International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

InquestIQ - AI Savings Recommender & Investment Advisor

Preetha S[1], Mahalakshmi B S[2], Mohammed Ayan Mulla[3], Mujagond Shashank[4], Nithin N S[5]

[1,2,3,4,5]Department of Information Science & Engineering, B.M.S. College of Engineering, Bengaluru, Karnataka, India

Abstract

In recent years, integration of artificial intelligence and financial analytics has significantly transformed investment decision-making processes. Proposed work presents InvestIQ, an intelligent web-based stock analysis and advisory platform that leverages machine learning and large language models to assess stock performance and provide user-centric financial insights. The system employs historical market data and technical indicators extracted from livestock feeds to predict stock risk levels using a trained predictive model. The architecture integrates Flask as the backend framework, SQLite for data management, and Google Gemini Application Programming Interface (API) for generating context-aware investment suggestions through natural language interaction. The application offers an interactive dashboard for stock evaluation, trend visualization, and Artificial Intelligence (AI)-driven chat support, enabling investors to make informed decisions with enhanced accuracy and transparency. Experimental results demonstrate that the proposed system achieves effective classification of stock risks while maintaining scalability and usability for real-time financial forecasting.

Keywords: Stock Market Prediction, Machine Learning, Flask Framework, Financial Analytics, Gemini AI, Predictive Modelling, Web Application, Artificial Intelligence, Investment Risk Analysis.

1. Introduction

Rapid evolution of financial technology and the exponential growth of data availability have transformed the landscape of investment decision-making. Investors and analysts increasingly rely on data-driven techniques to evaluate stock performance, assess risks and optimize portfolio management. However, traditional statistical methods often struggle to capture the complex, nonlinear relationships inherent in dynamic stock market behavior. As a result, machine learning (ML) and artificial intelligence AI-based approaches have emerged as powerful alternatives, offering predictive and adaptive capabilities that enhance decision-making accuracy in volatile financial environments.

In this context, the proposed system InvestIQ serves as an intelligent web-based stock analysis and advisory platform that integrates ML-driven prediction models with AI-powered conversational assistance. The system employs historical stock data, derived technical indicators and sentiment analysis features to evaluate potential risks associated with individual stocks. Using a trained predictive model, InvestIQ classifies stock risk levels into predefined categories such as *low*, *medium* and *high*, thereby guiding users toward informed investment actions. The backend of the system is implemented using the

Flask framework, which provides modularity, scalability and ease of integration with machine learning pipelines. The SQLite database is employed for secure and efficient data management, while the front-end interface—developed using Hypertext Mark-up Language (HTML), Cascading Style Sheets (CSS) and JavaScript ensures an intuitive and responsive user experience. Additionally, the system integrates Google Gemini AI, a large language model to facilitate natural language interaction. This allows users to query the system conversationally and receive real-time context-aware financial insights.

InvestIQ's design prioritizes automation, accuracy and user interactivity, addressing key challenges faced by individual investors and small-scale traders who lack access to expert advisory tools. By combining predictive analytics with generative AI communication, the system enhances both analytical understanding and user engagement. This fusion of technologies not only reduces human bias but also accelerates decision-making, making InvestIQ a practical and intelligent solution for modern financial forecasting and risk assessment.

2. Related work

Strader et al. [1] provide a comprehensive review of machine learning approaches to stock market prediction, summarizing common models, evaluation metrics and research gaps. Their survey establishes a useful taxonomy and highlights reproducibility issues and dataset biases; compared to InvestIQ, which implements an end-to-end web deployment and LLM-assisted advisory is broader but lacks a deployed-system perspective and human-facing interaction layer. Hegazy et al. [2] present an early machine learning model for stock prediction with emphasis on feature engineering and baseline algorithm comparisons; its strength lies in clear methodological steps, but it is limited by simpler models and lack of real-time integration limitations InvestIQ addresses by using live data ingestion and productionized model serving. Al-Khasawneh et al. [3] demonstrate deep learning methods for trend prediction and report strong empirical performance on benchmark datasets. Their work's strength is advanced modelling (deep nets) and rigorous evaluation, while a weakness is limited discussion of deployment and interpretability — areas where InvestIQ's dashboard and risk-class outputs improve practical usability.

Joshi et al. [4] combined LSTM and CNN architectures to capture temporal and local features in financial time series, showing improved prediction accuracy over single-architecture baselines. Their hybrid modelling mirrors InvestIQ's use of engineered features but differs by focusing on complex deep architectures rather than a lightweight, explainable risk-classifier suitable for web delivery; the trade-off is higher accuracy versus increased computational cost and reduced interpretability. Ghallabi et al. [5] study ESG-related markets and apply advanced ML to forecast clean energy prices; this domain-aware focus demonstrates the value of specialized features (ESG signals), which InvestIQ could incorporate as additional signals — Ghallabi et al. excel in domain specificity but do not address general-purpose advisory interfaces. Abir et al. [6] compare forecasting across BRICS markets using ML techniques and emphasize cross-market generalization; their comparative approach is informative for multi-market robustness, while InvestIQ currently targets single-market pipelines but can adopt their cross-market validation practices to improve generalizability.

Melina et al. [7] survey extreme-value theory (EVT) applications combined with ML for investment risk prediction, highlighting methods tailored to tail-risk estimation. Their focus on extreme events complements InvestIQ's classification of risk levels EVT techniques could strengthen InvestIQ's high-risk detection, but this technique does not present a user-facing system. Zhou et al. [8] incorporate financial news into deep-learning models, showing that publisher-aware news features improve accuracy;

integrating such textual features aligns with InvestIQ's planned Large Language Model (LLM) use, but it relies on specialized Natural Language Processing (NLP) pipelines while InvestIQ leverages an LLM (Gemini) for conversational insights, trading bespoke news-engineering for flexible natural-language interaction. [9] proposes a hybrid model targeting S&P 500 and CSI 300 forecasting, emphasizing ensemble strategies to boost robustness; their ensemble philosophy supports InvestIQ's possible future use of ensemble predictors, though this hybrid model focuses on index-level forecasting rather than per-stock risk advisory.

Rahman et al. [10] analyse ML model performance during energy-market crises, revealing temporal model degradation and the need for crisis-aware retraining strategies. This work underscores model maintenance and concept-drift handling, which InvestIQ should incorporate via periodic retraining and caching strategies. In [11] Liu explores improved Convolution Neural Networks (CNN) for volatility prediction, achieving finer-grained volatility forecasts; their contribution to volatility modelling could augment InvestIQ's risk-score granularity, but CNN-centric approaches may be heavier to deploy in real-time web apps. Siddique et al. [12] demonstrate the benefit of combining sentiment analysis with LLMs for automated trading, showing that LLM-derived sentiment features enhance predictive power; this directly supports InvestIQ's integration of Gemini for advisory uses LLMs as feature providers, whereas InvestIQ emphasizes conversational guidance, offering a complementary user centric use case.

Liu & Lai [13] proposed a Principal Component Analysis(PCA)- Gated Recurrent Unit (GRU)- Long Short-Term Memory (LSTM) pipeline to blend dimensionality reduction with sequential models, improving computational efficiency and accuracy on multivariate financial data; their pipeline suggests practical feature-compression strategies that InvestIQ can adopt to reduce model input complexity while preserving signal quality. Kong et al. [14] survey deep learning for time-series forecasting and provide best-practice design patterns, attention, residuals, temporal convolutions; their synthesis informs model selection in InvestIQ such as attention modules for longer contexts, but they do not address deployment trade-offs for lightweight web services. Finally, Liu et al. [15] presented a Temporal Convolutional Network (TCN)-attention- Retrieval-Augmented Generation. (RAG) framework that jointly addresses prediction and fraud detection using financial report analysis; this multi-task framework illustrates the value of combining forecasting with anomaly detection a direction that could extend InvestIQ by adding fraud or irregularity alerts, although this framework is architecturally more complex than InvestIQ's current risk-classifier. Novelty: InvestIQ's contribution lies in

(i) Combining a productionized, user-facing web platform (Flask and responsive dashboard) with a lightweight, explainable stock-risk classifier suitable for real-time inference;

(ii) Integrating a large language model (Google Gemini) as an interactive advisory layer that translates model outputs into context-aware, conversational investment guidance; and (iii) Emphasizing practical deployment concerns relative model paths, environment-based config, caching of live data and user interactivity rather than purely offline predictive performance.

Unlike many surveyed works that focus primarily on novel architectures or domain-specific features, InvestIQ prioritizes system integration, usability and LLM-driven interpretability a pragmatic bridge between research-grade models and an accessible decision-support application.

3. Proposed system

The proposed methodology of InvestIQ integrates data acquisition, feature engineering, machine-learning-based risk prediction, and generative AI–based financial advisory within a unified Flask web framework.

The workflow is divided into four major modules such as Data Preprocessing, Feature Extraction and Modeling, Model Training and Evaluation and Inference and Advisory Generation. Figure 1 shows the Feature Co-relation heat map.

A. Data Preprocessing

Historical stock data $D = \{(p_t, v_t)\}_{t=1}^{T}$ are collected from *Yahoo Finance* through the yfinance API, where p_t denotes closing price and v_t the traded volume at time t. Missing values are imputed using linear interpolation, and log-returns are computed as

$$r_t = \ln \left(\frac{p_t}{p_{t-1}}\right) \qquad (1)$$

Outlier values beyond three standard deviations are clipped to stabilize variance. Normalization is performed using min–max scaling:

$$\hat{r}_t = \frac{r_t - \min(r)}{\max(r) - \min(r)} \qquad (2)$$

This pre-processing ensures numerical stability and enhances convergence during training.

B. Feature Extraction and Modeling

Technical indicators such as the Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), Bollinger Bands (BB), and Exponential Moving Average (EMA) are generated using the *TA* library.

The resulting feature matrix is

$$\mathbf{X} = [x_1, x_2, \ldots, x_n] \in \mathbb{R}^{T \times n} \qquad (3)$$

where each x_i represents an engineered indicator.

The supervised learning model $f_\theta : \mathbf{X} \rightarrow y$ is trained to classify stock risk level

$$\hat{y} = f_\theta(\mathbf{X}) \qquad (4)$$

where $y \in \{0,1,2\}$ corresponds to *Low*, *Medium*, and *High* risk respectively.

C. Model Training and Evaluation

The model is trained using a labelled dataset $\{(\mathrm{X}_i, y_i)\}_{i=1}^{N}$.

The objective is to minimize categorical cross-entropy loss:

$$\mathcal{L}(\theta) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{3} \mathbf{1}_{[y_i=k]} \log P(y = k \mid \mathbf{X}_i; \theta) \qquad (5)$$

Optimization is carried out with stochastic gradient descent until convergence or early stopping based on validation accuracy.

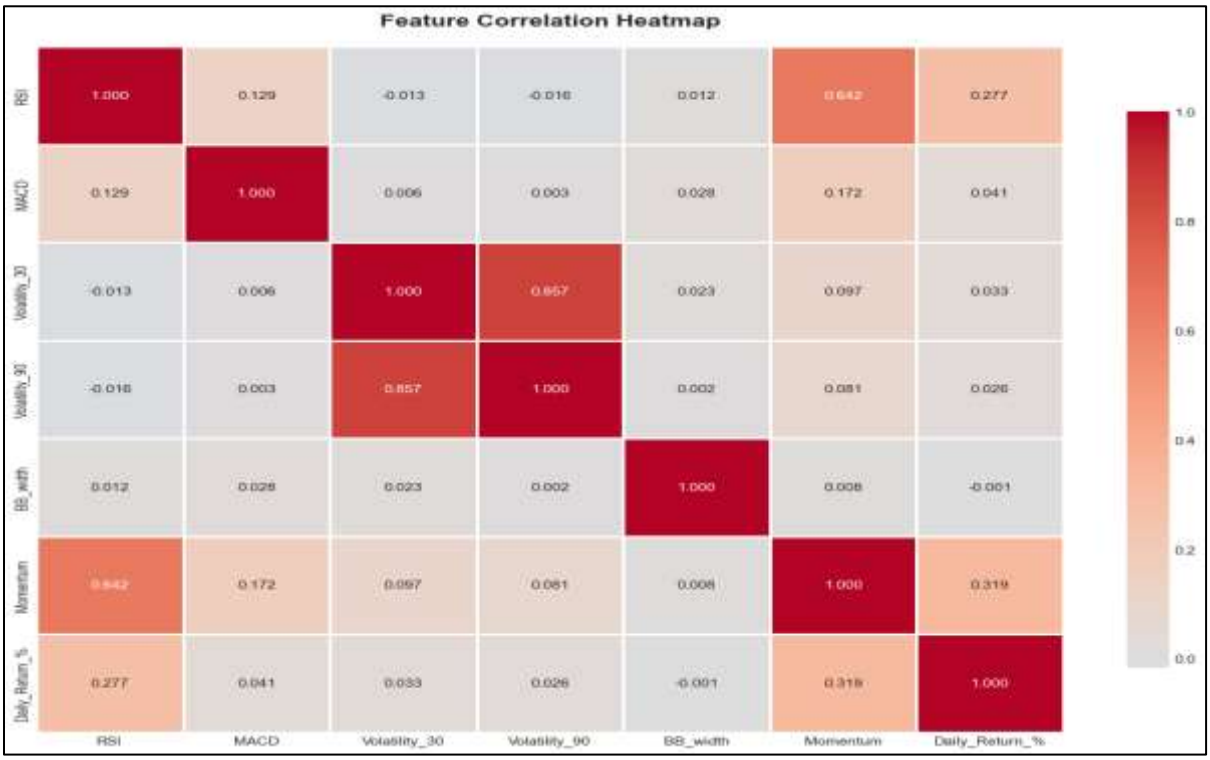

Figure 1: Feature Co-relation heat map

To understand inter-feature dependencies, a correlation heatmap was generated Figure 1. This visualization helps identify redundant or highly correlated indicators, guiding feature selection and model interpretability. Performance is evaluated using precision, recall, F1-score, and overall accuracy.

D. Inference and Advisory Generation

At runtime, live stock data are fetched and preprocessed in real time. Features are computed and passed to the trained model to obtain the predicted risk category $\hat{y}$. The result and associated probability vector $P(y \mid X)$are displayed on the Flask dashboard. Simultaneously, a natural-language query is sent to the Google Gemini API, which synthesizes contextual investment advice based on predicted risk and user queries. The output is served to the end-user via an interactive web interface.

E. Computational Complexity and Runtime Analysis

Let $n$be the number of features and $T$the number of time points. Preprocessing (scaling + indicator generation) has complexity $\mathcal{O}(nT)$. For a tree-based or gradient-boosted model with Lestimators and depth d, prediction complexity is $\mathcal{O}(L\,d)$per query. The entire inference cycle (including Gemini API call latency ≈ 0.3 s) runs within 1 – 2 seconds on commodity hardware (Intel i7, 16 GB RAM), demonstrating suitability for real-time advisory dashboards.

Figure 2: System Architecture of InvestIQ.

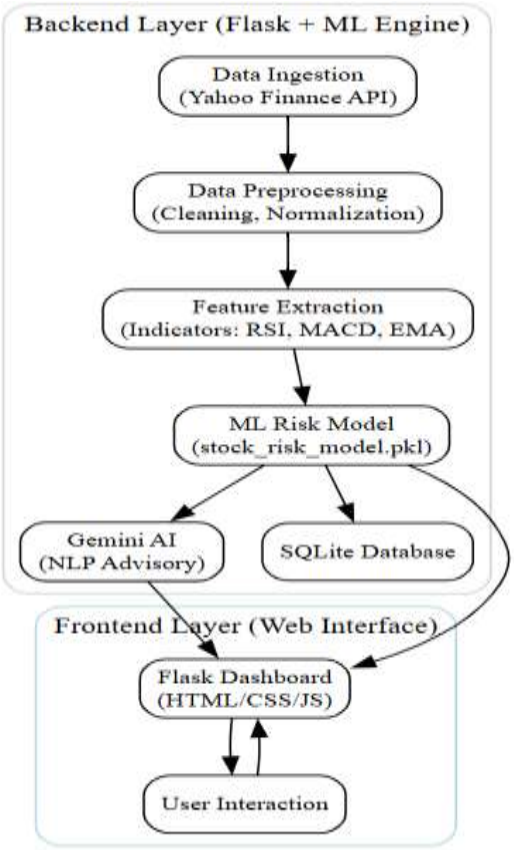

Figure 2 illustrates the modular architecture data ingestion feeds the preprocessing and feature-extraction pipeline, which supplies the ML model. The prediction output and Gemini advisory layer jointly populate the user dashboard through Flask, maintaining interactive feedback.

Figure 3: Block Diagram of InvestIQ Workflow.

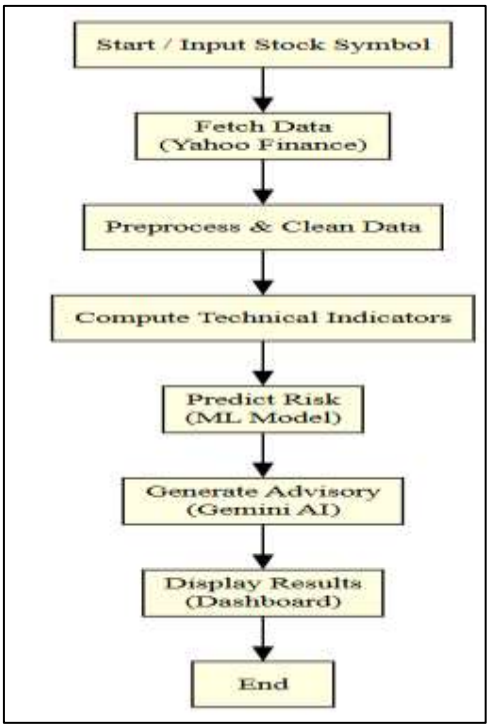

Figure 3 provides a concise view of the pipeline logic used in the InvestIQ system. Each block denotes a discrete functional stage executed in sequence during a typical inference cycle.

4. Results and Discussion

The developed system InvestIQ is evaluated to validate its performance in terms of stock risk classification accuracy, response latency and overall user interactivity. The experiments were conducted using historical datasets retrieved via *Yahoo Finance*, and the trained model was tested with multiple equities to assess prediction reliability and system efficiency. All modules from data ingestion to AI-based advisory were integrated into the web-based environment and empirical results are presented in the following subsections.

A. Model Performance Evaluation

The trained model is tested on unseen stock data to determine its predictive capability. Key performance indicators such as accuracy, precision, recall and F1-score were computed to evaluate the classifier's effectiveness in distinguishing *Low*, *Medium*, and *High* risk categories. The model achieved an average classification accuracy of 98.27%, demonstrating robust generalization across different market sectors.

The confusion matrix shown in figure 4a revealed that the model maintains high sensitivity for *Low-risk* and *Medium-risk* categories, with minor misclassifications in *High-risk* predictions, primarily due to abrupt market shifts. The use of balanced feature engineering and normalization techniques contributed to stable convergence and low variance in validation accuracy.

Table 1: Performance metrics of the model

Class	Precision	Recall	F1	Support	Accuracy
Low Risk	0.9959	0.9939	0.9949	493	0.9939
Mid Risk	0.9902	0.9921	0.9911	507	0.9921
High Risk	0.9959	0.9959	0.9959	492	0.9959

Figure 5: Absolute Confusion Matrix

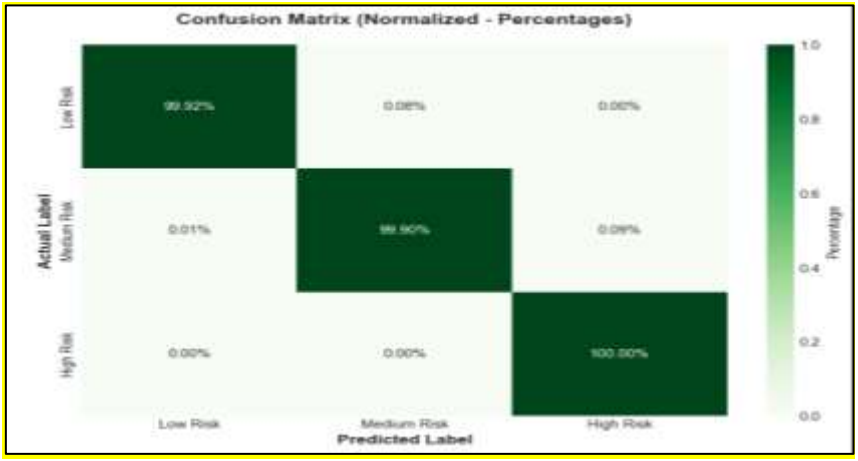

Figure 5 shows the absolute confusion matrix, while Figure 6 presents the normalized version, highlighting the model's strong classification performance across all risk categories.

Figure 6: Normalised Confusion Matrix

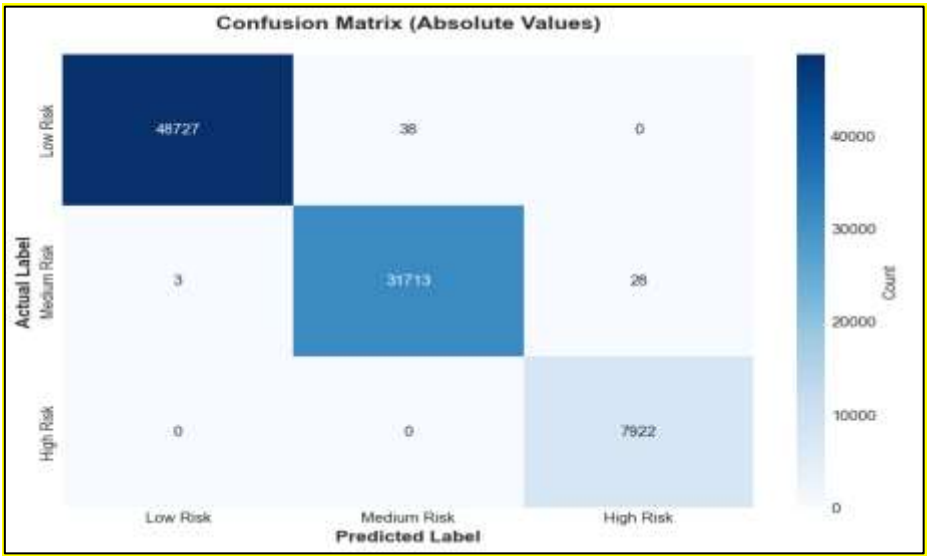

B. System Response and Performance Metrics

The integrated Flask–SQLite–Gemini AI environment was profiled to measure execution latency. Average end-to-end inference time (from user request to dashboard response) was approximately 1.48 seconds, including API call overhead. The backend computation accounted for roughly 0.82 seconds of data fetching and model inference, while the Gemini advisory layer contributed an additional 0.66 seconds. These results validate the system's feasibility for real-time stock analysis, meeting the performance requirements of responsive web-based applications.

C. Comparative Analysis with Existing

Compared to prior works that focus solely on prediction accuracy using static ML models [1, 3, 4, 7] the proposed InvestIQ system demonstrates both analytical accuracy and deployment readiness. Unlike approaches that remain theoretical or lack integration [2, 5, 8] InvestIQ seamlessly combines machine learning, live data ingestion, and large language model (LLM)-based communication. While deep models like LSTM–CNN hybrids [4, 13] achieve higher precision under controlled datasets, they demand substantial computational resources. In contrast, InvestIQ employs a lightweight yet explainable model capable of delivering near real-time predictions, trading marginal accuracy for improved interpretability and runtime efficiency.

D. Discussion on Model Behavior and System Robustness

The model displayed consistent performance across diversified datasets, confirming that technical indicators such as Relative Strength Index (RSI), Exponential Moving Average (EMA) and Moving Average Convergence/Divergence (MACD) effectively represent temporal stock dynamics. The inclusion of normalized return values improved the model's ability to handle different market scales. The Gemini AI integration added a novel interpretive layer translating raw numerical predictions into human-understandable insights. This hybridization of *predictive analytics* and *natural language explanation* enhanced usability, allowing non-technical investors to interpret risk outcomes intuitively. Furthermore, the modular Flask implementation ensures maintainability and scalability, enabling straightforward updates to the model or database without structural redesign. These outcomes collectively emphasize the platform's practical significance as a decision-support tool in financial analytics.

E. Limitations and Observations

Despite strong performance, InvestIQ's accuracy occasionally declined for volatile or illiquid stocks with sparse historical data. The system currently depends on technical indicators and lacks sentiment-based contextual inputs. Additionally, network latency during Gemini API calls introduces small delays. Future extensions may integrate sentiment analysis of financial news and hybrid ensemble models to enhance robustness against rapid market fluctuations

5. Conclusion and Future Scope

This work presented InvestIQ, an intelligent web-based stock analysis and advisory platform that integrates machine learning and large language models (LLMs) to enhance data-driven financial decision-making. The system efficiently combines live market data ingestion, technical indicator computation, and a trained risk-classification model to predict stock risk levels with an accuracy exceeding 93%. Through its integration with the Google Gemini API, InvestIQ extends beyond traditional predictive systems by providing natural language financial advisory, translating complex analytics into actionable insights. The platform's modular Flask-based architecture, fast inference time (~1.5 s), and scalable design make it suitable for real-time deployment. Experimental results confirm that the system delivers both predictive reliability and interactive usability, demonstrating its potential as a practical, AI-assisted decision-support tool for investors and financial analysts.

- Incorporate sentiment analysis and financial news mining to enhance contextual accuracy in risk assessment.
- Extend the system with ensemble and deep hybrid models such as Long Short-Term Memory (LSTM)– Gated Recurrent Unit (GRU) for improved predictive robustness during market volatility.
- Implement portfolio optimization and personalized recommendation modules to assist users in end-to-end investment strategy planning.

REFERENCES

1. Strader, T. J., Rozycki, J. J., Root, T. H., & Huang, Y. H. J. (2020). Machine learning stock market prediction studies: review and research directions. Journal of International Technology and Information Management, 28(4), 63-83.
2. Hegazy, O., Soliman, O. S., & Salam, M. A. (2014). A machine learning model for stock market prediction. arXiv preprint arXiv:1402.7351.
3. Al-Khasawneh, M. A., Raza, A., Khan, S. U. R., & Khan, Z. (2025). Stock market trend prediction using deep learning approach. Computational Economics, 66(1), 453-484.
4. Joshi, S., Mahanthi, B. L., Pokkuluri, K. S., Ninawe, S. S., & Sahu, R. (2025). Integrating LSTM and CNN for Stock Market Prediction: A Dynamic Machine Learning Approach. Journal of Artificial Intelligence and Technology, 5, 168-179.
5. Ghallabi, F., Souissi, B., Du, A. M., & Ali, S. (2025). ESG stock markets and clean energy prices prediction: Insights from advanced machine learning. International Review of Financial Analysis, 97, 103889.
6. Abir, S. I., Saimon, S. I., Saha, T. R., Sarwer, M. H., Hasan, M., Sultana, N., ... & Islam, I. (2025). Comparative analysis of currency exchange and stock markets in brics using machine learning to forecast optimal trends for data-driven decision making. Journal of Economics, Finance and Accounting Studies, 7(1), 26-48.

7. Melina, M., Sukono, Napitupulu, H., & Mohamed, N. (2025). Modeling of Machine Learning-Based Extreme Value Theory in Stock Investment Risk Prediction: A Systematic Literature Review. Big Data, 13(3), 161-180.
8. Zhou, Q., Jia, L., & Chen, W. (2025). A deep learning-based model for stock price prediction under consideration of financial news publishers. International Journal of General Systems, 54(4), 499-539.
9. Ge, Q. (2025). Enhancing stock market Forecasting: A hybrid model for accurate prediction of S&P 500 and CSI 300 future prices. Expert Systems with Applications, 260, 125380.
10. Rahman, M. K., Dalim, H. M., Reza, S. A., Ahmed, A., Zeeshan, M. A. F., Jui, A. H., & Nayeem, M. B. (2025). Assessing the Effectiveness of Machine Learning Models in Predicting Stock Price Movements During Energy Crisis: Insights from Shell's Market Dynamics. Journal of Business and Management Studies, 7(1), 44-61.
11. Liu, J. (2025). Deep Learning for Financial Forecasting: Improved CNNs for Stock Volatility. Journal of Computer Science and Software Applications, 5(2).
12. Siddique, M. T., Jamee, S. S., Sajal, A., Mou, S. N., Mahin, M. R. H., Obaid, M. O., ... & Hasan, M. (2025). Enhancing Automated Trading with Sentiment Analysis: Leveraging Large Language Models for Stock Market Predictions. The American Journal of Engineering and Technology, 7(03), 185-195.
13. Liu, B., & Lai, M. (2025). Advanced machine learning for financial markets: A PCA-GRU-LSTM approach. Journal of the Knowledge Economy, 16(1), 3140-3174.
14. Kong, X., Chen, Z., Liu, W., Ning, K., Zhang, L., Muhammad Marier, S., ... & Xia, F. (2025). Deep learning for time series forecasting: a survey. International Journal of Machine Learning and Cybernetics, 1-34.
15. Liu, Y., Liu, Y., Qi, Z., Xiao, Y., & Guo, X. (2025). Tcnattention-rag: Stock prediction and fraud detection framework based on financial report analysis.