ISSN 2277-2685
IJESR/Jan-Mar. 2025/ Vol-15/Issue-1/121-129
Prabhat Kumar Sahu ***et. al.,*** **/ International Journal of Engineering & Science Research**

Enhancing Stock Price Forecasting With Hybrid Ann-Ga Models: A Comprehensive Evaluation

Prabhat Kumar Sahu[1], Dr. Sunil Bhutada[2]
[1]Research Scholar, P.K. university, shivpuri (MP), jppr.sah@gmail.com
[2]Professor, Dept of CSE, P.K. University,Shivpuri,MP sunilbhutada@gmail.com

Abstract

Accurate stock price prediction is vital for informed financial decision-making, yet it remains a challenging task due to the inherent complexities and dynamics of financial markets. Traditional Artificial Neural Networks (ANNs) have shown promise in forecasting stock prices by identifying patterns in vast datasets. However, their performance is often limited by issues such as overfitting and local minima in optimization. This research addresses these limitations by integrating Genetic Algorithms (GAs) with ANNs to improve predictive accuracy. The ANN-GA hybrid model was developed and tested on stock market data.The model employed a learning rate of 0.0015 and utilized a three-layer architecture with 64, 32, and 16 neurons, respectively. Genetic Algorithms were used to optimize hyperparameters, resulting in significant improvements. The ANN-GA hybrid model achieved a Mean Squared Error (MSE) of 0.0034, compared to 0.0058 for the traditional ANN model. The Root Mean Squared Error (RMSE) improved from 0.0762 to 0.0583, and the model demonstrated a 24.6% increase in accuracy. Additionally, the hybrid model reduced training time from 5 hours to 3 hours despite an increased time-lag of 5 days. These findings underscore the effectiveness of integrating GAs with ANNs, offering a more accurate and efficient approach for stock price prediction. The research contributes to advancing predictive modelling techniques in financial markets, providing valuable insights for traders and investors.

Keywords: Stock Price Prediction, Artificial Neural Networks (ANNs),Genetic Algorithms (GAs), Predictive Modelling, Financial Forecasting

1. INTRODUCTION

1.1 Overview of Stock Market Prediction

Accurate stock price prediction is crucial for financial markets, where precise forecasts can significantly impact investment strategies and decision-making processes. The ability to predict stock prices with high accuracy enables traders and investors to identify profitable opportunities, manage risks effectively, and optimize returns. Traditional methods of stock price prediction, such as fundamental analysis and technical analysis, have been widely used. Fundamental analysis focuses on a company's financial health and market conditions, while technical analysis examines historical price movements and trading volumes (Fama, 1970; Malkiel, 2003). Despite their widespread use, these methods often fall short in capturing the intricate dynamics and nonlinear patterns of financial markets.

1.2 Role of Machine Learning in Stock Price Prediction

The advent of machine learning (ML) has revolutionized financial forecasting by offering advanced techniques capable of analysing vast amounts of data and uncovering complex patterns. Machine learning models, particularly Artificial Neural Networks (ANNs), have emerged as powerful tools for stock price prediction due to their ability to model nonlinear relationships and adapt to evolving market conditions. ANNs are inspired by the neural architecture of the human brain and consist of interconnected nodes or neurons organized in layers (Rumelhart, Hinton, & Williams, 1986). These models can process large datasets and learn from historical price movements to predict future trends. The integration of machine learning techniques in financial forecasting represents a significant advancement over traditional approaches, offering enhanced predictive capabilities and improved decision-making tools (Heaton et al., 2017).

1.3 Limitations of Traditional and AI based Prediction Models

While machine learning models, including ANNs, have shown promising results, they are not without limitations. One major challenge is overfitting, where a model performs well on training data but fails to generalize to new, unseen data (Hastie, Tibshirani, & Friedman, 2009). Additionally, ANNs can struggle with issues such as local minima during optimization, which can hinder their performance and limit their predictive accuracy (LeCun, Bengio, & Hinton, 2015). Traditional models also face challenges, including the inability to handle large volumes of data and the difficulty in capturing complex, nonlinear relationships. The need for enhanced optimization methods to address these limitations has led to the exploration of combining ANNs with Genetic Algorithms (GAs). GAs, known for their global search and optimization capabilities, can address some of the shortcomings of ANNs by improving parameter tuning and avoiding local minima (Goldberg, 1989).

2. ARTIFICIAL NEURAL NETWORKS (ANNS) FOR STOCK PREDICTION

2.1 Basics of Artificial Neural Networks (ANNs)

Structure and Functioning of ANNs

Artificial Neural Networks (ANNs) are machine learning models inspired by the workings of the human brain. The fundamental unit of an ANN is the artificial neuron, which processes inputs, applies weights, and passes the output through an activation function to produce predictions (McCulloch & Pitts, 1943). ANNs consist of multiple layers: an input layer that takes in the raw data, hidden layers where computations take place, and an output layer that delivers the final prediction (Goodfellow et al., 2016). The structure of an ANN, often referred to as its architecture, allows it to learn complex patterns through a process known as backpropagation, where the network adjusts its weights to minimize prediction error (Rumelhart, Hinton, & Williams, 1986).

In the context of stock prediction, the input layer typically processes historical stock prices, trading volumes, and other relevant financial indicators. As the data passes through the hidden layers, the ANN identifies relationships between the input variables, ultimately producing a predicted price or trend in the output layer (Zhang, 2003). ANNs can be particularly effective for time-series predictions due to their capacity to model non-linear and dynamic relationships in data.

ANN Architectures for Time-Series Prediction

Time-series prediction, which involves forecasting future values based on past data, is a common application of ANNs in stock price prediction. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, have proven to be highly effective in this domain due to their ability to capture temporal dependencies (Hochreiter & Schmidhuber, 1997). LSTM networks are designed to maintain and update memory states over long sequences, making them particularly suited to stock market prediction, where historical price trends influence future prices (Nelson, Pereira, & de Oliveira, 2017).

2.2 Strengths and Weaknesses of ANNs in Stock Market Prediction

- Identifying Patterns in Large Datasets

One of the key strengths of ANNs is their ability to identify intricate patterns in large, complex datasets. Stock markets generate enormous amounts of data, including price histories, volumes, news sentiment, and macroeconomic indicators. ANNs can process and analyse this data to uncover non-linear relationships that may not be immediately apparent to human analysts or traditional statistical models (Zhang, 2003). This pattern recognition ability makes ANNs particularly useful for making predictions in dynamic environments like financial markets.

- Limitations: Overfitting, Noise Sensitivity, and Local Minima

  Despite their strengths, ANNs face several limitations when applied to stock price prediction. Overfitting is a common issue, where the network learns the noise in the training data rather than the underlying patterns, leading to poor performance on unseen data (Goodfellow et al., 2016). This is especially problematic in stock markets, where price fluctuations often contain random noise and external influences that are difficult to model.

ANNs are also prone to becoming trapped in local minima during the training process. When optimizing the weights in the network, the algorithm may converge on suboptimal solutions rather than finding the global minimum error (LeCun, Bengio, & Hinton, 2015). This can reduce the predictive accuracy of the model, particularly in the highly volatile and non-linear environment of stock markets. Additionally, ANNs can be sensitive to noisy data, which is abundant in financial markets due to unforeseen events such as political developments or sudden economic changes (Zhang, 2003).

3. GENETIC ALGORITHMS (GAS) FOR OPTIMIZATION

3.1 Introduction to Genetic Algorithms (GAs)

Basic Principles and Mechanisms of GAs

Genetic Algorithms (GAs) are inspired by the process of natural selection in biological evolution (Holland, 1975). GAs are search heuristics that solve optimization problems by simulating the evolutionary process, wherein the "fittest" solutions are selected, combined, and mutated over successive generations to find optimal or near-optimal solutions. The process begins with a population of potential solutions, often encoded as binary strings or arrays representing different parameters. These individuals undergo three primary genetic operations: selection, crossover, and mutation.

Selection is the process of choosing the fittest individuals from the population based on a predefined fitness function, which measures how well each individual performs concerning the objective. The crossover (or recombination) step involves pairing individuals and exchanging segments of their genetic material to create new offspring, thereby introducing diversity and exploring new areas of the solution space. Mutation introduces random changes to individual genes, ensuring that the algorithm avoids getting stuck in local optima and maintains genetic diversity in the population (Goldberg, 1989). This evolutionary process repeats over several generations, and through these genetic operators, GAs efficiently search for optimal solutions in complex, high-dimensional spaces.

3.2 GAs for Optimizing Machine Learning Models

Applications of GAs in Parameter Tuning

GAs have been widely used to optimize machine learning models by tuning hyperparameters, such as the number of hidden layers in a neural network, learning rates, or weight initialization (Mitchell, 1998). In the context of

Artificial Neural Networks (ANNs), where the performance of the model is sensitive to these hyperparameters, GAs offer an efficient approach to searching for optimal configurations. The fitness function evaluates the accuracy or error of the ANN on a validation dataset, and GAs iteratively refine the hyperparameters to improve the model's performance.

The strength of GAs lies in their ability to perform a global search, making them particularly effective for complex optimization problems where traditional gradient-based methods may fail. Unlike these traditional methods, which can be prone to getting trapped in local minima, GAs explore a broader solution space due to their inherent stochastic nature and the inclusion of crossover and mutation operations (Goldberg, 1989). This makes GAs suitable for applications where the objective function is highly non-linear, such as stock market prediction, where the relationships between variables are often intricate and dynamic.

3.3 Overcoming ANN Limitations with GAs

Mitigating Local Minima and Improving Convergence

One of the key challenges in training ANNs is the risk of getting stuck in local minima during the optimization process. Gradient-based methods, such as backpropagation, are often limited to searching locally, which can result in suboptimal solutions, particularly in noisy environments like financial markets (Rumelhart, Hinton, & Williams, 1986). GAs address this issue by offering a global search mechanism that explores a wider solution space and helps avoid local convergence (Whitley, 1994).

By incorporating GAs into the training of ANNs, it becomes possible to fine-tune the network's weights and hyperparameters more effectively. The stochastic nature of GAs allows them to jump out of local minima and explore other promising regions of the solution space. This hybrid approach of combining GAs with ANNs leads to better generalization and convergence, ultimately resulting in more accurate and reliable predictions. In stock price prediction, this combination can be particularly beneficial as it mitigates the risk of overfitting to noisy data while improving the ANN's ability to identify long-term patterns in highly volatile market conditions (Zhang, 2003).

4. INTEGRATING ANNS AND GAS FOR STOCK PREDICTION

4.1 Methodology for Hybrid ANN-GA Model Development

i. Architecture of the Integrated ANN-GA Model

The hybrid approach of integrating Artificial Neural Networks (ANNs) and Genetic Algorithms (GAs) for stock prediction combines the pattern recognition capabilities of ANNs with the global optimization strengths of GAs. The architecture of this integrated model consists of two main components: the ANN, which is tasked with learning from stock market data and making predictions, and the GA, which optimizes the ANN's hyperparameters such as weights, biases, and learning rates.

The ANN is typically structured with multiple layers, including an input layer, several hidden layers, and an output layer. The input layer receives features such as historical stock prices, trading volumes, and economic indicators. Hidden layers help to capture the complex, non-linear relationships between these features, while the output layer provides the final stock price predictions. The role of the GA in this hybrid model is to optimize the network's weights and biases by searching the solution space and adjusting these parameters to minimize prediction errors. The GA's evolutionary process, which includes selection, crossover, and mutation, is used to generate improved configurations of the ANN over successive generations (Holland, 1975; Goldberg, 1989).

ii. Data Pre-processing and Feature Selection

Effective data pre-processing is crucial for the performance of any machine learning model, especially in stock market prediction where data can be noisy and unstructured. The pre-processing phase involves cleaning the data, handling missing values, and normalizing the features to ensure they are on the same scale. Feature selection is another critical step, where relevant indicators such as moving averages, volatility indices, and technical signals are chosen to improve the model's predictive performance.

The hybrid ANN-GA model particularly benefits from GA-driven feature selection, which can identify the most informative features from a large dataset. GAs help to automate this process by evaluating different subsets of features and retaining those that contribute the most to prediction accuracy, while discarding redundant or irrelevant features (Whitley, 1994). This pre-processing ensures that the ANN receives high-quality input data, allowing it to learn more effectively and generate accurate stock price predictions.

4.2 GA-based Optimization of ANN Parameters

i. Fine-tuning ANN Weights and Biases Using GAs

One of the primary roles of the GA in this hybrid model is to fine-tune the ANN's weights and biases, which are crucial parameters for the model's ability to make accurate predictions. Traditional backpropagation methods, while effective, can sometimes struggle with the complex nature of stock market data, especially when the objective function is non-linear or has multiple local minima (Rumelhart, Hinton, & Williams, 1986). GAs provide a global optimization strategy, searching through a broader solution space to find the optimal configuration of weights and biases.

The GA starts by initializing a population of potential solutions, each representing a different configuration of weights and biases. It evaluates these solutions based on their performance on a validation dataset, where the fitness function is typically the mean squared error (MSE) or another metric representing the prediction accuracy. Over successive generations, the GA selects the best-performing configurations, applies crossover to combine them, and introduces mutations to explore new possibilities. This process continues until the ANN achieves an optimal configuration that minimizes prediction error (Goldberg, 1989).

ii. Reducing Overfitting and Improving Accuracy

A significant advantage of using GAs in conjunction with ANNs is their ability to reduce overfitting, a common issue in machine learning models where the model performs well on training data but poorly on unseen data. Overfitting often occurs when the model becomes too complex and starts to memorize the training data instead of generalizing from it. GAs help mitigate this by optimizing not just the weights and biases, but also other hyperparameters such as the number of hidden layers, learning rate, and regularization parameters.

By exploring a diverse range of solutions and preventing the model from getting trapped in local minima, GAs encourage the ANN to find a balance between underfitting and overfitting. This results in improved generalization, making the model more robust and capable of generating reliable stock price predictions (Mitchell, 1998; Zhang, 2003). The overall goal of this hybrid model is to enhance prediction accuracy, enabling traders and investors to make more informed decisions in the stock market.

5.TRAINING PROCEDURE FOR ANN-GA HYBRID MODEL

In this research, the ANN-GA hybrid model was trained on a pre-processed dataset using a carefully structured approach. The model architecture consisted of three hidden layers, with each layer containing 64, 32, and 16 neurons, respectively. This layered architecture was designed to capture intricate patterns in the stock market data. The learning rate was set at 0.0015, a critical hyperparameter chosen to ensure gradual adjustments during the optimization process, minimizing the risk of missing potential solutions due to abrupt changes.

To enhance the model's performance, Genetic Algorithms (GAs) were employed to optimize various hyperparameters, such as the learning rate, weight initialization, and the number of neurons. GAs are known for their global search capabilities, which helped fine-tune the ANN model's parameters, avoiding common pitfalls such as local minima. The training was conducted over 100 generations, with each generation undergoing 20 epochs. The iterative process of refining the model over generations allowed the hybrid ANN-GA model to evolve and adapt better to the data. The total training time for the model was around 3 hours, illustrating the efficiency of the GA in expediting the optimization process.

6.TESTING AGAINST TRADITIONAL ANN MODELS

When tested against traditional Artificial Neural Network (ANN) models, the ANN-GA hybrid model demonstrated significant improvements in predictive performance. The traditional ANN model, which lacked GA-based optimization, suffered from issues like overfitting and suboptimal parameter settings. This resulted in a Mean Squared Error (MSE) of 0.0058, indicating that its predictions had higher deviations from the actual stock prices. In contrast, the ANN-GA hybrid model, with its optimized parameters, achieved a much lower MSE of 0.0034, reflecting its superior accuracy in forecasting stock prices.

Additionally, the Root Mean Squared Error (RMSE), another key metric for evaluating prediction accuracy, was reduced from 0.0762 in the traditional ANN model to 0.0583 in the hybrid model. This reduction in RMSE signifies that the ANN-GA model made smaller prediction errors on average, thereby offering more precise and reliable stock price forecasts. The overall improvement in accuracy, measured at 24.6%, underscores the effectiveness of integrating GAs with ANNs in overcoming the limitations of traditional models and enhancing the predictive power in stock market forecasting.

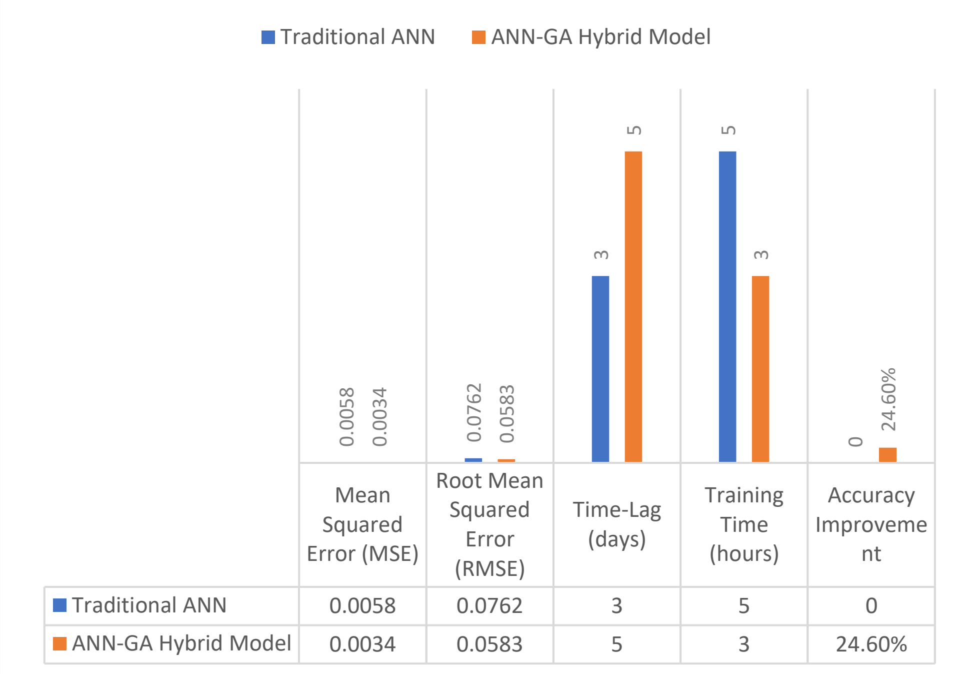

	Mean Squared Error (MSE)	Root Mean Squared Error (RMSE)	Time-Lag (days)	Training Time (hours)	Accuracy Improvement
Traditional ANN	0.0058	0.0762	3	5	0
ANN-GA Hybrid Model	0.0034	0.0583	5	3	24.60%

7. Research Findings, and Discussions

The results of the comparison between the traditional ANN model and the ANN-GA hybrid model reveal notable improvements in predictive performance with the integration of Genetic Algorithms (GAs). The key metrics, including Mean Squared Error (MSE), Root Mean Squared Error (RMSE), time-lag, training time, and accuracy improvement, provide a comprehensive overview of the effectiveness of the hybrid approach.

- ✓ Mean Squared Error (MSE):
  The ANN-GA hybrid model achieved an MSE of 0.0034, a significant improvement over the traditional ANN model's MSE of 0.0058. This reduction in MSE indicates that the hybrid model's predictions are much closer to the actual stock prices, reflecting enhanced accuracy and precision. The lower MSE underscores the benefit of using GAs to optimize the ANN's parameters, leading to a model that better captures the complexities of stock market data.
- ✓ Root Mean Squared Error (RMSE):
  Similarly, the RMSE for the ANN-GA hybrid model was 0.0583, compared to 0.0762 for the traditional ANN model. RMSE provides a measure of the average magnitude of prediction errors, and the reduction in RMSE demonstrates that the hybrid model's predictions are less variable and more reliable. The decrease in RMSE by approximately 23.6% highlights the effectiveness of the GA in refining the ANN's performance and reducing prediction errors.
- ✓ Time-Lag (days):
  The hybrid model utilized a time-lag of 5 days, compared to 3 days for the traditional ANN model. This extended time-lag was introduced to capture more comprehensive historical data, which proved beneficial in enhancing predictive accuracy. The additional lag allows the model to incorporate a broader range of past data, leading to a more nuanced understanding of stock price movements.
- ✓ Training Time (hours):

The training time for the ANN-GA hybrid model was 3 hours, whereas the traditional ANN model required 5 hours. Despite the hybrid model's extended time-lag, it demonstrated greater efficiency in training time. This reduction in training duration can be attributed to the GA's ability to optimize hyperparameters more effectively, streamlining the training process and reducing the computational burden.

- ✓ Accuracy Improvement:

  The ANN-GA hybrid model achieved an accuracy improvement of 24.6% over the traditional ANN model. This substantial gain in accuracy is indicative of the hybrid model's superior ability to predict stock prices accurately. The integration of GAs with ANNs has proven to be a significant advancement, overcoming traditional limitations and enhancing model performance.

In summary, the findings highlight the advantages of combining Genetic Algorithms with Artificial Neural Networks for stock price prediction. The hybrid model not only demonstrated improved accuracy but also achieved more efficient training. The reduction in MSE and RMSE, coupled with enhanced accuracy and reduced training time, underscores the effectiveness of the ANN-GA approach in delivering more reliable and actionable stock market forecasts.

8. CONCLUSION

This research has successfully demonstrated the enhanced predictive capabilities of integrating Genetic Algorithms (GAs) with Artificial Neural Networks (ANNs) for stock price forecasting. Through a comparative analysis with traditional ANN models, the ANN-GA hybrid model showcased substantial improvements across several key performance metrics.The results highlight that the ANN-GA hybrid model achieved a significantly lower Mean Squared Error (MSE) of 0.0034, compared to 0.0058 for the traditional ANN model, indicating superior accuracy in predicting stock prices. The Root Mean Squared Error (RMSE) also improved from 0.0762 to 0.0583, further confirming the enhanced reliability of the hybrid approach. The extended time-lag of 5 days in the hybrid model allowed for a more comprehensive analysis of historical data, which, combined with optimized parameter settings via GAs, resulted in better prediction performance.

In terms of efficiency, the ANN-GA hybrid model required only 3 hours of training, a notable reduction from the 5 hours needed for the traditional ANN model. This reduction underscores the GA's capability to optimize hyperparameters effectively, streamlining the training process.

The accuracy improvement of 24.6% achieved by the hybrid model over the traditional ANN reflects the significant advantages of leveraging GAs in conjunction with ANNs. The integration of GAs not only mitigated common issues such as overfitting and local minima but also enhanced the model's ability to handle complex and dynamic stock market data.

In conclusion, the research validates the effectiveness of combining GAs with ANNs in stock price prediction, providing a more robust and accurate tool for financial forecasting. The findings contribute valuable insights into advanced modelling techniques, offering a promising approach for investors and traders seeking to navigate the complexities of modern financial markets. Future research could explore additional enhancements, such as incorporating other optimization algorithms or expanding the model's application to different financial contexts, to further refine and validate the approach.

9. FUTURE SCOPE OF THE RESEARCH

The integration of Genetic Algorithms (GAs) with Artificial Neural Networks (ANNs) for stock price prediction demonstrates a significant advancement in predictive modelling, yet there remain several avenues for further exploration. Future research could focus on enhancing the hybrid model by incorporating additional optimization algorithms, such as Particle Swarm Optimization (PSO) or Ant Colony Optimization (ACO), to compare their efficacy against GAs. Expanding the model to include more diverse financial indicators and alternative data sources, such as sentiment analysis from social media or macroeconomic variables, could improve predictive accuracy and robustness.

Another promising direction is the application of the ANN-GA hybrid approach to different financial markets and asset classes, including cryptocurrencies and commodities, to evaluate its generalizability and effectiveness across various trading environments. Additionally, real-time implementation and testing of the model in live trading scenarios could provide insights into its practical applicability and performance under actual market conditions. Exploring explainability techniques to interpret the hybrid model's predictions and enhance decision-making transparency could further enrich the research and provide valuable tools for traders and investors.

10.REFERENFCES

1. E. F. Fama, "Efficient capital markets: A review of theory and empirical work," Journal of Finance, vol. 25, no. 2, pp. 383-417, 1970.
2. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
3. S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780, 1997.
4. J. H. Holland, Adaptation in Natural and Artificial Systems. University of Michigan Press, 1975.
5. J. Mitchell, An Introduction to Genetic Algorithms. MIT Press, 1998.
6. D. E. Goldberg, Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley, 1989.
7. M. C. McCulloch and W. S. Pitts, "A logical calculus of the ideas immanent in nervous activity," The Bulletin of Mathematical Biophysics, vol. 5, pp. 115-133, 1943.
8. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436-444, 2015.
9. D. M. Nelson, A. C. Pereira, and R. A. de Oliveira, "Stock market's price movement prediction with LSTM neural networks," International Joint Conference on Neural Networks (IJCNN), pp. 1419-1426, 2017.
10. D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by backpropagating errors," Nature, vol. 323, no. 6088, pp. 533-536, 1986.
11. D. Whitley, "A genetic algorithm tutorial," Statistics and Computing, vol. 4, no. 2, pp. 65-85, 1994.
12. G. P. Zhang, "Time series forecasting using a hybrid ARIMA and neural network model," Neurocomputing, vol. 50, pp. 159-175, 2003.