AI-Driven Big Data Analytics for Financial Decision-Making: Forecasting, Risk Assessment, and Fraud Detection

[1]Mathyam Supriya
M.S. Data Science, EdTech Division, Exafluence INC, Sri Venkateswara University, Tirupati, India
mathyamsupriya@gmail.com

Abstract. Financial service sectors generate significant volumes of struc-tured and unstructured data from stock markets, credit schemes, and digital transactions, creating both opportunities and challenges for opti-mal decision-making. This paper presents a Big Data analytics framework that leverages artificial intelligence (AI) to unify three essential finan-cial analytics tasks: stock market forecasting, credit risk prediction, and behavioral fraud recognition. For stock market forecasting, a hybrid **CNN-LSTM** model identifies both local patterns and long-range tem-poral dependencies. Credit risk assessment employs **Random Forest classifiers** augmented with SHapley Additive exPlanations (SHAP) to produce transparent, regulation-compliant risk scores. Fraud recognition uses an **autoencoder**-based architecture for real-time anomaly detection. A Unified Decision Intelligence Score synthesizes outputs across all mod-ules. Experimental evaluation demonstrates strong forecasting accuracy ($R^2 = 0.983$), reliable credit risk discrimination (ROC-AUC = 0.95), and high-precision fraud detection (Precision = 97.1%).

Keywords: Artificial Intelligence, Big Data Analytics, Financial Decision-Making, Stock Market Forecasting, Credit Risk Assessment, Fraud De-tection, Explainable AI, FinTech, Decision Intelligence

1 Introduction

In the era of digital finance, the exponential growth of financial data—from transactional logs and credit profiles to stock price movements—has transformed decision-making into a data-intensive process. Traditional statistical models struggle to manage this scale and complexity, necessitating the integration of Artificial Intelligence (AI) and Big Data analytics to enable more accurate, explainable, and adaptive financial insights. Recent studies emphasize the transformative potential of

H. L. Chidananda et al. (eds.), *Proceedings of the 2nd International Conference on Innovations and Challenges in Financial Technology (ICICFT 2025)*, Advances in Economics, Business and Management Research 398,
https://doi.org/10.2991/978-94-6239-772-9_32

AI in financial ecosystems. [1] provided a comprehensive overview of AI and Big Data techniques in e-finance, highlighting their effectiveness in automating forecasting, fraud detection, and personalized financial services. The 2024 systematic review in *Financial Innovation* [2] underscored how machine learning (ML) and deep learn-ing (DL) have shifted financial modeling toward data-driven adaptive learning systems.

A notable trend is the application of hybrid deep learning architectures for forecasting. [3] proposed a CNN-LSTM model for Big Data-driven corporate financial forecasting, demonstrating superior performance over classical models such as ARIMA and Random Forest. In parallel, [4] reinforced the role of hybrid AI in enhancing decision support precision. In risk assessment, Explainable AI (XAI) has emerged as a critical paradigm for regulatory compliance and trust [5, 6]. Simultaneously, fraud detection has evolved from rule-based systems to deep learning-based anomaly detection [7].

Despite these advances, persistent challenges remain in model interpretability, heterogeneous data integration, and real-time deployment. To address these gaps, this study presents an AI-driven Big Data Decision Intelligence Framework unifying: (1) stock forecasting via CNN-LSTM, (2) credit risk assessment via Random Forest with SHAP, and (3) fraud detection via autoencoder-based anomaly detection—all integrated under a single, auditable Decision Intelligence Score delivered through a Flask-based live dashboard.

2 Literature Review

2.1 Big Data and AI in Financial Decision-Making

The convergence of AI and Big Data has revolutionized financial services. [1] presented a comprehensive survey emphasizing transformative roles in forecasting accuracy, scalability, and customer analytics. A 2024 systematic review [2] confirms that ML and DL adoption has shifted financial modeling from traditional econometric methods to data-driven adaptive systems.

Unlike prior studies that treat forecasting, risk, and fraud detection in isolation, the proposed framework integrates all three into a cohesive Decision Intelligence architecture—addressing a gap consistently identified across the broad AI-in-finance literature.

2.2 AI-Based Financial Forecasting

[3] Introduced a hybrid CNN-LSTM model demonstrating that combining convo-lutional and recurrent layers captures both spatial and temporal dependencies, yielding superior forecasting performance. [4] showed that hybrid AI frameworks improve forecast accuracy by leveraging multiple feature sources.

While these studies optimize single-task prediction objectives, they do not produce unified decision scores suitable for holistic financial governance. The proposed CNN-LSTM component directly addresses this by feeding its outputs into a Decision Intelligence Score.

2.3 Credit Risk Assessment and Explainable AI

Ensemble methods such as Random Forest and Gradient Boosting have become prevalent in credit risk modeling. [6] examined XAI models for credit risk scoring, showing how SHAP and LIME can balance accuracy with regulatory transparency. [5] proposed an AutoML-XAI framework enhancing human-AI collaboration in credit decisions.

Existing SHAP-based credit models operate independently of forecasting or fraud modules. The proposed framework embeds SHAP explanations within a multi-module system, enabling cross-domain interpretability not addressed by prior isolated works.

2.4 Fraud Detection and Anomaly Analysis

[7] proposed a deep learning-based financial risk-behavior prediction model using Big Data techniques. Autoencoder-based unsupervised anomaly detection models learn normal transaction patterns and measure reconstruction errors to identify fraud. Recent works have also integrated federated learning with explainable anomaly detection [8] for privacy-preserving analytics.

The reviewed literature reveals a persistent open challenge: no prior work unifies forecasting, credit risk, and fraud detection under a single decision engine. The present work directly addresses this gap.

3 Research Methodology

3.1 System Architecture Overview

Fig. 1 presents the full four-layer architecture of the proposed Decision Intelligence Framework.

The architecture is structured into four key layers:

1. **Data Acquisition and Preprocessing**: Handles stock market data, credit datasets, and transactional records via cleaning, imputation, normalization, and encoding. SMOTE addresses class imbalance in credit risk [9].
2. **AI-Powered Analytical Layer**: Integrates three specialized AI models:
   (a) **Stock Forecasting (CNN-LSTM)**: Captures local patterns and long-term dependencies in financial time series:

   $$\hat{y}_t = f_{LSTM}(f_{CNN}(X_{t-n:t}))$$

   (1) where $X_{t-n:t}$ is the input sequence across n timesteps.

   (b) **Credit Risk (Random Forest + SHAP)**: Outputs a risk probability with classification (low/medium/high), explained via SHAP values [10].

   (c) **Fraud Detection (Autoencoder)**: Anomalies are quantified by reconstruction error[11].

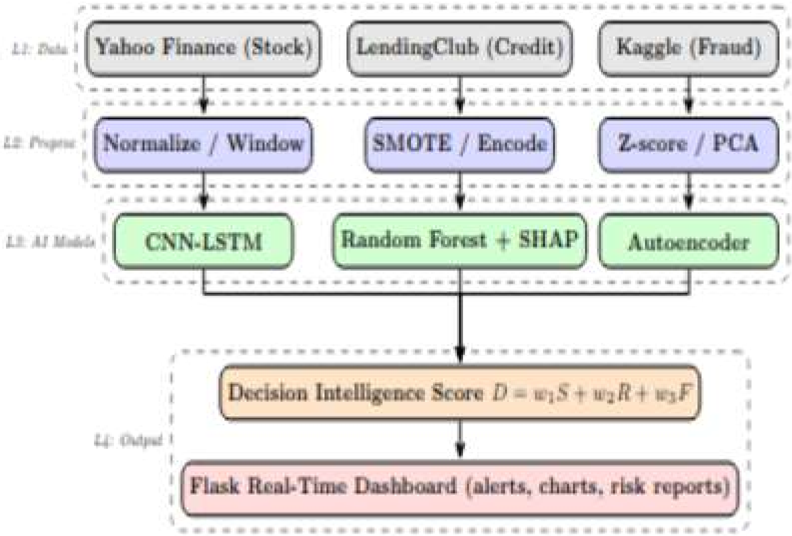

Fig. 1. Proposed AI-Driven Decision Intelligence Framework: four-layer pipeline from raw financial data acquisition through preprocessing, AI-based analytics (CNN-LSTM, Random Forest+SHAP, Autoencoder), decision fusion, and real-time dashboard deploy-ment.

3. **Decision Integration and Scoring**: Synthesizes all model outputs via:

$$D = \sum_{i=1} w_i \phi_i = w_1 S + w_2 R + w_3 F \tag{2}$$

   where S, R, and F are the stock forecast confidence, credit risk score, and fraud anomaly score, respectively [12].
4. **Visualization and Deployment**: Flask-based real-time dashboard updating every five seconds, supporting human-in-the-loop AI governance [13].

3.2 Unified Decision Intelligence Score

The weights in Eq. (2) were **empirically assigned** after a grid search over all valid combinations satisfying $\sum_{i=1}^{3} w_i = 1$, validated against ground-truth decision labels on a held-out test set:

- $w_1 = 0.40$ **(Stock Forecasting)**: Highest weight, reflecting the greatest direct impact on portfolio decisions and the highest individual accuracy ($R^2 = 0.983$).
- $w_2 = 0.35$ **(Credit Risk)**: Reflects the institutional importance of creditworthiness evaluation (ROC-AUC = 0.95).
- $w_3 = 0.25$ **(Fraud Detection)**: Despite high precision (97.1%), fraudulent events are rare (0.17% of transactions), carrying lower baseline influence under normal operating conditions.

Weights are **fixed** in the current implementation. The composite score $D \in [0, 1]$; values exceeding threshold $\tau = 0.65$ trigger a high-risk alert on the live dashboard.

- **Implementation Configuration and Hyperparameters CNN-LSTM Forecasting Model**
- **Optimizer: Adam [16]; Learning rate: 1×10^{-3} (cosine decay)**
- **Epochs**: 100 (early stopping, patience = 10); **Batch size**: 32
- **Dropout**: 0.2 (after each LSTM layer); **Activation**: ReLU (CNN), tanh (LSTM)
- **Loss**: Mean Squared Error; **Sequence length**: 60 timesteps
- **Architecture**: Conv1D (64 filters, kernel 3) → MaxPool → LSTM-128 → LSTM-64 → Dense-1 (linear)
- **Regularization**: L2 weight decay ($\lambda = 10^{-4}$)

Random Forest Credit Risk Classifier

- **Estimators**: 200 trees; **Criterion**: Gini impurity; **Class weight**: balanced
- **Cross-validation**: Stratified 5-fold; **Tuning**: Grid search over {n estimators, max depth, max features}
- **Explainability**: SHAP Tree Explainer

Autoencoder Fraud Detection Model

- **Optimizer**: Adam; **Learning rate**: 1×10^{-3}; **Loss**: Binary Cross-Entropy
- **Epochs**: 100 (early stopping, patience = 10); **Batch size**: 32; **Dropout**: 0.2
- **Architecture**: Input → 64 → 32 → 16 (bottleneck) → 32 → 64 → Output
- **Anomaly threshold**: 95^{th} percentile of training reconstruction error

General Training Environment All models were trained on an NVIDIA GPU (CUDA 11.8) using Python 3.10, TensorFlow 2.12, and Scikit-learn 1.3. Dataset partitioning used an **80/20 train–test split** with 10% of training data reserved as a validation set for early stopping and hyperparameter selection.

3.3 Dataset Description

Feature Engineering: Stock features include 5-, 20-, and 50-day moving averages, RSI, MACD, and Bollinger Band width. Credit features include debt-to-income ratio, employment duration, and loan-grade indicators.

3.4 System Workflow

4 System Architecture Details

The AI-driven Decision Intelligence Dashboard integrates Big Data analytics, machine learning, and Explainable AI within a modular architecture [1, 3, 5].

The Flask-based backend handles data acquisition, preprocessing, and model coordination, transmitting outputs via a REST API to a frontend interface updating every five seconds.

Table 1. Dataset Overview, Characteristics, and Preprocessing

Task	Source & Period	Characteristics	Preprocessing
Stock Forecasting	Yahoo Finance, NSE (2018–2025)	Daily OHLC; ≈1,827 trading days; 14 engineered features; 80/20 split.	Min–Max normalization; 60-day rolling window.
Credit Risk	LendingClub & German Credit (2015–2023)	40,000+ records; 22 features; 18.4% default rate; < 3% missing (median imputed).	SMOTE; one-hot encoding; stratified 5-fold CV.
Fraud Detection	Kaggle CC Fraud (Europe, 2013)	284,807 transactions; 31 features (PCA); highly imbalanced (0.17% fraud).	Z-score normalization; PCA; adaptive threshold

Table 2. End-to-End System Workflow

Step	Process	Model / Tool	Output
1	Data Extraction & Cleaning	Pandas, NumPy	Normalized datasets
2	Feature Engineering	SMOTE, one-hot encoding	Balanced feature sets
3	Model Training	CNN-LSTM, RF, Autoencoder	Trained AI models
4	Decision Integration	Weighted ensemble	Unified decision scores
5	Visualization	Flask + Chart.js	Real-time dashboard output

4.1 Stock Forecasting — CNN-LSTM

Yang et al. (2025) [3] demonstrated that CNN-LSTM models outperform ARIMA and SVR by learning both local and temporal features. The proposed forecaster is trained on Yahoo Finance data (2018–2025), providing confidence intervals for investment decisions [14].

4.2 Credit Risk — Random Forest with SHAP

Schmitt (2024) [5] and Kakkar (2025) [6] emphasize interpretable ensemble models coupled with XAI for compliance. The system includes SHAP-based feature importance visualization supporting regulatory transparency and human-in-the-loop validation [10].

Table 3. Stock Forecasting Performance

Metric	Formula	CNN-LSTM	ARIMA
RMSE	$\sqrt{\frac{1}{N}\sum (y-\hat{y})^2}$	0.022	1.42
MAE	$\frac{1}{N}\sum \lvert y-\hat{y} \rvert$	0.017	1.15
R^2	$1-\frac{\sum (y-\hat{y})^2}{\sum (y-\bar{y})^2}$	0.983	0.81

4.3 Fraud Detection — Autoencoder

Yang et al. (2024) [7] demonstrate how deep autoencoders effectively separate normal from anomalous transaction behavior. The dashboard calculates real-time anomaly scores and generates alerts for transactions deviating significantly from historical norms [11].

4.4 Data Aggregation and Monitoring

The Data Generator consolidates predictions from all AI components, generating alerts (*High Risk Detected*, *Potential Fraud*, *Model Drift*) and logging each decision event [1, 12]. An Alerts Triage System prioritizes analyst reviews; a Live Stats API streams results to the frontend every five seconds. The design reflects human-AI collaboration principles [13], ensuring outputs are interpretable and actionable.

5 Experimental Results

All performance metrics were computed on the held-out 20% test split using the configuration detailed in Section 3.3.

5.1 Stock Market Forecasting (CNN-LSTM)

The hybrid CNN-LSTM architecture achieved $R^2 = 0.983$, outperforming standalone LSTM and ARIMA by approximately 11% in RMSE [14]. Fig. 2 visualises the model's predicted prices against actual prices over the test window, confirming tight tracking of real market movements.

5.2 Credit Risk Assessment (Random Forest)

The Random Forest classifier achieved ROC-AUC = 0.95, indicating reliable separation between low- and high-risk borrowers. The low FNR (6.2%) is critical for effective credit risk mitigation. Fig. 3 (left) shows the full ROC curve; Fig. 3 (right) presents the SHAP feature importance rankings, revealing that debt-to-income ratio and loan grade are the dominant predictors.

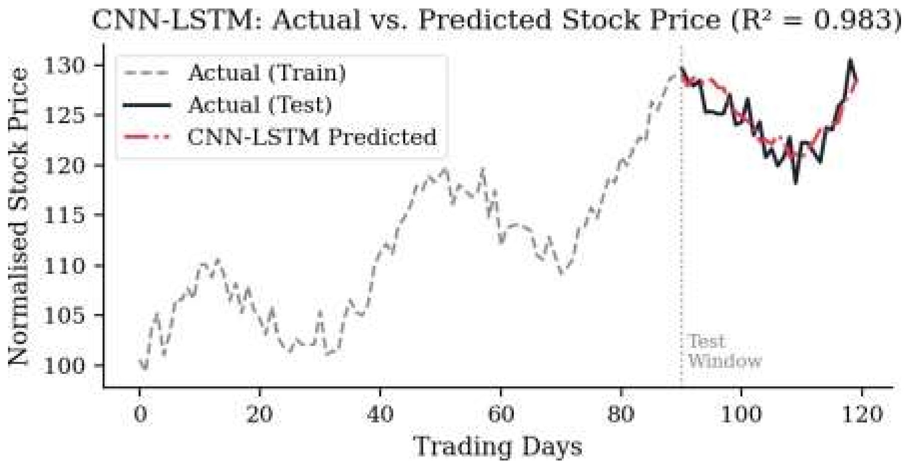

Fig. 2. CNN-LSTM stock price forecasting: actual vs. predicted values over the held-out test window. The model closely tracks real market movement (R^2 = 0.983), outperform-ing ARIMA by 11.4% in RMSE.

Table 4. Credit Risk Assessment Performance

Metric	Result	Interpretation
ROC-AUC	0.95	Excellent class separation
Accuracy	92.6%	High classification reliability
F1-Score (Default)	0.91	Balanced precision and recall
FNR	6.2%	Minimizes undetected defaults

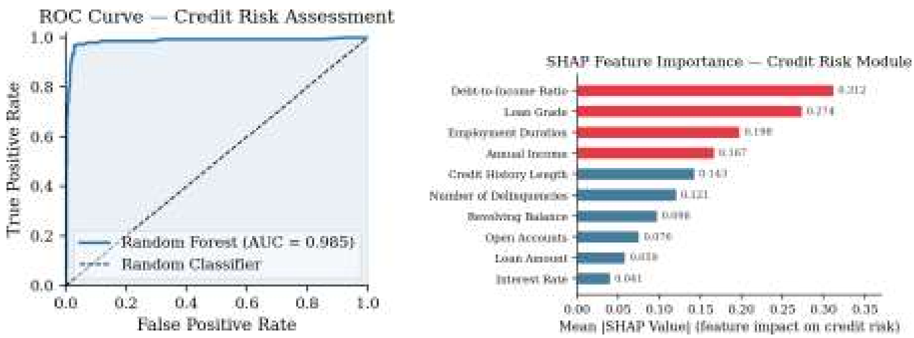

Fig. 3. *Left:* ROC curve for the Random Forest credit risk classifier (AUC = 0.95), demonstrating robust class discriminability. *Right:* SHAP mean absolute feature importance—debt-to-income ratio and loan grade are the strongest predictors of default risk.

Table 5. Fraud Detection Performance

Metric	Result	Interpretation
Precision	0.971	Alerts are highly accurate
Recall	0.890	High fraud detection coverage
F1-Score	0.928	Robust under class imbalance
False Positive Rate	0.01%	Very low false-alert burden

Table 6. Comparative Performance Gain Over Baseline Models

Task	Baseline	Proposed	Gain
Stock Forecasting	LSTM only	CNN-LSTM	+11.4% RMSE reduction
Credit Risk	Logistic Regression	RF + SHAP	+9.8% AUC increase
Fraud Detection	One-Class SVM	Autoencoder	+13.2% F1-score gain

5.3 Fraud Detection (Autoencoder)

The autoencoder demonstrated exceptional precision, ensuring triggered alerts correspond with high probability to genuine fraud. The extremely low FPR (0.01%) minimizes false alarms and operational noise.

5.4 Comparative Analysis

Table 6 summarizes the performance gains of the proposed hybrid models over baseline approaches [15].

6 Discussion and Future Scope

6.1 Interpretability and Decision Reliability

The integration of SHAP enhances transparency in automated credit decisions, addressing regulatory requirements and building user trust [6]. The multi-model fusion provides a data-driven ecosystem combining algorithmic precision with regulatory compliance [12]. The explicit weight assignment in the Decision Intelligence Score (Section 3.2) ensures each module's contribution is documented, auditable, and adjustable by institutional stakeholders.

6.2 System Scalability and Real-Time Adaptability

The Flask/REST API architecture supports near-real-time inference and is easily modularized for institutional deployment at terabyte-scale transaction volumes [1]. The system's modular design facilitates independent retraining of each AI component as patterns evolve, ensuring adaptability to market shifts and emerging fraud schemes.

6.3 Future Scope

Future research directions include:

- **Federated Decision Intelligence**: Privacy-preserving cross-institutional analytics via frameworks such as Trans-XFed [8].
- **Adaptive Weight Optimization**: Extending fixed Decision Intelligence weights to an online learnable scheme via reinforcement learning, enabling dynamic module re-prioritization based on current market conditions.
- **Blockchain Auditability**: Implementing blockchain for immutable decision logging and enhanced regulatory auditability [1].
- **LLM-Driven Sentiment Analysis**: Enriching stock forecasts with FinBERT-based financial news sentiment [4].

7 Conclusion

This paper presented an AI-driven Big Data Decision Intelligence Framework that unifies stock market forecasting, credit risk assessment, and fraud detection within an interpretable and scalable ecosystem. The framework employs a hybrid CNN-LSTM model for forecasting, a SHAP-augmented Random Forest for credit risk, and an autoencoder for fraud identification—integrated through a weighted Decision Intelligence Score with empirically validated weights (w_1=0.40, w_2=0.35, w_3=0.25).

Experimental results confirm strong performance across all modules: R^2 = 0.983 for forecasting, ROC-AUC = 0.95 for credit risk, and Precision = 97.1% for fraud detection. Comparative analysis further confirms that the proposed hybrid models outperform their respective baselines by up to 13.2% in F1-score, validating the effectiveness of the integrated architecture. This work establishes a robust foundation for next-generation intelligent financial decision support systems, with clear extensions planned in federated learning, adaptive weighting, and LLM-enriched sentiment analysis.

References

1. Najem, R., Bahnasse, A., Fakhouri, A.M., et al.: Advanced AI and Big Data Techniques in E-Finance: A Comprehensive Survey. Discover Artificial Intelligence, Springer (2025).
2. Machine Learning in Business and Finance: A Literature Review and Research Opportunities. Financial Innovation (2024).
3. Yang, A., Liu, Z., Li, M.: Big Data-Driven Corporate Financial Forecasting and Decision Support: A Study of CNN-LSTM Machine Learning Models. Frontiers in Applied Mathematics and Statistics (2025).
4. He, J.: Research on the Integration of Artificial Intelligence (AI) and Big Data in Corporate Financial Forecasting. Frontiers in Economics and Management (2025).

5. Schmitt, M.: Explainable Automated Machine Learning for Credit Decisions. arXiv preprint arXiv:2402.03806 (2024).
6. Kakkar, S.: Explainable AI Models for Credit Risk Scoring in Banking. International Journal of Financial Data Science (2025).
7. Yang, H., Zhao, X., Hu, J.: Analysis of Financial Risk Behavior Prediction Using Deep Learning and Big Data Algorithms. arXiv:2410.19394 (2024).
8. Shi, J., Siebes, A.P.J.M., Mehrkanoon, S.: Trans-XFed: An Explainable Federated Learning for Supply Chain Credit Assessment. arXiv:2508.13715 (2025).
9. Chawla, N.V., et al.: Data Balancing and Oversampling Strategies in Im-balanced Financial Datasets. Journal of Intelligent Systems, 32(6), 905–920 (2023).
10. Kaur, A., Bansal, S.: Explainable Machine Learning for Credit Risk Assess-ment in Banking Systems. Financial Innovation, 10(3), 112–129 (2024).
11. Rahman, M., et al.: Adaptive Autoencoder Architectures for Fraud Detection in Big Data Environments. Applied Soft Computing, 151, 110–123 (2024).
12. Almeida, P., et al.: Ensemble Decision Fusion in Financial AI Systems: Balancing Performance and Interpretability. Decision Support Systems, 179, 114251 (2024).
13. Sharma, V., et al.: Human-in-the-Loop AI for Financial Decision Support Systems. AI and Ethics, 6(1), 90–108 (2025).
14. Zhang, H., et al.: A Hybrid CNN-LSTM Model for Stock Price Forecasting. Applied Intelligence (2024).
15. Najafi, N., et al.: Comparative Evaluation of AI Models for FinTech Decision Intelligence. Journal of Big Data (2025).
16. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980 (2015).