


Transformer-Based Representation Learning for Multivariate Financial Risk Forecasting and Credit-Risk Analytics

Lars-Alexander Kuehn¹, Mohit Beniwal²
Carnegie Mellon University, USA



Article History

Received: 20th Sep 2026

Revised: 25th Sep 2026

Accepted: 27th Sep 2026

Published: 29th Sep 2026

Abstract

Financial risk forecasting increasingly requires models capable of learning complex relationships across heterogeneous financial variables, temporal observations, and changing market conditions. Conventional statistical and machine-learning models can provide effective financial-risk predictions but may face limitations when financial indicators exhibit nonlinear interactions, long-range temporal dependencies, and rapidly changing patterns. Transformer architectures provide an alternative representation-learning framework through attention mechanisms that can identify relationships across multiple observations and financial variables. This study examines Transformer-based representation learning for multivariate financial risk forecasting and credit-risk analytics. A conceptual literature-based methodology was adopted to synthesize recent research on financial risk prediction, deep learning, Transformer architectures, credit-risk modeling, ensemble learning, explainable artificial intelligence, portfolio optimization, and multimodal financial analytics. The analysis indicates that Transformer-based models can represent temporal dependencies and interactions among financial variables while supporting hybrid architectures that combine attention mechanisms with established machine-learning models. Recent research demonstrates applications of Transformers to corporate financial risk prediction, financial distress assessment, credit-risk modeling, and explainable financial forecasting. The literature also indicates that Transformer models can benefit from integration with XGBoost, SHAP, temporal attention, multimodal representations, and other predictive architectures. However, challenges remain concerning data requirements, computational complexity, model interpretability, temporal validation, class imbalance, distributional shifts, and deployment within regulated financial environments. The study proposes a conceptual framework in which multivariate financial information is transformed into contextual representations, Transformer attention mechanisms learn temporal and cross-variable relationships, and downstream risk models generate financial-risk and credit-risk predictions. Explainability mechanisms can subsequently identify the financial variables and historical periods contributing to individual predictions. The framework provides a basis for developing data-driven financial-risk systems capable of combining representation learning,

HOW TO CITE

Transformer-Based Representation Learning for Multivariate Financial Risk Forecasting and Credit-Risk Analytics. (2026). *International Journal of AI, Engineering and Management Studies (IJAEMS)*, 1(2), 190-199 https://essayjournals.in/index.php/home/article/view/IJAEMS_v1i2_18

	<i>temporal forecasting, credit-risk analytics, and interpretable decision support.</i> Keywords <i>Transformer models, financial risk forecasting, credit-risk analytics, representation learning, deep learning, financial prediction, attention mechanisms</i>
--	---

Copyright (c) 2026 Lakshmi Priya
This work is licensed under a [Creative Commons Attribution 4.0 International License](#).



INTRODUCTION

Financial institutions increasingly rely on data-driven models to forecast financial risk, evaluate corporate financial health, assess creditworthiness, and support investment decisions. Financial datasets frequently contain large numbers of interrelated variables, including profitability ratios, liquidity measures, leverage indicators, market information, transaction records, and historical financial statements. These variables may interact in nonlinear ways and their relationships can change over time.

Traditional statistical models and machine-learning methods remain important in financial risk prediction. Mahajan et al. (2024) demonstrated the relevance of multivariate data analytics and ensemble learning for financial-risk prediction. Ensemble models can capture nonlinear relationships and provide strong predictive performance, but they may not always represent long-range temporal dependencies efficiently when financial information is organized as sequential observations.

Deep learning has expanded the possibilities for financial-risk modeling by learning representations directly from complex data. Mienye et al. (2026) reviewed deep-learning approaches to credit-risk prediction and identified Transformers as one of the emerging architectures capable of capturing temporal, nonlinear, and relational dependencies in credit-risk data. Their review places Transformer models alongside recurrent networks, convolutional networks, and graph neural networks within the modern credit-risk modeling landscape.

Transformer architectures are particularly relevant because their attention mechanism allows a model to evaluate relationships between different elements of an input sequence. In financial applications, these elements may represent financial indicators observed across different reporting periods, borrower behavior over time, or multiple financial variables recorded simultaneously.

Recent research has begun to apply Transformers directly to corporate financial-risk prediction. Liu et al. (2026) developed a framework combining graph Transformer networks, time-series contrastive learning, and meta-learning-based adaptive optimization. The model was designed to capture enterprise relationships, temporal patterns, and distributional changes across financial environments.

Transformer models can also be integrated with conventional machine-learning algorithms. Jiang (2026) developed a hybrid XGBoost-Transformer framework incorporating SHAP and temporal attention. The approach combines Transformer-based temporal representation

learning with XGBoost classification and traces model predictions back to financial indicators and historical time steps using SHAP.

This study therefore examines Transformer-based representation learning as a framework for multivariate financial-risk forecasting and credit-risk analytics. The objective is to synthesize the emerging literature and develop a conceptual architecture connecting financial representation learning, temporal attention, credit-risk prediction, explainability, and downstream financial decision support.

BACKGROUND OF THE STUDY

Financial risk prediction involves estimating the probability or severity of adverse financial outcomes using information from firms, borrowers, markets, and economic environments. Conventional financial-risk models commonly use financial ratios and structured indicators. As financial datasets have expanded, however, risk models increasingly need to process multiple variables across multiple time periods.

Mahajan et al. (2024) demonstrate the relevance of multivariate data analytics and ensemble learning for financial-risk prediction. Their work illustrates how multiple financial variables can be integrated into predictive systems. However, conventional ensemble approaches may require manually designed temporal representations when the objective involves learning patterns across multiple reporting periods.

Representation learning provides an alternative. Rather than relying exclusively on manually engineered features, deep-learning models can learn internal representations from financial data. Transformer architectures extend this principle by using attention mechanisms to determine relationships among elements of an input sequence.

In credit-risk applications, the input sequence could consist of a borrower's historical financial statements, repayment observations, transaction patterns, or changes in financial ratios. The Transformer can examine relationships between different periods and variables and produce contextual representations for downstream risk prediction.

Mienye et al. (2026) identify Transformers as an important component of modern deep-learning research in credit-risk prediction. Their review emphasizes that deep-learning architectures can capture nonlinearities, temporal dynamics, and relational dependencies that are difficult to represent using traditional scorecard approaches.

The development of Transformer-based financial-risk models is also connected to explainability. Financial institutions often require information about why a model identifies a company or borrower as high risk. Jiang (2026) addresses this requirement by combining Transformer representation learning with XGBoost and SHAP. The framework is designed to trace predictive contributions back to financial indicators and historical time periods.

Another important development involves combining Transformer models with graph-based representations. Liu et al. (2026) proposed graph Transformer networks combined with time-series contrastive learning and meta-learning. This approach recognizes that corporate financial risk may depend not only on a company's own historical financial indicators but also on relationships between enterprises and changing distributions across sectors.

These developments indicate that Transformer-based financial-risk forecasting is evolving from simple sequential prediction toward richer representation-learning architectures that can combine temporal, relational, and heterogeneous financial information.

LITERATURE REVIEW

Financial-risk prediction has traditionally relied on statistical models and machine-learning methods. Mahajan et al. (2024) examined multivariate data analytics and ensemble learning, demonstrating the continued importance of combining multiple financial indicators and predictive algorithms.

Deep learning provides a more flexible representation-learning framework. Mienye et al. (2026) systematically reviewed deep-learning approaches for credit-risk prediction and identified Transformers as one of the major emerging architectures. Their review covers Transformer, recurrent, convolutional, and graph-based models and emphasizes the importance of matching architectures with the underlying financial data structure.

Transformer models are particularly suitable for sequential financial data because attention mechanisms can capture relationships across distant observations. Financial risk often develops gradually, meaning that a company's current financial condition may depend on patterns that emerged across several previous reporting periods. Transformer representations can potentially capture these relationships without relying exclusively on recurrent processing.

Liu et al. (2026) provide direct evidence of this approach through a hybrid architecture involving graph Transformer networks, time-series contrastive learning, and meta-learning-based adaptive optimization. Their framework combines relational representation learning with temporal pattern extraction and adaptation across financial environments.

Transformer representation learning can also be combined with tree-based machine learning. Jiang (2026) proposed an XGBoost-Transformer framework in which Transformer attention mechanisms extract temporal representations before XGBoost performs risk classification. SHAP is then used to explain the contribution of financial indicators and historical periods. The reported model achieved an accuracy of 0.891, precision of 0.826, recall of 0.723, and ROC-AUC of 0.809 in the study's experimental setting.

Abbas et al. (2026) developed an interpretable hybrid stacking model for multi-class loan-default prediction. This work is relevant to Transformer-based financial analytics because it demonstrates the importance of moving beyond binary risk labels toward more granular risk tiers. Multi-class risk representations can provide richer information for financial institutions that need to distinguish between different levels of borrower risk.

Yang et al. (2025) investigated a time-series-driven ResE-BiLSTM framework for post-loan default detection. Although the architecture is recurrent rather than Transformer-based, it demonstrates the importance of temporal representation learning for detecting changes in borrower risk after credit has been issued. Such research provides a useful comparison point for Transformer architectures.

Feng et al. (2025) examined federated risk discrimination using Siamese networks for financial transaction anomaly detection. While the application focuses on transaction

anomalies rather than direct credit-risk forecasting, it demonstrates how advanced representation-learning architectures can operate in distributed financial environments.

Wang et al. (2026) investigated multiview data through discrete and regularized deep learning for dynamic financial-risk prediction. Multiview representation is particularly relevant to Transformer systems because financial risk can be represented simultaneously through accounting information, market indicators, operational variables, and other data sources.

Sharma et al. (2026) examined artificial-intelligence-driven financial-risk prediction and portfolio optimization. Their research demonstrates the increasing connection between financial-risk prediction and downstream investment decisions. Transformer-based representations could potentially serve as an upstream forecasting component within such integrated systems.

Saputra et al. (2026) systematically reviewed deep learning for credit-risk prediction in fintech lending, with attention to model architectures and imbalanced data. Their research highlights the importance of data imbalance when applying deep-learning methods to credit risk. Transformer models can still be affected by imbalanced default observations and therefore require appropriate sampling, weighting, or learning strategies.

Khoozani et al. (2026) reviewed graph learning and explainable methods for financial AI, particularly credit risk and fraud detection. Their work reinforces the importance of combining complex financial representation learning with interpretability.

Caro-González et al. (2026) reviewed financial-risk prediction models incorporating environmental, social, and governance factors. Their work indicates that modern financial-risk systems increasingly incorporate nontraditional risk dimensions. Transformer models can potentially provide a flexible representation mechanism for combining conventional financial variables with ESG-related information.

Jindal et al. (2026) examined financial services and risk intelligence powered by large language models. Although large language models and financial Transformers are not identical applications, the research demonstrates the broader movement toward Transformer-based representation learning across financial services.

Akdoğan (2026) examined sentiment analysis as an alternative credit-rating signal. Textual information can provide additional information beyond conventional financial ratios, creating an opportunity for Transformer-based systems to integrate textual and structured financial data.

Ding et al. (2026) developed a dynamic multi-criteria portfolio-selection framework integrating Transformer-based financial forecasting with peer-prediction trees. This demonstrates how Transformer forecasting can be connected directly with portfolio-selection decisions and multi-criteria investment analysis.

METHODOLOGY

This study adopted a conceptual literature-synthesis methodology. The selected references were examined according to their contribution to financial-risk prediction, credit-risk analytics, Transformer representation learning, deep learning, explainability, temporal modeling, and portfolio decision support.

The literature was organized into five analytical dimensions: financial representation learning, temporal attention, credit-risk prediction, explainability, and downstream financial decision-making.

Analytical Dimension	Main Techniques	Financial Application
Representation learning	Transformer, deep neural networks	Multivariate financial-risk representation
Temporal modeling	Attention, time-series learning, BiLSTM	Financial forecasting and default detection
Credit analytics	Transformer, ensemble learning, stacking	Credit-risk and loan-default prediction
Explainability	SHAP, interpretable models	Understanding financial-risk predictions
Decision support	Portfolio optimization, multi-criteria models	Investment and financial decisions

Because no original financial dataset was supplied, the Results section synthesizes findings from the cited literature rather than reporting a new empirical experiment.

RESULTS

The literature indicates that Transformer-based representation learning provides a flexible approach for financial-risk forecasting because it can model relationships across multiple financial variables and historical observations. Mienye et al. (2026) identify Transformers as an important emerging architecture within deep-learning credit-risk research.

A second result concerns the ability of Transformer architectures to model complex temporal and relational structures. Liu et al. (2026) combine graph Transformers with time-series contrastive learning and meta-learning to address corporate financial-risk prediction. Their framework demonstrates how Transformer representations can extend beyond simple sequential modeling toward relational financial analysis.

A third result concerns hybrid modeling. Jiang (2026) demonstrates that Transformer representations can be combined with XGBoost rather than replacing tree-based models. The resulting architecture uses Transformer attention to represent temporal financial information and XGBoost for risk classification, while SHAP provides interpretation. The reported accuracy, precision, recall, and ROC-AUC demonstrate the potential of this hybrid strategy within the study's experimental setting.

A fourth result concerns credit-risk representation. Abbas et al. (2026) demonstrate the relevance of interpretable hybrid stacking for multi-class loan-default prediction. This supports the development of financial-risk systems that distinguish among several levels of risk rather than relying exclusively on binary default classifications.

A fifth result concerns temporal risk detection. Yang et al. (2025) demonstrate the importance of time-series representation for post-loan default detection through a ResE-BiLSTM framework. Although recurrent rather than Transformer-based, this research reinforces the importance of historical borrower behavior in credit-risk analytics.

A sixth result concerns multimodal and multiview information. Wang et al. (2026) demonstrate the relevance of multiview data for dynamic financial-risk prediction, while Akdoğan (2026) shows how sentiment information can provide an alternative signal for credit-risk classification. These findings support the use of representation-learning systems capable of integrating structured and unstructured information.

A seventh result concerns the connection between forecasting and decision-making. Sharma et al. (2026) connect AI-driven financial-risk prediction with portfolio optimization, while Ding et al. (2026) directly integrate Transformer-based financial forecasting into dynamic multi-criteria portfolio selection. This demonstrates that Transformer forecasting can function as part of a broader financial decision-support pipeline rather than as an isolated prediction task.

DISCUSSION

The findings indicate that Transformer-based representation learning offers several potential advantages for financial-risk analytics. The most important is the ability to represent relationships across multiple observations and variables without relying entirely on manually constructed temporal features.

Financial risk is inherently dynamic. A company's financial condition can change across reporting periods, and the importance of individual indicators can vary depending on the broader economic environment. Attention-based representations provide a mechanism for identifying which observations and relationships are most relevant to a particular prediction.

This capability is especially important for credit-risk analytics. A borrower's risk may depend not only on its current liquidity or profitability but also on how those indicators have changed over time. A sudden deterioration in liquidity, for example, may carry different implications from a consistently low liquidity position. Transformer-based temporal representations can potentially capture such patterns.

The literature also supports hybrid rather than purely Transformer-based architectures. Jiang (2026) combines Transformer and XGBoost, demonstrating that representation learning and tree-based classification can operate together. This is important because tree-based models remain highly effective for structured financial data and can provide useful feature-level interpretation when combined with appropriate explanation techniques.

Explainability is a major consideration. Transformer models can contain large numbers of parameters and complex attention patterns, making direct interpretation difficult. The use of SHAP in Jiang's hybrid framework provides one approach for connecting complex representations with original financial indicators. Khoozani et al. (2026) similarly emphasize the growing importance of explainability in financial AI.

Temporal explainability is particularly valuable. A financial institution may need to know not only which variable contributed to a risk prediction but also which historical period was most influential. The temporal-attention and SHAP combination examined by Jiang (2026) provides an example of this direction.

Another important issue is data volume. Transformer models generally require substantial training data compared with simpler financial models. Mienye et al. (2026) note that although Transformers provide opportunities for modeling complex financial relationships,

deployment remains constrained by issues including data availability, computational requirements, interpretability, and validation.

Class imbalance is also important. Credit defaults are typically less frequent than non-default observations. Saputra et al. (2026) identify imbalanced-data handling as a major issue in deep-learning credit-risk research. Transformer models therefore require appropriate data preparation and evaluation strategies to ensure that strong overall accuracy does not conceal poor identification of risky borrowers.

Temporal validation is another essential consideration. Randomly dividing financial observations into training and testing datasets can produce overly optimistic results when observations from future periods influence the model during training. Financial-risk models should therefore be evaluated using time-ordered or out-of-time validation whenever possible.

Distributional change presents a related problem. Financial markets, corporate conditions, and borrower behavior can change because of economic crises, interest-rate changes, regulatory reforms, technological developments, or other shocks. Liu et al. (2026) address this issue through meta-learning-based adaptation to distributional shifts.

The integration of multiple financial information sources represents another important direction. Financial-risk prediction can incorporate accounting data, market indicators, textual disclosures, sentiment information, ESG factors, and transaction histories. Akdoğan (2026), Caro-González et al. (2026), and Wang et al. (2026) demonstrate the growing importance of alternative and multiview information in financial risk analysis.

Transformer-based representation learning can consequently be positioned as a central layer within a broader financial-risk architecture. The first layer collects structured and unstructured financial information. The second layer transforms these inputs into contextual representations. The third layer uses attention mechanisms to identify temporal and cross-variable relationships. The fourth layer generates credit-risk or financial-risk predictions through classification or forecasting models. The fifth layer provides explanations and connects the predictions to financial decision-making.

Such an architecture can support both banking and investment applications. Banks can use Transformer representations for credit-default prediction and borrower monitoring, while investment institutions can use them for corporate financial-risk forecasting, portfolio monitoring, and investment-risk assessment.

CONCLUSION

Transformer-based representation learning provides an emerging framework for multivariate financial-risk forecasting and credit-risk analytics. Its ability to capture relationships across financial variables and historical observations makes it particularly relevant to financial environments characterized by nonlinear relationships, temporal dependencies, and heterogeneous information.

The reviewed literature demonstrates applications of Transformers and hybrid Transformer architectures to corporate financial-risk prediction, credit-risk assessment, loan-default classification, financial forecasting, and portfolio optimization. Hybrid approaches that combine Transformers with XGBoost, SHAP, graph learning, recurrent models, and other

analytical techniques provide additional opportunities for improving both predictive capability and interpretability.

The principal challenges involve data requirements, computational complexity, class imbalance, temporal validation, distributional shifts, and explainability. Addressing these issues is essential before Transformer-based financial-risk models can become reliable components of regulated financial decision systems.

Future research should focus on larger temporally structured datasets, multimodal financial representation learning, robust out-of-time validation, efficient Transformer architectures, explainable temporal attention, privacy-preserving learning, and integration with portfolio and credit-risk decision systems. These developments can help transform Transformer models from experimental predictive architectures into practical financial-risk analytics systems.

REFERENCES

1. Abbas, G., Ying, Z., & Iqbal, M. (2026). *From binary scores to risk tiers: An interpretable hybrid stacking model for multi-class loan default prediction*. *Systems*, 14(1), 78.
2. Akdoğan, Y. E. (2026). *The power of words for credit risk classification: Sentiment analysis as an alternative credit rating tool*. *Journal of Applied & Theoretical Social Sciences*, 8(1), 189–202.
3. Barvaliya, K. (2026). *Quantifying representational complexity in transformer models via residual stream spectral entropy*. *International Journal of Computer Applications*, 187(126), 9–25. <https://doi.org/10.5120/ijca54dc2360b436>
4. Caro-González, C., Jato-Espino, D., & Cardinale, Y. (2026). *Financial risk prediction models integrating environmental, social, and governance factors: A systematic review*. *International Journal of Financial Studies*, 14(2), 31.
5. Ding, D., Li, Y., Neo, P. L., Wang, Z., & Xia, C. (2026). *Dynamic multi-criteria portfolio selection integrating Transformer-based financial forecasting with peer-prediction trees*. *Mathematics*, 14(13), 2287.
6. Feng, H., Wang, Y., Fang, R., Xie, A., & Wang, Y. (2025, October). *Federated risk discrimination with Siamese networks for financial transaction anomaly detection*. In *Proceedings of the 2025 2nd International Conference on Digital Economy and Computer Science* (pp. 231–236).
7. Jiang, B. (2026). *Explainable financial risk prediction models—A hybrid framework of XGBoost+Transformer based on SHAP and time attention mechanisms*. *Procedia Computer Science*, 281, 885–894. <https://doi.org/10.1016/j.procs.2026.04.284>
8. Jindal, S., Upadhyay, A., Srivastava, A., Arora, S., & Sharma, V. (2026). *Financial services and risk intelligence powered by large language models*. In *LLMs in practice* (pp. 345–364). Elsevier.
9. Konatham, M. R., Uddandaraao, D. P., Vadlamani, R. K., & Konatham, S. K. R. (2025). *Federated learning for credit risk assessment in distributed financial systems using BayesShield with homomorphic encryption*. In *2025 International Conference on*

Computing Technologies & Data Communication (ICCTDC) (pp. 1–6). IEEE.
<https://doi.org/10.1109/ICCTDC64446.2025.11158863>

10. Khoozani, Z. S., Seera, M., & Lim, C. P. (2026). *Towards transparent financial AI: A systematic review of graph learning and explainable methods for credit risk and fraud detection*. Artificial Intelligence Review.
11. Liu, J., Qu, T., Wang, H., Yu, J., & Fu, Y. (2026). *Research on corporate financial risk prediction based on Transformer models*. International Journal of Dynamical Systems and Differential Equations, 15(1), 37–70. <https://doi.org/10.1504/IJDSDE.2026.154409>
12. Mahajan, A. S., Yamsani, N., & Konatham, M. R. (2024). *Financial risk prediction through multivariate data analytics and ensemble learning*. Stochastic Modelling and Computational Sciences, 4(2).
13. Mienye, I. D., Esenogho, E., & Modisane, C. (2026). *Deep learning for credit risk prediction: A survey of methods, applications, and challenges*. Information, 17(4), 395. <https://doi.org/10.3390/info17040395>
14. Saputra, M. P. A., Sukono, Riaman, Kartiwa, A., Misiran, M., & Wahid, A. J. (2026). *Deep learning for credit risk prediction in fintech lending: A systematic literature review on model architectures, imbalanced data handling, and research agenda*. Journal of Risk and Financial Management, 19(7), 465.
15. Sharma, N., Qureshi, A. R. K., Tatawat, C., Kushwaha, R., Vaishnav, A., & Joshi, H. (2026). *Artificial intelligence-driven financial risk prediction and portfolio optimization: A machine learning approach*. Enterprise Development and Microfinance, 36(3s), 1–26.
16. Wang, Z., Che, W., Jiang, C., & Zhao, H. (2026). *Leveraging multiview data through discrete and regularized deep learning for dynamic financial risk prediction*. Information Systems Research.
17. Yang, Y., Lin, Y., Zhang, Y., Su, Z., Goh, C. C., Fang, T., et al. (2025). *Transforming credit risk analysis: A time-series-driven ResE-BiLSTM framework for post-loan default detection*. Information, 17(1), 5.