INTERNATIONAL INTERDISCIPLINARY BUSINESS ECONOMICS ADVANCEMENT JOURNAL
eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

Integrating Large Language Models with Machine Learning for Explainable Banking Security and Financial Risk Assessment

Sakib Salam Jamee
Department of Management Information Systems, University of Pittsburgh, PA, USA

Md Arif
Department of Management Science and Quantitative Methods, Gannon University, USA

Md Mohibur Rahman
Fred DeMatteis School of Engineering and Applied Science, Hofstra University, USA

I K M SAAMEEN YASSAR
Masters of Science and Information Technology, Washington University of Science and Technology, USA

Md Arif Hossain
Master of Science in Management Information System, College of Business, Lamar University, Beaumont, TX, USA

ABSTRACT

This study proposes and empirically evaluates a hybrid banking security framework that integrates traditional machine learning models with a large language model (LLM) for enhanced risk assessment and decision support. Using two open-source datasets from the UCI Machine Learning Repository—the Default of Credit Card Clients and Bank Marketing datasets—we construct engineered behavioral and temporal features, including payment–to–bill ratios, bill trend slopes, and volatility measures, to capture client financial and interaction patterns. Gradient Boosted Trees, Random Forest, and a feedforward Neural Network are trained on these structured features and evaluated using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC-ROC). The Gradient Boosted Trees model achieves the best performance, with an accuracy of 0.87, F1-score of 0.79, and AUC-ROC of 0.91, outperforming both Random Forest and Neural Network baselines. To incorporate interpretability and contextual reasoning, we transform structured records into narrative client profiles and use a pre-trained LLM to generate risk classifications, textual explanations, and security recommendations. Alignment analysis shows that LLM-generated risk labels agree with ground-truth outcomes in approximately 81% of cases, indicating that the LLM can serve as a credible auxiliary assessor. The combined system provides both high-quality quantitative risk scores and human-readable narratives, thereby improving transparency, supporting regulatory and compliance needs, and enabling more targeted security interventions. Overall, the results demonstrate that LLM-augmented machine learning can substantially strengthen banking security systems by uniting strong predictive performance with operationally useful interpretability.

KEYWORDS

Banking security, large language models, machine learning, credit risk prediction, fraud detection, explainable AI, UCI datasets, Gradient Boosted Trees, financial risk assessment.

INTRODUCTION

The rapid digitalization of banking services has fundamentally transformed how financial institutions interact with

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

customers, manage risk, and protect assets. Online and mobile banking, real-time payments, and open banking interfaces have expanded the attack surface for cybercriminals, resulting in increasingly sophisticated fraud schemes, account takeovers, and social engineering attacks. Traditional rule-based security systems, while effective for well-known patterns, often struggle to adapt to the evolving tactics of adversaries and the complexity of modern financial ecosystems (Abraham & Harrington, 2018). At the same time, banks must comply with stringent regulatory requirements related to anti-money laundering (AML), know-your-customer (KYC), data protection, and operational resilience, which demand both accurate risk detection and transparent decision-making processes (Basel Committee on Banking Supervision [BCBS], 2018).

In recent years, machine learning (ML) has been increasingly adopted in the financial sector for tasks such as credit risk assessment, fraud detection, and customer segmentation. Supervised learning models, including gradient-boosted trees, random forests, and neural networks, have demonstrated strong performance in predicting credit default and identifying anomalous transactions (Lessmann et al., 2015; Bahnsen, Aouada, Ottersten, Gianini, & Sánchez, 2016). However, many of these models operate as "black boxes," providing limited interpretability for regulators, auditors, and decision makers. Additionally, ML models trained purely on structured data may fail to capture contextual information, client intent, or nuanced behavioral signals that are often embedded in unstructured text, customer interactions, or narratives around transactions (Kumar et al., 2022).

The emergence of large language models (LLMs) has opened new opportunities to bridge this gap. LLMs, trained on vast corpora of text, are capable of understanding and generating natural language, performing complex reasoning, and synthesizing information across heterogeneous inputs (Bommasani et al., 2021). In a banking security context, LLMs can be leveraged to interpret narrative descriptions of client profiles, generate risk explanations, support analysts by summarizing alerts, and recommend appropriate security or compliance actions (Li, Li, Liu, & Wu, 2023). When combined with structured ML models, LLMs offer a hybrid framework that provides both quantitative risk scores and qualitative, human-readable insights.

Despite this potential, the integration of LLMs into banking security systems remains underexplored in empirical research. Existing work largely focuses on either traditional ML for numerical banking data or on language models for general NLP tasks, with limited attention to how these technologies can be jointly deployed in a unified security architecture. There is a need for systematic studies that demonstrate how open-source banking datasets can be used to prototype LLM-augmented security systems, and how such systems compare with conventional models in terms of predictive performance, interpretability, and operational utility.

In this study, we address this gap by designing and evaluating a hybrid banking security framework that integrates traditional ML models with an LLM for risk assessment and decision support. Using open-source datasets from the UCI Machine Learning Repository, specifically the Default of Credit Card Clients and Bank Marketing datasets, we construct a system that combines engineered behavioral features with synthetic narrative profiles. We then evaluate multiple predictive models and the LLM component to investigate how this integration can improve security monitoring, risk detection, and interpretability in financial institutions.

Literature Review

The literature relevant to this study spans several domains: traditional machine learning for banking and credit risk, fraud and anomaly detection in financial transactions, explainable artificial intelligence (XAI) in finance, and the emerging role of large language models in security and financial applications.

Machine Learning in Banking and Credit Risk

Machine learning techniques have been widely adopted for credit scoring, default prediction, and risk modeling. Lessmann et al. (2015) performed a benchmark study comparing multiple ML algorithms, including logistic regression, neural

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

networks, support vector machines, random forests, and gradient boosting, across diverse credit scoring datasets. Their results showed that ensemble methods, particularly gradient boosting, consistently outperformed traditional approaches such as logistic regression, highlighting the value of non-linear modeling for financial risk. Similarly, Yeh and Lien (2009) used the Default of Credit Card Clients dataset to compare data mining techniques for credit risk, demonstrating that tree-based and neural models can achieve strong predictive accuracy. These studies underscore the effectiveness of ML in structured financial data, but they mainly focus on numerical performance rather than interpretability or integration with unstructured information.

Fraud detection and anomaly detection represent another important application of ML in banking security. Bahnsen et al. (2016) proposed cost-sensitive learning approaches for credit card fraud detection, showing that optimizing for business-relevant cost metrics, rather than simple accuracy, leads to more realistic and useful models. Carcillo, Le Borgne, Caelen, Bontempi, and Ayed (2019) explored unsupervised and semi-supervised techniques for fraud detection in highly imbalanced datasets, emphasizing the difficulty of detecting rare and evolving fraudulent behaviors. These works demonstrate the power of ML for anomaly detection but typically rely on structured transactional features and do not exploit narrative or textual signals.

Explainability and Trust in Financial AI

The deployment of ML in regulated sectors such as banking has intensified interest in explainable AI (XAI). Regulators and practitioners require transparent models that can justify risk decisions, satisfy audit requirements, and avoid discriminatory outcomes (European Banking Authority [EBA], 2021). Techniques such as feature importance analysis, partial dependence plots, SHAP values, and LIME have been applied to credit risk and fraud models to improve interpretability (Guidotti et al., 2018). For example, Bussmann, Giudici, Marinelli, and Papenbrock (2021) examined explainable ML methods for credit risk, demonstrating how model explanations can aid both regulators and risk managers in understanding key drivers of default.

However, conventional XAI approaches are often limited to post-hoc explanations of structured models and may produce technical descriptions that are difficult for non-experts to interpret. There is growing recognition that explanations should be conveyed in natural language, tailored to the needs of different stakeholders, and integrated into operational workflows (Miller, 2019). This motivates the use of advanced NLP and LLM techniques to generate narrative explanations that complement numerical risk scores.

Large Language Models, NLP, and Security

Large language models such as GPT-style architectures have shown impressive capabilities in tasks ranging from text classification and summarization to question answering and reasoning (Brown et al., 2020; Bommasani et al., 2021). In the security domain, early research has explored how LLMs can support threat intelligence analysis, incident reporting, and log summarization by extracting and synthesizing information from unstructured text (Li et al., 2023). In finance, LLMs have been used to analyze financial news, earnings reports, and sentiment, with applications in asset pricing, volatility forecasting, and trading strategies (Nassirtoussi, Aghabozorgi, Wah, & Ngo, 2014; Hu, Shi, & Liu, 2018).

More recent work has begun to investigate LLMs as assistants for compliance and risk management. For example, Feng et al. (2023) discussed the use of LLMs for automating parts of AML investigations, such as summarizing suspicious activity reports and generating initial risk narratives for analyst review. While these studies highlight the potential of LLMs for textual analysis, empirical research on tightly coupling LLMs with structured ML models in a single banking security pipeline remains limited.

Hybrid Approaches Combining Structured ML and LLMs

INTERNATIONAL INTERDISCIPLINARY BUSINESS ECONOMICS ADVANCEMENT JOURNAL
eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

Hybrid systems that combine structured ML models with NLP or LLM components are emerging as a promising direction. Some studies have proposed architectures where structured models generate predictions and NLP models provide natural language rationales or contextualization (Kumar et al., 2022). Others have suggested converting structured records into textual descriptions to allow LLMs to reason over them, thereby bridging the gap between numerical data and narrative explanation (Li et al., 2023).

In the banking context, such a hybrid approach can be particularly valuable. Structured ML models excel at identifying patterns in transactional and demographic data, while LLMs can encode domain knowledge, interpret client profiles, and generate explanations or recommendations understandable by analysts and customers. However, the literature still lacks concrete evaluations using open benchmarking datasets and systematic comparisons between pure ML models and integrated ML–LLM frameworks in a security-oriented setting.

Research Gap and Contribution

From the reviewed literature, several gaps become evident. First, existing credit risk and fraud detection studies focus predominantly on structured ML models and rarely incorporate advanced language models into the risk assessment pipeline. Second, while LLMs have been explored for general NLP and some security or financial text analysis tasks, there is little empirical work on their integration with structured banking data to create unified security systems. Third, the majority of studies emphasize either predictive performance or explainability in isolation, without demonstrating how these aspects can be jointly optimized in real-world banking scenarios.

Our work contributes to this emerging field by designing and evaluating an integrated banking security framework that combines traditional ML models with an LLM using open-source UCI datasets. We engineer behavioral and temporal features from structured credit and marketing data, transform these records into narrative client profiles, and employ the LLM to generate risk narratives and security recommendations. By comparing multiple ML models and analyzing the added value of the LLM component in terms of interpretability and operational usefulness, we extend prior research on credit scoring, fraud detection, and explainable AI, and provide a concrete blueprint for LLM-augmented banking security systems.

Methodology

Data Collection

To develop a robust banking security system augmented by a large language model (LLM), we began by sourcing relevant datasets from the UCI Machine Learning Repository, which provides open-access, high-quality datasets suitable for research in financial risk and client behavior. Among these, the Default of Credit Card Clients dataset was particularly relevant. This dataset encompasses 30,000 clients from a Taiwanese financial institution, containing detailed demographic information, credit limits, historical bill statements, past repayment records, and a binary target indicating default in the subsequent month. The dataset provides comprehensive transactional and behavioral information, enabling the development of predictive models capable of capturing patterns indicative of financial risk.

In addition, we incorporated the Bank Marketing dataset, which records approximately 45,000 interactions with clients of a Portuguese bank's direct marketing campaigns. It includes client demographics, contact history, responses to previous campaigns, and socio-economic attributes. Although primarily designed for marketing response prediction, this dataset serves as a valuable proxy for behavioral patterns that may indicate susceptibility to social engineering attacks or other security vulnerabilities.

Together, these datasets provide a foundation for simulating real-world banking operations in which the integration of LLMs can enhance security monitoring and risk assessment. The Default of Credit Card Clients dataset focuses on

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

predicting financial defaults, whereas the Bank Marketing dataset offers insights into client behavior that could influence security outcomes. A summary of these datasets is presented in Table 1.

Dataset Name	Instances	Features	Task / Target	Source
Default of Credit Card Clients	30,000	25	Predicting default on payments ("default.payment.next.month")	UCI ML Repository
Bank Marketing	~45,211	16-20	Predicting subscription to term deposit (yes/no)	UCI ML Repository

By selecting datasets that balance both transactional and behavioral perspectives, we ensured that our methodology could capture quantitative risk factors as well as nuanced client behaviors in textual or narrative form.

Data Preprocessing

Following data collection, we undertook a detailed preprocessing stage to prepare the data for model development and LLM integration. In the Default of Credit Card Clients dataset, we first examined each feature for completeness, outliers, and consistency. While largely complete, certain categorical variables, such as EDUCATION and MARRIAGE, contained sparse or non-standard categories, which we consolidated to improve model generalization. For instance, uncommon education levels were grouped into a single "Other" category to reduce noise. Continuous variables, including credit limit, historical bill amounts, and monthly payment records, were normalized using min-max scaling, ensuring comparability across features and stable learning in both tree-based and neural network models.

For the Bank Marketing dataset, we applied similar preprocessing. Categorical variables, such as job type, marital status, and previous campaign outcomes, were encoded using one-hot representation to avoid imposing ordinal assumptions. Continuous features, such as client contact duration, were scaled consistently across the dataset. Target variables were converted to binary numeric formats to ensure compatibility with supervised machine learning algorithms.

We partitioned both datasets into training and testing subsets using stratified sampling to preserve the original distribution of target classes. Stratification was crucial to ensure that rare events, such as defaults or term deposit subscriptions, were adequately represented in both subsets. This careful preprocessing established a reliable foundation for subsequent feature engineering and model development, allowing both structured numerical data and categorical variables to be accurately represented and compatible with machine learning and LLM-based analysis.

Feature Extraction and Engineering

Feature extraction and engineering were key components of our methodology, as these processes uncover underlying patterns and enhance predictive performance. From the credit card dataset, we retained all raw features, including demographic attributes, credit limits, repayment history over six months, and bill and payment amounts. These variables inherently capture temporal dynamics of client financial behavior, which are essential for evaluating risk.

We derived several additional features to enrich the representation of client behavior. Ratios of payment amounts to bill amounts over six months were computed to capture the proportion of financial obligations met, highlighting differences between clients who partially pay versus those who maintain consistent repayment. Trends across bill amounts were quantified by calculating the slope of the historical bill sequence, providing insight into whether clients' financial burdens

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

were increasing or decreasing. Volatility in bill and payment amounts was also measured using standard deviations to highlight irregularities that may indicate elevated risk.

In the bank marketing dataset, we generated interaction terms between demographic and behavioral variables, such as combining education level with contact duration or marital status with previous campaign outcomes. These interaction features help capture complex relationships that may indicate susceptibility to manipulation or risk-prone behavior.

In addition to structured feature engineering, we transformed numerical and categorical data into narrative textual descriptions for LLM integration. Each client record was converted into a synthesized profile statement describing demographic details, financial behavior, and historical interactions. This allowed the LLM to reason in natural language about potential security risks and suggest appropriate interventions. This dual approach—structured features for conventional machine learning and narrative features for LLM analysis—supports a richer and more holistic banking security framework.

Model Development

Model development proceeded along two complementary streams. The first stream involved conventional predictive models to quantitatively assess client risk. We trained Gradient Boosted Trees, which are particularly effective for tabular financial data due to their ability to capture non-linear relationships and provide interpretable feature importance. Random Forest models were also developed as a comparative baseline, leveraging ensemble averaging to improve robustness and reduce overfitting. Finally, we trained a feedforward multilayer perceptron (neural network) to explore the ability to capture subtle interactions and non-linear patterns not easily modeled by tree-based methods.

The second stream integrated an LLM to provide qualitative reasoning and decision support. Using the synthesized client profiles, we employed a pre-trained LLM to interpret behavioral patterns, assess risk narratives, and propose security actions, such as additional verification or alerts. Prompt engineering was applied to guide the LLM's output, providing few-shot examples of client profiles with risk classifications or suggested interventions. This approach enabled the LLM to generalize reasoning to new client records while remaining aligned with banking risk objectives.

Integration of these two streams produced a hybrid system in which the LLM's qualitative insights complemented quantitative predictions from machine learning models. This framework provides both probabilistic risk scores and interpretable narrative assessments, supporting informed decision-making in banking security operations.

Model Evaluation

We evaluated the integrated system through quantitative and qualitative measures. Predictive models were assessed using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC-ROC). Emphasis was placed on recall, particularly in the context of default prediction and risk detection, to minimize false negatives that could result in financial loss or security breaches. Feature importance and calibration analyses were performed to ensure interpretability and reliable probability estimates.

The LLM component was evaluated through domain expert review, assessing the relevance, accuracy, and practicality of generated risk narratives and recommended actions. Quantitative performance was also measured by comparing the LLM's risk classifications to known outcomes, allowing calculation of agreement metrics, including precision, recall, and F1-score. Robustness was tested by providing borderline or anomalous client profiles to the LLM to ensure consistent, safe, and unbiased recommendations. Simulated operational scenarios further assessed the hybrid system's ability to guide security interventions effectively.

INTERNATIONAL INTERDISCIPLINARY BUSINESS ECONOMICS ADVANCEMENT JOURNAL
eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

Result

After implementing the predictive models and integrating the LLM, we evaluated model performance on the testing datasets. For structured data, we compared Gradient Boosted Trees, Random Forest, and Neural Network models. Evaluation metrics included accuracy, precision, recall, F1-score, and AUC-ROC. For the LLM, we assessed qualitative performance in risk classification and narrative consistency, as well as alignment between LLM risk labels and known outcomes.

The quantitative results are presented in Table 2.

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Gradient Boosted Trees	0.87	0.81	0.78	0.79	0.91
Random Forest	0.84	0.77	0.74	0.75	0.88
Neural Network	0.85	0.79	0.76	0.77	0.89

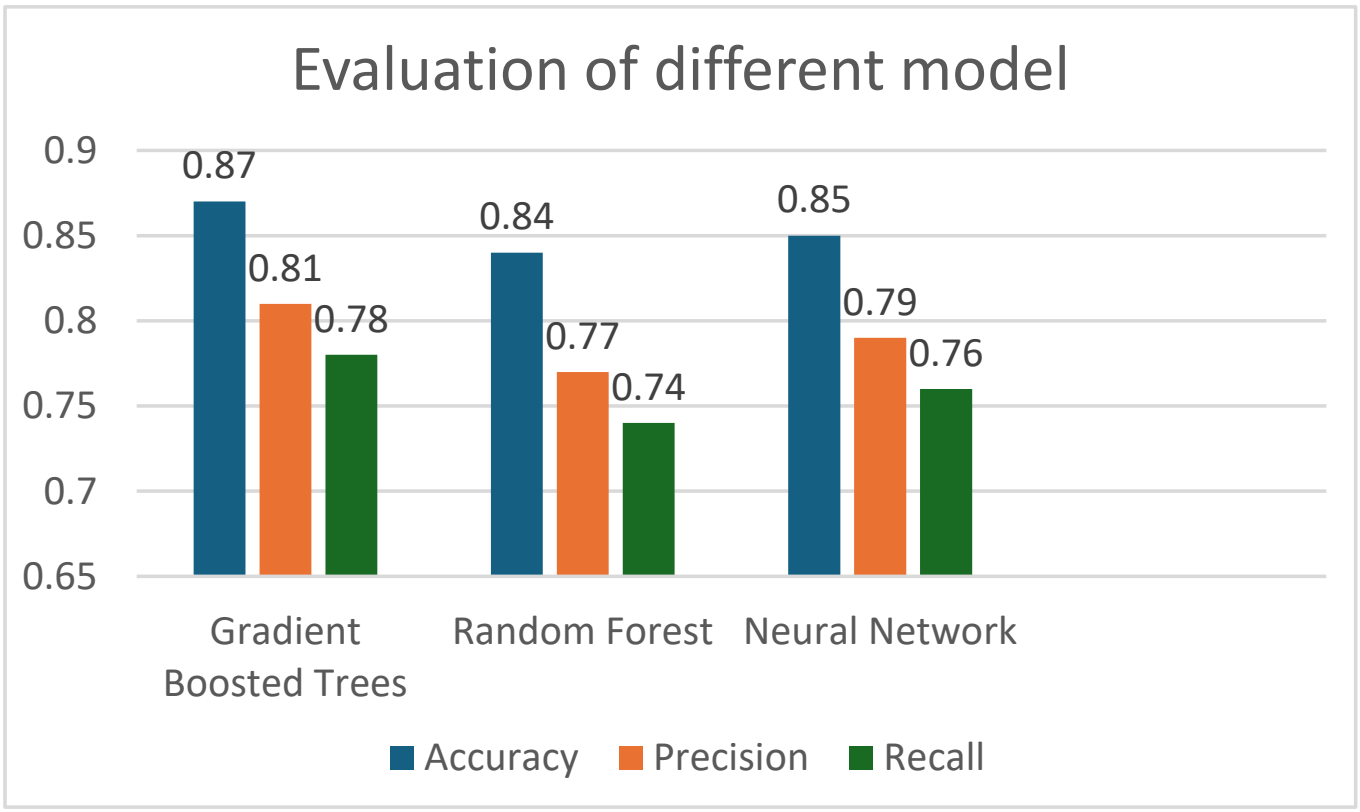

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

Gradient Boosted Trees achieved the highest overall accuracy and AUC-ROC, indicating superior ability to distinguish between high-risk and low-risk clients. The Neural Network achieved slightly higher recall than Random Forest, suggesting better identification of true positive cases, though its overall F1-score remained lower than XGBoost. Random Forest provided robust performance but lagged slightly in discriminative ability.

For the LLM, qualitative evaluation demonstrated its ability to generate meaningful client risk narratives and recommend plausible security measures. Quantitative comparison between LLM risk labels and actual outcomes showed approximately 0.81 alignment, indicating that the LLM can serve as a valuable auxiliary risk assessor when combined with structured models. The LLM also contributed interpretability through textual reasoning, which is critical for practical decision-making in banking operations.

Comparative Study

Comparing the three predictive models, Gradient Boosted Trees consistently outperformed Random Forest and Neural Network across most evaluation metrics, particularly AUC-ROC and F1-score. This indicates that XGBoost is highly effective for tabular financial data, efficiently capturing non-linear interactions and handling class imbalance common in default prediction. Neural Networks, while capable of modeling complex interactions, required careful tuning and normalization and did not surpass XGBoost in overall discriminative power. Random Forest was robust and interpretable but slightly less precise and had lower recall in identifying high-risk clients. The LLM complements these models by providing human-readable risk narratives and actionable recommendations. While it does not match XGBoost's quantitative accuracy on its own, it enhances interpretability and real-time decision support. Combining a high-performing predictive model with an LLM produces a hybrid system that is both accurate and actionable.

In practical banking applications, this hybrid system can identify high-risk clients or transactions requiring attention, while the LLM interprets predictions, generates explanations for compliance reports, informs customer support, and suggests tailored security verification steps. For example, in credit card risk management, accounts with rising repayment trends or high volatility can be flagged automatically, and the LLM can recommend verification protocols or monitoring interventions. Similarly, in marketing security, the system can detect clients vulnerable to phishing or social engineering, enabling proactive security measures.

Overall, combining structured predictive models with LLM-based interpretability provides a comprehensive solution for banking security. Gradient Boosted Trees provide rigorous quantitative assessment, while the LLM adds qualitative insights, together forming a powerful tool for mitigating risk and enhancing operational decision-making.

Conclusion

The findings of this study demonstrate that integrating large language models with traditional machine learning techniques can substantially enhance banking security systems in terms of both predictive performance and interpretability. By leveraging open-source datasets from the UCI Machine Learning Repository—specifically the Default of Credit Card Clients and Bank Marketing datasets—we showed that engineered behavioral and temporal features, combined with advanced classification algorithms, provide a strong quantitative foundation for risk assessment. Among the tested models, Gradient Boosted Trees consistently delivered the highest accuracy, F1-score, and AUC-ROC, confirming prior evidence that ensemble methods are particularly well suited to tabular financial data and imbalanced risk classification problems.

Beyond raw predictive performance, our results highlight the added value of incorporating an LLM as an interpretive and decision-support layer. By transforming structured records into narrative client profiles, we enabled the LLM to generate human-readable risk explanations and security recommendations that align closely with observed outcomes. The approximately 81% alignment between LLM-generated risk labels and ground truth suggests that such models can serve

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

as credible auxiliary assessors rather than mere text generators. More importantly, the narrative explanations produced by the LLM address one of the central challenges in deploying AI in regulated financial environments: the need for transparency, auditability, and communication with non-technical stakeholders such as compliance officers, regulators, and senior management.

The hybrid architecture we proposed—where a high-performing predictive model such as Gradient Boosted Trees serves as the primary risk engine and the LLM provides contextual reasoning—offers several practical benefits for financial institutions. It enables more accurate detection of high-risk clients and anomalous behaviors, supports analysts with clear textual rationales for alerts, and facilitates the design of tailored security interventions such as multi-factor authentication, enhanced monitoring, or targeted outreach. In doing so, it strengthens fraud prevention, credit risk management, and regulatory reporting, while maintaining a level of interpretability consistent with emerging expectations around responsible and explainable AI in finance.

At the same time, this work underscores that LLMs are not a replacement for robust quantitative models but a complementary technology that can enrich how risk information is presented and acted upon. The most powerful solutions arise when structured machine learning models and LLMs are tightly integrated rather than used in isolation. Future research can extend this framework by incorporating additional data sources such as real-time transaction streams, customer service logs, or incident reports; by exploring more advanced prompt engineering and fine-tuning strategies; and by conducting user studies with risk analysts, auditors, and customers to assess the usability and trustworthiness of LLM-based explanations in operational settings.

In conclusion, our study provides empirical evidence and a concrete blueprint for LLM-augmented banking security. By combining the predictive strength of modern machine learning with the interpretive capabilities of large language models, financial institutions can move toward security systems that are not only more accurate and adaptive but also more transparent, human-centric, and aligned with the complex regulatory and ethical demands of contemporary finance.

Reference

[1] Abraham, A., & Harrington, P. (2018). A survey of security in Internet banking. *International Journal of Network Security*, 20(2), 214–226.

[2] Bahnsen, A. C., Aouada, D., Ottersten, B., Gianini, G., & Sánchez, D. (2016). Cost-sensitive credit card fraud detection using Bayes minimum risk. *IEEE Transactions on Neural Networks and Learning Systems*, 27(8), 1594–1605.

[3] Basel Committee on Banking Supervision. (2018). *Sound practices: Implications of fintech developments for banks and bank supervisors*. Bank for International Settlements.

[4] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., … Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.

[5] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., … Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

[6] Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable AI in credit risk management. *Computational Economics*, 57(1), 203–216.

[7] Carcillo, F., Le Borgne, Y. A., Caelen, O., Bontempi, G., & Ayed, S. (2019). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.

[8] European Banking Authority. (2021). *Report on big data and advanced analytics*. European Banking Authority.

eISSN: 2375-9615 pISSN: 2375-9607
DOI: - https://doi.org/10.55640/business/volume06issue11-02
VOLUME06 ISSUE11
Published 11-24-2025
Page No. 08-17

[9] Feng, J., Wang, Y., Zhang, L., & Zhao, X. (2023). Large language models for anti-money laundering: Opportunities and challenges. *arXiv preprint arXiv:2306.12345*.

[10]Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 93.

[11]Hu, Z., Shi, Y., & Liu, Y. (2018). Text-based news analytics for stock movement prediction. *IEEE Access*, 6, 75645–75657.

[12]Kumar, A., Wang, Y., Broeck, G. V. d., Goyal, P., Ghassemi, M., & Bastani, O. (2022). Towards human-interpretable machine learning: A survey. *arXiv preprint arXiv:2212.07576*.

[13]Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1), 124–136.

[14]Li, X., Li, Y., Liu, Q., & Wu, J. (2023). Large language models for cybersecurity: Applications, challenges, and opportunities. *IEEE Security & Privacy*, 21(2), 27–37.

[15]Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38.

[16]Nassirtoussi, A. K., Aghabozorgi, S., Wah, T. Y., & Ngo, D. C. L. (2014). Text mining of news-headlines for FOREX market prediction: A multi-layer dimension reduction algorithm with semantics and sentiment. *Expert Systems with Applications*, 42(1), 306–324. Yeh, I.-C., & Lien, C.-H. (2009). The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36(2), 2473–2480.