

REVIEW

Stock Price Prediction: A Comprehensive Review of Methods and Trends

Qizhao Chen^{1,*} ¹Graduate School of Information Science, University of Hyogo, Japan

Abstract: Stock price prediction is an important problem in financial research. It is related to applications in investment, risk management, and algorithmic trading. However, accurate stock price prediction has a lot of challenges because stock price is affected by various factors such as market noise, non-stationarity, and external factors such as macroeconomic events, corporate news, and investor sentiment. In this review, we provide a comprehensive overview of traditional statistical models, machine learning approaches, emerging deep learning, and multimodal methods for predicting stock prices. We further discuss the recent application of large language models for sentiment extraction and direct stock price prediction. In addition, this review covers commonly used features in stock price prediction, including technical indicators, sentiment measures, and composite features, which are mathematical combinations of different basic features. Moreover, unconventional features such as environmental, social, and governance factors are included. We also discuss key challenges and open research directions, aiming to guide researchers in selecting suitable methodologies and identifying promising opportunities for future exploration.

Keywords: stock price prediction, LLM, deep learning, time-series analysis, multimodal models

1. Introduction

Forecasting stock movements remains one of the most persistent challenges in the field of finance and quantitative trading. The goal is to forecast future prices or price movements of financial instruments using historical data, market indicators, and additional information. Accurate predictions can provide significant advantages for investors, portfolio managers, and algorithmic trading systems, contributing to enhanced decision-making processes and more efficient risk management.

Nevertheless, stock markets exhibit high levels of noise and volatility, driven by complex interactions among macroeconomic variables, geopolitical events, and market sentiment. The efficient market hypothesis [1] posits that financial markets fully incorporate all available information into asset prices, which suggests that predicting price movements is theoretically impossible in an idealized efficient market. Despite this, empirical research [2] shows that short-term price patterns and correlations exist, providing opportunities for prediction through sophisticated models.

Traditional statistical models such as Autoregressive Integrated Moving Average (ARIMA) [3] and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) [4] have been widely used due to their interpretability and analytical foundations. With the rise of machine learning and deep learning, researchers have leveraged neural networks [5], ensemble methods [6], and transformers [7] to model nonlinear relationships and capture long-term dependencies in stock data. Recently, multimodal models that combine numerical price data, textual

sentiment analysis, and visual chart patterns have emerged as promising directions for improving predictive performance [8]. Figure 1 shows the evolution of stock price prediction methods.

This review aims to offer an in-depth examination of various approaches used for stock price prediction, from classical statistical approaches to modern AI-based techniques. We also discuss current challenges and future research directions.

In particular, although numerous reviews have summarized traditional statistical and deep learning models, few have systematically addressed the emerging generation of methods, including Transformer-based architectures, multimodal approaches that integrate numerical, textual, and visual data, and the application of large language models (LLMs) for financial forecasting. Furthermore, more and more investors' emphasis on sustainability has spurred interest in incorporating environmental, social, and governance (ESG) factors into predictive models, but their integration within advanced artificial intelligence (AI) frameworks remains underexplored in the literature. This review is motivated by the need to bridge this gap in the literature and provide a comprehensive, updated synthesis of these next-generation methods. By critically analyzing their capabilities, limitations, and potential synergies, we aim to give readers a clear understanding of the current state of the art and highlight avenues for future research in AI-driven stock price prediction.

2. Traditional Statistical Methods

Statistical approaches have historically served as the cornerstone of financial forecasting. These methods are valued for their interpretability, mathematical rigor, and analytical tractability. They are particularly useful when data is limited or when model

*Corresponding author: Qizhao Chen, Graduate School of Information Science, University of Hyogo, Japan. Email: af24o008@guh.u-hyogo.ac.jp

Figure 1
Methodological trends in stock price prediction

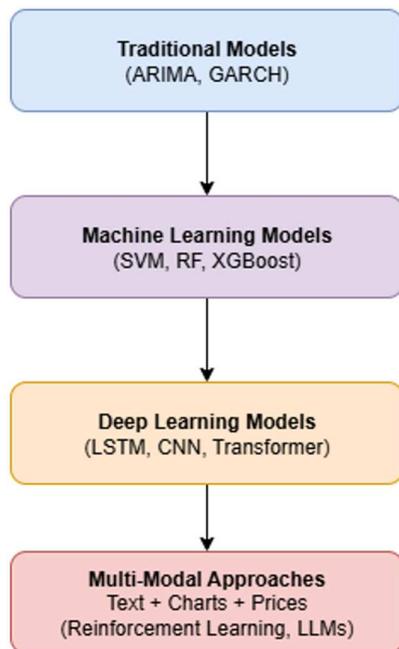


Table 1
Summary of traditional statistical models for stock prediction

Model	Input data	Strengths	Limitations
AR/ARIMA	Historical prices	Simple, interpretable	Cannot capture nonlinear patterns
GARCH	Price volatility series	Models volatility clustering	Sensitive to parameter assumptions
VAR	Multiple price series	Captures inter-stock dependencies	Assumes linear relationships

transparency is required. Table 1 shows the comparison between different traditional statistical models.

2.1. Autoregressive (AR) and ARIMA models

Autoregressive (AR) models estimate future stock prices based on a linear relationship with previous observations. ARIMA extends AR models by integrating differencing to achieve stationarity and adding a moving average component. ARIMA models are effective for capturing linear temporal dependencies but struggle with nonlinear patterns common in financial time series [9]. For example, Ariyo et al. [10] apply ARIMA models to stock price data from the New York Stock Exchange (NYSE) and Nigeria Stock Exchange (NSE), showing that ARIMA demonstrates strong potential for short-term prediction and performs competitively with other existing methods. Ho et al. [11] compare the performance of ARIMA with

LSTM and the neural network in the prediction of the closing price in Bursa Malaysia. The results show that LSTM has the best performance. Mahadik et al. [12] compare ARIMA and LSTM models for stock trend forecasting using historical price data, incorporating preprocessing techniques such as feature scaling and autocorrelation checks. The results indicate that both models achieve over 90% accuracy, with LSTM performing better on larger datasets with fewer missing values, while ARIMA provides higher accuracy but requires more computational time.

Nevertheless, the seemingly high accuracy figures in these research papers should be viewed with caution. These results might partly arise from overfitting to the historical dataset and can vary depending on the market conditions present during the testing period. Applying such models in actual trading can be difficult, as factors like market volatility, transaction fees, and unexpected events can lower their effectiveness. This underscores the importance of critically assessing reported performance instead of taking the numbers at face value.

2.2. GARCH models

GARCH models capture the phenomenon of volatility clustering commonly observed in financial markets. GARCH models predict the variance of returns rather than prices themselves, providing valuable insight into market risk [13]. For instance, Franses and Dijk [14] evaluate GARCH, Quadratic Generalized Autoregressive Conditional Heteroskedasticity (QGARCH) and GJsten–Jagannathan–Runkle (GJR) models for forecasting weekly stock market volatility, finding that QGARCH performs best in the absence of extreme events, while the GJR model is not recommended for prediction. Mutinda and Langat [15] introduce hybrid models that integrate GARCH with Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Transformer architectures to forecast the stock prices of Airtel, addressing the challenges of nonlinear and non-stationary financial data. Empirical results show that the hybrid models, particularly GARCH-LSTM, outperform their stand-alone counterparts, demonstrating improved accuracy and robustness in stock price forecasting. Caiado and Lucio [16] introduce a clustering approach based on forecast errors from asymmetric GARCH models to study how COVID-19 affects the US stock market industries, revealing that sectors like hotels and airlines were most affected, while pharmaceuticals and data processing were least impacted. Wang et al. [17] apply ARMA-GARCH, TGARCH, and EGARCH models to forecast volatility in the Shanghai Composite and Shenzhen Component indices, finding that ARMA(4,4)-GARCH(1,1) with t-distribution performs best for Shanghai, while ARMA(1,1)-TGARCH(1,1) is most effective for Shenzhen, providing insights for investors and policymakers.

In stock price prediction, the limitation of GARCH is obvious. A GARCH model is designed to forecast conditional volatility rather than the future level of stock prices. In practice, GARCH cannot directly predict stock prices, because it models the variance of returns rather than their mean. In addition, it should be emphasized that these volatility predictions rely on past data and the particular market context in which they were estimated. Unexpected shocks or structural shifts in the market could reduce their forecasting accuracy. Moreover, although these models are effective at capturing volatility clustering, they may fall short in representing nonlinear patterns or extreme events,

indicating that additional methods may be needed for practical applications.

2.3. Vector Autoregression (VAR)

Vector autoregression (VAR) models extend AR models to several interdependent time series. VAR is particularly useful for modeling correlations between multiple stocks or market indices, allowing analysts to capture the influence of one asset on another [18]. For example, Bessler and Luckoff [19] use a Bayesian vector autoregressive (BVAR) model, incorporating macroeconomic variables and past stock returns, to forecast returns of large German firms, showing that BVAR outperforms alternative time-series models, particularly over longer horizons. Khan et al. [20] propose a hybrid VAR model with smoothly clipped absolute deviation estimation to forecast US macroeconomic variables, demonstrating improved accuracy and efficiency over baseline models for multistep-ahead predictions. While these models are useful for short-term prediction and risk estimation, their reliance on linear assumptions limits their performance in highly complex markets. In addition, they can become computationally intensive when applied to a large number of time series.

3. Machine Learning Approaches

Machine learning methods allow for modeling nonlinear relationships and complicated relationships in financial data. They have become more and more popular in stock price forecasting because of their flexibility and performance potential. Table 2 demonstrates the comparison between some commonly used machine learning models in stock price prediction.

3.1. Supervised learning

Support vector machines (SVMs), Random Forests, and gradient boosting machines (GBM) (e.g., XGBoost) are widely used supervised learning algorithms that employ historical market

data along with engineered features to forecast subsequent price dynamics or directional tendencies.

3.1.1. Support Vector Machine

SVMs are commonly applied in stock forecasting tasks because they can effectively handle complicated and noisy financial data. In classification, SVM seeks the optimal hyperplane that separates stock movements (e.g., up or down) with the maximum margin, improving generalization. For regression, support vector regression (SVR) fits the data within a margin of tolerance, capturing nonlinear relationships between features such as historical prices, technical indicators, sentiment, and future stock prices. This makes SVM and SVR robust tools for both directional prediction and price forecasting.

Building on this, Hu et al. [21] apply SVM to predict stock performance using company-specific and macroeconomic factors, demonstrating that SVM effectively addresses model uncertainty and parameter instability, providing a robust tool for stock market forecasting. Mahmoodi et al. [22] integrate SVM with particle swarm optimization (SVM-PSO) to predict stock trading signals, showing superior performance with a 77.5% hit rate compared to SVM-Cuckoo Search and neural network models, demonstrating its effectiveness for short-term stock market forecasting. Purnama et al. [23] compare SVM and linear regression (LR) for predicting stock prices of PT. Vale Indonesia (INCO), finding that LR outperforms SVM, achieving a lower Root Mean Square Error (RMSE) of 42.82 compared to 64.32. Myilvahanan and Sundaram [24] develop a hybrid approach to forecast stock prices by combining LSTM and convolutional neural network (CNN) models optimized with Aquila Circle-Inspired Optimization and then fusing their outputs using SVM, achieving improved accuracy with an Mean Absolute Percentage Error (MAPE) of 0.378 and a normalized RMSE of 0.294. Henrique et al. [25] apply SVR to predict stock prices across multiple markets and frequencies, showing that SVR exhibits predictive power, particularly when models are periodically updated and during periods of lower volatility.

Table 2
Machine learning and deep learning models for stock price prediction

Model	Strengths	Limitations
SVM	Can model nonlinear relationships between features and price movements	Sensitive to parameter tuning and kernel choice; struggles with large datasets
Random Forest	Robust to noise and outliers; good for capturing feature importance	Less interpretable
Gradient Boosting	High predictive accuracy and flexibility; handles complex feature interactions	Prone to overfitting; less effective for sequential data without feature engineering
FNN/LSTM/GRU	Effectively captures temporal and sequential patterns in price data	Requires large datasets and long training time
CNN	Extracts local and spatial patterns from stock features or chart images (e.g., candlesticks)	May lose long-term dependencies; performance depends on input representation
Transformer	Captures long-range dependencies and attention across multiple stocks or modalities	Requires large data and high computational cost; sensitive to hyperparameters
Multimodal Approach	Integrates diverse data sources (text, images, numerical) to capture broader market signals	Complex model design and training; requires careful data alignment and fusion strategy

In terms of limitations, SVM is sensitive to parameter selection and kernel choice, which can lead to overfitting. Moreover, they struggle with large datasets, noisy financial data, and capturing complex nonlinear and temporal relationships.

3.1.2. Random Forest

The Random Forest algorithm integrates multiple decision trees, merging their outputs to deliver more accurate and robust predictions. In stock prediction, each tree may capture different aspects of market behavior, such as historical price patterns, technical indicators, or sentiment signals. By averaging predictions in regression and using majority voting in classification, Random Forest reduces overfitting and increases stability, which makes it a robust baseline model for stock movement and return forecasting. However, Random Forest can be computationally intensive and memory-demanding, especially with large datasets or many trees.

For example, Li [26] evaluates Random Forest for financial market forecasting, highlighting its capability to process data with a large number of features and reduce overfitting, while proposing integration with time-series methods like ARIMA and GARCH to address temporal dependencies and improve prediction accuracy. Du et al. [27] develop a stock market prediction model using Random Forest on historical trading data from Shanghai Stock Exchange (SSE) A-share stocks and Exchange-Traded Funds (ETFs), demonstrating that the approach achieves accurate predictions and provides valuable guidance for investors. Manojlovi and Štajduhar [28] use Random Forest to build 5- and 10-day-ahead stock market prediction models for the CROBEX index and selected Zagreb Stock Exchange companies, achieving average classification accuracies of 76.5% and 80.8%, respectively. Vijn et al. [29] use Artificial Neural Network and Random Forest models to predict the next-day closing prices of five companies with stock price features as inputs, showing strong performance through low RMSE and MAPE.

3.1.3. Gradient Boosting

Gradient boosting is an ensemble approach that builds models in sequence, where each successive tree addresses the mistakes of the previous ones. Unlike Random Forest, which trains trees independently, gradient boosting focuses on minimizing prediction errors by iteratively improving weak learners. In stock price forecasting, this allows the model to capture complex nonlinear relationships between historical prices, technical indicators, and sentiment features, often leading to higher accuracy but at the expense of greater computational complexity and risk of overfitting.

For instance, Yuvaraj et al. [30] use a GBM with EMA technical indicators to forecast adjusted closing stock prices, achieving outstanding performance with an R^2 of 0.99 and demonstrating its effectiveness in handling nonlinear patterns and mitigating autocorrelation in time-series data. Shahin et al. [31] propose a hybrid approach combining a Gradient Boosting Neural Network and SVR for feature selection, using technical indicators to predict cryptocurrency prices, and demonstrate higher accuracy compared to leading machine learning models. Deng et al. [32] use an explainable XGBoost model with sentiment features from institutional, individual, and foreign investors to predict the movement of Shanghai and Shenzhen composite indices, finding that institutional investor sentiment is the most influential and that the model achieves high forecasting accuracy. Reddy and Kumar [33] compare GBMs and Naive Bayes for stock price prediction, showing that GBM with a novel loss function achieves higher accuracy (92.3%) than Naive Bayes (87.7%). Mukhaninga et al. [34] apply

GBMs and principal component regression (PCR) to predict the JSE All-Share Index, finding that GBM consistently outperforms PCR in accuracy across multiple training–testing splits, with the superiority confirmed by the Diebold–Mariano test, highlighting GBM’s ability to capture nonlinear relationships in financial time series.

4. Deep Learning and Multimodal Approaches

Recent advances in deep learning and multimodal AI have further enhanced predictive performance in stock markets.

4.1. Neural networks

Deep learning approaches provide considerable improvements in stock price forecasting compared to conventional machine learning models like SVM, Random Forest, and gradient boosting. Although traditional models are proficient at capturing structured patterns and relationships, they often struggle with highly nonlinear and sequential dependencies inherent in financial time series [35]. Neural networks provide a powerful framework for capturing such nonlinearities, enabling them to learn complex feature interactions that may be difficult to engineer manually. Feedforward neural networks (FNNs) have been applied to short-term predictions, while recurrent neural networks (RNNs), including LSTM and GRU, excel at modeling temporal dependencies. LSTM and GRU networks, in particular, are adept at capturing long-term dependencies and mitigating vanishing gradient problems, making them especially suitable for dynamic and volatile stock market data.

For instance, Shinde et al. [36] assess the performance of LSTM models in forecasting long-term stock prices, showing that LSTM outperforms SVR, RNN, and other traditional models by better capturing stock price trends, providing a reliable tool for financial analysts and investors. Zhang et al. [37] propose a hybrid VMD–TMFG–LSTM model to predict stock price, combining Variational Mode Decomposition (VMD) for noise reduction, Triangulated Maximally Filtered Graph (TMFG) for feature selection, and LSTM for prediction. The model significantly outperforms ARIMA, Neural Network (NN), deep neural network (DNN), CNN, and other LSTM-based variants in forecasting multiple stocks, achieving substantial reductions in RMSE, Mean Absolute Error (MAE), and runtime, while improving R^2 , demonstrating enhanced accuracy, stability, and computational efficiency. Gao et al. [38] develop a stock price forecasting model incorporating technical indicators, investor sentiment, and financial data, using LASSO and Principal Component Analysis (PCA) for dimensionality reduction, and compare LSTM and GRU networks. Results show that both neural models perform efficiently, with LASSO-based models achieving better predictive accuracy than PCA-based models. Hafshejani and Mansouri [39] conduct a systematic review of 98 studies on LSTM networks applied to stock market prediction, highlighting their ability to model temporal dependencies in financial data. It also examines how LSTM networks can be integrated with approaches like sentiment analysis, providing insights for improving predictive accuracy and guiding future research and practical applications in financial forecasting. Rahmadyan and Mustakim [40] apply LSTM and GRU models to predict the stock price of Bank Rakyat Indonesia, focusing on the banking sector. The results indicate that the GRU model outperforms LSTM, achieving low Mean Squared Error (MSE), RMSE, and MAPE values, and predicts a decrease in stock prices in the following month. Das et al. [41] introduce GRU and ConvGRU models for

stock price prediction, leveraging historical data and feature extraction to capture temporal and spatial patterns. Experimental results show that GRU outperforms traditional models like LR and ARIMA, demonstrating its potential as a robust tool for forecasting stock prices, though predictions should be supplemented with broader market factors. Naeini et al. [42] compare two neural network architectures—feedforward multilayer perceptron (MLP) and Elman recurrent network—for stock price prediction. Results show that MLP better predicts the magnitude of stock value changes, while the Elman network and LR more accurately predict the direction of stock value movements. Kumarappan et al. [43] propose a Federated Learning-enhanced MLP–LSTM (Fed-MLP–LSTM) model for stock market prediction. By combining LSTM networks, which capture sequential dependencies, with Federated Learning, the approach allows multiple institutions to collaboratively train models on local CAC40 stock data while preserving privacy. Local MLP–LSTM models extract features and model sequences, and their parameters are aggregated centrally to create a global model. Evaluation using RMSE and accuracy demonstrates superior performance, achieving an RMSE of 0.0108 and 98.3% accuracy. The study highlights Fed-MLP–LSTM as a reliable, privacy-preserving solution for collaborative stock forecasting.

These studies indicate that RNNs, especially LSTM models, are capable of modeling temporal patterns and complex relationships in stock price series. However, their effectiveness often hinges on proper hyperparameter selection and the availability of sufficiently rich historical data. LSTMs can also be challenged by noisy or limited datasets, and their predictions may be prone to overfitting when tested on unseen data. Therefore, while they offer strong potential for sequential financial modeling, careful implementation and, in some cases, combination with other approaches may be necessary to ensure robust performance.

4.2. CNNs on stock charts

CNNs are frequently applied to analyze visual patterns in candlestick charts, price-volume charts, or heatmaps of correlations. Compared to numerical data, stock charts offer a richer representation by embedding temporal and structural patterns into a visual format. Candlestick charts, for instance, compactly convey information on price movement, volatility, and market sentiment that may be difficult to capture from raw numerical sequences. CNNs are particularly effective in this context because they automatically extract hierarchical features from visual inputs, reducing the reliance on manual feature engineering. By leveraging spatial feature extraction, CNNs can detect subtle shapes, textures, and local dependencies in stock charts that traditional numerical indicators or machine learning models might overlook. Transforming financial time series into images also enables researchers to apply advances from computer vision to uncover latent structures in stock movements, thereby enhancing predictive performance.

For example, Wojarnik [44] investigates the use of CNNs for analyzing stock market charts. Motivated by the growing popularity of deep learning in stock price prediction, the research builds a CNN capable of recognizing patterns from simplified stock chart images. Using TensorFlow and Keras, the model processes graphical data to extract relevant information. Experimental results demonstrate near-perfect efficiency, suggesting that CNNs have strong potential for analyzing stock data visually and capturing trends in share prices and other financial instruments. Khalid et al. [45] propose a convolutional DNN (2D-CNN) for predicting stock price trends by classifying images generated from

financial time-series data. Technical indicators computed over 21-day periods are transformed into images labeled as Sell, Hold, or Buy. Experimental results show that the 2D-CNN outperforms both LSTM models and one-dimensional CNNs, demonstrating its effectiveness in capturing complex patterns for stock trend prediction. Wu et al. [46] propose a hybrid model called Stock Sequence Array Convolutional LSTM (SACLSTM) for more accurate stock price prediction. The model combines CNN and LSTM networks by constructing a sequence array of historical stock data and leading indicators (e.g., options and futures). CNN extracts feature vectors from this array, which are subsequently input into LSTM models for forecasting. Experiments on 10 US and Taiwan stocks demonstrate that SACLSTM outperforms previous methods, achieving improved prediction accuracy.

4.3. Transformers and attention models

Through self-attention mechanisms, Transformers capture dependencies across both short and long ranges in sequential data. In contrast to LSTMs and GRUs, which process sequences one step at a time, Transformers operate on all time steps in parallel, making them more efficient and better adapted to large datasets. This parallelization allows Transformers to capture long-range dependencies without the vanishing gradient problem inherent in recurrent models. In the context of finance, Transformers have been applied to stock price sequences, sentiment series, and cross-asset correlations, where they excel at identifying subtle patterns across multiple modalities. Moreover, their ability to evaluate the importance of each time step and input feature enables them to perform well in multistep forecasting and to model complex interdependencies between markets. Compared to LSTMs, which are limited by sequential processing and memory bottlenecks, Transformers provide greater scalability, flexibility, and interpretability through attention weights, offering a significant advantage for stock price prediction.

For example, Gopali et al. [47] extend time-series forecasting to a multivariate setting, comparing models including LSTM, Bi-directional LSTM (Bi-LSTM), Temporal Convolutional Network (TCN), VAR, and Transformer-based Multi-Head Attention, and find that the Transformer model achieves superior performance for both stock and cryptocurrency data. Karthika et al. [48] develop a deep learning model to forecast stock prices using historical data, aiming to improve accuracy, stability, and generalization. The proposed BiLSTM-MTRAN-TCN-TFT model integrates Temporal Fusion Transformer (TFT) with MTRAN-TCN and Bi-LSTM, combining transformer and temporal convolutional networks for enhanced prediction. Using data from 14 Shanghai and Shenzhen firms and five benchmark stocks, the hybrid model consistently outperforms existing approaches across multiple evaluation metrics. Malibari et al. [49] also apply Transformer neural networks to predict stock price, employing self-attention mechanisms to identify nonlinear relationships in volatile time-series data. Their model predicts next-day closing prices using inputs from the Saudi Stock Exchange (Tadawul) and achieves over 90% predictive accuracy across four evaluation metrics. Yongchareon [50] investigates the performance of advanced deep learning models, such as Transformers, in comparison with conventional methods for predicting long-term stock market indices, employing data from the S&P 500, NASDAQ, and Hang Seng. The performance of 10 models is evaluated over multiple time horizons with both predictive accuracy and financial metrics such as returns, volatility, drawdown, and Sharpe ratio. Statistical tests confirm significant differences in performance. Results show

that transformer-based models (e.g., PatchTST) perform best in short-term forecasts, while simpler models provide more stable results in longer horizons. Hartanto and Gunawan [51] explore the application of the Temporal Fusion Transformer (TFT) for short-term stock market prediction. By integrating feature engineering and technical indicators and addressing multicollinearity with the variance inflation factor, the model captures complex temporal dynamics across multiple time series. Experimental results show that TFT significantly outperforms traditional statistical models and other transformer architectures. It achieves a remarkably low SMAPE of 0.0022, highlighting its ability to capture stock-specific patterns and improve forecasting accuracy.

4.4. Multimodal models

Multimodal models integrate multiple data sources, including numerical prices, textual news sentiment, or visual chart patterns. These models leverage complementary information to improve prediction accuracy. For example, Chen and Kawashima [52] propose a Dual Transformer model that incorporates news sentiment of related companies to improve stock price prediction. The architecture includes an Enhancement Transformer to strengthen inter-company correlations and a Forecast Transformer for price forecasting. Using polarity scores and historical prices from 2015 to 2024 for eight companies, the model predicts the closing price for the next day. Experiments show that considering related companies' sentiment improves accuracy, and the Dual Transformer outperforms Temporal Fusion Transformer, N-Beats, Informer, and LSTM in terms of MSE. Chen and Kawashima [53] investigate the use of LLMs for financial news sentiment analysis to enhance stock price forecasting with transformer-based models. Six sentiment models (GPT-4, Llama 3, Gemma 2, Mistral 7B, FinBERT, VADER) are compared, showing that modern LLMs outperform traditional models like FinBERT and VADER. Llama 3 is selected for classifying news sentiment of target companies. Stock prices are predicted using Informer, Transformer, TCN, LSTM, SVR, Random Forest, and Naive Forecast with various sliding windows. Results indicate that incorporating news sentiment improves prediction accuracy, with Informer achieving the best performance. An ablation study highlights the importance of Informer's generative-style decoder in enhancing predictions. Li et al. [54] address the challenges of stock price forecasting, noting that existing neural models only consider time-aligned stock correlations and assume static feature effectiveness. They propose MASTER (MArkert-guided Stock Transformer), a model that captures both short-term and cross-time stock correlations and leverages market information for automatic feature selection. MASTER achieves this by alternating between aggregating information within individual stocks and across different stocks, enabling it to model complex inter-stock relationships. Experimental results show that MASTER outperforms previous methods and provides interpretable insights into realistic stock correlations.

These studies collectively highlight the benefits of incorporating additional information, such as company interrelations, news sentiment, and market-driven feature selection, into transformer-based stock prediction models. They show that using data beyond historical prices can noticeably improve forecasting performance. At the same time, models like the Dual Transformer and MASTER tend to be complex, which can make them harder to interpret and more computationally demanding. Many of these approaches also rely on the assumption that correlations and sentiment patterns remain stable, which may not hold during sudden

market changes. This indicates that future work could explore more robust and adaptive models that balance accuracy with interpretability for practical use in financial markets.

4.5. Large Language Models

Recently, LLMs have proven effective in financial prediction by capturing subtle semantic and contextual information from large volumes of text. Originally, LLMs were primarily applied to sentiment analysis, where they outperformed traditional natural language processing methods in extracting market sentiment from financial news, social media, and reports.

Some notable research on the applications of LLMs is outlined as follows. Kirtac and Germano [55] evaluate OPT, BERT, FinBERT, and the Loughran-McDonald dictionary on over 965,000 US financial news articles (2010–2023), showing that the GPT-3-based OPT model achieves the highest performance with 74.4% prediction accuracy. Moreno and Ordieres-Meré [56] apply RoBERTa, FinBERT, and GPT to analyze sell-side equity analysts' reports from 2016 to 2022 on the IBEX 35 index. Findings show that LLM-extracted sentiment can effectively forecast stock price trends while mitigating bias in analysts' target prices, underscoring their value in supporting informed investment decisions. Mun and Kim [57] examine how LLMs can be applied to analyze financial news sentiment and inform investment strategies. Applying both discriminative models (BERT, FinBERT) and generative models (Llama 3.1, Mistral, Gemma 2) with advanced prompting techniques such as Super In-Context Learning and Bootstrapping, the results show that generative LLMs outperform discriminative ones, with long strategies yielding the best portfolio performance. The study also highlights explainability considerations and potential risks, demonstrating the value of LLMs for data-driven financial decision-making. Iacovides et al. [58] introduce FinLlama, a finance-specific LLM based on Llama 2 7B, designed to classify sentiment and quantify its strength in financial news. By fine-tuning on supervised financial data and employing parameter-efficient techniques, FinLlama provides nuanced sentiment analysis, improving portfolio returns and generating resilient investment strategies even in volatile market conditions.

More recently, researchers have begun to explore their direct application to stock price prediction by leveraging their generative and reasoning capabilities. This shift reflects a growing interest in using LLMs not only as feature extractors but also as predictive models that can integrate textual, numerical, and even multimodal inputs for forecasting. For example, Yan and Huang [59] propose MambaLLM, a hybrid framework combining state-space models and LLMs to integrate micro-level stock features with macroeconomic index information. Tested on six major US stocks, it reduces RMSE by up to 28.5% relative to traditional RNNs and MAMBA baselines, effectively capturing both asset-specific dynamics and broader market trends for stock price prediction. Koa et al. [60] propose the Summarize-Explain-Predict framework, which enables LLMs to generate explainable stock predictions autonomously. By combining a self-reflective agent with Proximal Policy Optimization (PPO), the framework lets the LLM teach itself to summarize, explain, and predict stock movements without requiring human-annotated explanations, achieving superior prediction accuracy and effectiveness in portfolio construction.

Despite the advantages of LLMs, the application of LLMs in finance faces several challenges. Computational costs remain high, and fine-tuning large-scale models requires significant resources. Domain adaptation is another critical issue, as generic

LLMs trained on general corpora may not capture the specific language of finance without targeted pre-training. Furthermore, the interpretability of LLM outputs and the potential risks of hallucination raise concerns for practical adoption in high-stakes financial decision-making. Nevertheless, the ongoing rapid development of LLMs implies they will increasingly contribute to financial forecasting and decision support, especially when leveraged alongside numerical and visual data within multimodal prediction frameworks.

4.6. Reinforcement learning for trading and portfolio optimization

Reinforcement learning (RL) approaches frame stock trading and portfolio optimization as sequential decision-making problems. An RL agent interacts with the market environment by choosing actions such as buy, sell, or hold, and obtains evaluative feedback as reward signals linked to portfolio performance. These rewards are often adjusted by risk metrics like the Sharpe ratio, ensuring that strategies balance both returns and risk. Unlike supervised models that passively predict prices, RL actively learns optimal trading and allocation strategies through trial and error, making it particularly suitable for dynamic portfolio management. RL methods are therefore highly promising for algorithmic trading, where adaptive, real-time decision-making is required.

For example, Chen and Kawashima [61] develop an RL framework using PPO to dynamically assign weights to formulaic alphas generated by LLMs for stock trading. Using DeepSeek to generate fifty alphas for five companies in different industries, the framework adaptively integrates signals and achieves higher returns and Sharpe ratios than equal-weighted portfolios and major market benchmarks. The results highlight the value of integrating LLM-generated signals with RL for more effective trading strategies. Jeon et al. [62] introduce FreQuant, a deep RL framework for optimizing portfolios by operating in the frequency domain through Discrete Fourier Transform techniques. By capturing both dominant and subtle market frequencies, FreQuant adapts more effectively to sudden market shifts than time-domain approaches. Experiments on real-world datasets show that FreQuant significantly improves profitability, achieving up to 2.1× higher annualized returns and 2.9× higher portfolio value compared to leading benchmarks. Choudhary et al. [63] introduce a Risk-Adjusted Deep Reinforcement Learning (RA-DRL) framework for portfolio optimization, integrating three DRL agents trained with different reward functions—log returns, differential Sharpe ratio, and maximum drawdown. A CNN fuses its actions into a single risk-adjusted policy, balancing multiple investment objectives. Tested on Sensex, Dow, TWSE, and IBEX data, RA-DRL consistently outperforms baseline DRL agents and benchmark methods across risk and return metrics. Enkhsaikhan et al. [64] propose a risk-constrained RL framework for portfolio optimization, integrating risk tolerance estimation with a Variational Autoencoder and Cost Network to manage epistemic uncertainty. The approach avoids unsafe actions, achieves zero constraint violations in testing, and shows strong potential for risk-averse investors.

These studies highlight the promise of combining RL with advanced signals, such as LLM-generated alphas or frequency-domain features, for portfolio optimization. Adaptive frameworks like PPO, FreQuant, and RA-DRL show improved returns and risk-adjusted performance. Still, their evaluation is largely based on historical data, and real-world challenges such as market shocks, transaction costs, and interpretability remain.

Future work could focus on enhancing robustness and practical applicability while keeping models computationally efficient.

5. Feature Engineering

In stock prediction, feature selection is essential because it has a direct impact on the performance and accuracy of predictive models. Historical stock prices, including open, high, low, close, and trading volume, are the most commonly used features, providing essential information on past market behavior and trends.

In addition to raw prices, technical indicators derived from historical data play a significant role in capturing market momentum and trend signals. These include measures such as Simple and Exponential Moving Averages, Relative Strength Index, Moving Average Convergence Divergence, and Bollinger Bands, which are widely used by traders and researchers to identify potential buy or sell signals. For example, Yuvaraj et al. [65] examine the ability of a Light Gradient Boosting Machine (LightGBM) to predict stock prices with Exponential Moving Averages (EMA_5) and Simple Moving Averages (SMA_5) as primary technical indicators. Fozap [66] introduces an LSTM-CNN deep learning model that leverages technical indicators to strengthen prediction performance for stock price movements.

Market sentiment is another valuable source of information and can be extracted from news articles, financial reports, or social media. Sentiment features allow models to incorporate behavioral and psychological factors that influence stock price movements, which are often not reflected in numerical data alone. For example, Maqbool et al. [67] propose a machine learning model that integrates financial news with historical stock data to predict short- and medium-term stock trends. Mamillapalli et al. [68] demonstrate that integrating sentiment analysis with machine learning improves stock price prediction, showing that the proposed GRUvader model outperforms stand-alone approaches and highlighting the strong correlation between sentiment indicators and market movements. Agrawal et al. [69] develop a hybrid model that combines insights from social media sentiment with technical indicators to improve stock market forecasting.

Finally, macroeconomic indicators such as interest rates, inflation, and GDP growth offer a broader economic context, helping models account for external factors that can significantly impact stock prices. Combining these diverse feature types enables machine learning models to capture a more comprehensive view of the market and enhance prediction performance. For example, Patsiarikas et al. [70] apply macroeconomic, technical, and sentiment indicators to forecast the S&P 500 index. Latif et al. [71] introduce a predictive model for S&P 500 returns that combines key macroeconomic indicators with 10 technical indicators to enhance forecasting accuracy.

Apart from basic features, researchers have also explored composite features to enhance stock prediction models. One notable example is formulaic alphas, which represent mathematically defined signals derived from various market inputs, including price and volume. Traditional approaches to mine these alphas include manual feature engineering [72], statistical analysis, and genetic programming. More recently, RL and deep learning-based generative frameworks have been applied to automatically generate and optimize sets of synergistic alphas. These alphas aim to capture subtle patterns and predictive signals that may not be evident from individual features alone, providing richer inputs for machine learning and deep learning models.

By integrating formulaic alphas, models can potentially improve forecasting accuracy and better inform trading strategies. For example, Yu et al. [73] address quantitative trading by focusing on formulaic alpha factors. Unlike traditional methods that mine alphas individually, their proposed framework prioritizes generating synergistic sets of alphas optimized for downstream combination models. Leveraging RL, the framework uses the contribution to portfolio returns as a reward to guide alpha generation. Experiments on real-world stock data demonstrate that this approach improves stock trend forecasting and achieves higher simulated investment returns compared to previous methods. Shi et al. [74] address the challenges of formulaic alpha factor mining in quantitative investment, where financial data is highly variable and noisy. Traditional methods with fixed alpha weights lack adaptability to dynamic markets. The proposed AlphaForge framework integrates a generative-predictive neural network for producing multiple alpha factors with a combination model that continuously adjusts factor weights based on their time-dependent effectiveness. Experiments on real-world datasets show that AlphaForge outperforms existing methods in alpha factor mining and enhances portfolio returns in quantitative investment. Zhao et al. [75] propose a novel RL algorithm based on REINFORCE, leveraging Monte Carlo policy gradient estimation with a dedicated baseline to reduce variance. Additionally, reward shaping with the information ratio encourages stable alpha factors adaptable to market volatility. Experiments on real-world data demonstrate improved correlation with returns and stronger excess return generation compared to existing alpha factor mining methods. Shi et al. [76] present a novel approach that leverages LLMs alongside Monte Carlo Tree Search (MCTS) to efficiently generate interpretable formulaic alpha factors. By leveraging LLM reasoning and MCTS-guided backtesting feedback, the approach discovers predictive alphas with improved accuracy, trading performance, and human interpretability compared to traditional methods. Tang et al. [77] introduce AlphaAgent, a framework that combines LLM-driven agents with regularization mechanisms to mine alpha factors resistant to decay. By enforcing originality, aligning factors with market hypotheses, and controlling complexity, AlphaAgent consistently generates predictive and durable alpha factors, outperforming traditional and LLM-based methods in both Chinese and US markets.

In addition to traditional financial indicators and textual data, ESG factors are increasingly being incorporated into stock price prediction models to account for sustainability considerations. ESG data can be used as numerical features, including scores or ratings from providers such as MSCI or Sustainalytics, or extracted via text-based sentiment analysis from company reports and news. These features can be integrated with conventional financial indicators in machine learning or deep learning models, including multimodal and Transformer-based approaches, allowing models to capture both market and sustainability-driven signals.

For example, Dincă et al. [78] investigate whether incorporating ESG factors improves financial forecast accuracy using machine learning models across 2548 firms from 98 countries. The results indicate that ESG scores generally do not enhance predictive performance, except marginally in the business services sector. Rosinus and Lansky [79] examine the impact of annual ESG scores on monthly stock return predictions for German DAX companies through multivariate LSTM models. Results show that incorporating ESG data does not enhance predictive accuracy,

suggesting that low-frequency ESG information offers little additional value beyond historical returns and highlighting the need for more timely data for investment decisions. Liang [80] investigates the impact of ESG performance on annual stock returns of Chinese A-share energy companies using an MLP model. Results indicate a significant positive relationship, with higher ESG scores increasing returns partly through improved profitability, demonstrating the potential of deep learning methods to leverage ESG metrics for financial forecasting.

Overall, these studies show mixed evidence regarding the usefulness of ESG factors in financial forecasting. These mixed results highlight several challenges in integrating ESG information into stock prediction. First, ESG data is often reported infrequently and lacks standardization across companies and regions, making it difficult to use in high-frequency forecasting. Second, ESG metrics may be noisy or subjective, reducing their reliability as predictive features. Finally, the impact of ESG factors can be sector- or market-specific, requiring careful model design and feature selection. Future research could focus on higher-frequency ESG data, better standardization, and adaptive modeling techniques to more effectively leverage ESG information in stock price prediction.

6. Challenges and Open Issues

While earlier sections discussed several specific challenges, this section synthesizes those points and further extends them to provide a broader perspective.

For stock forecasting, one of the most fundamental difficulties lies in the nature of financial markets, which are inherently noisy and often non-stationary. Market dynamics are influenced by a wide range of unpredictable factors, making it difficult for models to capture stable and consistent patterns over time.

Another challenge lies in the risk of overfitting, especially when dealing with complex neural networks. While such models may achieve impressive performance on training data, they can fail to generalize to unseen market conditions, reducing their reliability in practice. Closely related to this issue is the problem of data availability and quality. Financial datasets may be incomplete, inconsistent, or subject to biases, which can negatively impact model training and evaluation.

In addition, there exists a trade-off between interpretability and accuracy. Highly sophisticated models often achieve strong predictive performance, but their inner workings remain opaque to human analysts. This lack of transparency poses difficulties for practical adoption, especially in finance, where explainability is often essential. Finally, the issue of transferability remains a significant open problem. Models trained on one market or asset may not generalize well to other markets with different structures, regulations, or investor behaviors, limiting the broader applicability of prediction methods.

In addition to the challenges discussed above, data snooping remains a significant concern. Models may perform well on historical data but fail to generalize to unseen market conditions, leading to overly optimistic assessments. Another important issue is the lack of standardized benchmarks. Differences in datasets, preprocessing methods, evaluation metrics, and experimental setups make it difficult to compare models fairly or evaluate their real-world effectiveness. Addressing these issues is essential to improve the robustness, reproducibility, and reliability of financial prediction models.

7. Future Directions

Several promising directions can guide future research in stock prediction. One key area is the development of explainable AI methods, which enhance the transparency and trustworthiness of predictive models. For instance, researchers could apply SHapley Additive exPlanations (SHAP) values to a gradient boosting model to identify which features—such as trading volume, technical indicators, or news sentiment—most strongly influence a stock's predicted movement. By revealing the factors driving predictions, such techniques can help analysts better understand model behavior and increase confidence in using AI for real-world financial decisions.

Another promising avenue is multimodal learning, which integrates numerical, textual, and visual data to deliver a more complete understanding of market dynamics. For example, combining traditional stock features with news sentiment and candlestick chart analysis has the potential to capture richer and more diverse signals that can improve predictive accuracy. Related to this, LLMs represent a powerful tool for sentiment extraction and market analysis. Their ability to process and interpret vast amounts of unstructured text data enables a deeper understanding of market sentiment and its influence on asset prices.

Hybrid approaches that combine statistical models with AI-based methods also have considerable potential. Such methods can leverage the interpretability and robustness of traditional techniques while benefiting from the flexibility and predictive strength of modern machine learning models. Finally, risk-aware RL offers a promising pathway for developing trading strategies under uncertainty. By explicitly incorporating risk into decision-making, RL models can generate more reliable and practical strategies suited for real-world financial applications.

8. Conclusion

Prediction of stock prices has evolved from classical statistical methods to machine learning, deep learning, and multimodal approaches, each contributing to our understanding of market dynamics. Traditional models offered interpretability but were limited in capturing nonlinear patterns. Machine learning and neural networks enabled the modeling of complex relationships, revealing insights into patterns that classical methods could not detect. Multimodal deep learning models further integrate information from multiple sources, uncovering cross-domain dependencies and improving predictive performance. These advances have collectively expanded both the methodological toolkit and our scientific understanding of how financial, textual, and visual signals influence stock movements. Future research should focus on enhancing model interpretability, mitigating market noise, and developing approaches that generalize across markets and timeframes to further strengthen both practical forecasting and theoretical insights.

Ethical Statement

This study does not contain any studies with human or animal subjects performed by the author.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Author Contribution Statement

Qizhao Chen: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization.

References

- [1] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417. <https://doi.org/10.2307/2325486>
- [2] Seo, K., & Yang, J. (2024). Exploring the efficient market hypothesis for accurate stock movement prediction via feature-axis transformer. In *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*, 892–901. <https://doi.org/10.1145/3605098.3635928>
- [3] Kobiela, D., Krefta, D., Król, W., & Weichbroth, P. (2022). ARIMA vs LSTM on NASDAQ stock exchange data. *Procedia Computer Science*, 207, 3836–3845. <http://doi.org/10.1016/j.procs.2022.09.445>
- [4] Xu, Z., Liechty, J., Benthall, S., Skar-Gislinge, N., & McComb, C. (2024). GARCH-informed neural networks for volatility prediction in financial markets. In *Proceedings of the 5th ACM International Conference on AI in Finance*, 600–607. <https://doi.org/10.1145/3677052.3698600>
- [5] Turchenko, V., Beraldi, P., de Simone, F., & Grandinetti, L. (2011). Short-term stock price prediction using MLP in moving simulation mode. In *Proceedings of the 6th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems*, 2, 666–671. <http://doi.org/10.1109/IDAACS.2011.6072853>
- [6] Mozaffari, L., & Zhang, J. (2024). Predicting stock prices: Strategies of ensemble learning with transformer, ARIMA, and linear regression models. In *Proceedings of the 2024 7th International Conference on Machine Learning and Machine Intelligence*, 133–140. <http://doi.org/10.1145/3696271.3696293>
- [7] Sikarwar, S. S., Jaswal, S., Sharma, R., G, R., G., Ganguli, R., & Mahawar, M. K. (2024). Stock price prediction using transformers. In *2024 4th International Conference on Advancement in Electronics & Communication Engineering*, 243–248. <http://doi.org/10.1109/AECE62803.2024.10911285>
- [8] Nguyen, N. H., Nguyen, T. T., & Ngo, Q. T. (2025). DAS-Net: A multimodal framework for stock price forecasting with diffusion-based graph learning and optimized sentiment fusion. *Journal of Risk and Financial Management*, 18(8), 417. <http://doi.org/10.3390/jrfm18080417>
- [9] Hossain, M. L., Shams, S. N., & Ullah, S. M. (2025). Time-series and deep learning approaches for renewable energy forecasting in Dhaka: A comparative study of ARIMA, SARIMA, and LSTM models. *Discover Sustainability*, 6(1), 775. <https://doi.org/10.1007/s43621-025-01733-5>
- [10] Ariyo, A. A., Adewumi, A. O., & Ayo, C. K. (2014). Stock price prediction using the ARIMA model. In *2014 UKSim-AMSS 16th International Conference on Computer Modelling and Simulation*, 106–112. <https://doi.org/10.1109/UKSim.2014.67>

- [11] Ho, M. K., Darman, H., & Musa, S. (2021). Stock price prediction using ARIMA, neural network and LSTM models. *Journal of Physics: Conference Series*, 1(1), 012041. <https://doi.org/10.1088/1742-6596/1988/1/012041>
- [12] Mahadik, A., Vaghela, D., & Mhaisgawali, A. (2021). Stock price prediction using LSTM and ARIMA. In *2021 Second International Conference on Electronics and Sustainable Communication Systems*, 1594–1601. <https://doi.org/10.1109/ICESC51422.2021.9532655>
- [13] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- [14] Franses, P. H., & Van Dijk, D. (1996). Forecasting stock market volatility using (non-linear) GARCH models. *Journal of Forecasting*, 15(3), 229–235. [https://doi.org/10.1002/\(SICI\)1099-131X\(199604\)15:3<229::AID-FOR620>3.0.CO;2-3](https://doi.org/10.1002/(SICI)1099-131X(199604)15:3<229::AID-FOR620>3.0.CO;2-3)
- [15] Mutinda, J. K., & Langat, A. K. (2024). Stock price prediction using combined GARCH-AI models. *Scientific African*, 26, e02374. <https://doi.org/10.1016/j.sciaf.2024.e02374>
- [16] Caiado, J., & Lúcio, F. (2023). Stock market forecasting accuracy of asymmetric GARCH models during the COVID-19 pandemic. *The North American Journal of Economics and Finance*, 68, 101971. <https://doi.org/10.1016/j.najef.2023.101971>
- [17] Wang, Y., Xiang, Y., Lei, X., & Zhou, Y. (2022). Volatility analysis based on GARCH-type models: Evidence from the Chinese stock market. *Economic Research-Ekonomska Istraživanja*, 35(1), 2530–2554. <https://doi.org/10.1080/1331677X.2021.1967771>
- [18] Zivot, E., & Wang, J. (2003). Vector autoregressive models for multivariate time series. In *Modeling financial time series with S-Plus®* (1st ed., pp. 369–413). Springer New York. https://doi.org/10.1007/978-0-387-21763-5_11
- [19] Bessler, W., & Lückoff, P. (2008). Predicting stock returns with Bayesian vector autoregressive models. In *Data Analysis, Machine Learning and Applications: Proceedings of the 31st Annual Conference of the Gesellschaft für Klassifikation eV, Albert-Ludwigs-Universität Freiburg*, 499–506. https://doi.org/10.1007/978-3-540-78246-9_59
- [20] Khan, F., Iftikhar, H., Khan, I., Rodrigues, P. C., Alharbi, A. A., & Allohbi, J. (2025). A hybrid vector autoregressive model for accurate macroeconomic forecasting: An application to the US economy. *Mathematics*, 13(11), 1706. <https://doi.org/10.3390/math13111706>
- [21] Hu, Z., Zhu, J., & Tse, K. (2013). Stocks market prediction using support vector machine. In *2013 6th International Conference on Information Management, Innovation Management and Industrial Engineering*, 2, 115–118. <https://doi.org/10.1109/ICIII.2013.6703096>
- [22] Mahmoodi, A., Hashemi, L., Jasemi, M., Mehraban, S., Laliberté, J., & Millar, R. C. (2023). A developed stock price forecasting model using support vector machine combined with metaheuristic algorithms. *Opsearch*, 60(1), 59–86. <https://doi.org/10.1007/s12597-022-00608-x>
- [23] Purnama, I. P. C., Ginantra, N. L. W. S. R., Darma, I. W. A. S., & Udayana, I. P. A. E. D. (2023). Comparison of support vector machine (SVM) and linear regression (LR) for stock price prediction. In *2023 Eighth International Conference on Informatics and Computing*, 1–6. <https://doi.org/10.1109/ICIC60109.2023.10381982>
- [24] Karthick Myilvahanan, J., & Mohana Sundaram, N. (2025). Support vector machine-based stock market prediction using long short-term memory and convolutional neural network with aquila circle inspired optimization. *Network: Computation in Neural Systems*, 36(3), 1185–1220. <https://doi.org/10.1080/0954898X.2024.2358957>
- [25] Henrique, B. M., Sobreiro, V. A., & Kimura, H. (2018). Stock price prediction using support vector regression on daily and up to the minute prices. *The Journal of Finance and Data Science*, 4(3), 183–201. <https://doi.org/10.1016/j.jfds.2018.04.003>
- [26] Li, Y. (2025). Optimizing random forest models for stock market prediction with hyperparameter analysis. In *Proceedings of the 2nd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence*, 57–61. <https://doi.org/10.1145/3745238.3745250>
- [27] Du, S., Hao, D., & Li, X. (2022). Research on stock forecasting based on random forest. In *2022 IEEE 2nd International Conference on Data Science and Computer Application*, 301–305. <https://doi.org/10.1109/ICDSCA56264.2022.9987903>
- [28] Manojlović, T., & Štajduhar, I. (2015). Predicting stock market trends using random forests: A sample of the Zagreb stock exchange. In *2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics* (pp. 1189–1193). <https://doi.org/10.1109/MIPRO.2015.7160456>
- [29] Vijh, M., Chandola, D., Tikkiwal, V. A., & Kumar, A. (2020). Stock closing price prediction using machine learning techniques. *Procedia Computer Science*, 167, 599–606. <https://doi.org/10.1016/j.procs.2020.03.326>
- [30] Yuvaraj, S., Kumaran, J. C., Raja, R. V., & Jayanthi, G. (2025). Stock price prediction using gradient boosting machine with technical indicators. In *2025 International Conference on Intelligent and Innovative Technologies in Computing, Electrical and Electronics*, 1–6. <https://doi.org/10.1109/IITCEE64140.2025.10915343>
- [31] Shahin, T., Ballestar de las Heras, T. M., & Sanz, I. (2025). Enhancing stock market prediction using gradient boosting neural network: A hybrid approach. *Computational Economics*, 65(6), 3207–3235. <https://doi.org/10.1007/s10614-024-10671-9>
- [32] Deng, S., Huang, X., Zhu, Y., Su, Z., Fu, Z., & Shimada, T. (2023). Stock index direction forecasting using an explainable eXtreme gradient boosting and investor sentiments. *The North American Journal of Economics and Finance*, 64, 101848. <https://doi.org/10.1016/j.najef.2022.101848>
- [33] Reddy, P. V., & Kumar, S. M. (2022). A novel approach to improve accuracy in stock price prediction using gradient boosting machines algorithm compared with naive Bayes algorithm. In *2022 4th International Conference on Advances in Computing, Communication Control and Networking*, 695–699. <https://doi.org/10.1109/ICAC3N56670.2022.10074387>
- [34] Mukhaninga, M., Ravele, T., & Sigauke, C. (2025). Short-term forecasting of the JSE all-share index using gradient boosting machines. *Economies*, 13(8), 219. <https://doi.org/10.3390/economies13080219>
- [35] Liu, Z., Zhang, Z., & Zhang, W. (2025). A hybrid framework integrating traditional models and deep learning for multi-scale time series forecasting. *Entropy*, 27(7), 695. <https://doi.org/10.3390/e27070695>
- [36] Shinde, S., Wadhwa, L., Mohane, N., Pagar, V., Sherje, N., & Mane, S. (2023). Stock price prediction using lstm. In *2023 7th International Conference on Computing*,

- Communication, Control and Automation*, 1–7. <https://doi.org/10.1109/ICCUBEA58933.2023.10392023>
- [37] Zhang, Z., Liu, Q., Hu, Y., & Liu, H. (2025). Multi-feature stock price prediction by LSTM networks based on VMD and TMFG. *Journal of Big Data*, 12(1), 74. <https://doi.org/10.1186/s40537-025-01127-4>
- [38] Gao, Y., Wang, R., & Zhou, E. (2021). Stock prediction based on optimized LSTM and GRU models. *Scientific Programming*, 2021(1), 4055281. <https://doi.org/10.1155/2021/4055281>
- [39] Hafshejani, M. S., & Mansouri, N. (2025). Enhancing stock market prediction with LSTM: A review of recent developments and comparative analysis. *Archives of Computational Methods in Engineering*, 32(5), 1–98. <https://doi.org/10.1007/s11831-025-10370-0>
- [40] Rahmadeyan, A., & Mustakim. (2024). Long short-term memory and gated recurrent unit for stock price prediction. *Procedia Computer Science*, 234, 204–212. <https://doi.org/10.1016/j.procs.2024.02.167>
- [41] Das, H., Gourisaria, M. K., & Singh, J. P. (2024). A comprehensive study on stock prediction using GRU models. In *International Conference on Mathematical Modeling, Computational Intelligence Techniques and Renewable Energy*, 431–442. https://doi.org/10.1007/978-981-96-1449-3_34
- [42] Naeini, M. P., Taremiyan, H., & Hashemi, H. B. (2010). Stock market value prediction using neural networks. In *2010 International Conference on Computer Information Systems and Industrial Management Applications*, 132–136. <https://doi.org/10.1109/CISIM.2010.5643675>
- [43] Kumarappan, J., Rajasekar, E., Vairavasundaram, S., Kotecha, K., & Kulkarni, A. (2024). Federated learning enhanced MLP–LSTM modeling in an integrated deep learning pipeline for stock market prediction. *International Journal of Computational Intelligence Systems*, 17(1), 267. <https://doi.org/10.1007/s44196-024-00680-9>
- [44] Wojarnik, G. (2023). The potential of convolutional neural networks for the analysis of stock charts. *Procedia Computer Science*, 225, 941–950. <https://doi.org/10.1016/j.procs.2023.10.081>
- [45] Tatane, K., Rida, M., & Zaki, T. (2024). From time series to images: Revolutionizing stock market predictions with convolutional deep neural networks. *International Journal of Advanced Computer Science & Applications*, 15(1), 474. <https://doi.org/10.14569/IJACSA.2024.0150144>
- [46] Wu, J. M. T., Li, Z., Herencsar, N., Vo, B., & Lin, J. C. W. (2023). A graph-based CNN-LSTM stock price prediction algorithm with leading indicators. *Multimedia Systems*, 29(3), 1751–1770. <https://doi.org/10.1007/s00530-021-00758-w>
- [47] Gopali, S., Siami-Namini, S., Abri, F., & Namin, A. S. (2024). A comparative multivariate analysis of VAR and deep learning-based models for forecasting volatile time series data. *IEEE Access*, 12, 155423–155436. <https://doi.org/10.1109/ACCESS.2024.3447354>
- [48] Karthika, K., Balasubramanie, P., Harishmitha, S., Shanmugapriya, P., & Ramya, T. E. (2025). Deep learning based hybrid transformer model for stock price prediction. In *2025 International Conference on Multi-Agent Systems for Collaborative Intelligence*, 1603–1608. <https://doi.org/10.1109/ICMSCI62561.2025.10894439>
- [49] Malibari, N., Katib, I., & Mehmood, R. (2021). Predicting stock closing prices in emerging markets with transformer neural networks: The Saudi stock exchange case. *International Journal of Advanced Computer Science and Applications*, 12(12), 876–886. <https://doi.org/10.14569/IJACSA.2021.01212106>
- [50] Yongchareon, S. (2025). AI-driven intelligent financial forecasting: A comparative study of advanced deep learning models for long-term stock market prediction. *Machine Learning and Knowledge Extraction*, 7(3), 61. <https://doi.org/10.3390/make7030061>
- [51] Hartanto, S., & Gunawan, A. A. S. (2024). Temporal fusion transformers for enhanced multivariate time series forecasting of Indonesian stock prices. *International Journal of Advanced Computer Science and Applications*, 15(7), 140–148. <https://doi.org/10.14569/IJACSA.2024.0150713>
- [52] Chen, Q., & Kawashima, H. (2026). A novel sentiment correlation-based method with dual transformer model for stock price prediction. *International Journal of Data Science and Analytics*, 21(1), 39. <https://doi.org/10.1007/s41060-025-00932-7>
- [53] Chen, Q., & Kawashima, H. (2024). Stock price prediction using LLM-based sentiment analysis. In *2024 IEEE International Conference on Big Data*, 4846–4853. <https://doi.org/10.1109/BigData62323.2024.10825946>
- [54] Li, T., Liu, Z., Shen, Y., Wang, X., Chen, H., & Huang, S. (2024). Master: Market-guided stock transformer for stock price forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(1), 162–170. <https://doi.org/10.1609/aaai.v38i1.27767>
- [55] Kirtac, K., & Germano, G. (2024). Sentiment trading with large language models. *Finance Research Letters*, 62, 105227. <https://doi.org/10.1016/j.frl.2024.105227>
- [56] Moreno, A., & Ordieres-Meré, J. (2025). Predicting stock price trends using language models to extract the sentiment from analyst reports: Evidence from IBEX 35-listed companies. *Economics Letters*, 254, 112404. <https://doi.org/10.1016/j.econlet.2025.112404>
- [57] Mun, Y., & Kim, N. (2025). Leveraging large language models for sentiment analysis and investment strategy development in financial markets. *Journal of Theoretical and Applied Electronic Commerce Research*, 20(2), 77. <https://doi.org/10.3390/jtaer20020077>
- [58] Iacovides, G., Konstantinidis, T., Xu, M., & Mandic, D. (2024). Finllama: LLM-based financial sentiment analysis for algorithmic trading. In *Proceedings of the 5th ACM International Conference on AI in Finance*, 134–141. <https://doi.org/10.1145/3677052.369869>
- [59] Yan, J., & Huang, Y. (2025). MambaLLM: Integrating macro-index and micro-stock data for enhanced stock price prediction. *Mathematics*, 13(10), 1599. <https://doi.org/10.3390/math13101599>
- [60] Koa, K. J., Ma, Y., Ng, R., & Chua, T. S. (2024). Learning to generate explainable stock predictions using self-reflective large language models. In *Proceedings of the ACM Web Conference 2024*, 4304–4315. <https://doi.org/10.1145/3589334.3645611>
- [61] Chen, Q., & Kawashima, H. (2025). Adaptive alpha weighting with PPO: Enhancing prompt-based LLM-generated alphas in quant trading. *arXiv Preprint*: 2509.01393
- [62] Jeon, J., Park, J., Park, C., & Kang, U. (2024). FreQuant: A reinforcement-learning based adaptive portfolio optimization with multi-frequency decomposition. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1211–1221. <https://doi.org/10.1145/3637528.3671668>

- [63] Choudhary, H., Orra, A., Sahoo, K., & Thakur, M. (2025). Risk-adjusted deep reinforcement learning for portfolio optimization: A multi-reward approach. *International Journal of Computational Intelligence Systems*, 18(1), 126. <https://doi.org/10.1007/s44196-025-00875-8>
- [64] Enkhsaikhan, B., & Jo, O. (2024). Uncertainty-aware reinforcement learning for portfolio optimization. *IEEE Access*, 12, 166553–166563. <https://doi.org/10.1109/ACCESS.2024.3494859>
- [65] Yuvaraj, S., Kumaran, J. C., Benschwartz, R., & Pradeepa, M. (2025). Technical indicator-driven stock prediction: A LightGBM and ARIMA perspective. In *2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things*, 46–52. <https://doi.org/10.1109/IDCIOT64235.2025.10914899>
- [66] Fozap, F. M. P. (2025). Hybrid machine learning models for long-term stock market forecasting: Integrating technical indicators. *Journal of Risk and Financial Management*, 18(4), 201. <https://doi.org/10.3390/jrfm18040201>
- [67] Maqbool, J., Aggarwal, P., Kaur, R., Mittal, A., & Ganaie, I. A. (2023). Stock prediction by integrating sentiment scores of financial news and MLP-regressor: A machine learning approach. *Procedia Computer Science*, 218, 1067–1078. <https://doi.org/10.1016/j.procs.2023.01.086>
- [68] Mamillapalli, A., Ogunleye, B., Timoteo Inacio, S., & Shobayo, O. (2024). GRUVader: Sentiment-informed stock market prediction. *Mathematics*, 12(23), 3801. <https://doi.org/10.3390/math12233801>
- [69] Agrawal, S., Kumar, N., Rathee, G., Kerrache, C. A., Calafate, C. T., & Bilal, M. (2025). Improving stock market prediction accuracy using sentiment and technical analysis. *Electronic Commerce Research*, 25(5), 4103–4126. <https://doi.org/10.1007/s10660-024-09874-x>
- [70] Patsiarikas, M., Papageorgiou, G., & Tjortjis, C. (2025). Using machine learning on macroeconomic, technical, and sentiment indicators for stock market forecasting. *Information*, 16(7), 584. <https://doi.org/10.3390/info16070584>
- [71] Latif, S., Aslam, F., Ferreira, P., & Iqbal, S. (2024). Integrating macroeconomic and technical indicators into forecasting the stock market: A data-driven approach. *Economies*, 13(1), 6. <https://doi.org/10.3390/economies13010006>
- [72] Kakushadze, Z. (2016). 101 formulaic alphas. *Wilmott*, 2016(84), 72–81. <https://doi.org/10.1002/wilm.10525>
- [73] Yu, S., Xue, H., Ao, X., Pan, F., He, J., Tu, D., & He, Q. (2023). Generating synergistic formulaic alpha collections via reinforcement learning. In *KDD '23: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5476–5486. <https://doi.org/10.1145/3580305.3599831>
- [74] Shi, H., Song, W., Zhang, X., Shi, J., Luo, C., Ao, X., . . . , & Seco, L. A. (2025). AlphaForge: A framework to mine and dynamically combine formulaic alpha factors. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(12), 12524–12532. <https://doi.org/10.1609/aaai.v39i12.33365>
- [75] Zhao, J., Zhang, C., Qin, M., & Yang, P. (2025). Quant-Factor REINFORCE: mining steady formulaic alpha factors with variance-bounded REINFORCE. *IEEE Transactions on Signal Processing*, 73, 2448–2463. <https://doi.org/10.1109/TSP.2025.3576781>
- [76] Shi, Y., Duan, Y., & Li, J. (2025). Navigating the alpha jungle: An LLM-powered MCTS framework for formulaic factor mining. *arXiv Preprint: 2505.11122*
- [77] Tang, Z., Chen, Z., Yang, J., Mai, J., Zheng, Y., Wang, K., & Lin, L. (2025). AlphaAgent: LLM-driven alpha mining with regularized exploration to counteract alpha decay. In *KDD '25: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, 2813–2822. <https://doi.org/10.1145/3711896.3736838>
- [78] Dincă, M. S., Ciotlăuși, V., & Akomeah, F. (2025). Estimating the impact of ESG on financial forecast predictability using machine learning models. *International Journal of Financial Studies*, 13(3), 166. <https://doi.org/10.3390/ijfs13030166>
- [79] Rosinus, M., & Lansky, J. (2025). Predictive power of ESG factors for DAX ESG 50 index forecasting using multivariate LSTM. *International Journal of Financial Studies*, 13(3), 167. <https://doi.org/10.3390/ijfs13030167>
- [80] Liang, J. (2025). Deep multi-layer perceptron regression for ESG-driven stock return prediction in China's energy industry. In *BDAIE '25: Proceedings of the 2025 International Conference on Big Data, Artificial Intelligence and Digital Economy*, 75–80. <https://doi.org/10.1145/3767052.3767064>

How to Cite: Chen, Q. (2026). Stock Price Prediction: A Comprehensive Review of Methods and Trends. *FinTech and Sustainable Innovation*. <https://doi.org/10.47852/bonviewFSI62027630>