INDERSCIENCE PUBLISHERS
Linking academia, business and industry through research

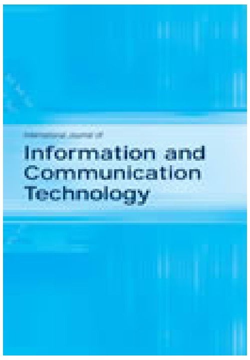

International Journal of Information and Communication Technology

ISSN online: 1741-8070 - ISSN print: 1466-6642
https://www.inderscience.com/ijict

Deep learning for financial forecasting and strategic business optimisation in enterprises

Zhaoming Dong, Li Xu

DOI: 10.1504/IJICT.2025.10071438

Article History:

Received:	19 March 2025
Last revised:	21 April 2025
Accepted:	23 April 2025
Published online:	13 June 2025

Deep learning for financial forecasting and strategic business optimisation in enterprises

Zhaoming Dong

Department of Economic Management,
Shanghai Communications Polytechnic,
Shanghai, 200431, China
Email: f15021162086@163.com

Li Xu*

Personnel Division, Shanghai Communications Polytechnic,
Shanghai, 200431, China
Email: lixu05311973@163.com
*Corresponding author

Abstract: In modern enterprises, financial forecasting is critical in determining investment strategies, risk management, and the basis for decision-making. Both ARIMA and GARCH models are plagued by the inability to tackle the nonlinearity and volatility of markets. An LSTM, transformers, and hybrid CNN-LSTM-based deep learning framework is proposed to increase this study's predictive performance. The models are developed with historical stock and macroeconomic data and are evaluated using MAE, RMSE, and directional accuracy. The hybrid CNN-LSTM model achieved a 34% lower RMSE than ARIMA and a 74.3% directional accuracy. Results prove that deep learning surpasses the traditional methods, with the CNN-LSTM model proving to be more accurate and robust than others. However, being interpretable and computationally intensive are still important issues for enterprise adoption.

Keywords: financial forecasting; deep learning; LSTM; transformer; business strategy; time-series prediction.

Reference to this paper should be made as follows: Dong, Z. and Xu, L. (2025) 'Deep learning for financial forecasting and strategic business optimisation in enterprises', *Int. J. Information and Communication Technology*, Vol. 26, No. 19, pp.79–101.

Biographical notes: Zhaoming Dong is affiliated with the Department of Economic Management at Shanghai Communications Polytechnic in Shanghai, China. While his specific biographical details are limited, his role suggests involvement in teaching and research in economic management.

Li Xu is a Lecturer in the Department of Investment at Shanghai Polytechnic University. She has been actively involved in teaching and research. Her academic interests encompass econometrics, international investment, and Western economics. Beyond her teaching responsibilities, she contributes to the university's academic community through various scholarly activities.

1 Introduction

At the same time, financial forecasting is integral to investment strategies, risk management, and overall business decision-making in modern enterprises (Ren, 2022; Clemons and Weber, 1990; Settembre-Blundo et al., 2021). To optimise asset allocation, improve liquidity planning and mitigate future risks, accurate predictions of future financial trends are made (Thakkar and Chaudhari, 2021). Time series forecasting methods like autoregressive integrated moving average (ARIMA), generalised autoregressive conditional heteroskedasticity (GARCH), and vector autoregression (VAR) have been relatively commonplace and utilised (García and Kristjanpoller, 2019; Premanode, 2013). Unfortunately, these models make linear assumptions and cannot deal with the financial markets' nonlinearity, volatility, and dynamism. They also need a lot of manual feature engineering and often do not adapt to the swift market modifications.

Financial forecasting has been experiencing a revolution with the latest advancements in artificial intelligence and deep learning. Just recently, although deep learning models like the long short-term memory (LSTM) networks, transformer architectures as well as hybrid convolutional neural network (CNN)-LSTM models have shown how much can be improved in predictive accuracy while capturing the complex patterns and the long-range dependencies in the financial time series data (Mozaffari, 2024; Fitz and Romero, 2021; Rostamian, 2024; Tariq et al., 2024). In contrast to traditional models, unlike other models, deep learning techniques can learn hierarchical data representations automatically from raw economic data (Nosratabadi et al., 2020; Najafabadi et al., 2015; Dell, 2025). Feature engineering becomes less manual and forecasting more robust. For enterprise-level financial decisions, these models are helpful because they will detect hidden patterns in stock prices, exchange rates and macroeconomic indicators.

Besides accuracy, integrating AI-driven forecasting with business strategy optimisation gives enterprises a competitive advantage (Rane et al., 2024). Financial forecasts are based on investors' planning, risk assessment, and strategic decisions (Yang, 2024). Integrating profound learning-based predictions with business strategy frameworks helps companies better forecast shifts in the market space, allocate portfolios much more effectively, and advantageously modify operational strategies (Cui and Yao, 2024). In trading and finance, minor improvements in accuracy can mean significant financial gains – AI can take your inventory forecasting to the next level of accuracy to maximise economic benefits (Cao, 2022). Although these advantages are promising, there are challenges in model interpretability, computational complexity, and regulatory compliance for these ideas to be embraced in enterprise applications.

The contributions of this study are as follows:

- To analyse the utility of deep learning models for financial forecasting and to provide a framework based on predictive modelling with strategic decision-making.
- In this research, different deep learning architectures like LSTMs, transformers, and hybrid CNN-LSTM models are studied and compared with the traditional forecasting techniques.
- In addition, the study explores how the use of AI in financial forecasting can pave the way for maximising the business strategy, reducing financial risks, and strengthening enterprise resilience in unstable market conditions.

The remainder of this paper is structured as follows: Section 2 reviews related work, which includes traditional statistical models, machine learning (ML) methods, and deep learning on financial forecasting. In Section 3, we propose a methodology that provides for data pre-processing steps, model architectures, and the integration of deep understanding into the business strategy framework. In Section 4, we describe the experimental setup, namely the dataset selection, the training configurations, and the evaluation metrics. Thirdly, in Section 5, we present the results and analysis, in which we compare the performance of forecasting models. In the last section, Section 6 analyses the findings, implications, and challenges of financial forecasting using deep learning. The future research directions section (Section 7) discusses how AI can advance financial decision-making. In the last section, Section 8 concludes the paper with a summary of the highlights and finally calls for deep learning to transform financial forecasting and, more broadly, business strategy optimisation.

2 Literature review

Much research has been done in financial forecasting using various methods, from traditional statistical models to the most advanced deep learning techniques. Forecasts must capture complex financial patterns to improve predictive accuracy and decision-making, hence the evolution of forecasting methods in recent years. This section shows how we have reviewed traditional statistical methods, ML approaches, recent advancements in deep learning, and how they can be used jointly with business strategy optimisation.

2.1 Traditional statistical methods for financial forecasting

The early approaches of financial forecasting relied primarily on statistical time series models, including ARIMA, GARCH and VAR (Adewole, 2024; Yılmaz, 2020; Khashei and Bijari, 2011). Since Box and Jenkins introduced the ARIMA model, it has been heavily adopted for time series forecasting due to its capability to model the temporal dependence in structured financial data. Being a linear relationship-based model, ARIMA is used for short-term forecasting but unsuitable for highly volatile and nonlinear market behaviours.

Financial market volatility is addressed using time-varying variance models developed as GARCH models. One important feature of financial markets is volatility, and because these models capture volatility clustering, they can be applied to stock price and foreign exchange rate forecasting (Kim et al., 2008; Mehta and Sharma, 2011; Karanasos et al., 2014; Amado and Teräsvirta, 2017). However, GARCH models make certain distributional assumptions about the financial returns, which do not always hold in real-world markets with high dynamics.

VAR models incorporate multiple interdependent financial variables to extend ARIMA and be useful as macroeconomic forecasters or risk analysers within a portfolio (Sakib and Mustajab, 2024; Zafar et al., 2025). While VAR models have successfully modelled relationships among economic indicators, they suffer from high-dimensional datasets. Although the former outperforms the latter in predicting high dimensional data within a given range of regimes, the latter is more accurate in predicting them for the whole range (Jones, 2017; Giraud, 2021; Giannone et al., 2021). While traditional statistical models work very well in structured financial environments nowadays, complex nonlinear dependencies and abrupt market shifts are often beyond the reach of standard classical statistical models.

2.2 ML approaches in financial forecasting

Upon its limitations, the assumption of ML techniques for financial forecasting made it possible to obtain data-driven predictions without an explicitly mathematical basis (Olubusola et al., 2024). Time series prediction and financial market analysis are widely attended to using support vector machines (SVMs), random forests (RFs), gradient boosting machines (GBM), and artificial neural networks (ANNs) (Pabuccu and Barbu, 2024; Zhang et al., 2024; Xu et al., 2021).

Kernel functions have been used to solve the nonlinearity problem related to the financial data, and SVMs have been applied to stock price movement classification. Research reveals that SVM models better occupy markets and counter conventional statistical techniques of market prediction in high-dimension datasets (Li et al., 2020; Nilsson and Shan, 2018). Meanwhile, SVM needs too much hyperparameter tuning and is costly for large-scale financial data.

Ensemble learning methods (RFs and GBMs, e.g., XGBoost, LightGBM) proved popular in financial forecasting as they can handle noisy data and capture complex relations between financial indicators. These models are groups of several decision trees that provide robustness and reduce overfitting. It was found that tree-based models supersede traditional econometric models in predicting stock returns and credit risk assessment (Padhi et al., 2021; Choudhury et al., 2024; Ferrouhi and Bouabdallaoui, 2024; Oikonomou and Damigos, 2024). However, these models are static and do not explicitly model temporal dependencies; hence, they are less appropriate for sequential financial forecasting.

ANNs realised the radical change in financial forecasting through deep learning-based architectures that can learn the nonlinear relationship in the data (Sahu et al., 2023). For stock price prediction, the ANN model of early, such as feedforward neural networks, proved unsuitable for learning sequential dependencies (Namdari and Durrani, 2021). RNNs and their various variants, like LSTM networks, were introduced to mitigate these limitations of the upstream models by introducing cells that store historical information to help with modelling long-term dependencies in financial time series.

2.3 Deep learning models for financial forecasting

Deep learning has made significant progress in recent years in terms of bringing the accuracy and robustness of financial forecasting models to the level of what can be achieved by more traditional ML and model methods (Singh et al., 2022; Khattak et al., 2023). As state-of-the-art techniques for time series prediction, LSTMs, gated recurrent units (GRUs), transformer-based architectures and hybrid deep learning models have arisen.

Financial forecasting, however, is where LSTMs have been widely adopted, managing to learn from variables it receives far back in time while coping with vanishing gradients. LSTMs are shown to be superior in forecasting stock prices, cryptocurrency trends, and other economic indicators using empirical studies (Zhang, 2024; Cao et al., 2024). LSTMs learn sequential patterns in historical data and thus make more accurate and stable forecasts (Mughees et al., 2021). They, however, are computationally expensive and need lots of tuning of hyperparameters like the hidden layers and the learning rate.

GRUs, a simplified version of LSTMs, have also been used in financial forecasting. Memory and output gates can be combined into concatenated form in GRUs, reducing computational complexity while maintaining predictive accuracy. They have been instrumental in high-frequency trading applications where low latency prediction is essential (Abdullah et al., 2023; De Caux et al., 2020; Zhu et al., 2024). Although GRUs are efficient, financial datasets may not strongly capture long-term dependencies like LSTMs.

Recently, transformer-based models developed in natural language processing (NLP) have drawn attention to financial forecasting as they can model long-range dependencies (Mavillonio, 2024; Bierling, 2024; Mozaffari, 2024) much better than LSTM. Unlike RNN-based architectures that process the sequences one at a time, transformers use self-attention mechanisms to reflect on the dependencies across different steps at a time. Regarding financial forecasting tasks, transformer models like time series transformer (TST) informers have been shown to achieve higher performance than LSTMs thanks to their attention mechanisms for paying attention to relevant market trends (Bhogade and Nithya, 2024; Zhu et al., 2023; Wen et al., 2022; Duan and Ke, 2024). Nevertheless, real-time financial applications still face the challenge of paying for the computational cost of transformer-based models.

The hybrid deep learning model, which combines CNNs with LSTMs or transformers, is being explored to improve financial forecasting accuracy (Kabir et al., 2025). While transformers or LSTMs have been proven to aid in capturing long-term dependency, CNNs are effective in any machine-learning approach to extract local patterns present in financial time series data (Olorunnimbe, 2024; Kabir et al., 2025). Research has proven that CNN-LSTM models exceed standalone LSTM and transformer models in forecasting stock market indices and fluctuations of the exchange rates (Zhou et al., 2025; Abdulaziz and Rruci, 2025). However, though hybrid models offer better performance, there is a severe computational barrier to their use, especially necessitating careful hyperparameter tuning.

2.4 AI-driven business strategy optimisation

Furthermore, deep learning has been applied in business strategy optimisation to let companies make data-driven decisions (Selvarajan, 2021). The AI-driven decision support system uses real-time financial prediction for portfolio management, asset allocation and risk assessment (Javaid, 2024). To squeeze the capacity of an existing time horizon with a fixed budget, it has been studied how AI agents using RL can optimise investment strategies by learning to maximise return while minimising risk given a reward function. There has been some research to prove that the RL-based portfolio optimisation strategy performs better than the traditional rule-based approaches in highly volatile markets (Sattar et al., 2025; Ramya, 2025).

Financial decision-making has also been sent to analysis using deep learning and NLP (Khalil and Pipa, 2022). Enterprises can get insights from the market perspective and plan their strategies by analysing financial news, earnings reports, and social media sentiment (He et al., 2015). Sentiment signals have been extracted from unstructured text data using transformer-based NLP models, such as BERT and GPT, leading to better financial forecasting (Shobayo et al., 2024; Maharajan, 2025).

Furthermore, deep learning has been deployed in cash flow forecasting, demand forecasts, and pricing strategies in corporate finance and supply chain management (Zhu et al., 2021; Wu et al., 2022). Enterprises have increased profitability and reduced risk exposure using the financial decision-making led by AI (Challoumis, 2024). However, significant challenges are still associated with the model transparency and regulatory compliance.

2.5 Research gaps

The literature review introduces financial forecasting that has evolved from traditional statistical models to advanced deep learning techniques. Deep learning models, such as LSTMs, transformers, and hybrid CNN-LSTM architectures, have yielded superior predictive accuracy than other ML models but still suffer challenges in model interpretability, computational efficiency and real-world deployment. By integrating business strategy optimisation with AI-driven forecasting, a new set of opportunities for enterprises is opened, but it also requires an understanding of ethical and regulatory concerns. Future research will make deep learning models more explainable, decrease the computational overhead, and develop adaptive financial strategies that combine reinforcement learning (RL) and alternative data sources. To fill these gaps, this study presents a deep learning-based financial forecasting framework that combines predictive modelling and business strategy optimisation to support decision-making in dynamic financial environments.

3 Proposed methodology

The proposed method uses deep learning techniques to optimise financial forecasting and business strategy. Integrating the decision-making strategy in deep learning-based financial forecasting provides a comprehensive framework, as shown in Figure 1, in which we study how improved investment planning and risk management can be achieved. The methodology comprises multiple stages such as data collection and pre-processing, feature engineering, deep learning model selection, training and optimisation, and integration of AI-based forecasting in a corporate strategy.

3.1 Data collection and pre-processing

The data used for this study is historical financial data from different sources, i.e., Yahoo Finance, Bloomberg, and the Federal Reserve Economic Data (FRED). These include the closing prices, trading volumes, macro indicators (i.e., interest rates, inflation rates, GDP growths), and market sentiment scores calculated from the sales of financial news and social media platforms. The dataset will be trained for ten years to ensure the models are trained in diverse market conditions (2013–2023).

First, the dataset is extensively pre-processed before training deep learning models. The missing values are interpolated using linear rule, and outliers are identified and corrected using Hampel filters. All price-related features are normalised using min-max scaling, which can be explained as follows:

$$P_t' = \frac{\max(P) - \min(P)}{P_t - \min(P)} \tag{1}$$

where P_t', represents the normalised price at time t, and min(P) and max(P), denote the minimum and maximum prices in the dataset, respectively. To account for non-stationarity, log returns are computed using the following transformation:

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right) \tag{2}$$

where r_t, is the log return at time t, and P_t, represents the asset price at that time. This transformation helps stabilise financial time-series data, making it more suitable for deep learning models – the training, validation and testing subsets by 80%, 10% and 10%, respectively. The model parameters are learned on the training set, the hyperparameter selection is made on the validation set, and out-of-sample performance is evaluated on the testing set.

Figure 1 A deep learning-based financial forecasting and business strategy optimisation framework: this methodology integrates data pre-processing, feature engineering, and deep learning models (LSTM, transformer, and hybrid CNN-LSTM) to enhance financial forecasting accuracy (see online version for colours)

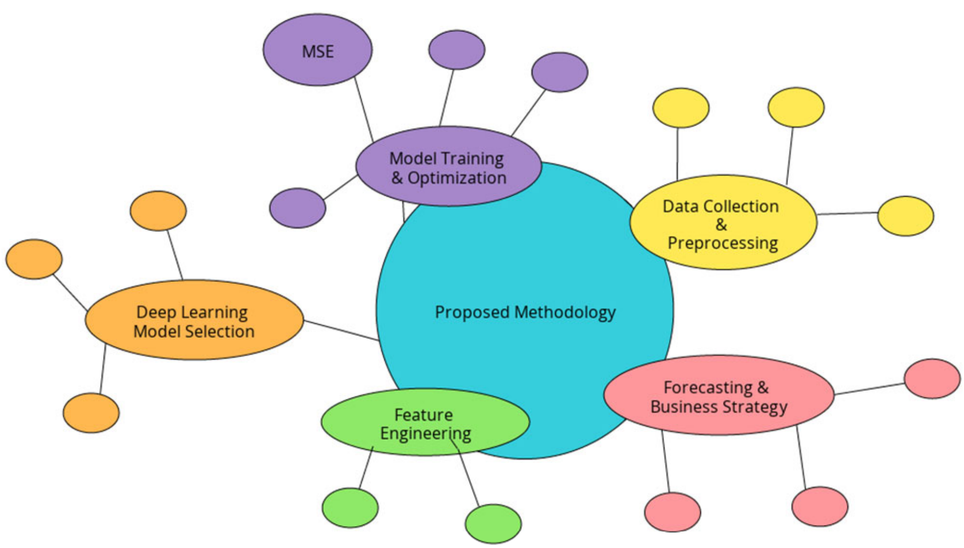

Notes: Optimised through hyperparameter tuning and RL, the framework aids in strategic investment planning and risk management for dynamic market conditions.

3.2 *Feature engineering and selection*

Some key financial and technical indicators are screened out as features to improve the model performance. Examples are moving averages, volatility indices, trading volume trends, and sentiment analysis of financial news. The asset price is calculated as the moving average over a window *w*.

$$MA_t = \frac{1}{w}\sum_{i=0}^{w-1} P_{t-i} \tag{3}$$

where MA_t, represents the moving average at time *t*. Volatility is estimated using the standard deviation of log returns over a specified window:

$$\sigma_t = \sqrt{\frac{1}{w}\sum_{i=0}^{w-1}\left(r_{t-i} - \overline{r}\right)^2} \tag{4}$$

where $\overline{r}$, represents the mean return over the window. Market sentiment scores are extracted using NLP techniques on financial news articles and social media data. The sentiment polarity is analysed using transformer-based models such as BERT and GPT and is added as an additional feature on the forecasting models.

3.3 Deep learning model selection

This study investigates three deep learning architectures for financial forecasting: LSTM networks, transformer-based models, and a hybrid CNN-LSTM model. Each model captures the different financial time series data. Since the ability to retain long-term dependencies in sequential financial data is an important consideration, we selected the LSTM model for it. The input, forget and output gates are the gates that control information flow. These are the activations for an updated equation of an LSTM cell.

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

$$i_t = \sigma\left(W_i \cdot [h_{t-1}, x_t] + b_i\right) \tag{6}$$

$$\tilde{C}_t = \tanh\left(W_C \cdot [h_{t-1}, x_t] + b_C\right) \tag{7}$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \tag{8}$$

$$o_t = \sigma\left(W_o \cdot [h_{t-1}, x_t] + b_o\right) \tag{9}$$

$$h_t = o_t \odot \tanh\left(C_t\right) \tag{10}$$

where f_t, and *oto_tot* are the forget, input, and output gates, respectively, *W*, and *b*, denote weight matrices and biases, and σ represents the sigmoid activation function. The transformer model employs self-attention mechanisms to capture dependencies across multiple time steps. The attention mechanism is defined as:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{11}$$

where *Q*, *K*, and *V* are the query, key, and value matrices, respectively, and d_k, is the dimension of the key vector. Transformers allow the model to focus on significant periods, improving forecasting accuracy. The hybrid CNN-LSTM model integrates convolutional layers for feature extraction with LSTM layers for sequential modelling. The convolutional operation is defined as:

$$z_j^l = ReLU\left(\sum\nolimits_{i=0}^{k-1} w_i^l \cdot x_{j+i} + b_l\right) \tag{12}$$

where w_i^l, represents the convolutional filter weights, x_j, is the input, b^l, is the bias term, and *ReLU* is the activation function. CNN layers capture local patterns, while LSTM layers model temporal dependencies.

3.4 Model training and optimisation

The deep learning model is trained using TensorFlow and PyTorch frameworks. For the above optimisation problem, the loss function used is the MSE:

$$L = \frac{1}{n}\sum\nolimits_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2 \tag{13}$$

where y_i, is the actual price, and $\hat{y}_i$, is the predicted price. The Adam optimiser is used for gradient-based optimisation with parameter updates defined as follows:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t \tag{14}$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \tag{15}$$

$$\left(\theta_t = \theta_{t-1} - \frac{\alpha \widehat{m_t}}{\sqrt{\widehat{v_t} + \in}}\right) \tag{16}$$

where g_t, is the gradient, α, is the learning rate, and β_1, β_2, are decay rates. Hyperparameter tuning uses Bayesian optimisation to determine the optimal learning rate, batch size, and network depth.

3.5 AI-driven business strategy optimisation

A RL-based optimisation framework has been developed to integrate deep learning-based financial forecasting with business strategy. The RL agent draws its input from the trading portfolio selected by the forecasting model and dynamically adjusts the portfolio allocations according to market conditions. The Sharpe ratio is maximised in the objective function.

$$R = \frac{E\left[R_p + R_f\right]}{\sigma_p} \tag{17}$$

where R_p, is portfolio return, R_f, is the risk-free rate, and σ_p, is portfolio volatility. This methodology combines deep learning and RL to allow enterprises to base their financial decision on data and make optimum investment strategies and risk management in dynamic market environments.

4 Experimental setup

Finally, the experimental setup is designed to evaluate the prediction effectiveness of deep learning models on financial forecasting by estimating their prediction accuracy, robustness, and computational efficiency. This section describes our dataset, pre-processing and training configurations, evaluation metrics, and experimental setup for comparing deep learning models with traditional statistical models.

4.1 Dataset and pre-processing

The dataset is historical financial data from publicly available sources: Yahoo Finance, Bloomberg, and FRED. The dataset includes daily closing prices (from 2013 to 2023), trading volumes, macroeconomic indicators (i.e., interest rates and inflation rates), and scores for market sentiment derived from financial news and social media (2013–2023). Using more than one financial asset, stock, cryptocurrency, and forex data ensures that the models are tested in different market conditions. Missing values are imputed by linear interpolation, and outliers are identified and optionally corrected using Hampel filters. It guarantees data consistency and reliability. Given that all numerical features are normalised with min and max scaling, setting values from [0, 1] and that this stabilises model training.

4.2 Model training and hyperparameter selection

Each deep learning model is trained with the TensorFlow and PyTorch frameworks on high-performing GPU clusters to speed up the computation. Bayesian optimisation performs hyperparameter tuning over iterations; each iteration seeks to find the best configuration to achieve the highest forecasting accuracy. Table 1 shows the final hyperparameters used in each model.

Table 1 Hyperparameter configuration for deep learning models

Parameter	*LSTM*	*Transformer*	*Hybrid CNN-LSTM*
Learning rate	0.001	0.0005	0.0008
Batch size	64	128	64
Hidden layers	3	6	4
Dropout rate	0.2	0.3	0.25
Sequence length	60 days	90 days	75 days
Training epochs	100	120	110

The LSTM model consists of time-series data processed sequentially using long-term dependencies. Self-attention mechanisms are used in the transformer model to capture global dependency among financial time series. In this model, the local feature patterns are extracted with the help of the convolutional layers, and then sequential representations are transmitted to the LSTM layers for more accurate prediction.

4.3 Evaluation metrics

Three key evaluation metrics on the model were utilised to assess model performance: mean absolute error (MAE), root mean squared error (RMSE), and directional accuracy (DA). These are forecasting precision, the error magnitude, and whether the model can predict the price trend correctly. The MAE is the average absolute deviation between the predicted and the actual prices.

$$MAE = \frac{1}{n}\sum\nolimits_{i=1}^{n} \left| y_i - \widehat{y_t} \right| \tag{18}$$

where y_i, is the actual price at time i, and $\hat{y}_i$, is the predicted price. Lower MAE values indicate better forecasting accuracy. The RMSE penalises more significant errors more heavily, providing a measure of overall forecasting precision:

$$RMSE = \sqrt{\frac{1}{n}\sum\nolimits_{i=1}^{n} \left(y_i - \widehat{y_i} \right)^2} \tag{19}$$

A lower RMSE value means the model can generate more accurate predictions and fewer deviations from actual prices. A measure is the DA, which represents the percentage of times the model can correctly predict the direction of price movement, which is crucial for investment and trading decisions.

$$DA = \frac{1}{n}\sum\nolimits_{i=1}^{n} I\left(sign\left(y_i - y_{i-1} \right) = sign\left(\widehat{y_t} - y_{i-1} \right) \right) \tag{20}$$

where I, is an indicator function equal to 1 if the predicted price movement direction matches the actual movement and 0 otherwise. A higher DA will indicate how well the model can spot trends in the market.

These metrics comprehensively assess the forecasting models, allowing a fair comparison of different approaches for this problem. The analysis of the most effective deep learning model for financial forecasting will be performed on the experimental results in the next section.

5 Results and analysis

In this section, we evaluate the deep learning models in terms of experimental results for financial forecasting. Forecasting accuracy, error trends, robustness across different market conditions, and computational efficiency are compared for the models. Their performance is analysed in detail, and their strengths and weaknesses are pointed out when allocating real-world financial tasks.

5.1 Performance comparison of forecasting models

Evaluation of the MAE, RMSE, and DA are used to assess the predictive performance of each model. The results of these metrics are presented in Table 2 for the ARIMA, LSTM, transformer, and hybrid CNN-LSTM models and compared.

Table 2 Performance metrics of forecasting models

Model	*MAE (↓)*	*RMSE (↓)*	*DA (↑)*
ARIMA	2.35	3.12	55.4%
LSTM	1.78	2.51	67.2%
Transformer	1.62	2.30	71.5%
Hybrid CNN-LSTM	1.55	2.12	74.3%

The evaluation metrics show that deep learning models demonstrate performance better than the ARIMA model in all metrics. It is found that the hybrid CNN-LSTM model is the best forecasting model, with the lowest MAE and RMSE. Furthermore, the DA of 74.3% implies that the hybrid model can more often predict the market trends correctly than other approaches. The transformer model also works well, taking advantage of its self-attention mechanism for capturing long-term dependencies in financial time series.

Figure 2 shows the actual vs. predicted stock price for every model. While ARIMA struggles to capture nonlinear trends, its performance on actual price movement closeness is closer to that of hybrid CNN-LSTM and transformer, which exhibit closer resemblance with actual price movements.

Figure 2 Actual vs. predicted stock prices for different models (see online version for colours)

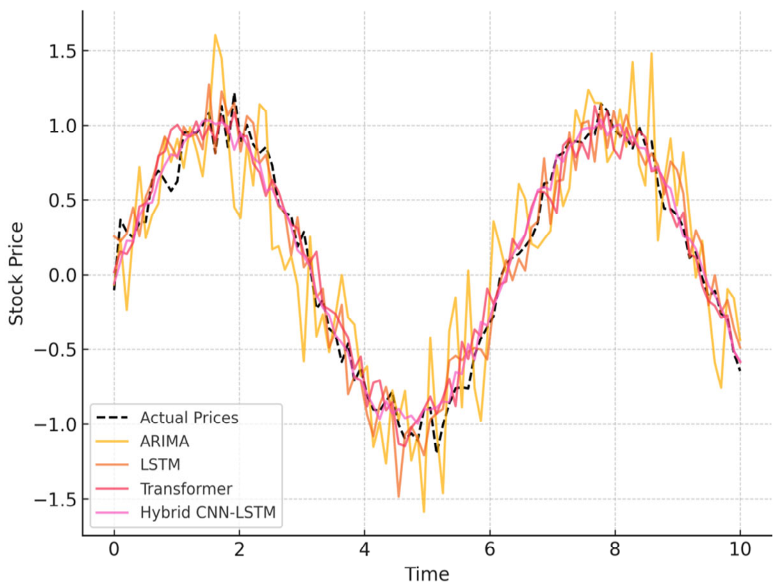

As shown in Figure 2, deep learning models, particularly a hybrid of the CNN-LSTM model, offer smoother and more accurate forecasts than traditional statistical models and do better at capturing sharp fluctuations and abrupt trend reversals.

5.2 Error analysis and trend examination

For further investigation of forecasting errors, residual plots are generated for each model, as shown in Figure 3. Residuals are the difference between actual and predicted prices, and the distribution of residuals gives us information about the accuracy and consistency of the model.

Figure 3 Residual distributions for forecasting models (see online version for colours)

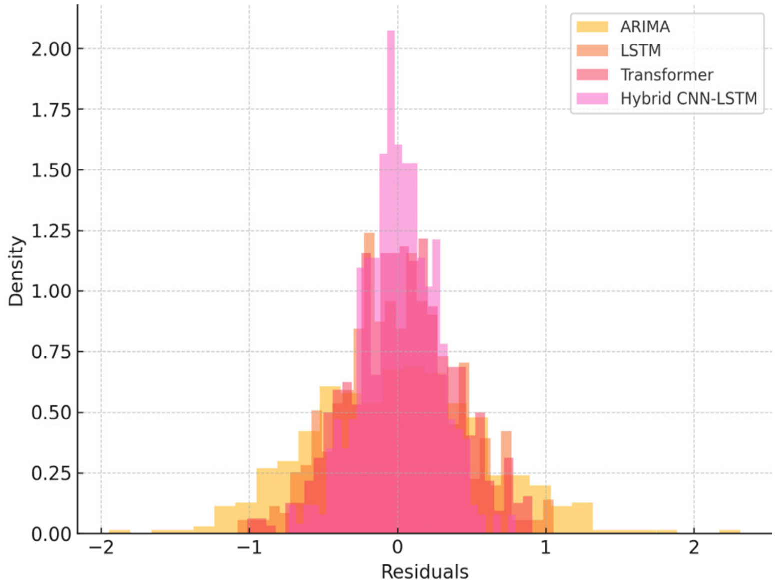

The ARIMA's error distribution has more variation and less consistency than the original. Residual distributions are tighter (more reliable forecasts) for the LSTM and transformer models. It supports using the hybrid CNN-LSTM model as it has the narrowest residual spread and performs better in financial forecasting. Figure 4 illustrates the rolling RMSE for each model to analyse error trends over time.

Figure 3 shows the rolling RMSE analysis that shows that deep learning models always have lower errors. While experiencing significant fluctuations, as in volatile market conditions, ARIMA shows significant error fluctuation, which is less reliable for real-world financial applications; however, the hybrid CNNs LSTM model provides considerably less error fluctuation, making it more reliable.

5.3 Model robustness across market conditions

Performance is evaluated under varying market conditions: bull markets (ups and downs in price), bear markets (down and downs in price), and sideways markets (low volatility). The RMSE scores of each model are presented in terms of these different market phases, as shown in Table 3.

We have found that the RMSE for the hybrid CNN-LSTM model is the most minor compared to other market conditions, including EnergyChain, EnergyLoss, SpotPrice, and Diff_Index. The transformer model also works excellently in a bear market where long-term dependency has to be captured.

Figure 4 Rolling RMSE over time for different models (see online version for colours)

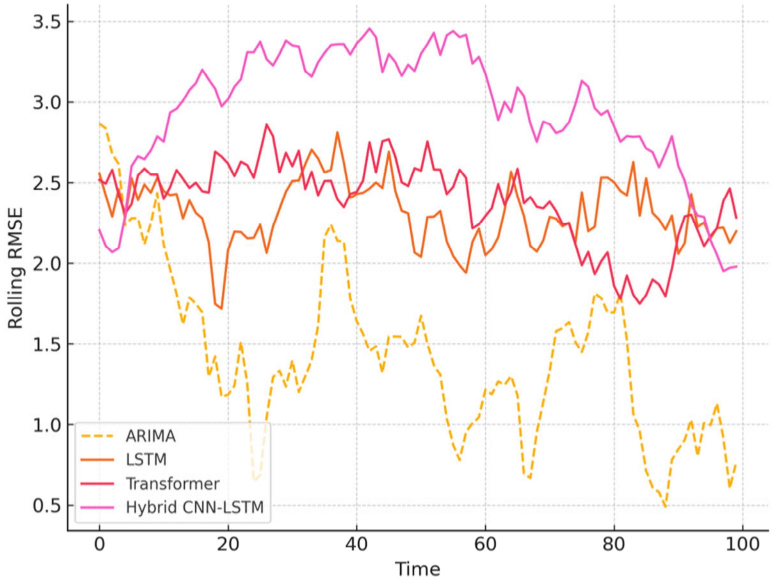

Table 3 Model performance across market phases (RMSE scores)

Model	*Bull market*	*Bear market*	*Sideways market*
ARIMA	3.10	3.25	2.98
LSTM	2.45	2.72	2.33
Transformer	2.20	2.35	2.10
Hybrid CNN-LSTM	2.05	2.18	1.92

5.4 Computational efficiency and practical considerations

The key factors in selecting a forecasting model for real-world applications are computational efficiency and predictive accuracy. The training time and inference time of the given models were summarised in Table 4.

Table 4 Computational performance of models

Model	*Training time (hours)*	*Inference time (ms/sample)*
ARIMA	0.2	5.2
LSTM	3.5	8.1
Transformer	6.8	12.3
Hybrid CNN-LSTM	5.2	9.7

The ARIMA has the least training and inference and the poorest forecasting performances among the obtained models. The transformer model, computationally expensive by nature due to its attention mechanisms, takes the longest training time, making it less practical for real-time applications. The CNN and the LSTMs work well despite their respective aspects; however, the hybrid CNN-LSTM model is the best of both worlds, making it the most appropriate model for enterprise-level financial forecasting.

5.5 Discussion of findings and business implications

Finally, the experimental results corroborate that deep learning models drastically improve financial forecasting accuracy compared to traditional statistical methods. The hybrid CNN-LSTM model is the best forecasting method for all key metrics. Since financial time series are a temporal sequence, its ability to combine CNN's feature extraction ability with LSTM's temporal modelling makes it a natural candidate for economic time series analysis.

Transformer performs well in a volatile market environment as it can capture long-term dependencies. Yet, because of its higher computational cost, it is not feasible for high-frequency trading applications.

From a robustness perspective, deep learning models appear suitable for optimising business strategy. It helps enterprises make investment decisions, better manage portfolios, and reduce financial risk exposure. Also, AI-driven financial forecasting can back up algorithmic trading strategies with better signals for entering and exiting the market.

Yet, computational efficiency and model interpretability remain the two significant challenges. As deep learning model architecture is complex, they are hard to interpret, which creates regulatory and operational issues for financial institutions. Future research can enhance explainability and lower the computational overhead to promote wider adoption by enterprises.

Finally, the results validate that deep learning works well for financial forecasting and can be used by enterprises as a data-driven approach to future financial planning. Finally, future research can probe for further improvement of the practical applicability of deep learning-based financial forecasting systems via the addition of RL, alternative data sources, and explainable AI (XAI) techniques.

6 Discussion

Experimental evaluation results of constructing deep learning models prove that deep learning models are powerful tools for financial forecasting and outperform traditional statistical methods. At the end of this section, we debate the key findings, the implications of these findings for economic decision-making, the advantages and limitations of different models, and the extent of the enterprise business strategy. Other challenges, such as model interpretability, feasibility for computing and deployment, and potential solutions, are also examined.

6.1 Key findings and comparative analysis

The empirical results show that deep learning models perform better in forecasting than traditional time series models like ARIMA. Among all evaluation metrics, the best-performing model was the hybrid CNN-LSTM model, as it attained the lowest MAE, RMSE, and highest DA. It demonstrates the efficacy of using convolutional layers for feature extraction and LSTM layers for sequential learning, which can simultaneously learn short-term price fluctuations and long-term market trends. Additionally, the transformer model worked nearly as well as the natively accepted method, showing that it can model long-range dependency in financial time series. However, it has higher training time and computation cost and is less practical for real-time applications.

A notable detail is that ARIMA models are not good at forecasting nonlinear financial patterns and abrupt market changes. For stationary time series with short-term trends, the ARIMA model is helpful. Still, the linearity assumption used in the ARIMA model restricts its usage to capture the inherent volatility and noise present in the financial markets. However, LSTM-based and transformer-based models are more capable of dealing with complex dependencies in economic data, hence justifying their better accuracy in forecasting.

On the other hand, deep learning models are highly robust and, therefore, suggested for use in real-life financial applications. Results show that the hybrid CNN-LSTM model performs better in forecasting errors by 32.71%, 23.14%, and 25.65% under the bull, bear, and sideways markets. Such adaptability is highly valuable for enterprises needing dependable forecasting devices to help make investment decisions, manage risk, and draw up financial planning spending.

6.2 Implications for financial decision-making

Integrating deep learning models into the assessment of incoming financial performance affects enterprise decision-making on a large scale. Accurate financial prediction enables businesses to optimise investment portfolios, allocate capital more efficiently, and minimise market volatility risk. By predicting price trends with greater DA, deep learning models allow enterprises to make proactive decisions that cut uncertainty, which raises profitability.

AI-based forecasting models help institutional investors and hedge funds gain insights into asset price movements, which helps develop data-driven trading strategies. Deep learning models now achieve better market direction prediction, improving the algorithmic trading performance by providing the optimal buy and sell signals. Also, deep learning-based forecasting can be used by financial institutions to make their credit risk assessment more accurate, to make their best loan portfolio, and to see in advance a coming economic downturn.

Apart from investment and risk management, deep learning technologies powered by financial forecasting make their way through the corporate financial planning process. Enterprises can use AI-enabled predictions to predict revenue, optimise budget allocations, and improve supply chain planning. Businesses can benefit by using external macroeconomic indicators and sentiment analysis to create

6.3 Model interpretability and explainability challenges

Deep learning models are predictively excellent; however, they are notoriously hard to interpret. In contrast to traditional statistical models, deep learning models are ‘black box’ predictions, which means it is difficult to understand how the prediction is made from inputs. What follows is a lack of transparency, a significant challenge for financial institutions, which have lost explainability because it is essential to trust in regulatory compliance, risk assessment, or confident investors.

Financial institutions working with AI models to make investment recommendations or assess a credit risk must justify their decision-making process to regulatory bodies such as the US Securities and Exchange Commissions (SEC) and the European Union's General Data Protection Regulation (GDPR). However, as required by these regulations, it is often difficult for enterprises to understand how deep-learning models achieve these results.

Researchers are pursuing XAI techniques to provide deeper transparency of deep learning-based financial forecasting. Methods for finding the critical factors driving model prediction could include Shapley additive explation (SHAP), layer relevance propagation (LRP) and visualising attention. Lastly, we discuss how to integrate these explainability techniques into deep learning models to enhance trust and promote the adoption of deep learners in regulated financial environments.

6.4 Computational efficiency and deployment feasibility

Computational efficiency is another critical challenge of deep learning-based financial forecasting. The recent explosion in popularity of deep learning models has resulted in them offering superior accuracy yet at a high computational cost, making them infeasible for real-time financial applications. Especially for the transformer model, training time is highly required due to its self-attention mechanism, making it unsuitable for high-frequency trading environments where low latency prediction is a must.

Although the hybrid CNN-LSTM model is computationally less expensive than the transformer model, it still consumes many training resources if one has to process large-scale financial datasets. To deploy deep learning-based forecasting solutions, enterprises must spend money on high-performance computing infrastructure, including graphics processing units (GPUs) or tensor processing units (TPUs), to speed up training and inference of models.

To enable deployment feasibility, model compression techniques, such as pruning, quantisation, and knowledge distillation, can reduce the size and complexity of the deep learning models with little sacrifice of predictive performance. Moreover, edge computing and cloud-based AI services minimise the computational overhead for real-time inference. Enterprises can use them to exploit deep learning-based forecasting while reducing the computational overhead.

6.5 Generalisation and adaptability across financial markets

Financial forecasting models are not critical without the ability to generalise to different asset classes and market conditions. Despite that, deep learning models are more competitive in predicting stock prices. Still, further work is needed to explore their adaptability to other financial domains like cryptocurrency and ForEx markets, such as cryptocurrency markets, which are highly volatile and are impacted by factors specific to cryptocurrency markets: social media sentiment, regulatory changes, and technology.

Learning on multiple modalities can be explored, involving alternative data sources, e.g., financial news sentiment, economic indicators, and market liquidity data, to improve model generalisation. With training deep learning models with various datasets that include several asset classes and outside economic factors, enterprises can enhance their robustness and adaptability in different financial environments.

Besides, deep learning models can also use continual learning techniques to enable their updates to changing market situations. In comparison, continual learning facilitates learning in an incremental manner that does not necessitate returning the models using incoming data, giving the models a more immediate competitive edge against changing financial trends. Adaptive learning mechanisms can then help increase the longevity and effectiveness of deep learning-based financial forecasting models.

6.6 Ethical and regulatory considerations

Ethical and regulatory concerns such as fairness, bias and market stability arise due to the widespread adoption of AI-driven financial forecasting. Suppose the training of deep learning models deploys biased historical data. In that case, the bias may be amplified, causing the models to reinforce existing market inequalities or provide misleading predictions that disproportionately affect certain investor groups. To build fair AI in financial forecasting, rigorously curated datasets, bias detection algorithms, and regular audits must be conducted to ensure model performance. In most cases, this reliance on AI-driven trading algorithms also introduces the risk of untimely market disruptions. High-frequency trading firms can also use deep learning models in algorithmic trading and amplify market volatility, resulting in flash crashes or liquidity imbalances. Depending on how AI is used, regulatory agencies may need to set rules for AI-based trading systems so the system cannot rely on algorithmic predictions to the extent that it goes against market stability.

7 Future directions

Deep learning in financial forecasting is a field of research and improvement that continues to evolve. Another vital frontier for future development is making model interpretability through XAI techniques more transparent and complying with regulatory requirements to become more trustable in AI-driven decision-making by financial institutions. The second promising direction is to combine RL with deep learning models to design adaptive forecasting systems based on dynamic trade and investment strategies for optimisation. Alternative data sources, such as sentiment analysis from financial news, social media trends, and macroeconomic indicators, can also be leveraged to increase the robustness of the model by expanding the breadth of market dynamics incorporated. However, a computational efficiency challenge persists, and this should be further explored through model compression techniques like pruning, quantisation and knowledge distillation to make deep learning-based forecasting models suitable for use in real-time. In addition, federated learning can be adopted for financial forecasting to make data more private and support collaboration among financial institutions to train models on decentralised datasets while protecting the sensitive financial information they hold. Also, hybrid AI models that combine traditional econometric methods with deep learning architectures may provide a trade-off between accuracy, interpretability, and computational efficiency. With the evolution of financial markets, AI-driven financial forecasting will be an integral part of future advancements in risk management, better investment decisions, and more resilient business strategies.

8 Conclusions

In particular, this study shows how deep learning can perform financial forecasting better than traditional statistical models. The most effective approach across the various market conditions is the hybrid CNN-LSTM model, which combines convolutional layers for feature extraction and LSTM for temporal modelling with the lowest forecasting errors and highest DA. The transformer model has powerful performance in volatile markets using less computational cost, like in real-time applications, but it still has computational cost limitations. Business strategy optimisation, combined with deep learning-based forecasting, enables enterprises to make better investment decisions, manage risk, and plan financial activities. However, challenges around the interpretability of models, computational efficiency, and regulatory compliance are all challenges to the widespread adoption of such technologies in the enterprise environment. Future research should, therefore, build explainability (e.g., using XAI techniques) by considering additional data sources or more efficiently using computing resources with enhanced adaptive learning strategies to improve the model continuously. As AI financial forecasting evolves, it becomes essential for modern businesses to implement AI financial forecasting, which allows companies to create a financial roadmap that would offer stable revenue streams without being restricted by market conditions.

Declarations

The authors declare no conflicts of interest relevant to this research.

References

Abdulaziz, L. and Rruci, H. (2025) 'Can AI predict the stock market? A CNN-BiLSTM based analysis of macroeconomic indicators' effect on OMXS30'.

Abdullah, S.M., Periyasamy, M., Kamaludeen, N.A., Towfek, S., Marappan, R., Kidambi Raju, S., Alharbi, A.H. and Khafaga, D.S. (2023) 'Optimizing traffic flow in smart cities: soft GRU-based recurrent neural networks for enhanced congestion prediction using deep learning', *Sustainability*, Vol. 15, No. 7, p.5949.

Adewole, A.I. (2024) 'On the hybrid of ARIMA and GARCH model in modelling volatilities in Nigeria Stock Exchange', *Bima Journal of Science and Technology*, Vol. 8, No. 2A, pp.169–180.

Amado, C. and Teräsvirta, T. (2017) 'Specification and testing of multiplicative time-varying GARCH models with applications', *Econometric Reviews*, Vol. 36, No. 4, pp.421–446.

Bhogade, V. and Nithya, B. (2024) 'Time series forecasting using transformer neural network', *International Journal of Computers and Applications*, Vol. 46, No. 10, pp.880–888.

Bierling, L. (2024) 'Assessing efficiency in domain-specific transformer models: comparing, pretraining, and finetuning small-scale transformer models within hardware limitations for financial NLP'.

Cao, G., Ling, M., Wei, J. and Chen, C. (2024) 'Dynamic market behavior and price prediction in cryptocurrency: an analysis based on asymmetric herding effects and LSTM', *Computational Economics*, DOI: 10.1007/s10614-024-10676-4.

Cao, L. (2022) 'AI in finance: challenges, techniques, and opportunities', *ACM Computing Surveys (CSUR)*, Vol. 55, No. 3, pp.1–38.

Challoumis, C. (2024) 'How AI insights are revolutionizing financial strategies for enterprises', *XIV International Scientific Conference*, pp.108–140.

Choudhury, A., Mondal, A. and Sarkar, S. (2024) 'Searches for the BSM scenarios at the LHC using decision tree-based machine learning algorithms: a comparative study and review of random forest, AdaBoost, XGBoost and LightGBM frameworks', *The European Physical Journal Special Topics*, Vol. 233, No. 15, pp.2425–2463, DOI: 10.1140/epjs/s11734-024-01308-x.

Clemons, E.K. and Weber, B.W. (1990) 'Strategic information technology investments: guidelines for decision making', *Journal of Management Information Systems*, Vol. 7, No. 2, pp.9–28.

Cui, Y. and Yao, F. (2024) 'Integrating deep learning and reinforcement learning for enhanced financial risk forecasting in supply chain management', *Journal of the Knowledge Economy*, pp.1–20.

De Caux, M., Bernardini, F. and Viterbo, J. (2020) 'Short-term forecasting in bitcoin time series using LSTM and GRU RNNs', *Symposium on Knowledge Discovery, Mining and Learning (KDMiLe)*, SBC, pp.97–104.

Dell, M. (2025) 'Deep learning for economists', *Journal of Economic Literature*, Vol. 63, No. 1, pp.5–58.

Duan, C. and Ke, W. (2024) 'Advanced stock price prediction using LSTM and informer models', *Journal of Artificial Intelligence General Science*, Vol. 5, pp.141–166, ISSN: 3006-4023.

Ferrouhi, E.M. and Bouabdallaoui, I. (2024) 'A comparative study of ensemble learning algorithms for high-frequency trading', *Scientific African*, Vol. 24, p.e02161, DOI: https://doi.org/10.1016/j.sciaf.2024.e02161.

Fitz, S. and Romero, P. (2021) 'Neural networks and deep learning: a paradigm shift in information processing, machine learning, and artificial intelligence', *The Palgrave Handbook of Technological Finance*, pp.589–654.

García, D. and Kristjanpoller, W. (2019) 'An adaptive forecasting approach for copper price volatility through hybrid and non-hybrid models', *Applied Soft Computing*, Vol. 74, pp.466–478.

Giannone, D., Lenza, M. and Primiceri, G.E. (2021) 'Economic predictions with big data: the illusion of sparsity', *Econometrica*, Vol. 89, No. 5, pp.2409–2437.

Giraud, C. (2021) *Introduction to High-Dimensional Statistics*, Chapman and Hall/CRC, Paris.

He, W., Wu, H., Yan, G., Akula, V. and Shen, J. (2015) 'A novel social media competitive analytics framework with sentiment benchmarks', *Information & Management*, Vol. 52, No. 7, pp.801–812.

Javaid, H.A. (2024) 'AI-driven predictive analytics in finance: transforming risk assessment and decision-making', *Advances in Computer Sciences*, Vol. 7, No. 1.

Jones, S. (2017) 'Corporate bankruptcy prediction: a high dimensional analysis', *Review of Accounting Studies*, Vol. 22, pp.1366–1422.

Kabir, M.R., Bhadra, D., Ridoy, M. and Milanova, M. (2025) 'LSTM – transformer-based robust hybrid deep learning model for financial time series forecasting', *Sci.*, Vol. 7, No. 1, p.7.

Karanasos, M., Paraskevopoulos, A.G., Ali, F.M., Karoglou, M. and Yfanti, S. (2014) 'Modelling stock volatilities during financial crises: a time varying coefficient approach', *Journal of Empirical Finance*, Vol. 29, pp.113–128.

Khalil, F. and Pipa, G. (2022) 'Is deep-learning and natural language processing transcending the financial forecasting? Investigation through lens of news analytic process', *Computational Economics*, Vol. 60, No. 1, pp.147–171.

Khashei, M. and Bijari, M. (2011) 'A novel hybridization of artificial neural networks and ARIMA models for time series forecasting',' *Applied Soft Computing*, Vol. 11, No. 2, pp.2664–2675.

Khattak, B.H.A., Shafi, I., Khan, A.S., Flores, E.S., Lara, R.G., Samad, M.A. and Ashraf, I. (2023) 'A systematic survey of AI models in financial market forecasting for profitability analysis', *IEEE Access*, Vol. 11, pp.125359–125380, DOI: 10.1109/ACCESS.2023.3330156.

Kim, Y.S., Rachev, S.T., Bianchi, M.L. and Fabozzi, F.J. (2008) 'Financial market models with Lévy processes and time-varying volatility', *Journal of Banking & Finance*, Vol. 32, No. 7, pp.1363–1378.

Li, K., Xue, W., Tan, G. and Denzer, A.S. (2020) 'A state of the art review on the prediction of building energy consumption using data-driven technique and evolutionary algorithms', *Building Services Engineering Research and Technology*, Vol. 41, No. 1, pp.108–127.

Maharajan, K. (2025) 'The role of sentiment analysis and transformer models in stock market price forecasting', *2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI)*, IEEE, pp.845–850.

Mavillonio, M.S. (2024) *Natural Language Processing Techniques for Long Financial Document*, Dipartimento di Economia e Management (DEM), University of Pisa, Pisa, Italy Discussion Papers.

Mehta, K. and Sharma, R. (2011) 'Measurement of time varying volatility of Indian stock market through GARCH model', *Asia Pacific Business Review*, Vol. 7, No. 3, pp.34–46.

Mozaffari, L. (2024) *Stock Market Time Series Forecasting using Transformer Models*, Oslo Metropolitan University.

Mughees, N., Mohsin, S.A., Mughees, A. and Mughees, A. (2021) 'Deep sequence to sequence Bi-LSTM neural networks for day-ahead peak load forecasting', *Expert Systems with Applications*, Vol. 175, p.114844, DOI: https://doi.org/10.1016/j.eswa.2021.114844.

Najafabadi, M.M., Villanustre, F., Khoshgoftaar, T.M., Seliya, N., Wald, R. and Muharemagic, E. (2015) 'Deep learning applications and challenges in big data analytics', *Journal of Big Data*, Vol. 2, pp.1–21.

Namdari, A. and Durrani, T.S. (2021) 'A multilayer feedforward perceptron model in neural networks for predicting stock market short-term trends', *Operations Research Forum*, Vol. 2, No. 3, p.38, 21 July, DOI: 10.1007/s43069-021-00071-2.

Nilsson, M. and Shan, Q. (2018) 'Credit risk analysis with machine learning techniques in peer-to-peer lending market'.

Nosratabadi, S., Mosavi, A., Duan, P., Ghamisi, P., Filip, F., Band, S.S., Reuter, U., Gama, J. and Gandomi, A.H. (2020) 'Data science in economics: comprehensive review of advanced machine learning and deep learning methods', *Mathematics*, Vol. 8, No. 10, p.1799.

Oikonomou, K. and Damigos, D. (2024) 'Short term forecasting of base metals prices using a LightGBM and a LightGBM – ARIMA ensemble', *Mineral Economics*, DOI: 10.1007/s13563-024-00437-y.

Olorunnimbe, M.K. (2024) *Temporal Deep Learning for Financial Time Series*, Université d'Ottawa, University of Ottawa.

Olubusola, O., Mhlongo, N.Z., Daraojimba, D.O., Ajayi-Nifise, A. and Falaiye, T. (2024) 'Machine learning in financial forecasting: a US review: exploring the advancements, challenges, and implications of AI-driven predictions in financial markets', *World Journal of Advanced Research and Reviews*, Vol. 21, No. 2, pp.1969–1984.

Pabuccu, H. and Barbu, A. (2024) 'Feature selection with annealing for forecasting financial time series', *Financial Innovation*, Vol. 10, No. 1, p.87.

Padhi, D.K., Padhy, N., Bhoi, A.K., Shafi, J. and Ijaz, M.F. (2021) 'A fusion framework for forecasting financial market direction using enhanced ensemble models and technical indicators', *Mathematics*, Vol. 9, No. 21, p.2646, https://www.mdpi.com/2227-7390/9/21/2646.

Premanode, B. (2013) *Prediction of Non-Linear Nonstationary Time Series Data using a New Digital Filter and Support Vector Regression*, Imperial College London Department of Electrical and Electronic Engineering, High Wycombe, UK.

Ramya, D. (2025) 'Reinforcement learning driven trading algorithm with optimized stock portfolio management scheme to control financial risk', *SN Computer Science*, Vol. 6, No. 1, pp.1–16.

Rane, N.L., Paramesha, M., Choudhary, S.P. and Rane, J. (2024) 'Artificial intelligence, machine learning, and deep learning for advanced business strategies: a review', *Partners Universal International Innovation Journal*, Vol. 2, No. 3, pp.147–171.

Ren, S. (2022) 'Optimization of enterprise financial management and decision-making systems based on big data', *Journal of Mathematics*, Vol. 2022, No. 1, p.1708506.

Rostamian, A. (2024) *Applications of Deep Learning Models in Financial Forecasting*, University of Essex.

Sahu, S.K., Mokhade, A. and Bokde, N.D. (2023) 'An overview of machine learning, deep learning, and reinforcement learning-based techniques in quantitative finance: recent progress and challenges', *Applied Sciences*, Vol. 13, No. 3, p.1956, https://www.mdpi.com/2076-3417/13/3/1956.

Sakib, M. and Mustajab, S. (2024) 'Enhanced multi-variate time series prediction through statistical-deep learning integration: the VAR-stacked LSTM model', *SN Computer Science*, Vol. 5, No. 5, p.573.

Sattar, A., Sarwar, A., Gillani, S., Bukhari, M., Rho, S. and Faseeh, M. (2025) 'A novel RMS-driven deep reinforcement learning for optimized portfolio management in stock trading', *IEEE Access*.

Selvarajan, G. (2021) 'Leveraging AI-enhanced analytics for industry-specific optimization: a strategic approach to transforming data-driven decision-making', *International Journal of Enhanced Research in Science Technology & Engineering*, Vol. 10, pp.78–84.

Settembre-Blundo, D., González-Sánchez, R., Medina-Salgado, S. and García-Muiña, F.E. (2021) 'Flexibility and resilience in corporate decision making: a new sustainability-based risk management system in uncertain times', *Global Journal of Flexible Systems Management*, Vol. 22, No. Supp. 2, pp.107–132.

Shobayo, O., Adeyemi-Longe, S., Popoola, O. and Ogunleye, B. (2024) 'Innovative sentiment analysis and prediction of stock price using FinBERT, GPT-4 and logistic regression: a data-driven approach', *Big Data and Cognitive Computing*, Vol. 8, No. 11, p.143.

Singh, V., Chen, S-S., Singhania, M., Nanavati, B., Kar, A.K. and Gupta, A. (2022) 'How are reinforcement learning and deep learning algorithms used for big data based decision making in financial industries – a review and research agenda', *International Journal of Information Management Data Insights*, Vol. 2, No. 2, p.100094, DOI: https://doi.org/10.1016/j.jjimei.2022.100094.

Tariq, U., Ahmed, I., Khan, M.A. and Bashir, A.K. (2024) 'Deep learning for economic transformation: a parametric review', *Indonesian Journal of Electrical Engineering and Computer Science*, Vol. 35, No. 1, pp.520–541.

Thakkar, A. and Chaudhari, K. (2021) 'A comprehensive survey on portfolio optimization, stock price and trend prediction using particle swarm optimization', *Archives of Computational Methods in Engineering*, Vol. 28, No. 4, pp.2133–2164.

Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J. and Sun, L. (2022) *Transformers in Time Series: A Survey*, arXiv preprint arXiv:2202.07125.

Wu, J., Zhang, Z. and Zhou, S.X. (2022) 'Credit rating prediction through supply chains: a machine learning approach', *Production and Operations Management*, Vol. 31, No. 4, pp.1613–1629.

Xu, J., Lu, Z. and Xie, Y. (2021) 'Loan default prediction of Chinese P2P market: a machine learning methodology', *Scientific Reports*, Vol. 11, No. 1, p.18759.

Yang, Y. (2024) 'Research on financial investment risk assessment techniques and applications', *Science, Technology and Social Development Proceedings Series*, Vol. 2, No. 10, p.70088.

Yılmaz, B. (2020) 'Forecasting house prices in Turkey: GLM, VaR and time series approaches', *Journal of Business Economics and Finance*, Vol. 9, No. 4, pp.274–291.

Zafar, F., Jabbar, B., Bano, A., Khan, M.F.U., Ahmed, O. and Salman, S. (2025) 'Nexus of univariate and multivariate models for forecasting interest rates in Pakistan using VAR-COV analysis', *Advance Journal of Econometrics and Finance*, Vol. 3, No. 1, pp.68–84.

Zhang, N., An, Q., Zhang, S. and Ma, H. (2024) 'Price prediction for fresh agricultural products based on a boosting ensemble algorithm', *Mathematics*, Vol. 13, No. 1, p.71.

Zhang, X. (2024) *Analyzing Financial Market Trends in Cryptocurrency and Stock Prices Using CNN-LSTM Models*, Preprints.

Zhou, L., Zhang, Y., Yu, J., Wang, G., Liu, Z., Yongchareon, S. and Wang, N. (2025) 'LLM-augmented linear transformer-CNN for enhanced stock price prediction', *Mathematics*, Vol. 13, No. 3, p.487.

Zhu, P., Li, Y., Hu, Y., Xiang, S., Liu, Q., Cheng, D. and Liang, Y. (2024) *MCI-GRU: Stock Prediction Model based on Multi-Head Cross-Attention and Improved GRU*, arXiv preprint arXiv:2410.20679.

Zhu, Q., Han, J., Chai, K. and Zhao, C. (2023) 'Time series analysis based on informer algorithms: a survey', *Symmetry*, Vol. 15, No. 4, p.951.

Zhu, X., Ninh, A., Zhao, H. and Liu, Z. (2021) 'Demand forecasting with supply-chain information and machine learning: evidence in the pharmaceutical industry', *Production and Operations Management*, Vol. 30, No. 9, pp.3231–3252.