

Article

Not peer-reviewed version

---

# Advanced Stock Price Prediction with xLSTM-Based Models: Improving Long-Term Forecasting

---

[Xiaojing Fan](#)<sup>\*</sup>, Chunliang Tao, Jianyu Zhao

Posted Date: 29 August 2024

doi: 10.20944/preprints202408.2109.v1

Keywords: Neural Networks; Stock Price Prediction; Long Short-Term Memory (LSTM); Extended Long Short-Term Memory (xLSTM); Time Series Forecasting



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

## Article

# Advanced Stock Price Prediction with xLSTM-Based Models: Improving Long-Term Forecasting

Xiaojing Fan <sup>\*,†</sup>, Chunliang Tao <sup>\*,†</sup> and Jianyu Zhao <sup>\*,†</sup>

New York University

\* Correspondence: xf435@nyu.edu (X.F.); ct1942@nyu.edu (C.T.); jz5246@nyu.edu (J.Z.)

† These authors contributed equally to this work.

**Abstract:** Stock price prediction has long been a critical area of research in financial modeling. The inherent complexity of financial markets, characterized by both short-term fluctuations and long-term trends, poses significant challenges in accurately capturing underlying patterns. While Long Short-Term Memory (LSTM) networks have shown strong performance in short-term stock price prediction, they struggle with effectively modeling long-term dependencies. In this paper, we propose an advanced stock price prediction model based on the Extended Long Short-Term Memory (xLSTM) algorithm, designed to enhance predictive accuracy over both short and long-term periods. We conduct extensive experiments by building and evaluating models based on xLSTM and LSTM architectures for multiple stocks. Our results demonstrate that the xLSTM model consistently outperforms the LSTM model across all stocks and time horizons, with the performance gap widening as the prediction period extends. The observations underscore the superior capability of the xLSTM-based model to capture long-term patterns in financial data, offering a promising approach for more accurate and reliable stock price predictions.

**Keywords:** neural networks; stock price prediction; Long Short-Term Memory (LSTM); Extended Long Short-Term Memory (xLSTM); time series forecasting

## 1. Introduction

Predicting stock prices has long been a central focus in financial modeling due to its profound impact on investment strategies, risk management, and market efficiency. The complexity and volatility of financial markets, driven by both short-term fluctuations and long-term trends, make accurate stock price prediction a particularly challenging task [1]. The unpredictable nature of stock prices is further compounded by numerous factors, including economic indicators, market sentiment, and external events, all contributing to the intricate dynamics of financial markets [2–4].

Over the years, various machine learning techniques, such as K-Nearest Neighbors (k-NN) [5], AutoRegressive Integrated Moving Average (ARIMA) [6], and XGBoost [7], have been employed to address the problem of stock price prediction. While these methods have provided valuable insights and demonstrated effectiveness in specific contexts, they often fall short of fully capturing the complex and dynamic nature of financial markets. The limitations of these models, particularly in dealing with non-linear relationships and long-term dependencies, have prompted researchers to explore more advanced approaches.

Recent advancements in deep learning have led to the development of more sophisticated models capable of capturing the non-linear relationships inherent in time series data [8]. Among these, Long Short-Term Memory (LSTM) networks have gained prominence for their ability to model temporal dependencies and effectively manage sequential data [9–11]. LSTM networks have demonstrated significant improvements in stock price prediction, particularly in capturing short-term market trends [12–14]. However, despite their effectiveness, LSTM models face considerable limitations in modeling long-term dependencies, which are crucial for accurate long-term predictions [15,16]. This limitation can result in reduced prediction accuracy, especially when forecasting stock prices over extended periods. A recent innovation, Extended Long Short-Term Memory (xLSTM) architecture, builds

upon the traditional LSTM by incorporating mechanisms to better manage and retain information over longer sequences and refining the gating mechanisms to allow for more effective processing of relevant information across sequences of various lengths[17]. xLSTM has demonstrated competitive performance in handling long sequences and large contexts compared to some prominent large language models [18]. However, its potential in financial time series forecasting remains largely unexplored.

In this work, we aim to develop and evaluate an advanced stock price prediction model based on the xLSTM architecture to address the limitations of current approaches in modeling long-term trends. While some studies have leveraged Large Language Models that incorporate news and market sentiment to enhance financial modeling [19–21], our focus in this paper is on non-text-based deep learning approaches. We conduct extensive experiments to compare the performance of xLSTM and LSTM models across multiple stocks and a range of prediction horizons. LSTM is chosen as the benchmark for comparison due to its established success in short-term stock price prediction in previous studies. The results demonstrate that our xLSTM model consistently outperforms the traditional LSTM model in both short-term and long-term stock price predictions. These findings underscore the strong potential of xLSTM in financial time series forecasting, offering more accurate and reliable predictions that can enhance decision-making processes in the financial industry.

## 2. Related Work

### 2.1. Traditional Methods

Traditional statistical and machine learning models have been extensively studied and applied to a wide range of prediction tasks in both financial [22–24] and non-financial domains [25–27]. For time-series forecasting tasks, including stock price predictions, Autoregressive Integrated Moving Average (ARIMA) [6], Generalized Autoregressive Conditional Heteroskedasticity (GARCH) [28], and Exponential Smoothing (ETS) [29] have been widely used. These models were effective in capturing linear relationships and volatility clustering, which are typical in financial time series data. However, their performance was limited by their inability to model complex non-linear patterns that are often present in stock market movements. Other machine learning methods such as Support Vector Machines (SVM) [30] and Random Forests [31] have also been applied for stock price predictions, with varying degrees of success, but they still face challenges in effectively capturing the temporal dynamics of financial data.

### 2.2. Deep Learning Approaches

Deep learning has revolutionized various fields by offering powerful models capable of capturing complex patterns in data. Its applications span numerous domains including but not limited to healthcare [32–36], retail[37], civil-engineering[38], psychology [39] and finance[40]. For stock price predictions, Recurrent Neural Networks (RNNs) have been particularly effective due to their ability to handle sequential data leveraging a hidden state that captures information from previous time steps. However, traditional RNNs suffer from issues such as vanishing and exploding gradients, which limits their ability to capture long-term dependencies in time series data [41]. Long Short-Term Memory (LSTM) networks are developed to overcome some of these limitations [42]. Specifically, the featuring memory cells and gating mechanisms of LSTM, enable it to retain information over longer sequences, making it particularly effective for time series forecasting. It has been shown that LSTM networks significantly outperform traditional models in predicting stock prices by effectively capturing complex temporal dependencies, especially in short-term forecasts [43,44].

### 2.3. Limitations of LSTM

Despite its success, LSTM is not without limitations, particularly when applied to long-term stock price prediction. Studies have shown that while LSTM models are robust for short to medium-term

predictions, their performance tends to degrade when predicting over longer horizons. [41]. To address the limitations of LSTM, several advancements and alternative methods have been introduced. Gated Recurrent Units (GRU), for example, offers a simplified architecture with fewer parameters compared to LSTM, making them easier to train while still capturing long-term dependencies [45]. Attention mechanisms, as seen in Transformer models, have also been applied to time series forecasting, allowing models to focus on the most relevant parts of the input sequence and improving performance in tasks like stock price prediction [46]. Among the advancements in RNNs, xLSTM has emerged as a promising extension of LSTM, designed to improve the model's capacity to handle complex, long-range dependencies through architectural modifications like enhanced gradient flow and multi-scale temporal processing [17]. In stock price prediction, short-term and long-term forecasts serve different strategic purposes. Short-term predictions are often driven by recent market movements and are critical for day trading and high-frequency trading strategies, whereas long-term predictions rely on capturing broader market trends, which are more relevant for investment strategies with extended time horizons. Given the fundamentally different drivers for short-term and long-term stock price changes, there is a clear need for prediction models that are robust for predictions across different time horizons.

### 3. Methodology

#### 3.1. xLSTM

The Extended Long Short-Term Memory (xLSTM) model extends the standard LSTM by incorporating additional mechanisms, such as enhanced memory cells and gating modifications. These improvements help the model better focus on relevant parts of the input sequence, making xLSTM particularly effective in scenarios where the relationship between past and future values is complex and non-linear. This is especially useful for stock price prediction across varying time horizons, where market conditions can shift rapidly. Some relevant innovations are explained below.

1. **Enhanced Memory Cells:** The xLSTM introduces modifications that allow memory cells to operate with more flexibility and efficiency. Specifically, xLSTM includes two variants of LSTM, sLSTM and mLSTM. The sLSTM variant uses scalar memory updates, which simplify the memory operations, allowing for faster and more efficient computation. On the other hand, mLSTM employs a matrix memory structure that can handle more complex dependencies across longer sequences. The matrix memory enables the model for parallel computations, significantly improving its ability to manage and retain information across extended periods, which is particularly beneficial in scenarios such as long-term stock price forecasting.
2. **Exponential Gating:** Gating is central to the operation of LSTM networks, controlling the flow of information through input, forget, and output gates. In xLSTM, these gates are modified to include exponential gating and advanced normalization techniques, as shown in Equation (1).

$$g_t = \exp(W_g \cdot x_t + U_g \cdot h_{t-1} + b_g) \quad (1)$$

where  $g_t$  represents the input or forget gate at time  $t$ ,  $W_g$  and  $U_g$  are weight matrices, and  $b_g$  is the bias. The exponential function  $\exp$  provides a more flexible and dynamic adjustment of gate values, allowing for a more stable gradient flow during training. Exponential gating helps manage the vanishing gradient problem more effectively by dynamically adjusting how much past information is retained or forgotten during the training process. The adjustment ensures that relevant information is preserved over long sequences, while less important data is filtered out, making the model more efficient in learning and predicting complex patterns. This is crucial in financial applications where retaining key market trends while discarding noise can significantly enhance prediction accuracy.



### 3.2. Dataset

**Data Collection** These enhancements make xLSTM particularly well-suited for scenarios involving complex, non-linear relationships and long-term dependencies, which are prevalent in stock markets.

This study employs historical stock price data sourced from Yahoo Finance [47] for three major companies: Apple, Johnson & Johnson, and Nike. The data spans from January 1, 2014, to December 31, 2023, covering a broad range of market conditions including bullish, bearish, and volatile phases. These companies are selected based on their significant market capitalization, industry leadership, and consistent trading volumes, which ensure that the data is representative of different sectors, i.e. consumer goods, healthcare, and technology, and is robust for time series analysis. We select the daily 'Close' price as the prediction target. The 'Close' price is a critical indicator of market sentiment and is widely used in financial forecasting due to its ability to encapsulate the outcome of daily trading activities.

**Data Pre-processing** Given the sequential nature of stock prices, the dataset is structured into sequences of 60 consecutive days. This sequence length is strategically chosen to capture sufficient historical context, balancing the need to model relevant market dynamics without introducing excessive noise. This approach is aligned with best practices in financial time series forecasting, where similar window lengths effectively capture meaningful patterns in stock price movements. Additionally, multiple future time points are selected in our work:  $t+1$ ,  $t+3$ ,  $t+5$ ,  $t+10$ , and  $t+15$  days to allow for a comprehensive analysis of predictive performance across both short-term and long-term horizons. Furthermore, to standardize the data and facilitate effective training of models, the 'Close' prices are normalized using a MinMaxScaler, with a range from 0 to 1. Normalization is crucial in time series modeling, especially when using neural networks, as it ensures consistent input data and improves model convergence and stability.

## 4. Experiments

### 4.1. Experiment Setup

We conduct a series of experiments across multiple stocks and prediction horizons to evaluate the effectiveness of the xLSTM model in comparison to the traditional LSTM model.

**Model Training** Both the xLSTM and LSTM models are trained using the pre-processed dataset described in the Methodology section. For each stock, the data was split into three sets: 60% for training, 20% for validation, and 20% for testing. Each model is trained separately on a per-stock and per-time horizon basis to ensure that the models are tailored to the unique characteristics of each stock and the specific prediction horizons. For example, the xLSTM and LSTM models are independently trained on the NIKE stock for each of the following prediction horizons:  $t+1$ ,  $t+3$ ,  $t+5$ ,  $t+10$ , and  $t+15$  days.

**Hyperparameter Tuning** To ensure a fair and thorough comparison between the xLSTM and LSTM models, we implement a hyperparameter tuning process that was carefully tailored to optimize each model individually. We choose a fixed value or range for some hyperparameters shared by both models across all prediction horizons to maintain fairness and consistency. Specifically, a consistent dropout rate of 0.1 is applied to both models to prevent overfitting, with early stopping based on validation performance to ensure that neither model is over-trained. Additionally, we set the learning rate uniformly at 0.001 across all training tasks to ensure stable and comparable training dynamics. The number of layers is explored in the range of 2 to 4, while the hidden size is varied between 32, 64, and 128. We also include key hyperparameters that are unique and critical to the xLSTM model, including the number of heads and kernel sizes for both the sLSTM and mLSTM components, as well as the mLSTM projection block sizes. We implement Random Search to fine-tune all hyperparameters to optimize the performance of both models over each stock and each prediction horizon. The

best-performing model configuration is selected based on its performance on the validation set for final evaluation.

**Model Evaluation** We evaluate the trained and tuned xLSTM and LSTM models by assessing prediction accuracy on the test set for each stock over each prediction horizon. This comparison provides a comprehensive understanding of how well each model performs on diverse stocks as well as in short-term and long-term predictions. Specific evaluation metrics are documented in the following section.

#### 4.2. Evaluation Metrics

To thoroughly evaluate the performance of the xLSTM and LSTM models, we employed three key metrics: Root Mean Squared Error (RMSE), R-squared ( $R^2$ ), and Net Error Reduction (NER). Each of these metrics provides a unique perspective on models' predictive accuracy and ability to capture both short-term and long-term trends in stock prices.

**RMSE** is a standard metric used to measure the average magnitude of prediction errors. RMSE is particularly useful for understanding the overall accuracy of the model, as it penalizes larger errors more heavily. A lower RMSE indicates higher predictive accuracy. RMSE is calculated as follows:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2} \quad (2)$$

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

**$R^2$**  is a statistical measure that represents the proportion of variance in the dependent variable (stock prices) that can be explained by the model.  $R^2$  indicates how well the model's predictions match the actual data, with values closer to 1 indicating a better fit. The  $R^2$  is calculated as follows:

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (3)$$

where  $\bar{y}$  is the mean of the actual stock prices.

**ER** is a metric introduced to quantify the net difference in predictive accuracy between the xLSTM and LSTM models across various stocks and prediction horizons. This metric is particularly useful for understanding how the performance gap between the two models changes as the prediction window increases. A positive ER indicates that the xLSTM model outperforms the LSTM model by reducing the error, while a negative ER suggests that the LSTM model has better predictive accuracy. By analyzing ER across different prediction horizons, we can gain insights into how the relative performance of the models evolves with the length of the forecast period. ER is calculated as:

$$\text{ER} = \text{RMSE}_{\text{LSTM}} - \text{RMSE}_{\text{xLSTM}} \quad (4)$$

## 5. Results

### 5.1. Stock-Specific Performance

Table 1 on page 4 reveals clear differences in the predictability of the three stocks due to different magnitudes in market volatility. For example, as illustrated in Figure 1 and Figure 2 on page 4, AAPL exhibits higher market volatility than JNJ from 2022 to 2023. For the results, AAPL consistently exhibits the highest RMSE and the lowest  $R^2$  across all models and prediction horizons, reflecting the challenges both LSTM and xLSTM face in forecasting highly volatile stocks. In contrast, JNJ, consistently shows the lowest RMSE and the highest  $R^2$ , indicating that both models are more effective at forecasting less volatile stocks where price movements are more stable. Despite the varying levels of volatility, xLSTM consistently outperforms LSTM across all stocks and time horizons. Even for the more volatile

AAPL, xLSTM reduces prediction errors more effectively, demonstrating its superior ability to handle complex market conditions. This consistent outperformance underscores xLSTM’s robustness and adaptability across different market environments.

Table 1. Overall Performance Across Time Horizons.

| STOCK MODEL |       | T = 1        |                  | T = 3        |                  | T = 5        |                  | T = 10       |                  | T = 15        |                  |
|-------------|-------|--------------|------------------|--------------|------------------|--------------|------------------|--------------|------------------|---------------|------------------|
|             |       | RMSE↓        | R <sup>2</sup> ↑ | RMSE↓        | R <sup>2</sup> ↑ | RMSE↓        | R <sup>2</sup> ↑ | RMSE↓        | R <sup>2</sup> ↑ | RMSE↓         | R <sup>2</sup> ↑ |
| AAPL        | LSTM  | 4.341        | 0.887            | 6.513        | 0.768            | 7.839        | 0.665            | 10.416       | 0.464            | 12.433        | 0.261            |
|             | xLSTM | <b>3.020</b> | <b>0.941</b>     | <b>5.155</b> | <b>0.831</b>     | <b>6.347</b> | <b>0.875</b>     | <b>8.750</b> | <b>0.742</b>     | <b>10.500</b> | <b>0.376</b>     |
| JNJ         | LSTM  | 2.381        | 0.941            | 3.613        | 0.882            | 4.340        | 0.834            | 5.429        | 0.712            | 6.413         | 0.601            |
|             | xLSTM | <b>1.716</b> | <b>0.972</b>     | <b>2.915</b> | <b>0.917</b>     | <b>3.605</b> | <b>0.875</b>     | <b>4.670</b> | <b>0.568</b>     | <b>5.616</b>  | <b>0.659</b>     |
| NKE         | LSTM  | 3.516        | 0.926            | 5.364        | 0.835            | 6.614        | 0.763            | 9.094        | 0.515            | 11.128        | 0.320            |
|             | xLSTM | <b>2.522</b> | <b>0.962</b>     | <b>4.334</b> | <b>0.886</b>     | <b>5.474</b> | <b>0.818</b>     | <b>7.817</b> | <b>0.625</b>     | <b>9.690</b>  | <b>0.419</b>     |
| AVG         | LSTM  | 3.413        | 0.918            | 5.163        | 0.828            | 6.264        | 0.754            | 8.313        | 0.564            | 9.991         | 0.394            |
|             | xLSTM | <b>2.419</b> | <b>0.958</b>     | <b>4.135</b> | <b>0.878</b>     | <b>5.142</b> | <b>0.812</b>     | <b>7.079</b> | <b>0.652</b>     | <b>8.602</b>  | <b>0.484</b>     |




Figure 1. AAPL 1-Day Stock Predictions by xLSTM and LSTM Models



Figure 2. JNJ 1-Day Stock Predictions by xLSTM and LSTM Models

5.2. Model Performance across Prediction Horizons

Table 1 also shows that both LSTM and xLSTM models experience a decline in predictive accuracy, reflected in increasing RMSE and decreasing  $R^2$ , as the prediction horizon extends from 1 to 15 days. This trend is consistent across all stocks and highlights the common challenge of maintaining accuracy over longer time horizons in time series forecasting. However, the xLSTM model consistently outperforms LSTM across all stocks and prediction horizons. The xLSTM’s lower RMSE and higher  $R^2$  demonstrate its enhanced capability to model complex temporal dependencies, particularly over extended periods. The superior performance is likely due to xLSTM’s advanced architecture, which includes improved gating mechanisms and memory management, enabling it to retain and utilize relevant information more effectively than LSTM as the forecast horizon lengthens.

5.3. Performance Gap across Time Horizons

Table 2 highlights the increasing performance gap between xLSTM and LSTM, measured by Error Reduction (ER), as the prediction horizon lengthens. The gap becomes more pronounced as the forecasting task grows more challenging with longer time horizons. While both models show increased RMSE as the horizon extends from 1 to 15 days, the rate of increase is notably slower for xLSTM. This suggests that xLSTM is more resilient to the typical degradation in predictive accuracy that accompanies longer forecasting periods. The widening gap underscores LSTM’s limitations in handling long-term dependencies, where its accuracy declines significantly. In contrast, xLSTM manages these dependencies more effectively, maintaining a more stable error profile across different prediction horizons.

Table 2. Error Reduction Across Time Horizons.

| STOCK | T = 1 | T = 3 | T = 5 | T = 10 | T = 15 |
|-------|-------|-------|-------|--------|--------|
| AAPL  | 1.321 | 1.358 | 1.492 | 1.666  | 1.933  |
| JNJ   | 0.665 | 0.698 | 0.735 | 0.759  | 0.797  |
| NKE   | 0.994 | 1.030 | 1.140 | 1.277  | 1.438  |
| AVG   | 0.993 | 1.029 | 1.122 | 1.234  | 1.389  |

5.4. Practical Implications

LSTM models have long been recognized for their strong performance in time series forecasting, especially in financial contexts, where they have outperformed traditional models. However, our study demonstrates that xLSTM not only builds on these strengths but also significantly enhances predictive accuracy, particularly as the forecast horizon extends. The practical implications of xLSTM’s superior performance are profound for financial analysts and investors. In scenarios where long-term forecasts are critical, xLSTM’s enhanced accuracy can lead to more reliable predictions, reducing uncertainty, improving decision-making, and potentially enhancing financial outcomes. Moreover, xLSTM’s robustness across different stocks and market conditions suggests it is a versatile and dependable tool for a wide range of financial forecasting applications. Whether dealing with volatile markets or stable trends, xLSTM’s ability to maintain lower prediction errors over extended periods positions it as a valuable asset in the financial industry.

6. Conclusions

In this work, we build an enhanced stock price prediction model using xLSTM to address the limitations of traditional LSTM models in managing long-term dependencies. Through extensive experiments conducted across three major stocks, Apple, Johnson & Johnson, and Nike, and various prediction horizons, we demonstrate that the xLSTM model consistently outperforms the standard LSTM. Furthermore, as the prediction horizon extends, the performance gap between xLSTM and LSTM becomes increasingly evident, underscoring xLSTM’s ability to maintain accuracy where traditional models tend to falter.

Despite the promising results, several limitations of our study warrant further exploration. For example, our work focuses exclusively on historical stock prices without incorporating external factors such as macroeconomic indicators, news sentiment, or market volatility indices. Including these variables could improve the model’s predictive accuracy, particularly in volatile markets. Future work should explore how the integration of such external factors can further enhance xLSTM’s performance. Meanwhile, our study focuses on specific prediction horizons of 1, 3, 5, 10, and 15 days. This approach may not fully capture the performance of xLSTM across other relevant time frames, such as intraday predictions or longer-term forecasts extending to months or years. Expanding the range of prediction horizons could provide deeper insights into the model’s applicability to different investment strategies. Lastly, our experiments are conducted on a small set of three stocks. While these stocks are chosen for their market significance, the findings may not generalize across a broader range of financial



instruments or different market conditions. Future research should validate these results using a more extensive and diverse dataset, encompassing stocks from various sectors and international markets.

## References

1. Feng Ma, Jiqian Wang, MIM Wahab, and Yuanhui Ma. Stock market volatility predictability in a data-rich world: A new insight. *International Journal of Forecasting*, 39(4):1804–1819, 2023.
2. Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. Temporal relational ranking for stock prediction. *ACM Transactions on Information Systems (TOIS)*, 37(2):1–30, 2019.
3. Malcolm Baker and Jeffrey Wurgler. Investor sentiment and the cross-section of stock returns. *The journal of Finance*, 61(4):1645–1680, 2006.
4. Muhammad Abubakr Naeem, Fiza Qureshi, Saqib Farid, Aviral Kumar Tiwari, and Mohamed Elheddad. Time-frequency information transmission among financial markets: evidence from implied volatility. *Annals of Operations Research*, 334(1):701–729, 2024.
5. Khalid Alkhatib, Hassan Najadat, Ismail Hmeidi, and Mohammed K Ali Shatnawi. Stock price prediction using k-nearest neighbor (knn) algorithm. *International Journal of Business, Humanities and Technology*, 3(3):32–44, 2013.
6. Supriti Khanderwal and Debasis Mohanty. Stock price prediction using arima model. *International Journal of Marketing & Human Resource Research*, 2(2):98–107, 2021.
7. Kyung Keun Yun, Sang Won Yoon, and Daehan Won. Prediction of stock price direction using a hybrid ga-xgboost algorithm with a three-stage feature engineering process. *Expert Systems with Applications*, 186:115716, 2021.
8. Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Deep learning for time series classification: a review. *Data mining and knowledge discovery*, 33(4):917–963, 2019.
9. S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997.
10. Martin Sundermeyer, Hermann Ney, and Ralf Schlüter. From feedforward to recurrent lstm neural networks for language modeling. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):517–529, 2015.
11. Haowei Ni, Shuchen Meng, Xieming Geng, Panfeng Li, Zhuoying Li, Xupeng Chen, Xiaotong Wang, and Shiyao Zhang. Time series modeling for heart rate prediction: From arima to transformers. *arXiv preprint arXiv:2406.12199*, 2024.
12. Md Arif Istiaque Sunny, Mirza Mohd Shahriar Maswood, and Abdullah G Alharbi. Deep learning-based stock price prediction using lstm and bi-directional lstm model. In *2020 2nd novel intelligent and leading emerging sciences conference (NILES)*, pages 87–92. IEEE, 2020.
13. Sidra Mehtab, Jaydip Sen, and Abhishek Dutta. Stock price prediction using machine learning and lstm-based deep learning models. In *Machine Learning and Metaheuristics Algorithms, and Applications: Second Symposium, SoMMA 2020, Chennai, India, October 14–17, 2020, Revised Selected Papers 2*, pages 88–106. Springer, 2021.
14. Abdul Quadir Md, Sanjit Kapoor, Chris Junni AV, Arun Kumar Sivaraman, Kong Fah Tee, H Sabireen, and N Janakiraman. Novel optimization approach for stock price forecasting using multi-layered sequential lstm. *Applied Soft Computing*, 134:109830, 2023.
15. Felix A Gers, Nicol N Schraudolph, and Jürgen Schmidhuber. Learning precise timing with lstm recurrent networks. *Journal of machine learning research*, 3(Aug):115–143, 2002.
16. Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.
17. Maximilian Beck, Korbinian Pöppel, Markus Spanring, Andreas Auer, Oleksandra Prudnikova, Michael Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. xlstm: Extended long short-term memory. *arXiv preprint arXiv:2405.04517*, 2024.
18. Musleh Alharthi and Ausif Mahmood. xlstmtime: Long-term time series forecasting with xlstm. *AI*, 5(3):1482–1495, 2024.

19. Haowei Ni, Shuchen Meng, Xupeng Chen, Ziqing Zhao, Andi Chen, Panfeng Li, Shiyao Zhang, Qifu Yin, Yuanqing Wang, and Yuxi Chan. Harnessing earnings reports for stock predictions: A qlora-enhanced llm approach. *arXiv preprint arXiv:2408.06634*, 2024.
20. Kangtong Mo, Wenyan Liu, Xuanzhen Xu, Chang Yu, Yuelin Zou, and Fangqing Xia. Fine-tuning gemma-7b for enhanced sentiment analysis of financial news headlines. *arXiv preprint arXiv:2406.13626*, 2024.
21. Wenjun Gu, Yihao Zhong, Shizun Li, Changsong Wei, Liting Dong, Zhuoyue Wang, and Chao Yan. Predicting stock prices with finbert-lstm: Integrating news sentiment analysis. *arXiv preprint arXiv:2407.16150*, 2024.
22. Yuqi Yan. Influencing factors of housing price in new york-analysis: Based on excel multi-regression model. In *Proceedings of the International Conference on Big Data Economy and Digital Management*, 2022. doi:10.5220/0011362000003440.
23. Siqiao Zhao, Zhikang Dong, Zeyu Cao, and Raphael Douady. Hedge fund portfolio construction using polymodel theory and itransformer. *arXiv preprint arXiv:2408.03320*, 2024.
24. Chao Yan, Yijie Weng, Jinyin Wang, Yang Zhao, Yuelin Zou, Zhuoying Li, and US Baltimore. Enhancing credit card fraud detection through adaptive model optimization. *10.13140/RG.2.2.12274.52166*, 2024.
25. Tiechuan Hu, Wenbo Zhu, and Yuqi Yan. Artificial intelligence aspect of transportation analysis using large scale systems. In *Proceedings of the 2023 6th Artificial Intelligence and Cloud Computing Conference*, pages 54–59, 2023.
26. Jianlong Chen, Wei Xu, and Jinyin Wang. Prediction of car purchase amount based on genetic algorithm optimised bp neural network regression algorithm. *Preprints*, 2024. <https://doi.org/10.20944/preprints202406.1304.v1>.
27. Shifan Zhao, Jiaying Lu, Ji Yang, Edmond Chow, and Yuanzhe Xi. Efficient two-stage gaussian process regression via automatic kernel search and subsampling. *arXiv preprint arXiv:2405.13785*, 2024.
28. Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007, 1982.
29. Everette S Gardner Jr. Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1):1–28, 1985.
30. Yuling Lin, Haixiang Guo, and Jinglu Hu. An svm-based approach for stock market trend prediction. In *The 2013 international joint conference on neural networks (IJCNN)*, pages 1–7. IEEE, 2013.
31. Luckyson Khaidem, Snehanishu Saha, and Sudeepa Roy Dey. Predicting the direction of stock market prices using random forest. *arXiv preprint arXiv:1605.00003*, 2016.
32. Chenwei Wu, Xiyu Yang, Emil Ghitman Gilkes, Hanwen Cui, Jiheon Choi, Na Sun, Ziqian Liao, Bo Fan, Mauricio Santillana, Leo Celi, et al. De-identification and obfuscation of gender attributes from retinal scans. In *Workshop on Clinical Image-Based Procedures*, pages 91–101. Springer, 2023.
33. Yihao Zhong, Yanan Liu, Erdi Gao, Changsong Wei, Zhuoyue Wang, and Chao Yan. Deep learning solutions for pneumonia detection: Performance comparison of custom and transfer learning models. *medRxiv*, pages 2024–06, 2024.
34. Qimin Zhang, Weiwei Qi, Huili Zheng, and Xinyu Shen. Cu-net: a u-net architecture for efficient brain-tumor segmentation on brats 2019 dataset. *arXiv preprint arXiv:2406.13113*, 2024.
35. Bo Dang, Danqing Ma, Shaojie Li, Zongqing Qi, and Elly Zhu. Deep learning-based snore sound analysis for the detection of night-time breathing disorders. *Applied and Computational Engineering*, 76:109–114, 07 2024.
36. Bo Dang, Wenchao Zhao, Yufeng Li, Danqing Ma, Qixuan Yu, and Elly Yijun Zhu. Real-time pill identification for the visually impaired using deep learning. *arXiv preprint arXiv:2405.05983*, 2024.
37. Lianghao Tan, Shubing Liu, Jing Gao, Xiaoyi Liu, Linyue Chu, and Huangqi Jiang. Enhanced self-checkout system for retail based on improved yolov10. *arXiv preprint arXiv:2407.21308*, 2024.
38. Han-Cheng Dan, Bingjie Lu, and Mengyu Li. Evaluation of asphalt pavement texture using multiview stereo reconstruction based on deep learning. *Construction and Building Materials*, 412:134837, 2024.
39. Panfeng Li, Mohamed Abouelenien, Rada Mihalcea, Zhicheng Ding, Qikai Yang, and Yiming Zhou. Deception detection from linguistic and physiological data streams using bimodal convolutional neural networks. *arXiv preprint arXiv:2311.10944*, 2024.
40. Chang Yu, Yongshun Xu, Jin Cao, Ye Zhang, Yinxin Jin, and Mengran Zhu. Credit card fraud detection using advanced transformer model. *arXiv preprint arXiv:2406.03733*, 2024.
41. Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
42. Jürgen Schmidhuber, Sepp Hochreiter, et al. Long short-term memory. *Neural Comput*, 9(8):1735–1780, 1997.

43. Thomas Fischer and Christopher Krauss. Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research*, 270(2):654–669, 2018.
44. Wei Bao, Jun Yue, and Yulei Rao. A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PloS one*, 12(7):e0180944, 2017.
45. Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. arxiv 2014. *arXiv preprint arXiv:1406.1078*, 2020.
46. A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
47. Yahoo Finance. Yahoo finance historical stock data. <https://finance.yahoo.com>. Accessed: July 15, 2024.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.