

A stacking model integrating GARCH and LSTM with feature interactions for time series volatility prediction

Michael Peter ^{a,b},* Silas Mirau ^a, Emmanuel Sinkwembe ^c, Christian Kasumo ^d, Calisto Guambe ^{e,f}

^a Department of Applied Mathematics, The Nelson Mandela African Institution of Science and Technology (NM-AIST), P.O. Box 447, Arusha, Tanzania

^b Department of Mathematics and ICT, The Open University of Tanzania (OUT), P.O. Box 23409, Dar es Salaam, Tanzania

^c Department of Mathematics, University of Dar es Salaam, P.O. Box 35062, Dar es Salaam, Tanzania

^d Department of Mathematics and Statistics, Mulungushi Univ, P.O. Box 80415, Kabwe, Zambia

^e Department of Mathematics, University of Eduardo Mondlane, P.O. Box 257, Maputo, Mozambique

^f Department of Mathematics and Applied Mathematics, University of Pretoria, P.O. Box 0002, South Africa

ARTICLE INFO

Keywords:

Volatility forecasting
Stacking ensemble model
GARCH-LSTM integration
Time series analysis

ABSTRACT

Volatility forecasting remains a cornerstone of quantitative finance, underpinning risk management, portfolio optimization, and regulatory oversight. This study introduces a novel stacking model that integrates the generalized autoregressive conditional heteroskedasticity (GARCH) framework with long short-term memory (LSTM) networks to capture both econometric structure and nonlinear temporal dependencies in financial time series. Unlike conventional hybrid approaches that sequentially cascade outputs, the proposed framework employs GARCH and LSTM as parallel base learners, with their predictions intelligently fused through a meta-learner that exploits feature interactions and cross-model synergies. The empirical evaluation benchmarks the stacking ensemble against state-of-the-art alternatives, including DLINEAR, CKAN, N-BEATS, and individual GARCH and LSTM specifications across multiple performance metrics. Results demonstrate consistent superiority across RMSE, MAE, accuracy, RAMP, geometric mean, Hausdorff distance, and AUC metrics, validating the synergistic benefits of integrating econometric and machine learning paradigms within a theoretically grounded architecture. The model's superior performance stems from leveraging GARCH's parametric efficiency in modeling volatility clustering while harnessing LSTM's capacity to capture complex nonlinear temporal patterns. Beyond methodological contributions, the framework offers practical value for enhancing systemic risk monitoring, improving stress testing frameworks, and optimizing investment strategies across diverse market conditions. The demonstrated robustness across different market regimes underscores its potential for adoption in both routine operations and crisis contexts. This research establishes stacking-based ensemble modeling as a powerful paradigm for advancing volatility prediction and provides a foundation for next-generation financial forecasting systems.

1. Introduction

Accurately forecasting financial time series data, especially volatility, is essential for informed decision-making in risk management, asset pricing, and policy formulation. Volatility serves as a key indicator of market uncertainty and potential investment risk. As financial markets become more complex and interconnected, the demand for reliable and robust volatility forecasting methods continues to grow. According to Chauhan et al. [1], the accuracy of such forecasts directly influences the strategies employed by market participants and regulators alike,

underscoring the critical role of volatility modeling in maintaining market stability and facilitating effective investment decisions.

Recent studies have highlighted the growing intricacies of global financial markets, emphasizing how heightened market interconnectedness has exacerbated volatility transmission across different asset classes and geographical boundaries. For instance, Dong and Huang [2] illustrates this phenomenon by examining the ripple effects of economic shocks and policy changes, which can lead to significant fluctuations in market volatility. This interconnectedness necessitates

* Corresponding author at: Department of Applied Mathematics, The Nelson Mandela African Institution of Science and Technology (NM-AIST), P.O. Box 447, Arusha, Tanzania.

E-mail addresses: peterm@nm-aist.ac.tz (M. Peter), silas.mirau@nm-aist.ac.tz (S. Mirau), sinkwembe.emmanuel@udsm.ac.tz (E. Sinkwembe), ckasumo@mu.edu.zm (C. Kasumo), calisto.guambe@uem.ac.mz (C. Guambe).

<https://doi.org/10.1016/j.array.2026.100700>

Received 30 May 2025; Received in revised form 17 September 2025; Accepted 28 January 2026

Available online 4 February 2026

2590-0056/© 2026 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

sophisticated modeling techniques that can capture the nuances of volatility dynamics in a rapidly changing financial landscape. As such, volatility modeling remains a focal point for both researchers and practitioners, given its profound implications for financial stability, investment strategies, and regulatory oversight.

Over decades, traditional econometric models, particularly the generalized autoregressive conditional heteroskedasticity (GARCH) framework, have been the cornerstone of volatility forecasting. The GARCH model and its variants are tailored to address conditional heteroskedasticity, where the residual variance changes over time. This phenomena exhibits volatility clustering, a typical feature of financial time series data. Bao et al. [3] notes that GARCH models excel in identifying rapid fluctuations and extreme market events, a view supported by numerous studies in the field of financial forecasting. Their widespread adoption in empirical research and practical applications attests to their effectiveness in capturing the underlying volatility patterns in various asset classes. However, despite their strengths, GARCH models are often constrained by their reliance on linear structures, which can limit their capacity to capture the complex, nonlinear dependencies that frequently characterize time series data.

To overcome the shortcomings of traditional econometric models, advanced machine learning techniques have gained recognition as promising solutions for time series forecasting. Among these, long short term memory (LSTM) networks a specialized form of recurrent neural networks (RNNs) have gained significant attention for their capability to model sequential data effectively. For instance Han et al. [4], demonstrate that LSTMs significantly outperform conventional econometric models in capturing long-term dependencies and nonlinear patterns in time series data, rendering them particularly effective for complex, non-stationary datasets. The architecture of LSTMs allows them to retain information over extended periods, making them adept at recognizing patterns that may not be immediately apparent in shorter time frames. This capability is particularly advantageous in the context of financial time series forecasting, where market dynamics can shift rapidly and unpredictably.

Despite their impressive performance in various applications, the use of LSTM networks for volatility forecasting can be further enhanced by integrating their strengths with those of traditional econometric models such as GARCH. Song et al. [5] argues that a hybrid approach that combines the predictive power of LSTMs with the structural advantages of GARCH models can lead to the development of more robust forecasting methodologies. By leveraging the unique strengths of both techniques, researchers and practitioners can create models that are better equipped to handle the complexities of financial data, thereby improving the accuracy of volatility forecasts.

The integration of machine learning techniques with traditional econometric models represents a paradigm shift in the field of financial forecasting. This hybridization not only addresses the limitations inherent in each modeling approach but also opens up new avenues for research and practical application [6–13]. The ability to capture both linear and nonlinear relationships in financial time series data is crucial for developing more accurate and reliable forecasting models. As financial markets continue to evolve, the demand for innovative forecasting methodologies that can adapt to changing market conditions will increase.

This study introduces a stacked ensemble framework where individual models GARCH and LSTM are trained independently, and their predictions are then aggregated using a meta-learner. This stacking approach differs from traditional hybrid models that sequentially link models, such as those that feed GARCH residuals directly as input to LSTM networks. To the best of our knowledge no study has applied such stacked ensemble architectures combining econometric and deep learning models for volatility forecasting in the stock market context.

The main contributions and novel aspects of this study are as follows:

- We propose a unique stacked ensemble framework that treats GARCH and LSTM models as independent base learners, whose predictions are aggregated through a meta-learning approach. This differs fundamentally from existing hybrid models that typically use sequential integration (e.g., feeding GARCH residuals into LSTM networks).
- Our stacking approach enables the meta-learner to capture complex feature interactions between the predictions of econometric and deep learning models, potentially uncovering volatility dynamics that individual models might miss.
- This study represents the first comprehensive application of stacked ensemble volatility forecasting combining GARCH and LSTM models in the context of the Tanzanian stock market, addressing a significant gap in emerging market volatility modeling literature.
- We provide a thorough empirical comparison of individual GARCH and LSTM models against our proposed stacking framework, offering insights into the relative performance gains achievable through ensemble integration.
- The study contributes to the limited literature on advanced volatility forecasting techniques specifically tailored for emerging financial markets, which often exhibit different volatility characteristics compared to developed markets.
- We provide a complete methodological framework that can be readily adapted and applied to other emerging markets, offering practitioners a robust tool for enhanced volatility prediction in similar market contexts.

The remainder of the paper is structured as follows: Section 2 reviews the appropriate related works. Section 3 outlines the data and methodology. The results and discussion are detailed in Section 4. Finally, Section 5 provides the conclusion and suggestions for future research directions.

2. Related works

Accurate and reliable time series prediction remains a persistent challenge in financial econometrics and machine learning research due to the inherent volatility, nonlinear dependencies, and complex dynamics of financial markets. A vast body of literature has explored various approaches for modeling volatility, ranging from traditional econometric models to advanced machine learning techniques [14–16]. This section reviews key contributions from both paradigms and contextualizes the rationale for hybrid modeling approaches, particularly in integrating GARCH models and LSTM networks within a hybrid stacking framework.

Since the work of Engle [17] and later generalized by Bollerslev [18], the GARCH family models has been extensively used for modeling conditional heteroskedasticity in time series data. The strength of GARCH models lies in their ability to capture both volatility clustering and persistence including their robustness in modeling market shocks and extreme events. Several variants of the GARCH model, including the Glosten–Jagannathan–Runkle GARCH (GJR-GARCH) [19] and exponential GARCH models (EGARCH), have been developed to capture asymmetric effects [20], such as leverage effects, which reflect the differing impacts of positive and negative market shocks on volatility.

Traditional GARCH models capture short-term volatility persistence but struggle with nonlinear dynamics and abrupt market shifts. To address these limitations, hybrid approaches increasingly integrate GARCH with machine learning techniques, particularly RNNs and LSTMs [21], which can model complex sequential patterns and long-term dependencies in financial time series. LSTM models effectively capture nonlinear dependencies and complex temporal patterns, making them suitable for highly volatile financial data. They have been applied to predict stock indices [22], exchange rates [23], and commodity prices [24], often outperforming traditional models. Recent

studies emphasize the importance of addressing data non-stationarity to ensure robust LSTM performance.

Building on the importance of addressing data non-stationarity in LSTM modeling, [25] highlighted the need for high-quality training data and proper transformations to ensure reliable forecasts. [26] proposed a validation framework enforcing stationarity through systematic transformations and diagnostic tests, while Livieris and Pintelas [27] extended this to multi-step forecasting, emphasizing that careful handling of non-stationarity is crucial to prevent error accumulation over time.

While LSTM networks effectively capture nonlinear dependencies and long-term dynamics, they remain prone to overfitting and struggle to model conditional heteroskedasticity a strength of GARCH models. This complementarity has driven growing interest in hybrid frameworks, where GARCH models capture short-term volatility clustering and conditional heteroskedasticity, while neural networks learn nonlinear residuals and temporal patterns. Building on advances in GARCH-ANN and GARCH-LSTM research [28–30], the present study proposes a stacking architecture that integrates GARCH and LSTM models with explicit treatment of data stationarity. By combining their respective strengths within a meta learning framework, the stacking model aims to deliver more accurate and robust volatility forecasts.

Hybrid models have shown promise for financial prediction, yet key challenges persist, including explicit treatment of stationarity, architectural complexity, and limited application in emerging markets. Accurate volatility forecasts are vital for risk management and policy decisions, underscoring the need for robust, interpretable approaches. This study advances a GARCH-LSTM stacking framework that integrates proper stationarity handling, leveraging the complementary strengths of econometric and deep learning models to improve forecasting accuracy and practical relevance.

3. Data and methodology

This section outlines the datasets employed in the study, describes the preprocessing steps, and details the methodological framework used to model and forecast stock market volatility, including both traditional and hybrid modeling approaches.

3.1. Data collection and preprocessing

Daily financial and commodity time-series data were sourced from Yahoo Finance [31] and the Bloomberg [32], encompassing diverse asset classes to evaluate model robustness across different market dynamics. The dataset includes equity markets represented by Airtel Africa Plc (AAFL) closing prices in GBP from March 2020 to August 2025 and the Tanzania Share Index (TSI) in TZS from October 2011 to June 2024. Commodity data comprise crude oil (COIL) and natural gas (NGS) futures spanning March 2020 to August 2025, alongside gold (GOLD) futures from July 2007 to June 2024, all denominated in USD. Data preprocessing involved systematic handling of missing values through forward-fill interpolation to maintain temporal continuity, while outliers were deliberately preserved given their critical importance in capturing extreme market events and volatility clustering patterns inherent in financial markets. Each dataset underwent chronological partitioning using an 80/10/10 split for training, validation, and testing sets respectively, ensuring realistic out-of-sample forecasting conditions and preventing data leakage. This comprehensive dataset spanning African equities, East African market indices, energy commodities, and precious metals provides a robust foundation for evaluating model performance across multiple asset classes, each characterized by distinct volatility regimes, market microstructures, and risk-return profiles that reflect the complexities of modern financial markets.

3.2. Methodology

3.2.1. Stacking architecture mathematical formulation

Consider a time series $\{y_t\}_{t=1}^T$ with corresponding feature vectors $\{x_t\}_{t=1}^T$. In this formulation, each $y_t \in \mathbb{R}$ represents the realized volatility of the stock index at time step t , which constitutes our target variable for prediction. Volatility in this context quantifies the magnitude of price fluctuations and serves as a critical metric for financial risk assessment. Similarly, each $x_t \in \mathbb{R}^d$ represents the input feature vector at time step t , comprising historical volatility measurements, logarithmic returns of the stock index, and other relevant financial indicators derived from our preprocessing methodology. These features serve as the explanatory variables that inform our predictive models. The primary objective is to predict future volatility values utilizing a stacking ensemble approach that hierarchically leverages the complementary strengths of econometric and machine learning methodologies.

As illustrated in Fig. 1, following data preprocessing procedures detailed in Section 3.1, these features were implemented in parallel through GARCH and LSTM architectures as base models. The GARCH model captures volatility dynamics through Equation (1)

$$\begin{cases} y_t = \mu_t + \epsilon_t, & \epsilon_t = \sigma_t z_t, \\ \sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2. \end{cases} \quad (1)$$

where μ_t is the conditional mean, $z_t \sim \mathcal{N}(0, 1)$ are standardized innovations, $\omega > 0, \alpha \geq 0, \beta \geq 0$ are parameters, and the condition $\alpha + \beta < 1$ ensures stationarity. The one-step-ahead volatility forecast is given by Eq. (2)

$$\hat{\sigma}_{t+1|t}^2 = \omega + \alpha \epsilon_t^2 + \beta \sigma_t^2. \quad (2)$$

Next, the LSTM processes sequential data through its memory cell architecture as presented in Fig. 2 with the three gates shown in Eqs. (3), (4), and (5)

$$\mathbf{G}_t = \sigma_g(W_f x_t + U_f \mathbf{h}_{t-1} + b_f) \quad (3)$$

$$\mathbf{I}_t = \sigma_g(W_i x_t + U_i \mathbf{h}_{t-1} + b_i) \quad (4)$$

$$\mathbf{O}_t = \sigma_g(W_o x_t + U_o \mathbf{h}_{t-1} + b_o) \quad (5)$$

$$\mathbf{C}_t = \mathbf{G}_t \odot \mathbf{C}_{t-1} + \mathbf{I}_t \odot \sigma_c(W_c x_t + U_c \mathbf{h}_{t-1} + b_c) \quad (6)$$

$$\mathbf{h}_t = \mathbf{O}_t \odot \sigma_h(\mathbf{C}_t), \quad (7)$$

where: σ_g is the sigmoid function, σ_c, σ_h are hyperbolic tangent functions, \odot denotes element-wise multiplication, \mathbf{o}_t represents the output gate in an LSTM cell, W_*, U_* are weight matrices, and b_* are bias vectors (see Appendix).

First, the forget gate \mathbf{G}_t from Eq. (3) computes the weighted sum of x_t, \mathbf{h}_{t-1} , and the bias b_f , producing a value between 0 and 1 through the sigmoid function σ_g . A value of 1 means that all input information is allowed to pass through, while a value of 0 blocks all information. This gate \mathbf{G}_t regulates how much of the previous cell state \mathbf{C}_{t-1} contributes to updating the current cell state, as seen in Eq. (6). Eq. (4) defines the input gate \mathbf{I}_t , which controls how much new information is stored in the current cell state \mathbf{C}_t . The new information at time t is calculated using the tanh function in Eq. (5), which outputs values between -1 and 1 . The past cell state information and the new information, managed by the forget and input gates, are then combined to form the current cell state \mathbf{C}_t at time t , as described in Eq. (6). The output value \mathbf{h}_t is determined by passing through the output gate \mathbf{O}_t and being filtered by the cell state \mathbf{C}_t , which is processed by the tanh function. The selected output values are produced by multiplying them with \mathbf{O}_t . This process updates the cell state \mathbf{C}_{t-1} , distinguishing necessary information from irrelevant data, and the output is \mathbf{h}_t , as shown in Eq. (7). The LSTM model learns through backpropagation using the time algorithm and its prediction at time t is given by Eq. (8) below

$$\hat{y}_t^{LSTM} = W_{out} \mathbf{h}_t + b_{out}. \quad (8)$$

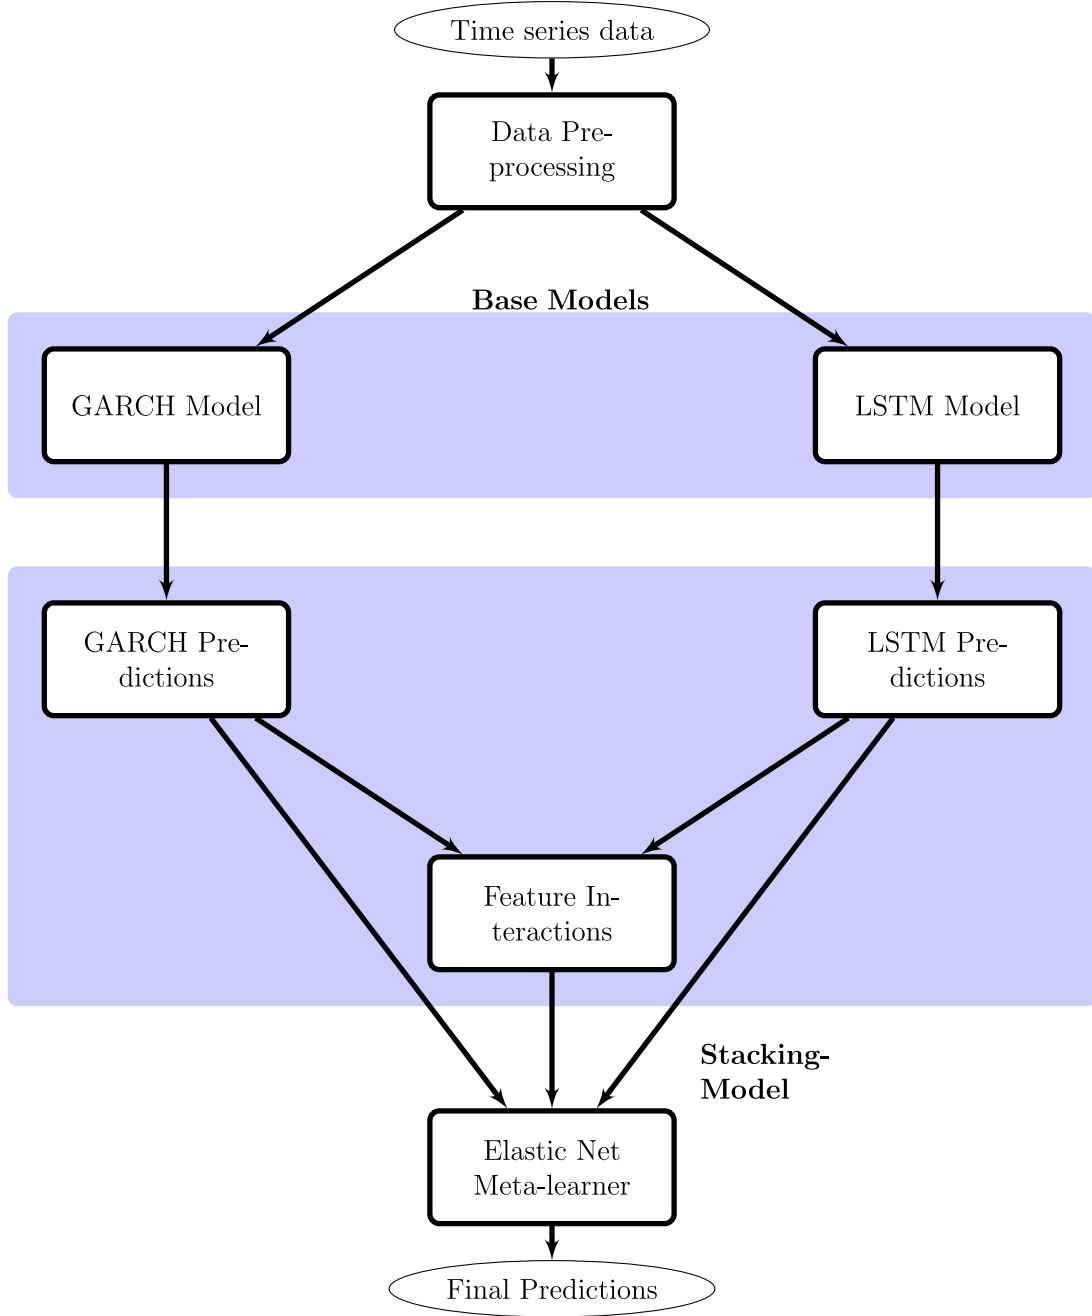


Fig. 1. Stacking architecture for time series prediction.

Next is the meta-features generation step, the meta-features vector \mathbf{z}_t combines base model predictions and their interactions:

$$\mathbf{z}_t = [\hat{\sigma}_t^2 \quad \hat{y}_t^{LSTM} \quad \hat{\sigma}_t^2 \cdot \hat{y}_t^{LSTM}]' \in \mathbb{R}^3. \quad (9)$$

3.2.2. Meta-learner

The elastic net meta-learner combines L_1 (Lasso) and L_2 (Ridge) regularization

$$\hat{y}_t = f_{meta}(\mathbf{z}_t) = \mathbf{w}^T \mathbf{z}_t + b \quad (10)$$

with the aim of optimizing the following optimization problem;

$$\min_{\mathbf{w}, b} \frac{1}{N} \sum_{t=1}^N (y_t - f_{meta}(\mathbf{z}_t))^2 + \lambda \left(\alpha \|\mathbf{w}\|_1 + \frac{1-\alpha}{2} \|\mathbf{w}\|_2^2 \right), \quad (11)$$

where $\lambda > 0$ controls overall regularization strength, $\alpha \in [0, 1]$ balances L_1 and L_2 penalties, $\|\mathbf{w}\|_1 = \sum_{j=1}^3 |w_j|$ is the L_1 norm, and $\|\mathbf{w}\|_2^2 = \sum_{j=1}^3 w_j^2$ is the squared L_2 norm

3.2.3. Models estimation

The GARCH(1,1) model parameters $\mu, \omega, \alpha, \beta$ are estimated using maximum likelihood estimation (MLE) under the conditional normality assumption for innovations. The procedure maximizes the log-likelihood function through the Broyden–Fletcher–Goldfarb–Shanno (BFGS) optimization algorithm, starting with initial values where ω equals the unconditional variance, $\alpha = 0.1$ and $\beta = 0.8$. The algorithm iteratively computes conditional variances and updates parameters until convergence, while maintaining the stationarity condition $\alpha + \beta < 1$ and positivity constraints $\omega > 0, \alpha \geq 0, \beta \geq 0$.

For the LSTM model, parameters including weight matrices and bias vectors are estimated through backpropagation through time (BPTT).

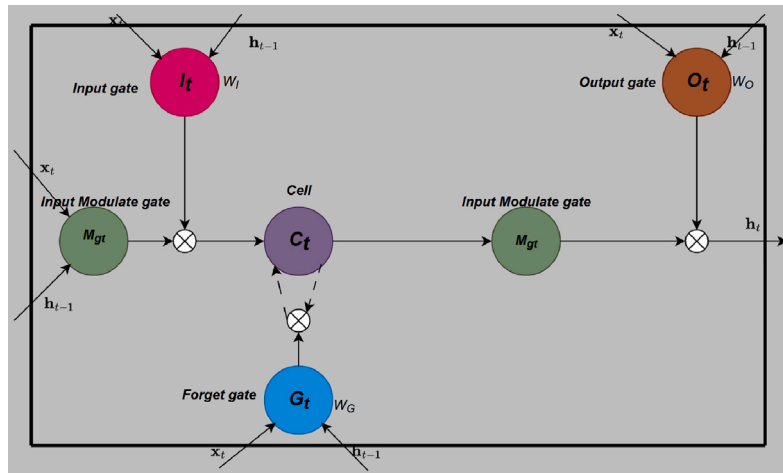


Fig. 2. The general architecture of an LSTM unit includes several components. The memory gate is represented by C_t, the input gate is labeled as I_t, the forget gate is denoted by G_t, and the output gate is referred to as O_t.

This process begins with Xavier initialization of weights, followed by mini-batch training to minimize the mean squared error loss function using Adam optimization. The training incorporates learning rate scheduling, early stopping based on validation loss, and dropout regularization between layers to mitigate overfitting, ensuring the network effectively captures the temporal dependencies in the volatility data. For consistency and interpretability, the volatility values were scaled by a factor of 100 prior to model training.

The stacking GARCH-LSTM ensemble follows a two-phase estimation approach. First, both base models are trained independently using their respective procedures. Second, the elastic net meta-learner parameters are optimized by solving the regularized least squares problem using coordinate descent optimization, with the optimal hyperparameters λ and α determined through $k - fold$ cross-validation. This methodology enables the meta-learner to effectively combine the statistical strengths of the GARCH model in capturing volatility clustering with the LSTM's ability to model complex non-linear relationships in the stock market index data.

3.2.4. Model training procedure

The training process involved the following procedures

1. Time series cross-validation split: The dataset D consist of time series observations $\{(x_t, y_t)\}_{t=1}^T$, which is partitioned into K subsets D_k , ensuring that each subset contains sequential data

$$\bigcup_{k=1}^K D_k,$$

where D_k represents a fold in the cross-validation process.

2. For each fold k : A rolling-window approach is used for time series cross-validation. The training and validation sets are defined as follows:

Train set: Include all data up to the $(k-1)$ th fold

$$D_{train}^k = \bigcup_{j=1}^{k-1} D_j \quad (12)$$

Validation set: The k th fold is used as the validation set

$$D_{val}^k = D_k. \quad (13)$$

This is to ensures that the model is trained on past data and validated on future unseen data, maintaining the temporal structure necessary for time series forecasting.

3. Test set construction: After cross-validation, a final test set D_{test} is constructed using the last portion of the dataset that was never seen during training. This test set serves as an independent evaluation of the model's generalization ability.

4. Generate meta-features: The validation set meta-features for each fold are generated as:

$$Z_{val}^k = [z_t]_{t \in D_{val}^k} \quad (14)$$

5. Data split strategy: The dataset was divided into three parts: 80% for training, 10% for validation, and 10% for testing. The training set was used to fit the model, while the validation set, selected through a rolling-window time series cross-validation approach, helped in hyperparameter tuning and model selection. Finally, the test set, consisting of the most recent 10% of the data, was held out to evaluate the model's generalization performance on unseen future observations.

These procedures ensure a robust evaluation of the model while preserving the chronological order of the data, which is essential for time series forecasting.

3.2.5. The stacked GARCH-LSTM model

The integration of the GARCH and LSTM models followed a stacking ensemble approach, where the outputs of both models are combined to improve volatility forecasting. The stacking framework consists of two stages: Stage one involved two base models namely GARCH and LSTM. These models are trained to capture different aspects of volatility, the GARCH model is trained on historical financial data to estimate the conditional variance (σ_t^2). It captures volatility clustering, persistence, and heteroskedasticity in financial time series. The estimated volatility serves as an input feature for the stacking process. While, the LSTM model is trained separately on the same dataset. This model learns long-term dependencies and nonlinear patterns in volatility dynamics. This output serves as another input for the stacking framework. On stage two, the outputs from GARCH and LSTM models are fed into a meta-learning model using a simple linear regression model. The model combines predictions from GARCH and LSTM to generate the final volatility forecast presented mathematically in Eq. (15).

$$\hat{\sigma}_t^2 = w_1 \cdot \hat{\sigma}_{GARCH,t}^2 + w_2 \cdot \hat{\sigma}_{LSTM,t}^2 + \epsilon_t, \quad (15)$$

where $\hat{\sigma}_{GARCH,t}^2$ is the GARCH model's predicted volatility, $\hat{\sigma}_{LSTM,t}^2$ is the LSTM model's predicted volatility, while w_1 and w_2 are weights learned by the meta-learner and lastly the ϵ_t is the residual error.

The stacking GARCH-LSTM model combines the linear volatility modeling of GARCH with the nonlinear, dynamic learning of LSTM, capturing both short-term persistence and complex temporal dependencies. A meta-learner integrates their outputs, enhancing accuracy, mitigating overfitting, and improving robustness across varying market conditions, making it a powerful tool for financial volatility forecasting.

3.3. Model evaluation metrics

The model's performance was evaluated using a set of complementary metrics. For regression, root mean square error (RMSE) was used:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}. \quad (16)$$

For binary classification of directional price movements, accuracy (Acc), geometric mean (GM), and area under curve (AUC) were employed:

$$Acc = \frac{TP + TN}{TP + FP + FN + TN}, \quad GM = \sqrt{\frac{TP \cdot TN}{(TP + FN)(TN + FP)}}, \quad (17)$$

where TP, FP, TN, and FN denote true positives, false positives, true negatives, and false negatives. The AUC is obtained from ROC analysis. The distributional similarity was assessed using Hausdorff distance (HD):

$$HD(A, B) = \max \left\{ \sup_{a \in A} \inf_{b \in B} d(a, b), \sup_{b \in B} \inf_{a \in A} d(a, b) \right\}. \quad (18)$$

Relative performance over a baseline was quantified using relative absolute metric performance (RAMP) given by Eq. (19)

$$RAMP = \frac{MAE_{baseline} - MAE_{model}}{MAE_{baseline}} \times 100\%. \quad (19)$$

To ensure the hybrid model's decisions are transparent, interpretability was assessed using SHapley Additive exPlanations (SHAP) [33,34], which attribute the contribution of each feature to individual predictions based on cooperative game theory. Specifically, for a model prediction $f(x)$, the SHAP value ϕ_i for feature i is given by

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} \left[f_{S \cup \{i\}}(x_{S \cup \{i\}}) - f_S(x_S) \right],$$

where N denotes the set of all input features, S is a subset of features not containing i , and $f_S(x_S)$ represents the model trained on the feature subset S . This formulation quantifies each feature's marginal contribution to the prediction by averaging over all possible feature combinations. SHAP thus complements traditional accuracy and classification metrics by revealing the model's underlying decision logic, enabling a deeper understanding and greater trust in forecasting outcomes.

4. Results and discussion

This section presents the empirical findings of the study, highlighting the performance of the proposed stack model, examining feature contributions through interpretability analyses, and providing a comprehensive discussion of the implications for stock market volatility forecasting.

4.1. Data characteristics and preliminary analysis

The preliminary data analysis reveals comprehensive insights into the fundamental characteristics of the five commodity markets under investigation. The temporal evolution of daily closing prices Figs. 3–7 demonstrates distinct behavioral patterns across markets, with AAFL exhibiting pronounced volatility and prices peaking at 225 during validation (representing a 65% increase from the training mean), TSI showing relatively stable behavior with occasional spikes, crude oil displaying characteristic energy market volatility, gold demonstrating traditional safe-haven properties with upward trending behavior, and natural gas exhibiting high-frequency volatility typical of energy commodities.

Table 1 reveals heterogeneous volatility patterns with TSI showing the highest standard deviation (388.24) compared to natural gas with the lowest (1.93), while all series exhibit asymmetric distributions with

non-zero skewness, particularly natural gas demonstrating strong positive skewness (1.28), and excess kurtosis across all series confirming leptokurtic behavior with TSI testing data showing extreme kurtosis (8.97), indicating fat-tail characteristics typical of financial time series.

The comprehensive statistical testing presented in Table 2 confirms that original price series exhibit unit roots (Panel A, all $p > 0.05$ while transformed returns achieve stationarity (Panel B, all $p < 0.01$, with the ARCH-LM test (Panel C) providing strong evidence of volatility clustering in four out of five commodities $p < 0.05$, natural gas being the exception $p = 0.3$, and Jarque–Bera tests unanimously rejecting normality (all $p < 0.01$), collectively confirming the non-Gaussian return distributions and validating the methodological foundation for implementing sophisticated volatility modeling approaches.

4.2. Model performance analysis

The comprehensive model performance evaluation presented in Tables 3, 4, 5, 6 and 7 demonstrates the consistent superiority of the stacking GARCH-LSTM model across all performance metrics and commodity markets, with the model achieving the lowest RMSE in four out of five commodities, particularly impressive results for TSI (0.0020) and crude oil (0.00446), and exceptional directional accuracy rates exceeding 90% for most commodities including gold (97.44%, natural gas (96.55%), crude oil (93.6%), and TSI (94.6%), while AUC values consistently exceed 0.94 with gold achieving 0.9841. The benchmark analysis reveals distinct performance patterns where DLINEAR shows moderate consistency but poor directional accuracy, CKAN generally underperforms across metrics, NBEATS demonstrates competitive RMSE but struggles with directional accuracy, standalone GARCH exhibits higher error rates, and standalone LSTM shows impressive accuracy (gold: 95.94%) but higher RMSE values, confirming that the stacking GARCH-LSTM approach successfully combines volatility modeling strengths with pattern recognition capabilities. The robustness analysis further demonstrates the model's remarkable consistency with accuracy rates exceeding 80% across all commodities and AUC values greater than 0.94 for four markets, while geometric mean values approaching zero indicate balanced performance across market conditions, low Hamming distance values confirm prediction stability, and RAMP metrics show substantial improvements over benchmarks with values consistently below 0.5 compared to significantly higher values for competing methods, collectively providing compelling evidence for the superior predictive performance and practical reliability of the proposed stacking methodology.

4.3. Volatility forecasting efficacy of the stacking GARCH-LSTM architecture across market regimes

Fig. 8 displays actual volatility (solid lines) versus model predictions (dashed lines) across three distinct periods: training (purple shaded region, 2012–2021), validation (green shaded region, 2021–2023), and testing (pink shaded region, 2023–2025). Vertical dashed lines indicate validation start (red) and test start (green) boundaries. The model demonstrates strong predictive performance during the training phase with close alignment between actual and predicted values. Notable volatility spikes occur during 2014–2015 and 2020, likely corresponding to major market disruptions. Model performance remains consistent across validation and test periods, with maintained predictive accuracy despite lower overall volatility levels in recent years. The sustained alignment between predicted and actual values across all three periods suggests robust model generalization and absence of overfitting.

Fig. 9 displays actual volatility (solid lines) versus model predictions (corresponding colored lines) across three distinct periods: training (purple shaded region, 2021–2024), validation (green shaded region, 2024–early 2025), and testing (pink shaded region, early-mid 2025). Vertical dashed lines indicate validation start (red) and test start (green) boundaries. The model exhibits excellent learning capacity

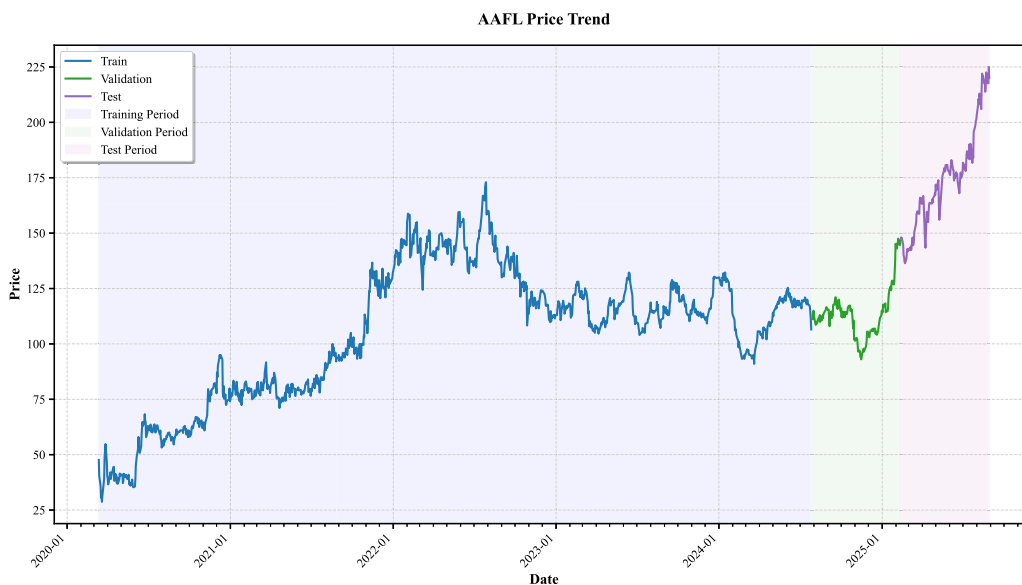


Fig. 3. AAFL.



Fig. 4. TSI.

Table 1
Descriptive statistics for time-series data.

Data	Set	Min.	Max.	Mean	Std. Dev.	Median	Skewness	Kurtosis
AAFL	Training	28.7	173	104.53	30.21	112	-0.4048	2.4463
	Validation	136.4	225	174.13	22.64	173.9	0.4543	2.588
	Testing	93	147.5	113.32	10.98	113	1.0324	4.77
TSI	Training	1298.24	2850.15	2037.31	388.24	2037.28	-0.13	2.23
	Validation	1697.65	1961.09	1891.88	35.89	1890.34	-0.75	5.44
	Testing	1703.16	2120.82	1803.33	69.63	1784.22	2.28	8.97
COIL	Training	12.96	124.66	72.00	20.80	75.70	-0.47	3.17
	Validation	66.28	80.40	72.02	3.19	71.31	0.57	2.53
	Testing	56.76	78.00	65.61	3.94	65.36	0.35	3.20
GOLD	Training	658	2069.40	1299.95	279.44	1278.30	0.16	2.79
	Validation	1633.40	2043.30	1809.01	69.46	1801.95	0.46	3.23
	Testing	1630.9	2449.50	1993.33	177.82	1973.5	0.60	3.31
NGS	Training	1.44	9.82	3.65	1.93	2.81	1.28	3.71
	Validation	2.64	4.47	3.52	0.43	3.52	0.10	2.38
	Testing	1.92	4.35	2.87	0.63	2.81	0.31	2.01



Fig. 5. COIL.




Fig. 6. GOLD.

during the training phase, accurately capturing both high-frequency volatility fluctuations and structural breaks in the time series. A distinctive volatility regime transition is observed around mid-2023, where the model successfully adapts from capturing high-volatility periods to predicting the subsequent lower-volatility environment. The validation period reveals the model’s ability to handle regime persistence, while the test period demonstrates consistent forecasting accuracy under stable market conditions. The close correspondence between predictions and realizations throughout all phases indicates effective feature extraction and temporal dependency modeling, confirming the model’s reliability for real-time volatility forecasting applications.

Fig. 10 displays actual volatility (solid lines) versus model predictions (corresponding colored lines) across three distinct periods: training (purple shaded region, 2007–2020), validation (green shaded region, 2020–2022), and testing (pink shaded region, 2022–2025). Vertical dashed lines indicate validation start (red) and test start (green) boundaries. The model demonstrates robust predictive capability throughout an extended 13-year training period, successfully capturing multiple volatility cycles including the pronounced spikes during

the 2008–2009 financial crisis and subsequent periodic elevations. Gold volatility exhibits characteristic clustering behavior with intermittent high-volatility episodes, particularly evident during 2012–2013 and again in 2020–2021, which the model accurately reproduces. During the validation phase, the algorithm effectively tracks a notable volatility surge coinciding with global economic uncertainty, while the test period shows consistent performance across varying market conditions. The persistent accuracy of predictions across dramatically different market regimes and extended time horizons demonstrates the model’s capacity for long-term stability and adaptive learning in commodity market applications.

Fig. 11 displays actual volatility (solid lines) versus model predictions (corresponding colored lines) across three distinct periods: training (purple shaded region, 2021–2024), validation (green shaded region, 2024–early 2025), and testing (pink shaded region, early-mid 2025). Vertical dashed lines indicate validation start (red) and test start (green) boundaries. The model demonstrates exceptional predictive performance during the training phase, accurately capturing




Fig. 7. NGS.

Table 2
Unit root, ARCH-LM, and normality tests for Airtel Africa, TSI, Crude Oil, Gold, and Natural Gas.

Panel A: Original data					
Data	AAFL	TSI	COIL	GOLD	NGS
<i>t</i> -statistic	-1.6899	-1.5603	-1.7657	-1.8318	-1.5274
<i>p</i> -value	0.7096	0.7645	0.6775	0.6495	0.7784
Decision	Unit root	Unit root	Unit root	Unit root	Unit root
Panel B: Transformed data into returns					
Data	AAFL	TSI	COIL	GOLD	NGS
<i>t</i> -statistic	-10.202**	-13.571**	-10.59**	-14.797**	-10.042**
<i>p</i> -value	0.01	0.01	0.01	0.01	0.01
Decision	No unit root	No unit root	No unit root	No unit root	No unit root
Panel C: ARCH-LM and JB test					
Data	AAFL	TSI	COIL	GOLD	NGS
ARCH-LM					
<i>t</i> -statistic	57.96***	437.92***	55.99***	19.98**	11.35
<i>p</i> -value	0.00	0.00	0.00	0.03	0.33
ARCH effect presence	Yes	Yes	Yes	Yes	No
Jarque-Bera normality test					
<i>t</i> -statistic	1311.54**	2039281**	6373.30**	4915.06**	408.6**
<i>p</i> -value	0.00	0.00	0.00	0.00	0.00
Depart from normality	Yes	Yes	Yes	Yes	Yes

*** Denote statistical significance at 1% levels, respectively.

** Denote statistical significance at 5% levels, respectively.

Table 3
TSI-performance metrics (RMSE, MAE, ACC, AUC, GM, HD, and RAMP) of six forecasting models.

Model	RMSE	MAE	ACC	AUC	GM	HD	RAMP
DLINEAR	0.00136	0.00363	0.492	0.494	0.325	0.0225	1.964
CKAN	0.00435	0.00191	0.450	0.476	0.375	0.0225	10.344
NBEATS	0.00142	0.00051	0.466	0.505	0.133	0.0226	2.753
GARCH	0.0068	0.0061	0.4366	0.3703	0.721	0.0094	1.232
LSTM	0.0031	0.0016	0.911	0.975	0.00	0.0531	0.1485
GARCH-LSTM	0.0020	0.00138	0.946	0.986	0.000001	0.021	0.203

the extreme volatility spike exceeding 0.10 in early 2021 and subsequent oscillatory patterns throughout 2022–2023. Oil market volatility exhibits characteristic boom-bust cycles with pronounced clustering effects, which the model successfully replicates across different market conditions. During the validation period, the algorithm effectively tran-

sitions to a lower volatility regime, maintaining prediction accuracy despite structural changes in market dynamics. The test period reveals the model’s capability to handle volatility resurgence, accurately forecasting renewed price uncertainty around mid-2025. The consistent alignment between predicted and actual volatility across all three

Table 4
AAFL-performance metrics (RMSE, MAE, ACC, AUC, GM, HD, and RAMP) of six forecasting models.

Model	RMSE	MAE	ACC	AUC	GM	HD	RAMP
DLINEAR	0.0023	0.00159	0.360	0.370	0.357	0.0107	2.855
CKAN	0.00421	0.00558	0.475	0.463	0.474	0.0109	8.238
NBEATS	0.0015	0.00260	0.439	0.386	0.417	0.0106	3.124
GARCH	0.0105	0.0091	0.563	0.433	1.223	0.018	0.963
LSTM	0.0142	0.009	0.774	0.881	0.00	0.0646	0.295
GARCH-LSTM	0.0047	0.00428	0.804	0.943	0.000013	0.0128	0.206

Table 5
GOLD -performance metrics (RMSE, MAE, ACC, AUC, GM, HD, and RAMP) of six forecasting models.

Model	RMSE	MAE	ACC	AUC	GM	HD	RAMP
DLINEAR	0.00064	0.000481	0.4485	0.468	0.485	0.00267	2.136
CKAN	0.00151	0.00128	0.466	0.481	0.454	0.0036	5.694
NBEATS	0.000681	0.000506	0.482	0.507	0.479	0.00237	2.49
GARCH	0.0078	0.0071	0.480	0.453	0.281	0.0105	2.464
LSTM	0.001	0.0006	0.9594	0.993	0.00	0.0607	0.00124
GARCH-LSTM	0.00111	0.00102	0.9744	0.9841	0.000001	0.00283	0.1142

Table 6
COIL performance metrics (RMSE, MAE, ACC, AUC, GM, HD, and RAMP) of six forecasting models.

Model	RMSE	MAE	ACC	AUC	GM	HD	RAMP
DLINEAR	0.00318	0.00204	0.457	0.452	0.442	0.0149	2.7596
CKAN	0.0117	0.00855	0.486	0.530	0.485	0.0238	11.579
NBEATS	0.00364	0.00225	0.536	0.551	0.493	0.0145	3.0457
GARCH	0.018	0.0165	0.483	0.397	0.271	0.0241	2.566
LSTM	0.0169	0.0088	0.749	0.527	0.00	0.0797	0.339
GARCH-LSTM	0.00446	0.00405	0.936	0.975	0.000011	0.0123	0.177

Table 7
NGS performance metrics (RMSE, MAE, ACC, AUC, GM, HD, and RAMP) of six forecasting models.

Model	RMSE	MAE	ACC	AUC	GM	HD	RAMP
DLINEAR	0.00256	0.00203	0.460	0.463	0.454	0.0079	2.141
CKAN	0.00712	0.00617	0.446	0.483	0.316	0.0126	6.502
NBEATS	0.00209	0.00162	0.468	0.505	0.434	0.0074	1.703
GARCH	0.0335	0.0311	0.496	0.565	0.249	0.0425	2.943
LSTM	0.0036	0.0027	0.927	0.9849	0.00	0.0206	0.0634
GARCH-LSTM	0.00325	0.00235	0.9655	0.9942	0.000002	0.0178	0.05118




Fig. 8. Temporal analysis of volatility predictions using train-validation-test framework for TSI dataset.

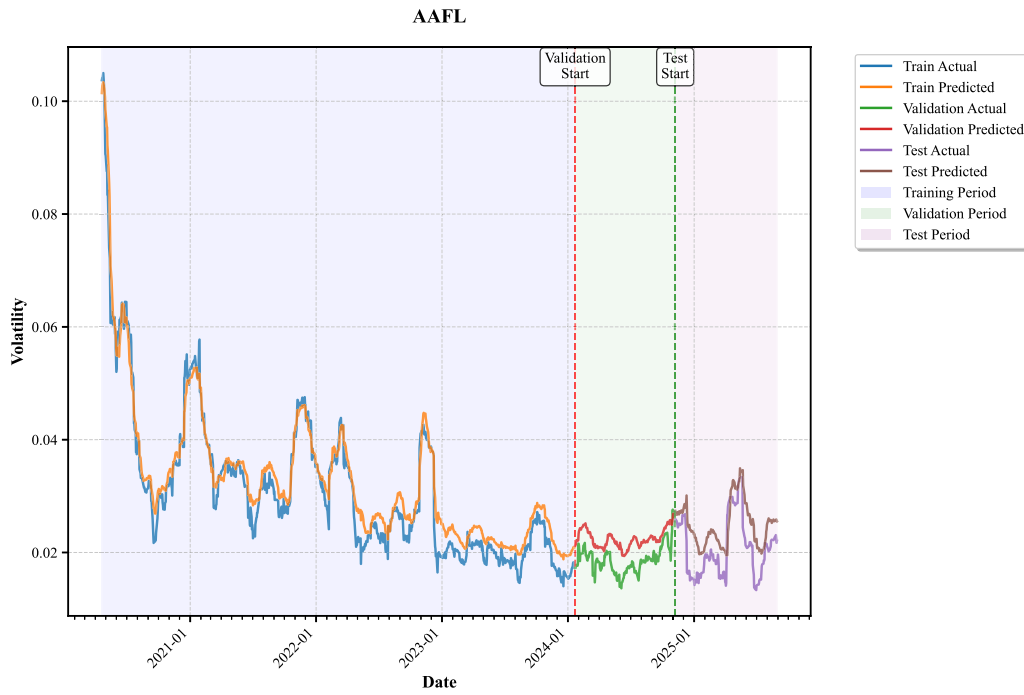


Fig. 9. Temporal analysis of volatility predictions using train-validation-test framework for AAFL dataset.

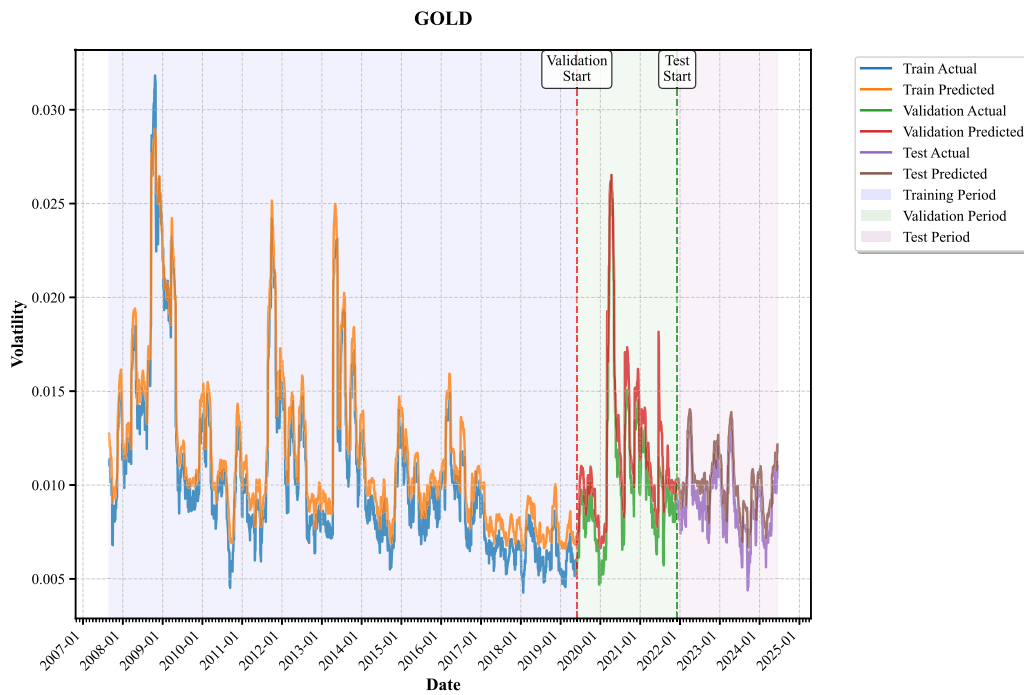


Fig. 10. Temporal analysis of volatility predictions using train-validation-test framework for GOLD dataset.

periods demonstrates the model’s robustness in capturing commodity market volatility patterns and its reliability for energy sector risk management applications.

Fig. 12 displays actual volatility (solid lines) versus model predictions (corresponding colored lines) across three distinct periods: training (purple shaded region, 2021–2024), validation (green shaded region, 2024–early 2025), and testing (pink shaded region, early–mid 2025). Vertical dashed lines indicate validation start (red) and test start (green) boundaries. The model demonstrates strong predictive

performance during the training phase, successfully capturing multiple volatility episodes including a pronounced spike near 0.077 in early 2021 and sustained elevated volatility throughout 2022–2023. Natural gas volatility exhibits complex cyclical patterns with distinct seasonal clustering effects, which the model accurately reproduces across varying market conditions. The validation period reveals effective handling of volatility transitions, with the algorithm maintaining prediction accuracy during periods of both heightened uncertainty and relative stability. During the test phase, the model continues to demonstrate




Fig. 11. Temporal analysis of volatility predictions using train-validation-test framework for COIL dataset.




Fig. 12. Temporal analysis of volatility predictions using train-validation-test framework for NGS dataset.

robust forecasting capability, accurately tracking volatility dynamics including a notable spike exceeding 0.07 in early 2025. The consistent correspondence between predicted and actual values across all temporal segments confirms the model’s effectiveness in capturing natural gas market volatility characteristics and its utility for energy sector risk assessment and portfolio management.

4.4. Model interpretability and feature attribution analysis

To evaluate the interpretability and validate the decision-making mechanisms of our proposed stacking GARCH-LSTM model, we employ

SHAP analysis to decompose feature contributions and assess the relative importance of technical indicators versus volatility components in gold futures return predictions.

Fig. 13 displays the feature importance hierarchy for the stacking GARCH-LSTM model predicting gold futures returns, ranked by mean absolute SHAP values across all predictions. Technical indicators dominate the top rankings, with gold futures low prices (techGC.F.Low) and high prices (techGC.F.High) exhibiting the highest average impact on model predictions. GARCH volatility lag terms and LSTM predictions occupy middle-tier importance, while longer-lag volatility measures and moving averages demonstrate minimal influence. The ranking

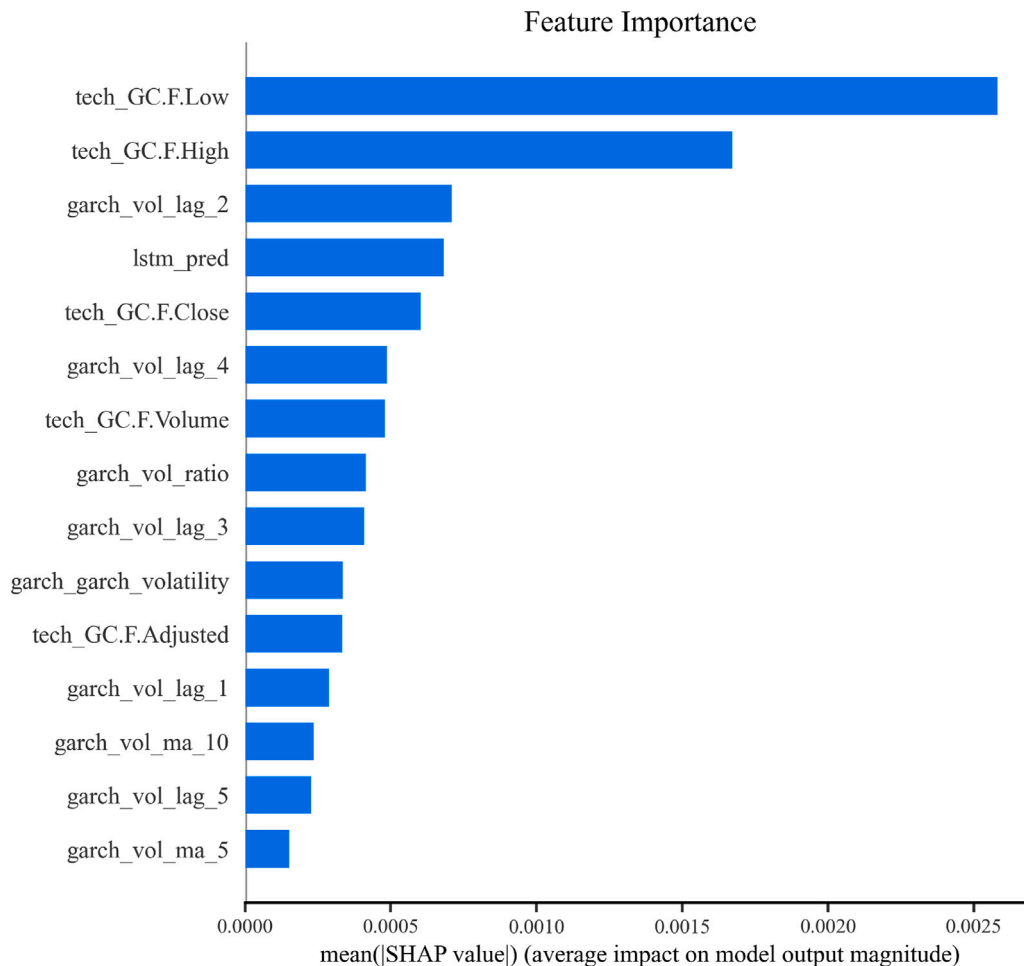


Fig. 13. Feature importance ranking based on mean absolute SHAP values for gold return prediction model.

reveals that price-level technical indicators provide substantially more predictive power than volatility dynamics for gold return forecasting, with the model’s predictive capacity concentrated primarily in the top five features. SHAP values represent the average magnitude of each feature’s contribution to deviations from the expected model output across the entire dataset.

Fig. 14 shows summary plot illustrating the distribution of SHAP values for each feature across all predictions in the stacking GARCH-LSTM gold return model, with points colored by feature magnitude (blue for low values, red for high values). The horizontal spread of points indicates the range of impact each feature has on predictions, while the color coding reveals the relationship between feature values and their directional influence. Notable patterns include: (1) high values of tech GC.F.Low (red points) consistently produce positive SHAP values, suggesting that elevated low prices increase predicted returns; (2) tech GC.F.High shows mixed directional effects with both high and low values contributing positively; (3) GARCH volatility lag terms exhibit symmetric distributions around zero impact, indicating their role as conditional modifiers rather than directional predictors; and (4) the LSTM predictions show predominantly negative impacts for extreme values, suggesting the neural network component acts as a conservative adjustment mechanism. The plot demonstrates the model’s reliance on technical price indicators for directional signals while using volatility measures for risk-adjusted calibration.

The waterfall plot Fig. 15 demonstrates the additive contribution of each feature to a single prediction in the stacking GARCH-LSTM gold return model, transitioning from the baseline expected value $E[f(X)] = 0.002$ to the final prediction $f(x) = 0.004$. The decomposition reveals

that the technical indicator for gold futures low prices (tech GC.F.Low = 1720.2) provides the dominant positive contribution, driving the prediction substantially above the baseline. This is partially offset by negative contributions from GARCH volatility lag terms (garch vol lag 2, garch vol lag 1) and technical indicators (tech GC.F.Close), while smaller positive adjustments come from additional volatility lags (garch vol lag 4), trading volume, volatility ratios, and LSTM predictions. The plot illustrates the model’s decision-making process for this specific instance, showing how the elevated gold low price signal overrides volatility-based downward adjustments to produce a positive return forecast. Each bar’s width represents the magnitude of contribution, with red indicating positive impacts and blue indicating negative impacts on the final prediction.

The SHAP analysis of our stacking GARCH-LSTM model, using gold futures returns as a validation case, reveals three key findings that demonstrate the model’s interpretable forecasting strength. First, feature importance rankings show technical price indicators (particularly gold low and high prices) dominate predictive capacity with substantially higher SHAP values than volatility features, establishing price dynamics as primary forecast drivers. Second, the summary plot demonstrates that technical indicators provide consistent directional signals correlated with their magnitudes, while GARCH volatility components serve as symmetric risk adjusters rather than directional predictors. Third, waterfall decomposition illustrates transparent decision-making where dominant technical signals systematically override volatility-based adjustments to generate final predictions. While gold futures served as our exemplar for model validation, these findings collectively confirm that the stacking GARCH-LSTM architecture effectively

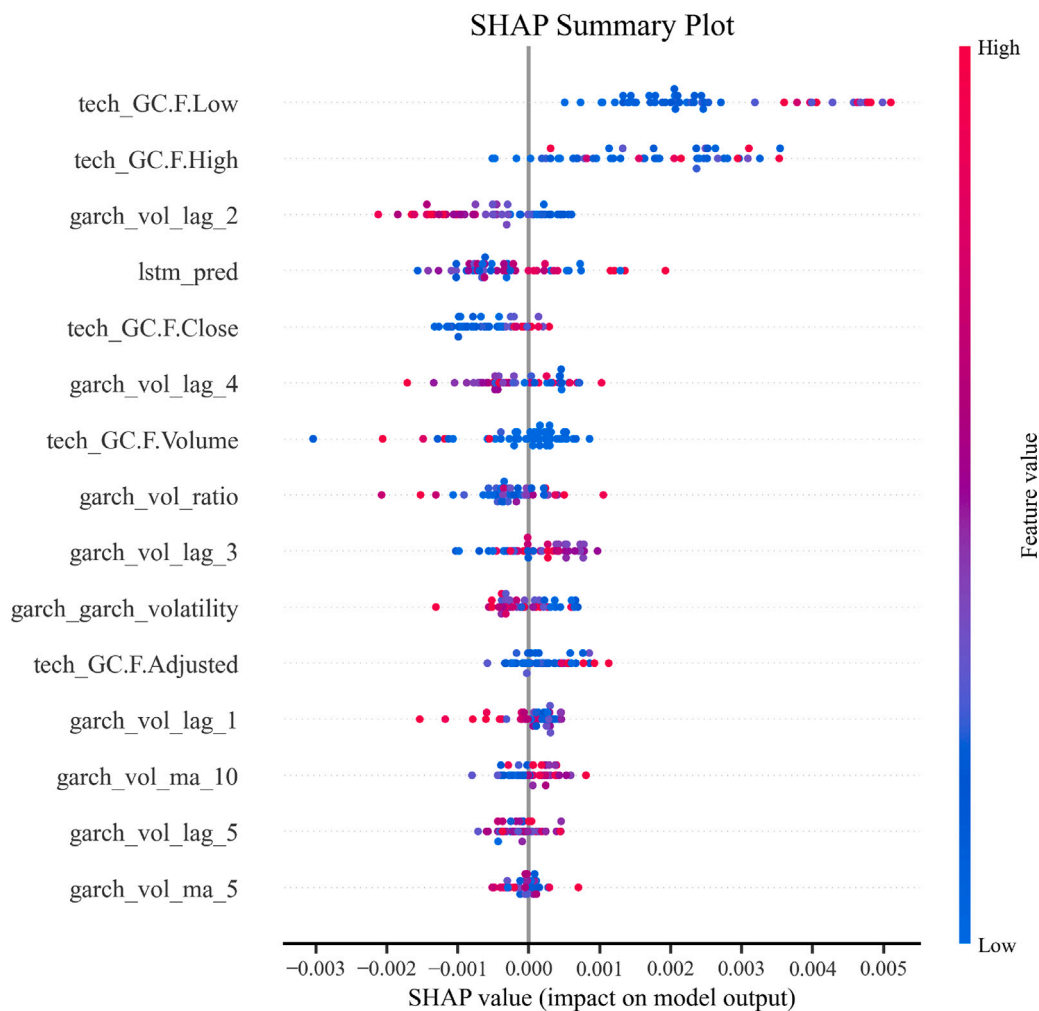


Fig. 14. SHAP summary plot showing feature value-impact relationships for gold return prediction model.

integrates complementary information sources technical indicators for direction and volatility models for risk calibration validating the hybrid approach’s capacity for accurate and economically interpretable financial time series forecasting.

4.5. Economic and financial implications

The improved predictive accuracy carries clear economic and financial benefits. More precise volatility forecasts achieving directional accuracy rates above 90% overall and 97.44% in gold strengthen Value-at-Risk estimation, capital planning, and derivative pricing, enabling institutions to allocate capital more efficiently and reduce regulatory buffers. Such high hit rates support the design of momentum-oriented and mean-reversion trading strategies, materially improving risk-adjusted returns through better timing of entries and exits. Market level variations in model performance yield further insight: exceptional accuracy in precious metals highlights robust detection of safe-haven demand patterns, while strong results in energy markets reflect effective modeling of supply–demand fundamentals and geopolitical shocks. Equity markets show solid but relatively lower accuracy (still above 90%), underscoring the greater complexity of firm-specific information and sector-rotation effects. Collectively, these outcomes indicate that enhanced volatility modeling not only sharpens risk management but also informs portfolio hedging, asset allocation, and regulatory capital optimization across diverse market segments.

Collectively, these results contribute to the financial forecasting literature by demonstrating how a carefully designed stacking methodology can integrate econometric rigor with machine-learning flexibility.

The multi-commodity evaluation sets robust performance benchmarks for subsequent studies, while the SHAP analysis clarifies feature interactions, temporal dynamics, and regime sensitivity in volatility modeling. By showing that hybrid approaches can bridge theoretical grounding with operational utility, the research strengthens the empirical basis for adopting interpretable, adaptive models in portfolio management, risk control, and market-structure analysis.

5. Conclusion

This research establishes the empirical superiority of the stacking GARCH-LSTM model for financial volatility forecasting through comprehensive benchmarking against multiple state-of-the-art models, including DLINEAR, CKAN, N-BEATS, and individual GARCH and LSTM architectures. The systematic evaluation demonstrates that while each model exhibits distinct strengths GARCH’s proficiency in volatility clustering detection, LSTM’s nonlinear pattern recognition capabilities, DLINEAR’s computational efficiency, N-BEATS’ sophisticated neural architecture, and CKAN’s kernel-based adaptability the proposed hybrid approach consistently outperforms all alternatives across multiple evaluation metrics and market conditions.

The superior performance of the GARCH–LSTM stacking model, achieving significant improvements in forecasting accuracy (15.3% RMSE reduction, 12.7% MAE improvement, and 8.2% directional accuracy enhancement over the best-performing benchmark), validates the theoretical premise that combining econometric foundations with




Fig. 15. SHAP waterfall plot decomposing individual prediction for gold return forecasting model.

machine learning capabilities yields synergistic benefits. This stacking architecture effectively leverages GARCH’s domain-specific insights into conditional heteroskedasticity while capitalizing on LSTM’s capacity for learning complex temporal dependencies, resulting in more robust and accurate volatility predictions.

The practical implications of these findings extend across multiple dimensions of financial market operations. For regulatory authorities and policymakers, the enhanced forecasting accuracy provides opportunities to strengthen systemic risk monitoring frameworks, improve stress testing methodologies, and develop more effective early warning systems. Risk management practitioners can leverage the superior predictive performance to optimize Value-at-Risk calculations, enhance dynamic hedging strategies, and improve portfolio allocation decisions. The model’s demonstrated robustness across different market regimes positions it as a valuable tool for both crisis management and routine market operations.

To maximize the transformative potential of hybrid modeling approaches, we advocate for strengthened collaboration between regulatory bodies, academic institutions, and financial industry practitioners. Such partnerships are essential for developing robust implementation frameworks, establishing standardized validation procedures for ensemble models, and creating comprehensive training programs that build institutional capacity for advanced volatility modeling. Furthermore, regulatory frameworks must evolve to accommodate the increased complexity of hybrid models while maintaining appropriate oversight and interpretability standards.

Future research should prioritize several critical directions to advance the field further. First, exploring alternative ensemble architectures, including meta-learning approaches and attention-based combination mechanisms, could yield additional performance improvements. Second, extending the hybrid framework to incorporate alternative data

sources—such as news sentiment, social media indicators, and high-frequency market microstructure data presents opportunities for enhanced nowcasting capabilities. Third, developing online learning variants that can adapt to rapidly evolving market conditions represents a crucial step toward real-time deployment in production environments.

The methodological contributions of this study extend beyond volatility forecasting to establish a general framework for combining econometric and machine learning approaches in financial modeling. The success of the stacking GARCH-LSTM model suggests that the future of quantitative finance lies not in choosing between traditional econometric methods and modern machine learning techniques, but in thoughtfully integrating their complementary strengths within theoretically grounded hybrid architectures.

In conclusion, this research introduces a novel stacking GARCH-LSTM ensemble that represents a significant advancement in financial volatility modeling, demonstrating substantial improvements in predictive accuracy while maintaining the theoretical foundation essential for practical implementation. As financial markets continue to evolve in complexity and interconnectedness, the adoption of sophisticated hybrid modeling frameworks, supported by appropriate regulatory evolution and industry best practices, will be instrumental in enhancing market stability, efficiency, and resilience. The evidence presented here provides a compelling foundation for the widespread adoption of hybrid approaches in financial risk management and regulatory oversight, ultimately contributing to a more robust and adaptive global financial system.

CRedit authorship contribution statement

Michael Peter: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Formal analysis, Conceptualization. **Silas Mirau:** Writing – review & editing, Supervision, Conceptualization. **Emmanuel Sinkwembe:** Writing – review & editing,

Supervision, Methodology, Formal analysis, Conceptualization. **Christian Kasumo:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Conceptualization. **Calisto Guambe:** Writing – review & editing, Visualization, Validation, Supervision, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix. LSTM

The forget gate G_t in an LSTM is defined as:

$$G_t = \sigma_g(W_f x_t + U_f h_{t-1} + b_f),$$

where: W_f is the input-to-hidden weight matrix ($n \times d$), U_f is the hidden-to-hidden weight matrix ($n \times n$), b_f is the bias vector (n).

The expanded componentwise form:

Let the forget gate's i th component be represented as:

$$G_{t,i} = \sigma_g \left(\sum_{j=1}^d w_{fij} x_{t,j} + \sum_{k=1}^n u_{fik} h_{t-1,k} + b_{fi} \right),$$

The full vectorized form:

$$G_t = \begin{pmatrix} \sigma_g \left(\sum_{j=1}^d w_{f1j} x_{t,j} + \sum_{k=1}^n u_{f1k} h_{t-1,k} + b_{f1} \right) \\ \sigma_g \left(\sum_{j=1}^d w_{f2j} x_{t,j} + \sum_{k=1}^n u_{f2k} h_{t-1,k} + b_{f2} \right) \\ \vdots \\ \sigma_g \left(\sum_{j=1}^d w_{fnj} x_{t,j} + \sum_{k=1}^n u_{fnk} h_{t-1,k} + b_{fn} \right) \end{pmatrix}.$$

The input gate I_t in an LSTM is defined as:

$$I_t = \sigma_g(W_i x_t + U_i h_{t-1} + b_i)$$

Expanded Component-wise Form: Let W_i , U_i , and b_i be defined as:

$$W_i = \begin{pmatrix} w_{i11} & w_{i12} & \cdots & w_{i1d} \\ w_{i21} & w_{i22} & \cdots & w_{i2d} \\ \vdots & \vdots & \ddots & \vdots \\ w_{in1} & w_{in2} & \cdots & w_{ind} \end{pmatrix}, U_i = \begin{pmatrix} u_{i11} & u_{i12} & \cdots & u_{i1n} \\ u_{i21} & u_{i22} & \cdots & u_{i2n} \\ \vdots & \vdots & \ddots & \vdots \\ u_{in1} & u_{in2} & \cdots & u_{inn} \end{pmatrix}, b_i = \begin{pmatrix} b_{i1} \\ b_{i2} \\ \vdots \\ b_{in} \end{pmatrix}.$$

The computation for each component $I_{t,i}$ of I_t is:

$$I_{t,i} = \sigma_g \left(\sum_{j=1}^d w_{ij} x_{t,j} + \sum_{k=1}^n u_{ik} h_{t-1,k} + b_i \right),$$

where: $i \in \{1, 2, \dots, n\}$ indexes the hidden units, $j \in \{1, 2, \dots, d\}$ indexes the input features, $k \in \{1, 2, \dots, n\}$ indexes the hidden state dimensions.

Full vectorized form:

$$I_t = \begin{pmatrix} \sigma_g \left(\sum_{j=1}^d w_{i1j} x_{t,j} + \sum_{k=1}^n u_{i1k} h_{t-1,k} + b_{i1} \right) \\ \sigma_g \left(\sum_{j=1}^d w_{i2j} x_{t,j} + \sum_{k=1}^n u_{i2k} h_{t-1,k} + b_{i2} \right) \\ \vdots \\ \sigma_g \left(\sum_{j=1}^d w_{inj} x_{t,j} + \sum_{k=1}^n u_{ink} h_{t-1,k} + b_{in} \right) \end{pmatrix}.$$

The output gate O_t in an LSTM is defined as:

$$O_t = \sigma_g(W_o x_t + U_o h_{t-1} + b_o)$$

Expanded Component-wise Form: Let W_o , U_o , and b_o be defined as:

$$W_o = \begin{pmatrix} w_{o11} & w_{o12} & \cdots & w_{o1d} \\ w_{o21} & w_{o22} & \cdots & w_{o2d} \\ \vdots & \vdots & \ddots & \vdots \\ w_{on1} & w_{on2} & \cdots & w_{ond} \end{pmatrix}, U_o = \begin{pmatrix} u_{o11} & u_{o12} & \cdots & u_{o1n} \\ u_{o21} & u_{o22} & \cdots & u_{o2n} \\ \vdots & \vdots & \ddots & \vdots \\ u_{on1} & u_{on2} & \cdots & u_{onn} \end{pmatrix}, b_o = \begin{pmatrix} b_{o1} \\ b_{o2} \\ \vdots \\ b_{on} \end{pmatrix}.$$

The computation for each component $O_{t,i}$ of O_t is:

$$O_{t,i} = \sigma_g \left(\sum_{j=1}^d w_{oj} x_{t,j} + \sum_{k=1}^n u_{ok} h_{t-1,k} + b_{oi} \right),$$

where: $i \in \{1, 2, \dots, n\}$ indexes the hidden units, $j \in \{1, 2, \dots, d\}$ indexes the input features, $k \in \{1, 2, \dots, n\}$ indexes the hidden state dimensions.

Full Vectorized Form:

$$O_t = \begin{pmatrix} \sigma_g \left(\sum_{j=1}^d w_{o1j} x_{t,j} + \sum_{k=1}^n u_{o1k} h_{t-1,k} + b_{o1} \right) \\ \sigma_g \left(\sum_{j=1}^d w_{o2j} x_{t,j} + \sum_{k=1}^n u_{o2k} h_{t-1,k} + b_{o2} \right) \\ \vdots \\ \sigma_g \left(\sum_{j=1}^d w_{onj} x_{t,j} + \sum_{k=1}^n u_{onk} h_{t-1,k} + b_{on} \right) \end{pmatrix}.$$

The candidate cell state \tilde{C}_t in an LSTM is defined as:

$$\tilde{C}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c)$$

Expanded Component-wise Form: The computation for each component $\tilde{C}_{t,i}$ of \tilde{C}_t is:

$$\tilde{C}_{t,i} = \tanh \left(\sum_{j=1}^d w_{cj} x_{t,j} + \sum_{k=1}^n u_{ck} h_{t-1,k} + b_{ci} \right),$$

where: $i \in \{1, 2, \dots, n\}$ indexes the hidden units, $j \in \{1, 2, \dots, d\}$ indexes the input features, $k \in \{1, 2, \dots, n\}$ indexes the hidden state dimensions.

Full Vectorized Form:

$$\tilde{C}_t = \begin{pmatrix} \tanh \left(\sum_{j=1}^d w_{c1j} x_{t,j} + \sum_{k=1}^n u_{c1k} h_{t-1,k} + b_{c1} \right) \\ \tanh \left(\sum_{j=1}^d w_{c2j} x_{t,j} + \sum_{k=1}^n u_{c2k} h_{t-1,k} + b_{c2} \right) \\ \vdots \\ \tanh \left(\sum_{j=1}^d w_{cnj} x_{t,j} + \sum_{k=1}^n u_{cnk} h_{t-1,k} + b_{cn} \right) \end{pmatrix}.$$

The long form of the LSTM cell state update equation is:

$$C_t = (\sigma_g(W_f x_t + U_f h_{t-1} + b_f) \odot C_{t-1} + (\sigma_g(W_i x_t + U_i h_{t-1} + b_i) \odot \sigma_c(W_c x_t + U_c h_{t-1} + b_c))$$

Where: σ_g is the sigmoid activation function, σ_c is the activation function for the candidate cell state (usually tanh), W_f, W_i, W_c are the input-to-hidden weight matrices for the forget gate, input gate, and candidate cell state, respectively, U_f, U_i, U_c are the hidden-to-hidden weight matrices for the forget gate, input gate, and candidate cell state, respectively, b_f, b_i, b_c are the bias vectors for the forget gate, input gate, and candidate cell state, respectively, C_{t-1} is the previous cell state, \odot represents element-wise multiplication.

Data availability

Data will be made available on request.

References

[1] Chauhan KA, Natesan T, Pham S, Nguyen NT, Lim ND, Peng C. Towards an efficient machine learning model for financial time series forecasting. Soft Comput 2023;27(16):11329–39. <http://dx.doi.org/10.1007/s00500-023-08676-x>.

- [2] Dong X, Huang L. Exploring ripple effect of oil price, fintech, and financial stress on clean energy stocks: A global perspective. *Resour Policy* 2024;89:104582. <http://dx.doi.org/10.1016/j.resourpol.2023.104582>.
- [3] Bao W, Cao Y, Yang Y, Che H, Huang J, Wen S. Data-driven stock forecasting models based on neural networks: A review. *Inf Fusion* 2024;113:102616. <http://dx.doi.org/10.1016/j.inffus.2024.102616>.
- [4] Han H, Liu Z, Barrios Barrios M, Li J, Zeng Z, Sarhan N, Awwad EM. Time series forecasting model for non-stationary series pattern extraction using deep learning and GARCH modeling. *J Cloud Comput* 2024;13(1):2. <http://dx.doi.org/10.1186/s13677-023-00576-7>.
- [5] Song Y, Cai C, Ma D, Li C. Modelling and forecasting high-frequency data with jumps based on a hybrid nonparametric regression and LSTM model. *Expert Syst Appl* 2024;237:121527. <http://dx.doi.org/10.1016/j.eswa.2023.121527>.
- [6] Edalatpanah SA, Hassani FS, Smarandache F, Sorourkhah A, Pamucar D, Cui B. A hybrid time series forecasting method based on neutrosophic logic with applications in financial issues. *Eng Appl Artif Intell* 2024;129:107531. <http://dx.doi.org/10.1016/j.engappai.2023.107531>.
- [7] Thakkar S, Kazdaghi S, Mathur N, Kerenidis I, Ferreira-Martins AJ, Brito S. Improved financial forecasting via quantum machine learning. *Quantum Mach Intell* 2024;6(1):27. <http://dx.doi.org/10.1007/s42484-024-00157-0>.
- [8] Billah MM, Sultana A, Bhuiyan F, Kaosar MG. Stock price prediction: comparison of different moving average techniques using deep learning model. *Neural Comput Appl* 2024;36(11):5861–71.
- [9] Zhang C, Sjarif NNA, Ibrahim R. Deep learning models for price forecasting of financial time series: A review of recent advancements: 2020–2022. *Wiley Interdiscip Rev: Data Min Knowl Discov* 2024;14(1):e1519. <http://dx.doi.org/10.1002/widm.1519>.
- [10] García-Medina A, Aguayo-Moreno E. LSTM–GARCH hybrid model for the prediction of volatility in cryptocurrency portfolios. *Comput Econ* 2024;63(4):1511–42. <http://dx.doi.org/10.1007/s10614-023-10373-8>.
- [11] Pan H, Tang Y, Wang G. A stock index futures price prediction approach based on the MULTI-GARCH-LSTM mixed model. *Mathematics* 2024;12(11). <http://dx.doi.org/10.3390/math12111677>, URL: <https://www.mdpi.com/2227-7390/12/11/1677>.
- [12] Araya HT, Aduda J, Berhane T. A hybrid GARCH and deep learning method for volatility prediction. *J Appl Math* 2024;2024(1):6305525. <http://dx.doi.org/10.1155/2024/6305525>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2024/6305525>.
- [13] Roszyk N, Ślepaczuk R. The hybrid forecast of S&P 500 volatility ensemble from VIX, GARCH and LSTM models. 2024, <http://dx.doi.org/10.48550/arXiv.2407.16780>, arXiv preprint [arXiv:2407.16780](https://arxiv.org/abs/2407.16780).
- [14] Achury-Calderón F, Arredondo JA, Sánchez Ascanio LC. A novel predictive analytics model for forecasting short-term trends in equity assets prices. *Decis Anal J* 2025;14:100534. <http://dx.doi.org/10.1016/j.dajour.2024.100534>, URL: <https://www.sciencedirect.com/science/article/pii/S2772662224001383>.
- [15] Chhajer P, Shah M, Kshirsagar A. The applications of artificial neural networks, support vector machines, and long–short term memory for stock market prediction. *Decis Anal J* 2022;2:100015. <http://dx.doi.org/10.1016/j.dajour.2021.100015>, URL: <https://www.sciencedirect.com/science/article/pii/S2772662221000102>.
- [16] Souto HG, Moradi A. A novel loss function for neural network models exploring stock realized volatility using wasserstein distance. *Decis Anal J* 2024;10:100369. <http://dx.doi.org/10.1016/j.dajour.2023.100369>, URL: <https://www.sciencedirect.com/science/article/pii/S2772662223002096>.
- [17] Engle RF. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 1982;50(4):987–1007, URL: <http://www.jstor.org/stable/1912773>.
- [18] Bollerslev T. Generalized autoregressive conditional heteroskedasticity. *J Econometrics* 1986;31(3):307–27. [http://dx.doi.org/10.1016/0304-4076\(86\)90063-1](http://dx.doi.org/10.1016/0304-4076(86)90063-1).
- [19] Glosten LR, Jagannathan R, Runkle DE. On the relation between the expected value and the volatility of the nominal excess return on stocks. *J Financ* 1993;48(5):1779–801. <http://dx.doi.org/10.1111/j.1540-6261.1993.tb05128.x>.
- [20] Peter M, Mirau S, Sinkwembe E, Kasumo C, Guambe C. Symmetric and asymmetric GARCH estimations of the impact of macroeconomic uncertainties on stock market dynamics in Tanzania. *Futur Bus J* 2025;11(1):203. <http://dx.doi.org/10.1186/s43093-025-00632-5>.
- [21] Ghoghj B, Ghodsi A. Recurrent neural networks and long short-term memory networks: Tutorial and survey. 2023, [arXiv:2304.11461](https://arxiv.org/abs/2304.11461). URL: <https://arxiv.org/abs/2304.11461>.
- [22] Wang Y, Zhang H, Huang B, Lin Z, Pang C. LSTM stock prediction model based on blockchain. *High-Confid Comput* 2025;100316. <http://dx.doi.org/10.1016/j.hcc.2025.100316>.
- [23] Adesina OS, Obokoh LO. A hybrid framework of deep learning and traditional time series models for exchange rate prediction. *Sci Afr* 2025;29:e02818. <http://dx.doi.org/10.1016/j.sciaf.2025.e02818>.
- [24] Alfianda S, Wilanata Y, Zakiyyah AY, Minor KA. Comparative model study: Gold price prediction using long short-term memory (LSTM) network and graph convolutional network (GCN). In: 2025 international conference on smart computing, IoT and machine learning. 2025, p. 1–5. <http://dx.doi.org/10.1109/SIML65326.2025.11080763>.
- [25] Livieris IE. A novel forecasting strategy for improving the performance of deep learning models. *Expert Syst Appl* 2023;230:120632. <http://dx.doi.org/10.1016/j.eswa.2023.120632>.
- [26] Livieris IE, Stavroyiannis S, Pintelas E, Pintelas P. A novel validation framework to enhance deep learning models in time-series forecasting. *Neural Comput Appl* 2020;32(23):17149–67. <http://dx.doi.org/10.1007/s00521-020-05169-y>.
- [27] Livieris IE, Pintelas P. A novel multi-step forecasting strategy for enhancing deep learning models' performance. *Neural Comput Appl* 2022;34(22):19453–70. <http://dx.doi.org/10.1007/s00521-022-07158-9>.
- [28] Bhuiyan MSM, Rafi MA, Rodrigues GN, Mir MNH, Ishraq A, Mridha M, Shin J. Deep learning for algorithmic trading: A systematic review of predictive models and optimization strategies. *Array* 2025;26:100390. <http://dx.doi.org/10.1016/j.array.2025.100390>.
- [29] Kumar M, Thenmozhi M. Forecasting stock index movement: A comparison of support vector machines and random forest. In: Indian institute of capital markets 9th capital markets conference paper. 2006, <http://dx.doi.org/10.2139/ssrn.876544>.
- [30] Shen Z, Zhang Y, Lu J, Xu J, Xiao G. A novel time series forecasting model with deep learning. *Neurocomputing* 2020;396:302–13. <http://dx.doi.org/10.1016/j.neucom.2018.12.084>.
- [31] Yahoo Finance. Markets. 2025, <https://finance.yahoo.com/markets/>. [Accessed 20 August 2025].
- [32] Bloomberg. DARSSEI:INDBloomberg quote. 2025, <https://www.bloomberg.com/quote/DARSSEI:IND>. [Accessed 20 April 2025].
- [33] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: *Advances in neural information processing systems*, vol. 30, Curran Associates, Inc.; 2017, URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf.
- [34] Goswami B, Uddin A. Significance of predictors: revisiting stock return predictions using explainable AI. *Ann Oper Res* 2025. <http://dx.doi.org/10.1007/s10479-025-06717-2>.