IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

A Systematic Literature Review of Machine Learning Methodologies in Stock Market Forecasting

Priyanka Ramchandra Mali and Dr. K. G. Sawarkar
Dept of Electronics and Telecommunication Engineering
MCT's Rajiv Gandhi Institute of Technology, Mumbai, India
priyanka.mali@rgit.ac.in and kg.sawarkar@rgit.ac.in

Abstract: *Stock market prediction remains one of the most compelling yet elusive challenges in computational finance due to the inherent volatility, non-linearity, and chaotic nature of financial time-series data. In recent years, the paradigm has decisively shifted from traditional statistical econometrics toward data-driven artificial intelligence. This paper presents a systematic literature review exploring the evolution, methodologies, and architectural frameworks of machine learning (ML) and deep learning (DL) applications in stock forecasting over the past decade. We establish a comprehensive taxonomy to evaluate existing research across three critical pillars: data input dependencies (technical, fundamental, and alternative textual data), algorithmic variations (shallow models, ensemble tree-based methods, sequential networks like LSTM, and advanced hybrid optimization systems), and performance evaluation metrics. Furthermore, this review shifts away from generic algorithmic summaries to critically examine persistent gaps in the literature that hinder the real-world deployment of these predictive frameworks. Specifically, we examine the systemic issues of data snooping and model overfitting, the "black box" interpretability dilemma within deep architectures, and the frequent omission of practical transaction costs and slippage parameters in backtesting environments. By synthesizing these distinct methodologies and highlighting current structural limitations, this paper provides a grounded roadmap for future research, emphasizing the necessary transition toward explainable AI (XAI) and resource-efficient, real-time trading pipelines..*

Keywords: Market Prediction; Machine Learning; Deep Learning; Financial Time-Series; Sentiment Analysis; Taxonomy Survey; Quantitative Finance; Overfitting

INTRODUCTION

The prediction of stock market movements has long stood as one of the most intriguing, hotly debated, and financially significant domains in computational finance. For centuries, investors, economists, and mathematicians have sought a reliable edge to navigate the complex, non-linear, and chaotic nature of financial time-series data. The profound allure of quantitative trading lies in its immense financial scale. With trillions of dollars in capital circulating daily across global exchanges, even a marginal increase in directional prediction accuracy can yield substantial economic returns.

Historically, financial markets were analyzed through a rigid paradigm of traditional statistical economics and classic financial theories. However, the modern financial landscape has undergone a massive transformation driven by unprecedented computational power, the explosion of multimodal digital data, and the rapid maturation of artificial intelligence (AI) pipelines. As standard linear econometrics models struggle to keep pace with structural macroeconomic shifts and highly volatile market micro-structures, the domain has decisively pivoted toward data-driven machine learning (ML) and deep learning (DL) architectures.

DOI: 10.48175/IJARSCT-36722

A. Theoretical Underpinnings: EMH vs. Predictability

To contextualize the evolution of machine learning in computational finance, one must evaluate the foundational academic conflict that defines the field: the Efficient Market Hypothesis (EMH) vs. Market Predictability.

Formulated comprehensively by Eugene Fama in 1970, the EMH asserts that financial markets are fully "informationally efficient". According to this framework, asset prices instantaneously incorporate and reflect all available and relevant information. Consequently, stock price changes are driven purely by new, unanticipated information or "news" which, by definition, occurs randomly and unpredictably. Fama split the EMH into three distinct operational forms:

- Weak Form: Asserts that all past trading data (historical prices and volume) is entirely reflected in current asset prices. This establishes the theoretical baseline, which holds that classical technical analysis is incapable of consistently generating risk-adjusted excess returns (*alpha*).
- Semi-Strong Form: Argues that prices adjust instantaneously to all publicly available information, including corporate earnings releases, balance sheets, macroeconomic indicators, and public announcements. This invalidates standard fundamental analysis as a tool for beating the broader market.
- Strong Form: Suggests that even private or insider information is fully absorbed by the market, making consistent outperformance structurally impossible without pure luck.

Directly adjacent to the EMH is the Random Walk Theory, popularized by Burton Malkiel, which holds that stock market price fluctuations are statistically independent and follow a random walk. Under this traditional economic lens, trying to predict the next price tick or closing candle is as mathematically futile as predicting successive, independent coin flips.

Conversely, the behavioral finance revolution and modern data science challenge the notion of absolute efficiency. Proponents of market predictability highlight real-world market anomalies, extended asset misalignments, asset bubbles, and human cognitive biases (such as herd behavior, loss aversion, and overreaction to news). These human and systemic realities introduce structural patterns, repetitive cyclical fluctuations, and non-linear dependencies into financial time-series data.

Machine learning operates on this exact counter-premise: while the market might be highly efficient in a macro sense, it displays transient pockets of predictability—micro inefficiencies—that can be mapped, extracted, and leveraged by advanced pattern-recognition algorithms before the broader market self-corrects.

B. The Methodological Shift: From Statistical Modeling to AI

Before the arrival of AI, time-series forecasting relied heavily on parametric, univariate statistical tools. The absolute gold standards for decades were:

1. Autoregressive Integrated Moving Average (ARIMA) frameworks.
2. Generalized Autoregressive Conditional Heteroskedasticity (GARCH) systems, deployed primarily to model volatile financial clustering.

While these tools are mathematically elegant and highly interpretable, they carry severe operational assumptions. They assume that financial time-series processes are strictly stationary (or can be easily differenced into stationarity) and that the underlying relationships between variables are inherently linear.

Real-world financial ecosystems constantly break these assumptions. Markets are dynamic, multi-variable environments where asset prices are simultaneously influenced by domestic politics, global supply chain shocks, social media hype cycles, and institutional algorithmic order flows. Forcing a highly nonlinear, multidimensional chaotic system into a rigid linear ARIMA framework creates substantial structural errors, causing these models to struggle during high-volatility shifts or black-swan market corrections.

The modern era of data-driven AI bypasses these parametric restrictions entirely. Machine learning architectures are fundamentally non-parametric, meaning they do not require rigid, preconceived assumptions about the statistical distribution or linearity of the input data. Instead of forcing data to fit a strict mathematical equation, ML algorithms

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 6, June 2026 Impact Factor: 8.2

adaptively learn internal representations, mathematical mappings, and intricate high-dimensional cross-features directly from the training data.

This evolution has fundamentally unlocked the ability to parse unstructured alternative data. While statistical econometrics could only process structured data like prices and volumes, modern natural language processing (NLP) and deep neural networks can invest millions of unstructured text inputs—such as regulatory filings, real-time financial news headlines, and social media feeds—converting public human sentiment directly into quantifiable trading features.

C. Scope, Objectives, and Structural Taxonomy of This Review

Given the monumental surge of research papers exploring artificial intelligence in finance over the past decade, the current literature has become highly fragmented. Researchers routinely report widely varying results using disparate datasets, completely different evaluation protocols, and contrasting hyperparameter setups. The primary objective of this systematic literature review is to synthesize this massive body of work and offer an organized architectural taxonomy that categorizes the field into clear, digestible components.

Rather than providing a passive summary of past papers, this review critically analyzes existing methodologies through a rigorous multi-pillar framework:

- Data Engineering and Feature Dependencies: We dissect how historical papers source and process input features, mapping the clear transition from basic technical price indicators to fundamental corporate ratios and unstructured textual sentiment analysis.
- Algorithmic Taxonomies: We group predictive systems into specific algorithmic tiers—contrasting classical shallow architectures (such as Support Vector Machines and Logistic Regression) with ensemble tree methods (XGBoost, Random Forest) and deep sequential neural networks (LSTM, GRU, and financial Transformers).
- Performance and Validation Metrics: We contrast how different authors validate their models, distinguishing purely mathematical error trackers (RMSE, MAE) from direct classification and risk-adjusted financial metrics (Directional Accuracy, Sharpe Ratio, Max Drawdown).

Most importantly, this paper acts as a critical evaluation of persistent literature gaps. We explore the structural limitations that currently prevent theoretical laboratory models from succeeding in live trading environments. Specifically, we focus on the widespread prevalence of data snooping and look-ahead biases, the profound interpretability challenges of neural "black boxes," and the catastrophic omission of market micro-structure costs (transaction slippage, execution latencies, and liquidation fees) that frequently erase paper profits during live deployment.

Ultimately, this literature review is designed to serve as a comprehensive, grounded roadmap for both computational researchers and quantitative financial practitioners, outlining the critical steps needed to transition financial AI from static, overfitted historical backtests into robust, explainable, and deployment-ready forecasting systems.

II. RELATED WORK

Kumar et al. [1] proposed a supervised machine learning framework utilizing ensemble methods, specifically Random Forest and XGBoost, to predict the closing prices of major stock indices. The methodology relied on extracting hand-crafted technical indicators, such as the Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Bollinger Bands, combined with lagged price features. The results demonstrated that ensemble tree-based models significantly outperformed simple classifiers such as Naïve Bayes and Decision Trees, achieving lower Root Mean Squared Error (RMSE) on high-noise index data. However, the critical research gap is that the framework lacks real-time computational scalability, failing to account for high-frequency execution latencies, transaction fees, and market slippage during live portfolio simulation.

Raza et al. [2] presented a machine learning model optimized specifically for volatile emerging equity markets using historical pricing segments. Their methodology integrated standard technical overlays with K-Nearest Neighbors (KNN), Support Vector Machines (SVM), and Gradient Boosting Regressors. The results showed that fine-tuning

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 6, June 2026 Impact Factor: 8.2

model hyperparameters on long-term data series minimized the Mean Absolute Error (MAE) and effectively protected simulated baseline portfolios from abrupt capital losses. Despite this, a major research gap lies in the complete omission of macroeconomic variables and alternative textual data feeds, which restricts the model's reliability during structural policy-driven market updates.

Aritonang et al. [3] proposed a novel conceptual architecture that integrates classical economic asset-pricing theory directly with a multivariate Long Short-Term Memory (LSTM) network. The methodology focused on validating statistical arbitrage signals using economic constraints rather than relying purely on data-driven patterns. The results highlighted a remarkable increase in directional classification accuracy and superior risk-adjusted portfolio optimization compared to traditional unconstrained Markowitz models. The identified research gap centers on the lack of an automated online adaptive learning kernel, meaning the network must be manually retrained whenever underlying macroeconomic regimes switch.

Al-Alawi et al. [4] presented an empirical review and architectural comparison of Deep Learning models applied specifically to time-series forecasting. Their methodology evaluated standard recurrent topologies against Gated Recurrent Units (GRUs) and standard LSTMs using multi-market price inputs. The results confirmed that deep sequential models were structurally superior at capturing long-term non-linear temporal dependencies that traditional linear models failed to represent. However, the prominent research gap is the "black box" nature of network outputs, which offer zero interpretability or explainability for institutional finance applications.

Siala et al. [5] proposed a multimodal sentiment-driven framework that leverages advanced large language model (LLM) architectures for financial trend classification. The methodology evaluated FinBERT, RoBERTa, and DeBERTa to extract contextual text features from corporate announcements and real-time public financial streams. The results indicated that transformer-based sentiment embedding models generated highly precise directional buy/sell signals, outperforming older dictionary-based NLP techniques. The underlying research gap is that the framework experiences severe feature-space explosion when processing multiple concurrent textual sources, leading to high processing latency that impairs real-time algorithmic trading.

Bukhori et al. [6] presented an advanced Heterogeneous Graph Neural Network (HGNN) framework engineered to track cross-stock market linkages. The methodology modeled individual stocks as nodes and industry sectors as relational edges, pairing graph spatial embeddings with an underlying temporal attention layer. The results demonstrated a dramatic reduction in forecasting errors across highly volatile sectors by effectively extracting structural market-ripple trends. Nevertheless, the research gap lies in the model's high sensitivity to dynamic edge changes, meaning that sudden, unmapped corporate restructurings or mergers significantly degrade graph-propagation accuracy.

Che et al. [7] proposed a hybrid deep architecture combining a Convolutional Neural Network (CNN) with a Bidirectional LSTM (BiLSTM) for concurrent spatial-temporal feature mapping. The methodology used 1D-CNN layers to filter short-term technical price momentum, while the BiLSTM blocks tracked long-term directional context. The results revealed a substantially superior Directional Accuracy (DA) metric for benchmark blue-chip equities compared with standard standalone networks. The major research gap is the high computational cost and overparameterization, which make the model highly prone to overfitting on low-volume equity classes.

Giantsidi et al. [8] presented a comprehensive systematic analysis tracking 187 standalone machine learning studies in financial forecasting. Their methodology grouped literature by technical data inputs and optimization kernels to establish performance benchmarks. The results verified that while LSTMs remain the dominant implementation, hybrid models leveraging meta-heuristic optimization frameworks consistently yield higher predictive accuracy. The clear research gap identified is the severe data-snooping bias present in the existing academic literature, as most analyzed papers used static, clean backtests that ignore real-world out-of-sample performance.

Wang et al. [9] proposed an integrated model that pairs an autoregressive Transformer architecture with a self-attention mechanism for multi-step stock index forecasting. The methodology leveraged multi-head attention blocks to prioritize historical trading windows during high-volatility shifts selectively. The results proved that the Transformer architecture significantly improved multi-day forecast horizons compared to traditional RNN variants. However, a distinct research gap remains in its inability to handle unstructured macro-environmental data flows, rendering the model vulnerable to unexpected geopolitical event shocks.

Zhang et al. [10] presented an automated machine learning (AutoML) framework that optimizes hyperparameter generation for stock trend classification. The methodology used Bayesian optimization to dynamically select optimal activation layers and dropout rates for deep neural networks. The experimental results demonstrated stable outperformance and a drastic reduction in the number of human-engineered tuning cycles. The research gap centers on the platform's complete disregard for execution microstructures, such as variable exchange transaction fees and order book liquidity limits.

Liu et al. [11] proposed an alternative data fusion model that uses an ensemble of gradient-boosted trees integrated with public web-scraping pipelines. The methodology processed real-time search volume indices and forum discussions to construct consumer retail demand features. The results confirmed that alternative internet metrics provide a powerful short-term momentum proxy for consumer-facing stocks. Despite this, the identified research gap is the feature set's high vulnerability to data manipulation and structural noise, such as internet bot activity, which can distort sentiment metrics.

Sharma et al. [12] presented a neuro-fuzzy inference network optimized via Genetic Algorithms (GA) to forecast stock price turnarounds. The methodology used linguistic fuzzy rules to capture fuzzy market sentiment values, with GA refining membership thresholds. The results showed that the model successfully minimized the Mean Absolute Percentage Error (MAPE) during market consolidation phases. The core research gap is that the framework suffers from the "curse of dimensionality," rendering it computationally too heavy when scaled across thousands of equities simultaneously.

Patel et al. [13] proposed a stacking ensemble model combining Support Vector Regression, Extreme Learning Machines, and Random Forests. The methodology applied an optimization layer to weight individual base-learner predictions depending on the asset's rolling historical volatility. The empirical results demonstrated highly robust forecasting capabilities in both bull and bear market regimes. However, the primary research gap lies in the static nature of the stacking weights, which fail to adapt instantly during sudden, high-impact black swan events.

Khan et al. [14] presented a wavelet-denoised LSTM framework tailored for equity index time-series data. The methodology used the Discrete Wavelet Transform (DWT) to split the raw price data into low-frequency approximation signals and high-frequency noise details before network training. The results showed that filtering out high-frequency market noise dramatically boosted the model's out-of-sample directional prediction consistency. The research gap is that the choice of the wavelet decomposition level remains highly subjective and requires continuous manual calibration.

Lee et al. [15] proposed a reinforcement learning (RL) framework based on deep Q-networks (DQN) for automated stock trading execution. The methodology coupled a price-forecasting LSTM module with an RL agent that directly generates trading decisions (Buy/Sell/Hold) to optimize cumulative returns. The simulation results delivered superior Sharpe ratios compared to standard buy-and-hold benchmarks. The critical research gap is the assumption of instant trade execution, completely ignoring the real-world impact of market slippage and partial order fills.

Zhao et al. [16] presented an Explainable AI (XAI) methodology for stock prediction using Shapley Additive exPlanations (SHAP) paired with an XGBoost core. The methodology focused on calculating the exact marginal contribution of individual fundamental financial metrics to the model's final prediction. The results successfully unmasked the "black box" pipeline, proving that debt-to-equity ratios dominate long-term forecasting outputs. The research gap is that the computed SHAP values are strictly static and fail to model real-time, dynamic variations in feature importance during highly chaotic market crashes.

Das et al. [17] proposed a multitask deep learning network to predict stock prices and corporate earnings surprises concurrently. The methodology used shared internal representation layers to exploit statistical correlations between historical price momentum and fundamental accounting periods. The results demonstrated that multi-task architectures

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

consistently beat single-task networks across both prediction domains. However, a major research gap is the reliance on synchronous data publishing, which makes the system struggle with non-synchronous, messy real-world reporting schedules.

Kim et al. [18] presented a transfer learning pipeline utilizing pre-trained deep temporal networks across vastly different international exchanges. The methodology fine-tuned a base model trained on high-liquidity US markets and deployed it directly on emerging, low-liquidity target exchanges. The results confirmed a significant acceleration in training convergence speeds and a drop in data-scarcity errors. The research gap involves the total neglect of changing foreign exchange risks and distinct regulatory trading rules between source and target markets.

Ahmed et al. [19] proposed a real-time anomaly detection and stock price prediction pipeline driven by Isolation Forests and GRU networks. The methodology filtered out extreme, unrepresentative spoofing spikes in order books before passing the clean sequences to the prediction layer. The results yielded a remarkably robust prediction accuracy during high-frequency manipulation attempts. The identified research gap is that the anomaly detection layer occasionally misclassifies genuine, high-impact fundamental news breaks as market anomalies, wiping out valid trading signals.

Mishra et al. [20] presented a multi-agent framework that uses Large Language Models as specialized forecasting, risk-management, and style-allocation agents. The methodology designed a dynamic self-reflection layer where individual agents review historical trading failures to adaptively fine-tune stop-loss parameters. The results delivered a highly resilient numerical reasoning engine that safely tracks market dynamics. However, a prominent research gap lies in the significant operational latency and API token costs, rendering the multi-agent design entirely impractical for live, fast-paced, high-frequency trading pipelines.

TABLE 1: COMPARATIVE ANALYSIS OF LIVER CANCER DETECTION AND CLASSIFICATION

Ref	Lead Author	Primary Architecture	Data Modality / Source	Key Evaluation Metrics	Primary Research Focus / Gaps
[1]	Kumar	Random Forest / XGBoost	Technical Analysis Indicators	RMSE	Fails to scale for high-frequency latencies and transaction fees.
[2]	Raza	KNN, SVM, Gradient Boosting	Historical Price Bars (OHLC)	MAE, Portfolio Loss	Completely omits macroeconomic changes and textual alternative data.
[3]	Aritonang	Multivariate LSTM	Theoretical Economic Pricing	Directional Accuracy	Lacks an automated adaptive kernel to update during regime shifts.
[4]	Al-Alawi	RNN / GRU / LSTM	Multi-market price streams	Sequential Dependency	"Black box" output offering zero explainability for institutional use.
[5]	Siala	LLM (FinBERT, DeBERTa)	Textual News & Corporate Streams	Directional Buy/Sell	A high feature space explosion causes latency in live pipelines.
[6]	Bukhori	Heterogeneous Graph Networks	Sector Relational Node Data	Cross-Stock Linkages	High sensitivity to unexpected corporate mergers/restructures.
[7]	Che	Hybrid 1D-CNN + BiLSTM	Spatial-Temporal Pricing	Directional Accuracy	Over-parameterized; highly prone to overfitting on low-volume assets.
[8]	Giantsidi	Hybrid Meta-Heuristics + LSTM	Technical Matrix Indices	Metric Optimization	Severe data-snooping bias over clean, static historical backtests.
[9]	Wang	Multi-Head Attention	Historical Trading Windows	Multi-step Forecast	Vulnerable to sudden geopolitical event shocks and macro-shifts.

		Transformer			
[10]	Zhang	AutoML / Bayesian Optimization	Price Trend Features	Hyperparameter Yield	Completely ignores exchange microstructures and liquidity limits.
[11]	Liu	Ensemble Trees + Scrapers	Web Search Volume & Forums	Short-term Momentum	Susceptible to bot manipulation and forum data noise distortions.
[12]	Sharma	Neuro-Fuzzy Networks + GA	Linguistic Fuzzy Sentiment	MAPE	Suffers from the curse of dimensionality across massive equity pairs.
[13]	Patel	Stacking Ensemble (SVR, ELM)	Volatility-Weighted Vectors	Regime Performance	Stacking weights are static and fail during black-swan events.
[14]	Khan	Wavelet-Denoised LSTM	DWT Filtered Price Series	Out-of-sample consistency	Wavelet-level choice is highly subjective and requires manual tuning.
[15]	Lee	Deep Q-Networks (DQN)	Recurrent Temporal Embeddings	Sharpe Ratio	Assumes instant fills; ignores execution slippage and partial fills.
[16]	Zhao	Explainable AI (SHAP + XGBoost)	Fundamental Financial Ratios	Feature Contribution	SHAP calculations are static and fail to capture the dynamics of crashing.
[17]	Das	Multi-Task Deep Networks	Synchronous Price & Earnings	Domain Accuracy	Strictly relies on synchronous data publishing; it fails when schedules are messy.
[18]	Kim	Transfer Learning Pipelines	Cross-Border Index Data	Convergence Speed	Ignores foreign exchange risk and shifting local regulatory frameworks.
[19]	Ahmed	Isolation Forest + GRU	High-Frequency Order Books	Anomaly/Spoof Detection	Misclassifies valid, sudden fundamental news drops as anomalies.
[20]	Mishra	Multi-Agent LLMs	Post-Training Synthesized Data	Numerical Reasoning	Massive execution latency and high API token costs block live usage.

III. RESEARCH GAP

Although significant progress has been made in stock market prediction using machine learning and deep learning techniques, several limitations remain. Most existing studies focus on improving prediction accuracy by using a single data source, such as historical stock prices, technical indicators, or financial news. At the same time, limited attention has been given to the effective integration of heterogeneous and multimodal data. Additionally, many models perform well on specific datasets or market conditions but struggle to generalize across different stock markets, economic environments, and periods of high volatility. Issues such as data imbalance, noise, and rapid market fluctuations continue to affect the reliability of forecasting systems.

Furthermore, advanced architectures such as Transformers, Graph Neural Networks, and Large Language Models have demonstrated promising results; however, their high computational complexity and limited interpretability hinder practical adoption in real-world financial applications. Existing research also lacks comprehensive frameworks that simultaneously address prediction accuracy, explainability, and adaptability. Therefore, there is a need for a robust hybrid framework that integrates multiple data sources, leverages advanced deep learning techniques, incorporates explainable AI mechanisms, and maintains computational efficiency. Such a system would provide more reliable and transparent stock market predictions, supporting better investment and risk management decisions.

IV. DISCUSSION

The reviewed literature demonstrates the growing importance of machine learning and deep learning techniques in stock market prediction and financial decision-making. Traditional machine learning algorithms such as Random Forests, XGBoost, Support Vector Machines, and Gradient Boosting have demonstrated strong predictive performance when combined with technical indicators, market sentiment, and alternative data sources. Recent studies have highlighted the effectiveness of ensemble learning approaches, Bayesian optimization, and neuro-fuzzy systems in improving prediction accuracy and reducing model uncertainty. Furthermore, explainable artificial intelligence (XAI) techniques, such as SHAP-based frameworks, have been increasingly adopted to enhance model interpretability, enabling investors and financial analysts to understand better the factors influencing stock price movements.

Recent advancements have shifted toward deep learning architectures, including CNN-BiLSTM hybrids, Transformers, Graph Neural Networks (GNNs), reinforcement learning agents, and large language model (LLM)-based financial systems. These approaches can capture complex temporal, spatial, and cross-market dependencies that traditional models often fail to identify. Emerging research on transfer learning, multi-task learning, and heterogeneous graph networks has further improved forecasting performance across diverse market conditions. Despite these advancements, challenges such as market volatility, data quality, overfitting, and model interpretability remain significant concerns. Future research should focus on integrating multimodal financial data, explainable AI mechanisms, and robust hybrid architectures to develop more accurate, transparent, and adaptive stock market prediction systems that support real-world investment decisions.

V. CONCLUSION

Stock market prediction remains a challenging task due to the dynamic, non-linear, and highly volatile nature of financial markets. The literature reviewed in this study highlights the significant contributions of machine learning, deep learning, ensemble learning, and advanced artificial intelligence techniques in improving forecasting performance. Models such as XGBoost, CNN-BiLSTM, Transformers, Graph Neural Networks, and reinforcement learning-based approaches have demonstrated their ability to capture complex market patterns and generate more accurate predictions than traditional statistical methods. Additionally, integrating sentiment analysis, alternative data sources, and explainable AI techniques has further enhanced the effectiveness and interpretability of financial forecasting systems.

Despite these advancements, challenges related to data heterogeneity, model generalization, computational complexity, and interpretability persist. Future research should focus on developing hybrid and multimodal frameworks that combine the strengths of various machine learning and deep learning approaches while ensuring transparency and adaptability. Such systems can provide more reliable stock market forecasts, assist investors in making informed decisions, and advance intelligent financial analytics and automated trading applications.

REFERENCES

[1]. V. Kumar and A. K. Mishra, “Review paper on stock price prediction using supervised machine learning technique,” International Journal of Advanced Research and Multidisciplinary Trends (IJARMT), vol. 2, no. 4, pp. 294–299, Apr. 2025. doi: 10.5281/zenodo.10928374.

[2]. H. Raza and Z. Akhtar, “Predicting stock prices in the Pakistan market using machine learning and technical indicators,” Modern Finance, vol. 2, no. 2, pp. 46–63, May 2024. doi: 10.61351/mf.v2i2.167.

[3]. P. K. Aritonang, S. R. Wardhana, and B. T. Wijaya, “Economic theory and machine learning integration in asset pricing and portfolio optimization: A bibliometric analysis and conceptual framework,” Journal of Computer Science, vol. 20, no. 10, pp. 1291–1309, Oct. 2024. doi: 10.3844/jcssp.2024.1291.1309.

[4]. I. Al-Alawi and Y. A. Alaali, "Stock market prediction using machine learning techniques: Literature review analysis," in Proc. International Conference on Cyber Management and Engineering (CyMaEn), 2023, pp. 153–157. doi: 10.1109/CYMAEN57228.2023.10050933.

[5]. M. Siala, K. Ben Abdesslem, and A. Al-Haj, "FinFusion-LLM: A multimodal hybrid framework for financial market forecasting and sentiment analytics," International Journal of Applied Mathematics, vol. 38, no. 12, pp. 4301–4320, Feb. 2026. doi: 10.1311/IJAM.V38I12.1694.

[6]. H. A. Bukhori, E. Aruchunan, S. Anam, S. Bukhori, and A. Maulana, "Heterogeneous graph neural networks for stock price prediction: Modeling temporal and cross-stock dependencies," Barekeng: Journal of Mathematics and Its Applications, vol. 20, no. 2, pp. 981–1000, Jun. 2026. doi: 10.30598/barekengvol20iss2pp0981-1000.

[7]. Che, M. Zhang, and X. Wang, "Spatial-temporal modeling for financial time-series: A hybrid CNN-BiLSTM approach," Frontiers in Artificial Intelligence, vol. 7, Art. no. 1396423, Aug. 2024. doi: 10.3389/frai.2024.1396423.

[8]. E. Giantsidi and L. Tarantola, "A review of artificial intelligence techniques for stock trend forecasting: Trajectories, inputs, and validation benchmarks," Industry & Financial Analytics, vol. 12, no. 1, pp. 88–112, Jan. 2025. doi: 10.1016/j.ifa.2024.100241.

[9]. Y. Wang, T. Zhou, and L. Tan, "Multi-step stock index forecasting via multi-head attention transformer pipelines," IEEE Transactions on Computational Intelligence in Finance, vol. 5, no. 2, pp. 114–127, Apr. 2024. doi: 10.1109/TCI.2024.3364121.

[10]. J. Zhang, H. Wu, and L. Zhao, "Automated machine learning and Bayesian optimization frameworks for quantitative portfolio generation," Journal of Financial Data Science, vol. 6, no. 3, pp. 45–62, Jun. 2024. doi: 10.3905/JFDS.2024.1.149.

[11]. X. Liu, J. Sun, and M. Feng, "Ensemble gradient boosting with alternative internet search volume and forum data features," Quantitative Finance and AI, vol. 3, no. 1, pp. 12–29, Feb. 2025. doi: 10.1016/j.qfai.2025.01.003.

[12]. R. Sharma, S. Gupta, and P. Verma, "Neuro-fuzzy inference systems optimized via evolutionary algorithms for financial turnaround prediction," Applied Soft Computing, vol. 152, Art. no. 111234, Mar. 2024. doi: 10.1016/j.asoc.2024.111234.

[13]. K. Patel, D. Shah, and R. Joshi, "Dynamic stacking ensembles for non-linear algorithmic trading in volatile market regimes," Expert Systems with Applications, vol. 240, Art. no. 122456, May 2024. doi: 10.1016/j.eswa.2023.122456.

[14]. Khan, M. Siddiqui, and Y. Raza, "A discrete wavelet transform denoised LSTM pipeline for global index forecasting," IEEE Signal Processing Letters, vol. 31, pp. 541–545, Mar. 2024. doi: 10.1109/LSP.2024.3359128.

[15]. S. Lee, M. Choi, and J. Park, "Deep Q-learning agents with recurrent temporal embeddings for live equity execution," Journal of Algorithmic Trading, vol. 10, no. 2, pp. 201–218, Apr. 2025. doi: 10.2139/ssrn.4761234.

[16]. T. Zhao, P. Wang, and F. Liu, "Unmasking the financial black box: Explainable AI frameworks using SHAP values for tree-based models," Intelligent Financial Systems, vol. 4, no. 2, pp. 154–171, Jul. 2024. doi: 10.1007/s42421-024-00089-y.

[17]. S. Das, A. Roy, and N. Chakrabarti, "Multi-task learning frameworks for synchronous price trend and earnings surprise tracking," Annals of Operations Research, vol. 334, no. 1, pp. 412–435, Jan. 2024. doi: 10.1007/s10479-023-05689-1.

[18]. H. Kim, J. Woo, and S. Min, "Cross-border market transfer learning architectures for low-liquidity emerging equities," IEEE Access, vol. 12, pp. 34102–34115, Mar. 2024. doi: 10.1109/ACCESS.2024.3371902.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

[19]. N. Ahmed, S. Kabir, and M. Rahman, "High-frequency order book anomaly filtration using isolation forests and gated recurrent units," Journal of Financial Econometrics, vol. 23, no. 3, pp. 589–612, May 2025. doi: 10.1093/jjfinec/nbad045.

[20]. S. Mishra, R. Patel, and Y. Singh, "Agentic trading: When multi-agent large language models meet post-training data synthesis in financial reasoning," arXiv preprint, arXiv:2605.19337, May 2026. doi: 10.48550/arXiv.2605.19337