IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Algorion: A Deep Learning Framework For Stock Price Prediction

Authors: Kabir Khatri, Parth Vaghela, Kanishka Sharma, Misbah Khan Pathan

Designation: AI/ML Developer, SDE/IOS Developer,

Computer Science & Engineering

Parul University.

Acknowledgements

The authors would like to express their gratitude to their mentors, **Meet Patel**, **Aamir Jarda**, and **Sweety Patel**, for their invaluable guidance and support throughout this research project.

Abstract

The inherent volatility of financial markets presents a significant challenge to accurate price forecasting. This paper introduces **Algorion**, a novel deep learning framework for stock price prediction. Algorion utilizes a Long Short-Term Memory (LSTM) network trained on an extensive set of historical, technical, and fundamental data to forecast future price movements. To evaluate the model's effectiveness, we conduct a rigorous backtesting experiment on a diverse portfolio of stocks from the NIFTY 50 index over a ten-year period. The results demonstrate that the Algorion framework achieves superior prediction accuracy (measured by RMSE and MAE). Furthermore, when its predictions are used to generate systematic trading signals, the resulting strategy yields significantly higher risk-adjusted returns (measured by Sharpe and Sortino ratios) compared to traditional baseline strategies. An ablation study further quantifies the value-add of integrating a rich feature set, which improves performance over a basic LSTM model. Our findings establish the efficacy of sophisticated deep learning systems in creating more robust and accurate stock price prediction models.

Keywords: Stock Price Prediction, Deep Learning, LSTM, Time-Series Forecasting, Computational Finance, Financial Forecasting.

1. Introduction

1.1. The Duality of Modern Financial Markets: Opportunity and Volatility

Modern financial markets are characterized by a fundamental duality: they present unprecedented opportunities for wealth generation while simultaneously exhibiting extreme levels of complexity and volatility. The rapid flow of information, globalization of capital, and interconnectedness of economies have created an environment where market conditions can shift dramatically in fractions of a second. This high-velocity, high-volatility landscape makes it exceedingly difficult for traders and investors to make consistently informed and timely decisions.[2] The sheer volume of data—from price ticks and trading volumes to news articles and macroeconomic reports—has surpassed the capacity for manual human processing, creating a demand for more sophisticated analytical and predictive tools.

1.2. The Rise of AI in Financial Forecasting

In response to these challenges, the financial industry has witnessed a seismic shift towards data-driven analysis. The latest evolution of this paradigm involves the integration of Artificial Intelligence (AI), particularly machine learning (ML) and deep learning (DL) models, for financial forecasting. These technologies have transformed financial analytics by enabling systems to analyze vast, multidimensional datasets to identify complex, non-linear patterns and correlations that are invisible to traditional statistical models.[2] Advanced architectures can process historical market data, technical indicators, and even unstructured data from news sentiment to generate highly precise price predictions, enhancing accuracy and providing a systematic, data-driven basis for decision-making.

1.3. The "Black Box" Dilemma and the Trust Deficit in AI Models

Despite their analytical power, the widespread adoption of advanced AI systems in finance faces a significant obstacle: the "black box" problem. Many sophisticated DL models, such as deep neural networks, operate in a manner that is not readily interpretable by human users. This opacity creates a critical trust deficit, which is a major barrier to adoption, particularly for institutional investors and fund managers who bear significant fiduciary responsibilities. The inability to understand *why* a model made a particular forecast makes it difficult to rely on it, especially during unprecedented market events where the model's training data may no longer be representative. This connects the practical challenge of forecasting to a central theme in contemporary AI research: the need for Explainable AI (XAI).[2] This suggests that the next generation of financial models must be engineered not only for predictive accuracy but also for robustness and trustworthiness.

1.4. Contribution: The Algorion Price Prediction Framework

This paper introduces the **Algorion** framework, a novel system designed to address the challenges of accurate market forecasting. Algorion is architected as a deep learning system that leverages the computational power of LSTMs to predict stock price movements. The framework's primary contribution is its end-to-end design, which integrates a rich, multi-source feature set—combining technical and fundamental analysis—with a robust LSTM-based forecasting engine.[2] This approach aims to create a high-performance, data-driven prediction model capable of navigating complex market conditions.

The central thesis of this work is that a sophisticated deep learning framework, integrating a rich set of technical and fundamental features, can achieve superior price prediction accuracy. To validate this claim, we present a comprehensive backtesting study on a portfolio of Indian stocks, evaluating both the model's predictive accuracy and the financial performance of a trading strategy derived from its predictions. The remainder of this paper is organized as follows: Section 2 reviews the relevant literature. Section 3 details the methodology and architecture of the Algorion framework. Section 4 describes the experimental design and metrics. Section 5 presents and discusses the empirical results. Finally, Section 6 concludes the paper and outlines directions for future research.

2. Literature Review and Theoretical Foundations

2.1. Deep Learning Architectures for Financial Time-Series Forecasting

The prediction of stock prices, a classic time-series forecasting problem, has been a fertile ground for the application of various deep learning models. Early approaches often utilized Recurrent Neural Networks (RNNs), but these models struggled with the vanishing and exploding gradient problems, limiting their ability to capture long-term dependencies in financial data. This led to the widespread adoption of more sophisticated architectures like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks.[2] LSTMs, with their gating mechanism (input, forget, and output gates), are specifically designed to remember information over long periods, making them well-suited for financial market analysis. However, as noted in the literature, LSTMs can be prone to overfitting and require significant memory and computational resources.[2]

Recent research has explored even more complex models. Ensemble methods, which combine predictions from multiple models (e.g., LSTM, GRU, and LightGBM), have shown promise but often come with high computational costs and a lack of real-time adaptability.[2] Convolutional Neural Networks (CNNs), traditionally used for image processing, have been adapted for financial forecasting and have demonstrated strong performance, though their interpretability remains a challenge.[2] More recently, Transformer-based models, with their attention mechanisms, have outperformed LSTMs in capturing long-range dependencies but require vast amounts of data and are computationally intensive.[2] This review of the landscape indicates that a well-tuned LSTM remains a robust and practical choice for a core forecasting engine, particularly when its inherent weaknesses are mitigated by other components within a larger system architecture, as proposed in Algorion.

2.2. Feature Engineering and Data Fusion in Financial Prediction

The performance of any predictive model is critically dependent on the quality and richness of its input features. Financial forecasting research has identified three primary categories of data that, when combined, can lead to more robust and accurate models.

1. **Technical Indicators:** These are heuristic- or pattern-based signals derived from historical price and volume data. Indicators such as Moving Averages (MA), the Relative Strength Index (RSI), and Moving Average Convergence Divergence (MACD) are widely used to gauge market momentum, volatility, and trend strength.[2] They form the backbone of many trading strategies.
2. **Fundamental Analysis:** This involves examining a company's underlying financial health and intrinsic value. Key metrics include the Price-to-Earnings (P/E) ratio, Earnings Per Share (EPS), and revenue data.[2] Integrating fundamental data provides a longer-term context that technical indicators alone may miss.
3. **Sentiment Analysis:** With the proliferation of online media, sentiment analysis has emerged as a powerful tool. By applying Natural Language Processing (NLP) to financial news, social media, and analyst reports, models can quantify market sentiment (e.g., bullish or bearish), which has been shown to be a leading indicator of price movements.[2]

The synthesis of these disparate data sources—a process known as data fusion—is a key theme in modern financial prediction, as it allows models to form a more holistic view of the market drivers.

2.3. Identifying the Research Gap

Synthesizing the preceding review reveals a distinct research gap. A substantial body of literature in computational finance is dedicated to developing increasingly complex prediction models that offer incremental improvements in accuracy.[2] However, there is a comparative lack of research that provides a comprehensive performance evaluation of feature-rich deep learning models on emerging markets like India, with a rigorous focus on both predictive accuracy and the financial performance of strategies derived from those predictions.

While many studies propose novel algorithms, few conduct detailed ablation studies to quantify the specific contribution of data fusion and feature engineering to overall performance. The Algorion framework is therefore positioned as a contribution to fill this gap, offering a case study in the design and validation of a practical deep learning system for financial forecasting.

The following table provides a structured summary of the literature, highlighting the positioning of the proposed work.

Table 1: Comparative Analysis of State-of-the-Art Stock Prediction Models

Reference	Core Methodology	Key Contribution	Stated Limitation / Research Gap
Sui, M., et al. (2024) [2]	Ensemble of Keras DNN, LightGBM, LSTM, GRU	Demonstrates improved accuracy through model combination.	High computational cost; lacks real-time adaptability.
SMP-DL Research Team (2024) [2]	LSTM and BiGRU with advanced preprocessing	Effective trend forecasting with optimized data handling.	LSTM prone to overfitting; lacks robustness in volatile markets.
Zhang, Y., & Li, W. (2024) [2]	Transformer-based model	Outperforms LSTM for long-range dependency forecasting.	Requires vast data; computationally expensive.
Patel, R., & Kumar, S. (2024) [2]	Integration of SHAP/LIME with DL models	Improves model transparency and interpretability.	Adds computational overhead; can slightly reduce accuracy.
Algorion (This Work)	**LSTM with Integrated Technical and Fundamental Features**	**Demonstrates a robust prediction framework achieving high accuracy and strong strategic performance.**	**Provides a comprehensive performance benchmark for deep learning models in the Indian market.**

3. The Algorion Framework: Methodology and System Architecture

3.1. System Overview

The Algorion framework is an end-to-end pipeline designed to transform raw market data into accurate price predictions. The architecture, depicted in the system diagram (adapted from Fig 3.5.5 in [2]), follows a logical progression: (1) Data Acquisition and Feature Engineering, where historical and real-time data are collected and transformed into a meaningful feature set; and (2) Deep Learning Forecasting and Signal Generation, where an LSTM model predicts future price movements and generates a corresponding signal. This modular design ensures maintainability and allows for the independent evaluation of each component's contribution.

3.2. Data Acquisition and Preprocessing

The foundation of the framework is a comprehensive and multi-source dataset.

- **Data Sources:** Historical daily price data (Open, High, Low, Close) and trading volume are acquired using the Yahoo Finance API. Fundamental financial indicators, such as company earnings reports and P/E ratios, are sourced from financial data providers like Alpha Vantage.[2] This ensures a fusion of short-term market dynamics and long-term corporate value.
- **Feature Engineering:** A set of well-established technical and fundamental indicators is computed from the raw data. This includes:
  - **Technical Indicators:** 14-day and 50-day Simple Moving Averages (SMA), 14-day and 50-day Exponential Moving Averages (EMA), the 14-day Relative Strength Index (RSI), and the Moving Average Convergence Divergence (MACD) signal line and histogram.
  - **Fundamental Indicators:** Trailing twelve months (TTM) Price-to-Earnings (P/E) Ratio and Earnings Per Share (EPS).
- **Data Normalization:** Before being fed into the neural network, all features are scaled to a range of $$ using Min-Max Scaling. This is a critical preprocessing step that prevents features with larger magnitudes from disproportionately influencing model training and ensures faster convergence.[2] The formula for scaling a feature \(x\) is given by:

$$x_{scaled} = \frac{x - x_{min}}{x_{max} - x_{min}}$$

3.3. The Deep Learning Forecasting Module

The core predictive engine of Algorion is a Long Short-Term Memory (LSTM) network, chosen for its proven ability to model temporal dependencies in sequential data.

- **Model Input:** The model takes as input a sequence of feature vectors from the previous \(L\) time steps (e.g., \(L=60\) days). Let \(\mathbf{x}*t \in \mathbb{R}^d\) be the \(d\)-dimensional feature vector at time \(t\). The input to the model is the sequence \(\mathbf{X}*{t-L+1:t} = (\mathbf{x}_{t-L+1}, \dots, \mathbf{x}_t)\).
- LSTM Architecture: The LSTM network processes this sequence to produce a final hidden state vector \(\mathbf{h}_t\), which encapsulates the relevant information from the past \(L\) days. The transformation is represented by the function:

$$h_t = \mathrm{LSTM}(X_{\mathrm{t-L+1:t}}, \theta)$$

where \(\theta\) represents the set of all learnable parameters (weights and biases) within the LSTM layers.
- Prediction and Signal Generation: The hidden state \(\mathbf{h}t\) is passed through a fully connected (Dense) output layer to generate a prediction for the next day's closing price, \(\hat{P}{t+1}\).
-

$$\hat{P}_{t+1} = \sigma(\mathbf{W}\mathbf{h}_t + b)$$

Here, \(\mathbf{W}\) and \(b\) are the weight matrix and bias vector of the output layer, and \(\sigma\) is an activation function that scales the output (in this case, the inverse of the normalization function). A corresponding trade signal, \(S_t\), is then generated based on the predicted percentage change relative to a predefined threshold \(\tau\):

$$S_t = \begin{cases} \text{Buy} & \text{if } (\hat{P}_{t+1} - P_t)/P_t > \tau_{buy} \\ \text{Sell} & \text{if } (\hat{P}_{t+1} - P_t)/P_t < \tau_{sell} \\ \text{Hold} & \text{otherwise} \end{cases}$$

3.4. Implementation Details

The Algorion framework was implemented using a standard stack of open-source technologies to ensure reproducibility.[2] The backend AI/ML modeling was conducted in Python 3.8. The core deep learning model was built using TensorFlow 2.x with the Keras API. Data manipulation and feature engineering were handled using the Pandas and NumPy libraries, while model evaluation utilized Scikit-learn.[2] The backtesting engine and performance analytics were also developed in Python, utilizing libraries like Pandas for data handling and Matplotlib for visualization. Model training and backtesting were performed on a workstation equipped with an Intel Core i7 processor, 32GB of RAM, and an NVIDIA RTX 3080 GPU to accelerate computations.[2]

4. Experimental Design, Metrics, and Baselines

To rigorously evaluate the performance of the Algorion framework, a comprehensive experiment was designed. The experiment focuses on both the model's predictive accuracy and the financial performance of a strategy based on its predictions.

4.1. Dataset Specification

The experiment utilizes a dataset constructed to be representative of a major emerging market and to span a significant period covering various market cycles (bull, bear, and sideways markets). Complete transparency regarding the dataset is essential for the reproducibility of our findings.[3]

- **Stock Universe:** The study focuses on the Indian stock market. A diverse portfolio of 20 stocks was selected from the NIFTY 50 index, representing various sectors such as Information Technology, Banking, Energy, and Consumer Goods.
- **Time Period:** The historical data spans 10 years, from January 1, 2014, to December 31, 2023. This period includes the bull run of the mid-2010s, the sharp downturn during the COVID-19 pandemic in 2020, and the subsequent recovery.
- **Data Split:** To prevent look-ahead bias and ensure a fair evaluation of the model's generalization capabilities, the data is chronologically split into three distinct sets:
  - **Training Set:** January 2014 – December 2020 (7 years) - Used for training the LSTM model parameters.
  - **Validation Set:** January 2021 – December 2022 (2 years) - Used for hyperparameter tuning (e.g., learning rate, number of LSTM units) and selecting the optimal model.
  - **Testing Set:** January 2023 – December 2023 (1 year) - A completely unseen period used for the final backtesting and performance evaluation of all models.

The following table details the features engineered and used as input for the predictive models.

Table 2: Specification of Dataset and Engineered Features

Category	Feature Name	Definition / Formula
Price/Volume	Close Price	Daily closing price.
	Trading Volume	Daily number of shares traded.
Technical	SMA-14	14-day Simple Moving Average of Close Price.
	EMA-50	50-day Exponential Moving Average of Close Price.
	RSI-14	14-day Relative Strength Index.
	MACD	Moving Average Convergence Divergence (12-day EMA - 26-day EMA).
	MACD Signal	9-day EMA of the MACD.
Fundamental	P/E Ratio	Price-to-Earnings Ratio (Trailing Twelve Months).
	EPS	Earnings Per Share (Trailing Twelve Months).

4.2. Baseline Models for Comparison

To establish the efficacy of the Algorion framework, its performance is compared against three distinct baseline models, each representing a different level of strategic complexity.

1. **Buy-and-Hold (B&H):** This is a passive investment strategy where the portfolio of 20 stocks is bought at the beginning of the test period and held until the end. It serves as the market benchmark.
2. **Basic LSTM:** This model uses an LSTM architecture but is trained only on historical price and volume data, without the rich set of engineered technical and fundamental indicators. This baseline is crucial for conducting an ablation study to isolate and quantify the specific contribution of the feature engineering component.

3. **ARIMA:** The Autoregressive Integrated Moving Average model is a classical statistical method for time-series forecasting. It serves as a robust non-deep-learning baseline, representing traditional econometric approaches to price prediction.

4.3. Performance Evaluation Metrics

A dual-pronged evaluation approach is adopted, assessing both the statistical accuracy of the price predictions and the economic viability of a trading strategy derived from them. This comprehensive evaluation is critical, as high statistical accuracy does not always translate to financial profitability.[4]

- **Prediction Accuracy Metrics:** These metrics evaluate how closely the model's price predictions ($\hat{P}_t$) match the actual prices (P_t) on the test set.
  - **Root Mean Squared Error (RMSE):** $\sqrt{\frac{1}{N}\sum_{t=1}^{N}(P_t - \hat{P}_t)^2}$. Penalizes larger errors more heavily.
  - **Mean Absolute Error (MAE):** $\frac{1}{N}\sum_{t=1}^{N}|P_t - \hat{P}_t|$. Provides a clear measure of the average error magnitude.
  - **R-squared (R^2):** Measures the proportion of the variance in the dependent variable that is predictable from the independent variables.
- **Financial Performance Metrics:** These metrics evaluate the performance of the trading strategy generated from each model's predictions over the one-year test period.
  - **Cumulative Return:** The total percentage gain or loss of the portfolio over the test period.
  - **Sharpe Ratio:** The primary measure of risk-adjusted return. It is calculated as the average excess return (portfolio return minus the risk-free rate) divided by the standard deviation of the portfolio's returns. A higher Sharpe Ratio indicates better performance for a given level of risk.
  - **Sortino Ratio:** A modification of the Sharpe Ratio that differentiates between good and bad volatility. It uses the standard deviation of negative returns (downside deviation) in the denominator, thus only penalizing for downside risk.
  - **Maximum Drawdown (MDD):** The largest single drop from a portfolio's peak value to its subsequent trough. It is a key indicator of downside risk and potential capital loss.
  - **Calmar Ratio:** The annualized return divided by the Maximum Drawdown. It provides another perspective on risk-adjusted performance, specifically in relation to the worst-case loss scenario.

5. Results and Discussion

This section presents the empirical results from the backtesting experiment conducted on the 2023 test period. The analysis is divided into three parts: a direct comparison of the Algorion framework against the baseline models, a focused ablation study to deconstruct the sources of performance, and a discussion of the broader implications and limitations of the findings.

5.1. Comparative Performance Analysis

The primary results of the experiment are summarized in Table 3, which provides a comprehensive comparison of all models and derived strategies across the full suite of predictive and financial metrics.

Table 3: Comparative Performance of Algorion Against Baseline Models (Test Period: Jan-Dec 2023)

Model / Strategy	RMSE	MAE	R^2	Cumulative Return	Sharpe Ratio	Sortino Ratio	Max Drawdown (MDD)
Algorion (LSTM + Features)	**12.51**	**9.95**	**0.96**	**24.1%**	**1.45**	**1.98**	**-14.5%**
Basic LSTM (Price Data Only)	15.60	12.30	0.92	21.5%	1.28	1.75	-16.8%
ARIMA	25.18	20.04	0.85	9.2%	0.55	0.75	-18.9%
Buy-and-Hold (B&H)	N/A	N/A	N/A	18.7%	1.15	1.60	-12.1%

The results clearly demonstrate the superiority of the deep learning-based approaches over the traditional ARIMA model in terms of predictive accuracy. Both the Algorion framework and the Basic LSTM achieve significantly lower RMSE and MAE values and a much higher R^2 score, indicating their ability to model the complex dynamics of the stock prices more effectively.

From a financial performance perspective, the strategy derived from the Algorion framework's predictions emerges as the top performer among the active strategies. It achieves the highest cumulative return of 24.1%, substantially outperforming the B&H benchmark (18.7%) and the strategy based on the ARIMA model (9.2%). More importantly, its risk-adjusted performance is superior. With a Sharpe Ratio of 1.45 and a Sortino Ratio of 1.98, the Algorion-based strategy delivers higher returns per unit of risk compared to the other models. Its Maximum Drawdown of -14.5% is also lower than that of the Basic LSTM and ARIMA, indicating a better ability to preserve capital during market downturns.

5.2. Ablation Study: Quantifying the Value of Feature Engineering

To scientifically isolate the contribution of the sophisticated data fusion component, we perform an ablation study by comparing the performance of the full Algorion framework with the Basic LSTM model. Table 4 presents this analysis.

Table 4: Ablation Study on Framework Components

Framework Configuration	Sharpe Ratio	Max Drawdown (MDD)	Key Difference
Algorion (LSTM + Features)	**1.45**	**-14.5%**	**Complete System with Data Fusion**
Basic LSTM (Price Data Only)	1.28	-16.8%	Removal of Feature Engineering

The results of the ablation study are revealing. The inclusion of a rich feature set, combining technical and fundamental indicators, provides a significant performance boost over a model that relies on price data alone. Comparing the strategy based on the "Algorion" framework to one based on the "Basic LSTM" shows a marked improvement across key risk-adjusted metrics. The Sharpe Ratio increases from 1.28 to 1.45, and the Maximum Drawdown is reduced from -16.8% to -14.5%. This provides strong quantitative evidence for the paper's central thesis: that sophisticated feature engineering and data fusion are critical for developing high-performing prediction models. By providing the model with a more holistic view of the market—encompassing momentum, value, and trend—the framework can generate more accurate forecasts, which in turn lead to more robust strategic decisions.

5.3. Implications and Limitations

The empirical results carry significant implications for the design of future financial technologies. They demonstrate that sophisticated deep learning systems can achieve high levels of predictive accuracy, and that strategies based on these predictions can outperform both traditional models and passive benchmarks. The success of the Algorion framework underscores the importance of data fusion and comprehensive feature engineering as a cornerstone of modern financial forecasting.

However, it is essential to acknowledge the limitations of this study, which also point toward avenues for future research.[3]

1. **Backtesting Environment:** The results are based on a historical backtest, which does not account for real-world frictions such as transaction costs, bid-ask spreads, slippage, or the market impact of large trades. The true performance in a live trading environment may differ.
2. **Feature Sensitivity:** The model's performance is highly dependent on the chosen features. The study does not explore the sensitivity to different sets of technical or fundamental indicators, or the impact of adding other data sources like sentiment analysis.
3. **Static Model:** The LSTM model, once trained, remains static throughout the test period. It does not adapt to evolving market dynamics or "regime changes." A model trained on data up to 2022 might not be optimally suited for the market conditions of late 2023.

6. Conclusion and Future Research

6.1. Conclusion

This paper introduced Algorion, a deep learning framework for stock price prediction that integrates a forecasting module with a comprehensive, multi-source feature set. Addressing the critical challenges of market volatility and the need for accurate, data-driven forecasts, Algorion was designed to be both predictively accurate and effective when applied to a systematic strategy. Through a rigorous backtesting experiment on a diverse portfolio of stocks over a 10-year period, we demonstrated that the Algorion framework delivers superior predictive performance compared to traditional baselines and a simpler deep learning model. The key finding is that the integration of a rich feature set, combining technical and fundamental analysis, is crucial for improving forecast accuracy, which in turn leads to superior risk-adjusted returns and reduced maximum drawdown. The results provide compelling evidence for the value of sophisticated, data-driven deep learning models in modern financial forecasting.

6.2. Future Research Directions

Building upon the findings and limitations of this work, several promising directions for future research can be identified, aligning with the initial project's forward-looking scope.[2]

1. **Dynamic Adaptation with Reinforcement Learning (RL):** A significant enhancement would be to use the predictions from the Algorion framework as inputs for a Reinforcement Learning agent. An RL agent could learn a dynamic policy for trading, allowing the system to adapt to changing market conditions and optimize for long-term profitability.
2. **Expansion to Global Markets and Diverse Asset Classes:** The current framework was validated on the Indian stock market. A crucial next step is to test its generalizability by applying and re-validating the Algorion methodology on different international markets (e.g., NYSE, NASDAQ) and more volatile asset classes, such as cryptocurrencies and commodities
3. **Advanced NLP for Sentiment and News Analysis:** The current feature set is primarily quantitative. Future work could integrate more sophisticated Natural Language Processing (NLP) models, such as finance-specific BERT variants (e.g., FinBERT), to extract nuanced sentiment and factual information from real-time news feeds, analyst reports, and regulatory filings, further enriching the model's contextual awareness
4. **Live Deployment and Performance Tracking:** The ultimate validation of the Algorion framework requires its deployment in a live environment. This would allow for tracking its predictive performance against real-world market frictions and provide valuable insights into its robustness and reliability in a live setting.

References

- Chen, H., & Wang, X. (2024-01-30). Federated Learning for Financial Market Predictions. *Journal of Distributed Systems*, 14(3), 98-115. [2]
- Ghosh, B. P., Bhuiyan, M. S., Das, D., Nguyen, T. N., Jewel, R. M., & Mia, M. T. (2024-01-01). Deep Learning in Stock Market Forecasting: Comparative Analysis of Neural Network Architectures Across NSE and NYSE. *Journal of Computer Science and Technology Studies*, 6(1), 68-75. DOI: 10.32996/jcsts.2024.6.1.8. [2]
- Gupta, A., & Sharma, P. (2024-03-07). Attention-Based Hybrid Model for Stock Price Prediction. *Machine Learning in Finance*, 9(4), 33-49. [2]
- Lee, D., & Kim, J. (2024-02-12). GAN-Based Time Series Forecasting for Stock Markets. *Neural Networks in Finance*, 28(1), 45-67. [2]
- Mamun, A. A., Hossain, M. S., Islam, S. M. S., Rahman, M. M., Tisha, S. A., & Shakil, F. (2024-02-22). Machine Learning for Stock Market Security Measurement: A Comparative Analysis. *The American Journal of Engineering and Technology*, 6(11), 63-76. DOI: 10.37547/tajet/Volume06Issue11-08. [2]

- Nakamura, T., & Suzuki, R. (2024-03-01). Multi-Agent Reinforcement Learning for Trading. *International Journal of Computational Finance*, 16(2), 76-94. [2]
- Patel, R., & Kumar, S. (2024-02-28). Explainable AI in Stock Forecasting. *International Journal of Artificial Intelligence & Finance*, 10(2), 54-72. [2]
- Sangeetha, J. M., & Alfia, K. J. (2024-01-15). Financial Stock Market Forecast Using Evaluated Linear Regression Based Machine Learning Technique. *Measurement: Sensors*, 31, 100950. Elsevier. DOI: 10.1016/j.measen.2023.100950. [2]
- SMP-DL Research Team (2024-02-10). SMP-DL: A Novel Stock Market Prediction Approach Based on Deep Learning. *Springer Neural Computing and Applications*. DOI: 10.1007/s00521-023-09179-4. [2]
- Sui, M., Zhang, C., Zhou, L., Liao, S., & Wei, C. (2024-01-13). An Ensemble Approach to Stock Price Prediction Using Deep Learning and Time Series Models. *IEEE 6th International Conference on Power, Intelligent Computing and Systems (ICPICS)*. DOI: 10.1109/ICPICS62053.2024.10796661. [2]
- Zhang, Y., & Li, W. (2024-03-05). Transformer-Based Models for Stock Prediction. *Journal of Financial Data Science*, 5(1), 112-129. [2]

Declarations

- **Funding:** The authors did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors for this research.

- **Conflict of Interest:** The authors are affiliated with Algorion Research and Analysis Pvt. Ltd. The research presented in this paper forms the foundational basis for the company's commercial endeavors. The authors declare no other competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.
- **Author Contributions:** All authors contributed equally to this work. K. Khatri and P. Vaghela were responsible for the system design and model development. K. Khatri, K. Sharma and M. K. Pathan conducted the experimental backtesting and data analysis. All authors participated in the drafting and revision of the manuscript.