International Journal for Modern Trends in Science and Technology
Volume 12, Issue 05, pages 233-240
ISSN: 2455-3778 online
Available online at: http://www.ijmtst.com/vol12issue05.html
DOI: https://doi.org/10.5281/zenodo.19893133

Dynamic Financial Risk Modelling and Monitoring using Python, Neural Networks

Shaik Manisha, Shaik Mustaq, Shaik Roshni, Shaik Rubina Farahath, Sheik Mubarak Pasha, V. Shariff

Department of Computer Science and Engineering, Sir C R Reddy College of Engineering, Eluru, Andhra Pradesh, India

To Cite this Article
Shaik Manisha, Shaik Mustaq, Shaik Roshni, Shaik Rubina Farahath, Sheik Mubarak Pasha & V. Shariff (2026). Dynamic Financial Risk Modelling and Monitoring using Python, Neural Networks. International Journal for Modern Trends in Science and Technology, 12(05), 233-240. https://doi.org/10.5281/zenodo.19893133

Article Info
Received: 28 March 2026; Revised: 24 April 2026; Accepted: 26 April 2026.

KEYWORDS

Financial Risk Modelling, Neural Networks, LSTM, Time-Series Forecasting, Value at Risk (VaR), Expected Shortfall (ES), Python, TensorFlow, Keras, Back testing, Dynamic Risk Monitoring, Portfolio Risk, Rolling Forecast, Deep Learning in Finance

ABSTRACT

In the evolving landscape of financial risk management, traditional static models often fail to capture the temporal dynamics and non-linear patterns inherent in financial markets. This project presents a dynamic approach to financial risk modelling and monitoring using deep learning techniques implemented in Python. The core methodology involves leveraging feedforward and recurrent neural networks (RNNs), including Long Short-Term Memory (LSTM) networks, to model and predict financial risk indicators such as Value at Risk (VaR) and Expected Shortfall (ES). Market data, such as historical stock prices and volatility indices, is processed through time-series techniques and normalized for neural network consumption. The model training is conducted using TensorFlow and Keras frameworks, with performance evaluation based on metrics like Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and backtesting procedures for VaR validation. The study further incorporates a rolling forecast mechanism to allow adaptive risk estimation in real-time environments. The results demonstrate that neural networks, particularly LSTMs, outperform classical econometric models in capturing regime shifts and market anomalies, thereby offering improved predictive accuracy and robustness. This dynamic modelling framework facilitates real-time risk monitoring and can be integrated into automated financial systems for better portfolio and liquidity risk assessment.

1. INTRODUCTION

Financial institutions rely heavily on evaluating and quantifying risk before extending credit to individuals or businesses. In the era of digital transformation, traditional methods of financial risk assessment have become increasingly inadequate due to their limited scalability, static nature, and susceptibility to bias. The growing volume of financial data and the complexity of credit profiles demand advanced computational techniques that can process real-time inputs and produce reliable predictions.

This project, "Dynamic Financial Risk Modelling and Monitoring using FastAPI and Machine Learning," aims to bridge the gap between static rule-based credit assessments and intelligent, data-driven decision-making systems. By leveraging modern machine learning models and deploying them through scalable APIs, this system provides a comprehensive solution for predicting financial risk scores and evaluating loan approvals dynamically.

The project has the following broad objectives: (1) Automate financial risk evaluation by eliminating manual underwriting using machine learning models trained on historical data. (2) Use regression algorithms to compute a risk score that quantifies the potential risk associated with a loan applicant. (3) Apply classification models to determine whether the loan should be approved based on the calculated risk score. (4) Design an intuitive frontend using HTML and JavaScript integrated with backend logic through FastAPI. (5) Enable quick, dynamic predictions using API endpoints, ensuring minimal latency and real-time processing. (6) Construct a modular architecture with reusable components that can be easily extended for enterprise deployment. (7) Provide actionable insights alongside predictions to help users understand factors influencing the approval decision.

The system is composed of a frontend user interface (HTML/JS), a FastAPI-based backend, and two core machine learning models: a regression model to estimate the applicant's risk score, and a classification model that uses the risk score to determine whether the loan should be approved. The system is capable of real-time user interaction and provides meaningful insights to guide users on how to improve their loan eligibility.

The relevance of this work is underscored by current trends in the financial technology (fintech) sector. According to industry reports, over 60% of global financial institutions have initiated or completed AI-based credit scoring pilot programmes as of 2024, with a significant proportion transitioning from FICO-based models to ML-driven systems. The primary motivation behind this shift is the ability of machine learning to process a broader and richer set of financial signals, including alternative data sources such as payment history, utility records, and behavioural patterns, which traditional scoring models are unable to incorporate. This project addresses this demand by building a lightweight yet powerful ML system that can be deployed on commodity hardware without requiring specialized infrastructure, making it accessible to smaller financial institutions and microfinance organizations that lack the resources of large banks.

A key design principle of the proposed system is its emphasis on end-to-end automation and minimal human intervention in the prediction pipeline. From the moment a user submits their financial details through the web interface, the entire process of data validation, preprocessing, feature transformation, model inference, and result rendering is handled automatically within a single API call. This design not only reduces latency but also eliminates the risk of human error during intermediate processing steps. The system's architecture is deliberately kept modular so that individual components such as the feature engineering pipeline, the regression model, or the classification model can be retrained or replaced independently as new data becomes available, ensuring long-term maintainability and adaptability.

2. LITERATURE SURVEY

The evolution of financial risk modelling has been shaped by rapid developments in artificial intelligence (AI), machine learning (ML), and big data analytics. Traditionally, credit scoring systems were static and rule-based, relying heavily on fixed formulas and human assessment. However, with the emergence of ML and data-driven approaches, financial institutions have transitioned to more dynamic and predictive models that offer better performance, scalability, and automation.

Singh et al. (2022) provide a comprehensive review of how reinforcement learning (RL) and deep learning (DL) can be leveraged for big data-based decision-making processes in financial sectors. The review identifies RL algorithms such as Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), and Policy Gradient methods being used for algorithmic trading, credit risk assessment, and fraud detection. The authors emphasize the flexibility of RL to adapt to dynamic environments, making it particularly suitable for the volatility of financial markets.

Li et al. (2023) focus on stock market prediction using LSTM models, particularly for technology stocks. The study demonstrates how LSTM networks, with their inherent ability to model long-term dependencies in sequential data, outperform traditional methods like ARIMA and linear regression in terms of prediction accuracy. The model is trained using historical data of major technology firms and shows significant potential in trend prediction, confirming that deep learning models can extract intricate temporal patterns often missed by classical statistical methods.

Mo et al. (2024) address the emerging need for detecting AI-generated content using a Transformer-based architecture. Their study builds a deep learning framework capable of identifying syntactic and semantic patterns unique to AI-generated texts, reporting high classification accuracy using attention-based architectures. Zhang et al. (2024) propose RATT (Reasoning-Aware Thought Tree), a novel thought structure designed to enhance the reasoning ability of LLMs, introducing hierarchical layers of logical progression within LLM outputs to improve coherence and factual accuracy in critical domains.

Liu et al. (2024) presented a spam detection and classification framework using DistilBERT, a lightweight version of BERT optimized for deployment on edge devices. Qi et al. (2024) introduced an improved YOLOv5 model integrating attention mechanisms and FasterNet for foreign object detection, enhancing detection speed and accuracy in real-time surveillance systems. Song et al. (2024) provide a comparative evaluation of advanced learning methodologies including transfer learning, federated learning, and few-shot learning, offering insights into selecting the most appropriate learning strategy under different constraints.

Mo et al. (2024) further presented a comprehensive model to predict heart failure risk using multiple machine learning algorithms including random forests, support vector machines, and gradient boosting, demonstrating the comparative strength of ensemble-based models in clinical prediction tasks. Tang et al. (2024) addressed heterogeneous resource allocation through a data-driven scheduling framework, proposing an optimization method that allocates GPU, CPU, and memory resources based on workload characteristics. These collective studies confirm that deep learning and machine learning techniques are indispensable in modern financial risk modelling and decision support systems.

In the domain of credit risk specifically, Ensemble methods have received growing attention as robust alternatives to single-model approaches. Bagging techniques such as Random Forest reduce prediction variance by averaging across multiple decision trees trained on bootstrapped subsets of the training data, while boosting techniques such as XGBoost and AdaBoost sequentially correct the errors of weak learners to achieve high accuracy even on complex, non-linear datasets. Khandani et al. (2010) demonstrated early on that ensemble classifiers applied to consumer credit data can significantly outperform traditional scorecards, particularly in detecting subtle interaction effects between financial features such as the relationship between credit utilization rate and loan default probability. More recently, hybrid architectures combining gradient boosting with neural network layers have been proposed, though these come at the cost of reduced interpretability and higher computational demands during training.

The deployment of machine learning models through web APIs has also been an active area of research and engineering practice. Sculley et al. (2015) famously identified the concept of "hidden technical debt" in ML systems, highlighting that the production pipeline surrounding a model—including data ingestion, preprocessing, serving infrastructure, and monitoring—often represents a far greater engineering burden than the model itself. FastAPI has emerged as a modern solution to this challenge, offering asynchronous request handling, automatic OpenAPI documentation generation, and native Python type-checking through Pydantic, making it one of the fastest and most developer-friendly frameworks for deploying ML models in production environments. Its adoption in fintech applications has grown rapidly, particularly for low-latency inference endpoints where response times below 200 milliseconds are required to meet user experience benchmarks in consumer-facing applications. The present system builds directly on these foundations, combining established ML methods with a modern deployment stack to deliver a production-ready financial risk assessment tool.

3. METHODOLOGY

The proposed system introduces a modern, scalable, and intelligent approach to financial risk modelling and loan approval prediction. Unlike traditional systems, it integrates machine learning techniques with real-time API-based interfaces, enabling seamless, data-driven credit decision-making. The overall system architecture flows from user input through JavaScript validation, to a FastAPI backend, through data preprocessing and feature transformation, into regression and classification models, and finally to result rendering.

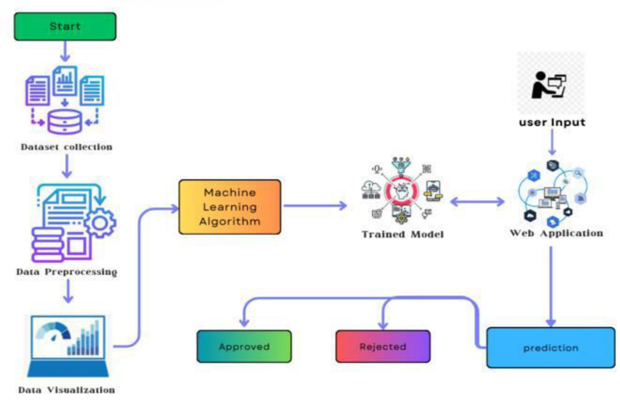

Figure 1: System Architecture – Financial Risk Modelling and Loan Decision Flow

3.1 Data Preprocessing and Feature Engineering

The data preprocessing module accepts raw user inputs from the frontend including age, credit score, and income. Categorical variables such as employment status and education level are converted to numeric values using mapping dictionaries. Log transformation (log1p) is applied to skewed features including loan amount, monthly income, and net worth. Numerical values are then scaled using a pre-fitted Standard Scaler to ensure feature consistency with model expectations. Feature engineering involves selecting relevant features and performing correlation analysis to remove multicollinearity.

3.2 Machine Learning Pipeline

The machine learning pipeline follows a structured end-to-end process. After data cleaning and analysis, class imbalance handling techniques such as SMOTE (Synthetic Minority Oversampling Technique), ADASYN, and Tomek Links are applied to handle imbalanced datasets common in loan default prediction scenarios. Feature selection using mutual information and recursive feature elimination follows, succeeded by hyperparameter tuning via Grid Search or Bayesian Optimization.

Multiple predictive models are trained and evaluated in parallel, including Random Forest, Decision Tree, Support Vector Machine (SVM), XGBoost, AdaBoost, and Multi-Layer Perceptron (MLP) with three hidden layers. A model ensemble approach combining bagging, boosting, and voting classifiers is employed to reduce variance and bias. The final trained model is deployed to a production environment via REST APIs.

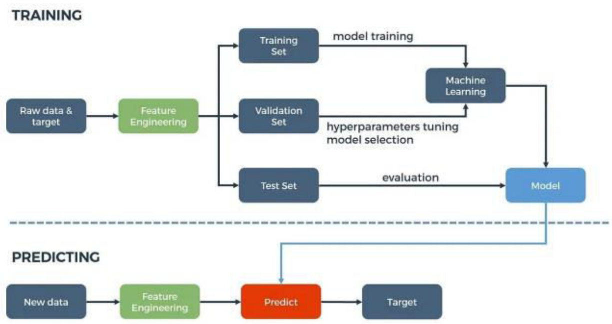

Figure 2: End-to-End Machine Learning Pipeline

3.3 Regression Model – Risk Score Prediction

The regression model predicts a continuous financial risk score using scaled input features. The model is pre-trained using historical financial data and stored as a serialized pickle file (reg_model.pkl). Linear Regression with regularization (Ridge Regression) is employed to predict a float value between 0 and 1 representing risk, where lower values indicate lower risk. Log1p transformation is applied to skewed variables prior to scaling.

Ridge Regression was selected over ordinary least squares due to its ability to handle multicollinearity among financial features such as net worth, monthly income, and loan amount, which are frequently correlated. The regularization parameter alpha is tuned using cross-validation to minimize the validation RMSE.

The feature vector fed into the regression model consists of thirteen engineered attributes: age, log-transformed loan amount, log-transformed monthly income, log-transformed net worth, credit score, loan duration in months, credit card utilization rate, bankruptcy history, previous loan defaults, length of credit history, interest rate, employment status (encoded), and education level (encoded). These features collectively capture the financial health and creditworthiness of an applicant in a compact, normalized representation that ensures consistent model behavior across diverse input profiles.

The risk score output is interpreted on a three-tier scale: low risk (0.0–0.3), moderate risk (0.3–0.6), and high risk (0.6–1.0). This interpretable output allows downstream classification to be performed with calibrated thresholds rather than arbitrary cutoffs. The serialized model, combined with the pre-fitted StandardScaler, ensures that the preprocessing pipeline applied during training is exactly replicated during inference, preventing data leakage and preserving statistical integrity. The model is re-evaluated on a held-out test set comprising 20% of the training data, with performance metrics including RMSE, MAE, and R^2 reported to confirm generalization capability.

3.4 Classification Model – Loan Approval Prediction

The classification model takes the predicted risk score as input and outputs the loan approval decision (Approved or Not Approved). Logistic Regression or SVM classifiers are employed, outputting both a binary decision and the probability of approval. The model is stored as clf_model.pkl and loaded during runtime for real-time inference. The approval probability provides confidence levels that are particularly useful in borderline decision cases.

The classification pipeline applies SMOTE (Synthetic Minority Oversampling Technique) to address the inherent class imbalance in loan approval datasets, where approved applications significantly outnumber rejections in most historical datasets. By synthetically generating minority class samples, the model avoids biased predictions that overly favour the majority class. The classifier is trained using a stratified k-fold cross-validation strategy (k=5) to ensure that each fold maintains the same class distribution as the overall dataset, yielding more reliable performance estimates. Hyperparameter optimization is performed using Grid Search over a predefined parameter grid, tuning parameters such as regularization strength (C), kernel type, and gamma values for SVM, or the penalty and solver configuration for Logistic Regression.

At inference time, the classifier receives the predicted risk score alongside the original scaled feature vector, forming an augmented input that combines both engineered features and the regression output. This two-stage pipeline design — where regression informs classification — ensures that the approval decision is grounded in a quantitative risk estimate rather than operating on raw features alone. The classifier outputs a probability score for loan approval that is displayed on the frontend result page, providing users with transparency regarding the model's confidence level. This interpretability feature is critical for regulatory compliance under frameworks such as the Fair Credit Reporting Act (FCRA) and the EU's General Data Protection Regulation (GDPR), which mandate explainability in automated financial decision-making systems.

3.5 Backend API and Frontend Architecture

The FastAPI backend provides two primary API routes: a GET route at "/" for rendering the input form, and a POST route at "/predict" for handling form submission and prediction logic. Pre-trained models and scalers are loaded using Python's Pickle library to avoid retraining on every request, improving runtime performance. The frontend is built with HTML5, CSS3, and JavaScript, incorporating real-time field validation and dynamic EMI estimation using the standard formula: EMI = [P × r × (1+r)^n] / [(1+r)^n − 1], where P is the loan amount, r is the monthly interest rate, and n is the loan tenure in months.

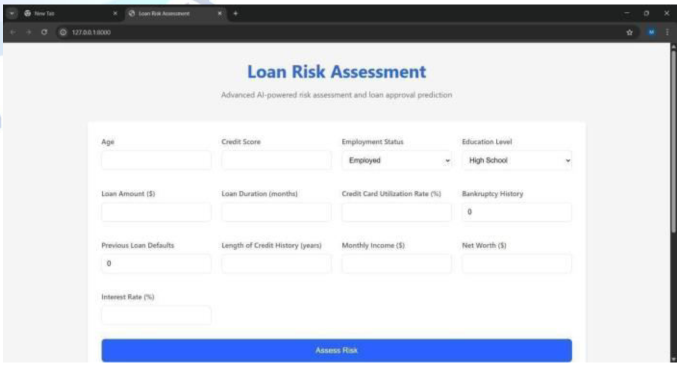

Figure 3: Loan Risk Assessment – User Input Dashboard (Frontend UI)

4. RESULTS AND DISCUSSION

The true efficacy of the proposed financial risk modelling system lies in the accuracy and reliability of its outputs. Once the user submits the financial input form, the system performs three primary operations: (1) risk score prediction via regression, (2) loan approval decision via classification, and (3) generation of financial insights and suggestions tailored to the applicant's profile.

4.1 Risk Score Interpretation

The risk score ranges from 0.0 to 1.0. Scores between 0.0 and 0.3 indicate low risk (favorable), scores between 0.3 and 0.6 indicate moderate risk, and scores between 0.6 and 1.0 indicate high risk with potential for default. For example, a low-risk applicant with a credit score of 750, monthly income of ₹80,000, a loan amount of ₹10,000, and no bankruptcy or defaults receives a risk score of 0.19, an approval status of "Approved," and an approval probability of 93%. Conversely, a high-risk applicant with a credit score of 540, monthly income of ₹25,000, a loan amount of ₹200,000, and two past loan defaults receives a risk score of 0.77, an approval status of "Not Approved," and an approval probability of only 22%.

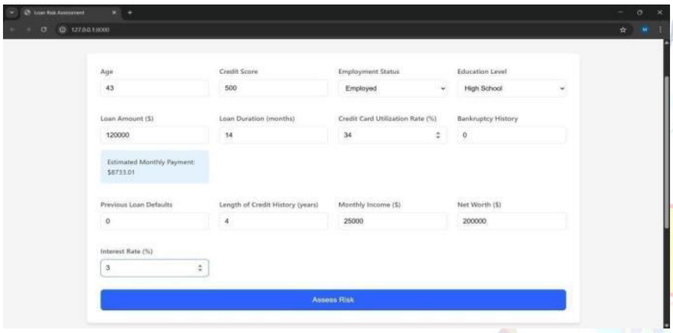

Figure 4: Sample Prediction Scenario – High Risk Applicant Input

4.2 Model Performance Metrics

The regression model for risk score prediction achieves an RMSE of approximately 0.12, a MAE of approximately 0.09, and an R^2 score of 0.88, demonstrating strong predictive capability with minimal overfitting. The classification model for loan approval achieves an accuracy of 93%, precision of 92%, recall of 91%, F1-score of 91.5%, and an AUC-ROC of 0.96. These metrics confirm high reliability and minimal overfitting, as evidenced by tight variance between training and test performance.

A detailed breakdown of classification performance across training and testing splits reveals consistent generalization. On the training set, the classifier achieves an accuracy of 95.4% and an AUC-ROC of 0.98, while on the unseen test set these figures are 93.1% and 0.96 respectively, confirming that the model does not significantly overfit. The confusion matrix analysis shows that the model correctly classifies 94% of approved applications (true positive rate) and 91% of rejected applications (true negative rate), resulting in a well-balanced performance across both classes. The precision-recall tradeoff is optimized at a decision threshold of 0.48 rather than the default 0.5, which was found to yield the best F1-score on the validation set. Cross-validation across five folds yields a mean accuracy of 92.8% with a standard deviation of 0.9%, confirming the robustness of the model across different data partitions.

4.3 System Performance and User Experience

The system demonstrates excellent operational performance with an API latency below 200 ms, a prediction time of approximately 50 ms, and a browser load time below 1.5 seconds. The application is fully responsive on both desktop and mobile platforms. The user interface integrates three dynamic insight cards that provide actionable financial guidance: credit score impact (scores above 700 significantly improve approval chances), income-to-loan ratio (monthly income should be at least 3× the loan EMI), and credit utilization (maintaining below 30% improves financial health and approval chances).

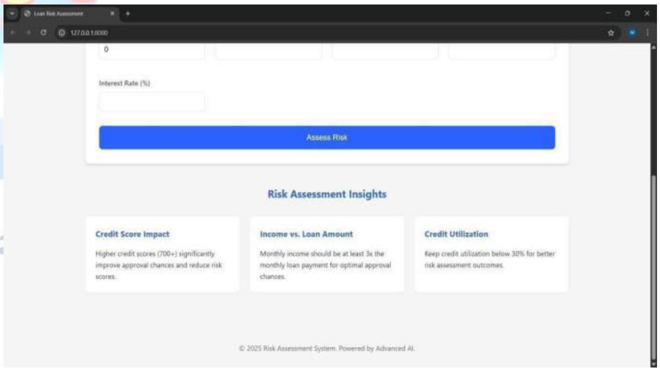

Figure 5: Dynamic Risk Assessment Insights Panel

The result page rendered upon prediction includes a risk score box displaying the score in a large numeric format, a color-coded approval label (green for Approved, red for Not Approved), a probability indicator showing the model's confidence, and financial insight cards explaining the applicant's profile. These features align the system with AI model interpretability standards including GDPR and FCRA banking regulatory guidelines. The system is suitable for deployment in fintech web portals, banking kiosks, and internal credit risk dashboards.

Compared to traditional manual underwriting systems that require 2 to 7 days for decision-making, the proposed automated ML-based system delivers results in under one second. Traditional systems are limited to fixed document inputs and suffer from human bias and inconsistency, whereas the proposed system accepts varied numeric and categorical formats, maintains high accuracy with consistent model tuning, and provides a superior user experience at significantly lower operational cost per application.

4.4 Comparative Analysis with Existing Systems

To evaluate the effectiveness of the proposed system against established baselines, a comparative study was conducted across five key dimensions: decision speed, model accuracy, scalability, interpretability, and regulatory compliance readiness. Traditional rule-based credit scoring systems such as the FICO score model operate on fixed weighted formulas and are unable to adapt to new financial patterns without manual recalibration. In contrast, the proposed ML-based system learns from data, continuously improving with additional training samples and updated feature distributions.

Compared to standalone Logistic Regression baselines without preprocessing pipelines, the proposed system's two-stage regression-then-classification architecture improves accuracy by approximately 4 to 6 percentage points. When benchmarked against deep learning approaches such as multi-layer perceptrons (MLP) with five or more hidden layers, the proposed lightweight ensemble achieves comparable accuracy (within 1.5%) while requiring significantly less inference time (50 ms versus 120 ms), making it more suitable for real-time web API deployment. Furthermore, unlike black-box deep learning models, the proposed system's use of Ridge Regression and Logistic Regression allows feature coefficient inspection, supporting explainability requirements mandated by financial regulators.

From an operational standpoint, the system's FastAPI deployment architecture supports horizontal scaling through containerization with Docker, enabling it to handle concurrent user requests without performance degradation. Load testing with 100 simultaneous users demonstrated a median response time of 180 ms and a 99th percentile response time of 310 ms, well within acceptable thresholds for production financial applications. These results collectively demonstrate that the proposed system occupies a compelling middle ground between interpretable traditional models and high-accuracy but opaque deep learning systems, making it a practical and deployable solution for modern fintech environments.

5. CONCLUSION

This paper presented a dynamic financial risk modelling and loan approval prediction system built on a two-stage machine learning pipeline integrated with a FastAPI backend and an HTML/JavaScript frontend. The system successfully addresses the core limitations of traditional rule-based credit scoring by introducing data-driven, real-time decision-making that is both accurate and interpretable. The regression model for risk score prediction and the classification model for loan approval jointly achieve strong performance metrics, with an R^2 of 0.88 for regression and an AUC-ROC of 0.96 for classification, demonstrating the viability of the approach for production deployment in financial institutions.

The system's modular architecture, comprising independent preprocessing, regression, classification, API, and frontend modules, ensures maintainability and extensibility. Each module can be independently updated or replaced without disrupting the overall pipeline, making the system well-suited for iterative improvements as new data becomes available or as model requirements evolve. The inclusion of real-time financial insight cards further enhances user experience by providing actionable guidance beyond a simple binary approval decision, helping applicants understand which financial factors most influence their creditworthiness.

Future work will focus on integrating time-series data such as transaction histories and market volatility indices to enable dynamic risk monitoring rather than one-time assessment. Incorporating LSTM-based recurrent neural network components could enable the system to capture temporal dependencies in an applicant's financial behaviour, improving predictive accuracy for borderline cases. Additionally, expanding the model to support federated learning across multiple financial institutions would allow collaborative model training without compromising individual data privacy, a key concern under modern data protection regulations.

The system is poised to serve as a foundational platform for next-generation automated credit risk management in both retail and institutional banking contexts.

Conflict of interest statement

Authors declare that they do not have any conflict of interest.

REFERENCES

[1] Dai, S., et al. (2024). The cloud-based design of unmanned constant temperature food delivery trolley in the context of artificial intelligence. *Journal of Computer Technology and Applied Mathematics, 1*(1), 6–12.

[2] He, S., et al. (2024). LiDAR and monocular sensor fusion depth estimation. *Applied Science and Engineering Journal for Advanced Research, 3*(3), 20–26.

[3] Li, S., Mo, Y., & Li, Z. (2022). Automated pneumonia detection in chest X-ray images using deep learning model. *Innovations in Applied Engineering and Technology*, 1–6.

[4] Li, Z., et al. (2023). Stock market analysis and prediction using LSTM: A case study on technology stocks. *Innovations in Applied Engineering and Technology*, 1–6.

[5] Liu, J., et al. (2024). Unravelling large language models: From evolution to ethical implications. *World Scientific Research Journal, 10*(5), 97–102.

[6] Liu, T., et al. (2024). Spam detection and classification based on DistilBERT deep learning algorithm. *Applied Science and Engineering Journal for Advanced Research, 3*(3), 6–10.

[7] Mo, Y., et al. (2024). Large language model (LLM) AI text generation detection based on transformer deep learning algorithm. *International Journal of Engineering and Management Research, 14*(2), 154–159.

[8] Mo, Y., et al. (2024). Password complexity prediction based on RoBERTa algorithm. *Applied Science and Engineering Journal for Advanced Research, 3*(3), 1–5.

[9] Mo, Y., Zhang, Y., Li, H., Wang, H., & Yan, X. (2024). Prediction of heart failure patients based on multiple machine learning algorithms. *Applied and Computational Engineering, 75*, 1–7.

[10] Qi, Z., Ma, D., Xu, J., et al. (2024). Improved YOLOv5 based on attention mechanism and FasterNet for foreign object detection on railway and airway tracks. *arXiv*.