STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **12**, July 2024, pp 1042–1060.
Published online in International Academic Press (www.IAPress.org)

Enhancing Volatility Prediction: Comparison Study Between Persistent and Anti-persistent Financial Series

Youssra Bakkali*, Mhamed EL Merzguioui, Abdelhadi Akharif

LMA, FSTT, Abdelmalek Essaadi University, Tetouan, Morocco

Abstract Predicting financial volatility is crucial for managing risks and making investment decisions. This research introduces a novel method for creating a prediction model that effectively handles the intricate dynamics of financial time series data. By utilizing the advantages of both time series models and recurrent neural networks, we present two hybrid models: Vanilla-RGARCH and LSTM-RGARCH. These models are designed to overcome the shortcomings of Realized GARCH (RGARCH) and HAR models in representing various stylized facts of financial data. While RGARCH models are proficient in capturing asymmetry, they fail to address long-term memory. Conversely, HAR models are adept at capturing long-term memory. The innovative model combines forecasted values from the RGARCH model with components from the HAR model, including daily, weekly, and monthly realized volatility, within a neural network framework. This combination helps to bypass the complexities involved in directly merging the HAR model with RGARCH. Through this method, our hybrid models provide a thorough depiction of the characteristics of financial data.
The proposed approach is evaluated on two distinct types of financial series; persistent and anti-persistent, to demonstrate its robustness and capacity to generalize in different contexts. The performance of hybrid models is compared to that of conventional RGARCH and HAR models, demonstrating their superiority in precise prediction of financial volatility and their ability to capture complex trends observed in real data. In addition, a principal component analysis (PCA) is used to visualize the results and facilitate their interpretation.

Keywords Recurrent Neural Networks, Realized GARCH, Long-memory, Asymmetry

AMS 2010 subject classifications 62M10, 62M45, 68T05

DOI: 10.19139/soic-2310-5070-2021

1. Introduction

Market participants employ a variety of strategies to forecast the volatility of economic indicators. The two main groups of these techniques are classical and neural networks. Although there has been numerous research comparing artificial neural networks (ANNs) to classical models, the results of prediction methods are conflicting as to whether the neural network approach is superior to time series or not. However, the hybrid ANN and conditional autoregressive models (GARCH type) [14], generally present advantages compared to ANNs or time series models. For example, Hajizadeh et al [4] found that the hybrid ANN and Exponential GARCH (EGARCH) can provide better volatility predictions for the S&P 500 index.
In [11], various econometric approaches for forecasting time series are explored, Specifically, it delves into the examination of econometric options, with a focus on the ARIMA and GARCH models, both extensively employed in the finance sector. These methods have the unique benefit of using less computer power when training models, which increases their usefulness in real-world scenarios. This work uses the discrete Fourier transform as a preprocessing step to improve the ARIMA model's prediction, resulting in the wavelet ARIMA method.

*Correspondence to: Youssra Bakkali (Email: youssra.bakkali2@etu.uae.ac.ma). Laboratory of Mathematics and Applications, FSTT, Abdelmalek Essaadi University, Tetouan, Morocco.

ISSN 2310-5070 (online) ISSN 2311-004X (print)

Unfortunately, because ARIMA is linear, it cannot capture the non-linear patterns that are present in volatility time series data. Two hybridization methods are suggested to overcome this restriction, combining GARCH and wavelet ARIMA models. By utilizing the advantages of each model, this methodology makes it possible to capture both linear and non-linear in the time series data.
This paper aims to capture the nonlinear relationships between stock market volatility, as measured by realized variance, and a set of historical variables using ANN and time series models. Traditional anticipation networks cannot be used for learning and predicting series that present long memory dependencies. A mechanism is needed to store past or historical information to predict future values. Recurrent neural networks or RNNs in abbreviated form are a variant of conventional forward propagating artificial neural networks, which can handle sequences of data and can be trained to retain knowledge about the past.
RNNs are particularly useful for predicting volatile financial variables exhibiting long-term non-linear dependence, such as stock prices, exchange rates, and realized volatility; Another way to deal with long-term dependencies are the Long short-term memory neural network (LSTM). LSTM can work where long memory effects are present in the time series structure. Consequently, there are many applications of LSTM models in finance [7, 10].
In the realm of volatility forecasting, recent research has made significant strides by applying machine learning models to predict intraday realized volatility (RV). A notable study [2] employs innovative techniques, such as pooling stock data together to exploit commonality in intraday volatility and incorporating a proxy for market volatility. The findings of this study reveal the superiority of neural networks over traditional linear regressions and tree-based models. Neural networks, distinguished by their capability to unveil and model complex latent interactions among variables, outperform other models in terms of performance.
It should be noted that hybrid models have demonstrated their predictive ability in the literature. For example, a recent study compared volatility prediction performance between two types of ANN and GARCH hybrid models. The results showed that the EGARCH-ANN model surpasses other models in predicting volatility in the Chinese energy market, highlighting the existence of significant leverage in this market [9].
In addition, [1] examines the prediction of the volatility of major cryptocurrencies using a variety of models, including GARCH, MLP, LSTM, and LSTM-GARCH hybrids. The results reveal that deep neuron network models outperform traditional GARCH models in terms of predictive accuracy, including MLP models. This research highlights the effectiveness of deep learning approaches in predicting volatility in cryptocurrency markets.
In the pursuit of developing an effective modeling approach for financial series, this article endeavors to leverage two critical features: volatility asymmetry and long-term memory of fluctuations. A novel strategy is adopted, integrating the predictive volatility from the RGARCH model proposed by Hansen [12] into the Recurrent Neural Network (RNN), alongside components from the Heterogeneous Autoregressive (HAR) model, encompassing daily, weekly, and monthly Realized volatilities.
The incorporation of the RGARCH model into our proposed framework is motivated by its aptitude in capturing volatility symmetry and dynamics. However, it is essential to acknowledge that the RGARCH model specification alone may not suffice in capturing long-term dependencies [16].
To address the challenges posed by the nonlinearity of input-output relationships and long-term dependencies, our proposed model combines the flexibility of RNN with the statistical robustness of traditional models. The selection of input variables crucially relies on insights gleaned from the statistical characteristics of the series under examination.
The predictive capabilities of the RGARCH model, as demonstrated in prior literature, substantiate its inclusion in our proposed framework. By integrating the predicted volatility from the RGARCH model and the daily, weekly, and monthly Realized volatilities into the RNN, our proposed model presents a more comprehensive framework adept at capturing both asymmetries and long-term dependencies inherent in financial series.
In this study, we applied the proposed model to two types of time series: persistent and anti-persistent series. This approach allows us to visualize and analyze the model's performance in different contexts where the persistence characteristics of the data vary. Persistent series are characterized by a tendency to follow a given direction over an extended period, while anti-persistent series exhibit a tendency to reverse or fluctuate randomly over the same period. By examining the model's performance on these two types of series, we can better understand its ability to capture underlying trends and patterns in diverse data environments. This comparative approach also enables us to

assess the robustness and generalizability of the model in realistic scenarios of financial time series.
The remainder of this article is segmented into four sections. Section 1 introduces the research context and objectives, setting the stage for the subsequent sections. It provides an overview of the importance of forecasting volatility in financial markets and outlines the motivation behind our study. Following the introduction, we delve into the methodology section 2, where we outline the framework for our hybrid modeling approach. This section details the integration of the RGARCH and HAR models within a neural network framework, highlighting the rationale behind this approach and the steps involved in model development.
Next, section 3 presents the empirical study, which comprises two main components: assessment of forecast accuracy and Principal Component Analysis (PCA) for visualizing the results. In the assessment of forecast accuracy, we evaluate the performance of our hybrid models against traditional RGARCH and HAR models, using various metrics to measure predictive accuracy. Subsequently, we employ PCA to visualize the results and provide insights into the underlying patterns and relationships captured by our hybrid models. Finally, Section 4 provides a comprehensive conclusion, summarizing the key findings of our study and discussing their implications.

2. Methodology

In this section, we delineate the methodology employed for forecasting stock market volatility. The RGARCH model is integrated as a fundamental component of the proposed model, leveraging its widespread utilization and proven efficacy in financial forecasting. By accounting for asymmetry and volatility dynamics, the RGARCH model provides valuable insights into the underlying dynamics of market volatility.
However, it is essential to acknowledge that the classical RGARCH model specification alone may be insufficient in capturing the long dependencies inherent in financial volatility [16]. To address this inherent limitation, we augment the RGARCH model with a HAR specification, incorporating it as input variables for the RNN (Recurrent Neural Network) model.
The selection process is based on a causality test [15]. This test helped to determine the variables that have a significant causal relationship with the target output. By identifying these influential variables, we ensured that our model incorporates inputs that are causally linked to the predicted outcome, thereby enhancing its predictive accuracy and reliability. The decision to integrate the RGARCH model and elements of the HAR model, encompassing daily, weekly, and monthly realized volatility, is grounded in their well-documented effectiveness in forecasting the volatility of economic variables characterized by asymmetry and long dependency features. By leveraging the complementary strengths of these models within our hybrid approach, we aim to enhance the accuracy and robustness of volatility forecasting in financial markets.

2.1. Realized Measures

The Realized variance (RV) over a time period n is constructed as the sum of the squares intraday return introduced by Andersen et al. (2001b). The variance is given by :

$$RV_t = \sum_{i=1}^{n} r_{i,t}^2, \tag{1}$$

where $r_{i,t}$ represente the i^{th} return on day t.

2.2. HAR model

We define the variables:

$$RV_t^w = 1/5 \sum_{i=1}^{5} RV_{t-i}, \tag{2}$$

and

$$RV_t^m = 1/22 \sum_{i=1}^{22} RV_{t-i}, \tag{3}$$

The HAR model of [5] is written in the form:

$$RV_t = \beta_0 + \beta_1 RV_{t-1} + \beta_5 RV_{t-1}^w + \beta_{22} RV_{t-1}^m + u_t. \tag{4}$$

where u_t represents the error term, while β_0, β_1, β_5, and β_{22} denote real parameters.

2.3. RGARCH model

The RGARCH model proposed by Hansen et al. (2012) is given by:

$$\begin{cases} \mathrm{r}_t = \sqrt{h_t} z_t, \; z_t \sim N(0,1), \\ \mathrm{h}_t = w + \sum_{i=1}^{q} \gamma_i x_{t-i} + \sum_{i=1}^{p} \beta_i h_{t-i}, \\ \mathrm{x}_t = \xi + \phi h_t + \eta_1 z_t + \eta_2 (z_t^2 - 1) + \epsilon_t, \; \epsilon_t \sim N(0,1) \end{cases} \tag{5}$$

where (r_t) represent the return, (h_t) is the conditional variance and (x_t) is the realized measure. Additionally, w, γ_i $(i = 1, .., q)$, β_i $(i = 1, .., p)$, ξ, ϕ, η_1, and η_2 are all real parameters.
To guarantee the stationarity and positivity in the RGARCH(1,1) model, the required conditions are as follows:

$$w + \gamma\xi \; < 0, \tag{6}$$

$$0 \; < \beta + \gamma\phi \; < 1 \tag{7}$$

Additionally, it is sufficient for all of the parameters w, β, and γ to be positive to ensure the positivity of each h_t.

2.4. Neural networks

ANN are computer models inspired by the human brain. The structure of the network can be modified to approximate a wide range of statistical and economic models. For this reason, ANNs have been widely used to predict time series in different domains, such as finance, engineering, and physics.
The empirical research indicates that ANNs are particularly well suited to predicting volatile financial variables that exhibit nonlinear behaviors, such as stock returns or stock market volatility [8].

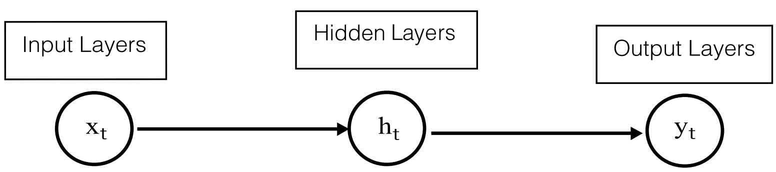

Figure 1. Forward Propagation Neural Network (FNN).

A FNN allows information to flow only in the forward direction from the input nodes, through the hidden layers. A simplified representation of the FNN is shown in Figure 1.
In a FNN, the hidden state h at time t is computed by applying a non-linear activation function to the sum of the input at time t (x_t) multiplied by its associated weight (W_x). A bias term b is also added. This process can be mathematically expressed as:

$$h_t = f(W_x * x_t + b), \tag{8}$$

Subsequently, the output y at time t is determined by applying a second linear transformation to h_t.

$$y_t = W_y * h_t + b_y \tag{9}$$

where W_y and b_y represent the weight and bias associated with y_t .

In a FNN, decisions are based on the current input. It does not store past data. FNN neural networks are used in general problems of regression and classification.
A recurrent neural network (RNN) is a type of artificial neural network that uses sequential or time series data. They are distinguished by their memory, they take information from previous entries to influence the current input and output.
A simplified presentation of a RNN (Figure 2):

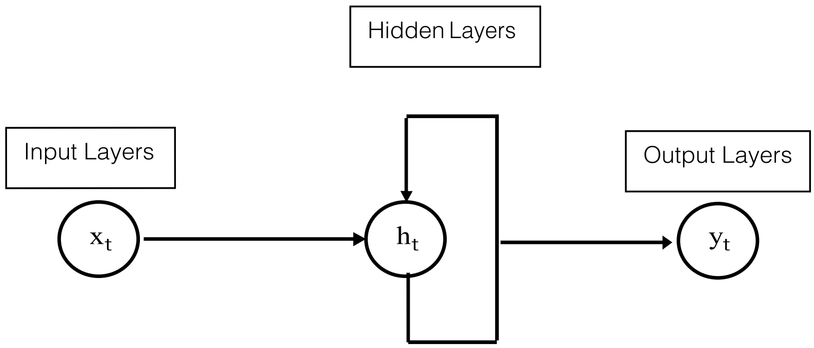

Figure 2. Standard Recurrent Neural Network (Vanilla-RNN).

The hidden state h at time t is determined by applying a non-linear activation function to the sum of the input at time (x_t) multiplied by its corresponding weight (W_x), and the previous time step's hidden state (h_{t-1}) multiplied by its corresponding weight (W_h). Additionally, a bias term b is included. This can be expressed mathematically as:

$$h_t = f(W_x * x_t + W_h * h_{t-1} + b), \tag{10}$$

The output y at time t is subsequently calculated by subjecting the hidden state h_t to a second linear transformation.

$$y_t = W_y * h_t + b_y \tag{11}$$

where W_y and b_y represent the weight and bias corresponding to y_t.
Vanilla RNNs suffer from the so-called "vanishing gradient problem". Long Short-Term Memory (LSTM) was introduced by Hochreiter and Schmidhuber [13] to alleviate the vanishing gradient problem using a mechanism based on memory cells.
LSTM introduces a memory unit to allow the capture of long dependencies in a sequence. Therefore, LSTM networks can selectively remember or forget information and can learn thousands of selectively remembered information through structures called cell states.

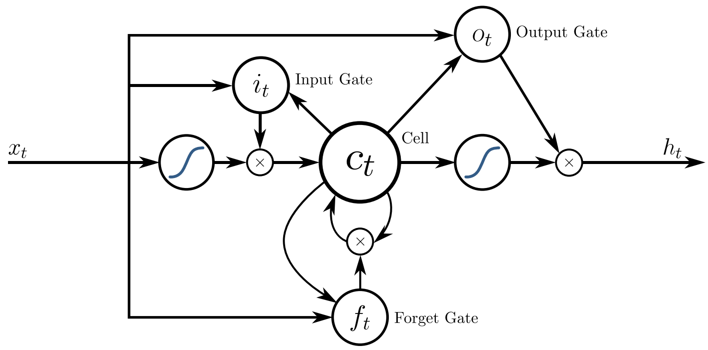

Figure 3. Basic LSTM memory cell.

Figure 3 presents the configuration of a basic LSTM memory block utilizing a single-cell architecture. In this depiction, the following notations are used:

- x_t signifies the vector of input variables at time t.
- c_t and c_{t-1} represent the cell state at time t and $t-1$ respectively.
- h_t and h_{t-1} denote the hidden state or output of the cell at time step t and $t-1$.
- i_t denote the input gate.
- f_t indicates the forget gate.
- o_t signifies the output gate.

Similar to traditional RNNs, both the input and output gates in LSTM serve similar functions. However, a notable addition is the f_t gate, known as the forget gate, responsible for eliminating unnecessary information from the cell state. At time t, the information represented by x_t and h_{t-1} undergoes evaluation through the forget gate f_t, which determines whether the information should be retained. This determination is made using a sigmoid function, where a response close to zero indicates discarding the information, while a response close to one implies retention. Simultaneously, the input gate processes the same information to contribute to the cell state c_t. Additionally, a non-linear layer, represented by $\phi = tanh$, generates a vector of candidate values $\tilde{c_t}$ to update the state of c_t. The output gate regulates the output values of an LSTM cell by employing a logistic function to filter the output. Subsequently, the final output of the memory cell, h_t, is computed by passing the cell state c_t through a $tanh$ layer and multiplying it by the value of the output gate. This entire process can be summarized by the following equations:

$$f_t = \sigma(W_f h_{t-1} + U_f x_f + b_f) \tag{12}$$
$$i_t = \sigma(W_i h_{t-1} + U_i x_f + b_i) \tag{13}$$
$$\tilde{c_t} = tanh(W_c h_{t-1} + U_c x_f + b_c) \tag{14}$$
$$c_t = f_t \circ c_{t-1} + i_t \circ \tilde{c_t} \tag{15}$$
$$o_t = \sigma(W_0 h_{t-1} + U_0 x_f + V_0 c_t + b_0) \tag{16}$$
$$h_t = o_t \circ tanh(c_t) \tag{17}$$
$$\tilde{y_t} = h_t \tag{18}$$

In this context, W_f, W_i, W_c, W_o, U_f, U_i, U_c, and U_o represent the weight matrices of the forget, input, memory cell state, and output gates respectively. V_c is the weight matrix of the cell state, $\tilde{y_t}$ is the output of the neural network, b_f, b_i, b_c, and b_o are the biases associated with the related gates. The function σ represents a sigmoid or logistic function, and $\circ$ denotes the Hadamard product function.

2.5. Hybrid model

Financial markets exhibit complex dynamics that necessitate sophisticated modeling approaches to accurately capture their behavior. Traditional models such as RGARCH and HAR each offer valuable insights into specific aspects of financial data, yet they may fall short in comprehensively addressing the diverse characteristics inherent in market volatility. In response to these challenges, this study proposes a novel hybrid modeling approach that integrates the strengths of both statistical models and neural networks to enhance the predictive capabilities of volatility forecasting.

The hybrid model presented in this study, denoted as Vanilla-RGARCH for the standard neural network-based component and LSTM-RGARCH for the long short-term memory (LSTM) network-based component, represents a unique fusion of statistical modeling techniques and advanced machine learning methodologies. By combining RGARCH and HAR models with the flexibility and complexity-capturing abilities of neural networks, the hybrid model aims to overcome the limitations of individual models and provide a more robust framework for volatility prediction.

At the core of the hybrid model lies the integration of predicted volatility from the RGARCH model and key elements of the HAR model specification, including daily, weekly and monthly realized volatility into the neural network architecture. This integration allows for a comprehensive analysis of volatility asymmetry and long-memory dependency, two critical factors influencing financial market dynamics. Leveraging the predictive power of the RGARCH model and the flexibility of neural networks, the hybrid model seeks to offer a more nuanced and accurate representation of financial data characteristics. The set of input variables used in RNN models is the following:

$\hat{y}_{t+1}$= $\{ \hat{RV}_{t+1}, RV_t, RV_t^w, RV_t^m \}$.

where $\hat{y}_{t+1}$ represents the forecasted volatility by the proposed hybrid model, and $\hat{RV}_{t+1}$ is the forecasted volatility by the Realized GARCH model.

The Figure 4 illustrates the descriptive diagram of this model.

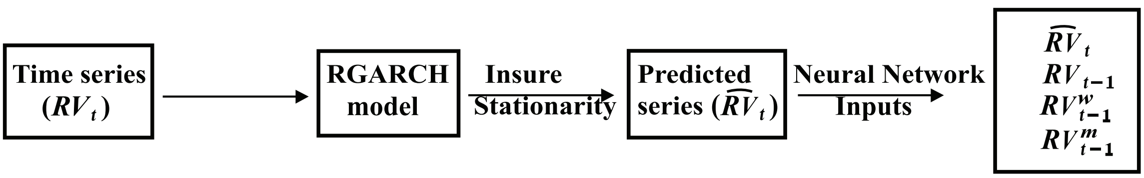

Figure 4. Schematic Representation of the Model.

In this study, our focus was on two types of neural networks, named Vanilla (Standard) recurrent networks and LSTM networks respectively. The first hybrid model will subsequently be referred to as Vanilla-RGARCH, while the second model will be designated as LSTM-RGARCH.

In the model development process, after identifying a suitable set of input variables, the subsequent step involves determining the optimal architecture of the neural network. This includes selecting the appropriate number of hidden layers and neurons within each layer. To refine the architecture and enhance performance, iterative adjustments are made by varying the number of hidden units and modifying network connections. These adjustments aim to optimize the network's ability to learn complex patterns and relationships within the data. Evaluation of different architectures is conducted by comparing the mean squared error (MSE) achieved by each configuration. Notably, there is no standardized method for determining the ideal number of hidden nodes. Therefore, we followed an empirical approach suggested by Tang and Fishwick (1993), where we systematically increased the number of hidden nodes until the training MSE reached its minimum value. This iterative process enables the neural network's structure to be fine-tuned for improved performance and predictive accuracy.

It is relevant to highlight that the existence of the hybrid model largely relies on the characteristics and stationarity conditions of the RGARCH model, whose predicted value is utilized as input for the neural network.

The RGARCH model, being a well-established model in financial volatility modeling, is designed to capture the stationarity properties of financial time series. The use of the predicted value from the RGARCH model as input for the neural network thus relies on the assumption that the underlying process is stationary, ensuring the stability of the hybrid model.
In our Simple Recurrent Neural Network (Vanilla-RNN) architecture, we utilized two hidden layers: the first hidden layer consists of 12 units and employs the Rectified Linear Unit (ReLU) activation function. Following this, we added another hidden layer comprising a fully connected Dense layer with 8 units, also employing the ReLU activation function. The network was trained using the Huber loss function, which is particularly robust to outliers, making it suitable for regression tasks like ours. For the LSTM model, we adopted a simplified structure with a single hidden layer comprising Long Short-Term Memory (LSTM) units. Similar to the Vanilla-RNN model, we utilized the Huber loss function as the loss function for training the network. During our experiments, we explored various network parameters and architectures to optimize performance. This iterative process allowed us to identify the most effective model configuration for achieving accurate predictions.
The Adam (Adaptive Moment Estimation) introduced by Kingma and Ba [3] has been used to train the Vanilla-RGARCH and LSTM-RGARCH architectures. The Adam algorithm has several benefits, including rapid convergence, and adaptability to various neural network topologies and datasets.
We utilized an early stopping technique, which automatically terminates the training process when the validation loss stops improving. This approach effectively prevents overfitting and ensures that the neural network is trained for the optimal number of epochs.
Thus, the existence of the hybrid model relies on the robustness of the RGARCH model and its stationarity conditions, while convergence is guaranteed by well-established optimization methods used in the neural network learning process. These combined elements enhance the theoretical robustness of the hybrid model and its ability to provide accurate predictions of financial volatility.

3. Empirical study

We employed data on the daily close-to-close returns and daily realized measures from the Oxford-Man Institute of Quantitative Finance. Which contains non-parametric daily measures of the volatility of financial indexes from 2000 to 2017 of the four indexes S&P 500, FTSE 100, Hang Seng, and All Ordinaries. We opted to utilize the realized variance (RV) as the realized measure due to its robustness against market microstructure noise. The sample period is from 28/05/2014 to 28/11/2017.
To establish an efficient neural network, the dataset used is divided into two subsets. we used 80% to train the model and 20% for testing model.
The realized variance of the S&P 500, FTSE 100, Hang Seng, and All Ordinaries stock indexes are shown in Figure 5. The realized volatility is highly persistent, as shown in Figure 5. The Hurst exponent is used as a measure of long-term memory of time series (Table 1).
Our findings revealed that the Hurst exponent values for stock indexes volatility S&P 500 and FTSE 100 were

Table 1. Hurst exponent test.

Stock Index	S&P 500	FTSE 100	Hang Seng	All Ordinaries
Empirical Hurst exponent	0.4157	0.4813	0.5398	0.77

below 0.5, suggesting a mean-reverting or anti-persistent behavior in the stock market volatility. This indicates that these indexes tend to fluctuate around a mean value and exhibit a tendency to revert to the average over time. In contrast, the Hurst exponent values for Hang Seng and All Ordinaries were above 0.5, indicating a trending behavior. This implies that these indexes demonstrate persistent and trending patterns over time, without significant mean reversion.
Financial returns volatility exhibits complex dynamics that are challenging to capture with traditional econometric models. These properties are known as stylized facts because they are common across a wide range of time

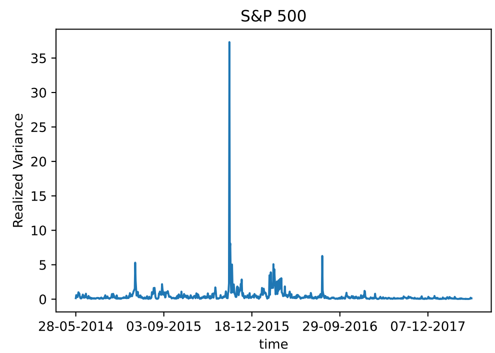

(A) S&P 500 Realized Variance

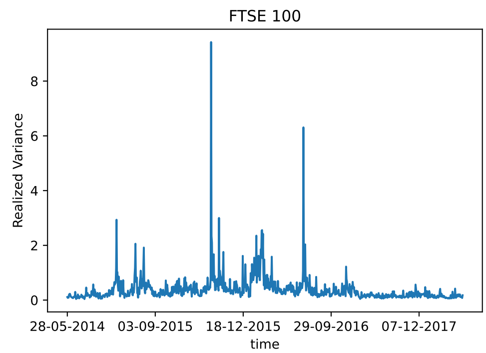

(B) FTSE 100 Realized Variance

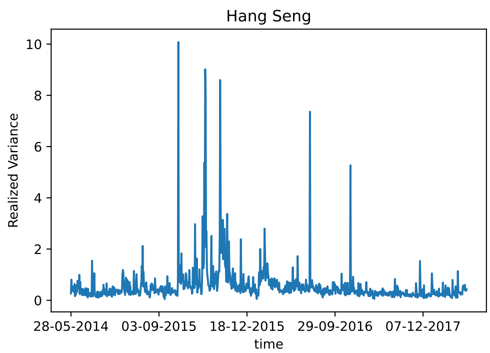

(C) Hang Seng Realized Variance

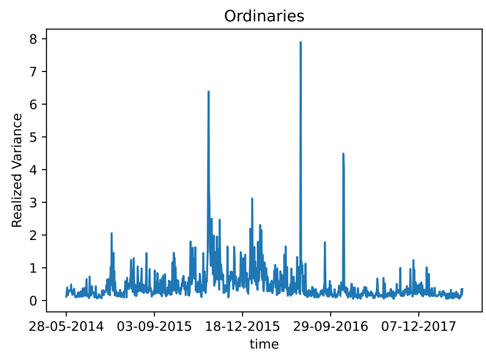

(D) All Ordinaries Realized Variance

Figure 5. Realized Volatility

series studied in finance. These stylized facts include the presence of asymmetry, volatility clustering, and long-term memory or persistence. These characteristics highlight the non-linear and interconnected nature of financial markets.

Firstly, the long memory detected in our time series represents past serial observations that are correlated with the current observation. This requires a certain class of models to describe this type of behavior. The most popular class is the HAR model.

Secondly, we highlighted the relationship between lagged returns and current market volatility (leverage). Several studies have confirmed that stock market volatility reacts asymmetrically to past returns (i.e. $\beta < 0$). Leverage describes a negative relationship between returns and volatilities; this was discussed by Black (1976) and Christie (1982).

In this context, we tested the "leverage" hypothesis. We estimated the realized volatility following [6], which presents the realized variation as follows:

$$RV_t = \alpha + \beta R_t + u_t, \tag{19}$$

where RV is the realized volatility and R is the return. Table 2 shows the impact of realized volatility on return. The negative coefficient suggests an inverse relationship between the realized volatility and the return.

In our study, we apply RGARCH and HAR models to four stock market indices to assess the extent to which the proposed Hybrid model captures the stylized facts compared to models RGARCH and HAR models, which are considered reference models.

Table 2. Estimation parameters.

Stock Index	Parameters	Coefficient	P-Value
S&P 500	$\hat{\alpha}$	0.5337	$\prec 2e-16$
	$\hat{\beta}$	-0.2045	1.72e-06 ***
FTSE 100	$\hat{\alpha}$	0.4006	$\prec 2e-16$
	$\hat{\beta}$	-0.0448	0.0166 *
Hang Seng	$\hat{\alpha}$	0.5337	$\prec 2e-16$ ***
	$\hat{\beta}$	-0.0843	0.000255 ***
All Ordinaries	$\hat{\alpha}$	0.3093	$\prec 2e-16$ ***
	$\hat{\beta}$	-0.0330	0.0156 *

The symbols "*", "**" and "***" denote statistical significance at the 0.05, 0.01, and 0.001 levels, respectively.

The estimation of the HAR model parameters for the in-sample data related to the S&P 500, FTSE 100, Hang Seng, and All Ordinaries stock indexes is provided in Table 3. Additionally, the RGARCH model, which aims to capture the characteristics of the return series, is estimated and presented in Table 4.

Table 3. Estimation of HAR model parameters.

Stock Index	Parameters	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_5}$	$\hat{\beta_{22}}$
S&P 500	Estimation	0.2201	0.2211	0.2053	0.1803
	P-Value	0.0396*	0.3571	0.9600	0.0292 *
FTSE 100	Estimation	0.1120	0.2187	0.3143	0.2177
	P-Value	0.0165 *	0.1608	0.0171 *	1.31e-05 ***
Hang Seng	Estimation	0.2532	0.3354	0.3112	0.0879
	P-Value	0.000455 ***	0.026290 *	0.961591	8.51e-05 ***
All Ordinaries	Estimation	0.3011	0.3586	0.2717	0.1351
	P-Value	0.002082 **	0.227888	0.006089 **	0.000151 ***

The notation "*", "**" and "***" signifies statistical significance at the 0.05, 0.01, and 0.001 levels.

Table 4. Estimation of RGARCH model parameters.

Stock Index	Parameters	$\hat{\omega}$	$\hat{\gamma}$	$\hat{\beta}$	$\hat{\xi}$	$\hat{\phi}$	$\hat{\eta_1}$	$\hat{\eta_2}$
S&P 500	Estimation	0.0239	0.2143	0.6886	-0.3544	0.6771	-0.0790	0.1932
	P-Value	0.4001	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000
FTSE 100	Estimation	-0.0038	0.2297	0.7139	-0.2170	1.0197	-0.0797	0.1241
	P-Value	0.8644	0.0000	0.0000	0.0056	0.0000	0.0000	0.0000
Hang Seng	Estimation	-0.0118	0.1944	0.7184	-0.1018	1.2064	-0.0511	0.1293
	P-Value	0.4834	0.0000	0.0000	0.1394	0.0000	0.0001	0.0000
All Ordinaries	Estimation	0.0678	0.1949	0.7459	-0.4901	1.0968	-0.0460	0.1963
	P-Value	0.0001	0.0000	0.0000	0.0056	0.0000	0.0064	0.0000

It is important to note that not all the values of the estimated parameters in the HAR and RGARCH models (Tables 3 and 4) are significant. Indeed, after estimating the models, it appeared that some of the estimates of the parameters did not reach an adequate level of statistical significance. This finding raises concerns about the robustness and reliability of the models in their ability to capture the underlying time relationships in the data. These results highlight the need for more in-depth analysis and the use of alternative methods to better understand and model the dynamics of the underlying financial data.

Before proceeding with the residual normality analysis, it is essential to understand the nature of these residues and their suitability to the model assumptions. Residues represent the differences between the observed values and the values predicted by the model. Ideally, these residues should be random and follow a normal distribution, which would indicate that the model captures data variability correctly and does not present systematic bias.
To assess the normality of residues, we will use several methods. First of all, we will look at the histogram of residues (Figures 6, 8, and 10) which will give us a visual representation of their distribution. A symmetrical, bell-shaped histogram would suggest a normal distribution. Next, we will look at the quantum-quantum chart (QQ plot), which compares the quantum of residues with those of a normal distribution (Figures 7, 9, and 11). A linear distribution of the points on the reference right would indicate a good correspondence to normality.
In addition to these graphical analyses, we will use the Kolmogorov-Smirnov test (Tables 5, 6, and 7) to statistically assess the adequacy of residues to a normal distribution. This test compares the empirical distribution of residues with a theoretical normal distribution and provides a measure of the gap between the two distributions.
In order to limit visual complexity and focus on a more targeted analysis, we chose to perform the normality test of residues on only two selected indices, representing persistent and anti-persistent and persistent characteristics, respectively. This approach will allow us to assess the normality of residues in different contexts while reducing the number of figures needed to present the results clearly and concisely.

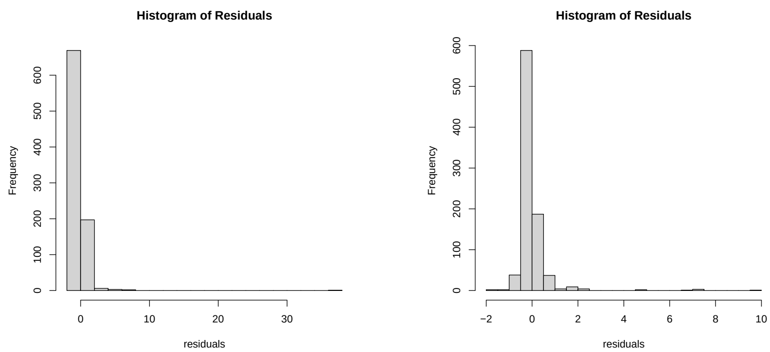

(A) S&P 500 (B) Hang Seng

Figure 6. Histogram of the HAR model residuals

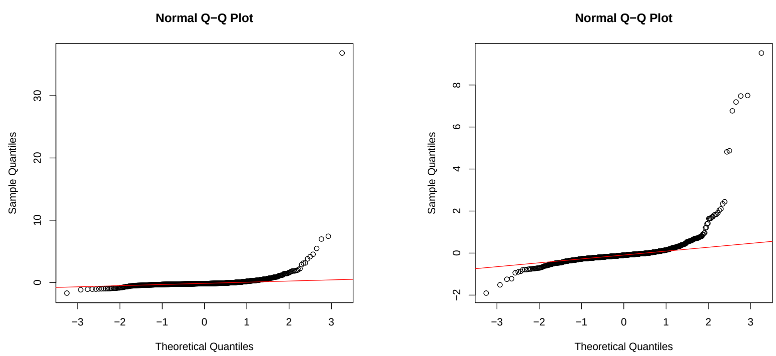

(A) S&P 500 (B) Hang Seng

Figure 7. QQ plot of the HAR model residuals

We will focus this analysis on the LSTM-RGARCH model, as it seems to capture long-term dependencies better than the Vanilla-RGARCH Model. By evaluating the normality of residues of this model, we can better understand

Table 5. The results of the normality tests of the HAR model residuals for each stock index.

Stock Index	S&P 500	Hang Seng
Kolmogorov-Smirnov Test ($p - value$)	$\prec 2.2e - 16$	$\prec 2.2e - 16$

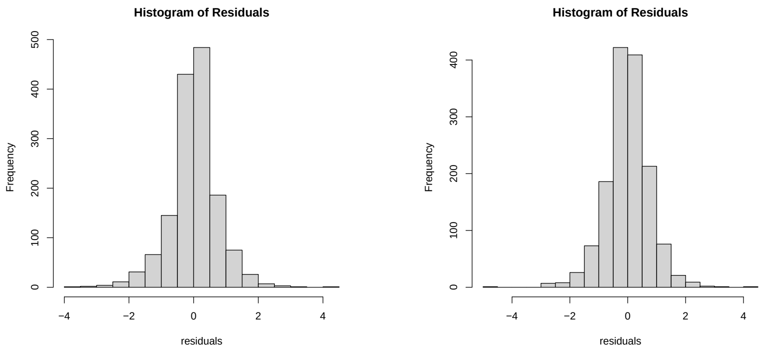

(A) S&P 500 (B) Hang Seng

Figure 8. Histogram of the RGARCH model residuals

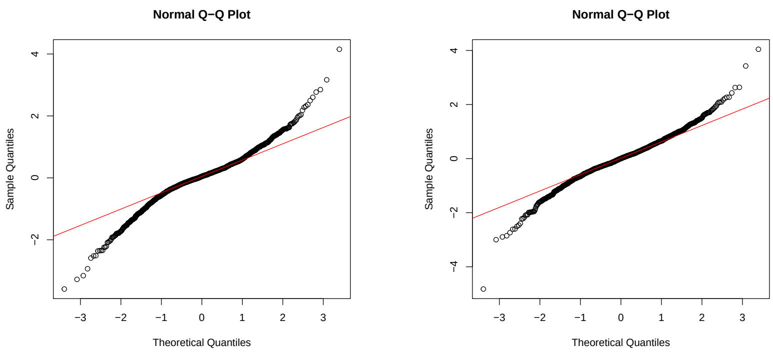

(A) S&P 500 (B) Hang Seng

Figure 9. QQ plot of the RGARCH model residuals

Table 6. The results of the normality tests of the RGARCH model residuals for each stock index.

Stock Index	S&P 500	Hang Seng
Kolmogorov-Smirnov Test ($p - value$)	$\prec 2.2e - 16$	$\prec 2.2e - 16$

its ability to adequately represent data variability and identify possible gaps in its performance. The results of this analysis will help to clarify the interpretation of the results of the hybrid model and guide any necessary improvements or adjustments.

Table 7. The results of the normality tests of the LSTM-RGARCH model residuals for each stock index.

Stock Index	S&P 500	Hang Seng
Kolmogorov-Smirnov Test ($p - value$)	0.05062	0.03042

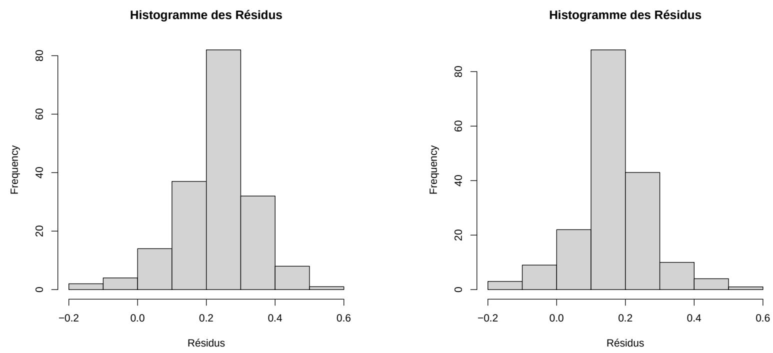

(A) S&P 500 (B) Hang Seng

Figure 10. Histogram of the LSTM-RGARCH model residuals

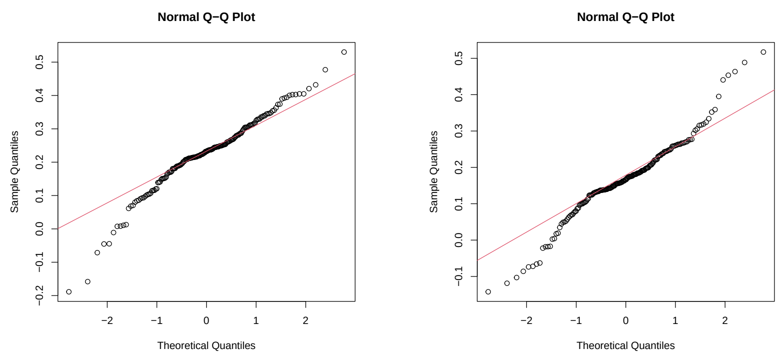

(A) S&P 500 (B) Hang Seng

Figure 11. QQ plot of the LSTM-RGARCH model residuals

For models HAR and RGARCH, the p-value is very close to zero ($p - value \prec 2.2e - 16$), indicating an extremely strong rejection of the null hypothesis of normality. However, for the LSTM-RGARCH model, the p-value of S&P 500 Stock Index is slightly higher than 0.05, suggesting there is less evidence to reject the null hypothesis of normality at the 0.05 significance level.

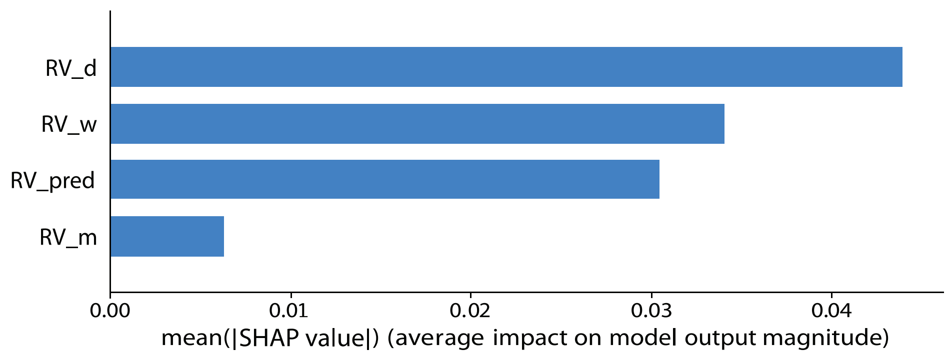

Figure 12. SHAP Values Analysis for LSTM-RGARCH Model of S&P500 Index.

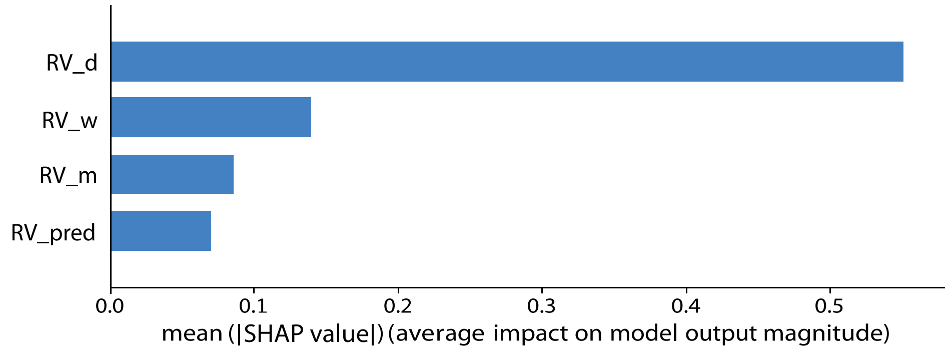

Figure 13. SHAP Values Analysis for LSTM-RGARCH Model of Hang Seng Index.

After conducting the normality test on the residuals of the LSTM-RGACH model, we proceeded with a SHAP (SHapley Additive exPlanations) analysis to assess the individual contributions of variables to the predicted value generated by the model. This approach allows us to quantify the importance of each variable in the prediction process, highlighting those that have a significant influence on the outcomes. By using SHAP values, we can pinpoint variables that substantially contribute to the variance of the predicted value, providing us with a deeper understanding of the internal workings of the proposed LSTM-RGARCH model. This information enables us to refine our interpretation of the results and identify potential areas for improvement or optimization of the model.
In Figure (12,13), the notations 'RV-d', 'RV-w', 'RV-m', and 'RV-pred' are used to represent the variables $RV_t, RV_t^w, RV_t^m, \hat{RV}_{t+1}$ respectively.
Analyzing the contributions of the variables 'RV-d', 'RV-w', 'RV-m', and 'RV-pred' in the context of stock indices, we observe distinct trends. The variables 'RV-d', 'RV-w', and 'RV-pred' seem to play a preponderant role in the case of the S&P 500 index, characterized by anti-persistence in its movements. On the other hand, in the case of the Hang Seng index, the variables 'RV-d', 'RV-w', and 'RV-m' show a strong contribution, which corresponds to the persistent nature of the movements observed in this market.

3.1. Assessment of forecast accuracy

In this study, we aim to harness the capabilities of neural networks and time series analysis to predict stock market movements. By incorporating these advanced techniques, we hope to capture and exploit the stylized facts prevalent in financial time series. We investigate the out-of-sample prediction accuracy of our model for predicting the volatility of the S&P 500, FTSE 100, Hang Seng, and All Ordinaries.
The predicted volatilities of statistical models and Hybrid models Vanilla-RGARCH and LSTM-RGARCH are shown in Fig.14.
The 180 out-of-sample observations covering the period 17 March 2017 to 28 November 2017 were used to generate forecasts for realized volatility. The percentage of out-of-sample observations was equal to 20% of the entire sample.
The mean squared error (MSE), mean absolute error (MAE) and mean absolute percentage error (MAPE), which are members of the family of loss functions resilient to a noisy volatility proxy, were used to evaluate the relative efficiency of the out-of-sample prediction accuracy.

$$MSE = 1/N \sum_{i=1}^{N} (RV_{iobs} - \hat{y}_i)^2 \tag{20}$$

$$MAE = 1/N \sum_{i=1}^{N} |RV_{iobs} - \hat{y}_i| \tag{21}$$

$$MAPE = 1/N \sum_{i=1}^{N} |RV_{iobs} - \hat{y}_i| / RV_{iobs} \tag{22}$$

where N is the number of out-of-sample observations, RV_{iobs} is the observed volatility at forecast period i, measured as the sum of the squares intraday return, and $\hat{y}_i$ is the forecast volatility.

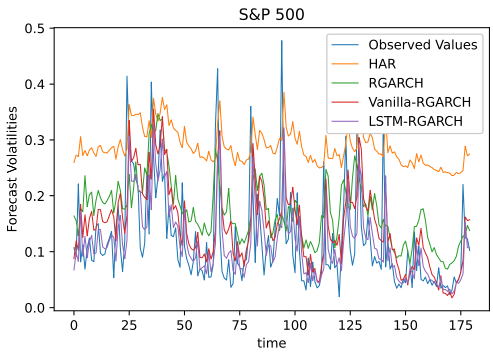

(A) S&P 500 predicting volatility

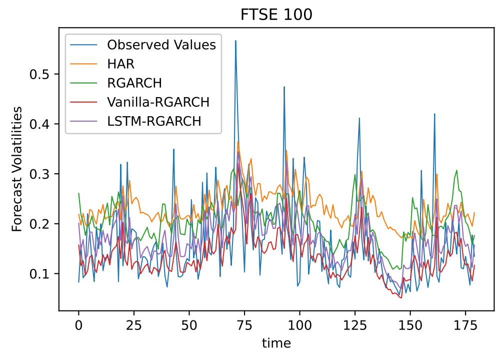

(B) FTSE 100 predicting volatility

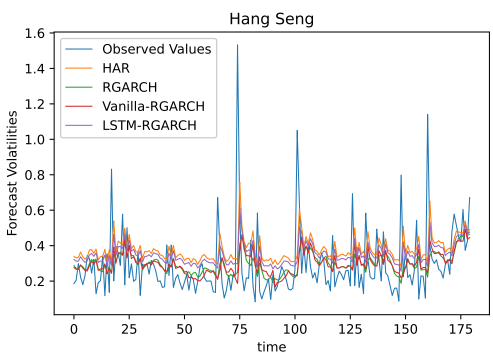

(C) Hang Seng predicting volatility

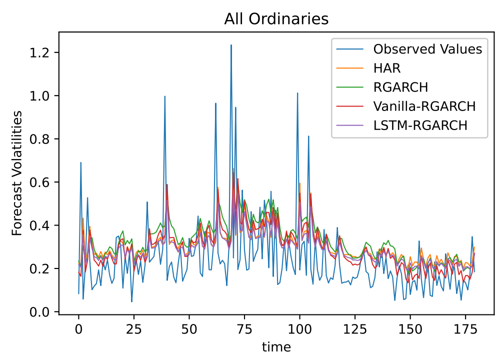

(D) All Ordinaries predicting volatility

Figure 14. Volatility predicted testing results.

Table 8. The results of the volatility forecast.

Stock Index	Model	MSE	MAE	MAPE
S&P 500	HAR	0.0314	0.1695	0.6073
	RGARCH	0.0078	0.0751	0.4503
	Vanilla-RGARCH	0.0057	0.0529	0.3389
	LSTM-RGARCH	0.0055	0.0493	0.3306
FTSE 100	HAR	0.0097	0.0859	0.7056
	RGARCH	0.0081	0.0739	0.5940
	Vanilla-RGARCH	0.0068	0.0647	0.4136
	LSTM-RGARCH	0.0062	0.0586	0.4380
Hang Seng	HAR	0.0448	0.1635	0.44
	RGARCH	0.0369	0.1209	0.4459
	Vanilla-RGARCH	0.0377	0.1299	0.4394
	LSTM-RGARCH	0.0373	0.1349	0.4401
All Ordinaries	HAR	0.0371	0.1395	0.4315
	RGARCH	0.0381	0.1495	0.4382
	Vanilla-RGARCH	0.0336	0.1243	0.43
	LSTM-RGARCH	0.0332	0.1231	0.4304

The out-of-sample predictions represented graphically in Figure 14 and Table 8 demonstrate that the proposed model outperforms statistical methods in forecasting volatility according to MSE, MAE, and MAPE.

3.2. Principal Component Analysis (PCA)

To analyze and visualize the obtained results, we employed principal component analysis (PCA), which allowed us to explore the relationships between the stock market indices (individuals) and the quantitative variables (MAPE of methods) used to describe them.
The table 9 comprises four stock market indices utilized in this study, which are considered as individual entities.
These indices are treated as individuals in the context of the Principal Component Analysis (PCA). Additionally, the table includes Mean Absolute Percentage Error (MAPE) values associated with each modeling approach as quantitative variables.
Furthermore, two supplementary variables are included in the table:
Average: This variable represents the average MAPE value calculated across all four models. It provides an aggregated measure of the MAPE for the dataset.
Amplitude: which is calculated as the difference between the largest and smallest MAPE values. It serves as an indicator of the range or variability in the MAPE across the different models.
Lastly, the table includes an additional qualitative variable labeled 'Hurst test'. This variable utilizes the notations 'AP' to denote Anti-Persitent and 'P' to denote Persistent characteristics. These qualitative labels offer insights into the persistence properties of the data.

we performed a Principal Component Analysis (PCA) to gain insights into the relationships between various modeling approaches for our dataset. Our results reveal interesting patterns among the different models. Notably, we observed correlations between the Vanilla-RGARCH and LSTM-RGARCH models, suggesting a similarity in their behavior, as well as correlations between the HAR and RGARCH models, indicating another distinct grouping. When examining the graph of variables (Figure 15), it becomes evident that the LSTM-RGARCH model closely aligns with Principal Component 1, while the RGARCH model aligns itself more with Principal Component 2. This suggests that the LSTM-RGARCH model captures variability primarily along the first principal component, while the RGARCH model is more influenced by the second principal component.
The cumulative percentage of variability associated with the axes in the representation space reached a satisfactory level of 99.99%. This high value indicates that the representation derived from the PCA is reliable and effectively

Table 9. PCA Data Table.

Stock Index	S&P 500	FTSE 100	Hang Seng	All Ordinaries
HAR	0.6073	0.7056	0.44	0.4315
RGARCH	0.5940	0.5940	0.4459	0.4382
Vanilla-RGARCH	0.4503	0.4136	0.4394	0.43
LSTM-RGARCH	0.3306	0.4380	0.4401	0.4304
Hurst test	AP	AP	P	P
Average	0.4317	0.5378	0.4413	0.4325
Amplitude	0.2767	0.2920	0.0065	0.0082

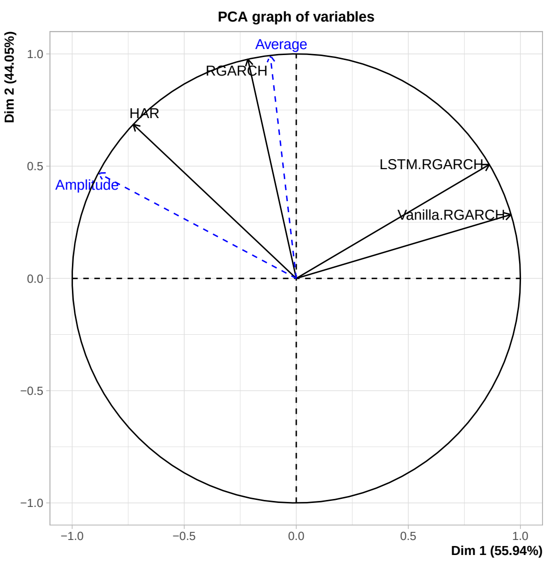

Figure 15: Graph of quantitative variables

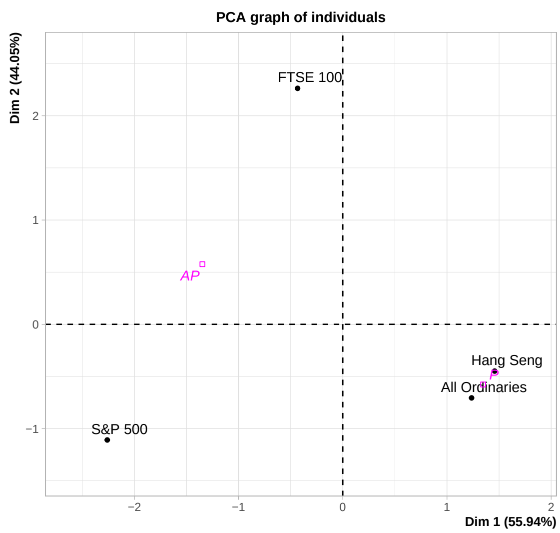

Figure 16: PCA graph of individuals.

captures the underlying structure of the data.
When examining the individual graph (Figure 16), a distinctive pattern emerges. Two indices, S&P 500 and FTSE 100 characterized by anti-persistent behavior are situated in the negative region along the x-axis and are notably distant from the origin along the y-axis. In contrast, two other indices, Hang Seng and All Ordinaries characterized by persistent behavior are positioned in the positive region along the x-axis and approach the origin along the y-axis. This divergence in the positioning of indices underscores a significant separation in their behavior.

4. Conclusion

In this study, volatility is predicted using a hybrid model that combines time series and recurrent neural networks, the article's major objective was to determine whether hybrid models could capture nonlinear interactions and offer more precise forecasts than conventional econometric techniques.
The most important feature of the stock market volatility series is the long memory. The Hurst exponent is employed as a measure of the long-term memory of four stock market volatility.
Our empirical findings demonstrate the varying performance of statistical models compared to the proposed model based on ANN. For the S&P 500 Stock index, the statistical models RGARCH and HAR show MAPE values of 45% and 60.73% respectively. Similarly, for the FTSE 100 Stock index, the MAPE values range between 59.40% and 70.56% for the RGARCH and HAR models. However, the proposed model based on RNN outperforms these

statistical models significantly, achieving MAPE values of 33.06% for the S&P 500 index and 41.36% for the FTSE 100 index.
On the other hand, we observed minimal differences in MAPE between the statistical models and the proposed model for the Hang Seng index, with MAPE ranging from 44% to 44.59% for the RGARCH and HAR models, while the proposed model achieves a slightly lower MAPE of 43.94%. Similarly, for the All Ordinaries index, the MAPE values are 43.15% and 43.82% for the RGARCH and HAR models respectively, while the proposed model achieves a lower MAPE of 43%.
Additionally, we implemented Principal Component Analysis (PCA) to provide optimal visualization of variables and data. In this context, the individuals represent the stock indexes, while the variables reflect the Mean absolute percentage error (MAPE) of different models.
This research highlights the potential of the proposed model as a reliable and efficient forecasting tool and the superior performance of the model in forecasting financial series with anti-persistent behavior, where notable enhancements are observed. Additionally, the model demonstrates modest improvements for series exhibiting persistent behavior. As it outperforms traditional statistical methods, the model's application holds promise for enhancing decision-making and risk management strategies in financial markets.
While our proposed hybrid modeling approach aims to improve volatility forecasting in financial markets, it's important to acknowledge the inherent complexity of factors influencing stock prices. Notably, financial news is recognized as a significant driver of fluctuations in stock prices, underscoring the importance of sentiment analysis in understanding investor reactions to news occurrences. Despite recent advancements in sentiment analysis, including considerations of functional relationships between words in a sentence, our study acknowledges the limitations in fully capturing the nuanced sentiments expressed in financial news. Therefore, while our model offers advancements in volatility forecasting, it's crucial to recognize the ongoing challenges and evolving nature of sentiment analysis in financial markets.

Conflicts of interest

The authors declare that they have no conflicts of interest regarding this research. They have not received any financial support or compensation from any organization. We confirm that this submission represents original work and is not currently under review by any other publication.

Acknowledgement

We would like to express our sincere appreciation to the reviewers for their valuable comments on our article. Their constructive feedback has greatly contributed to the improvement of the manuscript.

REFERENCES

1. A. García-Medina, and E. Aguayo-Moreno, *LSTM–GARCH Hybrid Model for the Prediction of Volatility in Cryptocurrency Portfolios*, Comput Econ 63, 1511–1542,2024.
2. C. Zhang, Y. Zhang, M. Cucuringu, and Z. Qian, *Volatility Forecasting with Machine Learning and Intraday Commonality*, Journal of Financial Econometrics, 2023.
3. D. Kingma, and J. Ba, *Adam: A method for stochastic optimization*, International Conference on Learning Representations, 2014.
4. E. Hajizadeh, A. Seifi, M. Zarandi, and I. Turksen, *A hybrid modeling approach for forecasting the volatility of sp 500 index return*, Expert Systems with Applications, 39, pp. 431–436, 2012.
5. F. Corsi, *A simple long memory model of realized volatility*, journal of financial econometrics, 7, pp. 174–196, 2009.
6. G. Bekaert, and G. Wu, *Asymmetric Volatility and Risk in Equity Markets*, The Review of Financial Studies, vol. 13, no. 1, pp. 1–42.
7. H. Y. Young, and C. H. Won, *Forecasting the volatility of stock price index: A hybrid model integrating lstm with multiple garch-type models*, Expert Systems with Applications, 103, pp. 25–35, 2018.
8. J. Maheu, and T. McCurdy, *Nonlinear features of realized fx volatility*, The Review of Economics and Statistics, 84, pp. 668–681, 2002.
9. X. Lu, D. Que, G. Cao, *Volatility Forecast Based on the Hybrid Artificial Neural Network and GARCH-type Models*, Procedia Computer Science, 2016.

10. L.D. Persio, and O. Honchar. *Recurrent neural networks approach to the financial forecast of google assets*, 2017.
11. L. Rubio, A.P. Pinedo, A. Castaño, and F. Ramos, *Forecasting volatility by using wavelet transform, ARIMA and GARCH models*, Eurasian Econ Rev, vol. 13, pp. 803–830, 2023.
12. P.R. HANSEN, Z. Huang, and H. H. SHEK, *Realized Garch: A Joint Model for Returns and Realized Measures of volatility*, vol. 27, no. 6, pp. 877–906, 2012.
13. S. Hochreiter, and J. Schmidhuber, *Long short-term memory*, Neural computation, 9, pp. 1735–80, 1997.
14. T. Bollerslev, *Generalized autoregressive conditional heteroskedasticity*, Journal of Econometrics, vol. 31, no. 3, pp. 307–327, 1986.
15. Y. Bakkali, M. El Merzguioui, A. Akharif, and A. Azmani, *Forecasting stock return volatility using the Realized GARCH model and an artificial neural network*, Bulletin of the South Ural State University. Ser. Mathematical Modelling, Programming & Computer Software (Bulletin SUSU MMCS), 16(4), 45-60, 2023.
16. Z. Huang, H. Liu, and T. Wang, *Modeling long memory volatility using realized measures of volatility: A realized har garch model*, Economic Modelling, 52, pp. 812–821, 2016.