scoring. In *Proc. ICMLA*, 263–269.

Bolton, R. J.; and Hand, D. J. 2002. Statistical Fraud Detection: A Review. *Statist. Sci.*, 17(3): 235–255.

Borrajo, D.; Veloso, M.; and Shah, S. 2021. Simulating and Classifying Behavior in Adversarial Environments Based on Action-State Traces: An Application to Money Laundering. In *Proceedings of the First ACM International Conference on AI in Finance*, ICAIF '20. New York, NY, USA: Association for Computing Machinery.

Boyd, K.; Eng, K. H.; and Page, C. D. 2013. Area under the Precision-Recall Curve: Point Estimates and Confidence Intervals. In *Proceedings of the 2013th European Conference on Machine Learning and Knowledge Discovery in Databases - Volume Part III*, ECMLPKDD'13, 451–466.

Chen, T.; and Guestrin, C. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, 785–794. New York, NY, USA: Association for Computing Machinery.

Cochran, W. G. 1977. *Sampling Techniques, 3rd Edition.* John Wiley.

Dal Pozzolo, A.; Caelen, O.; Johnson, R. A.; and Bontempi, G. 2015. Calibrating Probability with Undersampling for Unbalanced Classification. In *2015 IEEE Symposium Series on Computational Intelligence*, 159–166.

Davis, J.; and Goadrich, M. 2006. The Relationship between Precision-Recall and ROC Curves. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, 233–240. New York, NY, USA: Association for Computing Machinery.

Dervovic, D.; Hassanzadeh, P.; Assefa, S.; and Reddy, P. 2021. Non-parametric stochastic sequential assignment with random arrival times. *arXiv preprint arXiv:2106.04944*.

Elkan, C. 2001. The Foundations of Cost-Sensitive Learning. In *Proc. IJCAI*, 973–978.

Fan, G.; and Zhu, M. 2011. Detection of rare items with target. *Statistics and Its Interface*, 4(1): 11–17.

Hassanzadeh, P.; Dervovic, D.; Assefa, S.; Reddy, P.; and Veloso, M. 2021. Tradeoffs in streaming binary classification under limited inspection resources. In *Proceedings of the Second ACM International Conference on AI in Finance*, 1–9.

IEEE-CIS, I. C. I. S. 2019. IEEE-CIS Fraud Detection. https://www.kaggle.com/c/ieee-fraud-detection/datasets.

Le Borgne, Y.-A.; Siblini, W.; Lebichot, B.; and Bontempi, G. 2022. *Reproducible Machine Learning for Credit Card Fraud Detection - Practical Handbook.* Université Libre de Bruxelles.

Lucas, Y.; and Jurgovsky, J. 2020. Credit card fraud detection using machine learning: A survey. *arXiv preprint arXiv:2010.06479*.

Narayan, A.; Madhu Kumar, S. D.; and Chacko, A. M. 2023. A Review of Financial Fraud Detection in E-Commerce Using Machine Learning. In Bhateja, V.; Yang, X.-S.; Chun-Wei Lin, J.; and Das, R., eds., *Intelligent Data Engineering and Analytics*, 237–248. Singapore: Springer Nature Singapore.

Phua, C.; Lee, V.; Smith, K.; and Gayler, R. 2010. A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119*.

PwC. 2022. Global Economic Crime and Fraud Survey 2022. https://www.pwc.com/gx/en/services/forensics/economic-crime-survey.html.

Rigter, M.; Dervovic, D.; Hassanzadeh, P.; Long, J.; Zehtabi, P.; and Magazzeni, D. 2022. Optimal Admission Control for Multiclass Queues with Time-Varying Arrival Rates via State Abstraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 9918–9925.

Saito, T.; and Rehmsmeier, M. 2015. The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3): e0118432.

Saporta, G.; and Maraney, S. 2022. *Practical Fraud Prevention: Fraud and AML Analytics for Fintech and ECommerce, Using SQL and Python.* O'Reilly Media, Incorporated.

Shen, H.; and Kurshan, E. 2021. Deep Q-Network-Based Adaptive Alert Threshold Selection Policy for Payment Fraud Systems in Retail Banking. In *Proceedings of the First ACM International Conference on AI in Finance*, ICAIF '20. New York, NY, USA: Association for Computing Machinery.

Vimal, S.; Kayathwal, K.; Wadhwa, H.; and Dhama, G. 2021. Application of Deep Reinforcement Learning to Payment Fraud. In *Multi-Armed Bandits and Reinforcement Learning: Advancing Decision Making in E-Commerce and Beyond*, KDD '21.

Wagle, P. P.; and Manoj Kumar, M. V. 2023. A Comprehensive Review on the Issue of Class Imbalance in Predictive Modelling. In Shetty, N. R.; Patnaik, L. M.; and Prasad, N. H., eds., *Emerging Research in Computing, Information, Communication and Applications*, 557–576. Singapore: Springer Nature Singapore.