A Hybrid CNN-LSTM Model for Stock Price Prediction with Spatial and Temporal Dependencies

Yi Lan

School of Mathematics, Hefei University of Technology, Hefei, China
b11111085@office365.npust.edu.tw

Abstract: Stock price prediction is crucial in financial decision-making and investment strategies, significantly influencing investors' profitability and market stability. This paper aims to systematically review and evaluate Machine Learning (ML) and Deep Learning (DL) methodologies, primarily focusing on Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNNs) for stock price forecasting. A hybrid CNN-LSTM model is proposed to enhance predictive accuracy. Specifically, the CNN component initially extracts essential spatial features from historical financial data, while the LSTM network subsequently captures the temporal dependencies to produce accurate predictions of future stock prices. Additionally, the review briefly addresses other ML techniques, including Autoregressive Integrated Moving Average (ARIMA) Model and various hybrid approaches, highlighting their strengths and limitations. This study covers diverse market conditions on publicly available financial datasets. Results indicate that the hybrid CNN-LSTM model outperforms individual LSTM and CNN models, effectively capturing both rapid fluctuations and long-term price trends. Experimental results demonstrate the proposed hybrid model's superior accuracy and potential as an effective predictive tool for investors and traders in financial markets.

Keywords: Machine Learning, Deep Learning, Convolutional Neural Networks, Long Short-Term Memory.

1. Introduction

The financial industry is recognised as one of the most challenging sectors due to its significant profitability and commercial impact. Stock, currency, and mutual fund markets greatly influence national economic growth. Stock market prices, driven by diverse economic, psychological, and political factors, exhibit complex and nonlinear behaviours, making precise short-term or long-term predictions difficult. Given the high profitability of stock trading, accurate market forecasting is crucial and draws considerable attention from investors. Recent advancements in Artificial Intelligence (AI) have demonstrated superior performance in various fields, which offer a strong foundation for applying Deep Learning (DL) models to stock market prediction.

Traditional statistical forecasting methods, such as linear regression and Autoregressive Integrated Moving Average (ARIMA), often struggle with financial market complexity, volatility, and rapid informational shifts [1]. Their linear and stationary assumptions frequently fail to capture dynamic financial data and structural changes accurately. In contrast, ML techniques dynamically adapt, effectively managing nonlinear relationships and non-stationary behaviours, thus significantly

improving predictive accuracy. Supervised ML methods, including Support Vector Machines (SVMs), Random Forests (RFs), and Decision Trees (DTs), effectively handle high-dimensional data and resist overfitting [2]. SVMs, particularly noted for their reliability, efficiently delineate complex market boundaries. RFs and DTs offer interpretability and robustness against noise, making them beneficial for multifaceted financial analyses. Advancements in DL have notably enhanced forecasting capabilities. Recurrent Neural Networks (RNNs), specifically Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), excel at modelling temporal dependencies and mitigating gradient vanishing issues [3]. LSTMs use gating mechanisms to retain crucial historical information, significantly boosting the accuracy of stock market prediction. GRU provides computational efficiency and strong modelling performance, offering valuable alternatives. Initially intended for image processing, Convolutional Neural Networks (CNNs) have been successfully adapted for financial forecasting [4]. By extracting spatial-temporal features from financial indicators and historical prices, CNNs uncover market patterns often overlooked by traditional methods, enhancing their predictive capabilities. Hybrid and ensemble ML methods, combining diverse algorithms and financial indicators, consistently outperform single-model approaches. Recent innovations, such as attention mechanisms and transformer-based models from natural language processing, have further improved predictions by effectively utilizing historical market data [5]. Attention mechanisms enhance predictive performance by selectively emphasizing essential historical data points. Transformers efficiently capture long-range dependencies via self-attention mechanisms, providing deeper market insights. Reinforcement Learning (RL), inspired by strategic game successes, dynamically improves predictions through continuous adaptation to real-time market conditions, allowing predictive models to develop optimal trading strategies interactively [6]. However, interpretability remains challenging, as complex ML models often operate as "black boxes," hindering transparency and practitioner trust. Computational complexity and resource-intensive training processes pose additional practical hurdles. Addressing these issues through interpretable AI, simplified models, adaptive algorithms, and robust training methodologies is crucial for broader adoption and effective deployment.

Despite these advancements, challenges persist in effectively implementing ML-driven predictive models. Interpretability remains a key concern, as complex ML models often function as "black boxes," complicating decision transparency and practitioner trust [7]. Computational complexity and resource-intensive training processes pose additional hurdles, especially in real-time applications. Data non-stationarity arising from continuously evolving market conditions further complicates model reliability. Potential overfitting due to overly complex models can diminish generalisation capability, reducing prediction reliability. Addressing these challenges through research into model simplification, interpretable AI, adaptive algorithms, and robust training methodologies is essential for the broader acceptance and practical deployment of advanced forecasting techniques. This review systematically examines the evolution, current status, and emerging trends in ML applications for stock price prediction, evaluates mainstream and innovative methodologies, highlights ongoing challenges, and proposes future research directions to enhance predictive accuracy and practical relevance for investors and policymakers.

2. Methodology

2.1. Dataset description

In ML research for stock price prediction, three prominent data sources have become widely utilized due to their comprehensiveness, accessibility, and suitability for diverse financial forecasting tasks. The Yahoo Finance dataset, derived from the Yahoo Finance online portal, is extensively employed due to its broad availability and standardized format. It provides historical price information,

including opening, closing, high, low, volume, and adjusted close prices. It encompasses decades of data across numerous global companies, making it ideal for short-term trend prediction and long-term price forecasting tasks. The Alpha Vantage dataset, accessible via a publicly available API, offers a rich collection of intraday stock prices, technical indicators, and fundamental financial metrics across global markets. Its real-time data coverage and extensive technical analysis tools are particularly suitable for developing high-frequency trading algorithms and volatility modelling [8]. Finally, Quandl (recently rebranded as Nasdaq Data Link) aggregates financial datasets from multiple verified sources, featuring stocks, macroeconomic indicators, commodities, and derivatives. Its extensive historical depth and rigorous data curation make it valuable for sophisticated analyses, including portfolio optimisation, economic impact studies, and long-term forecasting tasks. These datasets collectively support comprehensive research into predictive modelling, market dynamics analysis, and algorithmic financial decision-making.

2.2. ML and DL techniques in stock price prediction

ML and DL have become powerful tools for predicting stock market trends and price fluctuations. This paper reviews several prominent techniques, summarising their applications and the datasets utilised. ARIMA, a traditional time-series forecasting method, has often been applied to index data such as historical data. It uses historical price sequences to fit parameters and forecast future price movements. Researchers typically assess ARIMA models based on metrics such as Root Mean Square Error (RMSE) and Mean Absolute Error (MAE), highlighting their effectiveness in short-term trend forecasting, especially when integrated with volatility models like Generalised Autoregressive Conditional Heteroskedasticity (GARCH).

SVMs have been extensively utilised in stock prediction tasks, primarily addressing non-linear classification and regression problems. Typically, SVM models employ historical price data and technical indicators from financial platforms. Researchers frequently select pertinent technical indicators as input features to predict short-term price fluctuations or classify stock movement directions. The evaluation of SVM performance generally involves classification accuracy and error metrics. Although SVMs offer consistent performance, they may underperform compared to ensemble methods when processing extensive datasets. RFs, a prominent ensemble learning method derived from DTs, have similarly been explored for stock prediction, leveraging comparable datasets. These models handle high-dimensional data and significantly mitigate overfitting, enhancing predictive reliability. Studies indicate RFs consistently surpass SVMs in prediction accuracy, particularly on large-scale datasets, as assessed through accuracy, RMSE, and MSE metrics. DTs commonly serve as baseline models due to their simplicity and interpretability. DTs facilitate transparent short-term forecasting by utilising historical datasets. Nonetheless, single DTs frequently experience overfitting and comparatively lower predictive performance than ensemble or neural network-based models. Advanced DL techniques, specifically RNNs such as LSTM and GRU, represent state-of-the-art approaches for handling sequential stock data. LSTM captures long-term temporal dependencies through sophisticated gating mechanisms, significantly improving forecasting accuracy. GRU, simpler yet computationally efficient, delivers comparable predictive performance, evaluated by RMSE and MAPE. CNNs, traditionally utilised in image processing, have effectively adapted to financial sequence prediction tasks, capturing short-term temporal patterns and demonstrating superior predictive sensitivity and lower errors compared to RNN-based models. Future research should explore integrated hybrid methods to exploit these complementary advantages.

2.3. Long short-term memory networks (LSTM) for stock market prediction

LSTM networks (shown in Figure 1), a specialised variant of RNNs, effectively address traditional RNN limitations concerning long-term dependencies. Utilising gating mechanisms—input, forget, and output gates—LSTMs control information flow, selectively retaining relevant historical patterns and mitigating the vanishing gradient problem common in standard RNNs. In stock market prediction, careful data preprocessing is crucial. Historical stock prices, typically sourced from platforms such as Yahoo Finance, cover extended periods (e.g., daily closing prices of Apple and Microsoft from 2010 to 2021). Data normalisation to a [0,1] scale using Min-Max normalization enhances training efficiency and accuracy. Subsequently, a sliding window technique transforms time-series data into sequences, where past n-day prices predict the next day's price. Datasets commonly employ an 80%-20% training-testing split. Empirical studies consistently report LSTM's superior predictive accuracy. For instance, predicting Apple stock prices yielded an RMSE of approximately 0.237 for LSTM, significantly outperforming ARIMA's RMSE of 3.261. Similarly, predicting Amazon's stock prices, the LSTM model achieved an MAPE of around 1.84%. Comparative analyses further highlight LSTM's advantages over linear models like ARIMA and CNNs, primarily due to its capacity to capture nonlinear sequential dependencies. Hybrid models combining CNNs and LSTM architectures enhance predictive performance by leveraging CNNs' local feature extraction with LSTM's sequential modelling capabilities. Furthermore, Bidirectional LSTM (Bi-LSTM), capturing temporal patterns from past and future contexts, consistently outperforms standard LSTMs in various stock prediction scenarios [9]. These findings confirm that LSTM networks substantially enhance stock market prediction with their distinctive gating mechanisms. Future research may explore further model optimisation through additional predictive variables and advanced hybrid architectures.
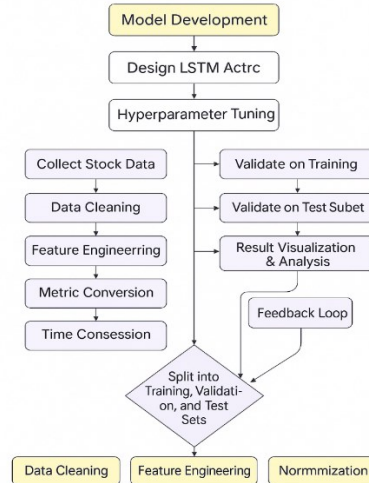

Figure 1: Flowchart of LSTM (photo credit: original)

2.4. CNNs in stock market prediction

CNNs, renowned for their effectiveness in image and natural language processing, have recently shown promising results in stock market prediction tasks (shown in Figure 2). CNNs leverage their local perception and weight-sharing properties, significantly reducing the number of parameters and enhancing computational efficiency and predictive accuracy. A typical CNN-based stock prediction model comprises an input, convolutional, pooling, flattening, and fully connected layers. The convolutional layers are crucial for extracting relevant temporal features from time-series stock data, while pooling layers reduce dimensionality and help control overfitting. Depending on model

complexity, the extracted features are flattened into a vector and passed to subsequent dense or recurrent layers. Empirical research supports CNN-based hybrid models' superior performance. Lu developed a CNN-Bi-directional Long Short-Term Memory (CNN-BiLSTM) model integrated with an attention mechanism (CNN-BiLSTM-AM) for predicting next-day stock prices. Their model outperformed standard methods (MLP, CNN, RNN, LSTM) and other hybrid approaches (CNN-LSTM, BiLSTM-AM), achieving the lowest MAE (21.952) and RMSE (31.694) [10]. Similarly, Rezaei introduced a hybrid CEEMD-CNN-LSTM model, employing Empirical Mode Decomposition to enhance CNN feature extraction, significantly reducing prediction errors compared to traditional CNN and LSTM models [11]. Additionally, researchers achieved notable improvements in forecasting indices the National Association of Securities Dealers Automated Quotations (NASDAQ), outperforming baseline methods [12]. Furthermore, applications such as Graph Convolutional CNN (GC-CNN) and CNN-based Bitcoin price predictions by Cavalli & Amoretti reinforce CNNs' versatility and robust predictive capabilities. These findings affirm CNN-based models' substantial potential in stock market forecasting.
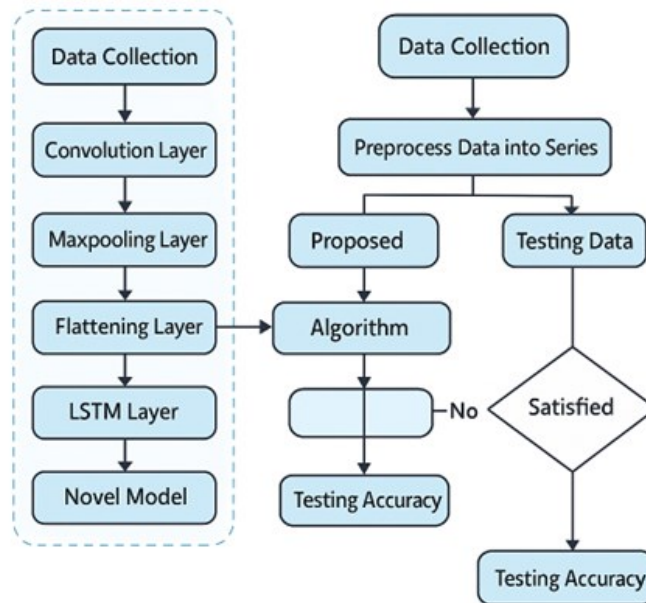

Figure 2: Flowchart of CNNs (photo credit: original)

3. Results and discussion

3.1. Results analysis

Each architecture demonstrates distinct advantages based on the collected experimental data evaluating LSTM and CNN models in stock price prediction tasks. LSTM models consistently exhibit superior performance in capturing temporal dependencies due to their ability to maintain long-term memory. They effectively recognize sequential patterns and trends over extended periods, as evidenced by higher R^2 values and lower error metrics such as MAE and RMSE on datasets like the Shanghai Composite index. LSTM's gating mechanisms offer robust predictions by selectively memorising relevant historical data in volatile market behaviour scenarios.

Conversely, CNN models excel at identifying local patterns within stock data due to their convolutional filters, resulting in efficient extraction of critical short-term features. CNN architectures often require significantly shorter training times than LSTM, as indicated by their rapid convergence and lower computational overhead, providing a notable efficiency advantage, as shown in Table 1. In

summary, LSTM is advantageous when long-term historical trends and temporal dependencies dominate market predictions, whereas CNNS offer strengths in computational efficiency and precise detection of localised market signals. Integrating both architectures, such as in CNN-LSTM hybrid models, often yields optimal predictive performance by leveraging their complementary strengths.

Table 1: Comparative performance of DL models on stock price prediction task

Datas	Model	MAE	MSE	RMSE	R^2	SOURCE
Apple Inc	CNN	1.7869		6.671	0.9979	scitepress.org
Apple Inc	LSTM	34.3995		43.4755	0.7315	scitepress.org
Nike Inc.	CNN	2,5883		11.6657	0.9244	scitepress.org
Nike Inc.	LSTM	7.3367		9.5071	0.2957	scitepress.org
Nike Inc.	CNN-LSTM	7.0301		9.0999	0.6696	scitepress.org
BBVA	CNN+LSTM	0.008967	0.000159	0.0126	0.9346	atlantis-press.com
BBVA	CNN+LSTM+CAM	0.007813	0.000106	0.0103	0.9565	atlantis-press.com

3.2. Discussion

Stock price prediction utilising DL models such as LSTM networks and CNNs has attracted significant interest due to their effectiveness in modelling complex financial dynamics. LSTMs, specialised RNNs, adeptly capture long-term temporal dependencies through gating mechanisms, align well with sequential financial data, and provide superior performance, especially for extended forecasting periods. However, their computational complexity and susceptibility to overfitting present practical challenges, particularly with limited or noisy datasets. Conversely, CNNs employ convolutional filters to identify local features within data sequences, efficiently capturing short-term market fluctuations and repetitive patterns. Their computational efficiency, faster training, and parallelisation scalability frequently enable CNNs to outperform LSTMs in short-term, high-frequency predictions. Nonetheless, CNNs inherently struggle with long-range temporal dependencies unless enhanced by advanced adaptations such as dilated convolutions. Hybrid CNN-LSTM models integrate CNNs' rapid feature extraction capabilities with LSTMs' proficiency in temporal modelling, consistently demonstrating improved predictive accuracy and robustness. Attention mechanisms enhance these models by dynamically highlighting significant historical events, increasing interpretability, and predicting performance. Despite these advancements, challenges persist, notably financial markets' non-stationarity, regime shifts, and inherent noise, limiting model generalisation. Practical mitigation strategies include rolling-window retraining, transfer learning, and integrating external contextual data, such as news sentiment. Model interpretability remains another critical challenge, highlighting the necessity for explainable AI frameworks. Future research should explore transformer-based architectures with self-attention mechanisms, graph neural networks to model inter-stock relationships, and multimodal approaches combining textual, economic, and sentiment data. Practical applications extend to high-frequency trading, portfolio optimisation, volatility forecasting, and regime detection, emphasising the continued need to address non-stationarity, interpretability, and noise reduction to advance stock market forecasting capabilities.

4.   Conclusion

DL algorithms significantly influence modern financial technology, particularly in developing stock price prediction models. This paper reviewed prominent models, including ARIMA, LSTM, and CNNs, and their hybrid approaches. The performance of these models was systematically quantified using error metrics and accuracy calculations. ARIMA models underperformed relative to DL models concerning prediction accuracy. LSTM and hybrid LSTM models demonstrated superior performance in forecasting future stock prices, whereas CNN and hybrid CNN models excelled in predicting stock trends. Extensive experiments confirmed that the hybrid CNN-LSTM model notably improved accuracy by effectively capturing temporal and spatial dependencies in stock data. The CNN-LSTM hybrid model emerged as highly accurate in forecasting future trends and price fluctuations, suitable for portfolio optimisation and intraday trading strategies. Future research will integrate sentiment analysis using Natural Language Processing (NLP) techniques. Social media sentiment from Twitter, Facebook, and financial news websites significantly influences stock trends, particularly in short-term trading. Incorporating sentiment data into deep neural networks promises further enhancement in predicting stock prices and trends, contributing substantially to decision-making and profitability in daily trading.

References

[1]   Mahadik, A., Vaghela, D., & Mhaisgawali, A. (2021). Stock price prediction using lstm and arima. In Second International Conference on Electronics and Sustainable Communication Systems, 1594-1601.

[2]   Illa, P. K., Parvathala, B., & Sharma, A. K. (2022). Stock price prediction methodology using random forest algorithm and support vector machine. Materials Today: Proceedings, 56, 1776-1782.

[3]   ArunKumar, K. E., Kalaga, D. V., Kumar, C. M. S. (2021). Forecasting of COVID-19 using deep layer recurrent neural networks (RNNs) with gated recurrent units (GRUs) and long short-term memory (LSTM) cells. Chaos, Solitons & Fractals, 146, 110861.

[4]   Kirisci, M., & Cagcag Yolcu, O. (2022). A new CNN-based model for financial time series: TAIEX and FTSE stocks forecasting. Neural Processing Letters, 54(4), 3357-3374.

[5]   Sen, J., Mehtab, S. (2022). Stock price prediction using deep learning and natural language processing. Machine Learning in the Analysis and Forecasting of Financial Time Series, 1-28.

[6]   Awad, A. L., Elkaffas, S. M., & Fakhr, M. W. (2023). Stock market prediction using deep reinforcement learning. Applied System Innovation, 6(6), 106.

[7]   Kumar, D., Sarangi, P. K., & Verma, R. (2022). A systematic review of stock market prediction using machine learning and statistical techniques. Materials Today: Proceedings, 49, 3187-3191.

[8]   Raut, K., Kasture, P., Gosavi, C. (2022). Stock market prediction using Alpha Vantage API and machine learning algorithm. International Research Journal of Engineering and Technology.

[9]   Sunny, M. A. I., Maswood, M. M. S., & Alharbi, A. G. (2020). Deep learning-based stock price prediction using LSTM and bi-directional LSTM model. In 2020 2nd novel intelligent and leading emerging sciences conference, 87-92.

[10] Lu, W., Li, J., Li, Y. (2020). A CNN-LSTM-based model to forecast stock prices. Complexity, 2020(1), 6622927.

[11] Rezaei, H., Faaljou, H., & Mansourfar, G. (2021). Stock price prediction using deep learning and frequency decomposition. Expert Systems with Applications, 169, 114332.

[12] Hoseinzade, E., & Haratizadeh, S. (2019). CNNpred: CNN-based stock market prediction using a diverse set of variables. Expert Systems with Applications, 129, 273-285.