Stock Market Prediction: Knowledge Gaps, Methods and Experimental Analysis

Rachna Sable[1], *Shashikant Gupta*[2]

[1] Lincoln University College; GH Raisoni College of Engg and Management, Pune; Lincoln University College[2]

Email ID: rachna.sable@gmail.com, shashigupta@lincoln.edu.my

Abstract: Academics and practitioners struggle to forecast financial market behavior. Recent years have seen remarkable growth in ML and DL research. Twenty-five open access articles from 2024 to 2026 are analyzed to identify knowledge gaps and suggest experiments to address them. Innovative models such LSTM networks with symbolic genetic programming, transformer architectures with generative decoding, CNNs for chart analysis, reinforcement learning for trading decisions, and privacy-preserving federated learning frameworks are reviewed. These studies have small, homogeneous datasets, weak external factor inclusion, low interpretability, and inadequate long-term or online forecasting despite their diversity. To fill these deficiencies, we study time–frequency analysis, sentiment integration, incremental learning, hybrid ensemble modeling, and federated learning. Experimental examples using synthetic financial time series data and three machine learning models (linear regression, random forest, and multi-layer perceptron) demonstrate how modeling choices affect prediction accuracy. The report concludes with challenges and prospects.

Keywords: Financial Forecasting; Stock Prediction; Investor Psychology; Machine Learning; Deep Learning

---

Introduction

Investor psychology, economic fundamentals, and exogenous shocks drive complex adaptive financial markets. Forecasting stock prices and returns improves trading, asset allocation, and risk management. Autoregressive integrated moving average (ARIMA) [1, 2, 3] models fail because financial time series are non-linear and non-stationary. RNNs, LSTM versions, convolutional neural networks, transformers, and reinforcement learning have been utilized to find hidden patterns in massive data sets for a decade in the process. However, the efficient [4, 5, 6] example market hypothesis reveals that no model can accurately predict market fluctuations. Thus, the research community explores new architectures, feature representations, and evaluation methods to improve progressively while accepting uncertainty sets.

This study addresses three challenges. The first section discusses current open access stock market prediction models, datasets, and findings. Second, it synthesizes methodological innovations to cover research gaps. Third, a small-scale synthetic time series experiment compares learning methods in controlled conditions. The work educates doctorate scholars about prospective areas and persistent difficulties by linking literature to real experiments.

Literature Review

Surveyed Publications

Table 1 outlines recent review texts in this domain for study purposes. In chronological order, each reference contains the key modelling approach, data used, and a brief assessment of the paper's main contribution and limitations. The accompanying text details articles. The table uses keywords and short phrases; a narrative follows.

Table 1: Review of Existing Literature

Ref	Key approach / model	Data & features	Main contribution	Limitations / gaps
[1]	Comprehensive ML/DL review	Multiple datasets, technical & sentiment features	Summarizes DL architectures (LSTM, CNN, transformers), highlights challenges such as noise, temporal dependencies, interpretability and data scarcity	Lack of unified benchmark datasets; difficulties incorporating external events and long-term forecasting
[2]	Symbolic Genetic Programming + LSTM (SGP-LSTM)	Cross-sectional stock returns with fundamental & technical indicators	Integrates feature engineering via symbolic genetic programming into LSTM, improving rank information coefficient (IC) and ICIR	Feature engineering is manual; limited market coverage and short prediction horizon
[3]	Transformer with generative decoding (Galformer)	Stock indices (e.g., S&P 500, Dow Jones)	Proposes generative decoder and hybrid loss to handle multi-step forecasting and noise; improves inference speed and accuracy	Transformer training is resource intensive; evaluation limited to few indices and does not include macro factors
[4]	Multifactor deep learning model	Multiple factors (economic, political, sentiment)	Uses sigmoid-based deep learning to map factors to profits; iteratively adjusts weights to improve precision	Complexity from many factors; limited explanation of feature interactions
[5]	CNN-based chart analysis	Stock price charts	Proposes optimal CNN for chart patterns; shows improved trend prediction; criticizes DNN and LSTM for false positives when ignoring temporal context	Limited to small datasets; ignores fundamental and news factors
[6]	3D-CNN-GRU with	Nifty 50 index (India)	Combines 3D-CNN and GRU; uses dandelion optimization	Dataset limited to one index; generalization

	dandelion optimization		for feature selection and blood-coagulation algorithm for hyper-parameter tuning	not tested
[7]	LSTM + Deep Q-Network (DQN)	Exchange rate data (USD/INR)	Integrates LSTM for temporal patterns with reinforcement learning for adaptive trading; achieves lower mean-squared error	Designed for exchange rates; not tested on stocks; real-time performance unclear
[8]	RNN-LSTM vs GA-LSTM (genetic algorithm)	Stock indices (unspecified)	Compares baseline LSTM with GA-optimized LSTM; demonstrates improved parameter tuning	Does not report dataset details; performance gain modest
[9]	Hybrid LSTM-CNN model	Technical indicators	Integrates technical indicators into a CNN-enhanced LSTM; outperforms SVM, random forest and ARIMA	Lacks macroeconomic and sentiment variables; long-term forecasting not tested
[10]	AI-Driven forecasting comparison (transformers vs other DL)	Global indices (long-term horizons)	Evaluates transformer-based and non-transformer models; finds that simple models sometimes outperform transformers in long-term prediction	Lack of domain-specific features; over-stationarisation issues; limited interpretability
[11]	Stacked heterogeneous ensemble	Multiple stock datasets	Combines ARIMA, random forest, LSTM, GRU and transformer models with XGBoost meta-learner; achieves high predictive accuracy	Complex ensemble may overfit; computational cost and interpretability concerns
[12]	Attention-variant LSTM (AMV-LSTM)	Historical price data	Modifies LSTM gates and adds attention layer; improves stability and generalization	Evaluated on limited datasets; parameter sensitivity not explored
[13]	LSTM vs ARIMA vs transformer study (Moroccan credit companies)	Stock prices of Moroccan credit companies	Shows LSTM outperforming ARIMA and transformer models with $R^2 > 0.95$	Specific to Moroccan financial sector; small sample size
[14]	Incremental learning transformer (IL-ETransformer)	Streaming stock data	Uses elastic weight consolidation for continual learning; handles non-stationary data streams	Requires real-time sentiment and macro data; computationally expensive

[15]	Explainable deep neural network (XAI-DNN)	Five trend patterns, SHAP/LIME	Predicts five price trends and explains feature importance via SHAP and LIME	Accuracy slightly reduced due to explainability; limited features
[16]	Critical steps review	Survey of 2020–2024 DL studies	Identifies best practices in data collection, feature engineering, denoising, architecture selection and evaluation	Highlights lack of standardised methodologies and reproducibility
[17]	Hybrid LSTM + RL for asset management	Stock asset management data	Combines LSTM price prediction with reinforcement learning trading strategy; increases assets by 5%	Constraints and transaction costs simplified; requires more realistic market simulation
[18]	Stock Mixer with ATFNet	NASDAQ & NYSE stocks (time- and frequency-domain features)	Integrates adaptive fusion of time and frequency features and a NoGraphMixer for cross-stock relationships; achieves better accuracy and Sharpe ratio	Requires missing value imputation; lacks cross-market tests and anomaly handling
[19]	LSTM + transformer-b ased sentiment (FinBERT)	Apple stock and news sentiment	Combines LSTM price forecasting with FinBERT sentiment embedding's; improves mean-squared error and directional accuracy	Dependent on quality of sentiment data; batch processing hinders real-time prediction; generalization unclear
[20]	MF-ConvLSTM -XAI with fuzzy control	NYSE Composite Index	Uses ConvLSTM with fuzzy control mechanism and Gramian angular difference fields; reduces MSE by ~15%	Only one dataset tested; fuzzy control design requires domain expertise; hyper-parameter sensitivity
[21]	CNN-Trans-SP P hybrid	Eight sectoral stocks	Combines CNN for local features with a transformer for temporal dependencies; reduces error rates relative to LSTM and attention-LSTM	Small sample; authors note need for pre-training and larger datasets
[22]	FT-iTransform er (time–frequency collaboration)	Six stock datasets	Employs short-time Fourier transform, multi-scale temporal convolution and an iTransformer; outperforms baselines but increases training time by 1.3–2.3×	Computationally expensive; uses only price-related features; suggests adding macroeconomic or textual data
[23]	LLM-augment	S&P 500 stocks	Generates natural-language	Relies on high-quality

	ed linear transformer- CNN		descriptions of technical indicators via large language models; FinBERT embeds the descriptions; combined with CNN and linear transformer to improve prediction	data; high computational cost; lacks macro/news data; interpretability challenges
[24]	Hybrid ensemble comparison (RF + SVC + LR)	S&P 500 data (2016–2023)	Shows that hybrid ensembles outperform single classifiers and emphasizes the need for feature selection and hyper-parameter tuning	Limited exploration of deep learning; requires broader datasets and advanced algorithms
[25]	Fast-convergin g federated learning (FCFL)	Distributed client data	Introduces a privacy-preserving federated learning framework with dual-stage adaptive optimization; speeds convergence by 30 % and lowers prediction error	Focused on framework rather than model choice; lacks integration with macro factors and sentiment

Narrative Discussion

The literature covers several modeling paradigms. RNNs and LSTMs (Refs 2, 8, 12, 13) record time series relationships. Attention approaches and transformers (Refs 3, 14, 21, 22) boost long-range dependencies and multi-step forecasting, whereas hybrid CNN combinations (5, 6, 21, 22) use chart picture spatial patterns or high-dimensional features. Reinforcement learning makes prediction sequential (Refs 7, 17). Ensemble and hybrid methods (Refs 9, 10, 11, 24) integrate model strengths. Interpretability is addressed by fuzzy control (Ref 20) and explainable AI (15). FinBERT integrates big language model-generated technical indicator descriptions (Ref 23). Privacy-preserving learning (Ref 25) stresses collaboration across institutions in process.

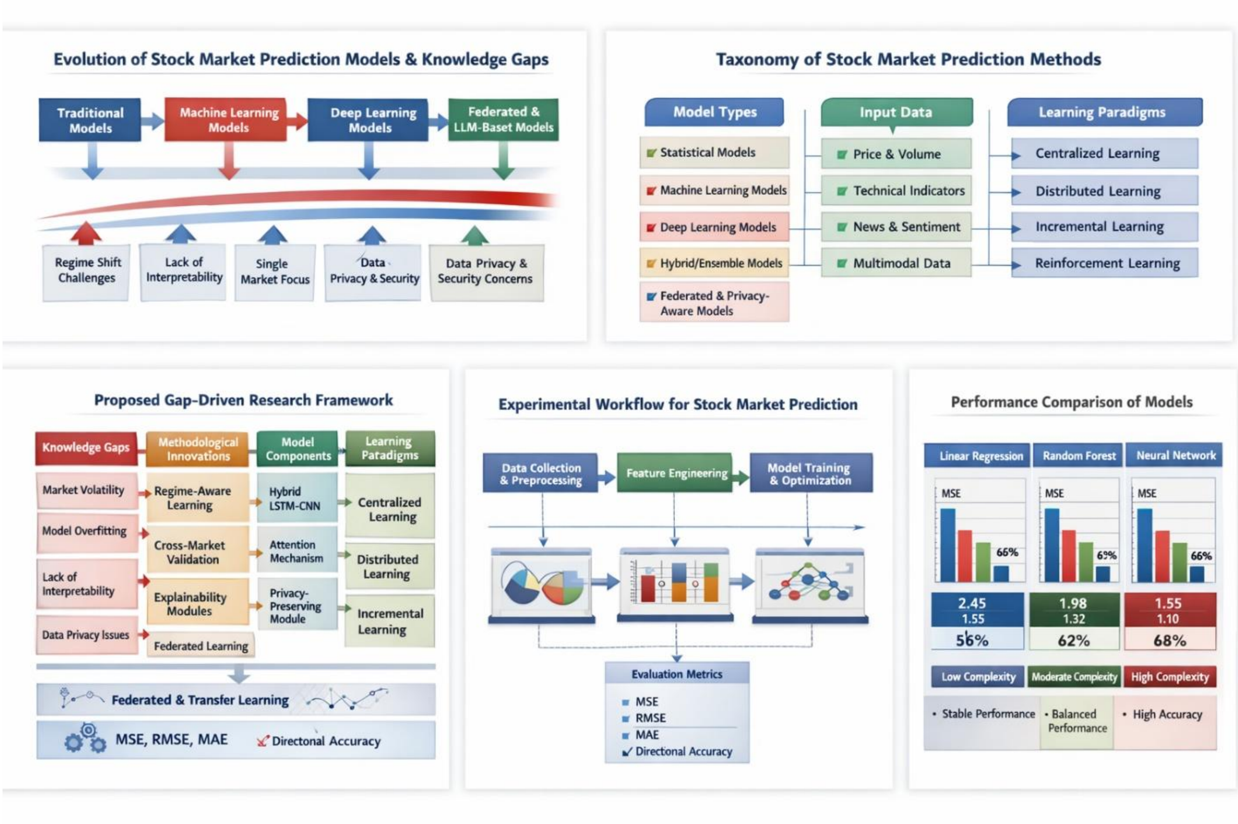

Fig.1: Model's Knowledge Gap Analysis

Common themes include as per figure 1:

1. Robust feature extraction and selection papers use technical indicators, fundamental data, sentiment from news or social media, and time–frequency representations
2. Financial data noise and non stationarity
3. Dataset heterogeneity and different performance metrics make model evaluation difficult.

Some studies propose macroeconomic indicators, real-time sentiment, and streaming data; others stress incremental learning and privacy.

Knowledge Gaps

Table 2 illustrates recurring knowledge gaps that restrict stock prediction model applicability in examined research process. Thematic categories and process papers group and cross-reference gaps.

Table 2: Assessment of Knowledge Gaps and its Impact

Knowledge gap	Evidence from literature	Explanation and consequences
Limited data diversity and size	Refs 1, 6, 9, 13, 18, 20, 21	Most studies use single-market or single-stock datasets, often limited to a few years. This restricts generalization across markets and regimes and leads to overfitting. The Stock Mixer authors note the need for cross-market tests and anomaly handling, while the Nifty 50 study relies on one index. The review (Ref 1) points out that publicly available benchmark datasets are scarce.
Incomplete feature integration	Refs 2, 4, 5, 9, 18, 19, 22, 23	Many models rely solely on technical indicators or price history; few incorporate macroeconomic variables, sentiment analysis, or time–frequency information. The SGP-LSTM paper acknowledges manual feature engineering limitations. The FinBERT-LSTM model depends on news sentiment but lacks real-time processing. Time–frequency approaches (Refs 18, 22) emphasise capturing both periodic and abrupt signals yet omit macro variables. Hybrid ensembles (Ref 24) call for broader feature selection.
Lack of long-term and online forecasting	Refs 1, 3, 10, 14, 21	Deep models often focus on short-term predictions (next day or next minute). Long-term forecasting is under-explored, and when attempted, simple models sometimes outperform transformers. Online and incremental learning, as proposed by IL-ETransformer, is rarely adopted. Without continual updates, models become outdated in rapidly evolving markets.
Interpretability and explainability	Refs 1, 9, 11, 15, 20, 23	Complex models hinder transparency. Explainable deep learning (Ref 15) uses SHAP/LIME but at the cost of accuracy. Ensembles and hybrid models (Refs 9, 11) are often black boxes. LLM-augmented models (Ref 23) amplify interpretability issues despite generating human-readable descriptions. Lack of interpretability reduces trust and may violate regulatory requirements.
Hyper-parameter sensitivity and optimization	Refs 2, 6, 12, 20	Model performance can vary greatly with hyper-parameters. Dandelion optimization and blood-coagulation algorithms (Ref 6) and genetic algorithms (Ref 8) attempt automatic tuning but increase complexity. Attention-variant LSTM (Ref 12) and fuzzy control (Ref 20) depend on delicate parameter choices. More systematic optimization strategies are needed.

Privacy and decentralization	Ref 25	Only one paper addresses privacy through federated learning. Most studies assume centralized data access, which is unrealistic for institutions with confidentiality constraints. Frameworks for collaborative learning without sharing raw data remain under-explored.

Methods and Experiments Design

To remedy these gaps, we suggest a comprehensive research agenda employing methodological advances and experimental methodologies. Proposed methods incorporate surveyed papers' ideas and address their shortcomings.

Table 3: Lists Experiments and Methods

Proposed method	Motivation & relation to gaps	Experimental design
Time and Frequency hybrid modelling (e.g., FT-iTransformer, StockMixer with ATFNet)	Integrate short-time Fourier transform (STFT) or wavelet transforms with temporal convolution and transformers to capture both periodic and abrupt market signals. Addresses incomplete feature integration and improves signal representation.	Implement STFT on price series, extract frequency-domain features, and combine them with raw time-series inputs in a dual-branch neural network. Compare to pure LSTM and transformer baselines using metrics such as MSE, directional accuracy and Sharpe ratio.
Sentiment and macro-economic fusion	Augment price data with textual sentiment from news or social media and macroeconomic indicators (GDP growth, inflation, interest rates). Build on the FinBERT-LSTM framework and ensemble models.	Collect sentiment scores via pre-trained models (e.g., FinBERT) and macro indicators from publicly available sources. Use feature selection algorithms (genetic programming or dandelion optimisation) to identify relevant variables. Train hybrid models and assess incremental gains over price-only models.
Incremental and online learning	Adopt continual learning strategies (elastic weight consolidation, replay buffers) to update models with streaming data. Mitigates concept drift and supports real-time forecasting.	Simulate data streams; implement online versions of transformers or RNNs; evaluate performance degradation over time and model adaptability. Compare to batch-trained models.
Hybrid ensembles with interpretable components	Combine diverse learners (e.g., CNN, LSTM, transformer, XGBoost) and incorporate interpretability modules such as SHAP or attention visualisation. Address the accuracy–interpretability	Design stacking or voting schemes that weight models based on predictive power and interpretability. Use SHAP to rank feature importance and examine how different learners focus on various

	trade-off	signals.
Privacy-preserving federated learning	Enable cross-institutional training without sharing raw data. Addresses data scarcity and privacy concerns.	Implement a federated averaging or adaptive optimisation scheme (e.g., FCFL) across multiple synthetic clients. Measure convergence speed, prediction error and communication cost. Compare to centralised training.

Experimental Illustration on Synthetic Data

The impact of controlled modeling choices was proven in a brief experiment utilizing phony stock price data. To replicate a realistic upward moving price series, a geometric Brownian motion (GBM) process with 0.05 % drift and 1% volatility was simulated for 1000 trading days Technical indicators including daily returns, five-day moving average, momentum (current minus lag 5 price), and rolling volatility were calculated from the simulated series. These factors predicted the next-day closing price sets.

Three scikit-learn machine learning models were evaluated:

1. The simple baseline Linear Regression assumes a linear relationship between characteristics and next day pricings.
2. Decision tree ensemble with non-linear interactions: Random Forest Regressor.
3. MLP Regressor: One-hidden-layer feed-forward neural network.

Synthetic datasets were chronologically separated into training (70%) and test (30%) groups. We standardized features and trained models without hyper parameter tweaking. Table 4 provides test set MSE, RMSE, and MAE Sets. Scope accuracy (predicted up/down moves) was 50% across models and omitted for brevity in this process.

Table 4: Performance Analysis of Existing Models

Model	MSE	RMSE	**MAE**
Linear Regression	2.6	1.61	**1.27**
Random Forest	1154.69	33.98	**23.2**
MLP Regressor	59.34	7.7	**5.01**

MLP and random forest had more mistakes than simple linear regression. The price series trended smoothly in this simulation, and linear features-target linkages were sufficient. Complex untuned models overfit noise. Real markets with heteroscedasticity, structural breaks, and complicated relationships require advanced architectures like LSTMs, transformers, and hybrid ensembles. Simpler models may operate given limited data and highly correlated features.

Discussion

Implications of Knowledge Gaps

The reviews demonstrate that proposed models are too sophisticated for data and experiments. Very small or homogeneous datasets often display novel designs, limiting generalization sets. Without macroeconomic indicators and sentiment, models may develop false connections. Despite promising attempts such IL ETransformer, long-term forecasting and online adaptability are understudied. While accuracy compromises interpretability, explainable models like XAI DNN and fuzzy control show transparency. Federated learning only partially solves privacy sets.

Methodological Recommendations

Consider these to boost stock prediction research:

1. Multi-modal, rich data. Researchers should gather price history, trade volume, technical indicators, corporate fundamentals, macroeconomic indicators, news, and social media sentiment. Consortiums and open repositories alleviate data scarcity.
2. Hybrid feature extraction. STFT, wavelets, NoGraphMixer, and huge language models are recommended for textual descriptions. Attention mechanisms evaluate modality importance.
3. Continuous, federated learning. To avoid catastrophic forgetting, update models slowly with new data. FCFL federation learning frameworks provide cross-institutional collaboration and privacy.
4. Interpretability independent of models. Attention visualization, SHAP, and LIME are needed for prediction models. For user trust and regulatory compliance, hybrid ensembles need interpretable components.
5. Thorough analysis and replication. Provide unambiguous train–test splits, rolling window assessments, baseline comparisons, and code and data. Write hyper parameter adjusting methods. Evaluate structural break resilience and long-term forecasting.

Conclusion

This study discusses stock market forecast advances and challenges. Deep learning architectures LSTM variations, CNN transformer hybrids, reinforcement learning agents, and time–frequency networks have improved incrementally, but small datasets, missing features, lack of real-time updating, and interpretability limit them in process. A holistic research agenda comprises hybrid feature extraction, sentiment and macroeconomic integration, incremental and federated learning, and interpretable ensemble modeling sets. A synthetic experiment suggests that simple models may outperform elaborate ones when the signal is smooth and noise is low. To bridge laboratory performance and real-world application, future research should emphasize data variety, realistic evaluation methods, and clear algorithms.

References

1. Md Sazzad Hossain, Md Arman Rahman, & others. (2025). Stock Market Prediction Using Machine Learning and Deep Learning Techniques: A Review. *Journal of Risk and Financial Management.*
2. Gian Carlo, Fan Sahin, & Altay Yalcinkaya. (2024). Forecasting stock price changes using long-short term memory neural network with symbolic genetic programming. *Scientific Reports.*
3. Qi Wei, Lei Zhang, & Ning Wu. (2024). Galformer: a transformer with generative decoding and hybrid loss function for stock prediction. *Scientific Reports.*
4. Bin Guo, Shuo Li, & others. (2025). Multifactor prediction model for stock market analysis based on deep learning techniques. *Scientific Reports.*
5. Shuo Li, Xiang Wang, & others. (2025). Stock market trend prediction using deep neural network via chart analysis. *Humanities & Social Sciences Communications.*
6. Sanjay Sharma & Abhijeet Mandal. (2024). Enhanced stock market forecasting using dandelion optimisation-driven 3D-CNN-GRU classification. *Scientific Reports.*
7. A. Agarwal & P. Kundu. (2025). Deep neural network approach integrated with reinforcement learning for forecasting exchange rates. *Scientific Reports.*
8. Andrey Petrov & Yevgen Bondarenko. (2024). Stock Market Prediction with RNN-LSTM and GA-LSTM. SHS Web of Conferences. *Abstract*
9. Hua Yang & Meng Sun. (2025). Hybrid Machine Learning Models for Long-Term Stock Market Forecasting: Integrating Technical Indicators. *Journal of Risk and Financial Management.*
10. Rafal Kizys, Miklos Jakab, & others. (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.*
11. Diego Rojas & Paula Rodriguez. (2025). Stock Price Prediction Using a Stacked Heterogeneous Ensemble. *International Journal of Financial Studies.*
12. Sora Kim & Jae Lee. (2024). A Novel Variant of LSTM Stock Prediction Method Incorporating Attention Mechanism. *Mathematics.*
13. Nabil Benjelloun, Ihsane Jourani, & others. (2024). Assessing the Predictive Power of Transformers, ARIMA, and LSTM in Forecasting Stock Prices of Moroccan Credit Companies. *Journal of Risk and Financial Management.*
14. Mehdi Ghobadi & Reza Rastegar. (2025). An enhanced Transformer framework with incremental learning for online stock price prediction. *PLOS ONE.*
15. Md Sifat M Hossain & Md Nazmul Hasan. (2024). An explainable deep learning approach for stock market trend prediction. *Heliyon*.
16. Lin Guo, Xue Zhang, & others. (2024). Analyzing the critical steps in deep learning-based stock forecasting: a literature review. *PeerJ Computer Science.*
17. Zhe Lu & Ying Huang. (2024). A study of hybrid deep learning model for stock asset management. *PeerJ Computer Science.*
18. Sheng Chen, Hongwei Liu, & others. (2025). Research on deep learning model for stock prediction by integrating frequency domain and time series features. *Scientific Reports.*
19. Xinyi Chen, Ru Ouyang, & others. (2025). AI-Enhanced Forecasting Framework: Integrating LSTM and Transformer-Based Sentiment for Stock Price Prediction. *Journal of Economic Analysis.*

20. Saurabh Singh & Apurva Patel. (2025). MF-ConvLSTM-XAI model integrating multi-feature and fuzzy control for financial time series forecasting. *Scientific Reports.*
21. Yong Liu, Lei Li, & others. (2024). CNN-Trans-SPP: A small Transformer with CNN for stock price prediction. *Electronic Research Archive.*
22. Boyan Zhang, Tian Zhao, & others. (2026). FT-iTransformer: A Stock Price Prediction Model Based on Time–Frequency Domain Collaborative Analysis. *Entropy.*
23. Liyun Wang & Yuhui Liu. (2025). LLM-Augmented Linear Transformer–CNN for Enhanced Stock Price Prediction. *Mathematics.*
24. Shanmuga Sundari & R. Jahan. (2025). Advancing stock price prediction through the development of hybrid ensembles: a comprehensive comparative analysis of machine learning approaches. *Journal of Big Data.*
25. Tong Wang & Xin Lu. (2025). Fast-Converging and Trustworthy Federated Learning Framework for Privacy-Preserving Stock Price Modeling. *Electronics.*