

Original Article

Hybrid Machine Learning Architectures for Short-Term Financial Forecasting Under Dynamic Market Conditions

Sri Veera

Research Scholar, India.

Received: 03-10-2026

Revised: 04-02-2026

Accepted: 04-07-2026

Published: 04-11-2026

ABSTRACT

Financial forecasting is difficult in the short run due to the nonlinear relationships, volatility clustering, structural changes, and unpredictability in financial markets. This paper presents a hybrid machine learning approach combining Long Short-Term Memory (LSTM) networks, attention mechanism based on transformer, gradient-boosted decision trees (GBDT) and statistical volatility modelling for enhanced forecast adaptability. It takes historical prices, trading volume, technical indicators, and other market variables to process and uncover temporal patterns and shifting market regimes. The key idea is to combine individual model predictions based on market conditions in an adaptive gating mechanism, thus capturing the sequential dependency, nonlinear interaction, and risk-related information. The methodology features partitioning of chronological data, walk forward validation, benchmarking and performance metrics such as mean absolute error, root mean square error and direction accuracy. Ablation and sensitivity analysis are suggested to determine contributions of individual components. The framework facilitates research in financial forecasting, with a focus on data integrity, computational efficiency and realistic assessment. It needs to be empirically tested on a variety of assets, market regimes and out-of-sample periods.

KEYWORDS

Financial Time-Series Forecasting, Hybrid Machine Learning, Long Short-Term Memory (LSTM), Transformer, Gradient-Boosted Trees, Market-Regime Detection, Volatility Forecasting, Ensemble Learning, Walk-Forward Validation, Nonstationary Data.

1. INTRODUCTION

1.1. Background

Investor expectations, economic indicators, liquidity, geopolitical and market events, market sentiment and information flows are all factors that affect the operation of financial markets. The short-term financial forecasting process is difficult due to the nature of the markets, which often shows nonlinear relationships, volatility fluctuations, and changing statistical patterns. Various other methods, such as autoregressive models [1] and generalized autoregressive conditional heteroskedasticity (GARCH) [2] are available to capture temporal dependence and estimate the conditional volatility. But when the financial relationships are highly non-linear or when the market conditions shift rapidly, these methods may be restricted. With the advancements of machine learning, some more flexible forecasting methods have also been developed, such as Long Short-Term Memory (LSTM) networks and Transformer architectures, which have been able to model complex temporal interactions and sequential dependencies, as seen in [3]–[5]. Combining these techniques with statistical models and gradient-boosted decision trees offers opportunities to improve predictive adaptability. Thus, hybrid forecasting architectures are becoming increasingly important for the design of more powerful financial prediction systems that are able to adapt to the dynamic nature of the market, while preserving the reliability of the validation and computational efficiency.

1.2. Key Challenges in Short-Term Financial Forecasting




Fig 1 - Key Challenges in Short-Term Financial Forecasting

1.2.1. Market Nonstationarity

Financial markets are dynamic as a result of economic developments, investor actions and unforeseen events. The return, trading volume and volatility characteristics of the product can therefore change over time. The performance of a trained model may be affected if the market conditions change drastically.

1.2.2. Complex Nonlinear Dependencies

The important components affecting the short term price movements are historical returns, momentum indicators, liquidity, volatility, and market sentiment. Those relationships may be

complex and can not be represented in traditional linear models. In this case, advanced machine learning techniques are needed to detect non-linear patterns and temporal dependencies.

1.2.3. Market Noise and Uncertainty

Often, financial time series have random fluctuations, irregular price movements and unpredictable reactions to new information. These features make it hard to tell if there is a real signal of interest or if it's just noise. Such forecasting models need to be stable in the face of uncertainty in the market and have minimal forecasting errors.

1.2.4. Heterogeneous Financial Information

Financial forecasting uses a variety of data, such as historical prices, trading volume, technical indicators, macroeconomic data and textual sentiment. These sources are not all available, have different time resolution, and are not all reliable. Properly integrating them demands proper pre-processing, time-stamping and feature selection.

1.2.5. Data Leakage and Evaluation Reliability

Data partitioning, feature building based on future information, and normalization done improperly may lead to information leakage in forecasting experiments. Such techniques can result in unrealistic performance projections which do not extrapolate to new markets. Therefore, splitting data chronologically and validating the model on a walk-forward basis are necessary for proper model evaluation [6].

1.2.6. Need for Adaptive Hybrid Forecasting

A single-model approach may not be reliable over market conditions, as individual forecasting models may vary. The combination of statistical modelling, deep learning and nonlinear feature learning can offer complementary representations that predict. An adaptive hybrid architecture is thus envisioned in order to enhance the robustness of forecasting, while simultaneously guaranteeing a high level of evaluation and computational efficiency.

1.3. Research Objectives

The main goal of this research is to build an accurate and reliable hybrid machine learning model with the ability to predict short-term financial variables in the dynamic market. The proposed framework is designed to address the drawbacks of any single forecasting model, by bringing together statistical methods, deep learning architectures, and adaptive ensemble mechanisms in a single predictive system. The first is to integrate the attention mechanism, inspired by the Transformer architecture, with LSTM networks, enabling the model to process both short-term and long-term temporal relationships in financial time-series data. The second goal is to combine the deep learning predictions with the gradient-boosted decision tree and statistical volatility estimation to capture the nonlinear interaction of features and a time-varying degree of market uncertainty. The third goal is to embed information about the market regime into an adaptive weighting method which weights the contribution of the individual models based on the observed characteristics of the

market and the past forecasting performance. The fourth goal is to develop a repeatable and leak-proof assessment technique based on chronological data partitioning, expanding window walk forward evaluation and uniform benchmark comparisons.

The fifth goal is to assess the proposed architecture with complementary performance metrics such as mean absolute error, root mean square error, directional accuracy, volatility estimation error, computational latency, and risk-sensitive indicators. The measures are designed to give an overall evaluation of the predictive accuracy, efficiency and reliability of forecasting. The sixth objective is to perform an ablation and sensitivity analysis to evaluate if any components make a measurable difference for various market conditions and forecasting periods. In addition, the study aims to investigate the effect of market nonstationarity, data quality, model complexity and hyperparameter optimization on generalization performance. Special focus is given to identifying those statistical forecasting enhancements that are not necessarily economic in terms of transaction costs, slippage, turnover and investment risk. The main concern is the one-step-ahead and short-horizon return forecasting, and the price-level prediction is regarded as an ancillary one. In general, these goals provide a structured approach to the design, assessment and testing of an adaptive hybrid forecasting system without the need to rely on the belief that the more complex the model, the more accurate the forecasts or profitable the service.

2. LITERATURE SURVEY

2.1. Statistical Time-Series Models

Financial forecasting is fundamentally driven by statistical time-series models, which can detect temporal dependencies, and help determine the dynamics of market fluctuations. In stationary time series, after differencing, autoregressive integrated moving average (ARIMA) models can be used to detect linear relationships. They are less likely to be able to capture nonlinear relationships, rapid shifts in markets, and intricate relationships between financial variables, however. The autoregressive conditional heteroskedasticity (ARCH) model provides an estimate of the time-varying volatility and the generalized ARCH (GARCH) model is an extension of the ARCH model that adds previous squared errors and conditional variance to enhance volatility estimation [1, 2]. These models can be used in financial risk assessment, uncertainty estimation and position sizing. However, it is important to consider their role in return prediction individually and not expect that they automatically increase the accuracy of the prediction.

2.2. Recurrent Neural Networks and LSTM Architectures

Recurrent neural networks (RNNs) can take the form of time series, such as a sequence of images, and encode information from previous time steps into hidden states. But for learning long-term dependencies, conventional RNNs can suffer from a vanishing gradient problem. The limitation is overcome by the use of gated memory mechanisms in Long Short-Term Memory (LSTM) networks, which control the addition, removal, and output of information [3]. LSTM networks can be used to learn the temporal relationship between the historical returns, the price, trading volume, and technical indicators in financial forecasting. There have also been attempts to use hybrid methods

that integrate both statistical volatility estimation and LSTM learning to enhance volatility prediction in volatile markets [4]. The performance of LSTM, however, is affected by the quality of the data, the length of the sequences, feature scaling and the training-window selection, making careful validation very important.

2.3. Transformer-Based Financial Forecasting

Transformer architectures employ self-attention layers to determine relationships between various positions in a sequence. This ability allows them to capture long-range temporal relationships and complex relationships between several financial variables without using a recurrent process exclusively [5]. As a result, Transformers are more and more studied to forecast market returns, trading volumes, volatility indicators, and many other market characteristics. Hybrid Transformer-based architectures can incorporate attention mechanisms along with recurrent networks or statistical volatility models to model temporal patterns and market uncertainty. The financial time-series data, however, might be noisy and provide few clues for prediction, thus adding to the risk of overfitting. Hence, the performance of the Transformer is compared with simpler models and the following factors are taken into consideration: dataset size, computational cost, forecasting horizon, and out-of-sample reliability.

2.4. Hybrid Learning, Market Regimes and Research Gaps

Hybrid Learning combines complementary modeling methods to provide a model to represent various aspects of financial data. LSTM networks can capture sequential dependencies, Transformers can capture longer-range dependencies, and gradient-boosted trees can capture nonlinear relationships in engineered features. Volatility estimates from the statistics can offer extra details on altering danger conditions and the weighting of each model could be changed dependent on market behavior as a result of the regime of the market [7] and [8]. While there are many benefits to hybrid architectures, they also come with some drawbacks in terms of model complexity, the need for redundant representations, and the extra computational power required. Ensemble weights can be ineffective when the market shifts, and data leakage in the pre-processing phase can give false results for the performance. In addition, it is not guaranteed that smaller forecast errors lead to profitable trading strategies once transaction costs and risk restrictions are taken into account. Thus, this research aims to explore the effectiveness of the proposed framework in different market regimes based on adaptive model weighting, component-wise ablation, and chronological out-of-sample evaluation.

3. METHODOLOGY

3.1. System Architecture and Forecasting Workflow

The suggested Hybrid Regime-Adaptive Financial Forecasting (HRAFF) approach unites the data preparation, features engineering, estimation of market regimes, parallel prediction, adaptive fusion, with the evaluation of forecast performance in a single forecasting process. In the first phase of the data acquisition stage, OHLCV (open, high, low, close, volume) observations, important market indicators and other financial data that can be verified with a certain availability timestamp

are collected. In the pre-processing stages, missing data is dealt with correctly, the time granularity is standardized, corporate actions are taken into account and normalization parameters are estimated solely based on training data, so as to avoid information leakage. Later on, the feature engineering produces informative representations based on information at every point of the forecast, such as historical returns, momentum indicators, attributes of trading volumes and measures of volatility. The framework then estimates market conditions, and processes the features prepared as described above through four complementary predictive components: a Long Short-Term Memory (LSTM) network for capturing sequential dependencies, a Transformer encoder for long-range temporal interactions and a gradient-boosted decision tree for capturing nonlinear feature relationships, as well as a statistical volatility estimator for capturing market uncertainty. They combine their outputs via a regime-aware fusion mechanism that reflects the characteristics of the market observed at the time and previous experience in validation.

Such an adaptive weighting scheme is designed to deal with the varying market conditions without assuming that any predictor will be the best predictor in all of the periods. The output from the fused is the desired short-term economic outlook along with appropriate uncertainty estimates or risk estimates (where supported by the component models). Finally, the evaluation stage is a chronological walk-forward validation to verify the forecasting performance on observations not included in the model training and selection. Evaluation metrics can be mean absolute error, root mean square error, directional accuracy, volatility-estimation error and computational latency. There is also a monitoring of model performance for deterioration when market characteristics change. All of the predictive components of the workflow should share the same information boundary for historical data, and feature selection, hyperparameters tuning, scaling, and estimation of the weight for the ensemble should not use future observations. This organized structure provides a replicable approach to studying the feasibility of increasing forecast reliability when using a combination of statistical-based, deep learning and tree-based forecast models under financial market dynamics that can be adapted to different regimes.

3.2. Data Preparation and Feature Engineering




Fig 2 - Data Preparation and Feature Engineering

3.2.1. Data Representation

Financial information is presented in terms of adjusted closing prices, trading volume and other market factors that are relevant and gathered in regular time intervals. Log returns are used to represent relative price changes, and to facilitate comparisons across price levels, historical price observations are converted into log returns. The forecasting target is the expected return over a given period of time in the future, for example, the next trading interval or the next trading day. Each forecasting horizon is evaluated separately, as the accuracy and uncertainty of the forecasts could differ for each horizon.

3.2.2. Feature Construction

Feature engineering includes generating new variables from financial observation that represent time series, market trends, trading activity, and volatility. Some of the features that are being considered include rolling averages of returns, standard deviations, relative strength indicators, price ranges, percentage change in volume, and estimates of realized volatility. Additional variables such as market-index returns, cross-asset measures, calendar measures and lagged macroeconomic variables can also be added where they are available at the forecast time and can be confirmed. All features are computed from data contained only at the forecast origin so that the forecast is not influenced by data from subsequent observations.

3.2.3. Preprocessing and Leakage Prevention

To maintain the time sequence of financial observations, the data set is split into training, validation, and testing sets chronologically. The missing value problem, feature selection and normalization, and other data transformations are fitted from the data used to train the model and applied to future data without relying on future information. Special consideration is given to overlapping forecast timelines, published economic data that might change, and adjustments to economic data that are backcasts. These preprocessing controls enhance the reliability, reproducibility and even fairness of the forecasting evaluation.

3.3. Hybrid Model Construction

The suggested framework is Hybrid Regime-Adaptive Financial Forecasting (HRAFF), which is a mixture of three complementary return-prediction models and a statistical volatility estimator that can be used to account for various features of financial time-series data. The Long Short-Term Memory (LSTM) network handles sequential feature vectors and has three gates: input, forget, and output, which help control the flow of information through the network and update the network's memory. This is a mechanism that allows historical returns, trading volume, momentum, and volatility features to be learned as local temporal dependencies within the model, enabling it to predict the returns from its learned hidden representation. Second, the Transformer encoder uses self-attention to capture relationships among observations at different points in the input sequence, enabling the model to capture longer-range temporal relationships and interactions between multiple financial variables. Positional information is used to maintain the sequence of observations, and temporal controls such as causal masking or the same type of control make the model unaware of

past or future information. Third, a gradient-boosted decision tree model learns the nonlinear relationship between engineered financial features by an ensemble of sequentially built decision trees. This component also offers an alternative tabular-learning view and checks if there is value in more complex sequence models than manually-designed features.

The fourth component is a statistical volatility estimator model, which is based on generalized autoregressive conditional heteroskedasticity (GARCH) framework, a framework that estimates conditional variance by previous innovations and volatility estimations. Its output does not necessarily represent a direct prediction of return direction, but rather a description of the shifting market risk that can be used to estimate the next regime or to weight the various models as a component of adaptive model weighting. Every component is trained with fixed chronological information boundaries and hyperparameters are tuned using the validation data, which comes before the final test period. Where necessary, regularization is applied, early stopping is used and features are preprocessed in order to minimize overfitting and generalize the results. The three return predictors produce individual forecasts, whereas the volatility predictor provides additional information on the uncertainty of the forecasts that is useful in the coordination phase. They, in turn, are supplied to a regime-sensitive fusion output model that calculates their contributions based on observable market conditions and performance validated against the model. This is a modular design that allows component-wise ablation and comparison, which will ensure that the study can determine if the incorporation of recurrent learning, attention based representation learning, tree based prediction, and statistical volatility estimation mechanisms improve the accuracy of forecasting, adaptability, and computational efficiency.

3.4. Regime-Aware Fusion, Training, and Evaluation

The framework proposed with HRAFF includes the estimation of the market regime, the use of adaptive ensemble weighting, regularized training and chronological evaluation to enhance the forecast adaptability in the face of changing financial conditions. The estimation of the market regimes is made with the help of historical returns, signals of volatility, and other market characteristics available in the market. The probabilities produced by the regime classifier can represent the likelihood of the various market states and not merely be based on strict classifications. To avoid information leakage, regime estimation only incorporates observations obtained at the state at which the forecast is made and not the volatility estimates for future periods nor retrospective labels. The adaptive fusion mechanism fuses the return forecasts from the LSTM, Transformer and gradient-boosted tree models with non-negative weights that sum to one. These weights are based on an economic gate that reflects some level of estimated market conditions and the performance of the models based on a set of indicators derived from past forecast results. Therefore, the importance of each predictor will also vary with market changes.

The mechanism for the fusion is learnt using an appropriate forecasting loss (e.g. Huber loss), and also regularization is used to lower the sensitivity of the mechanism to noise and overly complex models. The model parameters and ensemble weights are determined by the training and validation

data sets which are temporally separated. Performance assessment uses expanding-window or rolling-window, walk-forward validation, where models are trained on past observations and tested on future validation periods. Experiments are conducted using comparative models such as ARIMA, GARCH models for volatility estimation, standalone LSTM, Transformer, gradient-boosted trees and fixed weight ensemble. The forecast accuracy is evaluated by mean absolute error, root mean square error, and directional accuracy, whilst the computational training time, inference latency and stability under various market regimes are also noted. If trading simulations are provided, transaction costs, turnover, maximum drawdown and risk adjusted performance are taken into account. The same test periods and the same information boundaries are held in all models to make the comparisons fair. Finally, ablation experiments are conducted to evaluate the contribution of the different parts of the prediction, regime estimation and adaptive weighting. This evaluation design allows for a systematic study of performance gains in predictive accuracy, robustness, and computational efficiency, when using regime-aware fusion in dynamic market conditions.

4. RESULT AND DISCUSSION

4.1. Experimental Design and Evaluation Criteria

The experimental design of the proposed Hybrid Regime-Adaptive Financial Forecasting (HRAFF) framework is intended to address the question of whether the integration of features of recurrent learning, Transformer-based attention, gradient-boosted trees, and market-regime information is able to improve short-term financial forecasting under dynamic market conditions. However, due to the lack of empirical data and actual model results, the evaluation is shown as an outline of a research process, not as an actual experiment. The study can start with observations on a liquid equity market, for example, S&P 500 or NIFTY 50, and then be expanded to the foreign exchange or cryptocurrency market to see if it's generalizable. The data set should include adjusted prices, trading volume if present, and other market parameters (where available) that have been published with a known publication date. Forecasting can be done on a 1-day or 5-day basis; an initial historical input window of 60 observations may be used to begin with. The data is split in chronological order into training, validation, and testing sets, and expanding-window (EW) validation is used to maintain chronological order. The proposed HRAFF framework introduces regime-aware adaptive fusion in baseline models such as ARIMA, stand-alone LSTM, Transformer, gradient-boosted trees, and a fixed-weight ensemble.

Hyperparameters, preprocessing transformations and ensemble weights are chosen without using the final test period. Performance is assessed through mean absolute error (MAE), root mean square error (RMSE), and directional accuracy and by the volatility-forecasting loss and inference latency. Average absolute prediction errors are captured by MAE, while larger error in predictions are more strongly penalized by RMSE; directional accuracy will be evaluated by seeing if the predicted return sign is the one observed. These metrics are complementary, as a right direction prediction does not imply the right size prediction. In addition, performance is evaluated by the low volatility, high volatility, trending and transitional markets to determine the strengths and weaknesses in each regime. Use of statistical significance tests and uncertainty intervals should be

used to determine whether observed difference is significant. Where trading simulations are performed, transaction cost, turnover, drawdown and risk adjusted returns should also be considered. Any numerical comparison made before completion of these experiments should be expressed in such a manner that it clearly shows that it is an illustration of the comparison, not evidence. This methodology provides a repeatable framework for evaluation of the forecasting accuracy, robustness and computational efficiency of the proposed hybrid architecture, without making any assumption about the superiority of this architecture over simpler alternatives.

4.2. Illustrative Comparative Results

Table 1: Illustrative Comparative Performance of Forecasting Models

Forecasting Model	Directional Accuracy (%)	MAE Reduction (%)	RMSE Reduction (%)
ARIMA	52	0	0
Gradient-Boosted Trees	56.5	8	7.5
LSTM	58.2	12.5	11
Transformer	59.1	15	13.5
Fixed-Weight Hybrid	60.4	18	16.5
Proposed HRAFF	62.1	23	21




Fig 3 - Illustrative Comparative Performance of Forecasting Models

4.2.1. Comparative Model Performance

The comparative evaluation is used to compare the forecasting performance of ARIMA, gradient-boosted trees, LSTM, Transformer, fixed-weight hybrid and the proposed HRAFF model. The measure of the direction accuracy is the percentage of correct return directions and the relative reductions in MAE and RMSE are the impacts of the model on the prediction errors relative to the ARIMA model. The numbers given in Table I are totally hypothetical and are used to illustrate the reporting format only and not the actual experimental results. It is imperative that all models be

tested with the same chronological test intervals and the same preprocessing methods in a completed study.

4.2.2. Interpretation of Illustrative Findings

The illustrative comparison reveals how individual models may be compared to fixed and adaptive ensembles. The hypothetical results indicate that the addition of complementary forecasting elements may help to improve the accuracy of the direction and reduce the errors. The values are, however, merely hypothetical and do not prove the superiority of HRAFF as there is no actual data, model training or statistical testing. Only true conclusions should be based on measured results over several evaluation windows and market conditions.

4.2.3. Relative Error Reduction Analysis

Relative MAE reduction is a percentage reduction (increase) of the mean absolute error of a model relative to the chosen reference model. Likewise, relative RMSE reduction indicates the same reduction in the root mean square error, but is more sensitive to large prediction errors. Positive reductions mean errors lower than that of the reference, whereas negative reductions mean errors higher than in the reference. These measures should be derived from actual test predictions, not evaluated to discern the comparative forecasting effectiveness.

4.2.4. Directional Accuracy and Forecasting Reliability

When it comes to direction, don't take direction accuracy and MAE or RMSE together as they don't necessarily go hand in hand. Similarly, a model with smaller prediction errors does not necessarily yield better trading results when taking transaction costs, slippage, and portfolio risk into account. So one should combine error measures with directional performance, statistical significance measures with regime-specific analysis. All numerical results should be replaced with the verified experimental measurements prior to submission of manuscript to publication.

4.3. Discussion, Robustness, and Practical Implications

The Hybrid Regime-Adaptive Financial Forecasting (HRAFF) framework is a new hybrid of complementary predictive representations to cope with the complexity of short-term financial forecasting in changing markets. The Long Short-Term Memory (LSTM) network is used to model sequential relationships, the Transformer is used to model the context between historical observations, and the gradient-boosted trees are used to model nonlinear interactions between engineered financial features, while the statistical volatility model is used to provide more information about the changing market risk. But the usefulness of the integration relies on the fact that both the individual models provide enough different prediction information. If they are highly correlated, the errors of their predictions may be lessened by combining them. Thus, the respective contributions of LSTM, Transformer, volatility estimator, regime information and adaptive weighting mechanism should be analyzed in ablation experiments. When models' reliability varies during calm, volatile, trending, and transitional market situations, regime-aware fusion might enhance adaptability.

However, no market conditions are directly observable, and misclassification can result in weight adjustments or predictions that fluctuate. The gating mechanism should therefore contain proper temporal valid performance estimates, appropriate regularization and sensitivity analysis to avoid too many changes to the ensemble weights. Evaluating robustness should not be done solely based on aggregate accuracy, but rather through a number of chronological evaluation windows, several forecasting horizons, and individual market-regime analyses. Moreover, any gains in statistical forecasting need to be separated from the actual advantages to investments. While a lower MAE or higher directional accuracy may not necessarily mean profitability, the profitability of trading can be affected by transaction costs, bid-ask spreads, slippage, impact on the market, trading turnover, position restrictions and portfolio risk. There are types of forecast-comparison tests and the bootstrap confidence intervals that can be used to quantify the reliability of the observed difference in performance. If a trading simulation is performed, the net returns and risk-adjusted measures should be reported in addition to the prediction measures. The computational efficiency is also crucial since complex hybrid models might need more training time, memory, and inference resources.

5. CONCLUSION

The proposed Hybrid Regime-Adaptive Financial Forecasting (HRAFF) framework for short-term financial prediction in dynamic markets was presented in this paper. Long Short-Term Memory (LSTM) networks, Transformer-based attention, gradient-boosted decision trees, statistical volatility estimation, and regime-aware ensemble weighting are explored and incorporated into the architecture in order to investigate the benefit of using complementary predictive techniques in conjunction with one another. The LSTM part models temporal relationships between past observations on financial data, whereas the Transformer models a general temporal relationship between market variables. Gradient-boosted trees detect nonlinear relations between engineered features; statistical volatility estimation gives additional information about volatility of market risk. The adaptive fusion mechanism has to be designed to adapt the contribution of the single return predictors based on the estimated market situation and on the forecasting performance observed in the past. The methodology proposed in the project also focuses on ensuring reliable data preparation, chronological data partitioning, training only-based data preprocessing, and an expanding-window walk-forward validation approach to avoid information leakage and ensure credible out-of-sample performance estimates. In addition, benchmark comparisons, component-wise ablation, regime-specific analysis, statistical significance testing, and computational measurements are added to facilitate a holistic assessment of the predictive accuracy, robustness, and efficiency.

While the proposed architecture's merits are not yet proven in experiment, it was beyond doubt that the benefits of the proposed method would be demonstrated with the availability of a dataset, a trained model, and an evaluation. Therefore the illustrative comparative percentages used in Section IV should not be used as measured results, and must be changed to actual results prior to publication. Future studies should test the framework on various financial assets, forecasting horizons, and market regimes, to examine the generalizability and sensitivity of the framework to shifts in data distributions. Future work may focus on probabilistic forecasting, multimodal financial

information, online learning, algorithmic decision explainability, and efficient algorithms. Transaction cost, bid-ask spread, slippage, turnover, drawdown, and risk-adjusted returns should also be evaluated through practical evaluation to see if there are meaningful economic gains from statistical forecasting improvements. Overall, HRAFF provides a systematic research methodology for studying the integration of complementary machine learning models and the use of adaptive coordination and rigorous evaluation. Whether its ultimate contribution can be replicated, through experimentation, under realistic financial conditions in comparison to simpler forecasting alternatives remains to be seen.

REFERENCES

- [1] R. F. Engle, "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," *Econometrica*, vol. 50, no. 4, pp. 987-1007, 1982.
- [2] T. Bollerslev, "Generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, vol. 31, no. 3, pp. 307-327, 1986.
- [3] Bolla, S. R., Uddandaraio, D. P., Mokashi, S. P., Mistry, H., & Mavani, C. (2025, July). AI-Powered Fraud Detection in Real-Time Payment Systems Excellent Application of ML In Fintech-Critical and Impactful Use Case. In 2025 IEEE 4th World Conference on Applied Intelligence and Computing (AIC) (pp. 1-10). IEEE.
- [4] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed. Hoboken, NJ, USA: Wiley, 2015.
- [5] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
- [6] Brahmandam, L. M. K. (2023). Migrating Mission-Critical Enterprise Workloads from On-Premises VMware to AWS: An Empirical Study of a Multi-Account Landing-Zone Reference Architecture and the Seven Rs Decision Framework. *International Journal of Emerging Trends in Computer Science and Information Technology*, 4(4), 231-240. <https://doi.org/10.63282/3050-9246.IJETCSIT-V4I4P124>
- [7] A. Vaswani *et al.*, "Attention is all you need," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998-6008.
- [8] G. P. Zhang, "Time series forecasting using a hybrid ARIMA and neural network model," *Neurocomputing*, vol. 50, pp. 159-175, 2003.
- [9] Gantikota, S. (2024). Shift-Left Security for Decentralized Engineering Organizations: Embedding SAST, DAST, and Penetration Testing Throughout the Software Development Lifecycle in University and Research Computing Environments. *International Journal of Emerging Research in Engineering and Technology*, 5(4), 175-179. <https://doi.org/10.63282/3050-922X.IJERET-V5I4P118>
- [10] T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," *European Journal of Operational Research*, vol. 270, no. 2, pp. 654-669, 2018.
- [11] Gadam, H., Yadava, A. K., Arora, A. S., Yachamaneni, T., & Kotadiya, U. (2025, July). Revolutionizing Wealth Management with Temporal Attention Mechanism Based Convolutional Network for Algorithmic Stock Trading Strategies. In 2025 3rd International Conference on Data Science and Network Security (ICDSNS) (pp. 1-4). IEEE.
- [12] B. Lim and S. Zohren, "Time-series forecasting with deep learning: A survey," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 379, no. 2194, Art. no. 20200209, 2021.
- [13] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785-794.
- [14] Sunkara, R. (2023). Cost-Optimized Energy Compliance Testing for Smart TV Streaming Devices: Achieving Milliwatt-Precision Power Measurement at Sub-One-Thousand-Dollar per Setup. *American International Journal of Computer Science and Technology*, 5(6), 54-59. <https://doi.org/10.63282/3117-5481/AIJCSIT-V5I6P105>
- [15] J. D. Hamilton, "A new approach to the economic analysis of nonstationary time series and the business cycle," *Econometrica*, vol. 57, no. 2, pp. 357-384, 1989.
- [16] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 3rd ed. Melbourne, Australia: OTexts, 2021.

- [17] K. Olorunnimbe and H. Viktor, "Deep learning in the stock market – A systematic survey of practice, backtesting, and applications," *Artificial Intelligence Review*, vol. 56, pp. 2057–2109, 2023.
- [18] Sandra, K. (2024). THE REGULATED BANKING AI LAKEHOUSE. INDO-CONTINENTAL ACADEMIC PUBLISHERS.
- [19] D. H. Bailey, J. Borwein, M. López de Prado, and Q. J. Zhu, "The probability of backtest overfitting," *Journal of Computational Finance*, vol. 20, no. 4, pp. 39–69, 2017.
- [20] G. E. P. Box and G. M. Jenkins, "Time series analysis: Forecasting and control," *Journal of Time Series Analysis*, vol. 1, no. 2, pp. 97–100, 1980.
- [21] Paruchuri, J. K. (2021). Lakehouse architecture: Unifying data lakes and data warehouses. International Journal of Computer Techniques (IJCT), 8(1), 1-15.
- [22] R. S. Tsay, *Analysis of Financial Time Series*, 3rd ed. Hoboken, NJ, USA: Wiley, 2010.