ISSN: 2229-6948(ONLINE)
DOI: 10.21917/ijct.2026.0574

A COMPARATIVE STUDY OF LSTM AND TIME SERIES TRANSFORMER MODELS FOR BITCOIN PRICE FORECASTING WITH TECHNICAL INDICATORS AND GOOGLE TRENDS

N. Valliammal and P. Hemashree

Department of Computer Science, Avinashilingam Institute for Home Science and Higher Education for Women, India

Abstract

Cryptocurrency price forecasting is inherently challenging due to pronounced volatility, non-stationary behaviour, and strong sensitivity to market sentiment. In this work, a comparative analysis is conducted between the recurrent Long Short-Term Memory (LSTM) network and the attention-based Time Series Transformer (TST) for predicting Bitcoin prices. Daily Bitcoin price data are modelled using three feature configurations: (i) raw OHLCV variables, (ii) OHLCV combined with technical indicators, and (iii) an extended feature set incorporating public interest measured through Google Trends for the term Bitcoin. To examine temporal sensitivity, the TST model is evaluated across multiple lookback windows of 30, 60, and 90 days. Experimental results reveal that LSTM with OHLCV features achieves the lowest absolute forecasting error (MAE: 3,958; RMSE: 5,621), while feature enrichment degrades performance in both models. The TST demonstrates greater robustness across lookback window configurations, with its MAE varying within a narrower range (10,987–12,406) compared to LSTM under feature enrichment. These findings contribute nuanced insights into the conditions under which attention-based architectures offer practical advantages over recurrent models for cryptocurrency price prediction.

Keywords:

Bitcoin Price Forecasting, Time Series Transformer, Deep Learning, Google Trends, Cryptocurrency

1. INTRODUCTION

Forecasting cryptocurrency prices has gained increasing attention with the rapid expansion of the digital asset market. Bitcoin, as the first and most widely adopted cryptocurrency, is characterized by pronounced price volatility driven by factors such as market microstructure, speculative trading behavior, and public sentiment. These characteristics make accurate prediction particularly challenging, and traditional financial models often struggle to capture the resulting non-linear dynamics. In recent years, deep learning approaches, especially Long Short-Term Memory (LSTM) networks, have been widely implemented to time-series forecasting due to their ability to model sequential dependencies. Despite their success, recurrent architectures suffer from limitations such as gradient instability, sequential computation constraints, and reduced effectiveness in modeling long-range dependencies. To address these challenges, Transformer-based architectures have emerged as a promising alternative, offering parallel processing and attention mechanisms that can better capture long-term interactions. Among these, the Time Series Transformer (TST) adapts the Transformer encoder specifically for continuous numerical time-series data, making it well-suited for financial forecasting tasks. The main contributions of this study are threefold:

- A comparative evaluation of LSTM and Time Series Transformer models for Bitcoin price forecasting, highlighting the strengths and limitations of recurrent and attention-based architectures under identical experimental settings.
- An empirical analysis of feature enrichment, examining the impact of technical indicators and external sentiment information derived from Google Trends on forecasting accuracy.
- A systematic ablation study investigating the effect of different lookback window lengths on model performance, providing insights into how temporal context length affects attention-based forecasting performance.

2. LITERATURE REVIEW

The applications of innovative technologies in forecasting Bitcoin have rapidly risen over the past decade. The promising research works of technological evolutions are summarized,

Sheykhan et al. [13] provide a brief review of Machine learning methods specifically designed for Bitcoin price prediction, which compares the importance of advanced deep learning techniques, LSTM Long Short-Term Memory (LSTM) networks, with various advanced ML techniques in terms of their agility, fidelity, and compatibility. The results have provided better insights to predict the trend, which is greatly useful for market investors, analysts, and decision makers.

Ahmed et al. [14] introduce a novel comparative analysis of ensemble learning and deep learning forecasting models, such as MLP, LSTM, ADABoost, LightGBM, RNN, and GRU, with various cryptocurrencies like Bitcoin, Ethereum, Ripple, and Litecoin. Deep learning models such as GRU, simple recurrent neural network, and Light BGM perform better than other conventional approaches. The author suggests this approach will help investors make the optimum decision for investing in the trade market.

Michael Carter [15] emphasises the importance of transformer-based architectures compared with LSTM. The dataset is a combination of real-time tick-level and minute-level market data, curated from tick-level trade feeds, order book snapshots at multiple depth levels, minute-level aggregated OHLCV series, and public on-chain metrics relevant to Bitcoin, such as transaction volume and mempool size. The study incorporates technical indicators and market microstructure features to simulate real-time market dynamics. Controlled experiments are conducted to evaluate model performance in short-term forecasting tasks relevant to algorithmic trading and risk monitoring**.** The model is evaluated using performance metrics such as prediction accuracy, directional accuracy, root mean squared error, mean absolute error, latency, and robustness to distributional shift. The study evaluates the model under both rapid-response and strict policy scenarios, while also examining practical considerations such as processing latency and computational time associated with each resource.

Longchen Zheng et al. [16] introduce a novel TST to predict Bitcoin's short-term prices. Time series forecasting performs better than LSTM, since traditional models like LSTM often fail in handling linear dependencies and dynamic data. It leverages self-attention mechanisms that seemingly capture complex temporal patterns in historical Bitcoin data, which includes prices and trading volume.

Tripathi et al. [17] present a comprehensive overview of blockchain technologies, examining their evolution, core principles, applications, and associated challenges over time**.** Furthermore, the paper examines the adoption scenarios and practical use cases of blockchain technology in various domains such as finance, IoT, supply chain, and voting, which emphasize both its advantages and the barriers for widespread implementation. By leveraging ongoing innovations like AI and deep learning, we can overcome current limitations by improving the expandability through new architectures like DAG-based systems and strengthening security protocols.

Collectively, these studies establish that deep learning models particularly LSTM networks consistently outperform statistical baselines such as ARIMA for Bitcoin price forecasting, and that Transformer-based architectures show strong potential in high-frequency or feature-rich settings. However, several gaps remain unaddressed. Most existing studies either compare LSTM against statistical baselines or evaluate generic Transformers without considering purpose-built time-series architectures. Very few studies directly compare LSTM against a specialised Time Series Transformer under identical feature sets and controlled experimental conditions. Furthermore, the impact of external sentiment signals such as Google Trends on Transformer-based forecasting remains underexplored, as does the sensitivity of attention-based models to lookback window length. The present study addresses these gaps through a systematic ablation design that isolates the contribution of each feature group and temporal configuration across both architectures. Cryptocurrency prediction studies often rely on deep learning architectures due to their ability to capture non-linear dynamics. Several works have applied LSTMs to Bitcoin forecasting, showing improvement over statistical models such as ARIMA and GARCH. Transformer-based models are relatively less explored in cryptocurrency prediction. Some studies use generic Transformers but overlook specialized time-series architectures. Only a limited number of works incorporate external sentiment features such as Google Trends, and even fewer combine them with Transformer-based forecasting. This gap motivates the present study, which integrates:

- A specialized Transformer encoder (TST)
- A sentiment-driven external factor
- A detailed ablation study to isolate useful features

3. METHODOLOGY

In this study, Long Short-Term Memory (LSTM) and the Time Series Transformer (TST Transformer) were implemented to predict Bitcoin prices.

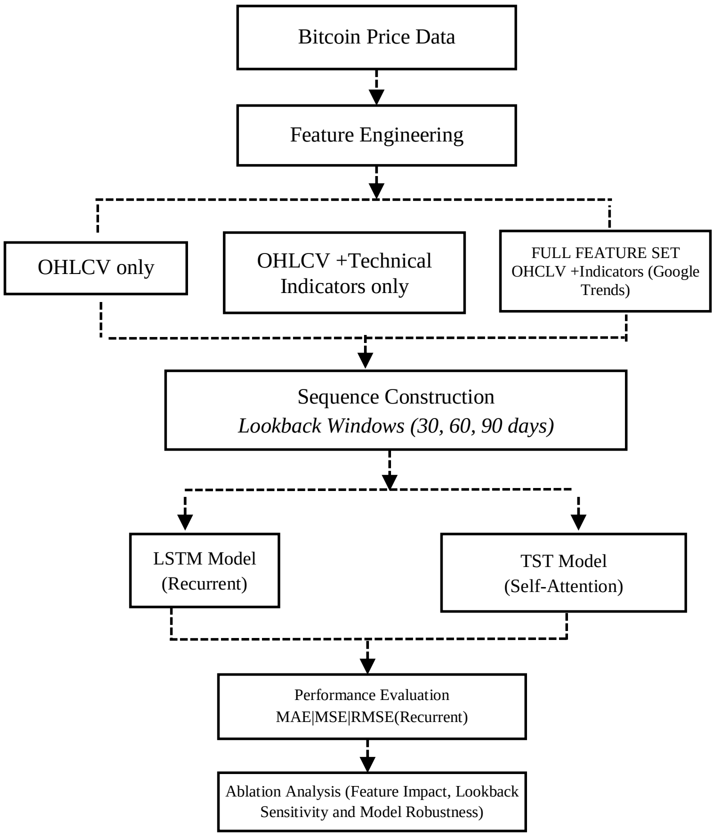

Fig.1. Workflow of forecasting Bitcoin

3.1 DATA COLLECTION

This study utilizes historical Bitcoin price data obtained from Yahoo Finance using the yfinance Python library. The dataset corresponds to the BTC–USD trading pair and spans from January 2020 to January 2025, covering five years of daily observations. This period was selected to ensure train and test data are drawn from the same market regime the post-institutional-adoption era of Bitcoin avoiding distribution shift introduced by the structurally different early market conditions of 2014–2019. Each record includes the standard OHLCV attributes: Open, High, Low, Close, and Volume. These variables represent the fundamental market dynamics and serve as the primary inputs for the forecasting models.

In addition to price-based data, an external sentiment-related variable was incorporated to enhance predictive performance. Public interest in Bitcoin was captured using Google Trends, which provides normalized search intensity scores for the keyword "Bitcoin". Google Trends data was collected over the same time period and resampled to daily frequency. This feature reflects collective market attention and sentiment, which is known to influence cryptocurrency price movements.

3.2 DATA PRE-PROCESSING

All numerical features were scaled using Min–Max normalisation to the range [0,1]to facilitate stable model convergence and to prevent variables with larger absolute magnitudes, such as price and trading volume, from disproportionately influencing the learning process. The scaler was adapted exclusively on the training data and subsequently implemented to the test data, thereby ensuring that no information from future observations was incorporated during feature normalisation. This procedure is essential to avoid data leakage, which would otherwise lead to artificially inflated performance estimates on the test set.

The dataset was partitioned chronologically into 80% training data (January 2020–December 2023, approximately 1,460 observations) and 20% testing data (January 2024–January 2025, approximately 366 observations). A time-ordered split was adopted instead of random sampling to preserve the temporal structure of the financial time series and to eliminate look-ahead bias. This approach more accurately reflects real-world forecasting conditions, where future price information is unavailable at training time. The training period spans both the 2020–2021 bull market and the 2022 bear market, exposing the models to diverse market regimes, while the test period captures the 2024 Bitcoin market cycle, including the ETF approval rally and post-halving price dynamics, thereby providing a rigorous out-of-sample evaluation.

To further mitigate data leakage arising from rolling-window technical indicators, which require an initial warm-up period of up to 30 days, all rows containing missing values generated during indicator computation were removed prior to the train–test split. Sequence windows for model inputs were constructed only after this cleaning step, ensuring that both the training and test sets contained fully defined indicators and that no partially initialised features were used during model training or evaluation.

3.3 FEATURE ENGINEERING

To capture market trends, momentum, and volatility, several widely used technical indicators were computed from the historical price series**.** These indicators are commonly applied in financial analytics, as they are guiding in identifying suitable key market dynamics such as overbought/oversold conditions, trend direction, and volatility [27].

Table.1. Technical Indicators

Indicator	Purpose	Key Parameters
SMA (Simple Moving Average)	Identifies overall trend direction and support/resistance levels	Window sizes: 3, 7, 15, 30, 50
EMA (Exponential Moving Average)	Captures recent price movements with higher sensitivity	Smoothing factor
MACD (Moving Average Convergence Divergence)	Measures momentum and trend reversals	EMA (12, 26), Signal EMA (9)
RSI (Relative Strength Index)	Detects overbought and oversold conditions	Period: 14
Bollinger Bands	Captures price volatility and potential reversals	SMA(20), ±2 standard deviations

Technical indicators are quantitative measures derived from historical price, volume, and open–close data of financial assets, such as stocks and cryptocurrencies. These curated indicators assist in identifying the market trends, momentum patterns, and potential reversal points in the market which enables a more detailed analysis of Bitcoin price. These techniques assist traders, investors, market analysts, and researchers in evaluating market strength and supporting informed trading decisions. They include:

- Simple Moving Average (SMA)
- Exponential Moving Average (EMA)
- Moving Average Convergence Divergence (MACD) and signal line
- Relative Strength Index (RSI)
- Bollinger Bands (upper, middle, and lower bands)

These indicators provide complementary information about short-term and long-term market behavior. Along with the OHLCV variables and Google Trends scores, they form the enhanced feature representation used in the experiments.

3.4 TIME SERIES FORECASTING

3.4.1 LSTM:

The Long Short-Term Memory (LSTM) network was employed as a recurrent baseline model for Bitcoin price forecasting due to its ability to address the vanishing gradient problem and capture long-range temporal dependencies through gated memory mechanisms. In this study, the LSTM model comprises multiple stacked LSTM layers with dropout regularization to mitigate overfitting, followed by a fully connected layer that outputs the predicted closing price**.** The model was implemented using the TensorFlow (Keras) framework and optimized using the Adam optimizer with Mean Squared Error (MSE) as the loss function.

The implementation process begins with data acquisition from Yahoo Finance, where historical OHLCV data are collected and enriched through feature engineering by computing technical indicators such as SMA, EMA, RSI, MACD, and Bollinger Bands. In addition to price-based features, an external sentiment-related variable derived from Google Trends search interest for the term "Bitcoin" is incorporated to capture public attention and market sentiment**.** The dataset is cleaned to address missing values introduced by rolling-window calculations, after which all features are stabilized using Min–Max scaling to stabilize the training process. The normalized time series is then transformed into overlapping sequential windows, where historical observations over a fixed lookback period serve as input features and the subsequent time step represents the prediction target, resulting in a three-dimensional input structure comprising samples, sequence length, and features. The data is chronologically divided into training, validation, and testing sets using an 80:20 split to preserve temporal integrity. All hyperparameters were fixed prior to training and held constant across all experimental configurations to ensure a controlled and reproducible comparison. The LSTM model comprises of two stacked LSTM layers 64 and 32 units respectively each followed by a dropout layer with the rate of rate=0.2 to reduce overfitting, and a fully connected output layer. Model training incorporates early stopping to prevent overfitting, and the trained LSTM is subsequently used to generate forecasts on unseen test data, with predictions inverse-transformed to the original price scale for performance evaluation and analysis.

3.4.2 Time Series Transformer:

To overcome the limitations of recurrent architectures such as LSTM in modeling long-range dependencies, a Time Series Transformer (TST) model was implemented for Bitcoin price forecasting. Unlike recurrent networks, TST leverages a self-attention mechanism that enables the model to directly capture relationships among all time steps within an input sequence, allowing efficient learning of both short-term fluctuations and long-term temporal patterns. The TST architecture adopts an encoder-only Transformer design, specifically tailored for continuous, multivariate numerical time series data, thereby avoiding the unnecessary decoder component used in vanilla Transformers. Multivariate input sequences are first segmented into fixed-length windows and linearly projected into higher-dimensional embeddings, after which positional encodings are added to preserve temporal ordering. Each encoder block consists of multi-head self-attention layers that model inter-temporal dependencies across the entire sequence, followed by feed-forward networks, residual connections, and layer normalization to ensure stable learning and effective gradient propagation. The final encoder representations are aggregated and passed through a fully connected regression head to generate future price forecasts. The TST model was trained using the Adam optimizer and optimized with loss functions such as Mean Squared Error (MSE), while evaluation was performed using the same metrics as the LSTM model to ensure a fair comparison. Owing to its attention-based global pattern recognition capability and efficient handling of high-dimensional inputs, the TST is designed to capture nonlinear dynamics and long-term trends across multivariate input sequences, making it a strong candidate for complex, feature-rich forecasting environments. The TST model uses three encoder blocks, each with 4 attention heads, a key dimension of 64, and a feed-forward dimension of 128, followed by a dropout layer (rate=0.1) and a two-layer dense regression head. Training for Both models were trained using the Adam optimiser with an initial learning rate of 0.001, a batch size of 64. The training process is limited to 80 epochs with early stopping (patience=10) based validation loss.

3.4.3 Ablation Study Design:

A structured ablation study was conducted to evaluate the contribution of different feature sets and temporal configurations. The experiments include:

1) **Feature ablation**
   a) OHLCV features only
   b) OHLCV + technical indicators
   c) OHLCV + technical indicators + Google Trends
2) **Model comparison**
   a) LSTM vs TST under identical experimental conditions
3) **Lookback window analysis (TST)**
   a) 30-day, 60-day, and 90-day input sequences
   b) This systematic evaluation enables a clear assessment of how attention mechanisms, sentiment features, and temporal context influence forecasting accuracy.

3.4.4 Performance Evaluation:

Model performance was assessed using Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE) to evaluate the accuracy and robustness of the forecasting models. MAE measures the average magnitude of prediction errors without considering their direction, providing an intuitive understanding of overall error. MSE assigns greater weight to larger errors by squaring the differences between predicted and actual values, making it particularly useful for penalizing large deviations. RMSE, as the square root of MSE, retains the same unit as the target variable and offers a balanced measure that reflects both variance and magnitude of errors. Lower values of these metrics indicate better forecasting performance across different experimental settings.

4. RESULTS AND DISCUSSION.

This section presents the experimental results of the LSTM and Time Series Transformer (TST) models across multiple feature configurations and lookback window lengths. The findings reveal several counter-intuitive patterns that challenge the assumption of Transformer superiority in time-series forecasting tasks on limited daily cryptocurrency data.

Table.2. Performance Comparison of LSTM and TST Models (60-Day Lookback)

Model	Feature Set	MAE	MSE	RMSE
LSTM	OHLCV	3958	31596881	5621
LSTM	Full (OHLCV + Indicators + GT)	10832	171727479	13104
TST	OHLCV	7161	137661763	11733
TST	Full (OHLCV + Indicators + GT)	12406	225246984	15008

The Table.2 presents the performance of LSTM and TST across OHLCV and full feature configurations using a 60-day lookback window. The LSTM with raw OHLCV features achieves the lowest error across all experiments, with an MAE of 3,958 and RMSE of 5,621.

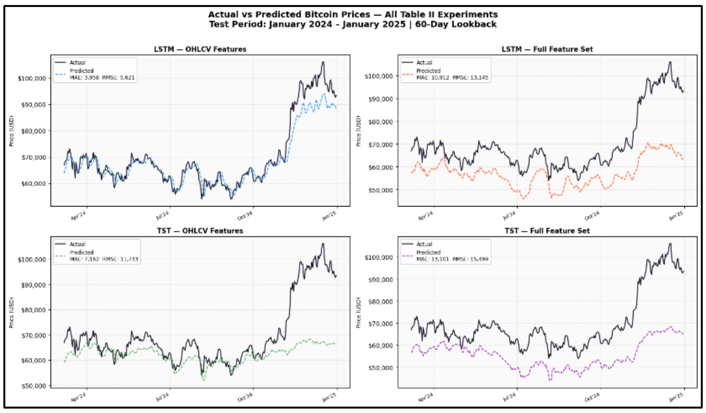

Fig.2. Actual vs Predicted Bitcoin prices during Jan 2024-Jan 2025

This result is consistent with established literature showing that recurrent architectures are data-efficient and well-suited to sequential datasets where local temporal dependencies dominate.

The training dataset spans approximately 1,460 daily observations a scale sufficient for LSTM to capture short-range price autocorrelation but insufficient for the TST's attention mechanism to fully leverage its global dependency modelling capability. The TST model using OHLCV features achieves an MAE of 7,161, which is competitive but not superior. This result indicates that, for the given data volume, the sequential inductive bias of the LSTM provides a measurable performance advantage. The Fig.2 indicates that the Actual vs Predicted Bitcoin prices during Jan 2024-Jan 2025.

A consistent yet counter-intuitive result is that enriching the feature set with technical indicators and Google Trends data degrades performance in both architectures. For the LSTM, the MAE increases from 3,958 to 10,832 a 174% rise when moving from OHLCV features to the full feature set. Similarly, the TST’s MAE increases from 7,161 to 12,406, representing an 83% relative increase. This degradation can be attributed to several factors. Most technical indicators, such as SMA, EMA, and Bollinger Bands, are derived directly from closing prices and therefore exhibit high collinearity with the target variable, contributing redundant rather than independent predictive information. In addition, Google Trends data are available at a weekly frequency and are forward filled to daily intervals, which introduces artificial smoothness and weak alignment with daily price movements. Notably, the smaller relative performance drop observed in the TST (83% versus 174% for the LSTM) suggests that its multi-head attention mechanism and dropout regularisation are more effective at mitigating the impact of redundant multivariate inputs. This indicates a potential architectural advantage of the TST in feature-rich settings.

Table.4. Lookback Window Analysis for TST Model (Full Feature Set)

Lookback Window (Days)	MAE	MSE	RMSE
30	10987	189148575	13753
60	12406	225246984	15008
90	11113	184797025	13594

The Table.3 reports the TST performance across 30-, 60-, and 90-day lookback windows using the full feature set. The results show limited variation across configurations, with MAE values ranging between 10,987 for the 30-day window to 12,406 for the 60-day window, yielding a total spread of 1,419. While the 30-day window achieves the lowest MAE (10,987), the 90-day window attains the lowest MSE (184,797,025), indicating that no single lookback length consistently outperforms the others across all evaluation metrics.

This relative stability across temporal settings reflects an important characteristic of the self-attention mechanism. By assigning weights to all timesteps in the input sequence simultaneously, the TST can dynamically adjust its effective receptive field rather than relying on a fixed sequential memory. In contrast, the LSTM exhibits a sharper performance degradation under feature enrichment, highlighting its greater sensitivity to increased input dimensionality. Consequently, the TST demonstrates stronger robustness across different lookback configurations, making it a more predictable and stable choice for practical deployment, even when it does not achieve the lowest absolute error under minimal feature conditions.

5. CONCLUSION

This study presents a comparative analysis of Long Short-Term Memory (LSTM) and Time Series Transformer (TST) models for Bitcoin price forecasting using daily data. The models are evaluated across multiple feature configurations and lookback window lengths to examine how feature enrichment and temporal context interact with model architecture to influence forecasting accuracy. The feature sets include raw OHLCV data, technical indicators, and an external sentiment proxy derived from Google Trends.

Three key findings emerge from the experimental results. First, the LSTM model using only OHLCV features achieves the lowest forecasting error across all configurations (MAE: 3,958; RMSE: 5,621), indicating that recurrent architectures retain a practical advantage for daily cryptocurrency data with moderate sample sizes. Second, enriching the feature set with technical indicators and Google Trends data leads to performance degradation in both models. However, the relative increase in error is substantially larger for the LSTM (174%) than for the TST (83%), suggesting that the TST’s attention mechanism is more effective at attenuating redundant or weakly informative multivariate inputs. Third, the TST exhibits greater robustness across different lookback window lengths, with MAE values confined to a narrower range (10,987–12,406), whereas the LSTM shows sharper degradation under full feature conditions.

Overall, these findings suggest that model selection for cryptocurrency price forecasting should be informed by dataset size and feature complexity rather than architectural preference alone. While the LSTM's sequential inductive bias provides a clear advantage on limited daily datasets, the TST's relative stability under feature enrichment and varying temporal contexts makes it a more predictable choice in feature-rich environments. Future work should explore TST performance on higher-frequency Bitcoin data, such as hourly or minute-level series, where larger training sets and longer effective sequences may better leverage the attention mechanism's capacity for modelling long-range dependencies. Additional extensions include the application of feature selection techniques, multi-asset forecasting frameworks, and the integration of on-chain blockchain metrics.

REFERENCES

[1] P. Kayal and P. Rohilla, “Bitcoin in the Economics and Finance Literature: A Survey”, *SN Business and Economics*, Vol. 1, No. 7, pp. 1-24, 2021.

[2] M. Vejacka, “Basic Aspects of Cryptocurrencies”, *Journal of Economy, Business and Financing*, Vol. 2, pp. 75-83, 2014.

[3] M. Seo and G.W. Kim, “Hybrid Forecasting Models based on the Neural Networks for the Volatility of Bitcoin”, *Applied Sciences*, Vol. 10, No. 14, pp. 47-68, 2020.

[4] H. Tayib and A.M. Abdulazeez, “A Review of Bitcoin Price Prediction based on Deep Learning Algorithms”, *Indonesian Journal of Computer Science*, Vol. 13, No. 2, pp. 1-15, 2024.

[5] A. Agarwal and Y. Agarwal, “AI-Driven Economic and Financial Forecasting: House Prices, Unemployment, Cryptocurrency and Business Stability in the USA”,

International Journal of Science and Social Science Research, Vol. 2, No. 4, pp. 228-237, 2025.

[6] M. Jimenez-Carrion and G.A. Flores-Fernandez, "Machine Learning Algorithms in Predicting Prices in Volatile Cryptocurrency Markets", *HighTech and Innovation Journal*, Vol. 6, No. 1, pp. 257-272, 2025.

[7] N.S. Venkatesh, N.B. Rashmitha, N.S. Manjunadha and N.M. Junaid, "A Survey on Cryptocurrency Price Prediction", *World Journal of Advanced Research and Reviews*, Vol. 22, No. 1, pp. 375-380, 2024.

[8] W. Li and K.L.E. Law, "Deep Learning Models for Time Series Forecasting: A Review", *IEEE Access*, Vol. 12, pp. 92306-92327, 2024.

[9] A. Lim and S. Zohren, "Time-Series Forecasting with Deep Learning: A Survey", *Philosophical Transactions of the Royal Society A*, Vol. 379, No. 2194, pp. 1-13, 2021.

[10] Q. Wen, T. Zhou, C. Zhang, W. Chen, Z. Ma, J. Yan and L. Sun, "Transformers in Time Series: A Survey", *Proceedings of International Conference on Artificial Intelligence*, Vol. 32, pp. 1-9, 2022.

[11] E. Mahdi, C. Martin-Barreiro and X. Cabezas, "A Novel Hybrid Approach using an Attention-based Transformer + GRU Model for Predicting Cryptocurrency Prices", *Mathematics*, Vol. 13, No. 9, pp. 1-11, 2025.

[12] M.J. Hamayel and A.Y. Owda, "A Novel Cryptocurrency Price Prediction Model using GRU, LSTM and Bi-LSTM Machine Learning Algorithms", *Artificial Intelligence*, Vol. 2, No. 4, pp. 477-496, 2021.

[13] P. Boozary, S. Sheykhan and H. GhorbanTanhaei, "Forecasting the Bitcoin Price using Various Machine Learning: A Systematic Review in Data-Driven Marketing", *Systems and Soft Computing*, Vol. 7, pp. 1-13, 2025.

[14] A. Bouteska, M.Z. Abedin, P. Hajek and K. Yuan, "Cryptocurrency Price Forecasting-A Comparative Analysis of Ensemble Learning and Deep Learning Methods", *International Review of Financial Analysis*, Vol. 92, pp. 1-7, 2024.

[15] M.J. Carter, E.R. Hayes, A.D. Morrison, R.L. Turner, M. Luke and C. James, "Evaluation of LSTM and Transformer-based Neural Networks for Real-Time Bitcoin Prediction", *Proceedings of International Conference on Artificial Intelligence*, Vol. 8, pp. 1-8, 2024.

[16] L. Zheng, "Time-Series Transformer for Predicting Bitcoin Price", *Highlights in Business Economics and Management*, Vol. 54, pp. 173-181, 2025.

[17] G. Tripathi, M.A. Ahad and G. Casalino, "A Comprehensive Review of Blockchain Technology: Underlying Principles and Historical Background with Future Challenges", *Decision Analytics Journal*, Vol. 9, No. 1, pp. 1-21, 2023.

[18] N.D.Y.R. Reddy and N.P. Gangothri, "Trend Analysis by using a Simple Moving Average", *International Journal of Latest Technology in Engineering Management and Applied Science*, Vol. 14, No. 8, pp. 177-191, 2025.

[19] M.B. Perry, "The Exponentially Weighted Moving Average", *Wiley Encyclopedia of Operations Research and Management Science*, Vol. 18, No. 4, pp. 1-16, 2010.

[20] M. Melda, M. Mira and N. Nurlina, "Analysis of the Effectiveness of RSI and MACD Indicators in Addressing Stock Price Volatility", *Wiga Journal of Economic Research*, Vol. 15, No. 1, pp. 71-79, 2025.

[21] N.F.I.P. Irwansyah, A. Jaya and N. Nasrullah, "Effectiveness of RSI and Bollinger Bands in Identifying Buy and Sell Signals", *Wiga Journal of Economic Research*, Vol. 15, No. 1, pp. 60-70, 2025.

[22] H. Wang, R. Yao, L. Hou, J. Zhao and X. Zhao, "A Methodology for Calculating the Contribution of Exogenous Variables to ARIMAX Predictions", *Proceedings of International Conference on Artificial Intelligence*, Vol. 12, pp. 1-12, 2021.

[23] M. Majka, "ARIMAX Time Series Forecasting with External Variables", *Proceedings of International Conference on Artificial Intelligence*, Vol. 21, pp. 1-9, 2024.

[24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez and I. Polosukhin, "Attention is all you Need", *Advances in Neural Information Processing Systems*, Vol. 30, pp. 1-11, 2017.

[25] S. Liu, H. Yu, C. Liao, J. Li, W. Lin, A.X. Liu and S. Dustdar, "Pyraformer: Low-Complexity Pyramidal Attention for Long-Range Time Series Modeling and Forecasting", *Proceedings of International Conference on Learning Representations*, Vol. 2, pp. 1-7, 2022.

[26] H.S. Jung, H. Lee and J.H. Kim, "Detecting Bitcoin Sentiment: leveraging Language Model Applications in Sentiment Analysis for Bitcoin Price Prediction", *Neural Processing Letters*, Vol. 57, No. 5, pp. 1-25, 2025.

[27] F.B. Oriani and G.P. Coelho, "Evaluating the Impact of Technical Indicators on Stock Forecasting", *Proceedings of International Symposium Series on Computational Intelligence*, Vol. 8, pp. 1-8, 2016.

[28] J. Ayala, T.M. Garcia, J.L.V. Noguera, V.F. Gomez and F. Divina, "Technical Analysis Strategy Optimization using a Machine Learning Approach in Stock Market Indices", *Knowledge-based Systems*, Vol. 225, pp. 1-9, 2021.

[29] A.K. Kang, "Improving MACD Technical Analysis by Optimizing Parameters and Modifying Trading Rules: Evidence from the Japanese Nikkei 225 Futures Market", *Journal of Risk and Financial Management*, Vol. 14, No. 1, pp. 3-7, 2021.

[30] S. Dongrey, "Study of Market Indicators used for Technical Analysis", *International Journal of Engineering and Management Research*, Vol. 12, No. 2, pp. 64-83, 2022.