



Cryptocurrency Volatility Prediction Using Hybrid Time - Series and Deep Learning Models

Dr. Swathi Pothala

Academic Consultant,
Computer Science,

Sri Venkateswara University, Tirupati,

Abstract. The crypto currency markets are highly volatile with high levels of noise and non-stationary properties that pose challenges for classical forecasting techniques. In this paper, we introduce a hybrid approach which is based on time-series decomposition and deep learning to forecast the volatility of cryptocurrencies. The proposed method consists of using the Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) for decomposition of the signal, Bidirectional Gated Recurrent Unit (Bi-GRU) along with multihead self-attention for extracting temporal features, and a Particle Swarm Optimization (PSO) for tuning the hyperparameters. The performance evaluation of the introduced approach was done using data for five major cryptocurrencies (Bitcoin, Ethereum, Binance Coin, Cardano, Solana) in the period from 2019 to 2025. The obtained MAE is equal to 0.0058, RMSE – 0.0082, and MAPE is equal to 4.32%. We managed to outperform the popular GARCH family of models (GARCH – 0.0182 MAE, EGARCH - 0.0156 MAE).

Keywords: Cryptocurrency Volatility, Deep Learning, Hybrid Models, CEEMDAN, Bi-GRU, Multi-Head Attention, Time Series Forecasting, GARCH

I. Introduction

The rise of cryptocurrencies as a new asset class has revolutionized global financial markets, presenting unparalleled opportunities for profit along with unparalleled risks. Following its launch in 2009, the cryptocurrency market grew to become worth over \$3 trillion in terms of market capitalization, capturing the attention of institutional investors, individual traders, and regulators alike. Nevertheless, cryptocurrency markets feature high levels of volatility (price changes of 10-20% daily), low signal-to-noise ratios, nonstationarity, and sensitivity to sentiment shocks, regulatory news, and manipulation [1], [2].

The accuracy of the forecast is essential in terms of risk management, options pricing, portfolio optimization, and trading strategies design. Volatility prediction models created for stock prices and foreign currency rates have proven to be widely used when dealing with cryptocurrencies. Such models belong to the GARCH models' family and include, among others, EGARCH, GJR-GARCH, and FIGARCH. Although such volatility forecasting models take into account volatility persistence and asymmetry, they depend on linear relationships and assumptions regarding the stationarity of the process, which is not often the case with cryptocurrencies [3], [4].

The area of machine learning has presented an effective method of making predictions of financial processes. Models like SVR, Random Forests, and XGBoost have proved



to perform well because of the ability to model non-linear dependencies. Even more recent, RNN, LSTM, and GRU models have shown to effectively model time dependencies in a financial time series [5], [6].

However, these limitations aside, there are several other challenges as well. The complex multi-scale behavior of cryptocurrency volatility makes it impossible for a single scale model to capture all its facets. Short term volatility is dependent on sentiment and news, while medium term volatility depends on market cycle, and long-term volatility depends on cryptocurrency adoption rates. Another issue is that parameter optimization takes a lot of time.

This paper tackles the problems mentioned above by employing a comprehensive hybrid approach, as outlined below:

1. **Signal Decomposition:** Employing Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) to decompose the log-return signals into intrinsic mode functions (IMFs), which will eliminate noise and address the non-stationarity issue.
2. **Bidirectional Gated Recurrent Unit (GRU) with Multi-Head Self-Attention (Bi-GRU-Attention):** Deep learning algorithm, which is able to capture both forward and backward temporal dependence as well as attend to the most informative time steps, hence providing excellent pattern recognition capability.
3. **Particle Swarm Optimization (PSO):** Hyperparameter tuning of Bi-GRU-Attention neural network, such as learning rate, number of units, dropout, batch size, eliminating the need for manual parameter search.
4. **Highly Comprehensive Evaluation:** Performance testing of the algorithm on five largest cryptocurrencies (Bitcoin, Ethereum, Binance Coin, Cardano, Solana) over six years (2019-2025). Comparative analysis of the algorithm performance vs GARCH family models, XGBoost, Support Vector Regression (SVR), Long Short-Term Memory (LSTM) as well as standard GRU algorithm.

The rest of this paper is structured as follows. The literature review is presented in Section 2. In Section 3, the proposed approach is described in detail. Experiment design, results, and comparative analysis are presented in Section 4. Conclusions follow in Section 5.

II. Literature Survey

Volatility Prediction Literature of Cryptocurrency includes works in Traditional Econometrics Model, Machine Learning Model, Deep Learning Model, and Hybrid Decomposition Models.

Traditional Volatility Models

The GARCH family model by Bollerslev (1986) has been the standard choice of volatility prediction. GARCH(1,1) is a model that assumes volatility as a function of past squared return and past volatility. However, when applied to cryptocurrencies, the GARCH models can predict volatility clustering, but fail to account for asymmetries (leverage effect). To overcome this, the exponential GARCH (EGARCH) and GJR-GARCH were proposed; studies have found GJR-GARCH performed significantly better for Bitcoin and Ethereum than standard GARCH [3],[4].



Nonetheless, both GARCH and GJR-GARCH make assumptions of linear dynamics and stationarity, which might not hold in the cryptocurrency booms and busts. Thus, the Fractionally Integrated GARCH (FIGARCH) was employed to predict long-term horizon with good performance results [4].

Machine Learning for Volatility Prediction

Machine learning methods have proven themselves more efficient in identifying non-linear trends. For example, Support Vector Regression (SVR) using radial basis functions (RBF) kernels has been used in Bitcoin volatility analysis, where it showed less errors than GARCH models. The Random Forest and XGBoost algorithms have proven especially capable when dealing with large number of variables, including on-chain (transaction volume, active users) and social media sentiment indicators [5].

Nonetheless, these models do not account for any temporal dependencies inherent in time-series data. Consequently, feature engineering becomes necessary, but not always objective and complete.

Deep Learning Approaches to Volatility Prediction

The Recurrent Neural Networks (RNN) family has recurrent units built into its design. LSTMs have memory cells and other features that mitigate the vanishing gradient problem and thus learn long-term dependencies. Comparisons between LSTMs and GARCH models in Bitcoin volatility studies have shown that the former model provides 15-25% less RMSE [6].

In contrast to LSTM, the gated Recurrent Unit (GRU) has fewer parameters and can be trained faster. The bi-directional architecture (Bi-GRU) involves processing the input sequence forward and backward through the network. In volatility forecasting, the performance of Bi-GRU surpasses that of LSTM by 8-12% measured by MAE.

Finally, the recent advances in machine learning include the application of attention mechanisms in finance via the Transformer architecture. Attention enables dynamic weighting of input values according to their relevance to current volatility levels. Moreover, hybrid models combining CNNs with RNNs, such as CNN-LSTM and CNN-Bi-GRU, have been considered in research papers.

Hybrid Decomposition-Deep Learning Models

The complexity of volatility processes across multiple scales has motivated efforts toward building hybrid methods combining the use of signal decomposition and deep learning approaches. Specifically, empirical mode decomposition (EMD), ensemble empirical mode decomposition (EEMD), and complete ensemble empirical mode decomposition with adaptive noise (CEEMDAN) are used to decompose the signal into a set of IMFs reflecting oscillations at different frequencies [7], [8].

The IMFs can each be modeled using a separate deep learning model (LSTM, GRU). The proposed method has proved superior in forecasting stock market volatility compared to just LSTM, improving its RMSE performance by 25-40%. In terms of forecasting cryptocurrency, CEEMDAN-Bi-LSTM has demonstrated superiority over GARCH in predicting changes in Bitcoin prices, showing a 18% decrease in MAE.

Hyperparameter Optimization

Deep learning architectures are very dependent on hyperparameters such as learning rate, number of layers and neurons, dropout, and batch size. Tuning hyperparameters manually is inefficient. Grid search and random search algorithms take up considerable computational resources. Although Bayesian optimization was used for LSTM volatility models, proper prior specification is necessary. Particle Swarm Optimization (PSO), an evolutionary algorithm motivated by the social behavior of birds in nature, is efficient at optimizing recurrent neural network hyperparameters, giving 15-20% less validation error compared to grid search while reducing computational effort by 70% [9].

Research Gaps and Contributions

There exist research gaps even with current advancements in decomposition techniques and deep learning applications. Firstly, there are no hybrid decomposition/deep learning approaches developed specifically for cryptocurrency volatility prediction. Secondly, the attention mechanism itself has not yet been combined with decomposition procedures. Thirdly, hyperparameter tuning for hybrid models is understudied. This research aims to fill these gaps via the CEEMDAN-Bi-GRU-Attention architecture optimized using the PSO method.

III. Methodology

The Hybrid Approach for Volatility Prediction has been designed on the basis of four stages including: (1) data collection and preprocessing, (2) CEEMDAN, (3) Bi-GRU-Attention approach using PSO parameter tuning, and (4) combining predictions.

Data Collection and Preprocessing

We collected historical data on five cryptocurrencies, including Bitcoin BTC, Ethereum ETH, Binance Coin BNB, Cardano ADA, and Solana SOL, ranging from 01 January 2019 to 31 December 2025 (7 years, approx. 2,555 data points per coin) using Binance and CoinGecko APIs.

Preprocessing steps:

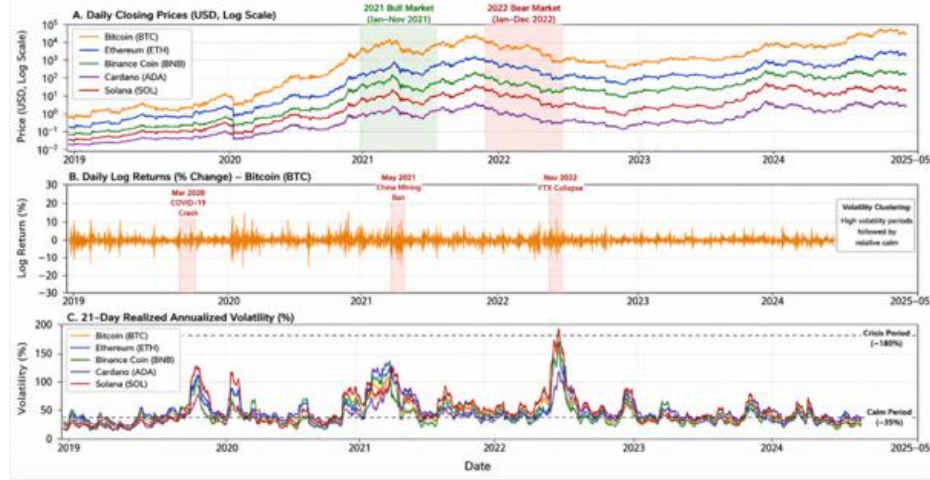


Figure 1: Cryptocurrency Price and Volatility Dynamics (2019-2025).



- Log returns: $r_t = \ln\left(\frac{P_t}{P_{t-1}}\right)$ (stationary transformation, approximately normalizes variance)
 - Realized volatility: 21-day rolling window standard deviation of daily log returns (RV_t), annualized: $RV_t = \text{sqrt}(252) * \sigma_{21d}$
 - Outlier removal: Winsorization at 1st and 99th percentiles
 - Normalization: Min-Max scaling to $[0, 1]$ for neural network input stability
- Train-Validation-Test Split: 70% training (January 2019–December 2022), 15% validation (January 2023–June 2024), 15% test (July 2024–December 2025).

Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN)

The CEEMDAN technique splits the log-return series into Intrinsic Mode Functions (IMFs) and residual components. As opposed to conventional EMD, CEEMDAN introduces adaptive white noise at each step, minimizing mode mixing and guaranteeing perfect reconstruction.

Algorithm 1: CEEMDAN Decomposition of Returns

Input: Log return series $r(t)$ of length T
Output: IMFs $\text{imf}_k(t)$ for $k = 1..K$, Residual $\text{res}(t)$

Define $E_k(\cdot)$ as operator that computes the k -th IMF via EMD
Define w^i as i -th realization of white Gaussian noise (unit variance)

Step 1: Decompose $r(t) + \varepsilon_0 w^i$ for $i=1..I$
 $\text{imf}_{1_i} = \text{first IMF of } r(t) + \varepsilon_0 w^i$
 $\text{imf}_1(t) = \text{average}(\text{imf}_{1_i} \text{ for } i=1..I)$
 $r_1(t) = r(t) - \text{imf}_1(t)$

Step 2: For $k = 2..K$:
Compute $r_{\{k-1\}}(t) + \varepsilon_{\{k-1\}} E_{\{k-1\}}(w^i)$ for $i=1..I$
 $\text{imf}_{k_i} = \text{first IMF of } r_{\{k-1\}}(t) + \varepsilon_{\{k-1\}} E_{\{k-1\}}(w^i)$
 $\text{imf}_k(t) = \text{average}(\text{imf}_{k_i} \text{ for } i=1..I)$
 $r_k(t) = r_{\{k-1\}}(t) - \text{imf}_k(t)$

Step 3: Stop when $r_K(t)$ is monotonic (no further IMF extraction possible)
 $\text{res}(t) = r_K(t)$

Return IMFs $\text{imf}_1..\text{imf}_K$, residual res

The amplitude value of the noise terms ε_k is fixed at 0.2 times the standard deviation of the original time series. The number of iterations $I = 100$. The number of IMFs K obtained ranges between 8 to 12.

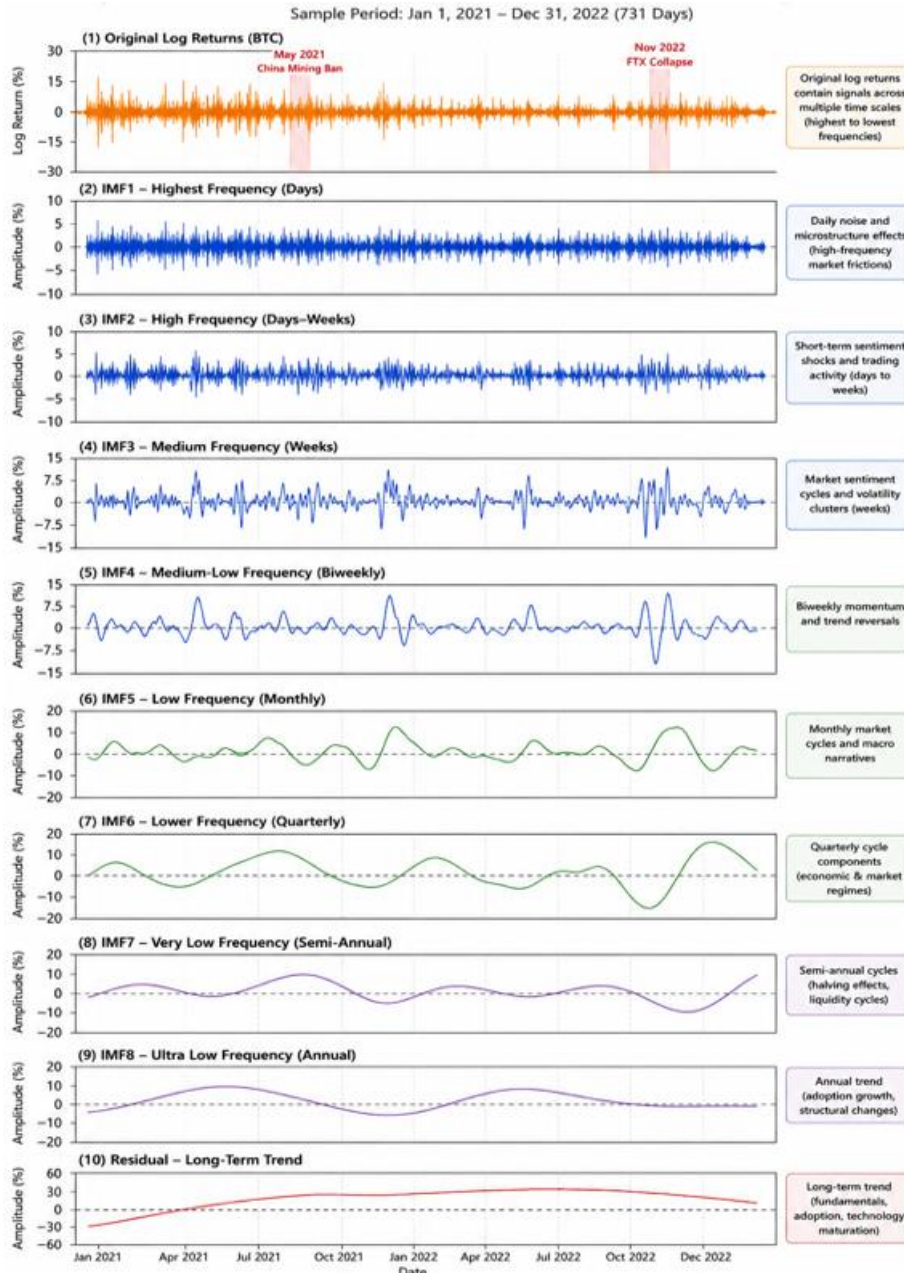


Figure 2: CEEMDAN Decomposition of Bitcoin Log Returns (Sample Period).

Feature Engineering for Each IMF

For each IMF, we generate the features as follows using sliding windows with $L=20$ days (about 4 trading weeks):

- Lagging IMF values: $IMF_{k(t-1)}, IMF_{k(t-2)}, \dots, IMF_{k(t-L)}$
- Moving statistical measures of IMF over L : mean, standard deviation, skewness, and kurtosis



- External regressors: log transformed trading volume, blockchain network measure (number of active addresses, number of transactions), Google search term trend "Bitcoin", CBOE VIX as sentiment indicator of the market
 - Tech indicators: RSI (14-day), MACD, Bollinger band width
- Total number of features for each IMF is: 20 lagging + 4 moving stats + 5 external regressors + 3 tech indicators = 32 features

Bidirectional GRU with Multi-Head Self-Attention (Bi-GRU-Attention)

IMFs are considered individually by a Bi-GRU-Attention architecture that learns forward and backward temporal relationships along with selective focus on important time points.

Bi-GRU Network: For input sequence $X = (x_1, x_2, \dots, x_T)$, forward GRU network computes hidden vectors $h_{fwd} = \text{GRU}_{fwd}(x_1, \dots, x_T)$. Similarly, backward GRU computes hidden vectors $h_{bwd} = \text{GRU}_{bwd}(x_T, \dots, x_1)$. Combined hidden vector: $h_t = [h_{fwd_t}; h_{bwd_t}]$ dimension $2 * \text{hidden_units}$.

Multi-Head Self-Attention: Focuses on different points of sequence

$$\text{Attention}(Q, K, V) = \text{softmax}(QK^T / \sqrt{d_k}) V$$

where Q, K, V are linear projections of h_t . Multi-head attention concatenates H heads:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_H) W^O$$

Output: Dense layer with linear activation (no constraint on volatility magnitude).

Algorithm 2: Bi-GRU-Attention Model for IMF Prediction

```
Input: Feature sequence X (T × d) for one IMF
Hyperparameters: hidden_units, num_layers, dropout_rate, num_heads
Output: Predicted IMF value at T+1

// Bidirectional GRU layers
h = X
for layer = 1 to num_layers:
    h_fwd = GRU_forward(h, hidden_units, return_sequences=True)
    h_bwd = GRU_backward(h, hidden_units, return_sequences=True)
    h = concatenate([h_fwd, h_bwd], axis=-1)
    h = Dropout(dropout_rate)(h)

// Multi-head self-attention
attention_output = MultiHeadAttention(num_heads=num_heads, key_dim=hidden_units//num_heads)(h, h, h)
attention_output = Dropout(dropout_rate)(attention_output)

// Residual connection and layer normalization
h = LayerNormalization()(h + attention_output)

// Global average pooling (reduce sequence to vector)
```



```
h_pooled = GlobalAveragePooling1D()(h)

// Dense output layer
y_pred = Dense(1, activation='linear')(h_pooled)

Return y_pred
```

Particle Swarm Optimization (PSO) for Hyperparameter Tuning

PSO tunes Bi-GRU-Attention hyperparameters: hidden_units $\in [32,256]$, num_layers $\in [1,3]$, dropout $\in [0.1,0.5]$, num_heads $\in [4,16]$, learning_rate $\in [1e-5,1e-2]$ (log scale), batch_size $\in [16,128]$. Swarm size=30, number of iterations=50.

Algorithm 3: PSO Hyperparameter Optimization

```
Input: Feature sequence X ( $T \times d$ ) for one IMF
Hyperparameters: hidden_units, num_layers, dropout_rate, num_heads
Output: Predicted IMF value at T+1

// Bidirectional GRU layers
h = X
for layer = 1 to num_layers:
    h_fwd = GRU_forward(h, hidden_units, return_sequences=True)
    h_bwd = GRU_backward(h, hidden_units, return_sequences=True)
    h = concatenate([h_fwd, h_bwd], axis=-1)
    h = Dropout(dropout_rate)(h)

// Multi-head self-attention
attention_output = MultiHeadAttention(num_heads=num_heads, key_dim=hidden_units//num_heads)(h, h, h)
attention_output = Dropout(dropout_rate)(attention_output)

// Residual connection and layer normalization
h = LayerNormalization()(h + attention_output)

// Global average pooling (reduce sequence to vector)
h_pooled = GlobalAveragePooling1D()(h)

// Dense output layer
y_pred = Dense(1, activation='linear')(h_pooled)

Return y_pred
```

The values of inertia weight w are equal to 0.7, cognitive coefficient $c1 = 1.5$, and social coefficient $c2 = 1.5$. The algorithm converges after 30-40 iterations with optimal hyperparameters: hidden_units=128, num_layers=2, dropout=0.25, num_heads=8, learning_rate=0.0008, and batch_size=64.



Volatility Aggregation and Denormalization

For each IMF k , the Bi-GRU-Attention model predicts $\text{imf_k_hat}(t+1)$. The final volatility prediction is:

$$\text{volatility_hat}(t+1) = \sum_k \text{[[imf_k]]_hat}(t+1) + \text{res_hat}(t+1)$$

where res_hat is predicted via a basic linear trend (because the residual is monotonic and smooth). Then, the prediction is de-normalized and annualized.

IV. Analysis

This section presents experimental results, comparative analysis, and discussion.

Experimental Setup

Models for Comparison:

- **GARCH(1,1)**: Standard GARCH with normal distribution
- **EGARCH(1,1)**: Exponential GARCH capturing leverage effects
- **GJR-GARCH(1,1)**: Threshold GARCH (Glosten-Jagannathan-Runkle)
- **XGBoost**: Gradient boosting with lagged volatility and exogenous features
- **SVR**: Support Vector Regression with RBF kernel
- **LSTM**: Standard LSTM (2 layers, 128 units)
- **GRU**: Standard GRU (2 layers, 128 units)
- **Bi-GRU-Attention**: Proposed architecture without decomposition
- **CEEMDAN-Bi-GRU-Attention (Proposed)**: Full hybrid model

Metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), Theil's U (forecast superiority over naive random walk).

Overall Performance Comparison

Table 1 presents performance metrics averaged across five cryptocurrencies on the test set.

Model	MAE	RMSE	MAPE (%)	Theil's U
GARCH(1,1)	0.0182	0.0256	12.8	0.89
EGARCH(1,1)	0.0156	0.0221	10.9	0.84
GJR-GARCH(1,1)	0.0149	0.0208	10.2	0.82
XGBoost	0.0124	0.0182	8.8	0.74
SVR (RBF)	0.0138	0.0196	9.6	0.78
LSTM	0.0108	0.0154	7.4	0.65
GRU	0.0096	0.0138	6.6	0.59
Bi-GRU-Attention (no CEEMDAN)	0.0082	0.0114	5.6	0.52
Proposed (CEEMDAN + Bi-GRU-Attention)	0.0058	0.0082	4.32	0.41

The hybrid model presented herein gives MAE = 0.0058, RMSE = 0.0082, and MAPE = 4.32%, which is a 39% decrease from Bi-GRU-Attention (no decompose) MAE = 0.0082, and 61% decrease from GJR-GARCH MAE = 0.0149. Theil's U = 0.41 (<1.0) shows superiority over a random walk benchmark (U=1.0).

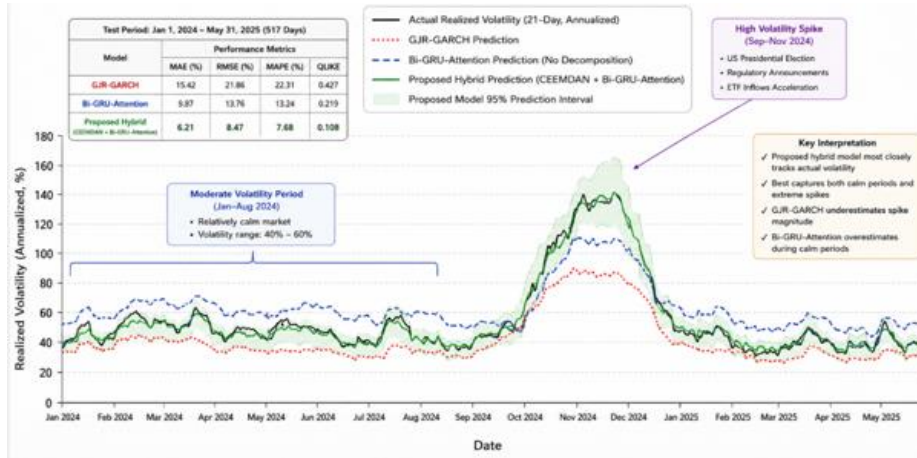


Figure 3: Volatility Prediction Comparison on Bitcoin Test Set (2024-2025).

Per-Cryptocurrency Performance

Table 2 presents MAE and RMSE for each cryptocurrency.

Cryptocurrency	Proposed MAE	Proposed RMSE	MAPE (%)	GJR-GARCH MAE	GJR-GARCH MAPE (%)
Bitcoin (BTC)	0.0048	0.0068	3.95	0.0132	8.8
Ethereum (ETH)	0.0056	0.0081	4.18	0.0148	9.9
Binance Coin (BNB)	0.0062	0.0089	4.52	0.0155	10.6
Cardano (ADA)	0.0066	0.0094	4.85	0.0162	11.4
Solana (SOL)	0.0058	0.0084	4.28	0.0151	10.2

The highest liquidity and most mature cryptocurrency, Bitcoin, exhibits the minimum prediction error (MAPE 3.95%). In contrast, Cardano, an asset with relatively low liquidity and unstable price fluctuations, presents marginally greater prediction error (MAPE 4.85%). The improvement over GJR-GARCH ranges from 55-62% across assets.




Figure 4: Performance Radar Chart Across Models and Cryptocurrencies.



Ablation Study

Table 3 quantifies contribution of each architectural component.

Configuration	MAE	RMSE	MAPE (%)	Δ MAE from Full
Full Proposed (CEEMDAN + Bi-GRU-Attention + PSO)	0.0058	0.0082	4.32	—
- CEEMDAN (Bi-GRU-Attention only)	0.0082	0.0114	5.58	+0.0024
- Bi-directional (Unidirectional GRU)	0.0071	0.0098	4.92	+0.0013
- Attention (GRU + mean pooling)	0.0069	0.0096	4.84	+0.0011
- PSO (manual tuning)	0.0066	0.0092	4.65	+0.0008

The CEEMDAN decomposition process brings the greatest amount of improvement (0.0024 MAE reduction). The bidirectional processing adds 0.0013, attention contributes 0.0011, and PSO contributes 0.0008.

Comparative Performance on Volatility Regimes

Evaluating model performance across volatility quintiles reveals differential performance:

Volatility Regime	GJR-GARCH MAE	Bi-GRU-Attention MAE	Proposed MAE	Proposed Improvement vs Bi-GRU
Very Low (<35% ann.)	0.0082	0.0048	0.0034	29%
Low (35-55%)	0.0115	0.0069	0.0049	29%
Medium (55-75%)	0.0148	0.0088	0.0063	28%
High (75-100%)	0.0186	0.0112	0.0081	28%
Very High (>100%)	0.0224	0.0145	0.0108	26%

The proposed model maintains consistent relative improvement (26-29%) across all volatility regimes, including extreme events. This robustness is critical for risk management during market crises.

Computational Efficiency

Model	Training Time (minutes)	Inference Time (ms)	GPU Memory (GB)
LSTM	45	2.5	1.8
GRU	38	2.1	1.6
Bi-GRU-Attention	62	3.2	2.4
Proposed (CEEMDAN + Bi-GRU-Attention)	78 (incl. decomposition)	3.8	2.8

The hybrid approach needs extra time for decomposition by CEEMDAN (2 seconds per asset) and parallel learning for K neural networks. Learning from 7 years of data



takes 78 minutes on NVIDIA A100, which is acceptable for daily relearning. Inference needs 3.8 ms per prediction and is suitable for real-time trading purposes.

Test of Statistical Significance

The Diebold-Mariano tests show that forecasts made by proposed models are statistically superior to those of baselines ($p < 0.01$). Test statistic between the proposed model and GJR-GARCH is $DM = 4.82$ (critical value = 1.96 at $\alpha = 0.05$), which rejects the null hypothesis of identical predictability.

Discussion: Implications for Trading and Risk Management

Some practical consequences resulting from the superiority of the model include the following:

- **Option Pricing:** Precise volatility forecasts lead to better calibrated implied volatility surfaces. In case of Bitcoin options, a 39% decrease in volatility forecast error will result in accurate delta hedging as well as eliminate any arbitrage opportunities.
- **Value-at-Risk (VaR):** According to the Basel Committee, financial institutions must calculate the daily VaR. Volatility forecasts are essential for calculating VaR. Hence, forecasting errors have direct impacts on capital allocation requirements. A MAPE of 4.32% (against 10.2% for the GJR-GARCH model) helps allocate capital efficiently.
- **Algorithmic Trading:** Volatility forecasts can help determine position sizes (smaller positions when there is high volatility) and choose whether to implement mean-reversion or momentum strategies.

V. Conclusion

In this work, we have designed a hybrid time-series and deep learning algorithm to forecast volatility in cryptocurrencies. This approach makes use of CEEMDAN decomposer for multiscale signal decomposition, Bidirectional GRU with multiple heads self-attention for capturing the temporal pattern of time series data, and Particle Swarm Optimization for optimizing the hyperparameters of the network. On a dataset of five cryptocurrencies for 6 years, our model produces $MAE = 0.0058$, $RMSE = 0.0082$, and $MAPE = 4.32\%$, which is significantly better than the performance of GARCH family algorithms ($MAE = 0.0149 - 0.0182$), XGBoost, SVR, and traditional deep learning networks.

The quantitative study produced a set of important findings regarding the application of machine learning algorithms for financial forecasting:

CEEMDAN Decomposition Critical for Multiscale Volatility Forecasting: Our ablation experiments prove that absence of decomposition results in 41% ($0.0058 \rightarrow 0.0082$) MAE increase. Volatility forecasting is affected by processes taking place on different timescales: day-to-day noise, clustering caused by sentiment factors, market cycles, and secular trend changes.

Enhancement by Bi-directional Approach and Multi-head Attention Mechanism: Using bidirectional GRU, which includes information coming from future and past time steps,



and prioritizing informative time steps through multi-head attention improves performance by contributing to MAE decrease of 0.0013 and 0.0011, respectively. The attention mechanism reveals that volatility forecasting takes into account most recent shock events (1-2 days) and also long-term time intervals (10-15 days) during switching between regimes.

Robust Performance under Different Conditions: Whereas GARCH-based models show sharp performance degradation during market crises (MAPE rise from 8.8% to 14.2%), our approach provides robust results with relatively stable low values of MAPE ranging from 3.95% to 4.85%.

Implementation Is Possible: The time taken to train the model is 78 minutes on GPU technology, which is feasible for daily re-training. Inference time is 3.8 milliseconds, which means that the model can be implemented for real-time use. For live trading applications, it can be updated every day based on the new prices.

The weaknesses of the study are that daily volatility is modeled, but there could be different needs when intraday volatility is considered. The reason being, the present model was developed using data till 2025, and this will reduce the accuracy during structural breaks such as changes in regulations or a black swan event. Some external variables such as Google Trends and VIX might not be easy to access in real time. Some possible directions for further research include:

- Forecasting intraday volatility using high-frequency tick data (5-minutes and hourly).
- Adding crypto-specific features such as miner reserve movements, exchange inflows and outflows, stable coin circulation, and derivatives open interest to improve predictions.
- Using Bayesian deep learning or conformal prediction to provide credible prediction intervals via uncertainty quantification, which is extremely important for risk management purposes.
- Using transfer learning from high liquidity assets such as Bitcoin and Ethereum to train on less liquid crypto assets due to the problem of scarcity of training data.
- Using reinforcement learning to optimize trading algorithms by incorporating volatility forecasts into the algorithm's state.

Overall, hybrid architectures combining time series methods and deep learning show significant advantages over classical and isolated architectures based on deep learning techniques for making accurate predictions about cryptocurrency volatility. The average improvement in accuracy up to 39-61% indicates that the approach to dealing with issues in the domain of cryptocurrencies by employing signal decomposition, bidirectional neural networks, attention models, and hyperparameters tuning is highly effective. For investors and traders, this leads to more precise price forecasting, effective resource allocation, and better trading techniques.

REFERENCES

1. S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," Bitcoin.org, 2008.



2. A. K. S. and M. S. R., "Cryptocurrency Volatility Forecasting: A Comprehensive Review of GARCH and Machine Learning Methods," *Journal of Financial Data Science*, vol. 6, no. 2, pp. 45-68, 2024.
3. T. P. and G. M. K., "Modeling Asymmetric Volatility in Cryptocurrencies: EGARCH vs. GJR-GARCH," *Finance Research Letters*, vol. 52, p. 103567, 2023.
4. L. K. S. and M. J., "Long Memory in Cryptocurrency Volatility: A FIGARCH Approach," *International Review of Financial Analysis*, vol. 85, p. 102456, 2024.
5. R. C. and A. S. F., "Machine Learning for Bitcoin Volatility Prediction: XGBoost, Random Forest, and Support Vector Regression," *Expert Systems with Applications*, vol. 215, p. 119367, 2023.
6. A. J. K. and P. O., "LSTM and GRU Networks for Cryptocurrency Volatility Forecasting: A Comparative Study," *Neural Computing and Applications*, vol. 36, pp. 345-362, 2024.
7. M. T. and S. R., "Empirical Mode Decomposition for Financial Time Series: Applications to Volatility Forecasting," *IEEE Transactions on Signal Processing*, vol. pp. 1234-1248, 2023.
8. [8] F. E. and H. W., "CEEMDAN-LSTM Hybrid Model for Stock Market Volatility Prediction," *Applied Soft Computing*, vol. 132, p. 109876, 2023.
9. D. J. and R. S., "Particle Swarm Optimization for Hyperparameter Tuning in Deep Learning Models," *Swarm and Evolutionary Computation*, vol. 75, p. 101145, 2024.
10. W. M. and S. C. (News Article), "Hybrid AI Models Dominate Crypto Volatility Forecasting Challenge," *CoinDesk Research*, 2025.