

---

## | RESEARCH ARTICLE

# Cryptocurrency Volatility Forecasting Using Transformer-Based Deep Learning Models and On-Chain Metrics

Mainuddin Adel Rafi<sup>1</sup>, S M Iftekhar Shaboj<sup>2</sup>, Iftekhar Rasul<sup>3</sup>, Md Kauser Miah<sup>4</sup>, Iftekhar Rasul<sup>4</sup>, Md Redwanul Islam<sup>5</sup>, Abir Ahmed<sup>6</sup>

<sup>1</sup>Master of Science Information System, Pacific State University, USA

<sup>2</sup>Master of Accountancy, University of Tulsa, Tulsa, Oklahoma, USA

<sup>3</sup>Information Technology Management, St. Francis College, USA

<sup>4</sup>Department of Computer and Information Science, Gannon University, PA, USA

<sup>5</sup>Department of Finance & Financial Analytics, University of New Haven, West Haven, CT, USA

<sup>6</sup>Department of Information Technology, Washington University of Science & Technology, VA, USA

**Corresponding Author:** Abir Ahmed, **E-mail:** [abira.student@wust.edu](mailto:abira.student@wust.edu)

---

## | ABSTRACT

Cryptocurrencies have emerged as highly dynamic digital assets, characterized by extreme price volatility and driven by both speculative behavior and network-level activities. Traditional volatility forecasting methods, including GARCH and LSTM-based models, often fall short in capturing the complex, nonlinear, and temporal dependencies inherent in crypto markets. This paper proposes a novel deep learning framework that leverages Transformer-based architectures originally designed for natural language processing to forecast short-term cryptocurrency volatility with enhanced accuracy and temporal sensitivity. By incorporating a comprehensive set of on-chain metrics (such as transaction volume, wallet activity, miner behavior, and token circulation), our model captures both market sentiment and blockchain-level dynamics. The proposed Transformer model is benchmarked against LSTM and GRU networks using Bitcoin and Ethereum datasets spanning multiple market cycles. Experimental results show that the Transformer-based model outperforms recurrent architectures in predicting both realized and implied volatility, particularly during high-turbulence periods. These findings suggest that attention mechanisms, combined with on-chain data, provide a powerful tool for managing risk and making informed decisions in the rapidly evolving digital asset ecosystem.

## | KEYWORDS

Cryptocurrency; Volatility Forecasting; Transformer; Deep Learning; Bitcoin; Ethereum; On-Chain Metrics; Attention Mechanism; Financial Time Series; Blockchain Analytics

## | ARTICLE INFORMATION

**ACCEPTED:** 04 February 2024

**PUBLISHED:** 20 February 2024

**DOI:** 10.32996/jefas.2024.6.1.12

---

## 1. Introduction

The rapid rise of cryptocurrencies such as Bitcoin (BTC) and Ethereum (ETH) has fundamentally transformed the landscape of global financial markets. These decentralized digital assets, built on blockchain technology, have attracted widespread attention from retail investors, institutional players, policymakers, and researchers alike. What differentiates cryptocurrencies from traditional financial instruments is not only their technological underpinnings but also their unique market behavior, especially in terms of extreme and often unpredictable price movements. Unlike conventional financial markets which are influenced by structured

trading hours, monetary policies, and economic cycles cryptocurrency markets operate continuously, 24 hours a day and 7 days a week, across decentralized exchanges with varying liquidity profiles. These characteristics result in exceptionally high volatility, often triggered by sudden regulatory news, protocol upgrades, or activity on the blockchain itself [1].

Volatility, defined as the statistical measure of the dispersion of returns, plays a central role in risk management, option pricing, and strategic asset allocation. In traditional finance, well-established models such as GARCH (Generalized Autoregressive Conditional Heteroskedasticity) and its extensions have been used for decades to forecast asset volatility. However, the cryptocurrency domain introduces new layers of complexity. The nascent nature of these assets, combined with their speculative and sentiment-driven market dynamics, renders many conventional econometric tools inadequate. In particular, traditional models often fail to capture the nonlinear dependencies, non-stationarity, and multi-scale temporal patterns observed in crypto price series [2]. Moreover, the volatility in cryptocurrency markets is not only a result of price action but also deeply tied to blockchain-level activities known as on-chain metrics. These include indicators such as transaction volume, miner inflow/outflow, active wallet addresses, gas fees (in Ethereum), and hash rates, all of which reflect the underlying health, usage, and sentiment of a particular cryptocurrency network. Ignoring these signals while attempting to forecast volatility is akin to analyzing equity markets without considering earnings reports or macroeconomic indicators. Therefore, integrating both market-level and on-chain data is critical to constructing accurate and reliable volatility forecasting systems. Recent advancements in artificial intelligence and deep learning, particularly in time series modeling, have introduced more sophisticated approaches for volatility prediction. Recurrent Neural Networks (RNNs), especially Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models, have shown promising results in capturing temporal sequences. However, these models are inherently limited in learning long-range dependencies due to their sequential nature. More recently, Transformer-based architectures, originally developed for natural language processing (NLP), have revolutionized sequential data modeling by replacing recurrence with self-attention mechanisms, allowing the model to capture dependencies across entire sequences regardless of distance [3]. Despite their demonstrated success in NLP and speech domains, Transformers have only recently begun to gain traction in financial time series forecasting. In particular, their application to cryptocurrency volatility prediction remains underexplored, and there is a substantial gap in the literature addressing the integration of on-chain metrics within these deep learning pipelines. Given the complexity, noise, and unpredictability of crypto markets, we posit that the attention mechanism of Transformers combined with rich on-chain data can significantly improve the accuracy and adaptability of volatility forecasting systems. This study thus emerges at the intersection of three powerful domains: blockchain analytics, financial forecasting, and deep learning innovation. By bridging these fields, our research aims to deliver a robust, interpretable, and forward-looking framework for forecasting short-term cryptocurrency volatility, with potential applications in portfolio risk management, algorithmic trading, and regulatory compliance.

### **1.1 Background and Motivation**

Cryptocurrency markets are fundamentally different from traditional financial systems. They operate 24/7, are highly speculative, and lack centralized regulatory oversight. This results in price movements that are both rapid and erratic, making standard econometric volatility forecasting models such as GARCH and its variants less effective in capturing complex temporal dependencies [2]. Moreover, traditional models often ignore on-chain metrics, which provide valuable insights into network activity, user behavior, and market sentiment from the blockchain itself.

Recent advances in deep learning have shown great promise in forecasting financial time series. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures, have been widely adopted for time series analysis due to their ability to model sequential dependencies [3]. However, these models suffer from limitations in handling long-term dependencies and parallelization. In contrast, Transformer-based models, which rely on self-attention mechanisms, have outperformed RNNs in various domains, particularly in natural language processing (NLP) and speech recognition tasks [4].

Motivated by the potential of Transformers to handle complex temporal relationships, this study investigates their application in cryptocurrency volatility forecasting, enriched by on-chain data. We hypothesize that combining the temporal modeling strength of Transformers with blockchain-derived features can significantly improve volatility prediction accuracy, especially in turbulent market periods.

### **1.2 Problem Statement and Research Gap**

While volatility forecasting in traditional finance is well-studied, the same cannot be said for the highly nonlinear and non-stationary world of digital assets. Existing deep learning studies in this domain primarily rely on price data, trading volume, and sentiment indicators, neglecting the rich trove of real-time on-chain metrics such as active addresses, hash rate, gas fees, and token supply dynamics [5]. Furthermore, despite the recent popularity of Transformers, their application in cryptocurrency volatility forecasting remains limited. Most current approaches either use LSTM-based architectures or ensemble statistical models, which may underperform in capturing long-range dependencies inherent in crypto price dynamics [6]. There is a pressing need to

evaluate whether Transformer-based models originally developed for NLP can be successfully adapted for forecasting high-frequency, high-volatility financial series in the presence of heterogeneous data inputs.

### **1.3 Objectives and Scope of the Study**

This paper aims to address the research gap by proposing and evaluating a Transformer-based deep learning framework for short-term cryptocurrency volatility forecasting. The objectives are fourfold:

- To design and implement a Transformer model tailored for financial time series.
- To integrate and preprocess a diverse set of on-chain and market-based features from Bitcoin and Ethereum networks.
- To benchmark the proposed model against established methods, including LSTM, GRU, and GARCH-family models.
- To analyze model performance during both stable and volatile market conditions.

The scope of this study is limited to short-term volatility forecasting (e.g., 1-day ahead realized volatility) using historical daily data. We focus on Bitcoin and Ethereum as representative assets due to their high liquidity and extensive on-chain activity.

### **1.4 Significance and Contributions**

This research contributes to the growing field of crypto-financial analytics in multiple ways. First, it demonstrates the applicability of advanced deep learning architectures like Transformers in forecasting financial volatility. Second, it highlights the predictive power of on-chain metrics, offering a new dimension to crypto market analysis beyond traditional trading data. Finally, this work lays the groundwork for real-time, explainable, and high-frequency crypto forecasting systems that can assist both retail and institutional participants in managing financial risk more effectively [7].

## **2. Related Work**

Forecasting financial volatility has been a long-standing problem in econometrics, yet the application of deep learning to cryptocurrency volatility is relatively new and evolving. This section reviews the relevant literature across three key strands: traditional models, deep learning models in cryptocurrency forecasting, and the emerging role of Transformer-based architectures. Each subsection provides insights into the current limitations and how our proposed work addresses these gaps.

### **2.1 Traditional Models in Volatility Forecasting**

Conventional volatility forecasting has largely relied on statistical models such as ARCH and GARCH. These models, and their extensions like EGARCH and TGARCH, have been widely applied in financial markets for decades due to their interpretability and ability to model time-varying volatility [8]. However, their linear assumptions and inability to capture complex, nonlinear relationships limit their effectiveness in cryptocurrency markets, which are known for their chaotic behavior and abrupt regime shifts. Several researchers have attempted to improve traditional models by integrating macroeconomic indicators or sentiment analysis into GARCH-type models. For instance, Chen et al. [9] incorporated Twitter sentiment into a GARCH framework and demonstrated modest gains in forecasting crypto volatility. Nonetheless, these models remain limited in adaptability and data flexibility, particularly when dealing with heterogeneous sources such as on-chain metrics.

### **2.2 Deep Learning Approaches in Cryptocurrency Prediction**

Deep learning, particularly Recurrent Neural Networks (RNNs), has shown considerable promise in capturing sequential patterns in financial time series. Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs) are frequently employed for cryptocurrency price prediction. Nelson et al. [10] demonstrated that LSTMs outperformed ARIMA models in forecasting Bitcoin price volatility, marking a pivotal shift toward deep learning applications in this space. Similarly, Kim & Kim [11] used GRUs to model short-term volatility in cryptocurrencies and reported improved prediction accuracy compared to standard econometric methods. However, despite their utility, RNNs and their variants suffer from drawbacks such as vanishing gradients, limited memory capacity for long sequences, and slow training times. More critically, these models typically rely only on price and volume data, thereby neglecting on-chain metrics—a data source that is uniquely available in blockchain environments.

### **2.3 Transformer-Based Architectures in Financial Time Series**

Transformers, introduced by Vaswani et al. [12] and later popularized through models like BERT [13], have redefined the state-of-the-art in natural language processing. Their core innovation self-attention mechanisms allows for parallelized computation and better handling of long-range dependencies, which are especially beneficial in noisy, high-frequency time series such as crypto markets. Recent studies have begun applying Transformers to financial prediction tasks. Qin et al. [14] proposed a Transformer-

based model for stock price prediction and demonstrated its superiority over LSTM and GRU baselines. Zhang et al. [15] extended this by incorporating hybrid attention into a CNN-LSTM framework for short-term volatility estimation. Liu et al. [16] took this a step further by integrating on-chain metrics with Transformer-based models for Ethereum volatility forecasting. Their findings suggest that incorporating blockchain-derived features such as gas fees, wallet activity, and smart contract volume significantly improves prediction accuracy. However, comprehensive frameworks combining realized volatility, Transformer encoders, and on-chain analytics for multiple cryptocurrencies remain scarce in the literature.

## 2.4 Summary and Research Positioning

Table 1 provides a consolidated overview of notable contributions in cryptocurrency volatility forecasting. While deep learning has shown potential in outperforming classical models, the integration of Transformer-based architectures and on-chain data is still in its infancy. Most prior studies are either limited to one asset class, rely only on price-based features, or use sequential models that are less efficient for long-range dependencies. Our work builds upon these foundations by proposing a robust Transformer-based deep learning model that fuses market-level and on-chain metrics for predicting short-term volatility in both Bitcoin and Ethereum markets. This hybrid approach positions our research at the frontier of crypto-financial forecasting and deep learning interpretability.

Table 1: Summary of Related Work in Cryptocurrency Volatility Forecasting

| Study                | Methodology/Model                         | Contribution                                                 |
|----------------------|-------------------------------------------|--------------------------------------------------------------|
| Nelson et al. (2017) | LSTM for Bitcoin price prediction         | Demonstrated LSTM superiority over ARIMA for BTC             |
| Kim & Kim (2019)     | GRU on cryptocurrency volatility          | Applied GRU to crypto volatility, limited feature set        |
| Chen et al. (2020)   | GARCH models with sentiment features      | Integrated sentiment data in volatility modeling             |
| Qin et al. (2021)    | Transformer-based time series forecasting | Proposed Transformer model for financial series              |
| Zhang et al. (2022)  | Attention-enhanced hybrid CNN-LSTM        | Hybrid attention model improved short-term accuracy          |
| Devlin et al. (2018) | BERT: Bidirectional Transformer (NLP)     | Introduced attention for sequence modeling in NLP            |
| Liu et al. (2023)    | Transformer + On-chain metrics for ETH    | Applied attention on-chain + market data for ETH forecasting |

## 3. Methodology

This section outlines the comprehensive methodological framework adopted in this study to forecast cryptocurrency volatility using Transformer-based models and on-chain metrics. The framework is divided into five phases: data collection and preprocessing, feature engineering, model architecture, training procedure, and evaluation metrics. Each phase is designed to ensure both predictive accuracy and robustness in real-world forecasting settings.

### 3.1 Data Collection and Preprocessing

The dataset used in this study encompasses daily historical data from January 2018 to March 2024 for two leading cryptocurrencies Bitcoin (BTC) and Ethereum (ETH). This period includes high-volatility cycles such as the 2021 bull market and the 2022 correction, ensuring that the model is tested across various regimes. The data is sourced from two main domains: market-level data and on-chain blockchain metrics. Market data, including OHLC prices, volume, and realized volatility, was retrieved from CoinMarketCap and Yahoo Finance. Meanwhile, blockchain-specific on-chain data such as active addresses, gas usage, hash rate, transaction

volume, miner flows, and token circulation was accessed through Glassnode, Etherscan, and Coin Metrics APIs [17]. Before model training, we performed rigorous preprocessing to ensure data integrity and compatibility. Missing values were filled using forward and backward fill techniques, and outliers were removed through interquartile filtering. Numerical features exhibiting heavy-tailed distributions—such as transaction fees and wallet activity—were log-transformed to stabilize variance. All features were then normalized using z-score standardization to harmonize input ranges across model inputs and facilitate gradient-based learning.

### 3.2 Feature Selection and Engineering

Feature engineering is critical for enriching the dataset with signals that reflect both market behavior and blockchain dynamics. From the price series, we derived conventional technical indicators including Exponential Moving Averages (EMA), Bollinger Bands, Relative Strength Index (RSI), and MACD. These indicators are commonly used by traders to capture momentum, trend, and volatility. In parallel, a robust set of on-chain metrics was constructed, such as the NVT ratio (Network Value to Transactions), active addresses, miner inflows, token age consumed, and gas fees. These features serve as proxies for blockchain congestion, user participation, and internal capital flow, all of which correlate with price dynamics and volatility [18, 35, 36, 37, 38]. The target variable realized volatility (RV) was calculated as the rolling 5-day standard deviation of log returns:

$$RV_t = \sqrt{\frac{1}{n} \sum_{i=1}^n (r_t - i)^2}, \quad (1)$$

Where  $r_t = \log\left(\frac{p_t}{p_{t-1}}\right)$  is the daily log return and  $n=5$ . This target was lagged to create a one-step-ahead prediction task, simulating daily forecasting in operational settings.

### 3.3 Model Architecture: Transformer Design

The primary model adopted in this study is based on the Transformer encoder architecture, adapted for multivariate financial time series prediction. Transformers were selected due to their capacity to model long-range dependencies through self-attention, which surpasses the memory limitations and sequential constraints of RNN-based models [19, 39, 40]. The model input includes time-lagged multivariate sequences of both market and on-chain features. An input embedding layer transforms these numerical sequences into dense vectors. To capture temporal order, positional encoding is added, allowing the model to distinguish the sequence of time steps despite the lack of recurrence. The encoder stack contains multiple multi-head self-attention layers, which allow the model to simultaneously consider relationships across all time steps and feature channels. Following the attention layers are feedforward neural networks activated by ReLU and stabilized by Layer Normalization and Dropout (rate = 0.2). The final output layer uses a single node to forecast the next-step realized volatility value. The model was implemented using TensorFlow and PyTorch, running on a system with an NVIDIA RTX 3090 GPU, 64 GB RAM, and an Intel i9 CPU for training acceleration.

### 3.4. Training and Optimization Strategy

For optimization, the Adam optimizer was employed with an initial learning rate of 0.001 and Mean Squared Error (MSE) as the loss function. Training was conducted using a batch size of 32 and a maximum of 100 epochs. Early stopping was applied with a patience of 10 epochs to avoid overfitting. The data was split into 80% training and 20% testing based on chronological order to maintain time consistency. To benchmark our Transformer model, we also trained traditional models including GARCH(1,1), LSTM, and GRU, each using the same lagged feature sequences and hyperparameter tuning via grid search [20]. We experimented with input lag sequences of 10, 30, and 60 days, analyzing the effect of short-, medium-, and long-term dependencies. The Transformer showed superior stability across these lags compared to recurrent models, especially in volatile periods.

### 3.5 Evaluation Metrics

To rigorously assess the performance of the proposed Transformer-based volatility forecasting model, this study employed a comprehensive set of quantitative evaluation metrics. These metrics were chosen to capture not only the accuracy of the predicted volatility values but also the model's ability to anticipate the direction and variability of market movements both of which are essential for practical applications in trading, risk management, and financial decision-making. The primary metric used for training and optimization was the Mean Squared Error (MSE). MSE quantifies the average of the squared differences between predicted

and actual values and penalizes large errors more severely, making it suitable for volatility forecasting where prediction errors can have amplified financial consequences. To enhance interpretability, especially for comparing across models, we also report the Root Mean Squared Error (RMSE), which is the square root of MSE and thus maintains the same scale as the target variable. These metrics are particularly useful for measuring the accuracy of volatility magnitude predictions, which is vital for derivatives pricing, portfolio hedging, and Value-at-Risk (VaR) computations.

In addition to squared-error-based metrics, Mean Absolute Error (MAE) was employed to provide a linear error perspective. Unlike MSE and RMSE, MAE treats all errors equally, offering a more balanced view when the dataset includes extreme volatility spikes or outlier events. MAE is especially valuable in financial markets where minimizing average error over time is often preferable to penalizing occasional extreme errors disproportionately. Another critical metric used in this study is the R-squared ( $R^2$ ) coefficient, also known as the coefficient of determination.  $R^2$  measures the proportion of variance in the observed volatility that is explained by the model's predictions. A higher  $R^2$  value indicates better explanatory power and is useful for understanding the overall model fit. However, due to the noisy and non-stationary nature of financial time series especially in the cryptocurrency domain  $R^2$  must be interpreted cautiously, as even moderate values can represent substantial improvements over baseline models like naive mean or historical average predictors.

Lastly, and perhaps most important for real-world trading scenarios, we included the Directional Accuracy (DA) metric. DA measures the proportion of times the model correctly predicts the direction of change in volatility whether it increases or decreases from one day to the next. This binary evaluation is crucial for many applications, such as options hedging or volatility arbitrage, where the direction of volatility matters more than its precise magnitude. For instance, predicting an increase in volatility may trigger a trader to buy options or hedge a position in anticipation of larger price swings. As such, DA serves as a practical benchmark for determining whether the model can support strategic financial decisions. All of the above metrics were computed for both the in-sample (training) and out-of-sample (testing) datasets to assess generalization. The Transformer model's performance was further benchmarked against LSTM, GRU, and GARCH models under identical conditions. Together, these metrics form a robust framework for evaluating not just the numerical fidelity of the model's outputs, but also its practical utility in real-world cryptocurrency trading and risk analysis.

### **3.6 Experimental Hardware Setup**

The computational demands of training Transformer-based deep learning models particularly for time series forecasting tasks involving multivariate inputs and long sequence lengths necessitate a robust and scalable hardware environment. To meet these requirements, all experiments in this study were conducted on a high-performance workstation equipped with an Intel Core i9-12900K CPU running at 3.20 GHz, 64 GB of DDR5 RAM, and an NVIDIA GeForce RTX 3090 GPU with 24 GB of dedicated video memory. This configuration ensured efficient handling of large datasets and enabled the parallel computation capabilities required for Transformer training, especially during backpropagation through multiple self-attention layers. The software environment was configured using Ubuntu 22.04 LTS (64-bit) as the operating system, ensuring stability and compatibility with machine learning libraries. All modeling and experimentation were implemented in Python 3.10, utilizing key open-source frameworks such as TensorFlow 2.12, Keras, and PyTorch 2.0. These libraries were selected for their support of GPU acceleration and high-level APIs for Transformer architecture implementation [31, 32, 33, 34].

Additionally, NumPy, Pandas, and Matplotlib were used for data manipulation and visualization, while Scikit-learn was employed for auxiliary tasks such as feature selection, normalization, and baseline model implementation. To manage experiments, hyperparameter tuning, and model version control, we adopted MLflow, a lightweight open-source platform that enabled us to log training metrics, track model performance across versions, and reproduce experiments. GPU utilization was optimized by enabling CUDA 11.7 and cuDNN 8.6 libraries, which ensured compatibility with the TensorFlow and PyTorch backends and maximized training speed. Training time varied depending on sequence length, model complexity, and input dimensionality. On average, the Transformer model required approximately 25 to 30 minutes to converge per training run for a 60-day input sequence over 100 epochs.

In contrast, LSTM and GRU models required slightly longer due to their sequential nature and lack of full parallelization. GARCH models, implemented using the arch package in Python, executed significantly faster but lacked the capacity to handle high-

dimensional, multivariate inputs. This robust experimental setup allowed for seamless training, evaluation, and comparative analysis of deep learning models under consistent computational conditions, ensuring the reliability and reproducibility of our findings.

#### 4. Results and Analysis

This section presents a detailed analysis of the experimental outcomes obtained from the Transformer-based model and its benchmark counterparts (GARCH, LSTM, GRU). We divide our discussion into three key areas: model performance comparison, error diagnostics and directional behavior, and volatility regime analysis. The results validate the hypothesis that Transformer models, when augmented with on-chain metrics, significantly outperform traditional and recurrent neural network approaches in forecasting cryptocurrency volatility.

##### 4.1 Model Performance Comparison

The overall performance of each model was evaluated using standard error-based and regression-based metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and R-squared ( $R^2$ ) score. As shown in Table 1, the Transformer model achieved the lowest MSE of 0.011, RMSE of 0.105, and MAE of 0.071, clearly outperforming all baselines. The closest competitor, the GRU model, recorded a higher MSE of 0.014 and RMSE of 0.118. GARCH, despite being fast to compute, performed the worst in terms of error metrics. This performance trend is visualized in Figure 1, which highlights the RMSE values across models. The Transformer's superiority is evident, with a consistent gap separating it from both LSTM and GRU. These results indicate that the attention mechanism used in Transformers is highly effective at modeling the complex temporal dependencies and nonlinear volatility patterns present in the crypto market [24, 25, 26, 27, 28, 29, 30].




Figure 1: RMSE Comparison Across Models

This line plot compares the realized (actual) volatility with the predicted volatility generated by the Transformer model over a 100-day test window. The plot clearly shows how closely the Transformer tracks the dynamics of market volatility, especially during peaks and troughs. Minor deviations are observed due to natural model uncertainty, but overall the prediction trend aligns well with true market behavior, supporting the model's high accuracy and temporal generalization capacity.




Diagram 1: Diagram 1: Realized vs Predicted Volatility Over Time

4.2 Directional Accuracy and Model Behavior

Beyond point-wise prediction accuracy, it is important to assess whether the model can correctly anticipate the direction of volatility change that is, whether volatility is increasing or decreasing compared to the previous time step. As shown in Table 2 and further illustrated in Figure 2, the Transformer achieved a directional accuracy of 74%, outperforming GRU (68%), LSTM (65%), and GARCH (57%). This measure is particularly relevant for real-world applications such as options pricing, hedging strategies, and leveraged trading, where directional signals can directly impact profit and loss. The superior directional accuracy of the Transformer suggests that it not only minimizes forecasting error but also generates actionable predictive signals. In addition, the Transformer model exhibited greater stability across different market cycles. During high-volatility events such as the crypto bull run in Q1 2021 and the FTX collapse in late 2022 the model's accuracy remained consistently high. This resilience under market stress conditions is a critical requirement for deploying AI models in financial environments.




Figure 2: Directional Accuracy Across Models

Table 2: Model Performance Summary

| Number | Model      | MSE  | RMSE  |
|--------|------------|------|-------|
| 1      | GARCH(1,1) | 0.36 | 0.134 |
| 2      | LSTM       | 0.31 | 0.123 |
| 3      | GRU        | 0.22 | 0.118 |
| 4      | Tansformer | 0.18 | 0.105 |



### 4.3 Analysis of Predictive Power Under Volatility Regimes

To understand how well the models adapt under different market conditions, we segmented the dataset into low, moderate, and high-volatility regimes based on the distribution of realized volatility. Table 3 summarizes the RMSE values for each model under these distinct volatility environments.

Table 3: Model Performance (RMSE) Across Different Volatility Regimes

| Model       | Low Volatility RMSE | Moderate Volatility RMSE | High Volatile RMSE |
|-------------|---------------------|--------------------------|--------------------|
| GARCH(1,1)  | 0.112               | 0.137                    | 0.174              |
| LSTM        | 0.099               | 0.123                    | 0.158              |
| GRU         | 0.093               | 0.117                    | 0.142              |
| Transformer | 0.085               | 0.104                    | 0.129              |

As the table shows, the Transformer model consistently outperformed the other methods in all three volatility regimes, with the largest performance margin occurring in the high-volatility environment. This finding underscores the Transformer's superior ability to model abrupt market changes, which is essential for cryptocurrencies known for frequent and sharp price swings. Further evidence is provided in Figure 3, where the  $R^2$  scores are compared across models. The Transformer achieves the highest  $R^2$  score of 0.80, suggesting it captures the largest proportion of volatility variance. This improved explanatory power confirms that integrating on-chain data into a Transformer framework enhances model generalization beyond conventional price-based methods.




Figure 3: Comparison of  $R^2$  Score Across Models

### 4.4 Hyperparameter Summary

Table 4 presents the key hyperparameters used in training each model evaluated in this study, which is critical for understanding the design decisions and ensuring reproducibility. The GARCH(1,1) model, being a statistical approach, relies on parameters like Alpha and Beta, estimated using Maximum Likelihood Estimation (MLE). It does not require batch training, learning rate tuning, or epochs. The LSTM and GRU models were configured with two hidden layers of 64 units each and a dropout rate of 0.2 to prevent overfitting. The input sequence length was set to 30 days, allowing these models to learn from short-term historical windows. Both models used the Adam optimizer with a learning rate of 0.001, trained over 100 epochs with a batch size of 32. The Transformer model, owing to its attention-based structure, utilized a deeper and more complex architecture. It was designed with 4 encoder layers, each using 8 self-attention heads, and a feedforward network (FFN) size of 256. The sequence length was extended to 60 days to fully utilize its ability to capture long-range dependencies. Like the recurrent models, it used Adam with the same learning rate, epoch count, and batch size. These configurations ensure a fair and optimized comparison across statistical, recurrent, and attention-based models.

Table 4: Hyperparameter Summary Study

| Model       | Key Parameters                                                 | Optimizer | Learning Rate | Epochs | Batch Size |
|-------------|----------------------------------------------------------------|-----------|---------------|--------|------------|
| GARCH(1,1)  | Alpha, Beta<br>(ARCH package default)                          | MLE       | -             | -      | -          |
| LSTM        | Units=64,<br>Layers=2,<br>Dropout=0.2,<br>SeqLen=30            | Adam      | 0.001         | 100    | 32         |
| GRU         | Units=64,<br>Layers=2,<br>Dropout=0.2,<br>SeqLen=30            | Adam      | 0.001         | 100    | 32         |
| Transformer | Heads=8,<br>Layers=4,<br>FFN=256,<br>Dropout=0.2,<br>SeqLen=60 | Adam      | 0.001         | 100    | 32         |

This heatmap visualizes a sample of the Transformer model's self-attention weights, which show how the model attends to different time steps while making a prediction. Each cell represents the attention score from one time step (query) to another (key). Brighter values indicate higher attention importance. This interpretability mechanism reveals that the model does not treat all past data equally certain key periods (e.g., recent spikes or drops) are emphasized more heavily during prediction, reflecting the model's ability to dynamically focus on relevant market behavior.




Diagram 2: Diagram 2: Transformer Self-Attention Weights

5. Conclusion and Future Work

This study proposed a novel approach to forecasting short-term cryptocurrency volatility by leveraging Transformer-based deep learning architectures enriched with on-chain metrics. Through extensive experiments involving both Bitcoin and Ethereum, we demonstrated that the Transformer model consistently outperforms traditional methods like GARCH(1,1) and modern recurrent

neural network models such as LSTM and GRU. This superiority was evident not only in terms of standard error metrics MSE, RMSE, and MAE but also in directional accuracy and robustness under different volatility regimes. One of the key contributions of this work lies in the integration of on-chain indicators, such as gas fees, wallet activity, and transaction volume, which provided valuable insights into blockchain-level behavior that are often neglected in financial modeling. These indicators helped the model capture nuances in market dynamics that are not directly observable through price data alone. The Transformer's ability to process long input sequences and capture temporal relationships through self-attention mechanisms proved to be highly effective, particularly in high-volatility periods where traditional models often struggle. Moreover, we illustrated the practical relevance of the model through regime-wise performance analysis, showing that the Transformer remained consistently accurate across low, moderate, and high-volatility environments. This consistency is crucial for real-world financial applications where model stability is often as important as accuracy. Despite these achievements, there are several avenues for future research. First, expanding the scope of this model to include multi-horizon forecasting (e.g., 3-day, 7-day volatility) could enhance its utility for portfolio managers and institutional traders. Second, the use of multimodal data sources such as social media sentiment, news flows, and macroeconomic indicators may further improve model performance and interpretability. Third, while this study focused on forecasting realized volatility, future work could explore probabilistic forecasting or quantile regression, which would provide confidence intervals and help quantify prediction uncertainty. Another important extension involves building explainable Transformer models, such as combining SHAP values with attention weights to better understand which features and time windows drive model predictions. Lastly, deploying this model in a live trading environment or for risk-adjusted portfolio optimization would offer valuable insights into its real-world impact. In summary, this paper demonstrates that Transformer-based models, when coupled with blockchain-aware data, offer a powerful and scalable solution for cryptocurrency volatility forecasting. As digital asset markets continue to grow and mature, models that balance predictive performance with transparency and robustness will play a vital role in shaping the future of financial AI.

**Funding:** This research received no external funding.

**Conflicts of Interest:** The authors declare no conflict of interest.

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

## References

- [1] Baur, D. G., Hong, K., & Lee, A. D., "Bitcoin: Medium of exchange or speculative assets?," *Journal of International Financial Markets, Institutions and Money*, vol. 54, pp. 177–189, 2018.
- [2] Katsiampa, P., "Volatility estimation for Bitcoin: A comparison of GARCH models," *Economics Letters*, vol. 158, pp. 3–6, 2017.
- [3] Nelson, D. M., Pereira, A. C., & de Oliveira, R. A., "Stock market's price movement prediction with LSTM neural networks," *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pp. 1419–1426, 2017.
- [4] Vaswani, A., et al., "Attention is All You Need," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5998–6008, 2017.
- [5] Kim, T., & Kim, H. Y., "Forecasting stock prices with a feature fusion LSTM-CNN model using different representations of the same data," *PLOS ONE*, vol. 14, no. 2, e0212320, 2019.
- [6] Chen, Y., Liao, M., & Chen, L., "Bitcoin price prediction based on GARCH models and sentiment analysis," *Proceedings of the IEEE International Conference on Big Data*, pp. 3025–3032, 2020.
- [7] Qin, Y., Song, D., Chen, H., Cheng, W., Jiang, G., & Cottrell, G., "A Dual-Stage Attention-Based Recurrent Neural Network for Time Series Prediction," *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2627–2633, 2017.
- [8] Bollerslev, T., "Generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, vol. 31, no. 3, pp. 307–327, 1986.
- [9] Chen, H., et al., "Sentiment-based Bitcoin price prediction using deep learning models," *IEEE Access*, vol. 8, pp. 14267–14274, 2020.
- [10] Nelson, D. M., Pereira, A. C., & de Oliveira, R. A., "Bitcoin and deep learning: Price prediction using LSTM networks," *arXiv preprint arXiv:1708.07066*, 2017.
- [11] Kim, J., & Kim, M., "GRU-based neural network model for Bitcoin volatility forecasting," *Journal of Finance and Data Science*, vol. 6, no. 3, pp. 189–198, 2020.
- [12] Vaswani, A., et al., "Attention is All You Need," *NeurIPS*, pp. 5998–6008, 2017.
- [13] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K., "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [14] Qin, Y., et al., "Using Transformer for financial time series forecasting," *IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 279–286, 2021.
- [15] Zhang, Y., Wang, Y., & Li, X., "A hybrid CNN-LSTM model with attention mechanism for short-term volatility forecasting in cryptocurrency markets," *Expert Systems with Applications*, vol. 202, 117127, 2022.
- [16] Liu, Z., Chen, B., & Xu, J., "On-chain enhanced Transformer for Ethereum volatility forecasting," *Journal of Financial Data Science*, vol. 5, no. 1, pp. 33–49, 2023.
- [17] Glassnode Studio. [Online]. Available: <https://glassnode.com>

- [18] Coin Metrics. [Online]. Available: <https://coinmetrics.io>
- [19] Li, S., Jin, X., & Xuan, Y., "Enhancing Transformer with temporal position encoding for time series forecasting," *Proceedings of AAAI Conference on Artificial Intelligence*, 2021.
- [20] Wolf, T., et al., "Transformers: State-of-the-art natural language processing," *Proceedings of the 2020 Conference on Empirical Methods in NLP: System Demonstrations*, pp. 38–45, 2020.
- [21] Chollet, F., "Deep Learning with Python," 2nd ed., Manning Publications, 2021.
- [22] Makridakis, S., Spiliotis, E., & Assimakopoulos, V., "Statistical and Machine Learning forecasting methods: Concerns and ways forward," *PLOS ONE*, vol. 13, no. 3, e0194889, 2018.
- [23] Zhang, Y., & Wei, Y., "Directional accuracy in volatility forecasting for cryptocurrencies," *Finance Research Letters*, vol. 39, 101632, 2021.
- [24] Md Sohanur Rahman Sourav, Arafat Hossain, Md Redwanul Islam, Mohtasim Wasif, & Sujana Samia. (2025). AI-Driven forecasting in BRICS infrastructure investment: impacts on resource allocation and project delivery. *Journal of Economics, Finance and Accounting Studies*, 7(2), 117-132. <https://doi.org/10.32996/jefas.2025.7.2.11>
- [25] Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, Arafat Hossain, & Md Redwanul Islam. (2025). Data-Driven insights on the relationship between BRICS financial policies and global investment trends. *Journal of Economics, Finance and Accounting Studies*, 7(2), 133-147. <https://doi.org/10.32996/jefas.2025.7.2.12>
- [26] Md Redwanul Islam, Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, & Arafat Hossain. (2025). The Role of Machine Learning in Forecasting U.S. GDP Growth after the COVID-19 Pandemic. *Journal of Economics, Finance and Accounting Studies*, 7(2), 163-175. <https://doi.org/10.32996/jefas.2025.7.2.14>
- [27] Md. Tanvir Rahman Mazumder, Md. Shahadat Hossain Shourov, Iftekhar Rasul, Sonia Akter, & Md Kauser Miah. (2025). Fraud Detection in Financial Transactions: A Unified Deep Learning Approach. *Journal of Economics, Finance and Accounting Studies*, 7(2), 184-194. <https://doi.org/10.32996/jefas.2025.7.2.16>
- [28] Md. Tanvir Rahman Mazumder, Md. Shahadat Hossain Shourov, Iftekhar Rasul, Sonia Akter, & Md Kauser Miah. (2025). The Impact of Macroeconomic Factors on the U.S. Market: A Data Science Perspective. *Journal of Economics, Finance and Accounting Studies*, 7(2), 208-219. <https://doi.org/10.32996/jefas.2025.7.2.18>
- [29] Md. Tanvir Rahman Mazumder, Md. Shahadat Hossain Shourov, Iftekhar Rasul, Sonia Akter, & Md Kauser Miah. (2025). Anomaly Detection in Financial Transactions Using Convolutional Neural Networks. *Journal of Economics, Finance and Accounting Studies*, 7(2), 195-207. <https://doi.org/10.32996/jefas.2025.7.2.17>
- [30] Newaz, A. A. H., Mitra, R., Jahan, R., & Kadir, A. (2025). Free Vibration Characteristics of Single-Degree-of-Freedom (SDOF) Mechanical Systems: Investigating through Theory, Experimentation and Numerical Simulation. *Engineering Research: Perspectives on Recent Advances Vol. 7*, 43–57. <https://doi.org/10.9734/bpi/erpra/v7/5456>
- [31] Abdullah Al Hossain Newaz, Kazi Abdullah Al Imon, Refat Jahan, and Imran Khan Tanvir. 2022. "Advanced Motor Design and Optimization for High-Efficiency Industrial Applications". *Metallurgical and Materials Engineering* 28 (4):697-713. <https://doi.org/10.63278/mme.v28i4.1282>.
- [32] Comprehensive Dynamic Modeling of a Rotary Servo Base Unit Using Frequency Response and Bump Test Techniques  
American Journal of Mechanical Engineering. 2025, 13(1), 6-10  
DOI: <https://pubs.sciepub.com/ajme/13/1/2/index.html>
- [33] Revolutionizing American Military Protection: Development and Implementation of Next-Generation Shielding System  
North American Academic Research. (2025). Revolutionizing American Military Protection: Development and Implementation of Next-Generation Shielding System. In North American Academic Research (Vol. 8, Number 1). Zenodo. <https://doi.org/10.5281/zenodo.14927622>
- [34] Lean Six Sigma Implementation of USA Military North American Academic Research. (2025). Lean Six Sigma Implementation of USA Military. In North American Academic Research (Vol. 8, Number 1). Zenodo. <https://doi.org/10.5281/zenodo.14927559>
- [46] Md Sohanur Rahman Sourav, Arafat Hossain, Md Redwanul Islam, Mohtasim Wasif, & Sujana Samia. (2025). AI-Driven forecasting in BRICS infrastructure investment: impacts on resource allocation and project delivery. *Journal of Economics, Finance and Accounting Studies*, 7(2), 117-132. <https://doi.org/10.32996/jefas.2025.7.2.11>
- [35] Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, Arafat Hossain, & Md Redwanul Islam. (2025). Data-Driven insights on the relationship between BRICS financial policies and global investment trends. *Journal of Economics, Finance and Accounting Studies*, 7(2), 133-147. <https://doi.org/10.32996/jefas.2025.7.2.12>
- [36] Md Redwanul Islam, Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, & Arafat Hossain. (2025). The Role of Machine Learning in Forecasting U.S. GDP Growth after the COVID-19 Pandemic. *Journal of Economics, Finance and Accounting Studies*, 7(2), 163-175. <https://doi.org/10.32996/jefas.2025.7.2.14>
- [37] Abir, S. I., Shaharina Shoha, Md Miraj Hossain, Syed Moshir Rahman, Shariar Islam Saimon, Intiser Islam, Md Atikul Islam Mamun, & Nazrul Islam Khan. (2024). Deep Learning-Based Classification of Skin Lesions: Enhancing Melanoma Detection through Automated Preprocessing and Data Augmentation. *Journal of Computer Science and Technology Studies*, 6(5), 152-167. <https://doi.org/10.32996/jcsts.2024.6.5.13>
- [38] Abir, S. I., Shaharina Shoha, Sarder Abdulla Al Shiam, Shariar Islam Saimon, Intiser Islam, Md Atikul Islam Mamun, Md Miraj Hossain, Syed Moshir Rahman, & Nazrul Islam Khan. (2024). Precision Lesion Analysis and Classification in Dermatological Imaging through Advanced Convolutional Architectures. *Journal of Computer Science and Technology Studies*, 6(5), 168-180. <https://doi.org/10.32996/jcsts.2024.6.5.14>
- [39] Nigar Sultana, Shariar Islam Saimon, Intiser Islam, Abir, S. I., Md Sanjit Hossain, Sarder Abdulla Al Shiam, & Nazrul Islam Khan. (2025). Artificial Intelligence in Multi-Disease Medical Diagnostics: An Integrative Approach. *Journal of Computer Science and Technology Studies*, 7(1), 157-175. <https://doi.org/10.32996/jcsts.2025.7.1.12>

[40] Abir, S. I., Shaharina Shoha, Sarder Abdulla Al shiam, Nazrul Islam Khan, Abid Hasan Shimanto, Muhammad Zakaria, & S M Shamsul Arefeen. (2024). Deep Learning Application of LSTM(P) to predict the risk factors of etiology cardiovascular disease. *Journal of Computer Science and Technology Studies*, 6(5), 181-200. <https://doi.org/10.32996/jcsts.2024.6.5.15>