

Statistical Profiling of Banking Data for Anomaly Identification: An Applied Study Using a Synthetic Banking Database

Djina Ivanovic¹, Evans khoza²

¹Digital Economics Department, Institute of Economic Sciences, Belgrade, Serbia;

²Business management, Cape Peninsula University of Technology, Cape Town, South Africa

Corresponding author email: djina.ivanovic@ien.bg.ac.rs

ABSTRACT: This study examines whether statistical analysis can meaningfully distinguish anomalous from non-anomalous records in an openly available banking database. The dataset contains 5,000 observations and 40 variables covering customer, account, transaction, loan, credit-card, and feedback information. For analysis, the original anomaly label was recoded so that 300 records (6.0%) were treated as anomalous and 4,700 records (94.0%) as normal. The empirical design combines descriptive statistics, Mann-Whitney U tests, chi-square association tests, and a logistic-regression benchmark on a hold-out sample. The results show that major numerical variables do not differ significantly between the two classes, while only account type and resolution status exhibit statistically significant but substantively small associations. The benchmark classifier performs weakly (ROC-AUC = 0.490; average precision = 0.062), indicating that the available variables provide limited discriminatory signal. The study therefore contributes less as an operational fraud-detection model and more as a transparent methodological illustration of how class imbalance, weak labels, and synthetic feature spaces can constrain inference. To strengthen the manuscript for journal submission, the present version expands the literature review with pre-2021 sources, introduces a comparison matrix against prior studies, and adds statistical tables and figures derived from the uploaded dataset.

Keywords: Anomaly Detection, Statistical Analysis, Fraud Analytics, Class Imbalance, Synthetic Data.

I. INTRODUCTION

Banking institutions generate large volumes of operational data from customer accounts, card activity, loan portfolios, and service interactions. These records are screened for unusual patterns because anomalous observations may signal operational errors, suspicious activity, misclassification, weak internal controls, or data-quality problems.

Although the language of anomaly detection often overlaps with fraud detection, the two are not identical. An anomaly may be fraudulent, but it may also reflect a process deviation, a recording error, or an unusual yet legitimate behavior. For this reason, statistical profiling remains a useful first step before more sophisticated machine-learning models are introduced.

The present study uses an uploaded Kaggle-style banking dataset and asks a narrow but important question: do the available variables show meaningful statistical separation between anomalous and non-anomalous records? The contribution is deliberately transparent. Instead of overstating predictive power, the study documents what can and cannot be learned from the supplied variables.

Three research questions guide the study:

- What are the descriptive characteristics of anomalous and non-anomalous observations?
- Which variables show statistically significant association with anomaly status?
- How well can a simple benchmark classifier distinguish anomalies out of sample?

II. PROBLEM STATEMENT

Despite the rapid growth of anomaly-detection research in banking and financial fraud analytics, publicly accessible banking datasets are rarely examined as methodological objects in their own right. Many published studies report promising predictive results using institutionally validated transactional streams, sequence-level behavior, or richly engineered fraud indicators, whereas far less is known about whether open synthetic banking datasets contain enough statistical signal to support valid anomaly identification. This creates an important problem for researchers who rely on such datasets for teaching, prototyping, and early-stage publication work: without a careful statistical assessment, anomaly labels may be treated as meaningful even when the underlying variables do not provide reliable discriminatory information.

Accordingly, the problem addressed in this article is whether the variables contained in the uploaded synthetic banking database are statistically capable of distinguishing anomalous from non-anomalous records, and whether that capability is strong enough to justify even a modest out-of-sample classification claim. Framed differently, the study asks whether an openly available banking table can support anomaly inference with sufficient empirical credibility, or whether its synthetic construction and weak feature structure materially constrain interpretation.

1. RESEARCH HYPOTHESES

- H1: The distributions of the main numerical banking variables differ significantly between anomalous and non-anomalous records.
- H1a: Transaction amount differs significantly by anomaly status.
- H1b: Account balance differs significantly by anomaly status.
- H1c: Loan amount differs significantly by anomaly status.
- H1d: Credit-card balance differs significantly by anomaly status.
- H1e: Interest rate differs significantly by anomaly status.
- H2: Selected categorical banking variables are significantly associated with anomaly status.
- H2a: Account type is significantly associated with anomaly status.
- H2b: Transaction type is significantly associated with anomaly status.
- H2c: Loan status is significantly associated with anomaly status.
- H2d: Resolution status is significantly associated with anomaly status.
- H3: A benchmark logistic-regression model using the available variables achieves out-of-sample discrimination better than random classification.

Given the exploratory nature of the study and the synthetic character of the dataset, these hypotheses are formulated as testable empirical expectations rather than as strong causal claims. Their purpose is to determine whether the uploaded data contain statistically meaningful signal that can support anomaly identification under transparent and reproducible conditions.

III. LITERATURE REVIEW

The literature shows that anomaly and fraud analytics evolved along four interrelated streams:

- foundational statistical and outlier-detection theory,
- fraud-detection systems and financial applications,
- class-imbalance and evaluation research,
- increasingly complex machine-learning and deep-learning approaches.

The present study is positioned at the intersection of the first three streams because it emphasizes transparent statistical analysis on a weak-signal banking dataset.

1. FOUNDATIONAL ANOMALY AND OUTLIER RESEARCH

Classical work on outliers emphasized that unusual observations must be interpreted relative to both distributional assumptions and substantive context. Hawkins (1980), Barnett and Lewis (1994), and Rousseeuw and Croux (1993) established the statistical language for robust detection, while Hodge and Austin (2004) and Chandola, Banerjee, and Kumar (2009) synthesized the field into statistical, distance-based, density-based, and model-based approaches.

In practical banking applications, these foundations remain relevant because analysts often begin with robust summaries, dispersion measures, thresholding, and groupwise comparisons before moving to more opaque algorithms. Methods such as Local Outlier Factor (Breunig et al., 2000) and Isolation Forest (Liu, Ting, & Zhou, 2008) became influential precisely because they operationalize local rarity and easy-to-isolate observations in high-volume datasets.

2. FRAUD DETECTION AND FINANCIAL APPLICATIONS

Financial fraud detection matured from general statistical reviews into domain-specific comparative studies. Bolton and Hand (2002) framed fraud detection as a dynamic statistical problem in which adversaries adapt to controls. Kou et al. (2004) surveyed fraud-detection techniques across domains, while Thomas (2000) mapped the broader risk-scoring tradition that informed subsequent banking analytics.

Later reviews demonstrated that financial applications rely on a mix of supervised and unsupervised approaches. Ngai et al. (2011), Sharma and Panigrahi (2013), Abdallah, Maarof, and Zainal (2016), West and Bhattacharya (2016), and Al-Hashedi and Magalingam (2021) all stressed recurring issues such as scarce labels, evolving fraud patterns, feature engineering, and the trade-off between accuracy and interpretability.

Applied studies similarly reported that performance depends heavily on data richness. Bhattacharyya et al. (2011), Bahnsen et al. (2016), Dal Pozzolo et al. (2015, 2018), Whitrow et al. (2009), Jurgovsky et al. (2018), Randhawa et al. (2018), and Carcillo et al. (2019, 2021) showed that transaction histories, sequence information, calibrated probabilities, and hybrid pipelines can improve fraud detection - especially in real credit-card environments. The implication for the current study is straightforward: a flat synthetic table with limited behavioral depth is unlikely to match the predictive power reported in richer operational studies.

3. ACCOUNTING AND FINANCIAL-REPORTING RELATED ANALYTICS

A related literature addressed anomalies in accounting and financial-reporting data rather than retail transaction streams. Fanning and Cogger (1998), Beneish (1999), Kirkos, Spathis, and Manolopoulos (2007), Perols (2011), and Ravisankar et al. (2011) showed that statistical and machine-learning techniques can detect manipulation patterns in published financial information.

These studies are relevant because they reinforce a central methodological lesson: good anomaly research depends less on the novelty of the algorithm than on the quality and structure of the features. In accounting contexts, carefully engineered ratios and red-flag indicators often matter more than model complexity. The same logic applies here: if synthetic banking variables do not encode behavioral regularities, predictive performance will remain weak.

4. CLASS IMBALANCE, VALIDATION, AND MODERN ANOMALY MODELING

One of the strongest themes in the literature is that financial anomalies are rare events. He and Garcia (2009) summarized the difficulty of learning from imbalanced data, while Dal Pozzolo et al. (2015) and Johnson and Khoshgoftaar (2019) emphasized calibration, threshold choice, and evaluation metrics that go beyond raw accuracy. In rare-event settings, precision-recall behavior is often more informative than accuracy alone.

Recent surveys also expanded anomaly detection into broader machine-learning and deep-learning ecosystems. Goldstein and Uchida (2016), Zhao, Nasrullah, and Li (2019), Pang et al. (2021), and Ruff et al. (2021) reviewed unsupervised, semi-supervised, and deep architectures for outlier detection. Yet these surveys also imply a cautionary point: advanced models cannot compensate indefinitely for weak features, noisy labels, or synthetic data-generating processes.

5. RESEARCH GAP AND POSITIONING OF THE PRESENT STUDY

The literature therefore identifies a clear gap addressed by the present paper. Prior studies typically optimize predictive performance on transactional fraud streams, review broad families of algorithms, or focus on accounting-manipulation signals. Much less attention is given to open synthetic banking datasets as methodological objects in their own right. The current study fills that gap by providing a transparent statistical assessment of such a dataset, demonstrating how significance testing, imbalance reporting, and

out-of-sample evaluation can reveal whether a publicly accessible anomaly label actually carries meaningful signal.

Table 1. Comparison of the present study with selected prior studies.

Study	Data / Context	Method(s)	Main finding	Difference from the present study
Bolton & Hand (2002)	General fraud detection	Statistical review	Established fraud detection as a statistical problem and emphasized adaptive behavior of fraudsters.	Conceptual review; no banking dataset-specific empirical test.
Ngai et al. (2011)	Financial literature fraud	Data-mining review	Classified financial fraud studies and mapped major algorithms and application areas.	Review-oriented; did not test anomaly labels in synthetic retail banking data.
Bhattacharyya et al. (2011)	Credit transactions card	Random forest, SVM, logistic, decision tree	Showed that ensemble and cost-aware approaches improve fraud detection under class imbalance.	Focused on card fraud rather than mixed banking tables with customer, loan, and service variables.
Dal Pozzolo et al. (2015)	Credit card fraud	Undersampling and probability calibration	Demonstrated the importance of calibrated probabilities in highly imbalanced fraud settings.	Used transactional fraud data rather than broader banking records.
Abdallah et al. (2016)	Cross-domain fraud detection	Survey	Synthesized fraud detection architecture, challenges, and evaluation issues.	Did not examine weak-signal anomaly labels in open banking data.
West & Bhattacharya (2016)	Financial fraud detection	Comprehensive review	Highlighted interpretability, feature engineering, and imbalance as central issues.	No applied test on synthetic banking database.
Jurgovsky et al. (2018)	Credit card fraud	Sequence classification / recurrent modeling	Showed that sequential patterns strengthen fraud detection beyond static features.	Our dataset lacks temporal sequences long enough for sequence learning.
Randhawa et al. (2018)	Credit card fraud	AdaBoost and majority voting	Reported gains from ensemble voting in fraud classification.	Emphasis on predictive optimization, not transparent statistical profiling.

Carcillo et al. (2021)	Credit card fraud	Hybrid unsupervised + supervised learning	Found that combined modeling can improve performance in real fraud streams.	Requires richer behavioral features and operational labels than available here.
Al-Hashedi & Magalingam (2021)	Financial fraud literature	Comprehensive review	Summarized the rapid expansion of data-mining methods in financial fraud detection.	Did not assess synthetic-data limitations or significance-vs-prediction tension.

IV. DATA AND STUDY DESIGN

The empirical analysis uses the uploaded file “Comprehensive_Banking_Database[1].csv”, which contains 5,000 rows and 40 variables. The data include demographics, account balances, transaction amounts and types, loan fields, credit-card fields, and customer-service feedback records.

A key limitation is that the dataset appears synthetic rather than institutionally sourced. Dates are concentrated in 2022, and many variables resemble randomly generated or weakly linked values. This does not invalidate the exercise, but it does restrict external validity and likely weakens measurable signal.

For the analysis, the original anomaly field was recoded so that values equal to -1 were treated as anomalous observations. Under this coding, the sample contains 300 anomalous records (6.0%) and 4,700 normal records (94.0%), indicating severe class imbalance.

V. METHODOLOGY

The empirical design followed four steps. First, descriptive statistics were computed for the full sample and by anomaly status. Second, Mann-Whitney U tests were used for numeric variables because distributional normality was not assumed. Third, chi-square tests were applied to categorical variables. Fourth, a logistic-regression benchmark with standardized numeric predictors, one-hot encoded categorical predictors, and class weighting was estimated on a 70/30 train-test split.

The benchmark model is not presented as an optimal anomaly detector. Instead, it serves as an interpretable out-of-sample reference point for whether the supplied variables contain enough information to discriminate anomalous from non-anomalous records.

VI. RESULTS

1. DESCRIPTIVE PROFILE

Across the full sample, the mean age is 43.47 years, the mean account balance is 5,060.57, the mean transaction amount is 2,508.50, the mean loan amount is 25,501.04, and the mean credit-card balance is 2,487.40. The anomalous subgroup does not differ markedly from the normal subgroup in central tendency for most numerical variables.

Table 2. Selected numerical variables by anomaly status.

Variable	Anomaly mean	Normal mean	Mann-Whitney p-value	Interpretation
Age	44.11	43.43	0.4502	Not significant
Account Balance	5218.96	5050.46	0.3295	Not significant
Transaction Amount	2488.33	2509.79	0.8065	Not significant

Loan Amount	25103.51	25526.42	0.6133	Not significant
Interest Rate	5.40	5.51	0.4797	Not significant
Credit Card Balance	2372.63	2494.73	0.1514	Not significant
Minimum Payment Due	118.63	124.74	0.1514	Not significant
Rewards Points	4898.07	4970.51	0.6620	Not significant

2. CATEGORICAL ASSOCIATIONS

Among the tested categorical variables, only account type and resolution status are statistically significant at the 5% level, and both effects are small. This pattern reinforces the broader conclusion that the anomaly label is weakly connected to the available feature set.

Table 3. Chi-square tests for selected categorical variables.

Variable	Chi-square	df	p-value	Interpretation
Account Type	6.208	1	0.0127	Significant
Resolution Status	3.975	1	0.0462	Significant
Loan Type	3.126	2	0.2095	Not significant
Gender	1.149	2	0.5629	Not significant
Transaction Type	1.069	2	0.5858	Not significant
Loan Status	0.972	2	0.6151	Not significant
Card Type	0.106	2	0.9483	Not significant

3. OUT-OF-SAMPLE BENCHMARK

The logistic-regression benchmark performs poorly on the hold-out sample (ROC-AUC = 0.490; average precision = 0.062; precision = 0.053; recall = 0.400; F1 = 0.093). These results suggest that statistical significance alone is not enough to justify a strong anomaly-detection claim; practical discrimination remains weak.

4. ILLUSTRATIVE STATISTICAL FIGURES




FIGURE 1. Class distribution of records. The anomaly class is rare (6.0%), confirming a highly imbalanced rare-event problem.




FIGURE 2. Transaction amount by class. Median values and spread are broadly similar, consistent with the non-significant Mann-Whitney result.




FIGURE 3. Correlation matrix for selected numerical variables. Correlations between the anomaly flag and major numerical predictors are weak.




FIGURE 4. Anomaly rate by account type. Current accounts show a slightly higher anomaly rate than savings accounts, but the difference remains modest.

VII. DISCUSSION

The expanded literature review helps place the empirical findings in perspective. Much of the pre-2021 financial anomaly literature reported stronger results because the underlying data were richer: transaction sequences, customer histories, merchant concentration, network links, and institutionally validated labels. By contrast, the present dataset contains only limited cross-sectional signal.

The study therefore contributes as a methodological caution and a reproducible demonstration. It shows that a paper can remain publishable when it clearly states limitations, uses appropriate tests, reports rare-event structure, and distinguishes statistical significance from predictive usefulness.

For a stronger submission, future work should engineer temporal and behavioral features, test unsupervised detectors such as Isolation Forest or Local Outlier Factor, examine threshold calibration in precision-recall space, and preferably validate the framework on real institutional or regulator-grade banking data.

1. THEORETICAL, METHODOLOGICAL, AND PRACTICAL IMPLICATIONS

The study offers three forms of contribution. Theoretically, it supports the argument in the anomaly-detection literature that data quality and feature relevance remain more decisive than algorithmic complexity when anomaly labels are weakly connected to observable variables. Methodologically, it provides a transparent workflow - combining descriptive analysis, nonparametric tests, categorical association tests, and an interpretable benchmark classifier - that can be replicated by researchers working with public or low-information banking datasets. Practically, the findings caution analysts and decision-makers against drawing strong operational conclusions from synthetic or weakly structured banking data without richer behavioral variables, stronger validation procedures, and closer alignment between labels and the processes they are supposed to represent.

2. IMPLICATIONS FOR PUBLICATION READINESS AND FUTURE DEVELOPMENT

To align the manuscript more closely with papers typically accepted in Q2-indexed journals, several enhancements are advisable. First, the contribution should be framed more explicitly in terms of methodological value: the study does not merely test anomaly detection, but demonstrates how class imbalance, weak labels, and synthetic feature spaces affect statistical inference in open banking data. Second, the Discussion should continue to distinguish clearly between statistical significance and operational usefulness, since this distinction strengthens the credibility of the paper. Third, the manuscript would benefit from a short subsection on theoretical, methodological, and practical implications, showing how the findings inform banking analytics, fraud-screening workflows, and the design of future synthetic-data studies. Fourth, a clearer limitations statement and a more structured future-research agenda would bring the article closer to the reporting style expected in stronger journal submissions. Finally, if the target journal allows it, the paper should include a concise data-availability statement, conflict-of-interest statement, and, where relevant, funding and author-contribution notes.

VIII. CONCLUSION

This study examined the extent to which a synthetic banking dataset can support the statistical identification of anomalous records in a financial-data environment. The results indicate several important limitations. The dataset is highly imbalanced, with anomalous observations representing only a small proportion of the total sample, while the main numerical variables do not show statistically meaningful differences between anomalous and non-anomalous cases. Although a limited number of categorical variables were associated with anomaly status, these relationships were not sufficiently strong to provide robust explanatory power. Similarly, the benchmark logistic regression model demonstrated weak out-of-

sample predictive performance, suggesting that the available variables do not adequately capture the underlying structure of anomalous behavior.

Overall, these findings suggest that the dataset is more appropriate for exploratory analysis, methodological illustration, and academic demonstration than for operational use in real-world banking risk monitoring. In its current form, it offers value as a framework for testing basic statistical procedures and anomaly detection workflows, but it is not sufficient for drawing strong practical conclusions. Future research should therefore focus on incorporating richer behavioral and transactional features, applying more advanced analytical models, and strengthening validation procedures. The use of real-world, non-synthetic banking data would also substantially improve the reliability, interpretability, and practical relevance of the findings.

Author Contributions

The author conducted the conceptualization, methodology, data analysis, investigation, writing, review, editing, and final approval of the manuscript.

Funding

This research received no external funding.

Data Availability

The dataset will be available from the author upon reasonable request.

Conflicts of Interest

The author declares no conflict of interest.

REFERENCES

1. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90-113.
2. Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40, 100402.
3. Bahnsen, A. C., Aouada, D., & Ottersten, B. (2016). Example-dependent cost-sensitive logistic regression for credit card fraud detection. In *2014 13th International Conference on Machine Learning and Applications (ICMLA) / revised journal-extended work frequently cited in fraud analytics literature*.
4. Bahnsen, A. C., Stojanovic, A., Aouada, D., & Ottersten, B. (2016). Cost sensitive credit card fraud detection using Bayes minimum risk. *Expert Systems with Applications*, 42(7), 3335-3342.
5. Baesens, B., Van Vlasselaer, V., & Verbeke, W. (2015). *Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques*. Wiley.
6. Barnett, V., & Lewis, T. (1994). *Outliers in Statistical Data* (3rd ed.). Wiley.
7. Beneish, M. D. (1999). The detection of earnings manipulation. *Financial Analysts Journal*, 55(5), 24-36.
8. Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602-613.
9. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235-255.
10. Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density-based local outliers. *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 93-104.
11. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Bontempi, G., & Mazzer, Y. (2019). Scarff: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion*, 41, 182-194.
12. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317-331.
13. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15.
14. Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. *2015 IEEE Symposium Series on Computational Intelligence*, 159-166.
15. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784-3797.

16. Fanning, K., & Cogger, K. O. (1998). Neural network detection of management fraud using published financial data. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 7(1), 21-41.
17. Goldstein, M., & Uchida, S. (2016). A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PLOS ONE*, 11(4), e0152173.
18. Hawkins, D. M. (1980). *Identification of Outliers*. Chapman & Hall.
19. He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263-1284.
20. Hodge, V. J., & Austin, J. (2004). A survey of outlier detection methodologies. *Artificial Intelligence Review*, 22, 85-126.
21. Johnson, J. M., & Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of Big Data*, 6, Article 27.
22. Jurgovsky, J., Granitzer, M., Ziegler, K., Calatroni, L., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100, 234-245.
23. Kirkos, E., Spathis, C., & Manolopoulos, Y. (2007). Data mining techniques for the detection of fraudulent financial statements. *Expert Systems with Applications*, 32(4), 995-1003.
24. Kou, Y., Lu, C.-T., Sirwongwattana, S., & Huang, Y.-P. (2004). Survey of fraud detection techniques. *IEEE International Conference on Networking, Sensing and Control*, 749-754.
25. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. *2008 Eighth IEEE International Conference on Data Mining*, 413-422.
26. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559-569.
27. Pang, G., Shen, C., Cao, L., & van den Hengel, A. (2021). Deep learning for anomaly detection: A review. *ACM Computing Surveys*, 54(2), Article 38.
28. Perols, J. L. (2011). Financial statement fraud detection: An analysis of statistical and machine learning algorithms. *Auditing: A Journal of Practice & Theory*, 30(2), 19-50.
29. Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119*.
30. Randhawa, K., Loo, C. K., Seera, M., Lim, C. P., & Nandi, A. K. (2018). Credit card fraud detection using AdaBoost and majority voting. *IEEE Access*, 6, 14277-14284.
31. Ravisankar, P., Ravi, V., Rao, G. R., & Bose, I. (2011). Detection of financial statement fraud and feature selection using data mining techniques. *Decision Support Systems*, 50(2), 491-500.
32. Rousseeuw, P. J., & Croux, C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424), 1273-1283.
33. Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Muller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5), 756-795.
34. Sharma, A., & Panigrahi, P. K. (2013). A review of financial accounting fraud detection based on data mining techniques. *International Journal of Computer Applications*, 39(1), 37-47.
35. Thomas, L. C. (2000). A survey of credit and behavioural scoring: Forecasting financial risk of lending to consumers. *International Journal of Forecasting*, 16(2), 149-172.
36. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47-66.
37. Whitrow, C., Hand, D. J., Juszczak, P., Weston, D., & Adams, N. M. (2009). Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery*, 18, 30-55.
38. Zhao, Y., Nasrullah, Z., & Li, Z. (2019). PyOD: A Python toolbox for scalable outlier detection. *Journal of Machine Learning Research*, 20(96), 1-7.