A Heteroscedastic Neural Network Framework for Volatility Prediction

Koushik Bera
Department of Mathematics Indian Institute of Technology Guwahati
Guwahati, India
koushik.bera@iitg.ac.in

Aryan Bhambu
Department of Mathematics
Indian Institute of Technology Guwahati
Guwahati, India
a.bhambu@iitg.ac.in

Prakash Raj
Department of Mathematics
Indian Institute of Technology Guwahati
Guwahati, India
prakashraj@iitg.ac.in

N. Selvaraju
Department of Mathematics
Indian Institute of Technology Guwahati
Guwahati, India
nselvaraju@iitg.ac.in

Abstract*—Volatility forecasting of financial market returns is vital for risk management and portfolio optimization. In the quest for accurate volatility modeling, this article introduces an Artificial Neural Network – Generalized Autoregressive Conditional Heteroscedasticity model (ANN-GARCH) that fuses the strengths of neural networks with the well-known GARCH framework to deliver superior volatility predictions. The approach integrates volatility captured by conventional GARCH models into a neural network framework, combining return and realized volatility series. Then, the ANN architecture is designed to learn and process patterns, ultimately generating enhanced volatility predictions. The proposed model consistently outperforms traditional volatility forecasting methods when applied to Dow Jones Index returns during the COVID-19 period and its aftermath. Across various error distributions, the ANN-GARCH model demonstrates superior forecasting accuracy, evidenced by significantly lower root mean square error (RMSE) and mean absolute error (MAE) compared to competing models. This underscores the model's reliability and effectiveness in volatility forecasting. This innovative model will help academics and finance practitioners make informed decisions when the market conditions are stressed.

***Keywords*— Volatility, GARCH, ANN, High-Frequency Data, Finance**

I. INTRODUCTION

Volatility is a trade-off between potential gains and uncertainties that could undermine them. Volatility helps investors assess risks and minimize losses, making it crucial for effective risk management. The generalized autoregressive conditional heteroscedasticity (GARCH)-family models are among the most widely employed tools for estimating volatility in time-series analysis. These models are grounded in statistical theory as they offer ease of use and straightforward interpretability of the model parameters. However, the assumption of stationarity in explanatory variables does not match the complex and nonlinear dynamics of real-world financial markets, which puts a question mark on the efficiency of GARCH-family models.

Machine learning (ML) models have recently gained traction in volatility forecasting [4]. ML models like neural networks (NNs) need fewer assumptions and offer flexibility to approximate nonlinear complex functions [7]. However, the complex architecture often poses difficulty in succinct interpretations of parameters and model outputs. This lack of interpretation presents a major challenge in the financial sector, where clarity and transparency are vital. Though ML models handle non-stationary data, they often struggle with generalization, particularly with noisy financial time series data. The data size is imperative as the ML models work on effectively learning the dataset. The limited dataset causes overfitting, whereas larger datasets introduce inconsistencies due to the non-stationarity of inputs, further complicating the learning process for NN models.

Recent studies [10] have shown that simple linear models can outperform more sophisticated transformation-based models in long-term time series forecasting. NN models [9] are generally designed as broad ML solutions for effective time series forecasting but lack appropriate optimization for capturing the stylized facts of volatility in time series data. As a result, statistical models like GARCH [2] remain dominant in volatility forecasting due to their simplicity, interpretability, and reliable performance.

Time-series volatility models and neural network approaches have proven their worth in their respective domains and offer unique advantages. However, they are often explored independently, with limited integration between the two fields. We bridge this gap by proposing a novel approach that leverages the established principles of GARCH volatility models within NN-based frameworks [11] to enhance prediction accuracy and model interpretability. Econometric research has identified that volatility, unlike other market variables, exhibits specific, predictable characteristics such as volatility clustering, asymmetry in market trends, and long memory effects [6]. Our empirical analysis reveals that integrating volatility into neural network models enhances their accuracy beyond that of leading neural network approaches.

The proposed methodology introduces an innovative approach that extends traditional GARCH-family volatility models into a comprehensive neural network framework for high frequency volatility time series forecasting. We achieve this by transforming conventional GARCH models into their neural network equivalents, integrating the volatility captured by these models within the neural network architecture. The model performs an in-depth comparative study of distributional assumptions namely, normal, student-t (st), skew-student-t (sst), and generalized error distributions (ged) across a family of GARCH models. The extracted insights are then processed within the neural network, resulting in the ANN-GARCH approach, which merges stochastic modelling with neural network techniques. This approach leverages the robust mathematical structures and statistical properties of GARCH models, well-established in econometrics. By building on this solid foundation with multiple distributions, our neural network-based GARCH model achieves a level of trustworthiness and interpretability comparable to traditional stochastic models, enhancing the reliability of volatility prediction beyond existing baselines.

II. PROPOSED MODEL

The GARCH family has long been the preferred method for capturing the dependence in volatility time series. Let p_t denotes the closing price of day t. The GARCH $(1,1)$ model is expressed as follows:

$$r_t = \mu_t + \epsilon_t,\ \epsilon_t \mid \psi_{t-1} \sim D(0, \sigma_t^2) \tag{1}$$

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

Here, $r_t = \log(p_t) - \log(p_{t-1}) * 100$ represents the log return at time t, D stands for the distribution assumption. And ϵ_t may be observed directly or as residuals from an econometric model, and ψ_{t-1} denotes historical information. This is an exponentially decaying function due to the autoregressive term, ensuring short memory in the model. The condition $\omega > 0, \alpha \geq 0, and\ \beta \geq 0$ ensures that the conditional variance is positive, and $0 < \alpha + \beta < 1$ guarantees the stationarity of the model.

To build the architecture, we have chosen the GARCH family of models: Generalized Autoregressive Conditional Heteroscedasticity (GARCH) [2], Exponential GARCH (EGARCH) [5], Asymmetric Absolute Value GARCH (AVGARCH) [1], Fractionally Integrated GARCH (FIGARCH) [3], Glosten, Jagannathan and Runkle-GARCH (GJR-GARCH) [8], and Threshold GARCH (TGARCH) [5], which offers a versatile toolkit for modeling volatility in financial time series, each addressing different market behaviors. The classical GARCH model effectively captures volatility clustering, where large/small return movements follow large/small return movements but struggle with asymmetry. EGARCH and AVGARCH improve on this by incorporating asymmetric responses to positive and negative shocks. These models can capture the leverage effect, which indicates the inverse relationship between return shocks and volatility. FIGARCH goes further by modeling long-memory effects, making it valuable for understanding persistent volatility over extended periods. TGARCH introduces threshold effects by adding an indicator variable that takes 1 when the returns go negative. It allows for modeling nonlinear responses to shocks, mainly when market reactions differ depending on the magnitude and direction of returns. GJR-GARCH refines by squaring the variables used in TGARCH, thus easing estimation. Together, these models provide a comprehensive approach to understanding and forecasting the complex dynamics of financial volatility. The ANN-GARCH model incorporates inputs such as realized volatility, returns, and outputs from the GARCH model through a structured neural network framework. The input layer is linked to subsequent hidden layers via connections, each assigned specific weights. In this process, key features and patterns are extracted at the input layer, and these are multiplied by their respective weights. The weighted values are then propagated to the neurons in the next hidden layer. At each neuron, a bias term is added, and an activation function is applied. A dense layer is then used to generate the output, where only activated neurons transmit their results. The neuron with the highest activation value determines the final output at the output layer. During the training phase, the model calculates the difference between predicted and actual outputs, and this error is backpropagated through the network to adjust the weights iteratively. This adjustment continues until the model achieves the output with the least possible error. Figure I illustrates the architecture of the ANN-GARCH model.

FIGURE I PROPOSED ANN-GARCH ARCHITECTURE

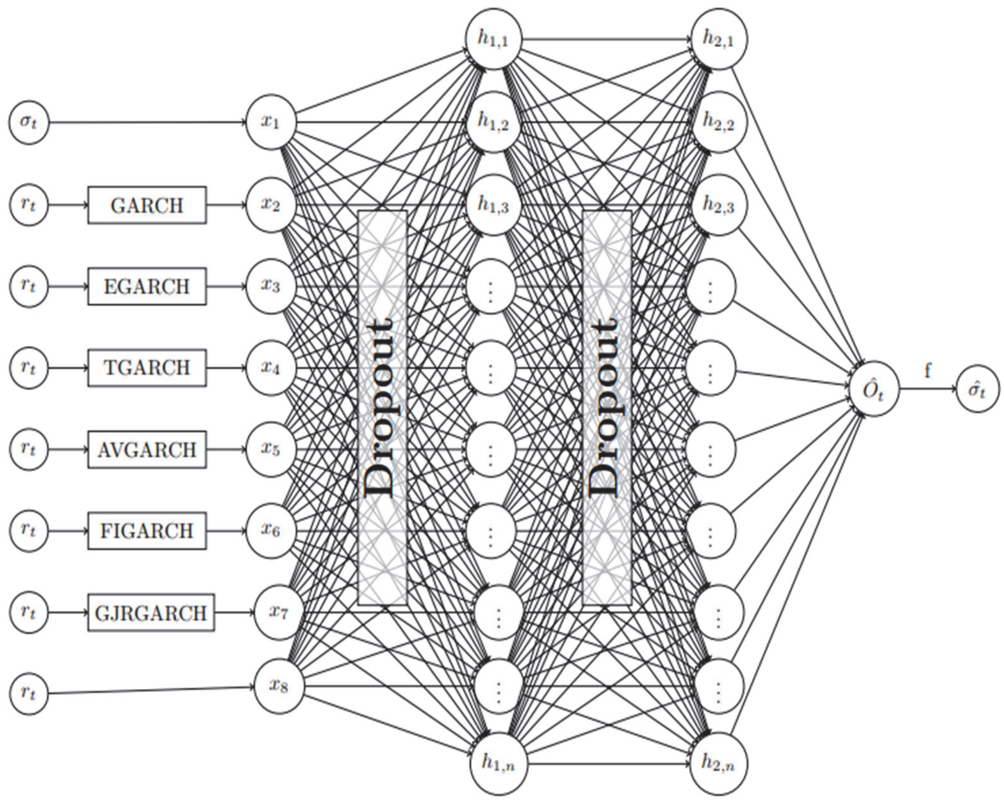

III. EMPIRICAL ANALYSIS

A. Data Description

This study utilizes the Dow Jones Industrial Average (DIA), a well-known indicator of the financial health of US markets. The dataset spans three years: 2020, 2021, and 2022, taken one period at a time. The raw data for the stock index are taken from FirstRateData (https://firstratedata.com/). The empirical investigation centers on high-frequency trading data from the NASDAQ exchange,

ISBN : 978-81-978522-2-0

specifically within regular trading hours (9:30 AM to 4:00 PM). We used 5-minute interval returns as inputs to GARCH models, complemented by realized variance calculated by summing squared returns from 1-minute intervals. This method allowed us to capture the intricate dynamics of financial instruments with greater precision. The DIA dataset consists of 19,622, 19,643, and 18,252 data points for the years 2020, 2021, and 2022, respectively. We allocated 80% of the data for training, 20% for testing, and further used 20% of the training data as a validation set.

B. Evaluation Metrics and Hyperparameter Setup

The volatility forecasting performance of the ANN-GARCH model against various comparison models is assessed using two of the most robust evaluation metrics:

Root Mean Square Error (RMSE) and Mean Absolute Error (MAE):

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n} (\hat{y}_i - y_i)^2}, \quad \textit{and} \quad \text{MAE} = \frac{1}{n}\sum_{i=1}^{n} |\hat{y}_i - y_i| \tag{3}$$

where $\hat{y}_i$ represents the predicted value, y_i is the actual value, and n is the number of observations.

Moreover, Bayesian optimization is employed to optimize the performance of the comparison models, namely ANN, Support Vector Regression (SVR), SVR-GARCH, and the proposed ANN-GARCH model. The Bayesian optimization efficiently explores the hyperparameter space to identify the optimal settings for each model. The key settings are fixed during the tuning process. The number of trials is set to 50 for all models. For the ANN, the batch size is 32, the learning rate is 0.001, the number of epochs is 100, and the number of layers is two. The hidden layer size for the ANN is tuned between 20 and 200 neurons, and the activation function is tuned between ReLU and sigmoid. For SVR and SVR-GARCH, the regularization parameter C is tuned from 10^{-3} to 10^{3} and ϵ insensitive loss parameter is adjusted from 10^{-1} to 10^{1}. The kernel is tuned among linear, polynomial, and RBF options.

C. Results and Discussions

The proposed model ANN-GARCH and the time-series volatility models such as GARCH, EGARCH, AVGARCH, FIGARCH, GJR-GARCH, and TGARCH used in the analysis are evaluated across various distributions, namely, normal, student-t, skew-student-t, and generalized error distribution.

Table I shows the volatility forecasting performance of all models for the DIA dataset regarding RMSE and MAE performance metrics. The forecasting errors highlight the effectiveness of different models under different distributions in volatility forecasting of financial time series. Across the years 2020, 2021, and 2022, the ANN-GARCH model consistently achieves the lowest RMSE and MAE values, demonstrating its superior ability to capture the complexities and volatilities inherent in financial data. Notably, the RMSE (MAE) values of 0.0008 (0.0005)/ 0.0013 (0.0005)/ 0.0199 (0.0023) for the years 2020/2021/2022 achieved by ANN-GARCH under the skew-student-t distribution results in the best performance, indicating its suitability for modeling the asymmetry and heavy tails observed in financial time series. Integration of the GARCH family with ANN significantly enhances predictive accuracy, as evidenced by the consistently lower error values achieved compared to competing standalone models.

FORECASTING ERRORS OF ALL THE MODELS FOR THE DIA DATASET

Year	Model Name	Distribution							
		RMSE				*MAE*			
		normal	*ged*	*st*	*sst*	*normal*	*ged*	*st*	*sst*
2020	ANN-GARCH	0.0014	0.0014	0.0013	0.0008	0.0009	0.0010	0.0009	0.0005
	ANN	0.0045	0.0045	0.0045	0.0045	0.0033	0.0033	0.0033	0.0033
	SVR-GARCH	0.6284	0.6284	0.6284	0.6282	0.6135	0.6135	0.6135	0.6131
	SVR	0.8237	0.8237	0.8237	0.8237	0.8223	0.8223	0.8223	0.8223
	GARCH	0.0362	0.0377	0.0377	0.0377	0.0295	0.0310	0.0310	0.0310
	AVGARCH	0.0336	0.0332	0.0330	0.0330	0.0262	0.0257	0.0253	0.0253
	EGARCH	0.0354	0.0352	0.0356	0.0356	0.0273	0.0270	0.0272	0.0272
	FIGARCH	0.0343	0.0326	0.0328	0.0328	0.0262	0.0245	0.0245	0.0246
	GJRGARCH	0.0365	0.0365	0.0365	0.0365	0.0298	0.0298	0.0298	0.0298
	TGARCH	0.0332	0.0328	0.0326	0.0326	0.0260	0.0255	0.0251	0.0251
2021	ANN-GARCH	0.0015	0.0013	0.0016	0.0013	0.0007	0.0008	0.0007	0.0005
	ANN	0.0034	0.0034	0.0034	0.0034	0.0024	0.0024	0.0024	0.0024
	SVR-GARCH	0.1202	0.1288	0.1301	0.1148	0.1196	0.1204	0.1195	0.1122
	SVR	0.1944	0.1944	0.1944	0.1944	0.1909	0.1909	0.1909	0.1909
	GARCH	0.0230	0.0242	0.0229	0.0230	0.0159	0.0171	0.0157	0.0160
	AVGARCH	0.0240	0.0235	0.0235	0.0235	0.0163	0.0168	0.0168	0.0168
	EGARCH	0.0233	0.0234	0.0238	0.0239	0.0171	0.0170	0.0173	0.0173
	FIGARCH	0.0233	0.0233	0.0234	0.0234	0.0167	0.0165	0.0165	0.0165
	GJRGARCH	0.0236	0.0232	0.0241	0.0236	0.0165	0.0160	0.0170	0.0165
	TGARCH	0.0249	0.0236	0.0236	0.0236	0.0173	0.0168	0.0167	0.0167
	ANN-GARCH	0.0204	0.0201	0.0203	0.0199	0.0022	0.0020	0.0026	0.0023
	ANN	0.0256	0.0256	0.0256	0.0256	0.0059	0.0059	0.0059	0.0059
	SVR-	0.1944	0.1944	0.1944	0.1944	0.1909	0.1909	0.1909	0.1909

ISBN : 978-81-978522-2-0

2022	GARCH								
	SVR	0.1318	0.1318	0.1318	0.1318	0.1303	0.1303	0.1303	0.1303
	GARCH	0.0679	0.0682	0.0685	0.0685	0.0329	0.0326	0.0321	0.0321
	AVGARCH	0.0708	0.0692	0.0692	0.0692	0.0326	0.0328	0.0324	0.0324
	EGARCH	0.0689	0.0687	0.0689	0.0689	0.0336	0.0326	0.0326	0.0326
	FIGARCH	0.0674	0.0678	0.0680	0.0667	0.0332	0.0326	0.0332	0.0341
	GJRGARCH	0.0684	0.0683	0.0682	0.0684	0.0321	0.0326	0.0324	0.0325
	TGARCH	0.0706	0.0692	0.0694	0.0694	0.0323	0.0328	0.0324	0.0324

The traditional GARCH models, though effective, generally yield higher error metrics than the ANN-GARCH model, suggesting they may be less capable of capturing intricate patterns in financial data. The SVR and SVR-GARCH models show the highest RMSE and MAE, indicating their limitations in this context. Our empirical findings underscore the importance of model selection and the incorporation of suitable distributional assumptions when forecasting financial time series. The ANN-GARCH model's consistent performance under different distributional assumptions and across various data periods highlights its robustness and adaptability. Thus, our proposed model is valuable for accurate volatility prediction in volatile financial markets.

IV. CONCLUSIONS

This article presents a comparative study using the ANN-GARCH model, which effectively estimates high-frequency volatility by combining inputs from GARCH-family models with an ANN framework. Existing literature often assumes a single error distribution for model estimation, limiting its practical applicability. For instance, the normal distribution is frequently used despite the well-known heavy-tailed nature of financial returns. In this article, we test our model across a range of error distributions, namely, normal, student-t, skew-student-t, and generalized error distribution. The study period spans the COVID-19 pandemic and its after-effects.

The empirical analysis demonstrates that the ANN-GARCH model consistently outperforms competing models on the most robust performance metrics, RMSE and MAE. Additionally, our model performs exceptionally well with the skew-student-t distribution, highlighting its capability to capture the complexities and volatilities inherent in financial data. The ANN-GARCH model shows better interpretability than machine learning models and improves predictive accuracy. The superior performance across COVID-19 underscores the importance of adopting adaptive models like the ANN-GARCH, especially in distressed market conditions.

REFERENCES

[1] Suhail Najm Abdulla and Heba Dhaher Alwan. Using APGARCH /AVGARCH models gaussian and non-gaussian for modeling volatility exchange rate. International Journal of Nonlinear Analysis and Applications, 13(1):30293038,2022.

[2] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3):307-327, 1986.

[3] Xuehui Chen, Hongli Zhu, Xinru Zhang, and Lutao Zhao. A novel timevarying FIGARCH

model for improving volatility predictions. Physica A: Statistical Mechanics and its Applications, 589:126635, 2022.

[4] Wenbo Ge, Pooia Lalbakhsh, Leigh Isai, Artem Lenskiy, and Hanna Suominen. Neural network-based financial volatility forecasting: A systematic review. ACM Computing Surveys (CSUR), 55(1):1-30, 2022.

[5] Yunpeng Luan, Shili Ye, Yanmei Li, Lu Jia, and Xiao-Guang Yue. Revisiting natural resources volatility via TGARCH and E GARCH. Resources Policy, 78: 102896,2022.

[6] Soheil Almasi Monfared and David Enke. Volatility forecasting using a hybrid GJR-GARCH neural network model. Procedia Computer Science, 36:246-253,2014.

[7] Intan Muchtadi-Alamsyah, Robin Viltoriano, Ferdinand Harjono, Martha Nazaretha, Martin Susilo, Ade Bayu, Bony Josaphat, Arief Hakim, and Khreshna Syuhada. Support vector regression-based heteroscedastic models for cryptocurrency risk forecasting. Applied Soft Computing, page 111792, 2024.

[8] Didit Budi Nugroho, D Kurniawati, LP Panjaitan, Z Kholil, Bambang Susanto, and LR Sasongko. Empirical performance of GARCH, GARCH-m, GJR-GARCH and log-GARCH models for returns volatility. In Journal of Physics: Conference Series, volume 1307, page 012003. IOP Publishing, 2019.

[9] Francisco Pérez-Hernández, Alvaro Arévalo-de Pablos, et al. A hybrid model integrating artificial neural network with multiple GARCH-type models and EWMA for performing the optimal volatility forecasting of market risk factors. Expert Systems with Applications, 243:122896, 2024.

[10] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In Proceedings of the AAAI Conference on Artificial Intelligence, volume 37, pages 11121-11128, 2023.

[11] Pengfei Zhao, Haoren Zhu, Wilfred Siu Hung Ng, and Dik Lun Lee. From GARCH to neural network for volatility forecast. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pages 16998-17006, 2024.