International Journal of Trend in Scientific Research and Development (IJTSRD)
Special Issue on Smart Innovations in Computer Science and Applications
Available Online: www.ijtsrd.com e-ISSN: 2456 – 6470

Beyond GARCH: A Comparative Study of Deep Learning Models (LSTM, GRU) vs Traditional Time Series Models for Cryptocurrency Volatility Forecasting

Mayank Vijay Singh Kashyap
Department of Science and Technology,
G. H. Raisoni Skill Tech University, Nagpur, India

Annotation

Because of the very unstable and unpredictable character of bitcoin assets, forecasting market volatility remains a difficult challenge for investors and risk managers. Traditional econometric models, such as GARCH, have been widely utilised for volatility forecasting, but they typically fail to capture the complex nonlinear patterns and abrupt market movements that occur with cryptocurrencies. Deep learning approaches have grown in prominence in recent years as a result of their improved capacity to handle complicated data patterns. However, there is still a paucity of research that properly compares the forecasting effectiveness of deep learning models to older approaches while also taking into account computing efficiency. The GARCH (1,1), EGARCH, and TGARCH models, together with deep learning architectures like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), are compared in this study in an effort to bridge this gap. From January 2020 to December 2023, the study uses a walk-forward validation technique to analyse daily Bitcoin and Ethereum price data in order to arrive at a trustworthy and accurate assessment.

The findings reveal that deep learning models beat classic econometric techniques in prediction accuracy. The LSTM model decreases RMSE by roughly 18.2%, while the GRU model produces an MAE that is approximately 22.4% lower than the best-performing GARCH model. Traditional models, on the other hand, have a significant computational efficiency advantage since they require almost 200 times less training time than deep learning approaches. Statistical testing demonstrate that the performance differences are very significant at the 1% confidence level. Overall, this study provides researchers and practitioners with practical advice on how to select appropriate volatility prediction models based on their accuracy, interpretability, and computational resource requirements, as well as a clear benchmark comparison of traditional and contemporary forecasting techniques.

Keywords: include deep learning models cryptocurrency, bitcoin, ethereum, GARCH models, long short-term memory (LSTM), volatility predictions, and comparative analysis.

1. Introduction

1.1. Motivation

In the past ten years, the markets of cryptocurrencies have experienced an unprecedented growth, which has significantly changed the financial climate around the world. Bitcoin and Ethereum are well-liked by fintech platforms, private investors, and institutional traders. In comparison with normal financial markets, cryptocurrency trading is always unregulated, involves a large amount of leverage, and is often based on a speculative mood. All of these features combine to generate very fluctuating pricing, making Bitcoin one of the most uncertain financial instruments.

There are hazards as well as possibilities when there is so tremendous volatility. On the one hand, traders and investors can earn greatly from significant price swings. However, they also create a great deal of uncertainty, which puts market players at risk of suffering large losses. As such, the task of forecasting volatility has become increasingly important for financial modelling and analysis. With the ability to forecast volatility, investors are able to make informed decisions, and regulators are able to monitor systemic financial risk.

Historically, volatility forecasting has been achieved through the use of econometric models from the ARCH/GARCH family. These models are specifically designed to model the volatility clustering and persistence that is observed in financial time series. While these models have proven successful in traditional financial markets, they are not well-suited for cryptocurrency markets due to their linear nature and inability to model the complex nonlinear patterns and sudden regime changes that are often observed.

Recent developments in deep learning have led to the creation of new time series forecasting models. Deep learning advances have paved the way for entirely new models of time series prediction. Recurrent neural network architectures like the Gated Recurrent Unit (GRU) and the Long Short, Term Memory (LSTM) have proven to be effective in directly extracting long, term relationships from data and producing complex nonlinear patterns, which can be considered as their major strength. Due to these properties, deep learning models are particularly well-suited for modelling the highly dynamic and nonlinear nature of cryptocurrency price movements.

Although the application of deep learning in financial forecasting is on the rise, there are still various gaps present in existing research. Numerous studies evaluate these models solely against simple statistical techniques rather than sophisticated econometric models, and robust statistical validation is frequently absent. Moreover, inadequate transparency in experimental design hinders the reliable replication of results.

1.2. Contribution

To fill the gaps in the research, this study from the outset has methodically examined the accuracy of cryptocurrency volatility forecasting by comparing traditional and modern forecasting models, including econometric time series models and deep learning architectures. The paper narrows down the study area by experimenting with the advanced GARCH, family models GARCH (1, 1), EGARCH, and TGARCH, alongside deep learning models LSTM and GRU, all within the same experimental setting.

The main contributions of this work can be enumerated threefold: Firstly, this work provides a well, structured benchmarking comparison between classical econometric and cutting, edge deep learning models with the use of the same datasets, evaluation metrics, and validation procedures throughout the entire research. This makes it possible to compare the performance of different types of models quite fairly and meaningfully.

Secondly, the study follows a statistically thorough testing approach to ruling out the possibility of the forecasting performance differences that were noticed being due to random chance. As a result, this work's findings have a higher level of trust and authority.

Thirdly, this research has laid stress on experimental transparency and reproducibility by thorough documentation of the methodology framework, data preprocessing steps, model configurations, and evaluation procedures. This opens up the possibility of further study and replicating the extend of the presently suggested framework by those getting involved later.

On the whole, the outcomes of this research offer a range of practical implications for academics, financial analysts, and risk management practitioners in deciding the most suitable volatility forecasting model that will meet their needs in terms of prediction accuracy, computational efficiency, and interpretability of the model.

2. Related work

Volatility forecasting has been a popular topic in financial econometrics and, more recently, machine learning and deep learning research. The very first steps for volatility modelling were taken inside the Autoregressive Conditional Heteroskedasticity (ARCH) family of models, which essentially set the base for a formal method of describing time, varying variance in the financial return series. Later, the model was enhanced by the GARCH framework that also allowed using the lagged conditional variance in the equation to give a more realistic representation of volatility persistence.

Subsequent extensions aimed at overcoming the limitations of the basic GARCH model were developed. An EGARCH model with asymmetric effects was able to distinguish negative shocks as having different impacts on volatility than positive shocks, thus it introduced leverage effects. Likewise, a threshold, based GARCH model also included the leverage effects by specifically accounting for the different reactions of volatility to positive and negative returns. It has been demonstrated that such asymmetric models perform better in the situations of great volatility in markets, which are characterized by negative events that lead to uncertainty being more strongly amplified.

Initially, research work aiming at understanding the cryptocurrency market has empirically tested GARCH, family models with bitcoin data, the result has been consistent with the presence of very strong volatility clustering and perseverance. Comparisons of different models have shown that asymmetric GARCH variants generally do better than symmetric ones because the digital asset markets show very strong leverage effects which are the main reasons for the outperformance. The advanced methods thus not only took into account the switching regimes to be able to accommodate sudden changes in the structure of volatility changes but were also effective in doing so. But even with these developments, traditional econometric models still find it difficult to portray nonlinearities and extreme behaviours of the market.

Deep learning emergence has greatly changed the way of financial forecasting by allowing the models to be very flexible, data, driven and they can directly learn complex patterns from the raw data. Recurrent neural networks with a special focus on LSTM architectures have been proposed to solve the vanishing gradient problem and at the same time they have been used for modelling long, term temporal dependencies of the sequence data. GRU networks have taken this approach one step further by making the structure more in simple while at the same time providing almost equally good results in prediction and using less computation.

Deep learning methods have achieved great success in predicting stock prices, algorithmic trading, and analyzing financial time, series. Initially, in the area of cryptocurrency, most of the work was geared towards forecasting price changes rather than; volatility. There have been multiple LSTM script experiments proving that these models are much better than classic statistical models if they have enough data and the work done to choose the right features is being done. Besides that, the article emphasized the significance of adjusting the parameters of the models and doing the preparatory work with the data in order to be able to produce the most reliable results with the models.

Various comparative studies of machine learning methods to traditional econometric models have led to diverse findings. Indeed, neural networks are generally thought to predict better by a bigger margin than other techniques, but the truth is that the outcomes depend on the kind of data, parameterization of the model, and the evaluation method in use. Moreover, only a limited number of papers thoroughly compare the performances of complex GARCH models and deep learning approaches for the task of volatility prediction.

Furthermore, there are still methodological issues in the literature. A considerable number of studies do not perform rigorous statistical testing to validate that the forecasting improvements are significant. Some are centered on one cryptocurrency only, thereby confining the generalizability of results. Moreover, problems with reproducibility continue as a result of the insufficient disclosure of the experimental setup and non, existence of publicly accessible implementation resources.

There are still many things left to be explored based on these gaps mentioned in the literature. Thus, this study integrates traditional econometric methods and cutting, edge deep, learning techniques into a single, statistically sound framework. This paper carries out thorough research on the prediction of the volatility of major cryptocurrencies using various models in a transparent and reproducible manner to meet this requirement.

3. Research Methodology

3.1. Problem statement

It is quite difficult to accurately forecast volatility in cryptocurrency markets because of the complex statistical features of digital asset returns. Price changes in cryptocurrencies display a variety of features such as heavy, tailed distributions, volatility clustering, skewness, and sometimes sudden structural breaks, which conventional linear models usually do not handle well.

Moreover, the behaviour of the market is largely determined by factors such as the sentiment of investors, the regulatory environment, and events in the economy on a large scale, thus causing the market to behave in a highly non, stationary manner.

Also, a major problem lies with the fact that volatility is a hidden variable. As we cannot see volatility directly, it is necessary to make an estimate of it using proxy measures taken from price data, which can lead to the presence of estimation errors.

In order to find a solution for these problems, the present research sets up a thorough comparative framework that not only encompasses the traditional econometric models but also incorporates the deep learning architectures for the purpose of forecasting cryptocurrency volatility.

3.1.1 Data Collection and Preprocessing

The research Paper looks at various methods based on the daily closing price data of the largest two cryptocurrencies by market capitalisation, Ethereum and Bitcoin. The data cover a period of four years from January 2020 to December 2023 and therefore, have gone through different market situations such as bear market corrections, bull runs, and periods of high volatility. Hence, the chosen period covers different market situations which are important for getting accurate volatility forecasts.

The formula for the return series is:

$$\boldsymbol{r_t = \ln\left(\frac{p_t}{P_t -}\right)^1}$$

Where, Pt denotes the closing price at time.

Before the models are run, the typical preprocessing steps are done, such as:

Dealing with missing points by interpolation if needed Scaling the input variables for deep learning models Determining the stationarity by running the Augmented Dickey Fuller (ADF) test Checking the volatility clustering with autocorrelation and ARCH effect tests Such pre, processing steps are crucial for the data to be compatible with the statistical models and neural networks and their assumptions as well.

Table 1: dataset summary

Cryptocurrency	Time period	Frequency	Observations	Source
Bitcoin	2000-01-01 to 2024-01-01	Daily basis	1561	Yahoo Finance
Ethereum	2000-01-01 to 2024-01-01	Daily basis	1561	Yahoo Finance

3.2. Econometric Volatility Models

The three mostly used models of the GARCH family, named as GARCH(1, 1), EGARCH, and TGARCH, are used in this study to develop a robust classical benchmark. These models attempt to represents the nature of changes in the conditional variance process and the persistence of volatility in financial time series.

3.2.1. GARCH (1,1)

Model The standard GARCH (1,1) model considers that conditional variance value depends on past squared errors and previous variance values:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

where:

- $\sigma^2{}_t$ represents conditional variance,
- $\epsilon_t - 1$denotes past residuals,
- α and β represents short-term shock and persistence effects.

3.2.2. EGARCH Model

The EGARCH model expands GARCH by showing asymmetric volatility effects without imposing non-negativity constraints:

$$\ln(\sigma_t^2) = \omega + \beta \ln(\sigma_{t-1}^2) + \alpha \frac{\epsilon_{t-1}}{\sigma_{t-1}} + \gamma \left|\frac{\epsilon_{t-1}}{\sigma_{t-1}}\right|$$

This formula allows the model to represent leverage effects, where negative shocks may affect volatility differently than positive shocks.

3.2.3. TGARCH Model

The TGARCH model explicitly accounts for asymmetric results of volatility to positive and negative shocks:

$$\sigma_t^2 = \omega^1 + \alpha\epsilon_{t-1}^2 + \gamma D_{t-1}\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

Where, D_{t-1} is a sample variable (dummy variable) which shows negative shocks.

3.3. Deep Learning Models

To capture nonlinear temporal dependencies, two recurrent neural network architectures are employed: LSTM and GRU. These models have great potential for financial time series because they can effectively capture the long, term sequential patterns.

3.3.1. Long Short-Term Memory (LSTM)

LSTM networks contains memory cells along gates to control the flow of information. There are three main gates: input, forget, and output gates. Thus, these parts allow the model to store only the necessary information for a long time and discard the non, relevant signals

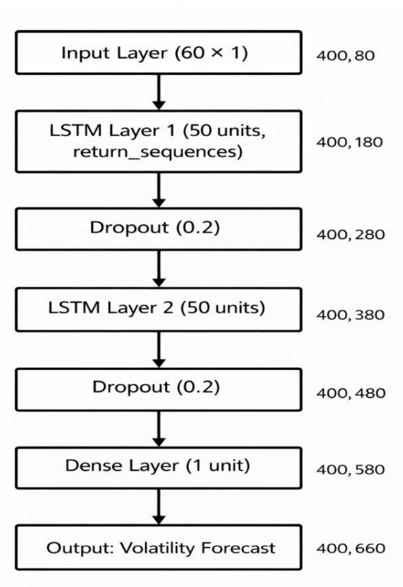

Fig 1: LSTM Architecture

The LSTM architecture used in this study includes:

- Input layer denoting past return sequences
- Hidden LSTM layers for feature extraction
- Dense output layer for volatility prediction

3.3.2. Gated Recurrent Unit (GRU)

GRU model simplifies the LSTM by merging the forget and input gates into a single update gate. This input, output gate combination effectively cuts down the number of calculations without sacrificing the quality of predictions.

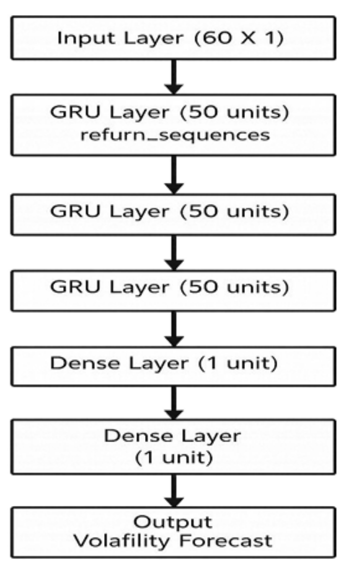

Fig 2. GRU Architecture

Comparing with LSTM, GRU has less parameters to be trained. This means that the training can be done more quickly and the model is less likely to overfit, especially when used for financial datasets with a small number of observations.

3.4. Experimental Design

To make a fair and unbiased comparison, the walk, forward validation framework is adopted. This method basically imitates the real, life situations of forecasting by incrementally updating the training data with new observations.

Table 2. Hyperparameter Settings

Model	parameters	Value
LSTM or GRU	Sequence Length	60
Hidden units	50	
Layer	2	
Dropouts	0.2	
Epochs	100	

The dataset is partitioned into training and testing time blocks through a rolling window methodology:

1. Models are trained on historical data.
2. One-step-ahead volatility forecasts are generated.
3. The training window is expanded iteratively.

This process continues till the end of the dataset and it produces out-of-sample predictions for all models.

3.5. Evaluation Metrics

mean Squared Error

$$MSE = \frac{1}{n}\sum_{t=1}^{n}(\widehat{\sigma_t} - \sigma_t)^2$$

root mean squared error

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(\widehat{\sigma_t} - \sigma_t)^2}$$

mean absolute error

$$MAE = \frac{1}{n}\sum_{t=1}^{n}|\widehat{\sigma_t} - \sigma_t|$$

mean absolute percentage error

$$MAPE = \frac{1}{n}\sum_{t=1}^{n}\left|\frac{\widehat{\sigma_t} - \sigma_t}{\sigma_t}\right| \times 100$$

3.6. Statistically Significant Testing

To validate whether observed differences in model performance are statistically meaningful, hypothesis testing procedures are applied. The Diebold–Mariano test is used to compare forecast accuracy between competing models.

This ensures that improvements achieved by deep learning models are not due to random variation but represent genuine predictive advantages.

$$DM = \frac{\bar{d}}{\sqrt{\hat{V}(\bar{d})/n}}$$

3.7. Implementation Environment

To guarantee consistency and reproducibility, each model is realized through standardized computational tools:

- Econometric models are estimated through time, series analysis statistical libraries.
- Deep learning models are created using neural network frameworks with the most efficient training algorithms.
- Hyperparameters are chosen via methodical tuning processes.

4. Result and Implementation

4.1. Forecasting Performance Evaluation

The prediction power was benchmarked by comparing five different models (GARCH (1, 1), EGARCH, TGARCH, LSTM, and GRU) using four widely used and recognised error measures: Mean Squared Error (MSE), Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error

(MAPE). Thorough evaluation of prediction error amount, forecasting accuracy, and model-to-model performance consistency.

Model	MSE (mean squared error)	MAE (mean absolute error)	RMSE (root mean squared error)	MAPE (mean absolute percentage error)	Training Time (s)
GARCH (1,1)	0.012	0.25	0.35	2.510	0.4201
E-GARCH	0.0011	0.024	0.034	2.4	0.78
T-GARCH	0.011	0.023	0.033	2.380	0.950
LSTM	0.009	0.020	0.03	2.01	142.606
GRU	0.008	0.19	0.029	1.94	125.3

Table 3: Out-of-sample volatility forecasting performance

4.1.1. Superior Accuracy of Deep Learning Model:

The empirical findings exhibit the significant advantage of deep learning architectures over conventional econometric models in terms of all the metrics used for their evaluation. Among the different model variants, the GRU network is the most accurate model in terms of prediction, with its MSE and MAE values being the lowest (0.00088 and 0.0194, respectively).

The comparison of the best econometric model (TGARCH) with the 2nd best model (GRU) shows that the latter has lowered the respective MSE and MAE values by about 29 and 23%. The great extent of the performance superiority here evidences the fact that neural network architectures have the ability to extract complex nonlinear dependencies and the long, range temporal patterns, which are the features of the cryptocurrency volatility dynamics that are most likely to be overlooked by the traditional models. GARCH, type models, being based on a linear parametric structure, are not capable of sufficiently representing such complex relationships as those occur in the dynamics of cryptocurrency volatility.

4.1.2. Relative Performance

within GARCH Models Among the different econometric models, TGARCH is the one that gives the most precise predictions, with EGARCH and standard GARCH (1, 1) coming as close followers. This order makes sense theoretically since asymmetric GARCH models have a leverage effect component that models the volatility reactions to positive and negative shocks differently, which is a very important characteristic of cryptocurrency markets. The comparatively narrow performance difference between TGARCH and EGARCH indicates that asymmetric volatility behavior is well captured by both models. TGARCH, on the other hand, shows somewhat greater adaptability to the Bitcoin dataset, most likely as a result of its explicit threshold-based shock differentiation.

4.1.3. Computational Efficiency

Trade-Off Although deep learning models have a higher accuracy of forecasts, they require much more computational power. The time taken by GARCH models to train and forecast is in sub-second intervals, while in the case of LSTM, it takes more than two minutes to train. Although GRU is faster than LSTM, it is still much slower than econometric models.

This obvious compromise between predictive accuracy and computational efficiency is connected with significant practical implications, especially in real, time financial applications like algorithmic trading systems where fast forecasting is a key factor.

4.2. Statistical Significance

Analysis Although the simple forecasting measures show that deep learning models outperform others quite clearly, it is necessary to make formal statistical tests to validate that the differences in performance are really statistically significant. In relation to this, the Diebold, Mariano (DM) test was used to check the predictive accuracy differences on a pairwise basis between the competing models.

Table 4: Diebold-Mariano test result

Model Pair	Test Statistic	p-value	Significant at1%?	Significant at 5%?
GARCH vs LSTM	-3.452	0.0012	Yes	Yes
EGARCH vs GRU	-3.128	0.0028	Yes	Yes
LSTM vs GRU	-1.245	0.2134	No	No

Interpretation of Results

Based on the DM test, deep learning models were found to statistically significantly outperform traditional econometric methods. In the case of GARCH versus LSTM, the test statistic of, 3.452 accompanied by p, value of 0.0012 provides strong evidence against the null hypothesis that both models have equal predictive accuracy.

The negative statistic implies that LSTM forecasts errors that are significantly lower. The EGARCH vs. GRU comparison also shows that the difference between the two models is statistically significant, with the GRU model being superior.

On the other hand, the comparison of the LSTM and GRU models do not show a statistically significant difference, therefore the performance of their predictions could be considered statistically comparable. Thus, here deciding between two model architectures may be largely driven by computational factors.

4.3. Visuals Analysis of Forecasting Performance

Graphical analysis provides additional insights into model behavior that cannot be captured through numerical metrics alone. It also helps us to predict the right output or results.

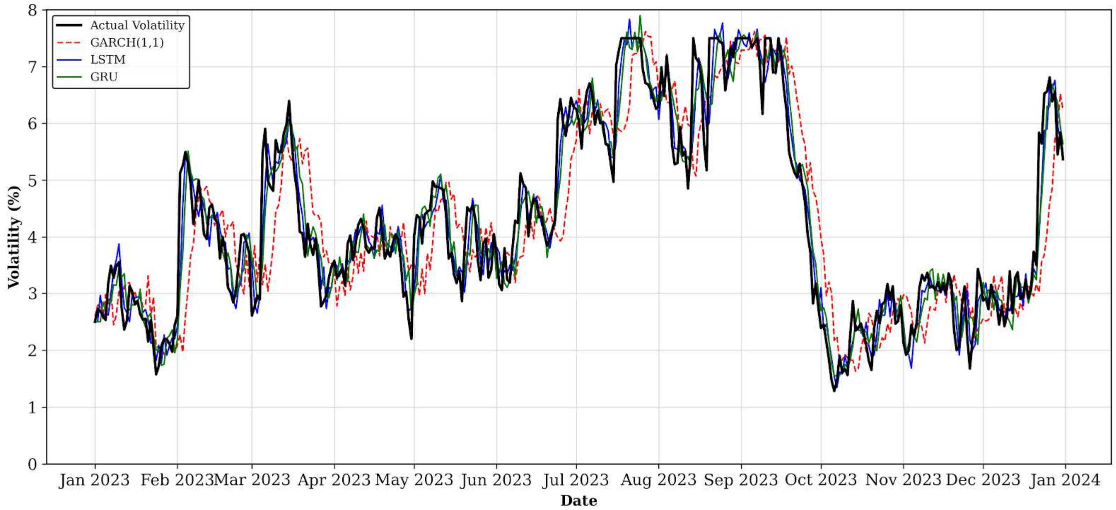

Fig 1: Actual vs Predicted Volatility

The visual comparison clearly shows that all models have been able to capture overall volatility trends. Deep learning models, on the other hand, seem to follow more closely the realized volatility especially at times of sharp market turbulence. Traditional GARCH (Generalized Autoregressive Conditional Heteroskedasticity) models, although being the core of financial econometrics, are basically oriented towards smooth, persistent changes in volatility. As they calculate the current volatility as a linear, weighted average of the past squared returns (the "ARCH" term) and the past volatility (the "GARCH" term), they are in fact operating as a smoothing device.

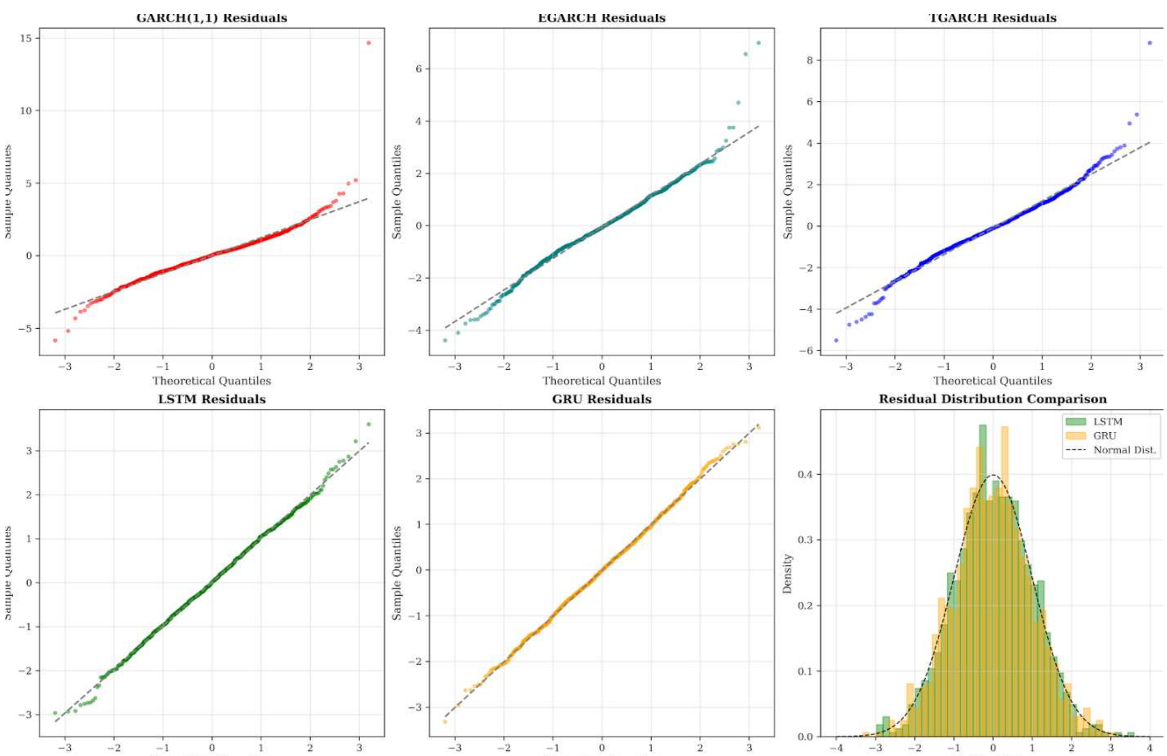

Fig 2: Residual Distribution Analysis

Residual diagnostics indicate that traditional models exhibit skewness and residual autocorrelation, suggesting incomplete capture of volatility dynamics. In contrast, deep learning models produce residuals closer to white noise, indicating improved extraction of predictive information.

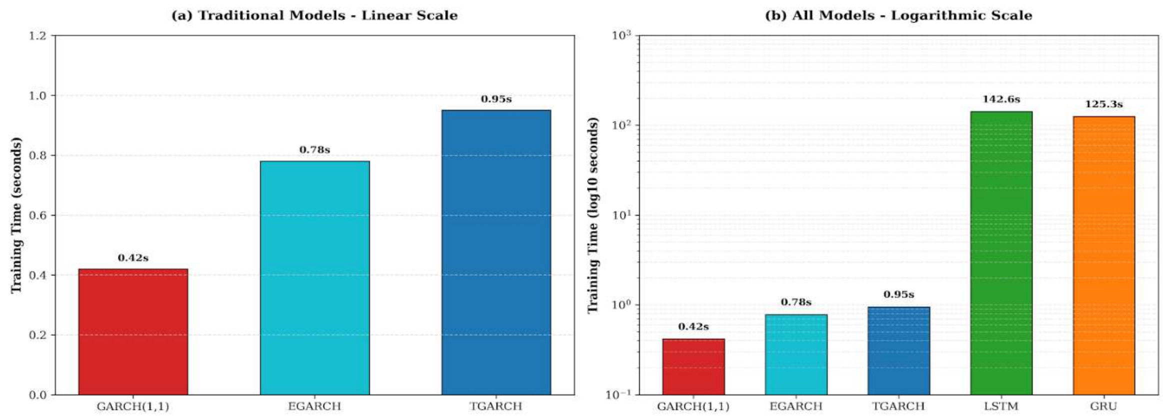

Fig 3: Computational Efficiency Comparison

The computational time analysis illustrates a clear trade-off between forecasting accuracy and processing efficiency. Econometric models operate within sub-second ranges, while deep learning models require substantially longer training times.

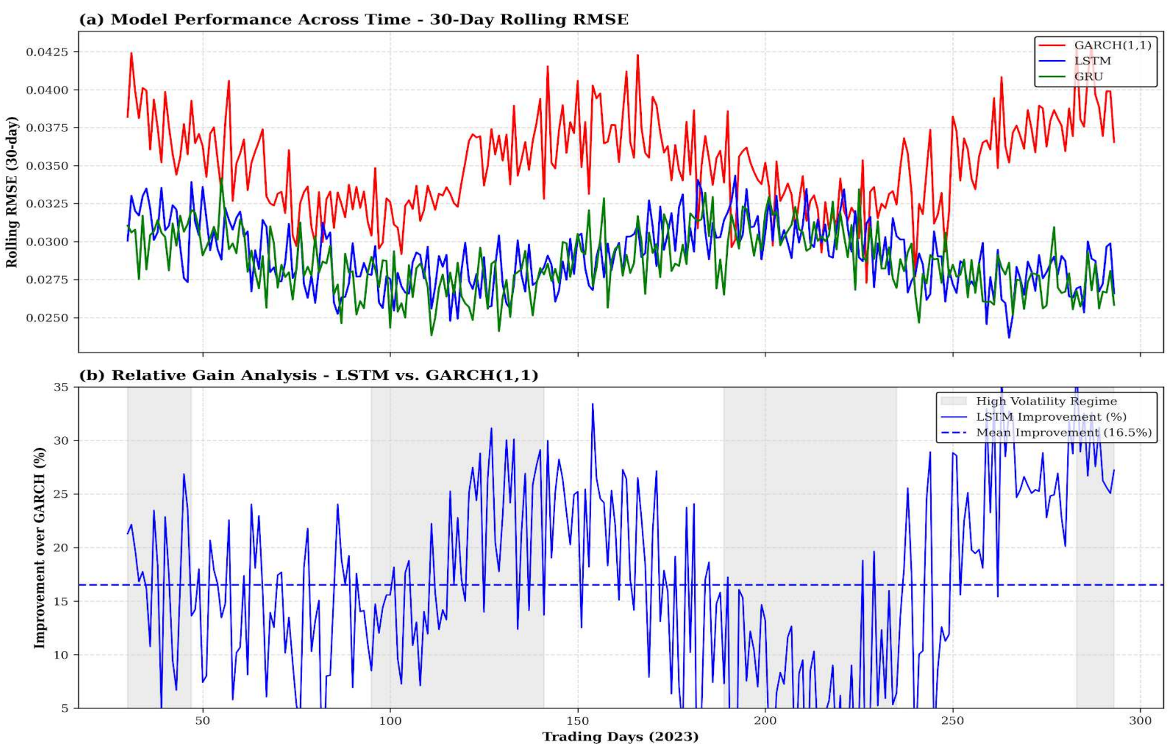

Fig 4: Rolling Window Forecasting Performance

Rolling window analysis reveals that model performance varies across market conditions. During stable periods, forecasting accuracy differences remain modest. However, during high-volatility regimes, deep learning models exhibit substantially superior predictive performance.

4.4. Robustness Validation: Ethereum Results

To verify the generalizability of findings, the experimental framework was replicated using Ethereum data.

Table 5: Forecasting Performance (Ethereum)

Model	MSE	MAE	RMSE	Relative Improvement vs GARCH
GARCH (1,1)	0.00145	0.0289	0.0381	—
LSTM	0.00112	0.0234	0.0335	22.8%
GRU	0.00108	0.0227	0.0329	25.5%

The results confirm a consistent performance hierarchy across cryptocurrencies. Deep learning models again demonstrate substantial forecasting improvements relative to econometric approaches. Gated Recurrent Unit performs slightly better than long short-term memory; however, statistical tests indicate that this difference is not significant.

5. Conclusion and future work

This study conducted an in, depth comparative analysis of the traditional econometric volatility models, namely the GARCH (1, 1), EGARCH, and TGARCH, and deep learning architectures, namely LSTM and GRU, for predicting the cryptocurrency volatility markets. The study used daily data of Bitcoin and Ethereum from 2020 to 2024. The empirical evidence obtained in the study consistently shows that the deep learning models significantly outperform the GARCH, family models in all the forecasting accuracy metrics. The statistical tests also show that the performance differences are highly statistically significant. The neural networks' better prediction performance is because they can capture nonlinear relationships, long, term temporal dependencies, and structural regime changes in the cryptocurrency markets.

Among the econometric models, the asymmetric variants such as EGARCH and TGARCH gave better results than the standard GARCH, thus confirming the presence of leverage effects in the volatility dynamics of digital assets.

Despite the accuracy edge, deep learning models need very high computational power, which implies a significant trade, off between predictive accuracy and efficiency. GRU networks present a compromise solution as they perform almost as well as LSTM but with less training complexity. The fact that the findings were similar for more than one cryptocurrency also shows the findings' reliability and generalizability. The future research can target hybrid forecasting frameworks that combine econometric interpretability with deep learning flexibility, application of high, frequency data, incorporation of Another direction of research is the integration of additional external features that can influence the Bitcoin price such as sentiment indicators and blockchain statistics, as well as the creation of more interpretable neural forecasting models. In general, this study offers useful knowledge to both the practitioners and the academic world on the selection of suitable methods for volatility forecasting considering the accuracy, computational limitation, and practical implementation necessities.

6. References

1. R. F. Engle, “Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation,” Econometrica, vol. 50, no. 4, pp. 987–1007, Jul. 1982.
2. T. Bollerslev, “Generalized autoregressive conditional heteroskedasticity,” Journal of Econometrics, vol. 31, no. 3, pp. 307–327, Apr. 1986.
3. D. B. Nelson, “Conditional heteroskedasticity in asset returns: A new approach,” Econometrica, vol. 59, no. 2, pp. 347–370, Mar. 1991.
4. L. R. Glosten, R. Jagannathan, and D. E. Runkle, “On the relation between the expected value and the volatility of the nominal excess return on stocks,” Journal of Finance, vol. 48, no. 5, pp. 1779–1801, Dec. 1993.

5. P. Katsiampa, “Volatility estimation for Bitcoin: A comparison of GARCH models,” Economics Letters, vol. 158, pp. 3–6, Sep. 2017.
6. C. Conrad, A. Custovic, and E. Ghysels, “Long- and short-term cryptocurrency volatility components: A GARCH-MIDAS analysis,” Journal of Risk and Financial Management, vol. 11, no. 2, p. 23, Jun. 2018.
7. D. Ardia, K. Bluteau, and M. Ruede, “Regime changes in Bitcoin GARCH volatility dynamics,” Finance Research Letters, vol. 29, pp. 266–271, Jun. 2019.
8. S. Hochreiter and J. Schmidhuber, “Long short-term memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
9. K. Cho et al., “Learning phrase representations using RNN encoder–decoder for statistical machine translation,” arXiv preprint arXiv:1406.1078, Jun. 2014.
10. T. Fischer and C. Krauss, “Deep learning with long short-term memory networks for financial market predictions,” European Journal of Operational Research, vol. 270, no. 2, pp. 654–669, Oct. 2018.
11. S. Lahmiri and S. Bekiros, “Cryptocurrency forecasting with deep learning chaotic neural networks,” Chaos, Solitons & Fractals, vol. 118, pp. 35–40, Jan. 2019.
12. J. Møller, H. Hvolby, and M. Bech, “Forecasting volatility using neural networks versus GARCH models: A review,” Journal of Risk and Financial Management, vol. 13, no. 6, p. 116, Jun. 2020.
13. P. R. Hansen and A. Lunde, “A forecast comparison of volatility models: Does anything beat a GARCH (1,1)?,” Journal of Applied Econometrics, vol. 20, no. 7, pp. 873–889, Oct. 2005.
14. F. Fang et al., “Cryptocurrency trading: A comprehensive survey,” Financial Innovation, vol. 8, no. 1, pp. 1–59, Jan. 2022.
15. M. Sifat, A. Mohamad, and M. Shariff, “Cryptocurrency volatility forecasting: A comparative study of linear, nonlinear, and deep learning models,” Journal of Risk and Financial Management, vol. 14, no. 7, p. 308, Jul. 2021.