forecasting

Article

Performance Evaluation of Advanced RNNs for Accurate Prediction of Adjusted Closing Gold Prices

Thabang Molefi, Tshegofatso Botlhoko * and Tlhalitshi Volition Montshiwa

Department of Business Statistics & Operations Research, North West University, Mmabatho 2745, South Africa; 44085427@mynwu.ac.za (T.M.); volition.montshiwa@nwu.ac.za (T.V.M.)
* Correspondence: tshegofatso.botlhoko@nwu.ac.za

Highlights

What are the main findings?

- The GRU algorithm consistently outperformed LSTM across the overall sample and across the pre-COVID-19 and post-COVID-19 regimes when forecasting South African adjusted closing gold prices.
- During the COVID-19 regime, LSTM marginally outperformed the GRU, indicating that model performance varies under periods of extreme market volatility and structural disruption.

What are the implications of the main findings?

- The superior overall performance of the GRU suggests that it is a more efficient and reliable RNN architecture for modelling long-term gold price dynamics in the South African context, particularly under stable and post-shock conditions.
- The regime-dependent performance differences imply that adaptive or hybrid forecasting frameworks, potentially combining GRUs, LSTM, and traditional econometric models, may be more effective in capturing market behaviour during periods of heightened uncertainty, such as global crises.

Abstract

This study aimed to compare RNN algorithms and select the best-performing one between the GRU and LSTM for forecasting South African adjusted closing gold prices. The study used weekly secondary data sourced from Yahoo Finance and partitioned into three regimes, pre-COVID-19, COVID-19, and post-COVID-19, as well as the overall sample. The results indicated that the GRU algorithm consistently outperformed the LSTM algorithm across all evaluation periods based on the selected metrics, except during the COVID-19 period, where LSTM exhibited slightly better performance. Consequently, the GRU algorithm was identified as the best-performing algorithm for the South African adjusted closing gold price series. The relative effectiveness of GRU and LSTM algorithms in financial time series forecasting was clarified by the results. By integrating GRU-based forecasts into development finance frameworks, stakeholders can strengthen resilience against global shocks, improve financial planning, and foster more stable pathways for economic development. The authors recommended that future studies explore the performance of the GRU and LSTM with other advanced algorithms like Transformer architectures, hybrid algorithms, or traditional statistical methods to further enhance the forecasting robustness.

Academic Editors: Sonia Leva and João Paulo Ramos Teixeira

Received: 20 January 2026
Revised: 6 March 2026
Accepted: 22 March 2026
Published: 18 April 2026

Keywords: recurrent neural networks (RNNs); long short-term memory (LSTM); gated recurrent unit (GRU); adjusted closing gold price; forecasting; deep learning

1. Introduction

Gold is a critical global commodity in South Africa, valued as a hedge against inflation and economic uncertainty. The weekly adjusted closing gold price data exhibit temporal dependencies, trends, and potential volatility, which are known characteristics inherent to time series data that can make reliable forecasting difficult. Despite the widespread use of gated recurrent unit (GRU) and long short-term memory (LSTM) models, limited evidence exists of their comparative performance across distinct market regimes within the South African gold market context. Accordingly, this study conducts a comparison analysis of deep learning (DL) algorithms, particularly the GRU and LSTM algorithms, using sequential data for time series forecasting. It evaluates the two algorithms based on performance metrics such as the root mean squared error (RMSE), mean absolute percentage error (MAPE), coefficient of determination (R^2), mean squared error (MSE), and mean absolute error (MAE) tailored to financial time series. Additionally, a forecast graph is produced to compare the actual values with the forecasts generated by the best-performing model, the GRU.

The effectiveness of recurrent neural network designs for gold price forecasting has been compared in recent research. For example, the study [1] compared the performance of LSTM and GRUs in forecasting gold prices using historical Yahoo Finance data. The study focused on overall forecasting performance, without explicitly examining how model behaviour varies across different market regimes or economic conditions. While it provides useful insights into their comparative accuracy, limited attention is given to contextual factors such as regime shifts or the specific dynamics of the South African gold market. By comparing the performance of the GRU and LSTM models in the South African gold market and across several market regimes, the current study, on the other hand, expands on this line of inquiry and offers a contextual and regime-aware viewpoint that goes beyond comparisons of aggregate accuracy.

Additionally, the study [2] proposes a hybrid framework that combines LSTM networks with autoencoder structures to enhance feature extraction and predictive performance. While such hybrid models demonstrate improved accuracy through architectural complexity, they primarily emphasise performance enhancement rather than comparative analyses between standard recurrent architectures. In contrast, the current study focuses on a structured comparison between two widely used recurrent models, the GRU and LSTM, within the South African gold market context. Rather than introducing additional architectural layers, this study evaluates the robustness and forecasting performance of these foundational RNN models across distinct market regimes.

The main objective of this study is to develop GRU and LSTM models for weekly South African adjusted closing gold price forecasting and assess their robustness and comparative performance across various market regimes using common financial accuracy criteria. In doing so, the study contributes to the regime-aware application of advanced deep learning models to South African gold price data, addressing a gap in the use of such methods within South African financial statistics. The results indicated that the GRU model demonstrated superior performance compared to the LSTM model across all evaluation periods based on the selected metrics, except during the COVID-19 period, where LSTM exhibited slightly better performance, and the forecast plot for the GRU further showed that the GRU predictions closely followed the actual gold price movements. By analysing the adjusted closing gold price in a market of significant national economic importance, the study deepens the understanding of gold price temporal dynamics while providing practical forecasting insights relevant to investors, policymakers, and market analysts.

The study is structured into the following sections: Section 2 covers the related works on gold price forecasting using machine learning; Section 3 covers the materials and methods; Section 4 presents the results; Section 5 presents the discussion; Section 6 covers the limitations and implications for further research; and, lastly, Section 7 concludes the article.

2. Related Works on Gold Price Forecasting Using Machine Learning, Deep Learning, and Statistical Models

In this section, the latest studies on the evaluation of LSTM and GRU algorithms in predicting stock price time series data are reviewed, with a focus on establishing the dominant model and trends in the variables to which these algorithms are applied and identifying the gap(s) that can be addressed by the current study.

In a study focused on stock price prediction, Xiao, Deng [3] conducted a comparative analysis of LSTM, GRU, and Transformer algorithms for stock price prediction using the Tesla stock data from 2015 to 2024, comprising a total of 2274 observations. The objective of the study was to evaluate the accuracy of these algorithms in predicting stock price trends and to provide investors with better tools for making informed decisions in the stock market by leveraging advanced deep learning techniques for stock price prediction. Comparison metrics such as the R-squared, MAE, MSE, and RMSE were used, and the results of the study revealed that the LSTM algorithm achieved the highest accuracy of 94% in predicting Tesla's stock price trends compared to the GRU and Transformer algorithms. The LSTM algorithm was also found to be the most suitable for stock time series analysis due to its strong ability to capture long-term dependencies and its adaptability to different market conditions. On the other hand, the Transformer algorithm had the worst prediction performance among the three algorithms tested. Similarly, the current study compares the GRU and LSTM using comparison metrics such as the R-squared, MAE, MSE, RMSE, and MAPE.

In another study, Wang [4] explored the application of RNN, LSTM, and GRU algorithms in stock market forecasting and comparatively analysed the performance of these algorithms to determine the one that could best capture the complex stock market fluctuations. The authors used stock trading data from 28 June 2010 to 10 November 2017 for Tesla, Inc. The authors used the MSE as the comparison metric, and the results of the study revealed that the GRU algorithm exhibited superior performance compared to the LSTM and RNN algorithms in long-term stock price prediction. Moreover, the LSTM algorithm showed significant performance degradation after 500 training cycles, revealing potential shortcomings in handling long-term training tasks. Unlike the study by Wang [4], the current study compares the GRU and LSTM using multiple comparison metrics at different COVID periods (pre-COVID-19, during COVID-19, post COVID-19 and the whole sample).

The study by Sudiatmika, Putra [5] sought to determine the best GRU for forecasting gold prices. Using a dataset from Yahoo Finance for the period of 14 December 2017 to 14 March 2024, the authors searched for the optimum GRU network structure by focusing on the batch size and epochs, and the MSE and MAE were used to evaluate the differently structured GRU algorithms. The authors concluded that the GRU with a batch size of 64 and 100 epochs was the most accurate, since it yielded the lowest values for the model evaluation metrics. The current study aims to further explore the accuracy of the GRU in forecasting gold prices by applying it to a different dataset, comparing its performance to that of LSTM, and evaluating the algorithms with more than two model evaluation metrics.

In their study, Aljohani, Alshamrani [6] examined the effectiveness of ARIMA, LSTM, GRU and simple linear regression (SLR) models in forecasting monthly gold prices in Indian Rupees. The authors used monthly historical gold prices for the period of January 2001 to December 2019, sourced from the Kaggle website. They evaluated the algorithms using the RMSE and MAPE. The results of this study showed that LSTM had the lowest RMSE but the highest MAPE; SLR had the lowest MAPE but the highest RMSE; the GRU obtained the second-lowest values for both the RMSE and MAPE; and ARIMA had the third-lowest values for both the RMSE and MAPE. Therefore, in general, the GRU had the lowest error margin based on both evaluation metrics, followed by SLR and LSTM, whereas, based on both the RMSE and MAPE, the ARIMA model performed relatively poorer. The GRU and LSTM, as DL algorithms, have been proven to be very competitive in forecasting the gold prices in Indian Rupees over traditional models (ARIMA and SLR); hence, they are considered in the current study and are further evaluated in a different context by comparing their effectiveness in forecasting the South African adjusted closing gold price.

In another study, Baradaran, Bohlouli [7] evaluated the basic GRU and a hybrid model combining a convolutional neural network (CNN) and LSTM to form CNN-LSTM. The authors used the daily prices of 18-carat gold in Iran for the period of April 2013 to January 2024, which were sourced from the Rahvard365 website. The two algorithms were used to predict the closing gold price of the succeeding day based on the price data of the past 150 days. The algorithms were compared using the MAE, MAPE, and RMSE, and the results showed that the hybrid CNN-LSTM was more accurate than the basic GRU. These results suggest that more studies are needed on hybrid algorithms to improve the accuracy of standalone DL algorithms. Therefore, the current study further extends the scope of the study by Baradaran, Bohlouli [7] by using a different dataset, namely the South African adjusted closing gold prices.

The research conducted by Dewi, Ginantra [8] compared the performance of GRU and LSTM algorithms using different batch sizes and epochs in predicting the price of gold, and the algorithms were compared using the RMSE. By comparing the actual and predicted gold prices in a plot, the authors found that both LSTM with a batch size of 32 and epoch number of 50 and the GRU with a batch size of 8 and epoch number of 50 were accurate in predicting the gold price. However, the GRU yielded a lower RMSE than LSTM, implying that it was the best model for forecasting gold prices. The study by Dewi, Ginantra [8] used data from Yahoo Finance that were collected from January 2010 to June 2023. The current study uses data from the same source and for a similar period, except that it ends on 31 January 2024 and the data used are for South Africa. Moreover, the current study examines different periods (pre-COVID-19, during COVID-19, post-COVID-19, and the whole sample), differing from the study by Dewi, Ginantra [8], which only examined the whole period.

In a separate study, Pirani, Thakkar [9] conducted a comparative analysis of ARIMA, GRU, LSTM, and bidirectional long short-term memory (BiLSTM) algorithms in forecasting financial time series data using stock data from Yahoo Finance—specifically the closing prices of IBM stocks from July 2009 to 2019. The aim of the study was to determine if training data in the opposite track—that is, right to left—in addition to traditional data training—that is, left to right—could improve the forecasting accuracy and to assess algorithm performance with fewer gates than LSTM-based algorithms. The study used the RMSE to compare the algorithms. Pirani, Thakkar [9] reported that using an additional layer of training (BiLSTM) improved the forecast accuracy by 11.86% on average compared to the unidirectional LSTM algorithm. Moreover, the GRU algorithm, with one fewer gate than the LSTM algorithm, improved the accuracy by 20.32% over the BiLSTM algorithm. The

https://doi.org/10.3390/forecast8020033

GRU algorithm was quicker to train than both the LSTM and BiLSTM algorithms, despite the BiLSTM algorithm having higher forecast accuracy. Generally, the GRU algorithm outperformed the other algorithms in terms of accuracy and efficiency. The current study expands the scope of the study by Pirani, Thakkar [9] by comparing the GRU and LSTM using comparison metrics such as the R-squared, MAE, MSE, RMSE, and MAPE during different COVID-19 periods (pre-COVID-19, during COVID-19, post-COVID-19, and the whole sample).

In their study, Primananda and Isa [10] implemented LSTM and the GRU to forecast gold prices in Rupiah for a multivariate case. The daily dataset was for the period of 1 January 2001 to 31 December 2020 and was sourced from the World Gold Council (gold.org) and investing.com websites. The eight features consisted of gold prices in other countries (in the currencies of those countries) and were used to predict the gold price in Rupiah. The forecasting performance of the algorithms was evaluated for different time intervals ranging from three months to 20 years. The algorithms were evaluated using the RMSE and MAPE, and the results showed that the GRU was the best model to forecast gold prices when the time interval was under three years, whereas, when the time interval was above three years, LSTM exhibited higher accuracy. Unlike the study by Primananda and Isa [10], the current study uses univariate data but it also evaluates the algorithms for different intervals—pre-COVID-19, during COVID-19, post-COVID-19 and the whole sample—focusing on the adjusted closing gold prices in South African Rands.

The research conducted by Gao, Wang [11] focused on stock prediction based on the optimised LSTM and GRU algorithms. The authors aimed to develop a new algorithm for stock price forecasting that incorporates technical indicators and financial data, performed dimension reduction on the input features using the Least Absolute Shrinkage and Selection Operator (LASSO) and principal component analysis (PCA), and compared the performance of the LSTM and GRU algorithms using the LASSO and PCA reduced feature sets to determine the optimal forecasting approach. The dataset used in the study was based on the Shanghai Composite Index and included a variety of technical indicators and investor sentiment indicators. The comparison metrics used were the MSE, RMSE, and MAE. The results of the study revealed that both the LSTM and GRU algorithms could effectively predict stock prices, with none of the algorithms being more efficient than the others. Additionally, the prediction results of the two neural networks using LASSO dimension reduction (LASSO-LSTM and LASSO-GRU) were mostly better than those using PCA dimension reduction data (PCA-LSTM and PCA-GRU) in terms of prediction accuracy. This study evaluates the performance of GRU and LSTM models using metrics such as the R-squared, MAE, MSE, RMSE, and MAPE across different COVID-19 periods (pre-COVID-19, during COVID-19, post-COVID-19, and the entire dataset). The dataset, consisting of adjusted closing gold prices, is sourced from Yahoo Finance.

The literature shows that DL algorithms, especially LSTM and GRU algorithms, have been applied to various financial time series data, and their performance varies according to the data used. However, the researchers are not aware of any previous study that has compared these algorithms on the adjusted closing gold prices; this is a gap in the literature that the current study intends to fill.

3. Materials and Methods

This section describes the dataset, preprocessing steps, and methodological framework employed in the study.

3.1. Data Preparation and Preprocessing

3.1.1. Data

This study used weekly secondary data that are publicly available on the Yahoo Finance website (https://finance.yahoo.com/quote/GC%3DF/, accessed on 28 October 2025). The data range from 28 August 2000 to 27 October 2025. The variable of interest is the South African adjusted closing gold price. Both statistical qualities and economic relevance motivate the focus on gold prices in the South African context. Due to the sensitivity of gold prices to local economic conditions, global shocks, and structural changes, the South African gold market offers a useful environment for evaluating how GRU and LSTM models react. Since gold is traded globally in US dollars (USD) and is closely correlated with changes in the South African Rand (ZAR), local prices are impacted by both global variations and exchange rate volatility. Although production has decreased, the mining industry is still commercially significant and continues to affect gold prices. In the past, South Africa was the world's greatest producer of gold. Financial crises, mining disruptions, labour strikes, production expenses, and problems with energy supply are only a few examples of the domestic and foreign structural factors that the market is vulnerable to. These events can cause structural breaks and regime shifts in the price series. It is more volatile and sensitive to foreign capital movements because it is a developing economy. The South African gold market is an ideal case in which to assess how well GRU and LSTM models capture regime shifts and produce reliable forecasts in a complex financial environment, since these interconnected factors produce intricate nonlinear patterns and temporal dependencies.

3.1.2. Data Split

The data used in the study were partitioned into three periods—pre-COVID-19, COVID-19, and post-COVID-19, as well as considering the overall sample—to perform an in-depth empirical analysis of the performance of the GRU and LSTM in predicting the adjusted closing gold price. By investigating whether the algorithms' performance was impacted throughout these periods, this study seeks to assess how resilient the algorithms are to significant market shocks. Specifically, the whole sample includes these shocks, but, once the data are split into pre-COVID-19, COVID-19, and post-COVID-19, the shocks between the periods are not included. Each sample was split into two subsets, with 70% for training and 30% for testing.

3.1.3. Creation of Sequences

A sliding window approach was employed to convert the univariate series into a supervised learning problem, using sequences of seven consecutive observations as inputs and the subsequent value as the prediction target.

3.1.4. Data Normalisation

The weekly gold price series was normalised using min–max scaling to the range [0, 1] to enhance the numerical stability.

3.1.5. Model Architecture

The GRU model included a fully connected output layer for one-step-ahead forecasting, two stacked GRU layers with 50 hidden units each, and dropout regularisation (rate = 0.2) to reduce overfitting. Performance was assessed on the testing set after the model was trained over 10 epochs with a batch size of 32 using the Adam optimiser, with the mean squared error as the loss function. Meanwhile, a deep recurrent architecture with two stacked LSTM layers, each with 50 memory units, was used to implement the LSTM model. In order to facilitate hierarchical temporal feature extraction, the first LSTM layer was set up to return complete sequences, whereas the second layer generated a single output vector. To lessen overfitting, dropout regularisation was conducted after each LSTM layer at a rate of 0.2. Weekly gold price forecasts were produced one step ahead of time using a fully connected dense layer with a single neuron. Using the same training and validation methodology as for the GRU model for consistency, the model was trained using the Adam optimiser, with the mean squared error as the loss function.

3.1.6. Software

Data analysis was conducted using the Python software, version 2022.3.3. The methodology followed in this study is outlined in Figure 1, which presents the flow chart for the analysis.

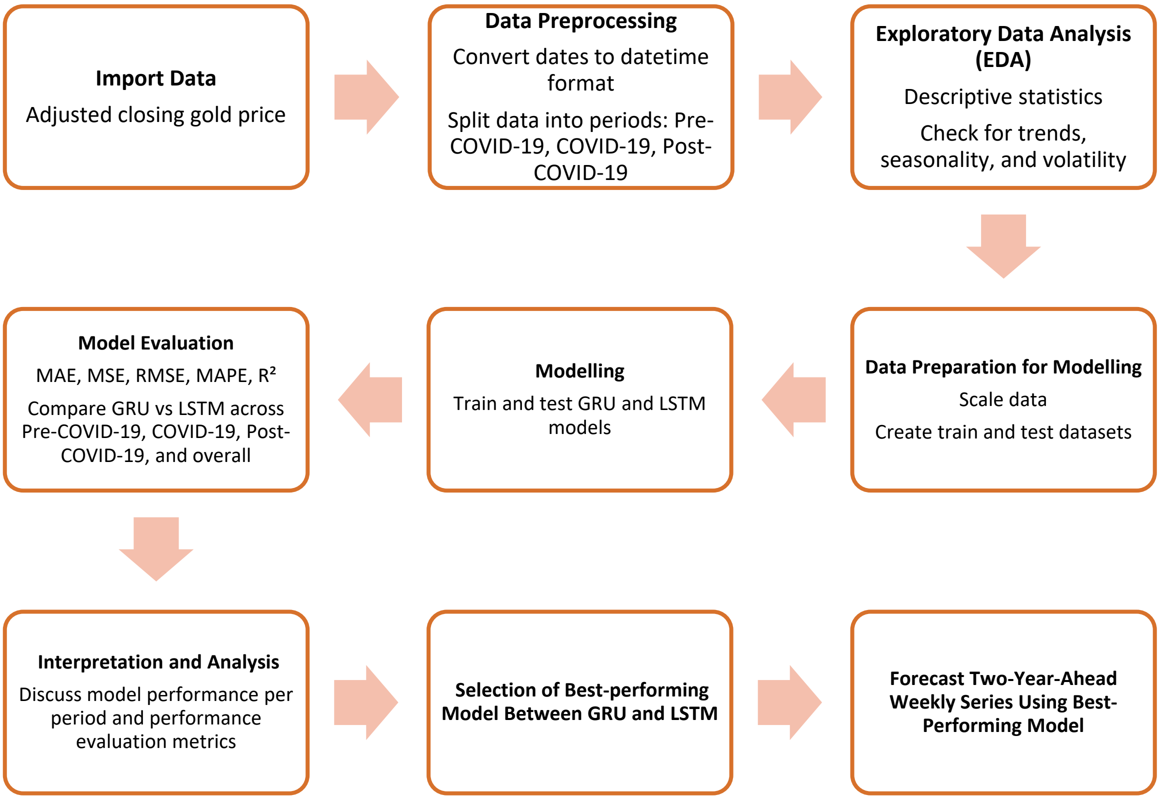

Figure 1. Flow chart for the analysis.

According to Figure 1, the process begins with the importing the weekly adjusted closing gold prices into the Python software. The data undergo preprocessing to convert the date formats and segment the series into pre-COVID-19, during COVID-19, and post-COVID-19 periods. EDA is then conducted to summarise key statistical properties and identify underlying trends, seasonality, and volatility. Subsequently, the data are scaled and structured into suitable input sequences for model training. Two DL algorithms, the GRU and LSTM, are implemented to forecast the adjusted closing gold prices. The models are trained on historical data and tested on out-of-sample observations to simulate an ex ante forecasting setting, ensuring that predictions reflect real-world conditions. Model behaviour is evaluated using standard statistical metrics such as the MSE, RMSE, MAPE, and R^2. Finally, the results are interpreted across the regimes to evaluate the robustness and adaptability of the models. The details of the algorithms are discussed in the following subsections.

3.2. GRU

The GRU is an algorithm developed by Kyunghyun Cho and colleagues in 2014 that is considered to be a cornerstone in sequence algorithm tasks that balances simplicity and performance. The algorithm is popular in the deep learning community, especially for tasks involving sequential or time series data. The algorithm is formed by four main components, namely the rest gate, the update gate, the candidate hidden state, and the current hidden state.

The reset gate (r_t) component of the GRU controls how much of the past state to forget when computing the candidate hidden state. The reset gate is expressed mathematically using the following formula, which is adopted from Mateus, Mendes [12]:

$$r_t = \sigma(x_t W^r + h_{t-1} U^r + b_r) \tag{1}$$

where r_t is the reset gate, σ represents the sigmoid activation function, W^r and U^r denote the reset gate weight matrices, x_t is the input vector at time t, h_{t-1} is the hidden state of the previous time step, and b_r is the bias vector for the reset gate. If r_t is close to 0, the network ignores most of the previous hidden state, effectively resetting its memory for that step. If r_t is close to 1, more of the past information is retained.

Secondly, the update gate (z_t), which determines how much of the past state to retain and how much of the new candidate state to use, is mathematically expressed using the following formula, which is also adopted from Mateus, Mendes [12]:

$$z_t = \sigma(x_t W^z + h_{t-1} U^z + b_z), \tag{2}$$

where W^z and U^z denote the update gate weight matrices, and b_z is the bias vector for the update gate. If z_t is close to 1, the GRU mostly retains the previous hidden state (h_{t-1}), minimising updates from the new candidate state. If z_t is close to 0, the GRU heavily relies on the new candidate state $(\widetilde{h_t})$ and updates its memory more significantly.

Thirdly, the candidate hidden state $(\widetilde{h_t})$, which is the new information computed based on the reset gate's output, is computed mathematically using the following formula, which is again adopted from Mateus, Mendes [12]:

$$\widetilde{h_t} = \tanh\,(r_t \times h_{t-1} U + x_t W + b), \tag{3}$$

where *tanh* is the hyperbolic tangent activation function, and U and W denote the reset gate weight matrices. The candidate hidden state is calculated using the reset gate (r_t), which controls how much of the previous hidden state (h_{t-1}) influences the candidate computation. Moreover, by resetting irrelevant parts of the past, it focuses only on the most useful information for the current step.

Lastly, the current hidden state (h_t), which is a combination of the previous hidden state and the candidate hidden state, weighted by the update gate, is mathematically expressed as follows, according to Mateus, Mendes [12]:

$$h_t = (1 - z_t) \times \widetilde{h_t} + z_t \times h_{t-1}, \tag{4}$$

The current hidden state is effective in ensuring a smooth transition between past and new information, avoiding abrupt changes in memory, and balancing long-term dependencies h_{t-1} and short-term relevance $(\widetilde{h_t})$. The architecture of the GRU is illustrated in Figure 2.

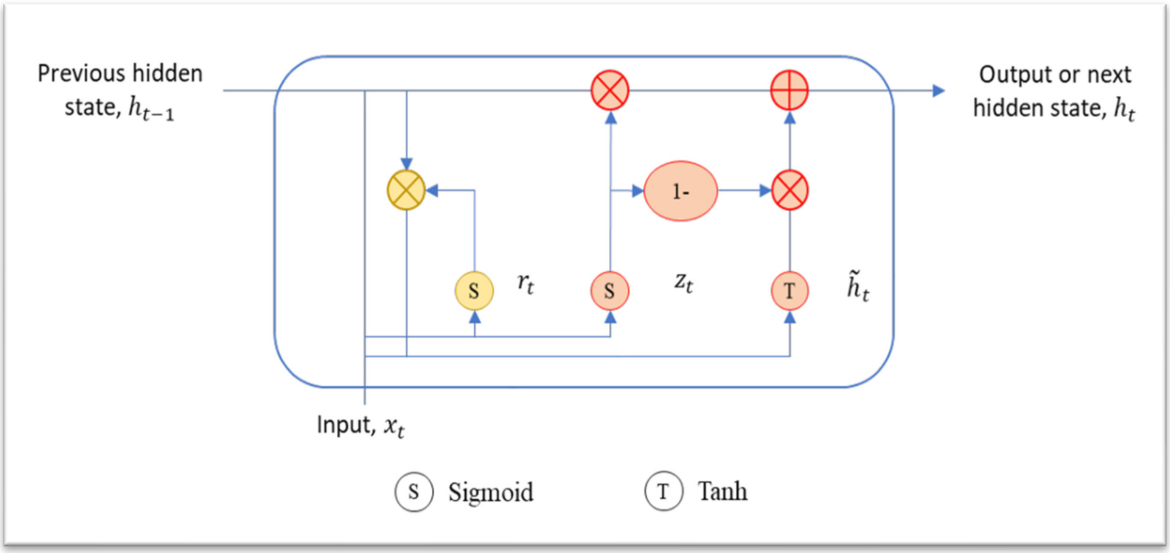

Figure 2. GRU architecture. Source: Ng, Lim [13].

3.3. LSTM

LSTM is a popular and widely used DL technique, introduced by Sepp Hochreiter and Jürgen Schmidhuber in 1997. It is known for its ability to effectively learn both short-term and long-term dependencies in sequences by mitigating the issues of standard RNNs, such as vanishing and exploding gradients and selective memory [14]. This technique consists of two key features: the memory cell and the gates (the forget gate, input gate, and output gate).

The forget gate is an important component of the LSTM architecture that is responsible for deciding which information to remove from the cell state. It filters out irrelevant or outdated information to ensure that only useful data are retained for future time steps. The forget gate is mathematically represented using the following formula, which is adopted from Mateus, Mendes [12]:

$$f_t = \sigma\left(x_t W_f + h_{t-1} U_f + b_f\right), \tag{5}$$

where f_t denotes forget gate activation; σ is the sigmoid activation function, which generates numbers between 0 and 1; W_f is the weight of the matrix; x_t is the input vector at time t; h_{t-1} is the hidden state from the previous time step; U_f is the weight matrix for the previous hidden state h_{t-1}; and b_f is the bias term of the forget gate. If the forget gate outputs a value close to 1 for a specific part of the cell state, the information is retained. If the output is close to 0, the corresponding information is discarded. The forget gate also determines which parts of the cell state should be forgotten.

Meanwhile, the input gate controls how much new information from the current input (x_t) and the previous hidden state (h_{t-1}) is added to the cell state. It determines which parts of the new information are relevant enough to update the memory. The input gate consists of two main components: the input gate vector and the candidate memory. The input gate vector (i_t) decides which parts of the new candidate memory ($\widetilde{C}_t$) should be added to the cell state. The input gate is mathematically calculated using the following formula, which is again adopted from Mateus, Mendes [12]:

$$i_t = \sigma(x_t W_i + h_{t-1} U_i + b_i), \tag{6}$$

where σ is the sigmoid activation function, which generates numbers between 0 and 1; W_i and U_i are the weight matrices for the input and previous hidden states, respectively; x_t is the input vector at time t; h_{t-1} is the hidden state from the previous time step; and b_i is the bias term of the forget gate.

The second component is $\widetilde{C}_t$, which proposes new information to be added to the cell state. Mathematically, it is presented as follows, according to Mateus, Mendes [12]:

$$\widetilde{C}_t = tan(x_t W_C + h_{t-1} U_c + b_C), \tag{7}$$

where tan is the tangent activation function, which generates numbers between -1 and 1; W_C and U_c are the weight matrices for the input and previous hidden state; and b_C is the bias term. The input gate further updates the cell state as follows, according to Mateus, Mendes [12]:

$$\mathrm{C}_t = \sigma\,(f_t \times \mathrm{C}_{t-1} + i_t \times \widetilde{c}_t), \tag{8}$$

If i_t is close to 1, more of $\widetilde{C}_t$ is added to the cell state, but, if i_t is close to 0, little or none of $\widetilde{C}_t$ is added to the cell state.

The primary role of the output gate is to determine how much of the updated cell state contributes to the current hidden state. The hidden state is then used to generate the output of the LSTM model at the current step and is passed to the next step. The output gate uses the input and the previous hidden state to decide how much of the cell state to reveal. The output gate is mathematically presented as follows, as adopted from Mateus, Mendes [12]:

$$o_t = \sigma(x_t W_o + h_{t-1} U_o + b_o), \tag{9}$$

where $\mathrm{o_t}$ is the output gate vector.

The output gate is applied to the updated cell state to restrict its values to between -1 and 1 for suitability of usage in the hidden state. The hidden state is calculated as follows, as adopted from Mateus, Mendes [12]:

$$\mathrm{h_t} = \tanh(\mathrm{C_t})\ \mathrm{o_t}, \tag{10}$$

where $\tanh(\mathrm{C_t})$ is the activated cell state. If $\mathrm{o_t}$ is close to 1, more information from $\tanh(\mathrm{C_t})$ is included in $\mathrm{h_t}$, but, if $\mathrm{o_t}$ is close to 0, the contribution of $\tanh(\mathrm{C_t})$ to $\mathrm{h_t}$ is minimal. The architecture of LSTM is illustrated in Figure 3.

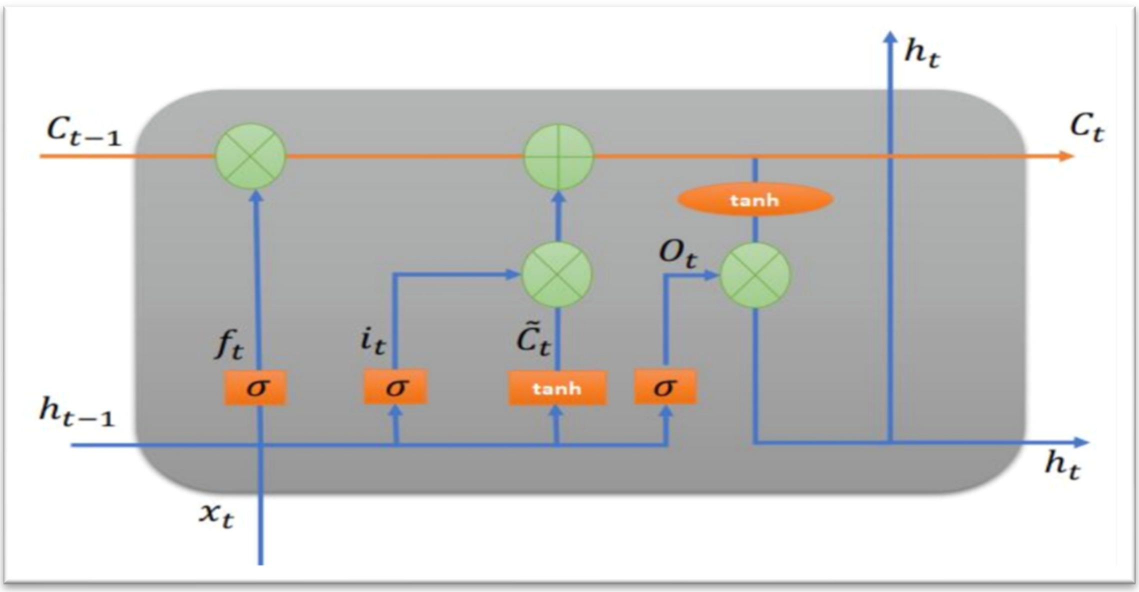

Figure 3. LSTM architecture. Source: Ng, Lim [13].

3.4. Model Evaluation Metrics

In this study, several model evaluation metrics are used to compare the performance of the two algorithms. The metrics used are the MAE, MSE, RMSE, MAPE, and R^2. Each metric provides unique information about the correctness and dependability of the model. Although its squared units make direct interpretation less straightforward and it is sensitive to outliers, the MSE is useful in penalising greater prediction errors. While both the MAE and RMSE assess prediction errors in the same units as the data, the RMSE emphasises larger deviations, while the R^2 evaluates the percentage of variance in the actual data that the model can account for, indicating the quality of fit as opposed to the direct error magnitude. Although it can be unstable as real values approach zero, the MAPE, on the other hand, expresses forecast errors as percentages, making it easier to evaluate models. By combining various metrics, forecasting behaviour can be evaluated more fairly and meaningfully, avoiding biases that could result from depending only on one indicator, such as the MSE. These metrics are computed using the following equations, adopted from Ningsih, Sarwido [15]:

$$\mathrm{MAE} = \frac{1}{\mathrm{n}} \times \sum_{\mathrm{i}=1}^{\mathrm{n}} |\mathrm{y_i} - \hat{\mathrm{y}}_\mathrm{i}|^2 \tag{11}$$

$$\mathrm{MSE} = \frac{1}{\mathrm{n}} \sum_{\mathrm{i}=1}^{\mathrm{n}} (\mathrm{y_i} - \hat{\mathrm{y}}_\mathrm{i}) \tag{12}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{\mathrm{n}} \sum_{\mathrm{i}=1}^{\mathrm{n}} (\mathrm{y_i} - \hat{\mathrm{y}}_\mathrm{i})} \tag{13}$$

$$\mathrm{MAPE} = \frac{1}{\mathrm{n}} \sum \frac{|\mathrm{actual} - \mathrm{forecast}|}{|\mathrm{actual}|} \times 100 \tag{14}$$

$$\mathrm{R}^2 = 1 - \frac{\mathrm{SS\ error}}{\mathrm{SS\ total}} = 1 - \frac{\sum_{\mathrm{i}=1}^{\mathrm{n}} (\mathrm{y_i} - \hat{\mathrm{y}}_\mathrm{i})^2}{\sum_{\mathrm{i}=1}^{\mathrm{n}} (\mathrm{y_i} - \overline{\mathrm{y}})^2}, \tag{15}$$

where n is the number of observations, $\mathrm{y_i}$ denotes the actual adjusted closing gold price, and $\hat{\mathrm{y}}_\mathrm{i}$ denotes the predicted closing price.

4. Results

This section provides an analysis of the results of the study. The results are presented for four samples: pre-COVID-19, during COVID-19, post-COVID-19, and the overall sample. To assess the behaviour of the algorithms, the MAE, MSE, RMSE, MAPE, and R^2 were computed, and the results are summarised in Table 1.

Table 1. Descriptive statistics of the weekly adjusted closing gold prices per regime.

Regime	Mean	Min	25%	50%	75%	Max	Standard Deviation
Pre-COVID-19	957.512	257.900	439.600	1109.100	1300.050	1873.700	456.080
COVID-19	1787.122	1524.900	1730.600	1793.500	1853.600	2001.200	103.652
Post-COVID-19	2321.401	1639.600	1884.875	2015.650	2656.250	4189.900	604.954
Overall sample	1230.741	257.900	644.550	1245.400	1665.500	4189.900	692.384243

Table 1 presents the descriptive statistics of the weekly adjusted closing gold prices. Across the different regimes, the descriptive statistics reveal noticeable structural variations in the adjusted gold prices. During the pre-COVID-19 period, the gold prices show a moderate average level, with a mean of 957.51, reflecting relatively stable market conditions;

however, the large standard deviation of 456.08 and the wide range between the minimum (257.90) and maximum (1873.70) indicate noticeable price fluctuations even before the pandemic. The median (1109.10) exceeding the mean suggests a left-skewed distribution, where occasional lower prices pull the mean downwards, pointing to cyclical movements driven by traditional macroeconomic and financial factors. In contrast, the COVID-19 regime is characterised by a sharp upward shift in gold prices, with the mean increasing to 1787.12, highlighting gold's role as a safe-haven asset during heightened uncertainty. This period exhibits remarkably low volatility, as evidenced by the lowest standard deviation (103.65), a narrow interquartile range (1730.60 to 1853.60), and close alignment between the mean (1787.12) and median (1793.50), indicating a highly concentrated and stable price distribution supported by sustained investor demand. In the post-COVID-19 regime, the gold prices reach their highest average level, with a mean of 2321.40, reflecting prolonged uncertainty, inflationary pressures, and evolving monetary policies. This period also shows substantial volatility, with a standard deviation of 604.95, alongside extreme price movements, as reflected by the wide range from 1639.60 to 4189.90. The 75th percentile (2656.25) being considerably higher than the median (2015.65) suggests right skewness, indicating occasional price surges that elevate the upper tail of the distribution. When considering the overall sample, the mean gold price is 1230.74, but this aggregate measure conceals significant regime-specific heterogeneity, with the highest overall standard deviation (692.38) capturing the combined effects of pre-pandemic stability, pandemic-induced price consolidation, and post-pandemic volatility, thereby reinforcing the necessity of a regime-based analytical framework rather than a single pooled modelling approach. Figure 4 presents the original time series plot of the adjusted closing gold prices.

Figure 4. Original time series plot.

The time series plot for the adjusted closing gold prices in Figure 4 shows an overall upward trend from 2000 to 2025. For the period of 2000 to 2005, the model shows low volatility. This could be due to the relatively stable global economy, where no major global financial shocks occurred during this period to push investors towards gold as a safe haven. For the period of 2006 to 2012, the model shows a steep increase in the adjusted closing gold price, with peak prices observed at around the year 2011. Market volatility rose at that time due to the global financial crisis, a surge in commodity prices, and increasing

https://doi.org/10.3390/forecast8020033

investor demand for gold as a safe haven. The series is seen to stabilise with moderate fluctuations from 2013 to 2018. The series experienced renewed growth from 2019 to 2021. Global financial concern and extensive economic disruptions were brought about by the COVID-19 pandemic. As a result, gold became a more popular safe-haven asset among investors, increasing the demand and helping to increase the gold prices further. From 2022 to 2025, the adjusted closing gold price experienced a significant surge, where the steepest increase in prices was observed for the entire series. Following 2021, markets stabilised as the COVID-19 pandemic's shocks subsided, leading to more modest increases in the price of gold. Despite a minor decline in volatility relative to the pandemic peak, periods of elevated volatility continued because of inflationary pressures and geopolitical events like the conflict between Russia and Ukraine. Overall, the graph shows an increase in volatility over time, as the early periods show smooth changes, while the later periods show steep fluctuations. This indicates that the variance of the model is not constant; therefore, the model is non-stationary.

The rolling mean and standard deviation were computed to assess the stability of the series over time. As illustrated in Figure 5, both measures vary across the period of 2000 to 2025. The rolling mean shows a fluctuating upward trend, reflecting the changes in the central tendency of the data. Meanwhile, the rolling standard deviation remains relatively stable for most of the period; however, it increases sharply after 2024. These observations confirm that the adjusted closing gold price series is non-stationary, as the mean and variance are not constant over time. The authors acknowledge that non-stationarity may still affect forecasting performance, even though GRU and LSTM models can handle non-stationary time series because of their capacity to capture intricate temporal connections. Predictive accuracy was measured using standard performance criteria, and the models were tested on out-of-sample data to guarantee resilience. This shows that the models' predictions hold up well in a range of market conditions, even though differencing might not be strictly required. Table 2 reveals the results of the compared models.

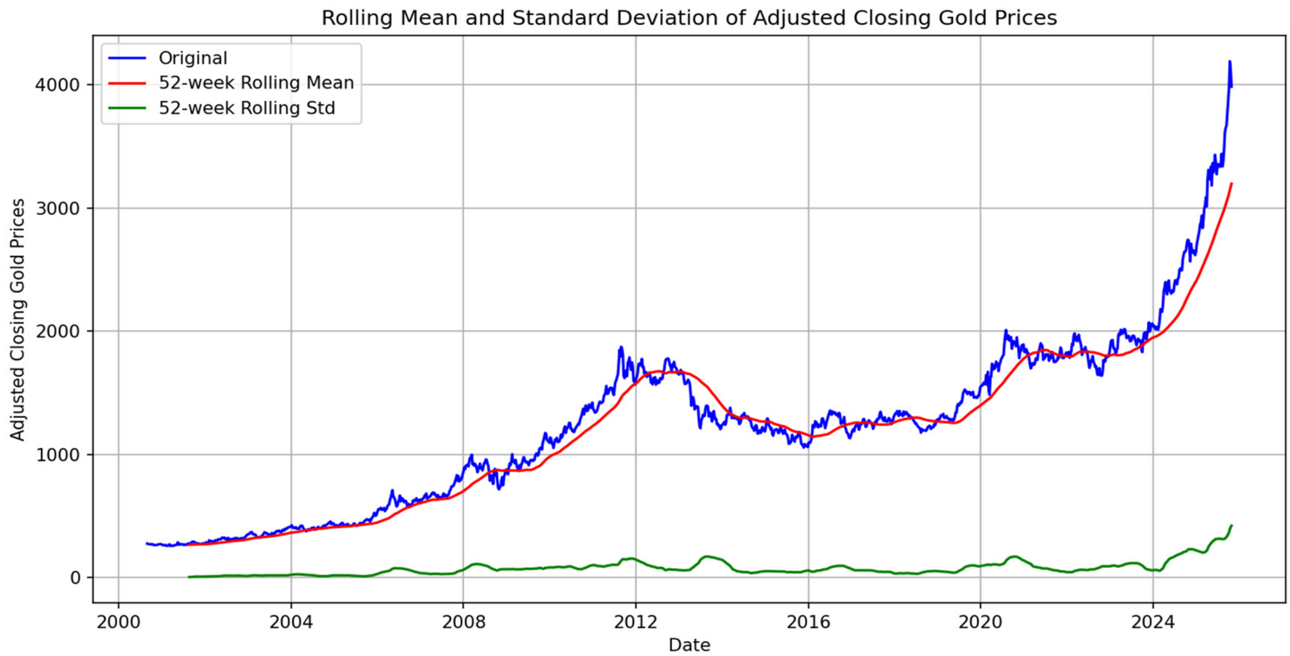

Figure 5. Rolling mean and Standard deviation of the original plot series.

Table 2. Comparison of GRU and LSTM for the testing dataset.

	GRU	LSTM
	Pre-COVID-19	
MAE	0.0165	0.0222
MSE	0.0004	0.0008
RMSE	0.0212	0.0289
MAPE (%)	2.68	3.56
R^2 (%)	87.07	75.99
	COVID-19	
MAE	0.0554	0.0459
MSE	0.0049	0.0044
RMSE	0.0702	0.0661
MAPE (%)	9.93	9.12
R^2 (%)	−55.50	−37.89
	Post-COVID-19	
MAE	00436	0.0909
MSE	0.0030	0.0122
RMSE	0.0551	0.1105
MAPE (%)	7.79	15.46
R^2 (%)	88.93	55.51
	Overall sample	
MAE	0.0168	0.0261
MSE	0.0006	0.0017
RMSE	0.0245	0.0406
MAPE (%)	3.51	5.09
R^2 (%)	97.37	92.80

As shown in Table 2, the performance of the GRU and LSTM models was evaluated across the pre-COVID-19, COVID-19, post-COVID-19, and overall sample periods using the MAE, MSE, RMSE, MAPE, and R^2 as performance metrics. During the pre-COVID-19 period, the GRU model outperformed LSTM across all performance metrics. The GRU achieved an MAE of 0.0165, MSE of 0.0004, RMSE of 0.0212, MAPE of 2.68%, and R^2 of 87.07%, whereas LSTM exhibited higher error values (MAE = 0.0222, MSE = 0.0008, RMSE = 0.0289, MAPE = 3.56%) and a lower R^2 of 75.99%. This indicates that, under steady market conditions, the GRU was more successful in capturing the underlying dynamics of the weekly adjusted closing gold prices because of its gating mechanism, which enables it to effectively model short- to medium-term dependencies without overfitting to noise—a feature that is especially useful when prices evolve gradually and predictably. During COVID-19, the performance of both models dropped due to increased volatility. The GRU showed a higher MAE (0.0554) and RMSE (0.0702) compared to LSTM (MAE = 0.0459, RMSE = 0.0661). However, both models exhibited negative R^2 values (GRU = −55.50%, LSTM = −37.89%), indicating that they both performed poorly relative to the mean of the series. This draws attention to a key drawback of deep learning models: although they can capture intricate temporal patterns under normal market circumstances, they might not be able to do so in times of severe structural upheaval. Unprecedented volatility, sudden changes in investor behaviour, and quickly shifting market dynamics brought about by the pandemic likely exceeded the models' capacity to extrapolate from historical trends. Such performance deterioration aligns with findings in the broader financial forecasting literature, where models often show reduced predictive accuracy during periods of financial crisis or sudden market shocks. LSTM outperformed the GRU. This is because of its more complex memory cell structure, which enables it to capture longer-term dependencies and adjust to sudden changes in market dynamics. It is deemed to performed better during this tumultuous period. Economically speaking, the COVID-19 pandemic led to previously unprecedented swings in gold prices, which were fuelled by investor anxiety, stimulus plans, and uncertainty. These extremely non-linear patterns favoured LSTM's capacity to remember and analyse data over longer time periods. In the post-pandemic period, the GRU again demonstrated superior performance relative to LSTM. In contrast to the peak of the pandemic, the gold price swings became more moderate as the markets stabilised following the severe shocks of COVID-19. While LSTM's ability to handle long-term dependencies was less advantageous in this calmer regime, the GRU's architecture, which effectively captures short- to medium-term dependencies, enabled it to better simulate these more predictable patterns. The GRU achieved an MAE of 0.0436, MSE of 0.0030, RMSE of 0.0551, MAPE of 7.79%, and R^2 of 88.93%, while LSTM's errors were substantially higher (MAE = 0.0909, MSE = 0.0122, RMSE = 0.1105, MAPE = 15.46%, R^2 = 55.51%). Economically, a more stable environment was produced by the reduction in pandemic-related uncertainty and slow market recovery, underscoring the fact that the GRU is especially well suited for forecasting at times of stable market conditions, whereas LSTM tends to underestimate or misrepresent post-crisis fluctuations.

Across the entire sample, the GRU model consistently adapted more effectively than the LSTM model across all evaluation metrices. The GRU recorded a lower MAE value of 0.0168, compared to 0.0261 for the LSTM model, signifying that the GRU's average forecast deviations from the actual observations were smaller. Similarly, the GRU's MSE (0.0006) and RMSE (0.0245) values were lower than those of LSTM (MSE = 0.0017; RMSE = 0.0406), confirming that the GRU produced more precise forecasts with fewer large deviations. Lower error values across these metrics collectively indicate that the GRU achieved a closer fit to the observed data and showed a higher level of reliability in capturing temporal patterns. In terms of percentage-based accuracies, the MAPE for the GRU (3.51%) was considerably lower than that of LSTM (5.09%), indicating that, on average, the GRU forecasts deviated from the actual values by approximately 3.5%, while the LSTM's predictions deviated by more than 5%. The R^2 value of the GRU was 97.37%, compared to 92.80% for LSTM, which further substantiates the findings. The marginally higher R^2 score of the GRU demonstrates its stronger capacity to capture the linear and non-linear dependencies inherent in the dataset. This reinforces the conclusion that the GRU was more effective in minimising relative forecasting errors. Therefore, the GRU was selected as the best-performing model. Figure 6 illustrates the actual values and the forecasts generated by the GRU.

Figure 6 shows the two-year-ahead weekly forecast produced by the GRU model, alongside the historical adjusted gold prices, revealing that the model effectively captures the long-term upward trend and nonlinear dynamics observed in the series. The forecast projects a continued increase in gold prices over the forecast horizon, with values rising from approximately 4100at the start of the forecast period to nearly 5000 by the end, consistent with the strong momentum evident in the most recent post-COVID-19 observations. The smooth evolution of the forecasted path indicates that the GRU model prioritises dominant structural patterns and long-term dependencies while filtering out short-term noise, resulting in stable and economically plausible projections. Although the forecast

https://doi.org/10.3390/forecast8020033

exhibits reduced volatility relative to the historical series, suggesting limited sensitivity to abrupt shocks, this smoothing behaviour enhances its robustness for medium- to long-term forecasting. Overall, the results reinforce the GRU model's suitability as the preferred forecasting framework for the adjusted gold prices within this regime-based analysis.

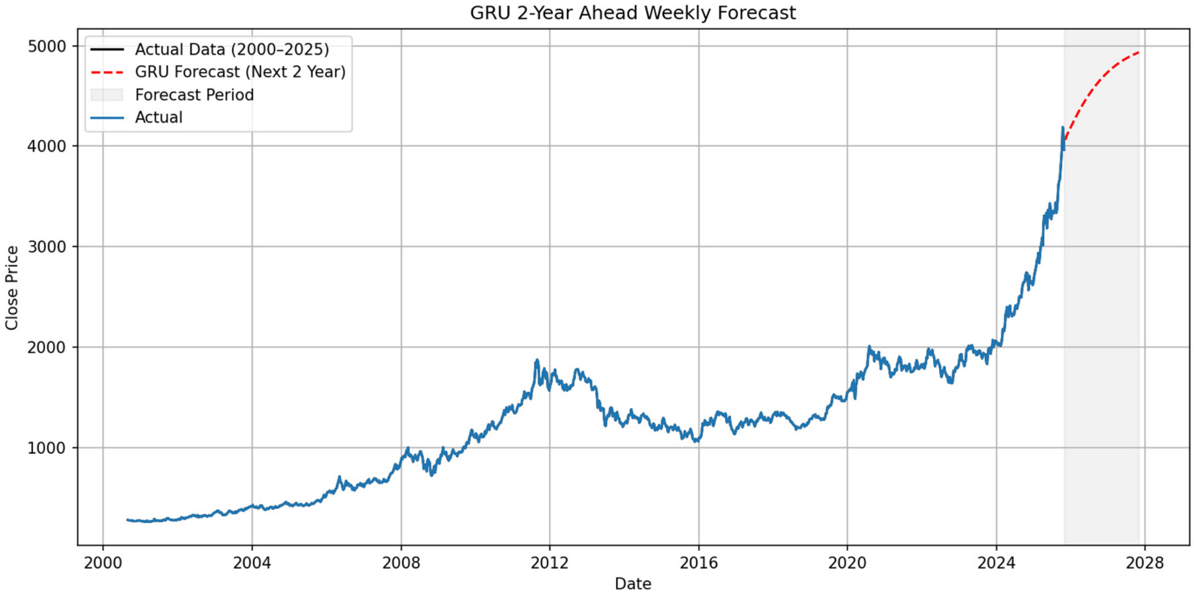

Figure 6. Actual values and the forecasts generated by the GRU.

5. Discussion

This study modelled the South African gold adjusted closing price using advanced RNN algorithms, namely the GRU and LSTM. The study aimed at comparing these algorithms and selecting the best-performing one between the GRU and LSTM in a financial time series forecasting task.

Descriptive statistics showed that the gold prices varied significantly across the three regimes. Pre-COVID-19 prices were moderate but highly variable, COVID-19 prices increased sharply with more stability, and post-COVID-19 prices reached their highest levels, with greater volatility. Overall, the full sample reflected substantial long-term variability in gold prices across all periods. The time series of gold prices from 2000 to 2025 showed an overall upward trend with increasing volatility, from smooth changes in early years to steep fluctuations in later periods. This pattern indicates that the series is non-stationary, with periods of rapid growth and heightened variability. The rolling mean and standard deviation were computed to assess the stability of the series over time. They confirmed that the adjusted closing gold price series was non-stationary, as the mean and variance were not constant over time. Differencing in this case was not necessary as the GRU and LSTM can handle non-stationary data.

As per the literature reviewed, similar studies—such as those conducted by Pirani, Thakkar [9]; Gao, Wang [11]; and Yamak, Yujian [16]—have used the same RNN algorithms that were evaluated in the current study, also in the financial sector. The results of the current study revealed that the GRU algorithm outperformed the LSTM algorithm across all analysed periods except during the COVID-19 period, where LSTM exhibited slightly better performance across all performance metrics. Similar results were also evident in the studies by Wang [4]; Pirani, Thakkar [9]; Mateus, Mendes [12]; Kumar, Hussain [17]; and Chaka, Semie [18]. However, there are historical instances where LSTM outperformed the GRU, such as in the studies by Saini, Kumar [19] and Xiao, Deng [3]. Additionally, both algorithms were selected as best-performing ones in the study of Gao, Wang [11], and they were outperformed by traditional algorithms such as ARIMA in the study of Yamak, Yujian [16]. This implies that, although several previous studies that were reviewed support the dominance of the GRU over LSTM, the performance of these algorithms may vary from one dataset to another. Therefore, it was important to conduct the current study, which focused on evaluating these two algorithms in a different context. As shown in Figure 6, the GRU forecast aligns more closely with the actual gold price movements compared to the LSTM forecast, particularly during periods of rapid upward trends. This visual evidence reinforces the empirical finding that the GRU provides a more accurate representation of the underlying temporal dynamics, while LSTM tends to lag slightly behind during sharp market fluctuations. The forecast comparison further strengthens the assessment of the relative performance of these deep learning architectures on sequential financial data characterised by both stability and volatility.

This research bridges the gap between advanced statistical methods and real-world financial industry applications in South Africa's financial sector. It lays the foundation for future research on the implementation of DL algorithms in financial and economic modelling. By implementing and comparing DL methods for time series forecasting and providing insights into their performance metrics, this study advances statistical modelling. It also contributes to the growing body of knowledge on the use of advanced RNNs in modelling financial time series, providing a robust methodology for evaluating algorithm accuracy and efficiency using metrics like the RMSE, MAPE, MAE, MSE, and R^2. In the South African context, this study addresses a gap in the application of advanced statistical and DL techniques to South African financial statistics—specifically the gold adjusted closing price. It also provides a localised perspective by analysing South Africa's gold market, a vital component of the national economy, deepening the understanding of its temporal dynamics. In the financial context, this study provides a useful financial forecasting tool that will benefit investors, policymakers, and market analysts in improving the accuracy of gold price predictions. DL methods show the potential of advanced computational techniques to model and predict market trends, supporting data-driven decision-making in the financial sector.

6. Limitations and Implications for Further Research

A limitation of this study is that the algorithms relied solely on historical adjusted closing prices, without incorporating external factors such as macroeconomic indicators and market sentiment, which have an influence on gold prices. This implies that future studies should be conducted to include these important variables. Additionally, the current study was restricted to the South African gold adjusted closing prices and does not generalise to other commodities, financial instruments, or global gold markets. Therefore, future research may be conducted to include adjusted closing gold prices from other countries and consider whether the algorithms' performance varies across different economic or market conditions, and the results may be compared to those found in the current study. Future studies could also explore the performance of the GRU and LSTM in comparison with other advanced algorithms like Transformer architectures, hybrid algorithms such as ARIMA-GRU and ARIMA-LSTM, or traditional statistical algorithms such as ARIMA- and GARCH-type models.

7. Conclusions

This study examined the predictive capabilities of LSTM and GRU algorithms in predicting the weekly adjusted closing price of gold in South Africa throughout the pre-COVID-19, COVID-19, and post-COVID-19 eras. With lower error metrics (MAE, MSE, RMSE, and MAPE) and stronger explanatory power (R^2), the results showed that the GRU model consistently outperformed the LSTM model in most assessment periods, especially before and after the COVID-19 pandemic. The LSTM model, however, showed marginally superior adaptation to abrupt structural changes during the turbulent COVID-19 period, underscoring the subtle variations in how each model captures temporal dependencies under various market situations. To the authors' knowledge, the use of weekly gold price data is rarely investigated in deep learning-based financial forecasting. This highlights the novelty of this work. The work provides a temporal robustness assessment of the GRU and LSTM by examining model performance throughout several economic stages, offering a more comprehensive understanding of these algorithms' behaviour in stable and crisis-driven conditions. This adds a new perspective to the existing body of research, which tends to ignore structural gaps in the time series in favour of focusing on daily or monthly data. Furthermore, because the GRU achieves competitive accuracy with fewer parameters than LSTM, the comparison methodology used here highlights the added value of using GRUs in situations where computational efficiency and stability are crucial. Practically, these results are important for financial analysts, investors, and policymakers looking for trustworthy instruments for risk management and medium-term gold price forecasts. The findings imply that GRUs may model sequential economic data with mild temporal dependencies more effectively than LSTM. Overall, this study highlights the significance of choosing an architecture that is in line with the data's features and the economic environment, in addition to validating the efficacy of deep learning models in capturing intricate temporal patterns in financial data.

Author Contributions: T.M. contributed to the conceptualisation, methodology, formal analysis, and writing—original draft preparation. T.B. contributed to the conceptualisation, methodology, writing—review and editing, and supervision. T.V.M. contributed to the methodology and writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The study employed secondary data obtained from the Yahoo Finance website. The authors agree to make the data and materials supporting the results or analyses presented in the paper available upon reasonable request.

Acknowledgments: The authors of this research acknowledge North-West University (NWU) for providing its resources to support this research.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Sudiatmika, I.P.G.A.; Putra, I.M.A.W. Comparison of LSTM and GRU Models Performance in Forecasting Gold Prices: A Case Study Using Historical Data from Yahoo Finance. *ARRUS J. Eng. Technol.* **2024**, *4*, 157–165. [CrossRef]
2. Saini, A.; Singh, R.K.; Sinha, P. Forecasting gold price using hybrid deep neural network LSTM-autoencoder. *Discov. Artif. Intell.* **2025**, *5*, 281. [CrossRef]
3. Xiao, J.; Deng, T.; Bi, S. Comparative analysis of LSTM, GRU, and transformer models for stock price prediction. In *Proceedings of the International Conference on Digital Economy, Blockchain and Artificial Intelligence*; Association for Computing Machinery: New York, NY, USA, 2024.
4. Wang, Q. Stock Prediction Under GRU, LSTM and GNN. In *Proceedings of the 2024 IEEE 2nd International Conference on Image Processing and Computer Applications (ICIPCA)*; IEEE: Piscataway, NJ, USA, 2024.
5. Sudiatmika, I.P.G.A.; Putra, I.M.A.W.; Artana, W.W. The Implementation of Gated Recurrent Unit (GRU) for Gold Price Prediction Using Yahoo Finance Data: A Case Study and Analysis. *Brill. Res. Artif. Intell.* **2024**, *4*, 176–184. [CrossRef]
6. Aljohani, H.; Alshamrani, S.; Aljojo, N.; Tashkandi, A.; Alsahfi, T. Predicting monthly gold prices in indian rupees using ARIMA, LSTM, GRU, and Simple Linear Regression models. *Rom. J. Inf. Technol. Autom. Control* **2024**, *34*, 35–48. [CrossRef]
7. Baradaran, A.H.; Bohlouli, M.; Motlagh, M.R.J. Decoding Tomorrow's Gold Prices: A Comparative Study of GRU and CNN-LSTM in the Iranian Market. In *Proceedings of the 2024 10th International Conference on Web Research (ICWR)*; IEEE: Piscataway, NJ, USA, 2024.

8. Dewi, N.P.J.R.; Ginantra, N.L.W.S.R.; Darma, I.W.A.S.; Indrawan, I.G.A. Application of Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) Algorithm in Gold Price Prediction. In *Proceedings of the 2023 11th International Conference on Cyber and IT Service Management (CITSM)*; IEEE: Piscataway, NJ, USA, 2023.
9. Pirani, M.; Thakkar, P.; Jivrani, P.; Bohara, M.H.; Garg, D. A comparative analysis of ARIMA, GRU, LSTM and BiLSTM on financial time series forecasting. In *Proceedings of the 2022 IEEE International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)*; IEEE: Piscataway, NJ, USA, 2022.
10. Primananda, S.B.; Isa, S.M. Forecasting gold price in rupiah using multivariate analysis with LSTM and GRU neural networks. *Adv. Sci. Technol. Eng. Syst. J.* **2021**, *6*, 245–253. [CrossRef]
11. Gao, Y.; Wang, R.; Zhou, E. Stock prediction based on optimized LSTM and GRU models. *Sci. Program.* **2021**, *2021*, 4055281. [CrossRef]
12. Mateus, B.C.; Mendes, M.; Farinha, J.T.; Assis, R.; Cardoso, A.M. Comparing LSTM and GRU models to predict the condition of a pulp paper press. *Energies* **2021**, *14*, 6958. [CrossRef]
13. Ng, Y.N.; Lim, H.Y.; Cham, Y.C.; Abu Bakar, M.A.; Ariff, N.M. Comparison Between LSTM, GRU and VARIMA in Forecasting of Air Quality Time Series Data. *Malays. J. Fundam. Appl. Sci.* **2024**, *20*, 1248–1260. [CrossRef]
14. Mienye, I.D.; Swart, T.G.; Obaido, G. Recurrent neural networks: A comprehensive review of architectures, variants, and applications. *Information* **2024**, *15*, 517. [CrossRef]
15. Ningsih, R.D.; Sarwido, S.; Wibowo, G.W.N. Comparative Analysis of Linear Regression, Decision Tree and Gradient Boosting for Predicting Stock Price of Bank Rakyat Indonesia. *J. Dinda Data Sci. Inf. Technol. Data Anal.* **2024**, *4*, 98–104. [CrossRef]
16. Yamak, P.T.; Yujian, L.; Gadosey, P.K. A comparison between arima, lstm, and gru for time series forecasting. In *Proceedings of the 2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence*; Association for Computing Machinery: New York, NY, USA, 2019.
17. Kumar, S.; Hussain, L.; Banarjee, S.; Reza, M. Energy load forecasting using deep learning approach-LSTM and GRU in spark cluster. In *Proceedings of the 2018 Fifth International Conference on Emerging Applications of Information Technology (EAIT)*; IEEE: Piscataway, NJ, USA, 2018.
18. Chaka, M.D.; Semie, A.G.; Mekonnen, Y.S.; Geffe, C.A.; Kebede, H.; Mersha, Y.; Anose, F.A.; Benti, N.E. Improving wind speed forecasting at Adama wind farm II in Ethiopia through deep learning algorithms. *Case Stud. Chem. Environ. Eng.* **2024**, *9*, 100594.
19. Saini, U.; Kumar, R.; Jain, V.; Krishnajith, M. Univariant Time Series forecasting of Agriculture load by using LSTM and GRU RNNs. In *Proceedings of the 2020 IEEE Students Conference on Engineering & Systems (SCES)*; IEEE: Piscataway, NJ, USA, 2020.