Published by Nigerian Society of Physical Sciences. Hosted by FLAYOO Publishing House LTD

Proceedings of the Nigerian Society of Physical Sciences

Journal Homepage: https://flayoophl.com/journals/index.php/pnspsc

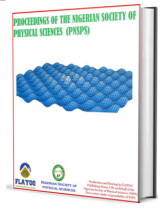

SmoteGAN and TabNet: a hybrid framework for detecting pump-and-dump schemes in cryptocurrency markets

Umar Faruq Abdulrazaq[a], Muhammad Nazeer Musa[b,*]

[a]*Department of Computing, University of Stirling, Stirling, Scotland*
[b]*Department of Cyber Security, Nigerian Defence Academy, Kaduna, Nigeria*

ABSTRACT

This paper describes a new method for detecting pump-and-dump (P&D) schemes in cryptocurrency markets, an important cybersecurity problem because of the large financial losses suffered by investors as a result of market manipulation and cyber-enabled attacks. Current P&D detection methods often fail to keep pace with changing manipulation strategies, particularly because transaction data are highly imbalanced. This study proposes a hybrid approach that combines generative adversarial networks (GANs) and TabNet to address these limitations. In the proposed framework, a GAN variant, SmoteGAN, is used to create synthetic P&D transaction samples and augment the original training data for the TabNet classification model. This mitigates class imbalance and allows TabNet to learn feature relationships sequentially through its attention mechanism. Evaluation on 25-second intervals of cryptocurrency P&D transaction data shows that the developed model achieved a precision of 98%, recall of 83% and F1-score of 90%, outperforming several existing state-of-the-art methods for detecting pump-and-dump schemes. The findings provide a hybrid method for improving cybersecurity in cryptocurrency markets by enhancing the detection of market manipulation and supporting a safer trading environment.

Keywords: Cryptocurrency fraud detection, Pump-and-dump schemes, SmoteGAN augmentation, TabNet.

DOI:10.61298/pnspsc.2026.3.257

1. INTRODUCTION

The rise of cryptocurrencies, which are viewed as a potential game changer for global finance, has also led to a significant increase in sophisticated financial crimes that exploit the decentralized and pseudonymous features of blockchain ecosystems [1]. Many of these criminal activities involve pump-and-dump (P&D) schemes, which have long been used in traditional markets. The characteristics of cryptocurrency markets, including volatility, regulatory arbitrage and information asymmetry, have enabled criminals to use P&D schemes to defraud investors [2, 3]. In these schemes, criminals create hype around a particular cryptocurrency and artificially inflate its price before selling their holdings, resulting in large losses for retail investors and reduced confidence in the cryptocurrency market [4, 5].

Despite their socioeconomic impact, P&D schemes are difficult for law enforcement agencies to detect because perpetrators rapidly change their strategies, labelled data are scarce and transaction patterns in decentralized marketplaces are difficult to model accurately [6]. Traditional fraud-detection methods,

*Corresponding Author Tel. No.: +234-816-6046-507.

such as rule-based systems and shallow statistical models, have had limited success because they cannot easily adapt to changing fraud strategies and often fail to model nonlinear relationships in high-dimensional cryptocurrency data [7, 8]. Machine learning (ML) techniques may detect subtle signals of P&D manipulation, but existing ML approaches remain limited by class imbalance, weak generalizability across cryptocurrencies and inadequate handling of tabular transaction data [7, 9].

To address these limitations, this study proposes a hybrid approach that combines synthetic data augmentation using SmoteGAN [10] with TabNet [11], a deep learning architecture designed for tabular data. The approach generates high-fidelity synthetic P&D samples to reduce data imbalance and uses TabNet's attention-based feature selection and interpretability to improve detection performance. To develop and validate the proposed model, the study uses historical cryptocurrency transactions and confirmed P&D events [12, 13]. The three main objectives are to: (1) demonstrate the effectiveness of SmoteGAN in generating realistic synthetic P&D data; (2) combine GAN-augmented data with TabNet to enhance detection performance; and (3) evaluate the model against benchmark methods using precision, recall and F1-score.

The improved model can provide regulators with a scalable tool for market surveillance and can support proactive risk mitigation by investors. This study also contributes to financial fraud detection research by demonstrating the value of GAN-based synthetic financial data generation and the usefulness of TabNet for modelling complex tabular relationships. The adaptability of the method suggests possible extension to other assets vulnerable to manipulation, such as decentralized finance (DeFi) tokens and non-fungible tokens (NFTs) [14].

2. LITERATURE REVIEW

P&D schemes have been present for a long time in traditional finance, and they have become an increasing concern in cryptocurrency markets. A P&D scheme occurs when an individual or group artificially increases the price of a cryptocurrency by disseminating false information and hype to stimulate buying interest, then sells its position at the inflated price, leaving other purchasers with the loss [4]. As shown in Figure 1, unsuspecting investors who buy at or near the peak may experience substantial losses when the scheme ends [2].

Generally, a P&D scheme occurs in three stages [2]. First, the organisers accumulate a large position in a particular cryptocurrency, typically an illiquid and low-cost asset, without attracting attention. Next, during the pump phase, they disseminate misinformation, hype and false promises through social media, forums and private messaging platforms to create demand and increase the asset price. Finally, during the dump phase, the organisers sell their holdings at the inflated price. The resulting selling pressure causes the price to collapse, and investors who purchased at inflated prices suffer substantial losses.

A typical cryptocurrency P&D scheme is illustrated in Figure 1. During the accumulation period, the perpetrator establishes a large position in a target cryptocurrency by acquiring a large amount at a low price while minimizing exposure to price movements. In the second stage, false and misleading information is disseminated through social media and online forums to influence new investors and encourage large purchases. This buying pressure quickly increases the price and builds momentum. In the final stage, the perpetrator sells the cryptocurrency at the inflated price to new or existing investors, and the resulting selling pressure causes the price to fall sharply.

The rapid rise in P&D schemes in cryptocurrency markets can be attributed to several market limitations. The marketplace is still developing, and many governments and regulatory bodies do not have clearly defined or consistent frameworks for regulating cryptocurrencies and related businesses [1]. P&D schemes exploit the volatility of cryptocurrency prices and the extent to which trades and price fluctuations may be driven by speculation or social media popularity [14]. As shown in Figure 2, social media platforms such as Telegram are often used to coordinate promotional activity around low-market-cap cryptocurrencies. The pseudonymity of cryptocurrency transactions and the decentralized nature of many exchanges also make it difficult for law enforcement agencies to identify and prosecute perpetrators [15]. P&D schemes create downside risk for investors, especially new investors, and may reduce trust in the cryptocurrency market [16].

2.1. TRADITIONAL DETECTION METHODS

Conventional methods for detecting P&D activity in financial markets are primarily rule-based, statistical or based on technical indicators. Rule-based systems generate alerts using predefined parameters, such as sudden spikes in trading volume [13]. Statistical analysis is used to determine whether anomalous trading behaviour has occurred [17, 18]. Technical indicators, such as the moving average convergence/divergence and relative strength index, can also assist in identifying suspicious activity, although they are often better suited to longer-term cryptocurrency trends than to short-lived P&D activity [19].

These traditional methods face major challenges in cryptocurrency markets. High volatility can produce many false-positive alerts [8]. Rule-based systems may be evaded by perpetrators who adapt their strategies to avoid detection [2]. Statistical methods may also fail to capture the complex interactions among market forces surrounding P&D events [20]. These limitations motivate the development of more sophisticated approaches, especially machine learning methods for cryptocurrency P&D detection.

2.2. MACHINE LEARNING TECHNIQUES

Because of the limitations of existing detection methods, there is increasing interest in using ML techniques to identify complex patterns and anomalies in cryptocurrency data [21]. Supervised learning uses labelled datasets for model training and classification [22]. Unsupervised learning, including clustering and anomaly detection, can identify anomalous patterns without labelled data [23], although it is prone to false positives. Semi-supervised learning combines labelled and unlabelled datasets for modelling.

Several studies have evaluated ML models such as random forests, k-nearest neighbours and decision trees for identifying P&D schemes in cryptocurrencies, with random forest methods showing promising results [24]. Other studies have examined logistic regression, random forest and AdaBoost at different

Figure 1. The three phases of pump-and-dump events [2].

Figure 2. Images of a Telegram P&D group: (a) getting ready for announcement and (b) countdown [23].

temporal scales to improve true-positive rates and reduce false-positive rates [7]. Deep learning techniques such as convolutional long short-term memory (C-LSTM) and Anomaly Transformer models have also outperformed traditional ML and statistical methods by capturing long- and short-term dependencies in P&D time-series data [25]. The integration of social media signals, including sentiment and tweet volume, has also improved detection accuracy in some models [26]. Nevertheless, class imbalance remains a major challenge because legitimate transactions substantially outnumber manipulative events [7]. Other key issues include feature selection, optimal analysis-window length and generalizability across cryptocurrencies.

2.3. SYNTHETIC DATA AUGMENTATION FOR IMBALANCED DATA

Class imbalance is a widespread problem in P&D detection. Models trained on imbalanced datasets may show high accuracy but low F1-scores because they fail to identify a sufficient number of minority-class P&D events. To address this problem, researchers have considered synthetic data generation methods to balance class distributions and improve model performance.

Common resampling methods include random undersampling (RUS), random oversampling (ROS), the synthetic minority oversampling technique (SMOTE) and random oversampling examples (ROSE). RUS reduces the number of majority-class samples, ROS replicates minority-class samples, SMOTE uses interpolation to generate synthetic minority-class samples and ROSE combines aspects of the other approaches. These methods are illustrated in Figures 3–5. Although they can improve model performance, they have limitations. RUS may discard useful information, ROS may increase noise and outlier influence, and SMOTE may perform poorly when minority-class samples are close to majority-class samples [27, 28]. Recent studies also show that class rebalancing can reduce performance or distort model interpretation in some settings [29, 30]. These limitations justify the exploration of advanced generative methods.

2.4. RELATED WORKS

Recent literature on P&D detection can be grouped into three main approaches: traditional ML techniques, imbalanced-data methodologies and deep learning methods. Ensemble methods, especially random forest (RF) and XGBoost, have been success-

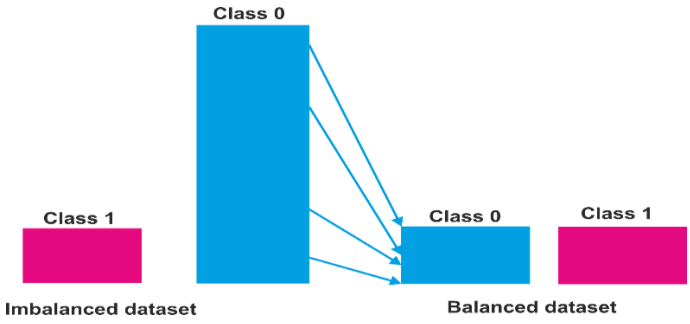

Figure 3. Random undersampling approach [28].

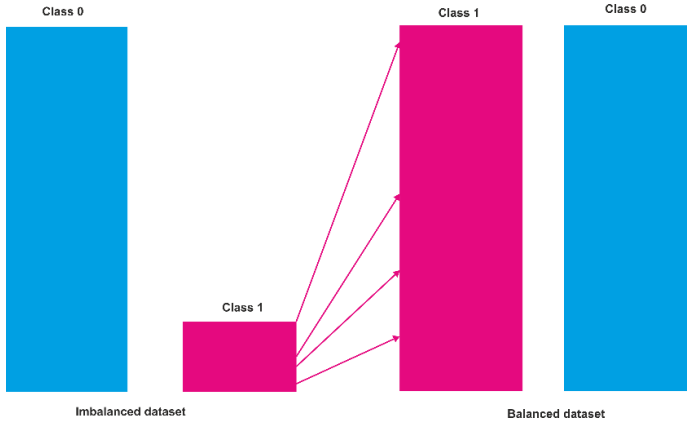

Figure 4. Random oversampling approach [28].

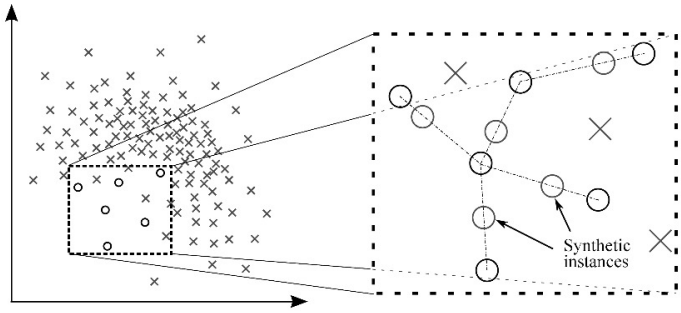

Figure 5. Generation of synthetic examples using SMOTE [27].

ful in supervised ML approaches. La Morgia *et al.* [4, 11] provided a baseline benchmark by monitoring Telegram groups for three years and detecting about 900 events. Their RF model identified schemes using 15-second chunks and achieved a precision of 91.3%, recall of 84.4% and F1-score of 87.7% based on five-fold cross-validation. Building on this work, Alfano *et al.* [26] showed that combining social media and financial features improved performance; their RF model on 25-second chunks achieved a precision of 97.6%, recall of 94.3% and F1-score of 95.9%.

Other studies have focused on different markets and time periods. Victor and Hagemann [6] used XGBoost on Binance and Telegram data and reported an area under the curve of 0.995, although their approach focused on known patterns and did not optimize chunk sizes for early P&D detection. Shao [24] used bagging and stacking with RF for Dogecoin P&D activity during the COVID-19 pandemic and reported 99% test accuracy and an F1-score of 84%. However, restricting the study to one asset limits its generalization and may increase overfitting risk.

Class imbalance remains a central challenge because normal trades greatly outnumber anomalous pumping events. Traditional balancing methods have produced inconsistent results. Fantazzini and Xiao [7] found that threshold-probability adjustment on the original imbalanced dataset produced superior performance in some cases, with 99% area under the curve and a 72% F1-score using 30-second chunks. Xu and Livshits [20] addressed imbalance by adjusting normal-data volumes in RF and generalized linear models, achieving 94% area under the curve. Victor and Hagemann [6] reduced imbalance effects through ten-fold cross-validation on a reduced dataset.

Standard oversampling methods such as SMOTE and ROSE often provide little advantage and may reduce performance, indicating the need for more powerful augmentation techniques. Deep learning methods remain underused, despite their suitability for complex sequential data. Chadalapaka *et al.* [25] applied Anomaly Transformer and C-LSTM models to the La Morgia dataset. The Anomaly Transformer achieved a precision of 88.4% and an F1-score of 89.2% on 25-second chunks, while the C-LSTM model achieved 94.2% precision and an F1-score of 89.3%. Market and social signals have also been used to support P&D detection [31].

Two gaps remain. First, traditional resampling techniques are common, but there is limited research on using advanced generative models, especially GANs, to create realistic minority-class instances for P&D detection. Second, deep learning models designed specifically for high-dimensional tabular cryptocurrency trading data, such as TabNet, have not been extensively explored. This study addresses these gaps by using GAN-based augmentation and TabNet to establish new benchmarks against ML [11] and deep learning [25] methods.

3. RESEARCH METHODOLOGY

The methodology of this research is shown in Figure 6. It describes the process by which data are collected, preprocessed, augmented with GANs, used to build the TabNet model and evaluated.

3.1. INPUT DATASET

The raw transaction dataset used in this research comes from the Binance cryptocurrency exchange and was derived from the analysis of La Morgia *et al.* [11] and Chadalapaka *et al.* [25]. It consists of documented P&D operations conducted on various cryptocurrency exchanges through public instant-messaging applications such as Telegram [25]. La Morgia *et al.* [11] collected the dataset from P&D-operated Telegram groups by recording the timestamps of official pump signals released by group administrators over a two-year period. These timestamps were then matched with Binance API data, and transaction data were collected for all pumped cryptocurrencies within one week before and after each pump event. This process produced a dataset containing 343 P&D occurrences [25].

After gathering the data, the authors created three aggregated datasets from the raw transaction data. Each aggregated dataset contains 15 features, as shown in Table 1. This study uses the 25-second chunk size of the transaction data.

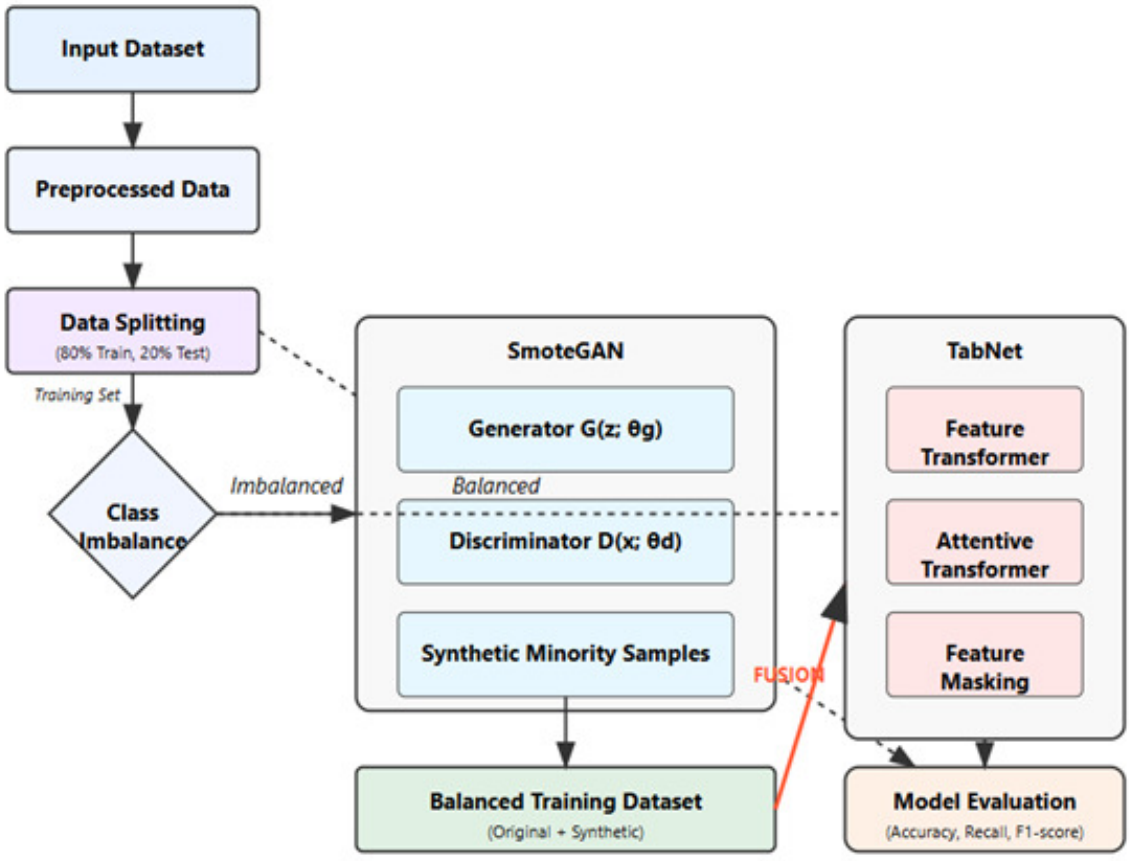

Figure 6. Research methodology workflow.

Table 1. Dataset features.

S/N	Feature	Type	Description
1	Date	String	Date of the chunk
2	HourSin, HourCos, MinuteSin, MinuteCos	Float	Positional encoding of the hour and minute
3	PumpIndex	Integer	Index of the pump
4	Symbol	String	Cryptocurrency symbol
5	StdRushOrder, AvgRushOrder	Float	Standard deviation and average percent change of rush orders
6	StdTrades	Float	Standard deviation of trades
7	StdVolume, AvgVolume	Float	Standard deviation and average percent change of volume
8	StdPrice, AvgPrice, AvgPriceMax	Float	Standard deviation, average percent change and average maximum percent change of price

3.2. DATA PREPROCESSING

After data collection, a preprocessing pipeline was created to make the data suitable for model training. Three transformations were applied. First, the Symbol feature was converted to a numerical representation using label encoding. Second, the Date feature was removed because it was redundant with other time-based features. Third, all numerical features were normalized to provide a consistent scale. This step is important because some algorithms, including neural networks and support vector machines, are sensitive to feature scaling and may assign excessive weight to variables with larger numerical ranges.

3.3. DATA SPLITTING

The preprocessed data were divided so that resampling and training were performed only on the training portion. The splitting strategy is shown in Figure 7. Before training, the model checks whether the data are balanced; when imbalance is detected, the training data are resampled.

3.4. DATA AUGMENTATION WITH GENERATIVE ADVERSARIAL NETWORKS

This study used the SmoteGAN architecture to address the inherent imbalance in the cryptocurrency P&D dataset. The purpose was to create realistic artificial P&D samples, increase training-set diversity and improve the model's ability to identify rare P&D events. GANs are generative models comprising a generator and a discriminator. Standard GANs produce artificial data, conditional GANs generate data based on predetermined labels and Wasserstein GANs improve training stability and output quality. SmoteGAN upsamples the minority class using SMOTE when creating artificial training samples [32, 33]. The GAN architecture is shown in Figure 8.

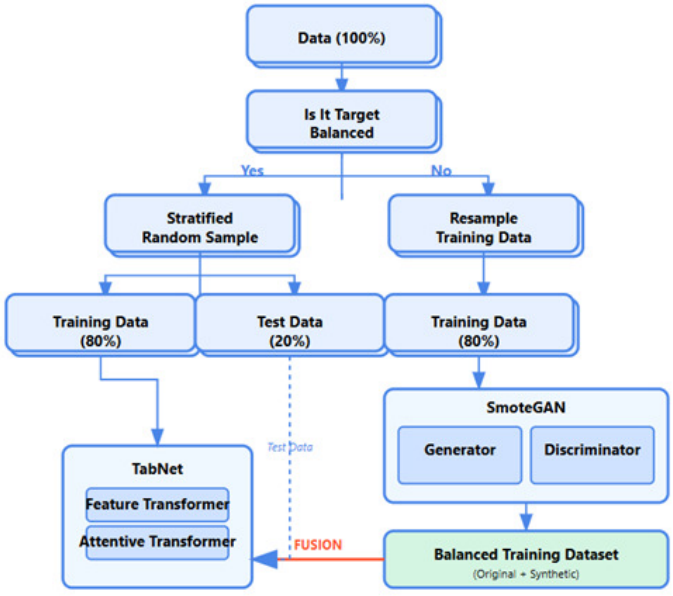

Figure 7. Train/test splitting strategy.

A discriminator D and a generator G form the two main components of the GAN architecture. The generator creates artificial P&D samples by mapping a latent-space vector z to the data space. It can be expressed as

$$G(z;\theta_G), \tag{1}$$

where θ_G represents the generator parameters. The discriminator is a binary classifier that differentiates between generated and real samples. It can be expressed as

$$D(x;\theta_D), \tag{2}$$

where θ_D indicates the discriminator parameters.

The generator and discriminator compete in an adversarial training procedure. During training, D is optimized to maximize the likelihood of correctly identifying generated and real data, while G is optimized to minimize

$$\log\left(1 - D(G(z))\right). \tag{3}$$

After training, G is used to generate new synthetic samples of the minority class by entering random noise vectors z into the generator. The basic dataset and the synthetic samples produced by SmoteGAN are then combined to form an augmented dataset with a more balanced class distribution. Similar to SMOTE, this method addresses class imbalance by producing realistic synthetic samples of P&D events.

3.5. BUILDING THE TABNET MODEL ON THE AUGMENTED DATASET

TabNet is a deep learning model developed for tabular data, such as cryptocurrency P&D events, and combines interpretability similar to decision-tree methods with the predictive power of neural networks. TabNet [10] was applied because it can model complex financial prediction problems through sequential attention processes that select relevant features along the algorithmic path. As shown in Figure 9, TabNet uses three main architectural components: a feature transformer for nonlinear transformation, an attentive transformer for determining feature focus and feature masking for sparse feature selection. These components improve interpretation and reduce overfitting. The model was trained on the augmented sample to address class imbalance in the original data.

Table 2. TabNet hyperparameters used.

Hyperparameter	Value
Number of decision steps	3
Feature dimension	8
Attention dimension	8
Number of independent GLU layers	2
Number of shared GLU layers	2
Learning rate	2×10^{-2}
Batch size	512
Virtual batch size	128
Optimizer	Adam
Learning-rate scheduler	StepLR
Step size	4
Gamma	0.9

3.6. MODEL CONFIGURATION

The TabNet hyperparameters used in this study are presented in Table 2. They include the number of decision steps, feature and attention dimensions, numbers of independent and shared gated linear unit (GLU) layers, learning rate and batch size.

The TabNet hyperparameters in Table 2 were determined through systematic tuning using a validation split of the training data. A grid search was performed over key parameters, including the number of decision steps (2–5), feature and attention dimensions (4–16) and learning rate (10^{-3}–10^{-1}). The final configuration—three decision steps, feature and attention dimensions of eight and a learning rate of 2×10^{-2}—was selected by maximizing the F1-score on the validation set while maintaining training stability and avoiding overfitting. The remaining hyperparameters, including the number of GLU layers, batch size and learning-rate scheduler, were adopted from the original TabNet implementation and refined through empirical validation to ensure convergence and generalization on the augmented dataset.

The TabNet training procedure can be summarized as follows. The model is initialized with the specified configuration. For each epoch, the model performs a forward pass to generate batch predictions, a backward pass to compute gradients and update parameters, learning-rate updating through the StepLR scheduler and validation based on balanced accuracy. Training was performed with a maximum of 50 epochs and an early-stopping criterion of 10 epochs based on balanced accuracy.

3.7. MODEL EVALUATION

After training, the TabNet model was assessed using several metrics, as shown in Table 3.

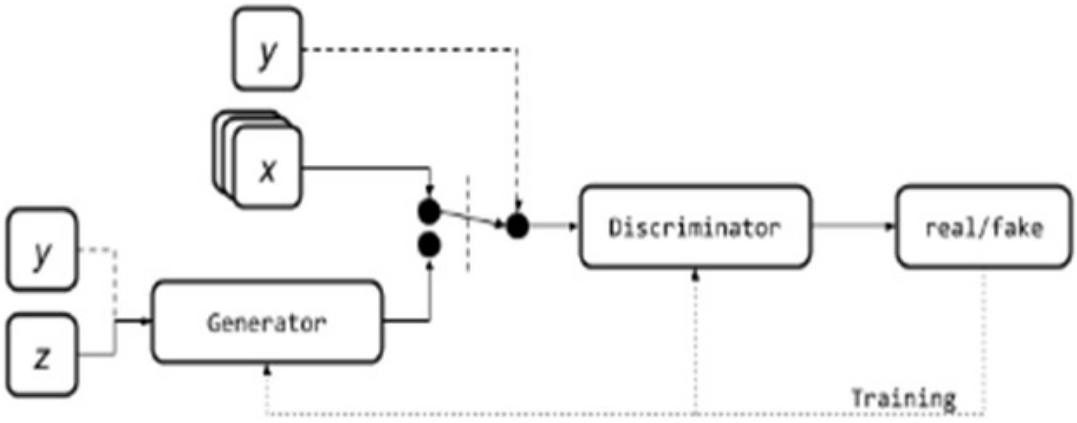

Figure 8. Generative adversarial network architecture.

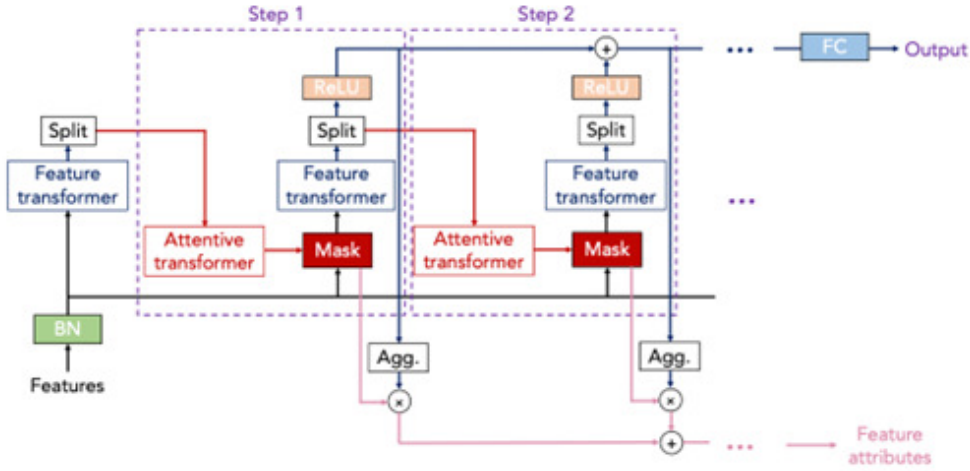

Figure 9. TabNet architecture [10].

Table 3. Evaluation metrics.

Metric	Formula	Description
Precision	TP/(TP + FP)	Percentage of positive forecasts that are actually positive
Recall (sensitivity)	TP/(TP + FN)	Percentage of true positives correctly predicted as positive
F1-score	2(precision × recall)/(precision + recall)	Harmonic mean of precision and recall
Overall accuracy	(TP + TN)/(TP + TN + FP + FN)	Percentage of accurate forecasts

In Table 3, TP denotes true positive, TN denotes true negative, FP denotes false positive and FN denotes false negative.

4. RESULTS AND DISCUSSION

This section presents and discusses the findings obtained using the SmoteGAN method and TabNet model to identify P&D events. The results are examined in relation to the research objectives and previous studies.

4.1. DATA AUGMENTATION RESULTS

The SmoteGAN method produced additional synthetic P&D samples for the minority class to balance the training data and produce a more even class distribution. As shown in Figure 10, the quality of the generated samples was assessed based on their distributional similarity to the initial minority-class samples. The similarity demonstrates that the SmoteGAN-generated samples support the original P&D sample distribution and indicates the effectiveness of the technique in producing realistic synthetic data.

4.2. TABNET MODEL PERFORMANCE

Compared with the baseline model trained on the imbalanced dataset, the TabNet model trained on the SmoteGAN-augmented dataset showed a notable improvement in P&D event detection. The results in Table 4 show that SmoteGAN and TabNet jointly addressed class imbalance and improved detection performance.

The integration of SmoteGAN with TabNet substantially improved model performance for P&D schemes across several metrics. Recall increased from 0.52 to 0.83, indicating that the

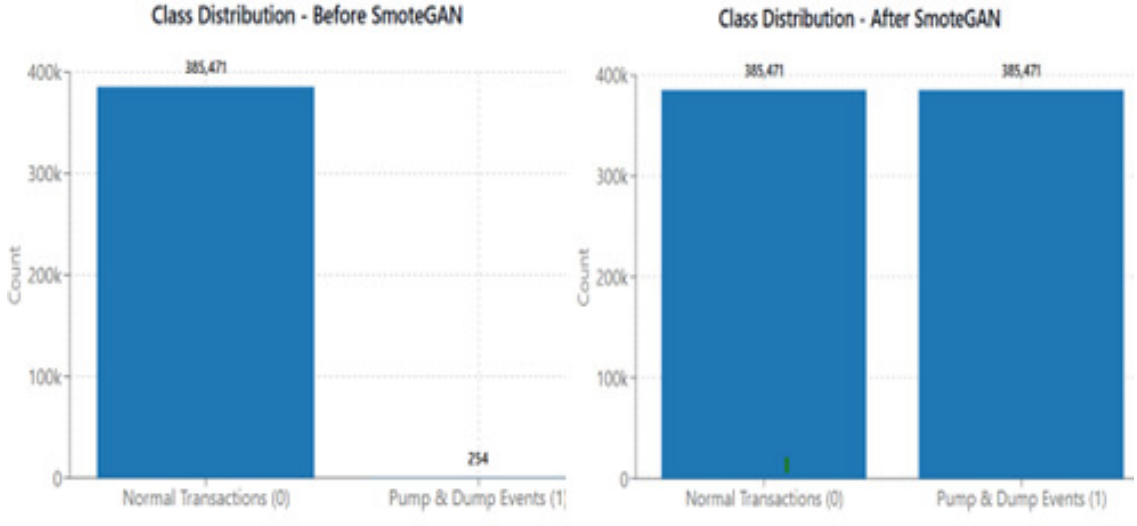

Figure 10. Distribution comparison of original and GAN-generated samples.

Table 4. Classification performance comparison.

Metric	Baseline model	SmoteGAN + TabNet
Accuracy	0.9985	0.9998
Precision (P&D class)	1.00	0.98
Recall (P&D class)	0.52	0.83
F1-score (P&D class)	0.69	0.90

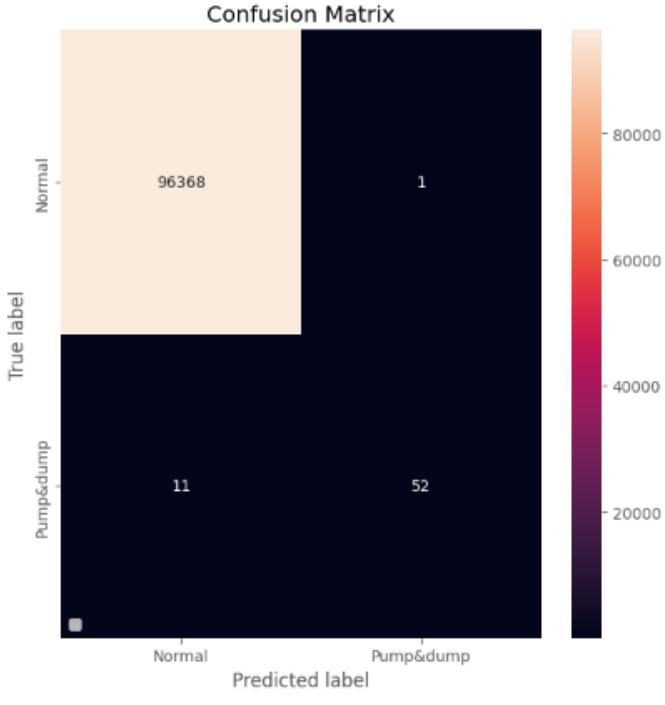

Figure 11. Confusion matrix for the proposed model.

model identified more actual P&D events. Although precision decreased slightly from 1.00 to 0.98, the F1-score increased from 0.69 to 0.90, showing a better balance between false positives and true positives.

The confusion matrix in Figure 11 shows the model's ability to identify P&D schemes. From 96,432 cryptocurrency exchange transactions, the model correctly classified 96,368 legitimate transactions as non-fraudulent true negatives and identified 52 anomalous P&D transactions as true positives. Misclassifications were low: one normal transaction was classified as an anomaly, and 11 P&D transactions were not detected. These results indicate that the proposed model can distinguish between legitimate transactions and P&D schemes in the cryptocurrency network with high overall accuracy.

4.3. MODEL VALIDATION AND INTERPRETATION

The SmoteGAN + TabNet model was tested on normal transactions and hidden P&D examples. Figure 12 shows one normal transaction example, including input features, the true label (Normal), the predicted label and the predicted class probabilities. The model correctly classified this example as Normal with a probability exceeding 99%. This high confidence indicates that the model can identify legitimate transactions while reducing false alarms in practical P&D detection settings.

One P&D instance tested by the model is shown in Figure 13. The true label (Fraud) aligned with the model prediction, indicating that the model can detect subtle but important anomalies in the minority class. The model assigned probabilities close to 100% to the P&D label, indicating successful separation between the two classes.

The successful identification of both classes suggests that P&D schemes produce measurable high-intensity signals. Features 2–7 were the most important for classification, which is consistent with established indicators of market manipulation. This suggests limited overlap in the decision boundary between classes. TabNet's built-in feature-selection capability also provided transparency by highlighting economically meaningful indicators, such as volatility spikes, as key classification drivers, which is valuable for regulatory applications.

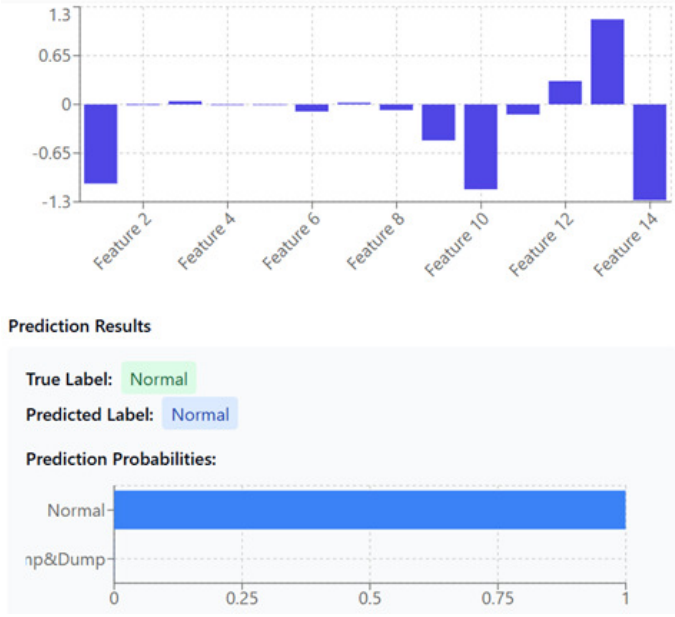

Figure 12. Simulation of a sample of normal transactions from unseen data.

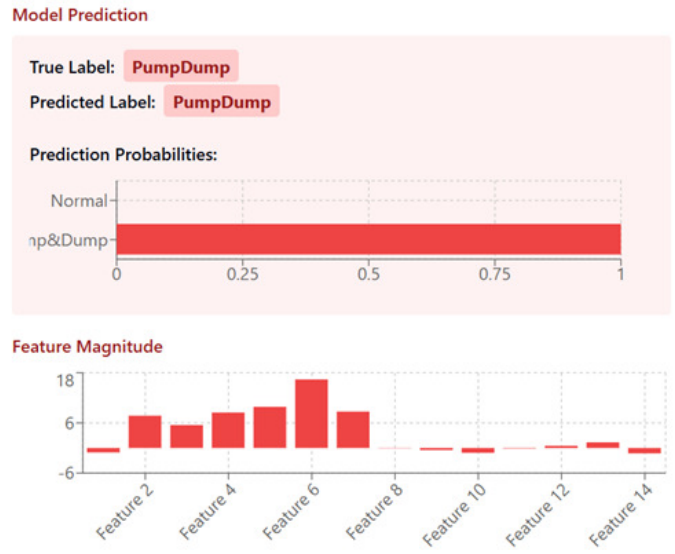

Figure 13. Simulation example of P&D transactions from unseen data.

Table 5. Comparison with previous studies using three performance metrics.

Study	Method	Precision	Recall	F1-score
Current study	SmoteGAN + TabNet	98%	83%	90%
Chadalapaka *et al.* [25]	C-LSTM	94%	85%	89.3%
Chadalapaka *et al.* [25]	Anomaly Transformer	88%	90%	89.2%
La Morgia *et al.* [11]	Random forest	95%	84%	88.8%

4.4. COMPARISON WITH PREVIOUS STUDIES

The performance of the proposed SmoteGAN + TabNet model was compared with previously published methods for identifying P&D events. The comparison used precision, recall and F1-score, as shown in Table 5.

Based on Table 5, the SmoteGAN + TabNet model achieved the strongest overall performance among the compared methods, with 98% precision and a 90% F1-score. The C-LSTM, Anomaly Transformer and random forest models did not exceed the proposed model in F1-score. Although the proposed model's recall of 83% was slightly lower than that of the Anomaly Transformer (90%), its precision–recall balance, as reflected in the F1-score, confirms its usefulness for addressing class imbalance and detecting P&D events. The high precision also indicates that the model substantially reduced false-positive detections, which is important for real-time applications where false alarms must be minimized.

5. CONCLUSION AND FUTURE WORK

This research addressed the challenge of detecting P&D schemes in cryptocurrency markets under extreme class imbalance and complex transaction patterns. The results support the use of SmoteGAN for synthetic data generation with TabNet's interpretable deep learning architecture. First, SmoteGAN produced 385,662 P&D instances from 63 original instances. By significantly increasing the size, diversity and representativeness of the training dataset, this augmentation improved recall from 0.52 to 0.83 and the F1-score from 0.69 to 0.90 for P&D detection. Second, the model achieved 98% precision and a 90% F1-score and outperformed existing techniques, including C-LSTM (94% precision and 89.3% F1-score) and random forest (95% precision and 88.8% F1-score). Third, the model exceeded 99% confidence for P&D predictions and demonstrated its ability to distinguish subtle and critical fraud signals, especially those associated with price volatility and trading anomalies.

The interpretability of TabNet is also useful because it highlights features that support the analytical comparison and evaluation of trading-volume spikes and price anomalies. Investors and regulators can therefore use the model to support the investigation of potentially unethical P&D schemes.

Several limitations remain. The current evaluation was based on previously collected data and does not demonstrate real-time performance. In addition, the validation used a limited set of cryptocurrencies and 25-second data intervals, so the results may not generalize to all digital assets or shorter temporal intervals such as 5- or 15-second windows. Future research should evaluate the model in real-time settings across multiple cryptocurrency marketplaces and integrate temporal indicators that capture the evolving nature of P&D schemes. Another area for development is to combine TabNet with graph-based or temporal models to improve detection effectiveness. The accuracy and utility of synthetic data should also be further evaluated for economic validity and potential bias using larger historical datasets.

DATA AVAILABILITY

The data will be available on request from the corresponding author.

ACKNOWLEDGMENT

The authors express their gratitude to the Department of Computing, University of Stirling, for providing the necessary support to conduct this research.

References

[1] S. Foley, J. R. Karlsen & T. J. Putniņš, "Sex, drugs, and bitcoin: how much illegal activity is financed through cryptocurrencies?", The Review of Financial Studies **32** (2019) 1798. https://doi.org/10.1093/rfs/hhz015.

[2] J. Kamps & B. Kleinberg, ‘‘To the moon: defining and detecting cryptocurrency pump-and-dumps’’, Crime Science **7** (2018) 18. https://doi.org/10.1186/s40163-018-0093-5.

[3] J. T. Hamrick, F. Rouhi, A. Mukherjee, M. Vasek, T. Moore & N. Gandal, ‘‘Analyzing target-based cryptocurrency pump and dump schemes’’, in Proc. 2021 ACM CCS Workshop on Decentralized Finance and Security, 2021, pp. 21. https://doi.org/10.1145/3464967.3488591.

[4] M. La Morgia, A. Mei, F. Sassi & J. Stefa, ‘‘The doge of Wall Street: analysis and detection of pump and dump cryptocurrency manipulations’’, ACM Transactions on Internet Technology **23** (2023) 1. https://doi.org/10.1145/3561300.

[5] A. Dhawan & T. J. Putniņš, ‘‘A new wolf in town? Pump-and-dump manipulation in cryptocurrency markets’’, Review of Finance **27** (2023) 935. https://doi.org/10.1093/rof/rfac051.

[6] F. Victor & T. Hagemann, ‘‘Cryptocurrency pump and dump schemes: quantification and detection’’, in Proc. 2019 International Conference on Data Mining Workshops (ICDMW), 2019, pp. 244. https://doi.org/10.1109/ICDMW.2019.00045.

[7] D. Fantazzini & Y. Xiao, ‘‘Detecting pump-and-dumps with crypto-assets: dealing with imbalanced datasets and insiders’ anticipated purchases’’, Econometrics **11** (2023) 22. https://doi.org/10.3390/econometrics11030022.

[8] M. Bartoletti, S. Carta, T. Cimino & L. Saia, ‘‘Dissecting Ponzi schemes on Ethereum: identification, analysis, and impact’’, Future Generation Computer Systems **102** (2020) 259. https://doi.org/10.1016/j.future.2019.08.014.

[9] A. Dal Pozzolo, O. Caelen, R. A. Johnson & G. Bontempi, ‘‘Calibrating probability with undersampling for unbalanced classification’’, in 2015 IEEE Symposium Series on Computational Intelligence, 2015, pp. 159. https://doi.org/10.1109/SSCI.2015.33.

[10] S. O. Arik & T. Pfister, ‘‘TabNet: attentive interpretable tabular learning’’, Proceedings of the AAAI Conference on Artificial Intelligence **35** (2021) 6679. https://doi.org/10.1609/aaai.v35i8.16826.

[11] M. La Morgia, A. Mei, F. Sassi & J. Stefa, ‘‘Pump and dumps in the bitcoin era: real-time detection of cryptocurrency market manipulations’’, in Proc. 2020 29th International Conference on Computer Communications and Networks (ICCCN), 2020, pp. 1. https://doi.org/10.1109/ICCCN49398.2020.9209660.

[12] SystemsLab-Sapienza, ‘‘Pump and dump dataset’’, GitHub, 2021. Available online: https://github.com/SystemsLab-Sapienza/pump-and-dump-dataset.

[13] R. K. Aggarwal & G. Wu, ‘‘Stock market manipulations’’, The Journal of Business **79** (2006) 1915. https://doi.org/10.1086/503652.

[14] M. T. Tran, N. Sohrabi, Z. Tari & Q. Wang, ‘‘How to cook the fragmented rug pull?’’, arXiv preprint arXiv:2511.15463, 2025. https://doi.org/10.48550/arXiv.2511.15463.

[15] N. Gandal, J. T. Hamrick, T. Moore & T. Oberman, ‘‘Price manipulation in the Bitcoin ecosystem’’, Journal of Monetary Economics **95** (2018) 86. https://doi.org/10.1016/j.jmoneco.2017.12.004.

[16] T. Li, D. Shin & B. Wang, ‘‘Cryptocurrency pump-and-dump schemes’’, Journal of Financial and Quantitative Analysis **60** (2025) 3622. https://doi.org/10.1017/S0022109025000201.

[17] R. Mitchell & R. Chen, ‘‘Behavior-rule based intrusion detection systems for safety critical smart grid applications’’, IEEE Transactions on Smart Grid **4** (2013) 1254. https://doi.org/10.1109/TSG.2013.2258948.

[18] P. Priyadarshi & P. Kumar, ‘‘A comprehensive review on insider trading detection using artificial intelligence’’, Journal of Computational Social Science (2024) 1. https://doi.org/10.1007/s42001-024-00284-5.

[19] Z. Chen, S. Wang, D. Yan & Y. Li, ‘‘Research and implementation of bank credit card fraud detection system based on reinforcement learning and LSTM’’, in Proc. 2023 3rd International Conference on Mobile Networks and Wireless Communications (ICMNWC), 2023, pp. 1. https://doi.org/10.1109/ICMNWC60182.2023.10435890.

[20] J. Xu & B. Livshits, ‘‘The anatomy of a cryptocurrency pump-and-dump scheme’’, in Proc. USENIX Security Symposium, Santa Clara, CA, USA, Aug. 2019, pp. 1609. Available online: https://dl.acm.org/doi/10.5555/3361338.3361450.

[21] T. M. Mitchell, *Machine Learning*, 1st ed. McGraw-Hill, New York, NY, USA, 1997. Available online: https://dl.acm.org/doi/10.5555/541177.

[22] T. Hastie, R. Tibshirani, J. H. Friedman & J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, New York, NY, USA, 2001. https://doi.org/10.1007/978-0-387-21606-5.

[23] M. J. Rajaei & Q. H. Mahmoud, ‘‘A survey on pump and dump detection in the cryptocurrency market using machine learning’’, Future Internet **15** (2023) 267. https://doi.org/10.3390/fi15080267.

[24] S. Shao, ‘‘The effectiveness of supervised learning models in detection of pump and dump activity in Dogecoin’’, in Proc. Second IYSF Academic Symposium on Artificial Intelligence and Computer Engineering **12079** (2021) 356. https://doi.org/10.1117/12.2622877.

[25] V. Chadalapaka, K. Chang, G. Mahajan & A. Vasil, ‘‘Crypto pump and dump detection via deep learning techniques’’, arXiv preprint arXiv:2205.04646, 2022. https://doi.org/10.48550/arXiv.2205.04646.

[26] D. Alfano, R. Abbruzzese & D. Parente, ‘‘Pump and dump cryptocurrency detection using social media’’, in Proc. International Conference on Data (DATA), 2023, pp. 235. https://doi.org/10.5220/0012059300003541.

[27] N. V. Chawla, K. W. Bowyer, L. O. Hall & W. P. Kegelmeyer, ‘‘SMOTE: synthetic minority over-sampling technique’’, Journal of Artificial Intelligence Research **16** (2002) 321. https://doi.org/10.1613/jair.953.

[28] G. Menardi & N. Torelli, ‘‘Training and assessing classification rules with imbalanced data’’, Data Mining and Knowledge Discovery **28** (2014) 92. https://doi.org/10.1007/s10618-012-0295-5.

[29] C. Tantithamthavorn, A. E. Hassan & K. Matsumoto, ‘‘The impact of class rebalancing techniques on the performance and interpretation of defect prediction models’’, IEEE Transactions on Software Engineering **46** (2020) 1200. https://doi.org/10.1109/TSE.2018.2876537.

[30] T. Wongvorachan, S. He & O. Bulut, ‘‘A comparison of undersampling, oversampling, and SMOTE methods for dealing with imbalanced classification in educational data mining’’, Information **14** (2023) 54. https://doi.org/10.3390/info14010054.

[31] H. Nghiem, G. Muric, F. Morstatter & E. Ferrara, ‘‘Detecting cryptocurrency pump-and-dump frauds using market and social signals’’, Expert Systems with Applications **182** (2021) 115284. https://doi.org/10.1016/j.eswa.2021.115284.

[32] R. Sauber-Cole & T. M. Khoshgoftaar, ‘‘The use of generative adversarial networks to alleviate class imbalance in tabular data: a survey’’, Journal of Big Data **9** (2022) 98. https://doi.org/10.1186/s40537-022-00648-6.

[33] A. Sharma, P. K. Singh & R. Chandra, ‘‘SMOTified-GAN for class imbalanced pattern classification problems’’, IEEE Access **10** (2022) 30655. https://doi.org/10.1109/ACCESS.2022.3158977.