INTERNATIONAL REAL ESTATE REVIEW

2025 Vol. 28 No.4: pp. 475 – 504

Forecasting Real Estate Performance during the COVID-19 Pandemic Crisis: A Comparison of Statistical and Machine Learning Models

Damilola Aboluwodi*
University of KwaZulu-Natal, Email: Oludamilola123@gmail.com

Kazeem O. Isah
University of KwaZulu-Natal, Email: IsahK@ukzn.ac.za

Peter B.D. Moores-Pitt
University of KwaZulu-Natal: Email: MooresPittP@ukzn.ac.za

Paul F. Muzindutsi
University of KwaZulu-Natal: Email: MuzindutsiP@ukzn.ac.za

Machine learning (ML) methods, such as long short-term memory (LSTM) models, are increasingly proposed as alternatives to traditional statistical approaches for time series forecasting. However, given the speed of the real estate industry in providing data that reflect economic climates, there are few comparisons of ML techniques with statistical methods in the context of real estate data during crisis periods. The study investigates the predictive accuracy of the autoregressive integrated moving average (ARIMA) and LSTM models by using daily data from the Financial Times Stock Exchange/Johannesburg Stock Exchange South Africa Listed Property Index. Through a comprehensive analysis of 1628 observations from January 2, 2015, to July 8, 2021, the study finds that the ARIMA models produce fewer forecasting errors compared to the LSTM models during the COVID-19 crisis. These findings suggest that traditional ARIMA models may be more efficient for forecasting volatile real estate data in crisis periods, although the results could vary with larger and more complex datasets. This research is crucial as it provides insights into the comparative performance of statistical and ML models, thus emphasizing the need for context-specific model selection in economic forecasting.

* Corresponding author.

DOI: https://doi.org/10.53383/100411

Keywords

COVID-19 crisis, Real estate, Machine learning, Time series, LSTM and South Africa

1. Introduction

Over the years, time series forecasting has proven relevant as a predictive tool for improving decision making across various sectors such as the health, business, government, political, and technology sectors, and so on and so forth (Makridakis et al., 2018). Where sufficient data are available, time series analysis forecasting is beneficial for systematically and objectively predicting variations in data collected over time (Rhif et al., 2019). Despite these benefits, traditional forecasting time series data are beleaguered by several challenges, such as incomplete data, generalizations where conclusions are liberally inferred from single studies, challenges related to the choice of the right model and relevant performance measures to accurately predict a given dataset. In addition, where the data contain significant outliers, the accuracy or predictive power of time series analysis models is notably reduced. Finally, it is well known that traditional time series forecasting models are inadequate when calculating complex, nonlinear data (Cerqueira et al., 2019). Such methods therefore may not be sufficient when forecasting data during crisis periods. There is a constant need to improve predictive accuracy and reduce market volatility-induced forecasting errors in times of crises and uncertainties. To this end, machine learning (ML) techniques have been often presented as modern alternatives to traditional statistical methods, but whether they are an improvement of the former, is hotly debated in the forecasting circles.

The real estate sector is characterized by its speed and elasticity in responding to market changes, and therefore its embodiment of the prevailing state of the economy (Zheng et al., 2024). Volatility in the economy is often reflected in the real estate sector and fluctuations in the latter are often an indication of the same in the former (Nazlioglu et al., 2016). However, while frequent, daily real estate data are relatively accessible, the same cannot be said for gross domestic product (GDP) data as a traditional indicator of economic activity (Nielsen, 2019). Therefore, forecasting of real estate data is strategic in that it provides insight into broader economic data, and somewhat of an economic forecast by proxy (Alola, 2021). This study is useful in forecasting the real estate sector more accurately and also providing more timely and accurate indications of broader economic activity than would be the case with more traditional indicators. The primary goal of this research is to comparatively evaluate and determine which forecasting technique is more suitable for predicting real estate performance during COVID-19 as a market crisis period, based on higher forecasting accuracy and fewer forecasting errors. This paper contributes to the ongoing debate on the effectiveness of ML models as superior alternatives to traditional statistical models in times of crisis, and in particular, the long short-term memory (LSTM) model versus the autoregressive integrated moving average (ARIMA) model. While much related work exists on the comparison of the LSTM and ARIMA models in forecasting, fewer studies are found on this same comparison during periods of economic crisis, and even less in the context of South Africa (SA).

The traditional econometric techniques for forecasting time series data are the popularly known autoregressive (AR) and moving average (MA) models. These have evolved over time into several variants for forecasting univariate and multivariate time series data (Hamilton, 2020). For univariate data, variants of the AR and MA models include: the autoregressive moving average (ARMA), ARIMA, seasonal ARIMA (SARIMA), ARIMAX, and seasonal ARIMAX (SARIMAX) models (Guidolin and Pedio, 2018). Likewise, for multivariate econometric time series models, the AR and MA model frameworks are extended through models such as the vector autoregression (VAR) and vector autoregressive moving average (VARMA) models (Guidolin and Pedio, 2018). These models are especially useful for forecasting and understanding the dynamic interrelationships among multiple time series, under the assumption that historical linear dependencies among variables can inform future outcomes (Vishwas and Patel, 2020). For example, investors of real estate may base their future purchases on the recent past. In such a case, higher house or rental prices (or similar trends in the supporting data) may trigger additional investment into real estate, whereas lower prices may trigger the opposite. In this case, the MA and AR models would prove to be ideal forecasting tools (Samadani and Costa, 2021).

ML methods tend to be built on statistical techniques so a complete split or differentiation between the two is often infeasible (Cerqueira et al., 2019). Nonetheless, ML forecasting techniques rely on the ability of computers to learn from data by developing algorithms through trial and error, and then forecast depending on this process (Makridakis et al., 2018). However, despite the recent attention, such methods are still not well established in the forecasting literature (Cerqueira et al., 2019). These models are particularly beneficial for the analysis of complex and non-linear data owing to their neural network design; a simulation of the human brain which can handle the same complex process of logic and computation (Pal and Prakash, 2017; Brownlee, 2018).

The main objective of this research paper is to comparatively evaluate and determine which forecasting technique is more suitable for predicting the performance of the listed real estate market of SA during periods of economic crisis, specifically the COVID-19 pandemic. Within this context, performance refers to return levels derived from market prices rather than volatility or risk measures. While real estate market volatility is a commonly studied indicator of uncertainty, this paper focuses on the directional forecasting of real estate returns. To achieve this, we compare the predictive accuracy of the ARIMA model, a classical time series technique, and the LSTM model, a deep learning method designed for capturing complex temporal patterns. By clarifying the distinction between volatility modelling and return forecasting, we position the study within the broader discussion on selecting appropriate forecasting tools during high-uncertainty periods in emerging property markets. To this end, the research study is structured into the following sections. Section 2 presents the literature review and discusses related work. Section 3 explores empirical studies on the ARIMA and LSTM models. In Section 4, an introduction and a description of the data are provided, and the methodology is discussed. In Section 5, the results are presented, while the findings are summarized in Section 6. Finally, the conclusion and directions for future research are found in Section 7.

This study contributes by advancing the field of real estate forecasting in three key ways: (1) a rare empirical comparison between the ARIMA and LSTM models is done to examine listed real estate performance during a crisis period, specifically the COVID-19 pandemic in SA, an underrepresented context in the literature; (2) high-frequency daily data are leveraged to improve predictive resolution of the models, thus addressing calls in the literature for more granular real estate modelling; and (3) the often-assumed superiority of ML models is challenged by demonstrating that ARIMA can outperform LSTM in settings with relatively limited complexity and sample size. These contributions help to refine model selection considerations for researchers and practitioners who are working with financial and real estate time series, particularly in emerging markets and crisis conditions.

2. Literature Review

As a modelling technique, time series analyses have a proven track record of use in different professions and disciplines, such as science, business, economics, finance health, media, meteorology, and even military studies (Nielsen, 2019). The wide acceptance of this modelling tool is largely due to its relevance and versality in practical applications as well as its reliance as a scientific basis for predictions. As a veritable modelling tool, the end goal of time series analyses is borne out of the question of causality, which seeks to establish how past data trends and patterns can predict the future direction of a given data set (Nielsen, 2019; Vishwas and Patel, 2020). This ultimate concern and general applicability of time series models explicate the criticality of the forecasting accuracy of time series models. Hence, it is no surprise that several time series analysis models have evolved over the years in an attempt to constantly improve forecasting accuracy.

Until recently, forecasting time series data have been synonymous with econometric models with roots that can be traced to the AR and MA models. Notably, among these models is the ARIMA model that has long dominated the forecasting domain of the econometrics and statistics discipline. While the ARIMA model has shown laudable predictive prowess in forecasting simple and linear time series problems, there has been empirical evidence that shows the limitations of the ARIMA model. For instance, several studies have noted the shortcomings of the ARIMA model in non-linear time series problems (Wang et al., 2018; Nassar et al., 2020). Likewise, some studies have observed that ARIMA models are inadequate in predicting time series data over a long run horizon (Brownlee, 2018; Kazemzadeh et al., 2020). These constant concerns have queried the long-term relevance of traditional econometric models especially in the light of advancements in artificial intelligence and ML forecasting techniques.

In recent times, ML techniques, especially deep learning algorithms, have been developed to address the shortcomings of traditional econometric forecasting models such as the ARIMA. Notable among these techniques is the LSTM model, which is an advanced form of recurrent neural networks (RNNs) used in deep learning (Rosinus, 2025). As a modern alternative forecasting model, LSTM networks are remarkably efficient in dealing with a number of forecasting problems as they are capable of avoiding vanishing gradient issues, thus recalling information over a long sequence of time series, learning long term dependencies and incorporating feedback connections (Jadon et al., 2021). However, the LSTM model suffers from critical drawbacks such as long computational time, higher memory usage for training, sensitivity to weight initialization and the tendency to overfit data (Jadon et al., 2021). These challenges of the LSTM model have raised doubts on its viability as a preferred substitute for traditional and econometric models.

As a result of these doubts, several empirical studies have sought to evaluate and compare the forecasting accuracy of both techniques by using different time series datasets. Miswan et al. (2014) find that the ARIMA model is most suitable when comparing the performance between the ARIMA and generalized autoregressive conditional heteroskedasticity (GARCH) models to model and forecast the volatility of Malaysian market properties and shares. However, Adebiyi et al. (2014) note that ML techniques provide even better forecasting accuracy for predicting American stocks when compared to the ARIMA. Particularly during the COVID-19 pandemic, researchers appear to have increasingly opted for ML techniques in analyzing financial and real estate data. They include, among other researchers, Saravagi et al. (2021), who employ the LSTM model to predict the stock prices of selected companies during 2020 in India; as well as Grybauskas et al. (2021) who use ML techniques to analyze real estate market big data for Vilnius and Lithuania during the first wave of the COVID-19 pandemic. This is in comparison to the few who have opted to forecast financial and real estate data through traditional time series methodologies during COVID-19 (Zheng et al., 2024).

There appears to be limited related work in SA, in particular where domestic real estate data are forecasted with both models. However, data from other sectors may have similar characteristics to real estate data during crisis periods, such as fluctuation and non-linear trends (Zhang et al., 2022) and therefore, may shed some light. Zambezi (2021) forecasts social unrest incidents in SA by comparing ARIMA and LSTM modelling, and daily non-linear and irregularly fluctuating social data between 2002 and 2017. The study finds that complex multivariate social data on social unrest incidents is even better predicted by using LSTM models. Similarly, Essa et al. (2021) use 20 million lightning observations and find that the LSTM model is more effective at forecasting lightning data in SA than the ARIMA model. Lightning data are also characteristically nonlinear and dynamic in nature. The LSTM model appears to perform more efficiently in data with similar characteristics in the region. The model also appears to be more effective with more voluminous (big) data.

Existing research shows that there is no real consensus in whether ML techniques or traditional time series modelling are better for forecasting data, which leads to the proposition that the choice of the model may need to be determined by context, data size, and the nature of the data. For example, Makridakis et al. (2018) contend that there is little objective evidence available on the relative performance of ML techniques as a standard forecasting tool compared to traditional statistical methods. This conclusion is based on a forecast of 1045 monthly time series used in the M3 Competition, in which eight traditional methods, including the ARIMA and seven other statistical models are compared to the LSTM and nine other ML models by way of forecasting accuracy. They conclude that traditional statistical methods systematically outperform ML methods in this regard.

In response, Cerqueira et al. (2019) contend that the relative predictive performance of ML methods improves with larger sample sizes. They argue that the analysis in Makridakis et al. (2018) is biased as their dataset is too small for ML models to perform at their optimum. Therefore, the conclusions of the study are deemed inconclusive for generalization purposes.

Siami-Namini and Namin (2018) show that an increase in data size enhances the accuracy of the results when they compare the accuracy of the ARIMA model to LSTM model in forecasting twelve financial and economic time series, which range from 368 to 1,698 observations (one being real estate or housing data). The results greatly point to the superiority of the LSTM model over the ARIMA model. However, Siami-Namini and Namin (2018) only test the results by using the root mean square error (RMSE) as the performance metric (and no other metric) to determine the accuracy of the prediction and evaluate the forecasts. Even so, not all ML methods appear to be equal. In analyzing the performance of LSTM neural networks for nowcasting (where current or near current variables are estimated) during the COVID-19 crisis, Hopp (2021) notes that LSTM neural networks are more effective as they contain a time-based element, which traditional ML techniques (traditional feed-forward networks such as back-propagation neural networks (BPNNs) and Elman's RNN (ERNN) lack. He finds that when compared to a dynamic factor model (DFM) in nowcasting global merchandise export values and volumes and global services exports, LSTM performs better. He also finds LSTM to have had more gradual forecasting changes. Consequently, LSTM models produce superior predictions in nowcasting.

Aladag et al. (2009) use a hybrid model of ARIMA and ERNN to find an optimal model to model the non-linear data of Canadian Lynx. They find that the proposed hybrid approach has the highest forecasting accuracy, compared to forecasting by using either method alone. Likewise, Merh et al. (2010) present a comparison between the hybrid approaches of ANN and ARIMA for Indian stock trend forecasting with many instances of the ARIMA predicted values shown to be better than those of the artificial neural network (ANN) predicted values in relation to the actual stock value. Similarly, Lee et al. (2008) compare the performance of ARIMA versus ANN models in forecasting the Korean Stock Price Index and find the former to generally provide more accurate forecasts than the BPNN model.

Table 1 provides a summary of the most related studies. Where the ARIMA and LSTM models are directly compared with analyzed economic and financial data, more studies are in favor of the use of the LSTM model by virtue of the results of the forecast error estimators. These include Siami-Namini and Namin (2018), and Hopp (2021), (2022). Two other studies compare other ML methods to the ARIMA model and find them superior including Adebiyi et al. (2014) and Cerqueira et al. (2019).

Table 1 An Overview of the Most Related Works

Reference	Method	Data	Region	Contribution
Hopp (2022)	LSTM vs DFM	Global merchandise export values and volumes, and global services exports (2020 Q2 to2021 Q2). 9 quarterly datasets with 4 observations/46 monthly datasets with 15 observations (estimated)	Global	LSTM models are better at nowcasting and gradual forecasting changes than DFM models when considering mean absolute and root mean square errors. LSTM models are more effective as they contain a time-based element, when compared to traditional feed forward networks (traditional neural networks).

(*Continued…*)

(Table 1 Continued)

Reference	Method	Data	Region	Contribution
Hopp (2021)	LSTM vs DFM	Global merchandise and services trade (2016 - 2019). 101 monthly datasets with 48 observations/13 quarterly datasets with 16 observations (estimated)	Global	LSTM models are better at nowcasting and gradual forecasting changes than DFM models when considering mean absolute and root mean square errors. LSTM models are more effective as they contain a time-based element, when compared to traditional feed forward networks.
Cerqueira et al. (2019)	ARIMA and seven other statistical models vs. LSTM and nine other ML models	90 univariate time series of 1000 observations each in various sectors including healthcare, physics, economics. Various frequencies, incl. daily, monthly, etc.	N/A	ML methods improve their relative predictive performance as the sample size increases.
Makridakis et al. (2018)	ARIMA and seven other statistical models vs. LSTM and nine other ML models	1045 monthly (micro, industry, macro, finance, demographic, other) time series used in the M3 Competition with an average, minimum, and maximum number of observations per time series of 118, 66, and 144, respectively.	N/A	Traditional statistical methods systematically outperform ML methods

(*Continued*...)

(Table 1 Continued)

Reference	Method	Data	Region	Contribution
Soy Temür et al. (2018)	ARIMA, LSTM and hybrid models	Monthly housing sales in Turkey, from January 2008 – April 2018. 124-observations.	Turkey	The hybrid model has better predictive power than ARIMA and LSTM models separately, however the ARIMA model produces better results than the LSTM model.
Siami-Namini and Namin (2018)	ARIMA vs LSTM (RMSE)	Financial and economic timeseries monthly (incl. housing data), ranging from 368 to 1,698 observations (Housing data 368 observations) (estimated)	USA, Hong Kong, Japan	Deep learning-based algorithms such as LSTM outperform traditional-based algorithms such as ARIMA model.
Adebiyi et al. (2014)	Machine learning (ANN) ARIMA	Daily NYSE stock prices (Dell stock index) from 18 August 1988 – 25 February 2011. 5679 observations.	USA	Machine learning techniques provide better forecasting accuracy for American stocks
Merh et al. (2010)	Hybrid approaches of ANN versus ARIMA		India	Many instances of ARIMA predicted values shown to be better than those of the ANN predicted values in relation to actual stock value
Lee et al. (2008)	ARIMA versus BPNN		Korea	The ARIMA model generally provides more accurate forecasts than the BPNN model used.

No conclusive distinction can be made in respect of the data size between the three financial and economic studies in favor of LSTM models versus the remaining two in favor of the ARIMA model, as those in favor of LSTM generally range from four observations over the COVID-19 period (Hopp, 2022) to 1,698 in more stable times (Siami-Namini and Namin, 2018), while the two in favor of ARIMA modelling generally range from 66 observations to 144. However, both Cerqueira et al. (2019) and Adebiyi et al. (2014) include observations in excess of 1,000 (1,000 and 5,679, respectively) and both favor ML methods. This may suggest, as Cerqueira et al. (2019) advocate, that ML methods improve the relative predictive performance as the sample size increases. Nonetheless, all studies with observations in excess of 500 appear to be in favor of ML forecasting in general, and in one case (Siami-Namini and Namin, 2018), LSTM in particular. Studies in favor of ARIMA modelling appear to have smaller datasets; however, more research is required to reach a conclusion.

Even less conclusive is the fact that few studies exist that compare both methods (LSTM and ARIMA) to determine whether the superiority of the model can be influenced by other factors such as the region where the data are sourced, its sector, or even whether the data are sourced from periods of economic crisis. However, studies in other sectors (social data, scientific data) appear to suggest that LSTM models are more suited for forecasting more complex, multivariate data (Essa et al., 2021; Zambezi, 2021a), while ARIMA models are best suited for predicting univariate linear time series data.

In summary, the jury is still out on the superiority of ML techniques when compared to traditional statistical models for the estimation of financial and real estate data; however, the literature and existing empirical studies appear to support several mixed conclusions. The first is that of all the traditional statistical models, the ARIMA model appears to be most suitable for volatile data (Miswan et al., 2014). Moreover, the ARIMA model has performed more efficiently than traditional feed forward neural networks such as the ERNN and BPNN models (not LSTM models) (Lee et al., 2008).

Secondly, the literature points to the superiority of neural networks over traditional statistical models when estimating nonlinear data (such as the financial and real estate data used in this study) (Adebiyi et al., 2014). Of all the neural networks, LSTM models also appear to have a better performance due to their temporal component, which other traditional feed forward networks do not have (Essa et al., 2021; Zambezi, 2021a). Therefore, of all the neural networks, the literature points to LSTM as the most efficient for the study in question. Moreover, empirical studies point to increasingly better performance of LSTM models as datasets increase in size. As sample sizes decline, traditional statistical models provide a better forecasting performance (Cerqueira et al., 2019).

Thirdly, empirical studies and the literature suggest that the type of data may influence the suitability of the model. Fewer errors have been produced in empirical studies with LSTM over ARIMA models when forecasting multivariate, volatile data, whereas fewer errors have been produced by ARIMA over LSTM models when forecasting univariate data. Finally, hybrid models, those which are a cross between ML and traditional statistical models, are more efficient than either alone, given that they both have limitations. Notably, there appears to be negligible related work (apart from Hopp (2021)) based on a comparison of the performance of both models particularly during crisis periods in particular. Zhang et al. (2022) confirm the uniqueness of crisis data; however, Hopp (2021) finds no difference in the preference for LSTM over the DFM whether pre- or post- COVID-19 crisis period.

Therefore, the literature and related empirical work have not reached a consensus. However, there is some evidence that suggests the ARIMA models provide the best statistical tools for estimating volatile real estate data, while the LSTM models appear to provide the best ML tools for forecasting the same, and ultimately, the latter perform relatively better as the sample size increases.

3. Methodology

3.1. Data and Sample Period Description

The study uses the daily prices of the Financial Times Stock Exchange/Johannesburg Stock Exchange (FTSE/JSE) South Africa Listed Property Index (J253). Specifically, the use of daily data is justified as it enables the model to capture abrupt, short-term shifts in market behaviors that are often smoothed over in lower-frequency datasets. This is particularly important in periods of crisis where price sensitivity and market reactivity are heightened. Furthermore, this dataset comprises low, high, opening, and closing prices as well as the volume of daily prices. For the purpose of this study, the closing prices of the South African listed property index are adopted as the benchmark measure of the listed property index performance to ensure consistency. While listed properties are unique in their business and investment objectives, as they focus on real estate and a steady stream of income, the SA listed property index (J253) is particularly critical as it enlists only the top 20 liquid real estate companies on the JSE.

More so, the FTSE/JSE SA Listed Property Index (J253) is selected as the focus of this study because it serves as a key benchmark for the performance of the listed real estate sector in SA. This index includes the most actively traded and liquid real estate investment trusts (REITs) and property stocks on the Johannesburg Stock Exchange, thus offering a representative and consolidated view of investor sentiment and market trends within the property asset class. As a subcomponent of the broader equity market, the J253 is particularly sensitive to macroeconomic shocks and interest rate changes, thus making it an ideal candidate for assessing the responsiveness and forecasting accuracy of financial models during crisis periods. Additionally, although the South African listed property market remains well-established and institutionally mature, it still operates within an emerging market context which is characterized by volatility, regulatory shifts, and unique structural dynamics. This duality enhances the value of the index as a test case for forecasting models intended to inform both investment decisions and policy-making in similar economic environments. Furthermore, the data are retrieved from the Integrated Real-time Equity System (IRESS) Research Domain database. The sample period for this study is from January 2, 2015, to July 8, 2021. This timeframe covers both the pre-COVID-19 and COVID-19 periods, thus providing sufficient and robust data to evaluate the impact of the pandemic in our analysis. Hence, a total of 1,628 observations are obtained and considered for the purpose of this research inquiry. The subsequent section clarifies the adopted methodology for the research study.

3.2. Model Specifications

3.2.1 ARIMA Model

The ARIMA model, which is also known as the Box-Jenkins model, is a generalized form of the ARMA model that includes integrated components (Pal and Prakash, 2017). These integrated components are crucial when dealing with datasets that are non-stationary, as the integrated component of the ARIMA model helps to transform non-stationary data into a stationary form. This is accomplished by the ARIMA model which applies differencing on the time series data one or more times to eliminate the non-stationarity effect (Pal and Prakash, 2017; Nielsen, 2019).

The ARIMA model is denoted by (p, d, and q), which represents the AR, MA and the differencing components. The d component which distinguishes the ARMA model from the ARIMA model, aims to de-trend the signal to make time series data stationary and suitable for forecasting purposes (Pal and Prakash, 2017). This is why the ARIMA model is more widely preferred and utilized for forecasting and academic research than other time series models such as the AR, MA and ARMA models (Nielsen, 2019).

Using the ARIMA model for forecasting, the future value of a defined variable can be determined via the use of a linear function that contains both past observations as well as random errors as expressed below:

$$y_t = c + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \cdots + \alpha_p y_{t-p} + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \cdots + \theta_q \varepsilon_{t-q} + \varepsilon_t \tag{1}$$

or equivalently by using:

$$\left(1-\sum_{j=1}^{p}\alpha_j \boldsymbol{L}^j\right)(1-\boldsymbol{L})^d y_t = \left(1+\sum_{i=1}^{q}\theta_i \boldsymbol{L}^i\right)\varepsilon_t \tag{2}$$

where L is a lag operator, y_t is the actual value (i.e. price of listed property) at time t, ε_t is a random error at t, c is the intercept or constant, θ_i ($\mathrm{i} = 1, 2, \ldots .\mathrm{q}$), α_{j} ($\mathrm{j} = 1, 2, \ldots \mathrm{p}$) are the model parameters, and q and p are integers and often referred to as the MA and AR orders of the model, respectively. The assumption regarding the random errors ε_t is that they are identically and independently distributed with a mean zero and constant variance of σ^2.

Steps to ARIMA Modelling

Using the ARIMA methodology, ARIMA modelling follows a three-step iterative approach which involves identification of the model variables, estimation of the parameters and application of diagnostic tests to evaluate and ascertain the optimal parsimonious model from a variety of different ARIMA models (Olson and Wu, 2017).

Stationarity Test

In the context of this study, we first visualize the time series dataset to determine whether the dataset is stationary in order to commence the ARIMA modelling process We then perform preliminary tests for both seasonality and stationarity of our times series data by using natural logarithm transformation and differencing techniques. If the dataset is non-stationary, we apply differencing to the time series to enforce stationarity. This is illustrated in Figure 1.

Autocorrelation and Partial Autocorrelation Functions

Upon stationarity of the dataset, we determine the optimal parameters to be estimated by our model. To this end, we apply both the autocorrelation function (ACF), and partial autocorrelation function (PACF) to determine the best order of our ARIMA model. This is presented in Figure 2, which shows the natural logarithm transformed dataset after first differencing and the suitability of the data for time series modeling.

Model Diagnostic Checks

Thereafter, we express the different series of our ARIMA model as well as their estimated parameters (α, θ and σ) by using the maximum likelihood method and the assumption that the error terms are independent and normally distributed. Thus, our log-likelihood function is expressed as:

$$\log L\,(\alpha, \theta, \sigma^2\,) = -\frac{n}{2}\log(2\pi\sigma^2) - \sum \varepsilon_t^2(\,\alpha, \theta)/2\sigma^2 \tag{3}$$

Thereafter, we build our time series model and make predictions based on the optimal ARIMA model.

Figure 1 **Closing Price Changes (July 1, 2019, to July 1, 2020): A) before Differencing, and B) after Natural Log Transformation and Differencing**

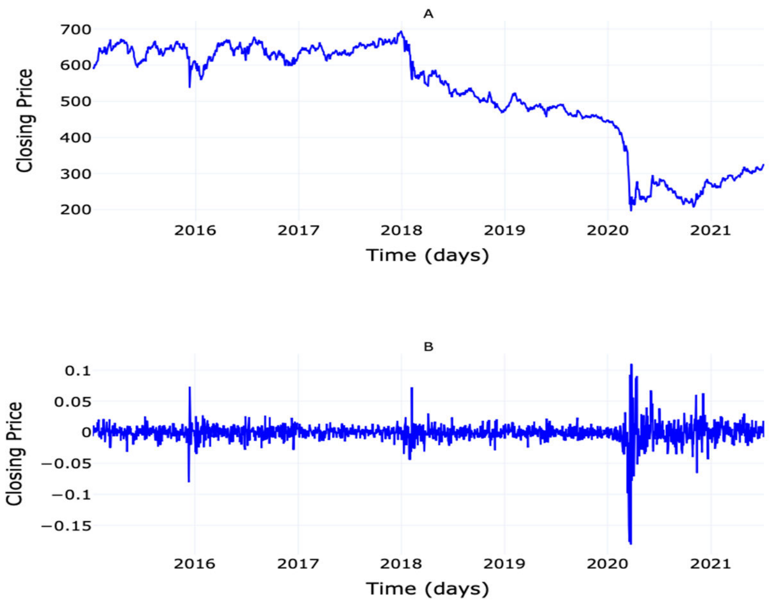

Figure 2 **ACF and PACF of Lagged Time Series after First Order Differencing**

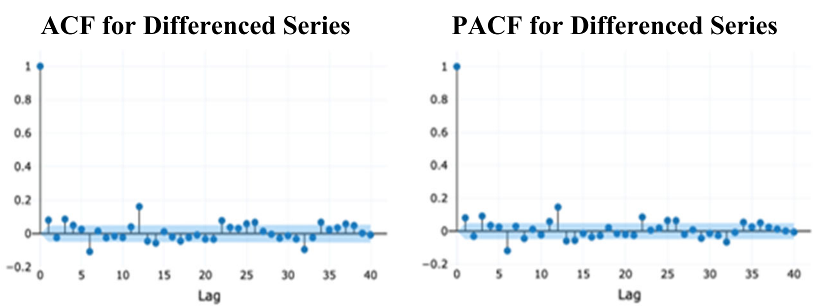

As presented in Table 2, model diagnostic checks are performed to determine the most parsimonious- as opposed to complex- model, as well as the goodness of fit of each one by estimating the relative prediction accuracy or error for each ARIMA model variation, with the use of the six measures as listed. In each case, where the value of the measure is the smallest, this is considered the best fitting and most parsimonious variation of the ARIMA model. All but ARIMA (1,1,3), which was selected twice, were selected by only one measure if at all, and therefore ARIMA (1,1,3) is selected as the model for forecasting.

Furthermore, the choice of the ARIMA (1,1,3) specification is informed by a rigorous model selection process based on widely accepted criteria such as the Akaike information criterion (AIC) and Bayesian information criterion (BIC), both of which favor this configuration. Additionally, this model shows superior performance across most forecasting error metrics, including the mean absolute error (MAE) and mean absolute percentage error (MAPE). While several alternative ARIMA configurations are estimated, the (1,1,3) model consistently outperforms the others in both in-sample fit and out-of-sample forecasting accuracy. This approach aligns with that of Stevenson (2007) who emphasizes the importance of balancing parsimony with predictive accuracy in ARIMA model selection, particularly in real estate applications where overfitting can distort inference. Accordingly, the chosen model reflects both empirical robustness and theoretical soundness.

Table 2 Evaluation of ARIMA Models

Model	AIC	BIC	ME	RMSE	MAE	MAPE
ARIMA (1, 1, 1)	-7889.31	-7873.51	-0.250	6.287	4.256*	0.869*
ARIMA (1, 1, 2)	-7888.72	-7867.65	-0.255	6.284*	4.261	0.869
ARIMA (1, 1, 3)	-7919.55*	-7893.21*	-0.224	6.371	4.333	0.881
ARIMA (1, 1, 4)	-7914.73	-7883.13	-0.234	6.369	4.331	0.880
ARIMA (2, 1, 1)	-7907.16	-7886.09	-0.257*	6.318	4.315	0.876
ARIMA (2, 1, 2)	-7907.74	-7881.40	-0.230	6.317	4.314	0.877
ARIMA (2, 1, 3)	-7917.22	-7885.62	-0.225	6.371	4.340	0.882
ARIMA (2, 1, 4)	-7918.45	-7881.58	-0.223	6.374	4.348	0.883

***Notes*:** * indicates the lowest value in each row.

3.2.2 LSTM Model

LSTM networks are a type of RNN designed to handle the vanishing gradient problem that plagues traditional RNNs (Nielsen, 2019). LSTM has memory cells that allow information to be remembered over a longer period of time, thus making it suitable for tasks such as language modeling, speech recognition, and sequential prediction. An estimated LSTM model is a pre-trained LSTM model that has been fine-tuned on a specific task with a dataset. The parameters of the model have been estimated or trained from data to minimize a loss function that measures the difference between the predictions and the actual output. This fine-tuning process allows the model to make predictions for the specific task with high accuracy (Gridin, 2021).The choice of using the LSTM model over other ML models is largely due to its strong capacity to learn and retain long-term dependencies in sequential data, a property particularly relevant to real estate time series where historical price patterns often influence future trends. Compared to other deep learning architectures such as convolutional neural networks (CNNs), which excel in spatial data tasks, or gated recurrent units (GRUs), which are a simplified version of LSTM at the cost of nuanced memory retention, LSTM networks provide a more robust framework for modelling temporal dynamics. Given the research aim to forecast time-dependent financial performance, LSTM is deemed the most suitable among the neural network alternatives for capturing the intricate, time-lagged relationships found in daily real estate market data.

In addition to this, there are several reasons why an LSTM model may be selected for time series forecasting. First, an LSTM model is designed to handle long-term dependencies in time series data, where the current value of the time series depends on values from many time steps ago (Bianchi et al., 2017). Secondly, an LSTM model can handle multiple input features, which makes the model well-suited for time series data that involve multiple related variables (Gridin, 2021). Thirdly, an LSTM model can handle missing or noisy data by using gates to control the flow of information in the network (Pal and Prakash, 2017). Lastly, an LSTM model is capable of modeling non-linear relationships between the input and target variables, which make this model well-suited for time series data that exhibit non-linear patterns or trends (Lazzeri, 2020).

Steps Involved in LSTM Model Estimation and Training

To build an ML (LSTM) model for forecasting with Python, a programming language, this study follows a 5- step approach that involves: data preprocessing, building the model architecture, and compiling, training, and evaluating the model (Brownlee, 2018; Korstanje, 2021).

Data preparation

To commence the analysis, the real estate data are preprocessed by using a normalization technique that utilizes the scikit-learn preprocessing library and a classical data transformer to scale each data feature into transformed intervals between 0 and 1 (Brownlee, 2018; Korstanje, 2021).

$$\tilde{Y}_t = \frac{Y_t - Y_{min}}{Y_{max} - Y_{min}} \tag{4}$$

Where

$$Y_{min} = \min_{t=1,2\ldots N} Y_t$$

$$Y_{max} = \max_{t=1,2\ldots N} Y_t$$

$\tilde{Y}_t \in [0,1]$ are the scaled values of Y_t, N is number of observations, and Y_t is the original data.

This process helps to rescale datapoints into manageable weight ranges that are stable and less prone to ML errors. To effectively utilize the LSTM algorithm for prediction, the study utilizes real estate market data features including opening, high, low and closing prices, volume and returns that are considered over a 5-window period from January 5 to 9, 2015. The last 5 values of the data features are used to predict the next closing price and squared return. The data before normalization are listed in Table 3 and post normalization in Table 4.

Table 3 Data before Normalization

Date	Opening Price	High Price	Low Price	Closing Price	Volume	Returns
2015-01-05	593.73	594.18	589.51	589.99	11,411,264	-0.00630
2015-01-06	589.99	594.29	587.74	593.63	14,818,142	0.00617
2015-01-07	593.63	594.39	590.17	592.85	15,017,394	-0.00131
2015-01-08	592.85	600.82	592.22	599.19	21,304,403	0.01070
2015-01-09	599.19	601.56	597.29	600.74	18,151,556	0.00258

Table 4 Data post Normalization

Date	Open	High	Low	Close	Volume	Returns
2015-01-05	0.79745	0.78745	0.80801	0.78995	0.00306	0.56446
2015-01-06	0.78995	0.78768	0.80445	0.79726	0.00619	0.60871
2015-01-07	0.79726	0.78789	0.80934	0.79569	0.00637	0.58216
2015-01-08	0.79569	0.80105	0.81347	0.80842	0.01214	0.62476
2015-01-09	0.80841	0.80256	0.82368	0.81153	0.00924	0.59599

Model architecture design

Building the architecture of the model involves determining the number of layers and hidden units, type of activation functions, etc. that will be used in the LSTM model (Brownlee, 2018; Korstanje, 2021). This study considers a 6-2 LSTM network topology, where each number corresponds to the number of neurons in a given layer and the last layer corresponds to the prediction (we predict the closing price and squared returns at the same time). We also use a rectified linear unit (ReLU) as an activation function (for cell and hidden states) and sigmoid as a recurrent activation function (for activating the input/forget/output gate) in each layer. The ReLU is a popular activation function used in deep learning models, including LSTM models. The ReLU activation function is defined as $y = max(0, x)$, where x is the input for the activation function and y is the output. The motivation for using the ReLU activation function is that it introduces non-linearity into the model, thus allowing the model to learn more complex relationships between the inputs and outputs.

Also, to minimize overfitting errors, we use the dropout technique and bias regularizer. Using the drop-out technique, we define the dropout level = 0.12 (i.e., it drops out the cell and hidden states in LSTM) for each layer and recurrent dropout level = 0.02 (i.e., it drops out the input/update gate in LSTM). Likewise, we also use the bias regularizer in the L2 form, i.e., the cost function which contains a penalty term $0.2 \times \sum b_i^2$, where b_i is bias. The purpose of bias regularization is to prevent overfitting and improve generalization by adding a penalty to the loss function during training.

Model compilation

The compilation of the LSTM model involves specifying the optimizer, loss function, and evaluation metrics to be used during training (Castro et al., 2022). The loss function measures the difference between the actual and predicted values and is used as a guide for updating the model parameters during training. Common loss functions for time series forecasting include the mean square error (MSE), mean absolute error (MAE), and mean absolute percentage error (MAPE). However, only the MSE is considered in this study as it is a commonly used loss function for regression problems, including time series forecasting with LSTM models. This is because the MSE measures the average squared difference between the model predictions and the actual output.

Furthermore, the LSTM optimizer updates the model parameters based on the gradient of the loss function. While common optimizers for LSTM models include the stochastic gradient descent (SGD), and adaptive moment estimation (Adam), and adaptive gradient algorithm (Adagrad) optimizers, this study adopts the Adam optimizer as it calculates adaptive learning rates for each weight and bias in a neural network, which helps the model to converge faster and more effectively compared to traditional gradient descent optimization (Siami-Namini et al., 2018; 2019).

Lastly, the LSTM model is evaluated on several metrics to determine the performance of the model on the validation or test data. The selected metrics include the MAE, RMSE, and MAPE. Evaluating the model performance across several benchmarks offers a robust assessment of the LSTM model for forecasting purposes.

Model training

Training the model involves inputting the preprocessed data into the model and training it over multiple epochs. During each epoch, the model processes the data, makes predictions, and updates its parameters based on the loss and optimizer (Korstanje, 2021; Castro et al., 2022). In a typical LSTM model training process, the loss function, which in this case is the MSE, is calculated after each epoch and used to update the model parameters. Over time, the model continues to learn from the data and improves its predictions, thus leading to smaller MSEs. However, it is important to monitor the validation loss, which is calculated by using a separate validation set. The validation loss is used to evaluate the performance of the model on new data and ensures that it is not overfitting to the training data. If the validation loss increases while the training loss continues to decrease, this may indicate that the model is overfitting to the training data, and further steps may need to be taken to improve its generalization performance. In general, a good model should have a decreasing training loss and a stable or slightly decreasing validation loss as the training progresses. This suggests that the model is improving its predictions on the training data while still maintaining a good level of performance on new data.

In the context of this study, the loss function is determined after 13 epochs, which is the smallest MSE that is closest to the validation loss as presented in the loss function graph in Figure 3**.** During this training process, a total of 326 parameters (312 input parameters and 14 output parameters) are trained using a batch size of 10. The root mean square propagation (RMSprop) algorithm is used as the learning algorithm for the data at a learning rate of 0.01, and a momentum rate of 0.0. Thus, the optimal LSTM model after 13 epochs is further evaluated across several performance metrics.

Model evaluation

Evaluating the model involves using the validation set to evaluate the performance of the model and ensure that it is not overfitting to the training data (Gridin, 2021). This may involve calculating metrics such as accuracy, precision, recall, etc. To evaluate the performance of the LSTM model, we split all of the data into 2 sets: 90% of the data comprise the training set, while 10% the testing set (we choose this split since in this case, the end point of the training data is November 12, 2020, which means that the COVID-19 crisis information is contained in the data). As presented in Table 5, the optimal

LSTM model with the smallest MSE after 13 epochs is selected based on several performance evaluation criteria, including the AIC, BIC, ME, RMSE, MAE and MAPE.

Figure 3 LSTM Model Configuration

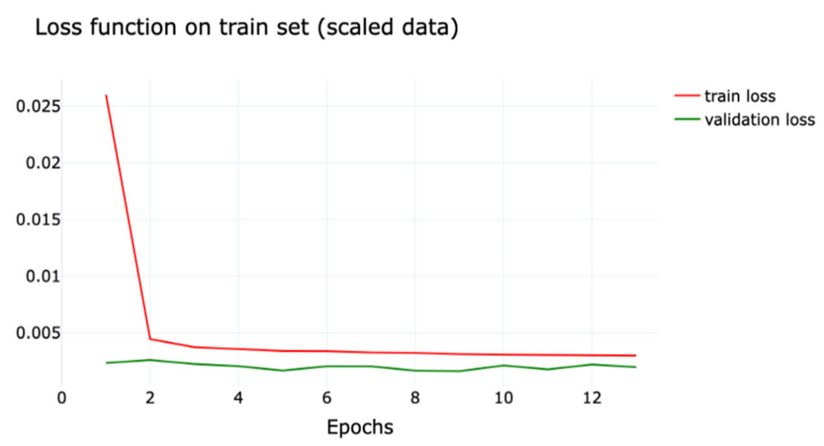

Evaluation of LSTM Model

Table 5 Performance Evaluation of LSTM Model

AIC	BIC	ME	RMSE	MAE	MAPE
8410.43	23,710.66	4.4	5.91	6.23	2.14

4. Results and Discussion of Findings

Forecasting is performed with two different models, including the ARIMA (1,1,3) as dictated by the AIC and BIC, as well as the LSTM model. To compare these forecasts between models, the first 90% of the real estate data set are reserved for training the models and the last 10% are used as the test data. This split allows for inclusion of the initial onset of the COVID-19 crisis within the training sample. Both models are trained with 90% of the data observations and subsequently tested with the remaining data. Thereafter, the residuals of both models are analyzed prior to comparing the forecasting performance of the evaluated model.

Model Evaluation using Residual Analysis

The residual analysis is a common method for evaluating the performance of time series models. The residuals are defined as the difference between the observed and predicted values from the model. A good time series model should have residuals that are not correlated and have constant variance and a zero mean. As presented in Figures 4.1 and 4.2, after selecting a potential model and estimating its parameters, diagnostic checks are performed with the basic assumption that the residuals have white noise - are randomly distributed, a mean of zero, and a constant variance. Furthermore, the Ljung-Box test is used to check whether the residuals are not correlated. Upon testing, our Ljung-Box statistics reveal that the chi-squared statistics=1.938, and the p-value = 0.7873, thus, the null hypothesis of white noise being present is rejected. Also, it can be observed that there is no residual correlation in the time series data set as all of the series are within the boundaries. This is further depicted in the normal Q-Q graph as all residuals approximately fall along the line. Therefore, this implies that both our ARIMA (1,1,3) and LSTM models are good and statistically fit for forecasting the real estate time series data.

Figure 4.1 Analysis of Residuals from ARIMA Model

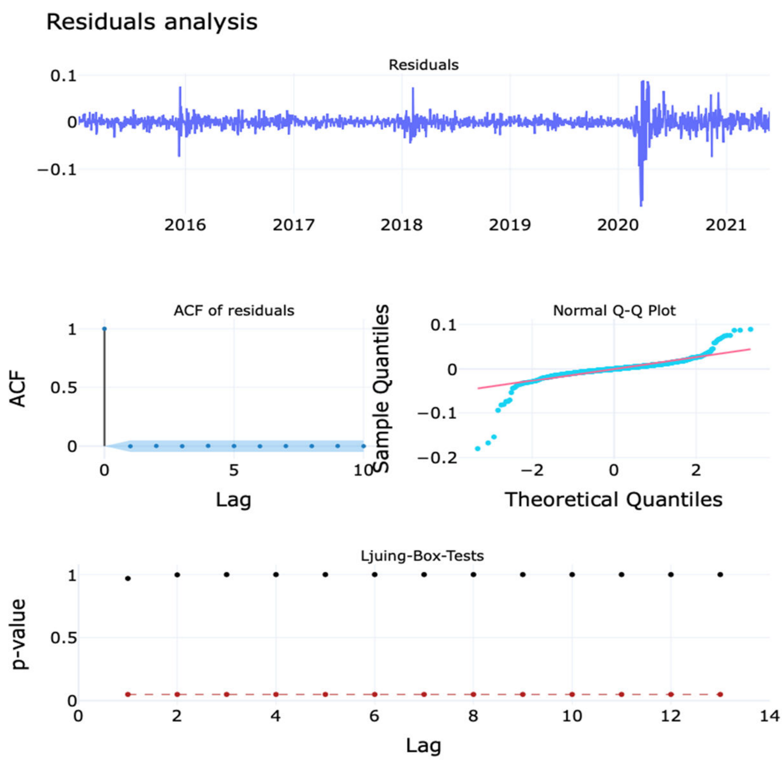

Figure 4.2 Analysis of Residuals from LSTM Model

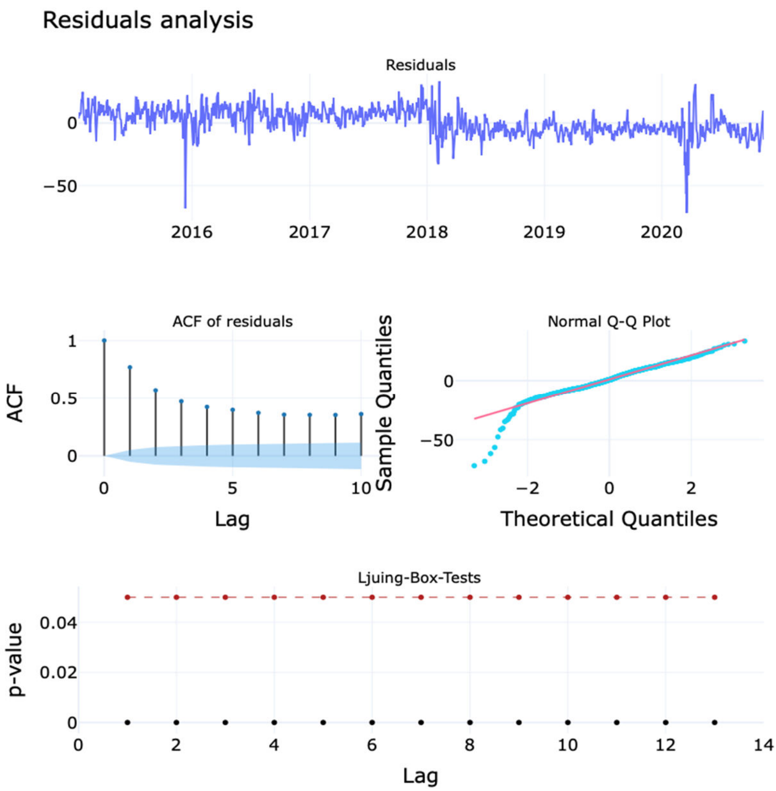

Forecasting Using Evaluated Models

Likewise, we evaluate and compare the ability of our optimal ARIMA (1,1,3) and LSTM models to forecast South African real estate during the COVID-19 crisis by using a variety of performance metrics including the MSE, RMSE, MAE and MAPE. The results of our forecasts are summarized in Table 6 and graphically presented in Figure 4.3.

Table 6 ARIMA vs. LSTM model evaluation with forecasting performance evaluation metrics

Model	MSE	RMSE	MAE	MAPE
ARIMA (1, 1, 3)	-0.224	6.371	4.333	0.881
LSTM	4.40	5.91	6.23	2.14
Reduction in error rates (LSTM-ARIMA) / LSTM	105.1%	-7.8%	30.4%	58.8%

Table 6 shows that for all of the metrics except the RMSE, the ARIMA (1,1,3) model performs more parsimoniously than the LSTM model, with values of -0.224, 4.333 and 0.881 for the MSE, MAE and MAPE respectively. At these values, the ARIMA model provides a 105.1%, 30.4%, and 58.8% reduction in error rate for the ME, MAE and MAPE respectively, when its forecasting performance is compared with those of the LSTM model. Conversely, the LSTM model performs better according to the RMSE metric, at a value of 5.91 compared to the 6.371 of the ARIMA model, or a 7.8% reduction in error rate by the LSTM compared to the ARIMA model. Based on 3 of the 4 metrics, it can be concluded that the ARIMA model outperforms the LSTM model at medium to high margins during the COVID-19 pandemic period.

Figure 4.3 Forecasting Using Evaluated Models

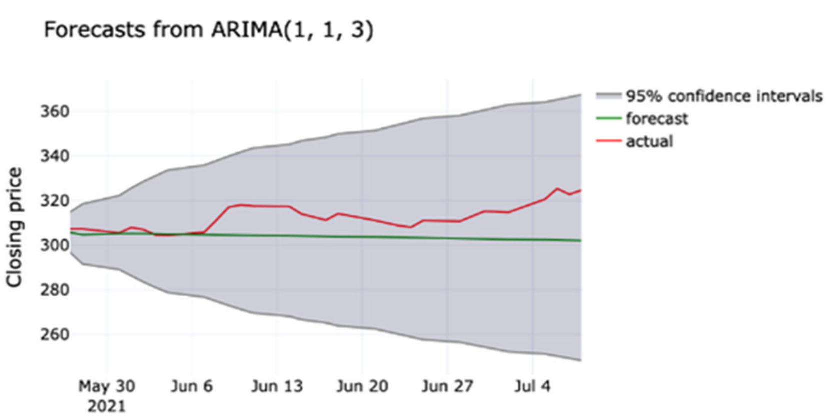

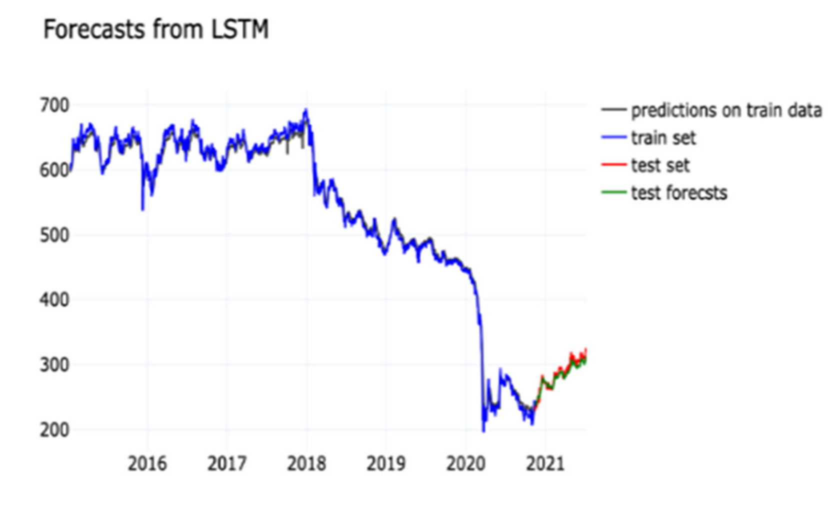

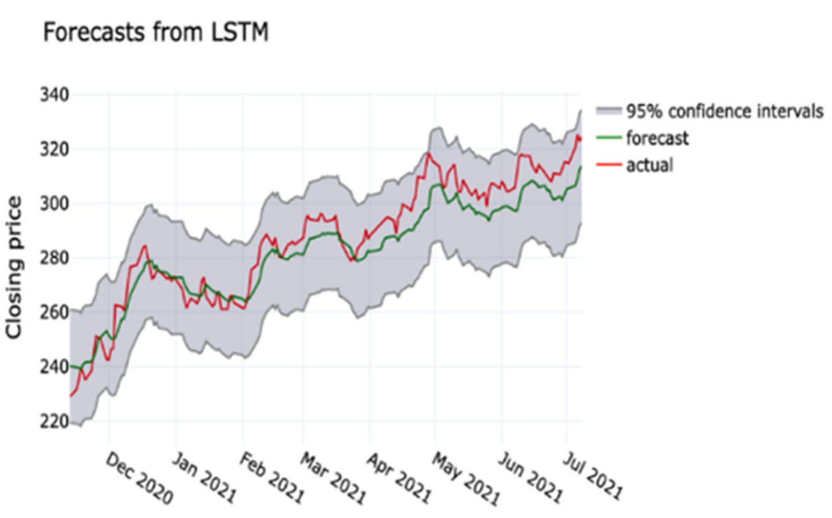

The study estimates real estate performance by using closing prices as the baseline. Estimations are repeated numerous times in order to reach realistic values. The data appear to have significant enough reductions in error rates when the ARIMA model is compared to the LSTM model, in order for the superiority of the former to be established over the latter for the data in question.

The superior performance of the ARIMA (1,1,3) model over the LSTM model in this study aligns with several empirical investigations that suggest traditional statistical models often outperform ML approaches under specific conditions. For example, Lee et al. (2008), Merh et al. (2010) and Makridakis et al. (2018) report that ARIMA models tend to deliver more accurate forecasts when applied to datasets with shorter time spans, less complexity, and predominantly linear characteristics. These subsisting conditions are present in this study, which favors a univariate, moderate-sized dataset during the relatively brief COVID-19 crisis period. In contrast, LSTM models are generally more effective in handling large, multivariate datasets with complex nonlinear patterns and longer temporal sequences, as evidenced by the findings of Adebiyi et al. (2014), Cerqueira et al. (2019) and Hopp (2021), among others.

However, the mixed conclusions in the literature suggest that model superiority is highly context-dependent and not generalizable across asset classes, sectors, or data structures. Cerqueira et al. (2019) emphasize the inherent difficulty of determining the precise volume of data needed for predictive tasks and note that optimal model performance is influenced by the interplay among data complexity, problem difficulty, and learning algorithm design. Further supporting this premise, studies from other South African sectors, such as Essa et al. (2021) and Zambezi (2021), indicate that LSTM models tend to outperform ARIMA models in contexts with greater data complexity and larger volumes of data. Accordingly, while the ARIMA model outperforms the LSTM model based on the specific conditions of this study, it is important to note that this may not hold across all forecasting scenarios or asset classes.

5. Conclusion

Based on the findings of our research, we do not find that the ML technique as represented by the LSTM model has more accurate predictive capabilities, despite expectations to the contrary from the literature and empirical studies, particularly given the size of the data set in this study. Our findings indicate that a traditional predictive technique as represented by the ARIMA model is more robust, and a better predictive technique for forecasting the South African real estate sector during the COVID-19 crisis. Hence, we conclude that traditional statistical techniques can better forecast the South African real estate sector during crises/ unforeseen shocks such as the COVID-19 pandemic due to their robust nature. This finding indicates that the traditional econometric models still have a role to play in forecasting financial data, despite the recent development of ML models.

Given the findings in related work, however, we propose that the findings of the study may have been different, given a larger data size or more complex (e.g. multivariate) data. It is therefore recommended that the study is repeated with a larger, more complex set of data, in order to stress-test the conclusion. Also, it would be interesting to compare the forecasting performance of hybrid models such as LSTM-ARIMA with non-hybrid models.

Lastly, a notable limitation of this study is that the LSTM model is implemented by using a single architecture configuration, without comprehensive hyperparameter tuning. This decision is guided by the exploratory objective nature of the study, and computational resource constraints, which require extensive tuning beyond the scope of this initial investigation. Moreover, similar approaches have been adopted in related studies, such as Adebiyi et al. (2014) and Essa et al. (2021), which also employ fixed LSTM configurations to benchmark performance against traditional models in financial and economic forecasting tasks. While basic adjustments such as modifying the dropout rate and number of training epochs through preliminary experimentation are implemented, it is necessary to emphasize that systematic optimization techniques, such as grid and random searches, or Bayesian optimization, are not applied. Hence, given that deep learning models are highly sensitive to hyperparameters, including the number of layers, neurons per layer, learning rate, batch size, and activation functions, alternative configurations may have produced improved forecasting accuracy. As such, the limited tuning in this study may have contributed to the underperformance of the LSTM model relative to the ARIMA model. Future research should apply robust hyperparameter optimization strategies to more fully exploit the predictive potential of LSTM models and explore the comparative performance of alternative deep learning architectures such as GRUs and CNN-LSTM hybrids, particularly in the context of real estate time series data.

Declaration of Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Adebiyi, A.A., Adewumi, A. and Ayo, C. (2014). Comparison of ARIMA and artificial neural networks models for stock price prediction. *Journal of Applied Mathematics*, (1): 1-7. https://doi.org/10.1155/2014/614342

Aladag, C.H., Egrioglu, E. and Kadilar, C. (2009). Forecasting nonlinear time series with a hybrid methodology. *Applied Mathematics Letters,* 22(9): 1467-1470. https://doi.org/10.1016/j.aml.2009.02.006

Alola, A.A. (2021). Evidence of speculative bubbles and regime switch in real estate market and crude oil price: Insight from Saudi Arabia. *International Journal of Finance & Economics* 26(3): 3473-3483. https://doi.org/10.1002/ijfe.1971

Bianchi, F.M., Maiorino, E., Kampffmeyer, M.C., Rizzi, A., and Jenssen, R. (2017). *Recurrent Neural Networks for Short-Term Load Forecasting: An Overview and Comparative Analysis.* Springer Publishing Company, Incorporated.

Brownlee, J. (2018). *Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs and LSTMs in Python.* Machine Learning Mastery.

Castro, R., Pineda, I., Lim W, Morocho-Cayamcela, M.E. (2022). Deep Learning Approaches Based on Transformer Architectures for Image Captioning Tasks. *IEEE Access,* 10: 33679-33694.

Cerqueira, V., Torgo, L. and Soares, C. (2019). Machine Learning vs Statistical Methods for Time Series Forecasting: Size Matters. *arXiv preprint arXiv:1909.13316*. https://doi.org/10.48550/arXiv.1909.13316

Essa, Y., Hunt, H.G. and Ajoodha, R. (2021). 'Short-term Prediction of Lightning in Southern Africa using Autoregressive Machine Learning Techniques', *2021 IEEE International IOT, Electronics and Mechatronics Conference (IEMTRONICS).* 2021 April. IEEE, pp. 1-5.

Gridin, I. (2021). *Time Series Forecasting using Deep Learning: Combining PyTorch, RNN, TCN, and Deep Neural Network Models to Provide Production-Ready Prediction Solutions (English Edition).* BPB Publications.

Grybauskas, A., Pilinkienė, V. and Stundžienė, A. (2021). Predictive analytics using Big Data for the real estate market during the COVID-19 pandemic. *Journal of Big Data*, 8(1): 1-20.

Guidolin, M. and Pedio, M. (2018). *Essentials of Time Series for Financial Applications.* Academic Press.

Hamilton, J.D. (2020). *Time Series Analysis.* Princeton University Press.

Hopp, D. (2021). Performance of LSTM Neural Networks in Nowcasting during the COVID-19 Crisis. *UNCTAD Research Paper No. 74.* Available at: https://unctad.org/system/files/official-document/ser-rp-2021d17_en.pdf

Hopp, D. (2022). Performance of long short-term memory artificial neural networks in nowcasting during the COVID-19 crisis. *arXiv preprint arXiv:2203.11872*.

Jadon, S., Kanty, J. and Patnakar, A. (2021). Challenges and Approaches to Time series forecasting: A Survey. *arXiv preprint arXiv:2101.04224*.

Kazemzadeh, M.-R, Amjadian, A. and Amraee, T. (2020). A hybrid data mining driven algorithm for long term electric peak load and energy demand forecasting. *Energy,* 204(6).

Korstanje, J. (2021). *Advanced Forecasting with Python With State-of-the-Art-Models Including LSTMs, Facebook's Prophet, and Amazon's DeepAR.* Apress.

Lazzeri, F. (2020). *Machine Learning for Time Series Forecasting with Python.* John Wiley & Sons.

Lee, K.J., Chi, A.Y., Yoo, S., and Jin, J.J. (2008). FORECASTING KOREAN STOCK PRICE INDEX (KOSPI) USING BACK PROPAGATION NEURAL NETWORK MODEL, BAYESIAN CHIAO'S MODEL, AND SARIMA MODEL. *Academy of Information and Management Sciences Journal*, 12(2): 32-35.

Makridakis, S., Spiliotis, E. and Assimakopoulos, V. (2018). Statistical and Machine Learning forecasting methods: Concerns and ways forward. *PloS ONE* 13(3): e0194889.https://doi.org/10.1371/journal.pone.0194889

Merh, N., Saxena, V.P. and Pardasani, K.R. (2010). A comparison between Hybrid Approaches of ANN and ARIMA for Indian Stock Trend Forecasting. *Business Intelligence Journal,* 3(2): 23-43.

Miswan, N.H., Ngatiman, N.A., Hamzah, K., and Zamzamin, Z. (2014). Comparative performance of ARIMA and GARCH models in modelling and forecasting volatility of Malaysia market properties and shares. *Applied Mathematical Sciences*, 8(137): 7001-7012. https://doi.org/10.12988/ams.2014.47548

Nassar, L., Okwuchi, I.E., Saad, M., Karray, F. (2020). 'Deep Learning Based Approach for Fresh Produce Market Price Prediction', *2020 International Joint Conference on Neural Networks (IJCNN).* IEEE, 1-7.

Nazlioglu, S., Gormus, N.A. and Soytas, U. (2016). Oil prices and real estate investment trusts (REITs): Gradual-shift causality and volatility transmission analysis. *Energy Economics,* 60: 168-175. https://doi.org/10.1016/j.eneco.2016.09.009

Nielsen, A. (2019). *Practical Time Series Analysis: Prediction with Statistics and Machine Learning.* O'Reilly Media.

Olson, D.L. and Wu, D. (2017). *Predictive Data Mining Models*. Springer.

Pal, A. and Prakash, P. (2017). *Practical Time Series Analysis: Master Time Series Data Processing, Visualization, and Modeling using Python.* Packt Publishing Ltd.

Rhif, M., Ben Abbes, A., Farah, I.R., Martinez, B., and Sang, Y. (2019). Wavelet Transform Application for/in Non-Stationary Time-Series Analysis: A Review. *Applied Sciences,* 9(7): 1345. https://doi.org/10.3390/app9071345

Rosinus, M. (2025). Comparison of Classical Arima Forecasting Methods to the Machine Learning LSTM Method: a Case Study on DAX® 50 ESG Index, *International Acta VŠFS,* 19(1): 32-52. RePEc:prf:journl:v:19:y:2025:i:1:p:32-52

Samadani, S. and Costa, C.J. (2021). 'Forecasting real estate prices in Portugal: A data science approach', *2021 16th Iberian Conference on Information Systems and Technologies (CISTI).* IEEE, 1-6.

Saravagi, D., Agrawal, S. and Saravagi, M. (2021). Indian stock market analysis and prediction using LSTM model during COVID-19. *International Journal of Engineering Systems Modelling and Simulation,* 12(2-3): 139-147.

Siami-Namini, S., and Namin A.S. (2018). Forecasting economics and financial time series: ARIMA vs. LSTM. https://doi.org/10.48550/arXiv.1803.06386

Siami-Namini, S., Tavakoli, N. and Namin, A.S. (2018). 'A comparison of ARIMA and LSTM in forecasting time series', *2018 17th IEEE international conference on machine learning and applications (ICMLA).* IEEE, 1394-1401.

Siami-Namini, S., Tavakoli, N. and Namin, A.S. (2019). A Comparative Analysis of Forecasting Financial Time Series Using ARIMA, LSTM, and BiLSTM. *arXiv preprint arXiv:1911.09512*.

Stevenson, S. (2007). A comparison of the forecasting ability of ARIMA models. *Journal of Property Investment & Finance*, 25(3): 223-240. https://doi.org/10.1108/14635780710746902?urlappend=%3Futm_source%3Dresearchgate

Temür, A.S., Akgün, M., and Temür, G. (2019). Predicting housing sales in Turkey using ARIMA, LSTM and hybrid models. *Journal of Business Economics and Management,* 20(5): 920-938. https://doi.org/10.3846/jbem.2019.10190.

Vishwas, B. and Patel, A. (2020). *Hands-on Time Series Analysis with Python.* Springer.

Wang, Q., Li, S. and Li, R. (2018). Forecasting energy demand in China and India: Using single-linear, hybrid-linear, and non-linear time series forecast techniques. *Energy,* 161: 821-831. https://doi.org/10.1016/j.energy.2018.07.168

Zambezi, S. (2021). *Predicting social unrest events in South Africa using LSTM neural networks*. Master thesis. University of Cape Town.

Zhang, Q., Zhang, P. and Zhou, F. (2022). Intraday and interday features in the high-frequency data: Pre-and post-Crisis evidence in China's stock market. *Expert Systems with Applications*, 209: 118321. https://doi.org/10.1016/j.eswa.2022.118321

Zheng, M., Song, H.S. and Liang, J. (2024). Modeling the Volatility of Daily Listed Real Estate Returns during Economic Crises: Evidence from Generalized Autoregressive Conditional Heteroscedasticity Models. *Buildings,* 14(1):182. https://doi.org/10.3390/buildings14010182