Journal of Applied Research and Technology

www.jart.icat.unam.mx

Journal of Applied Research and Technology 24 (2026) 93-110

Original

A comprehensive model to support investment decisions based on deep learning and evolutionary algorithms

Xochitl Segura Lozano[1] • Reimundo Moreno-Cepeda[1]
Juan Antonio Granados Montelongo[2] • José Daniel Corona Flores[2]
Juan Antonio Álvarez–Gaona[1] • Carlos Tolentino*[3]

[1]*Universidad Autónoma de Coahuila, Saltillo, Mexico*
[2]*Universidad Autónoma Agraria Antonio Narro, Saltillo, Mexico*
[3]*Tecnológico Nacional de México, Saltillo, Mexico*

Received: 20 11 2024; Accepted: 01 23 2025
Available: 28 02 2026

Abstract: This paper presents and assesses a comprehensive approach to stock price forecasting and portfolio selection that integrates advanced computational intelligence techniques with fundamental and technical analyses. Several outstanding forecasting methods are compared to identify the most accurate model. Subsequently, differential evolution is used to optimize a stock portfolio, leveraging the results of the selected forecasting method along with key technical and fundamental indicators. Experiments show that the proposed method consistently yields higher returns and better risk management than several benchmarks. Statistical validation confirms the model's superior performance, highlighting its potential as a robust tool for optimizing investment portfolios.

Keywords: **Stock selection, price forecasting, computational intelligence, technical analysis, fundamental analysis, differential evolution.**

*Corresponding author.
E-mail address: carlos.th@saltillo.tecnm.mx (C. Tolentino).
Peer Review under the responsibility of Universidad Nacional Autónoma de México.

1. Introduction

Accurate stock price forecasting and effective stock selection are important concerns for financial management, especially within institutions that offer investment services. Additionally, portfolio optimization supports financial decisions that maximize returns while minimizing risk. These activities are associated with overall economic growth and job creation. Investment portfolio optimization, especially involving stock portfolios, has developed as one of the most important areas of financial research and practice. Traditionally, this was done through a process of portfolio construction and judicious stock selection. Harry Markowitz, in a seminal contribution in 1952, laid the foundation for Modern Portfolio Theory (MPT), introducing the concept of portfolio optimization based on trade-offs between expected return and risk (Markowitz, 1952). Markowitz's framework, which focused on achieving the best risk-adjusted return through minimum variance, has long been the bedrock of portfolio optimization. Despite its apparent theoretical robustness, MPT faced several limitations in practical applications that motivated the search for alternative models and techniques. More recent works tend to focus on only one aspect of portfolio management, neglecting that comprehensiveness is crucial to generating high performance (Solares, De-León-Gómez, et al., 2022). Great technological advancements, combined with market volatility, have brought about changes in the financial sector over the last couple of decades. New challenges and opportunities arise in investment service institutions. This tradition in forecasting and selection methods is increasingly being scrutinized for effectiveness in such a dynamic environment.

The advent of deep learning and evolutionary computing in computational intelligence opens promising avenues to enhance the precision and efficacy of stock price forecasting and selection. Deep learning, a subset of machine learning, has demonstrated remarkable success in various domains, including image recognition, natural language processing, and, more recently, financial forecasting. Recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and gated recurrent units (GRUs) are among the advanced deep learning architectures employed for time-series forecasting (Fischer & Krauss, 2018; Wong & Luo, 2018). These models are very good at learning from complex patterns and interdependencies in sequential data, making them suitable for stock price forecasting. It provides more accurate and reliable forecasts based on historical price data and market indicators through the deep learning model. Deep learning techniques are complemented by other evolutionary computing techniques, such as differential evolution, which provides sound methods of optimization in stock selection. Differential evolution is defined as a population-based optimization algorithm that iteratively enhances candidate solutions based on a fitness criterion (Das & Suganthan, 2010; Solares, Salas, et al., 2022). It will be the first time that, in the context of portfolio optimization, differential evolution is able to facilitate real efficiency in making its way through the intricate landscape of possible stock combinations and identifying the most suitable stocks that return and minimize risk.

This, therefore, makes the integration of deep learning with evolutionary computation in financial management a significant improvement over the more conventional methods. It represents a superior hybrid approach to increasing the accuracy of stock price predictions and optimizing stock selection to generate better-performing portfolios. This will be achieved through their combined strengths in leveraging advanced computational techniques to improve the decision-making process of investment services institutions for their clients and add value to overall economic stability. To this end, the study's general problem statement is to establish how a financial management model with integrated advanced forecasting techniques can strengthen decision-making in investment selection within the investment services industry.

The specific questions to be addressed are: Which of the forecasting methods can promise more accurate and reliable results? How could such techniques be integrated into existing portfolio management systems for better performance? What is the real benefit of such improved models on the operational practice and performance of investment services institutions?

This paper describes an integrative financial management model that employs deep learning and stock price prediction as well as evolutionary computation for stock selection. The proposed model should, by all accounts, address the complexity of the new financial landscape and provide a robust, reproducible methodology to ensure the practice of investment service institutions. The main objectives of this research will be twofold: finding and comparing various superior forecasting methods in order to eventually find an accurate model that does the job in stock price forecasting, and applying the differential evolution method to make an optimized stock portfolio based on a selected forecasting model, along with technical and fundamental indicators.

The rest of this paper is organized as follows: Section 2 presents a review of related literature on stock price prediction and portfolio optimization, including their development and shortcomings. Section 3 discusses methodology, data sources, data preprocessing, and implementation techniques using deep learning and differential evolution. Results from experiments comparing different forecasting models and presenting the efficacy of our proposed approach for stock selection optimization are discussed in the results section. Finally, the discussion and conclusion provide a view on the implications of our findings, how a model can be applied, and avenues for future research.

The remainder of this paper is organized as follows: Section 2 reviews the literature on stock price prediction and portfolio optimization, including their development and shortcomings. Section 3 discusses methodology, data sources, data preprocessing, and implementation techniques using deep learning and differential evolution. Results from experiments comparing different forecasting models and evaluating the effectiveness of our proposed approach for stock selection are presented in the results section. Finally, the discussion and conclusion provide an insight into the implications of our findings, how a model can be applied, and avenues for future research.

2. Literature review

2.1 Forecasting methods in stock selection

There are different schools in the literature showing that historical returns can be used to estimate future returns. The conventional techniques include those based on time series techniques (Sezer et al., 2020) and those based on computational intelligence techniques (Andriosopoulos et al., 2019; Bustos & Pomares-Quimbaya, 2020).

Time series analysis has been considered significant for a long period of time in stock price predictions. Traditional methods include the ARMA (AutoRegressive Moving Average) model (Wei et al., 2023), the ARIMA (AutoRegressive Integrated Moving Average) model (Wati et al., 2023), the ARCH (AutoRegressive Conditional Heteroskedasticity) model (Engle, 1982), and the GARCH (Generalized AutoRegressive Conditional Heteroskedasticity) model (Bollerslev, 1987). Recently, these models have been combined into hybrid frameworks, as in Adewole (2024) and Zolfaghari & Gholami (2021). For an overview of recent developments in financial time series forecasting see (Sezer et al., 2020). Unlike ARIMA, SARIMA (Seasonal AutoRegressive Integrated Moving Average) includes an additional set of autoregressive and moving-average terms to account for seasonality. Those correspond to the seasonal frequency and, hence, the SARIMA model intends to capture seasonal patterns in data but also the non-seasonal ones. Despite the wide use of these traditional methods for general purpose forecasting, their application in financial time series is limited by at least one of the following limitations: struggle to model complex, nonlinear relationships that are common in financial data (Zhang, 2003); requiring the data to be stationary, meaning constant mean and variance over time (William & Wei, 2006); highly sensitive to the selection of lag orders for AR and MA terms, making parameter estimation challenging (Nelson, 1991; Tsay, 2005); fail to capture the impact of rare but extreme market events (Black Swan events), which can lead to underestimation of risk (Taleb, 2010).

Recently, much emphasis has been placed on the application of computational intelligence techniques in stock price forecasting (Andriosopoulos et al., 2019; Bustos & Pomares-Quimbaya, 2020). The main algorithms that have been used in this field include support vector machines (SVMs), text analysis, and fuzzy logic systems. The leading bio-inspired methods are artificial neural networks and evolutionary algorithms. They are mainly commended for their ability to provide efficient forecasts and cope with complicated nonlinear dynamics in financial data. Artificial Neural Networks (ANNs) have now, for some time, enjoyed a leading position among the different computational intelligence techniques applied to classify, cluster, recognize patterns, and predict various scientific and technological disciplines (Abiodun et al., 2018).

The ANN is a versatile tool in computational intelligence that finds applications in all fields of study, including exact sciences, engineering, and social sciences. According to (Abiodun et al., 2018), ANNs were designed to mimic the architecture and function of the human nervous system. This model consists of interconnected neurons, the basic units of signal processing. These connections, as in a biological system, are weighted, and the weights are fine-tuned by training algorithms. This training enables ANNs to learn from data that can be used to model complex systems or phenomena effectively. The architecture of an ANN can vary greatly in complexity: it may have only one layer with one neuron. There are deeper ones with multiple layers that house more neurons. Neurons are arranged in layers with parallel association for inputs and outputs. Networks in which connections flow only in one direction, from input to output, are called feedforward networks. When the signal can flow back, they are called feedback networks (Abiodun et al., 2018).

Deep learning is a branch of ANNs characterized by particularly complex, multi-layered structures (Albawi et al., 2017). These networks typically feature more neurons and more intricate connections between layers than traditional ANNs, allowing them to capture more nuanced patterns in data. Examples of deep learning architectures include convolutional neural networks (CNNs), which are commonly used in image processing, recursive neural networks for hierarchical data, and recurrent neural networks (RNNs) for sequential data. The most representative RNN models are the Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures. LSTMs are designed to maintain information over very long sequences; therefore, they can handle tasks that require time-series data, such as handwriting recognition, speech recognition, and machine translation. Unlike other models, such as hidden Markov models, LSTMs do not require large amounts of input data (Albawi et al., 2017). GRUs, introduced by Chung and colleagues in 2014, work similarly to LSTMs but are simpler because they have a forget gate but lack an output gate (Chung et al., 2014). This results in fewer parameters, which can lead to faster training times while still achieving strong performance on sequence modeling tasks (Albawi et al., 2017). These advanced network architectures highlight the growing capabilities of ANNs to handle increasingly complex data and tasks. Advanced network architectures such as these illustrate how current ANN capabilities can be extended to handle increasingly complex data and tasks.

Several authors use a multi-layer perceptron with one input layer, several hidden layers, and one output layer, together with backpropagation for good forecasting performance (Hu et al., 2018), (Coyne et al., 2018) and (Zhong & Enke, 2017). Other authors propose a radial basis function network which performs very well in forecasting (Dash & Dash, 2016) and (Guo et al., 2017). A combination of fuzzy logic with a multi-layer perceptron to tackle the challenge of forecasting financial time series in (Khuat & Le, 2017). The Extreme Learning Machine (ELM) is a powerful algorithm designed for fast and efficient training of feedforward neural networks (Yang et al., 2019). Unlike traditional methods, ELM bypasses the need to iteratively adjust the parameters of hidden neurons, significantly speeding up the training process. Because of this fact, ELM can perform very well while making fewer computational costs. It has already been applied up to date, performed in financial markets, such as in stock prices' prediction, with outstanding performance shown (Huang & Li, 2017). ELM offers fast model training without the burden of time-consuming parameter tuning; this fact makes the technique quite plausible in applications where there is a need for speed towards arriving at a decision. Kyriakou et al. (2021) developed a stock return prediction-based optimization approach for both the short and long-term using a predictive non-parametric regression model. To reduce the variance in the short term, the authors decrease the long-term variance. In the case of short-term investments, noise is reduced because the predictive regressions for two different horizons are combined. On the other hand, Kaczmarek & Perez (2022) apply the random forest method of cross-sectional stock return predictions to Markowitz's quadratic portfolio optimization technique while Ma et al. (2020) apply three deep neural networks, that is DMLP, LSTM, and CNN in producing portfolio optimization strategies by means of predictions while retaining deep learning advantages and portfolio theory.

Yang et al. (2019) proposed a hybrid stock selection model that consists of the prediction step by using an ELM algorithm as an extremely efficient technique to train ANNs (Huang et al., 2006; Huang & Siew, 2004; Patel et al., 2015b). Patel et al. (2015a) empirically compared four predictive models such as ANN, support vector machines, random forest, and naive-Bayes by using two different approaches. They first extracted ten technical indicators from open, high, low, and close prices of stock trading data and then interpreted those indicators as deterministic trend data. Predictive accuracy for each model were tested using historical data of two Indian stock market stocks under both methodologies. In the case of feature selection and transformation, Chong et al. (2017) discussed the review of principal component analysis, *AutoEncoder*, and the restricted Boltzmann machine, using one of them to predict the performance of assets with some machine learning algorithm. Another relevant contribution that is worthwhile mentioning is that by Fischer & Krauss (2018), who attained some results in price forecasting using deep learning techniques.

2.2 Technical analysis

Technical analysis assesses price trends and patterns to guide the most opportune timing for buying and selling stocks. This type of approach may always notice a regular pattern in time-series stock prices (Han et al., 2024; Gorgulho et al., 2011a; Nazário et al., 2017; Nti et al., 2020). Among others, Gorgulho et al. (2011b) have proposed mechanisms to predict stock movements for the next period, in time to make investment decisions. Other authors merge several technical indicators using investor preferences with the theory of multi-criteria decision aid to set together an optimal portfolio, using non-hierarchal criteria structure (Fernandez et al., 2019, Fernandez et al., 2020).

Technical analysis bases its reasoning on market tendencies, demand, and supply of stocks (Achelis, 2001). In this way, price information becomes one factor to consider when making rules that can be manipulated to the decision-maker's benefit when selecting a list of stocks that meet these rules. For example, once a technical indicator suggests that the price of a particular stock is about to rise, the decision-maker should buy it, since it can be sold later at a higher price to create portfolio returns.

Technical analysts study price charts and use various indicators such as moving averages, relative strength index (RSI), moving average convergence-divergence (MACD), and Bollinger bands (BB) to plot trends and possible turning points. These indicators help in judging momentum, volatility, and possible reversals in the price of stocks. Volume also plays a very important role in technical analysis. Analysts look at volume patterns to confirm trends and any possible reversals. High volume during an uptrend or downtrend also shows strong momentum and may continue for some time longer. Head and shoulders, double tops and bottoms, and triangles are used to predict future movement. These are formed as stock prices continually change over time, and from these, future trend reversals and continuations are predicted.

Recently, several publications have extended technical analysis to include much more sophisticated computer techniques. Traditional technical analysis methodologies complement tools such as artificial intelligence and machine learning algorithms to improve the latter. These advanced methods can process large volumes of information to identify complex arrangements that might not be detectable using conventional techniques. Techniques such as SVMs, random forests, and neural networks have been employed to perform technical analysis (Kara et al., 2011; Patel et al., 2015b). Combining sentiment analysis with technical indicators has shown increased accuracy in the study. If more qualitative information such as news and social media is considered, so-called sentiment analysis is an additional layer of information that serves to support conventional technical analysis (Nassirtoussi et al., 2014). Technical analysis is widely used in high-frequency trading, where algorithms execute trades at extremely high speeds based on minute price changes. High-frequency trading firms leverage technical indicators and machine learning models to make split-second trading decisions (Aldridge, 2013).

Therefore, technical analysis remains an indispensable part of the ammunition that financial analysts and traders employ. Although it has moved with the times by embracing state-of-the-art computational techniques, the core principles of observing price patterns and trade volumes continue to play a significant role in leading investment decisions.

2.3 Fundamental analysis

Fundamental analysis is a technique for estimating the intrinsic value of a stock, bond, or any other financial instrument. While technical analysis is based on the study of historical price patterns and trading volumes, fundamental analysis is based on economic, financial, and qualitative factors that generally influence the value of such an asset. Therefore, the investor needs such analysis to identify undervalued or overvalued investments and make informed long-term decisions. The main objective of fundamental analysis is to determine an asset's intrinsic value using economic and financial variables. Comparing this intrinsic value with the market price will help the investor determine whether to buy or sell. If the intrinsic value is higher than the market price, the asset is considered undervalued; when the opposite is true, it is considered overvalued (Xidonas et al., 2009).

Fundamental analysts study macroeconomic factors, including inflation, interest rates, and fiscal and monetary policies. These types of ratios usually shed light on the overall state of the company's economic context (Marasovic et al., 2011). Analysts assess a company's financial health by analyzing its financial statements, cash flow, profitability, debt, and capital structure. Financial ratios help evaluate a company's ability to generate future profits. According to Xidonas et al. (2009).They help investigate the quality of the company's management, competitive advantage, industry conditions in which the company operates, and its market position through qualitative analysis using fundamental indicators. These factors, although somewhat harder to quantify, are important for assessing the company's potential long-term growth.

The fundamental documents comprising the balance sheet, income statement, and cash flow statement paint a clear and wide-reaching picture of a company's financial health and effectiveness in its operations. Some key ratios for comparison that help identify trends, assess efficiency, and evaluate profitability include P/E, ROA, ROE, debt-to-equity, and profit margin. Relative valuation compares a firm's financial ratios with those of its peers. It also evaluates the macroeconomic environment, comprising economic growth, inflation, and government policies, for their influence on the company's performance.

Fundamental documents, including the balance sheet, income statement, and cash flow statement, paint a clear, comprehensive picture of a company's financial health and operational effectiveness. Some of the key ratios for comparisons that will help identify the trend and assess efficiency and profitability include P/E, ROA, ROE, debt to equity, and profit margin. Relative valuation compares a company's financial ratios to those of its peers. It also assesses the macroeconomic environment comprising economic growth, inflation, and government policies to determine their influence on the company's performance.

It is also of particular importance in sustainable investing, where one can assess, using environmental, social, and governance criteria, a firm's financial health and performance. Such ESG analysis, in addition to fundamental analysis, therefore offers fuller insights into a company's intrinsic value and its capabilities for long-term value creation. Through sustainability reports, fundamental analysis can identify companies using responsible and sustainable practices by using ESG indicators. This risk involves considering issues such as environmental practices and regulatory compliance, which are important to investors when making decisions and reducing risks. Besides, such companies in sustainable practice are most likely to capture the long-term growth opportunities as more demands are realized for responsible products and services. It helps identify these opportunities and assesses growth potential in light of sustainability. This encourages better decision-making, whereby an investor can build a portfolio that, in addition to providing financial returns, contributes to sustainable development.

It is also of particular importance in sustainable investment, whereby the financial health and performance of the company can be found through environmental, social, and governance criteria. This ESG analysis, in addition to fundamental analysis, therefore offers a more complete view of a company's intrinsic value and its capabilities in terms of creating long-term value. Through sustainability reporting, fundamental analysis can identify companies that use responsible and sustainable practices through the use of ESG indicators. This risk involves consideration of issues such as environmental practices or regulatory compliance, which would be important for investors when making decisions and reducing some risks. Furthermore, companies that practice sustainability are more likely to capture long-term growth opportunities as demand for responsible products and services grows. Fundamental analysis helps to identify these opportunities and assess the growth potential in terms of sustainability, which supports better decision-making as it allows the investor to build a portfolio that, in addition to providing financial returns, contributes to sustainable development.

Fundamental analysis is being highly used for the selection of useful stocks. Although the data obtained in fundamental analysis can be qualitative, most of it takes the form of numerical values derived from companies' financial statements. Various studies aggregate these indicators into an overall evaluation index through a subjective process that often reflects the decision maker's policy (Xidonas et al., 2009). Aggregation itself poses its own set of challenges. It is also known that fundamental analysis differs between different companies based on their business activities, that is, between financial and non-financial firms (Marasovi'c et al., 2011). It is, thus, necessary to properly identify and utilize the most applicable indicators for conducting fundamental analysis (Xidonas et al., 2009).

Fundamental analysis is widely used for useful stock selection. Although the data obtained in fundamental analysis can be qualitative, most of it usually takes the form of a set of numerical values from the financial statements of companies. Various studies aggregate these indicators into an overall evaluation index, using a subjective process that often reflects the policy of the decision-maker (Xidonas et al., 2009). Aggregation itself poses its own set of challenges. It is also known that fundamental analysis differs between different companies based on their business activities, that is, between financial and non-financial firms (Marasović et al., 2011). It is, thus, necessary to properly identify and utilize the most applicable indicators for conducting fundamental analysis (Xidonas et al., 2009).

3. Materials and methods

The proposed financial management model aims to optimize stock portfolios through an integrated approach that combines advanced stock price prediction techniques and a robust stock selection methodology. Therefore, the model comprises two main components: stock price prediction and stock selection.

3.1 Stock price prediction

The first component of the model focuses on predicting future stock prices using various deep learning techniques. The main objective is to identify the most accurate prediction model among the most common ones by comparing their performance using historical stock price data. The canonical version of each of these algorithms is used in the procedure described below. The steps to carry out this procedure are as follows:

1. Identify the most used algorithms and tools to create predictive models based on time series and computational intelligence.
2. Gather historical data from companies listed on stock exchanges.
3. Perform simulations of the predictive models.
4. Compare the simulated predictive models.
5. Select the predictive models that best estimate stock prices.

For the first step, the literature review presented in Section 2 led us to four forecasting tools:

1. SARIMA: a generalization of an autoregressive moving average that considers seasonality.
2. RNN: Unlike the more common feedforward neural networks, RNNs are bidirectional artificial neural networks that allow for the processing of data in both forward and backward directions.
3. LSTM: an RNN intended to address the vanishing gradient problem present in traditional RNNs.
4. GRU: Special activation mechanism in recurrent neural networks.

The objective is to predict the stock's performance for the upcoming period:

$$\mathrm{R}_{(\mathrm{i},\mathrm{t}+1)} = \frac{\mathrm{P}_{(\mathrm{i},\mathrm{t}+1)} - \mathrm{P}_{(\mathrm{i},\mathrm{t})}}{\mathrm{P}_{(\mathrm{i},\mathrm{t})}} \tag{1}$$

Where $R_{(i,t+1)}$ denotes the return for stock i in the following period, with $P_{(i,t+1)}$ and $_{(i,t)}$ representing closing prices for the upcoming and current periods. The equation is expressed as follows to denote prediction rather than actual values:

$$\widehat{\mathrm{R}}_{(\mathrm{i},\mathrm{t}+1)} = \frac{\widehat{\mathrm{P}}_{(\mathrm{i},\mathrm{t}+1)} - \mathrm{P}_{(\mathrm{i},\mathrm{t})}}{\mathrm{P}_{(\mathrm{i},\mathrm{t})}} \tag{2}$$

$\widehat{\mathrm{R}}_{(\mathrm{i},\mathrm{t}+1)}$ and $\widehat{\mathrm{P}}_{(\mathrm{i},\mathrm{t}+1)}$ represent predicted return and anticipated closing price for the upcoming period, respectively. Thus, the equation can be rearranged to determine the predicted closing price if needed. Certain methods, like those proposed by (Yang et al., 2019), forecast the closing price directly. Nevertheless, it is not that convenient to compare forecasted close prices for different stocks. However, comparing predicted closing prices across different stocks is not straightforward. On the other hand, comparing the projected returns of various stocks is more direct and meaningful. In many ways, forecasting returns is preferable over forecasting the closing prices, as the former are usually stationary and hence more suitable for modeling, while the latter are non-stationary and prone to market noise. Returns are scale-invariant; therefore, it is easier to compare them across assets, they directly feed into investment decisions and risk management. Financial models like Sharpe ratio require the input of returns rather than raw prices. Therefore, return forecasting generally provides more reliable, interpretable, and actionable insights for financial analysis and portfolio optimization.

The forecast must be carried out individually for each stock; hence, in the case of neural networks, a preparation dataset is used for each stock, which is further divided into a training dataset and a testing dataset. A customized configuration of the network is developed for each stock to forecast its performance in period .

The preparation dataset from which the configuration of the network for any given stock is determined is defined as follows:

$$\mathrm{C} = \left\{ \left[\mathrm{t}_{(\mathrm{k}+1)}, \mathrm{v}_{(\mathrm{i},\mathrm{k})} \right] \middle| \mathrm{i} = 1{,}2\ldots,\mathrm{m}; \mathrm{k} = \mathrm{t}, \mathrm{t}-1, \ldots, \mathrm{t}-\mathrm{n} \right\} \tag{3}$$

$t_{(k+1)}$ represents the target value for the stock in the upcoming period, while $v_{(i,k)}$ denotes the value of input variable i for the stock during period k. m indicates the total number of input variables. The data is partitioned into a training dataset $C_{training}$ and a testing dataset $C_{testing}$. After the datasets have been prepared, normalization is applied to each time window of each variable to facilitate their use in the training algorithms, though the target values themselves remain unnormalized. Following the training of the algorithms, both the training error and the testing error (also referred to as accuracy) must be computed. A lower testing error indicates better predictive performance of the algorithms. However, because neural networks use stochastic procedures to determine weights and biases, results can vary across runs. Therefore, the algorithm is executed *n* times, and the iteration that produces the most favorable results is selected.

The results of performing the five steps mentioned above are described in Section 4.

3.2 Input information from technical analysis

The technical indicators to be used in the proposed model, according to the literature review and expert recommendations, include the following:

- Exponential Moving Average (EMA),
- Double Crossover (DC),
- Rate of Change (ROC),

- Relative Strength Index (RSI),
- Moving Average Convergence/Divergence (MACD),
- On-Balance Volume (OBV),
- Bollinger Bands (BB), and
- True Strength Index (TSI).

The EMA is one of the simplest technical indicators, assigning higher weights to more recent data. The EMA for the *i*-th stock in period t, EMA_t^i, is defined as follows (Macedo et al., 2017):

$$\mathrm{EMA}_t^i(\mathrm{ws}) = \left[\mathrm{p}_t^i - \mathrm{EMA}_{t-1}^i(\mathrm{ws})\right] \cdot \mathrm{w} + \mathrm{EMA}_{t-1}^i(\mathrm{ws})$$

where *ws* is the number of periods considered to compute the average. Usually, *ws* = 12 (Gorgulho et al., 2011a). The guideline linked to this indicator suggests that if the price line moves above the EMA line, the stock should be considered for buying.

The DC indicator relies on two MAs and indicates when one (shorter, usually five periods) crosses above the other (larger, usually twenty periods).

The ROC measures the difference, in percentage terms, between the current price of a stock and its price *h* periods earlier (Armano et al., 2005), with positive values indicating favorable conditions (Gorgulho et al., 2011a).

The RSI intends to measure the circumstances of a company regarding its stock price condition and if it is under- or over-valued. The RSI suggests buying the stock if its value rises above 30% and the current price exceeds that of the previous period (Gorgulho et al., 2011a; Macedo et al., 2017).

The MACD merges two EMAs to assess the appropriateness of purchasing a stock by comparing it to a signal line. Typically, this indicator is set up using two EMAs with 12 and 26 periods of historical data, referred to as $EMA(12)$ and $EMA(26)$, to create the $MACD(12,26)$ (Armano et al., 2005; Gorgulho et al., 2011a):

$$\mathrm{MACD}_t^i(12, 26) = \mathrm{EMA}_t^i(12) - \mathrm{EMA}_t^i(26) \quad (4)$$

The literature often plots another moving average that does not depend on stock prices but rather on the MACD indicator. This new moving average, MM_t^i, is used to create a momentum signal line over price movements. The signal line is created as a nine-period EMA of the MACD_t^i. The common strategy associated with this indicator states that when the value of $\mathrm{MACD}_t^i(12, 26)$ crosses above $\mathrm{MM}_t^i(9)$, there is evidence that the stock price will increase, making it advisable to invest in the current period (Armano et al., 2005; Gorgulho et al., 2011a).

The OBV considers that an increase in trading volume can signal an upcoming rise in stock price. If the OBV value is above the stock price, it suggests a strong upward trend.

BBs are strategies known for delivering solid net positive outcomes (Macedo et al., 2017). A BB functions as a volatility indicator, consisting of bands formed by taking a moving average of prices and subtracting two standard deviations of price fluctuations over the same period (Macedo et al., 2017).

The TSI uses two moving averages, which helps reveal the trend while also identifying overbought and oversold conditions (Gorgulho et al., 2011a). In practice, an EMA of the TSI is often used as a signal when it enters a specific zone, suggesting that the stock price is undervalued. The TSI guideline indicates that a stock should be considered for support if its TSI moves above the trigger point within the oversold region.

3.3 Input information from fundamental analysis

Investors who use fundamental analysis thoroughly review auditors' reports, profit-and-loss statements, quarterly balance sheets, dividend histories, and the policies of the companies they are monitoring. They assess sales figures, management effectiveness, production capacity, and the competitive landscape. Furthermore, they consult bank and treasury reports, production indices, price statistics, and crop forecasts to assess the broader business environment. After considering these aspects, the decision maker assesses whether the stock is currently undervalued; if so, it is considered a good alternative to buy.

It is well known that fundamental analysis varies by company type; for example, the analysis criteria required for financial institutions differ from those for non-financial institutions (Marasović et al., 2011). Thus, discussing the most appropriate indicators is a first step in any fundamental analysis. Some fundamental indicators that can be used regardless of the sector of operation are:

1. Return on assets: This is calculated as earnings before interest and taxes divided by total assets. It is a measure of efficiency in the use of assets in generating income.
2. Return on equity: Determined by net income divided by shareholder equity. In other words, how well and effectively a firm is using invested capital to generate profits.
3. Earning per share: The net income is divided by the number of outstanding shares after deducting preferred dividends on it.
4. Dividend yield: The annual dividends per share are divided by the price per share.

5. P/E ratio: This is calculated by dividing the market value per share by the earnings per share. This helps an investor to identify how much the market places a company's valuation in relation to earnings.
6. P/B ratio: Calculated through the division of the stock price by total assets minus intangible assets and liabilities. It is used to determine whether a stock is overvalued or undervalued by considering its net asset value.
7. Price-to-sales ratio: The share price is divided by the revenue per share. This helps in assessing the value the market places on each dollar of a company's sales.
8. Price-to-Cash flow ratio: Share price divided by cash flow per share is the ratio. This provides a measure of the stock's valuation relative to its cash-generating ability.

3.4 Differential evolution for portfolio optimization

Differential evolution is an effective method for selecting the most viable stocks by weighing their impact on relevant factors (predicted prices, and fundamental and technical indicators) to assign ratings to the stocks. The best stocks are subsequently picked to buy. This selection technique involves establishing weighting factors. The predicted prices as well as the technical values are standardized and merged to identify promising stocks. A linear combination is used to merge these factors, following the recommendations of Yang et al. (2019). Once the factors are normalized and the optimal weights for each factor are determined, the function indicated by Yang et al. (2019) is used to score the stocks, as shown in the following equation:

$$\mathrm{S(a)} = \sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{w_i f_i(a)} \tag{5}$$

f_1 represents the projected return $\widehat{\mathrm{R}}_{(\mathrm{i,t+1})}$ of stock *a* for the upcoming period, while f_j, j = 2, ..., N corresponds to the effects of this stock on various fundamental and technical indicators. The parameter w_i, i = 1, ..., N, refers to the weight assigned to the aspect *i*.

Higher scores indicate a stronger justification for buying stocks in the upcoming period. Consequently, a standard method for selecting promising stocks consists in ordering them focusing on a limited number of the highest-ranked companies.

Performance measure

A stock selection approach that utilizes a formula derived from stock scores and their past performance was introduced by Becker et al. (2007). The main idea in such an approach consists of giving more weight to the most significant factors so the higher-rated stocks would generate the greatest returns, while the lower-rated ones would yield the least. Thus, the objective is to maximize the gap between the current average return of the group of "selected stocks", $R_t(G_{\text{selected}})$, and the average return of the group of "not selected stocks", $R_t(G_{\text{discarded}})$, by adjusting the weights accordingly. These returns should be calculated as the mean returns of the stocks within each group. As a result, the objective is to identify the stock group that optimizes the following performance measure function (Yang et al., 2019):

Maximize

$$\frac{1}{\mathrm{P}} \sum_{\mathrm{p}=1}^{\mathrm{P}} (\mathrm{R_t}(\mathrm{G_{selected}}) - \mathrm{R_t}(\mathrm{G_{discarded}})) \tag{6}$$

subject to

$$\mathrm{l_{selected}} \leq \mathrm{card}(\mathrm{G_{selected}}) \leq \mathrm{u_{selected}}$$

where *P* refers to the number of iterations (i.e., the frequency with which G_{selected} is generated), G_{selected} denotes the group of selected stocks, and l_{selected} and u_{selected} represent the minimum and maximum sizes of G_{selected}, respectively.

Decision variables

The independent variables that influence the resolution of the problem include:

- The weights applied in Equation (1).
- The size of G_{selected} within the performance measurement function.

Algorithm for problem solving

Differential evolution has been recognized for its superior performance compared with other optimization methods on similar problems, particularly single-objective optimization tasks involving continuous variables. Consequently, differential evolution is employed here to determine the optimal parameter values for the proposed stock selection strategy.

Differential evolution works through a process of generational improvement that emulates natural selection. At any given time, a set of potential solutions, known as individuals or agents, is considered, collectively forming the population. The main parameters of differential evolution include the population size, $P_{\text{size}} \geq 4$, differential weight, $F \in [0,2]$, and crossover probability $CR \in [0,1]$. Individuals are described as $z = [z_1, z_2,..., z_m]^T$ of real numbers, where each of its components is a decision variable.

Let $y \in \mathbb{R}^m$ be a candidate solution (individual). The differential evolution algorithm follows these procedures:

1. Begin by randomly initializing individuals across the search space, assigning a random value to each decision variable while adhering to any relevant constraints.
2. Iterate the following steps until the stopping criterion is reached, which, in this case, is determined by a specified number of iterations, $N_{iterations}$.
3. For each individual z in the population:
   - Randomly pick three distinct individuals a, b, c from the population, ensuring that they are all unique.
   - Randomly choose an index $r \in \{1, ... , m\}$.
   - For every index i within $\{1, ... , m\}$:
     - Generate a random number u within the range $[0,1]$.
     - If $u < CR$ or $i = r$, set $y_i = a_i + F(b_i - c_i)$; otherwise, set $y_i = z_i$.
4. If the score $S(z)$ is less than or equal to $S(y)$, then z is replaced by y in the population (as defined in Equation (1)).

Standard parameter values are recommended for this algorithm

- $CR = 0.9$,
- $F = 0.8$,
- $P_{size} = 200$,
- Number of iterations = 100.

Finally, the top 5% of the companies considered should be selected.

3.5 Experimental setup and data

The experimental setup involves using historical stock price data from major companies and applying the aforementioned techniques to predict future prices and optimize stock selection. Data preprocessing includes normalizing the stock prices, splitting the data into training and testing sets, and implementing cross-validation to ensure the robustness of the models.

The analysis was conducted using Python 3.9 in Jupyter Notebook. Key libraries and frameworks include NumPy 1.21.2, Pandas 1.3.3 for data handling; TensorFlow 2.6.0 and Keras 2.6.0 for deep learning; and DEAP 1.3.1 for implementing differential evolution. Visualizations were created using Matplotlib 3.4.3 and Seaborn 0.11.2. Historical financial data was retrieved via the yfinance 0.1.63 library. The experiments were run on an Intel Core i7 with 16 GB DDR4 RAM and Windows 11.

The data used in this study consists of historical stock price information collected over a substantial period to ensure a comprehensive and reliable analysis. The dataset covers stock price data from January 1, 2010, to April 15, 2023. This period of over 13 years allows the models to capture various market conditions, including bullish and bearish trends, economic cycles, and market volatility. On the other hand, the stock price data was recorded on a daily basis, considering closing prices. The overall time span and periodicity were chosen to allow the models to learn from historical trends, seasonal patterns, and market anomalies, ensuring robustness in the prediction and stock selection processes.

Data sources: Historical stock prices are obtained from financial databases such as Yahoo Finance and Bloomberg.
Preprocessing: The data is normalized to ensure consistent scaling across different stocks.
Training and testing: The data is split into training (70%) and testing (30%) sets, with cross-validation applied to validate the models.
Evaluation metrics: The performance of the models is evaluated using MAPE, RMSE, and statistical tests (ANOVA).

4. Results

This section presents the outcomes of the stock price prediction models and the differential evolution-based portfolio optimization.

4.1 Stock price prediction performance

This section gives an overview of the results of the stock price prediction phase. It aims to identify the most appropriate method for our integrated approach by comparing ARIMA, RNN, LSTM, and GRU. The implementation and results are discussed in the following sections.

4.1.1 Implementation of prediction methods

Every element of the training set is a tuple, whose first element is a tensor of the five most recent known historical stock prices - excluding the last known price - and the objective value to be estimated is the last known price.

In order to create enough tensors along with corresponding objectives, a sliding window was used on the historic stock prices. For instance, the first tensor would be formed from prices between January 1, 2010 and

January 5, 2010, while the target is the price for January 6, 2010. After that, the window shifts one period forward to form the next tensor for prices ranging from January 2, 2010, up to January 6, 2010, with the target price for January 7, 2010. Of course, this process would continue until all periods in the training dataset are used up.

For this work, the historical stock prices from January 1, 2010, up to April 15, 2021, constitute the training set, while the prices from April 16, 2021, going forward until April 15, 2023, are considered the test set.

4.1.2 Performance measurement

The effectiveness of each prediction method is measured using the Mean Absolute Percentage Error (MAPE), the Mean Squared Error (MSE) and the Root Mean Squared Error (RMSE) defined as:

$$\mathrm{MAPE} = \frac{1}{\mathrm{n}} \sum_{\mathrm{t}=1}^{\mathrm{n}} \left| \frac{\mathrm{p_t} - \widehat{\mathrm{p}}_\mathrm{t}}{\mathrm{p_t}} \right|$$

$$\mathrm{MSE} = \frac{1}{\mathrm{n}} \sum_{\mathrm{t}=1}^{\mathrm{n}} \left(\mathrm{p_t} - \widehat{\mathrm{p}}_\mathrm{t} \right)^2$$

$$\mathrm{RMSE} = \sqrt{\mathrm{MSE}}$$

where p_t is the actual stock price at time t and $\widehat{\mathrm{p}}_\mathrm{t}$ is the predicted price.

4.1.3 Discussion of results

The generated models after training of the forecasting methods were used to forecast the stock prices present in the test set. The predicted values were compared with the real values. Table 1 presents a detailed comparison of the performance of the four predictive models—GRU, RNN, LSTM, and SARIMA—using the three standard regression metrics described in the previous subsection across various stocks.

The analysis of the MAPE values reveals that the GRU model consistently achieves the lowest error rates across most stocks, indicating superior predictive accuracy. For example, in the case of the AAPL stock, GRU reports a MAPE of 0.0256, outperforming RNN (0.0451), LSTM (0.0580), and SARIMA (0.2385). This trend holds across the majority of stocks, demonstrating GRU's robustness. RNN and LSTM models perform similarly but slightly worse than GRU. In contrast, SARIMA consistently yields the highest MAPE values, reflecting its poor ability to model the non-linear and volatile behavior of stock price movements. For instance, SARIMA's performance with the AAP stock shows a high MAPE of 0.2643, and for ZBH, it rises to 0.6252, confirming its inadequacy.

In terms of MSE, which measures the average squared differences between actual and predicted values, GRU continues to outperform the other models. For stock AAL, GRU has an MSE of 0.57, significantly lower than SARIMA's 30.23. SARIMA's high MSE values across various stocks, such as 3977.64 for AAP and 22477.34 for ZBRA, highlight its inability to capture complex market behaviors, leading to larger prediction errors. RMSE, the square root of MSE, provides error measurements in the same unit as the stock prices, making it easier to interpret. The GRU model consistently records the lowest RMSE values, reaffirming its prediction accuracy. For example, the RMSE for stock ZION is 1.99 with GRU, compared to 4.74 with SARIMA.

Table 1. Performance of different prediction methods measured through MAPE, MSE and RMSE.

Stock	MAPE				MSE				RMSE			
	GRU	RNN	LSTM	SARIMA	GRU	RNN	LSTM	SARIMA	GRU	RNN	LSTM	SARIMA
A	0.0282	0.0268	0.0558	0.2499	24.76	19.47	92.64	2343.89	4.98	4.41	9.63	48.41
AA	0.045	0.0515	0.0592	0.2499	11.16	14.65	24.37	93.47	3.34	3.83	4.94	9.67
AAL	0.0352	0.0349	0.0409	0.1019	0.57	0.56	0.77	30.23	0.76	0.75	0.88	5.5
AAP	0.029	0.0281	0.0435	0.2643	49.42	51.79	110.88	3977.64	7.03	7.2	10.53	63.07
AAPL	0.0256	0.0451	0.058	0.2385	24.44	62.42	108.68	1597.46	4.94	7.9	10.43	39.97
ABBV	0.0301	0.0528	0.0551	0.2236	28.07	76.4	88.15	3639.77	5.3	8.74	9.39	60.33
...	...	...	...	...	...	...	...	...	...	...	...	...
XYL	0.0246	0.0207	0.0313	0.3629	10.17	7.83	17.31	1871.54	3.19	2.8	4.16	43.26
YUM	0.017	0.0143	0.0368	0.4288	6.57	4.82	26.97	3996.92	2.56	2.19	5.19	63.22
ZBH	0.0183	0.0291	0.0217	0.6252	8.86	19.3	12.69	8647.69	2.98	4.39	3.56	92.99
ZBRA	0.0288	0.0232	0.0462	0.2272	220.61	144.96	721.12	22477.34	14.85	12.04	26.85	149.92
ZION	0.0278	0.0232	0.0356	0.0633	3.98	2.76	6.82	22.47	1.99	1.66	2.61	4.74
ZTS	0.0208	0.0166	0.0281	0.1328	23	14.72	48.36	894.94	4.8	3.84	6.95	29.92

Table 2. Descriptive Statistics for MAPE Values.

Statistic	GRU	RNN	LSTM	SARIMA
Number of Observations	559	559	559	559
Sum of Values	17.6967	19.0197	26.3933	175.7536
Mean Value	0.0317	0.0340	0.0472	0.3144
Sum of Squares	5.9561	3.4427	9.8449	101.0849
Sample Variance	0.0097	0.0050	0.0154	0.0821
Sample Standard Deviation	0.0983	0.0708	0.1241	0.2866
Standard Error of the Mean	0.0042	0.0030	0.0053	0.0121

Therefore, the results in Table 1 demonstrate the superiority of the GRU model in stock price forecasting. Its ability to minimize errors across all evaluated metrics confirms its effectiveness in capturing temporal dependencies and market dynamics. RNN and LSTM models also show competitive performance but are slightly less effective than GRU. SARIMA consistently underperforms, highlighting its limitations in handling the complex and non-linear nature of financial time series data. These findings strongly support the use of advanced deep learning models, particularly GRU, for more accurate and reliable stock price predictions. Table 2 supports these findings with descriptive statistics, highlighting the consistency and reliability of the deep learning models (GRU, RNN, LSTM) compared to SARIMA. The standard deviation and variance for SARIMA are much higher, suggesting greater prediction volatility and reduced reliability.

Table 3 presents the results of a One-Way ANOVA test comparing the mean MAPE values across the four models. The p-value is below the significance threshold of 0.05, indicating that there are significant differences in predictive performance among the models. The F-statistic (38.5319) confirms that at least one model performs differently, primarily pointing to SARIMA's poor performance.

Table 3. One-Way ANOVA of MAPE Values.

Source	Sum of Squares	Degrees of Freedom	Mean Square	F Statistic	p-Value
Treatment	32.1949	3	10.7316	38.5319	1.11e-16
Error	62.6170	2232	0.0281		
Total	94.8118	2235			

To verify the differences between means of all prediction methods, we conducted post hoc tests using Tukey HSD and multiple comparison tests (Scheffé, Bonferroni, and Holm) as shown in Tables 4 and 5. These tests identify which pairs of methods are significantly different.

Table 4. Tukey HSD Test Results.

Pair of Treatments	Statistic	p-Value	Inference
GRU vs RNN	0.3341	0.8999947	Insignificant
GRU vs LSTM	2.1961	0.4074583	Insignificant
GRU vs SARIMA	39.9125	0.0010053	p < 0.01
RNN vs LSTM	1.8620	0.5462276	Insignificant
RNN vs SARIMA	39.5784	0.0010053	p < 0.01
LSTM vs SARIMA	37.7164	0.0010053	p < 0.01

Table 5. Scheffé Multiple Comparison Test Results.

Pair of Treatments	Statistic	p-Value	Inference
GRU vs RNN	0.2362	0.9965508	Insignificant
GRU vs LSTM	1.5529	0.4916676	Insignificant
GRU vs SARIMA	28.2224	1.1102e-16	p < 0.01
RNN vs LSTM	1.3166	0.6295785	Insignificant
RNN vs SARIMA	27.9861	1.1102e-16	p < 0.01
LSTM vs SARIMA	0.2362	0.9965508	Insignificant

Tables 4 and 5 consistently show no significant difference in performance among the deep learning models, confirming that these models are statistically equivalent in predictive accuracy. However, SARIMA is statistically inferior to all three deep learning models, with p-values < 0.01, reinforcing its inadequacy for stock price forecasting. These results validate that the neural network-based models—GRU, RNN, and LSTM—were significantly more effective for stock price forecasting than traditional models like SARIMA in the context of the experiments addressed here.

After all, the results indicate that, while being strongly outperformed by SARIMA, the neural network methods GRU, RNN, and LSTM give more comparable and accurate predictions of stock prices. As one reviewer pointed out, this is probably because the SARIMA model is an extension of ARIMA models, which often violate the assumption of homoscedasticity on the residuals when applied to this type of data. Therefore, we no longer consider such a model in the following sections.

A further analysis was conducted to assess the sensitivity of these approaches. After establishing the baseline performance of the models without any perturbations, the models were evaluated using the original bias configurations on the test dataset. A Gaussian noise with a mean of 0 and a standard deviation of 0.01 was added to each weight and bias. This process simulates slight inaccuracies or fluctuations in the model parameters that can arise from training instability or real-world deployment noise. The baseline performance metrics were taken as the averages of the GRU columns in Table 1. After introducing perturbations, the models were re-evaluated on the same test dataset. The results of this analysis are shown in Table 6.

The results demonstrate that small perturbations had minimal impact on the GRU and RNN models, as indicated by only slight increases in MAPE, MSE, and RMSE. For the GRU model, the MAPE only increased by 0.08%, suggesting that it is highly robust to small parameter changes. The RNN model showed similar robustness. In contrast, the LSTM model exhibited a more significant performance drop, with a 0.23% increase in MAPE and a notable rise in both MSE and RMSE. This indicates that LSTM is more sensitive to perturbations, potentially due to its complex architecture and longer memory paths, which may be more vulnerable to small parameter fluctuations.

4.2 Performance of the stock selection

The second part of our model involves choosing the most promising stocks to invest in, using the predicted prices from the LSTM model (GRU or RNN can be used interchangeably), along with fundamental and technical analysis. We evaluate the performance of the selected stocks and compare it to the performance of various benchmark indices.

Accordingly, there are two fundamental benchmarks that can be adopted to measure the efficacy of a similar stock investment strategy. The first benchmark is the performance of the strategy against the market. This can be achieved by looking at the average return of stocks in a particular index or taking the index as a comparative standard for the investment strategy. A market index, like the S&P 500, aggregates the values of the stocks in it to provide an indication of the performance of the market segment that the index is meant to represent. These indexes are often applied to benchmark the performance of stock portfolios. Generally, it may not be appropriate to compare a portfolio's performance to a market index if the portfolio includes stocks that are not part of the index. However, since this research only covers stocks in the S&P 500 index, the comparison is both fair and relevant.

Accordingly, there are two fundamental benchmarks that can be adopted to measure the effectiveness of a similar stock investment strategy. The first benchmark is the performance of the strategy against the market. This can be achieved by looking at the average performance of stocks in a particular index or by taking the index as a comparative standard for the investment strategy. A market index, such as the S&P 500, aggregates the values of the stocks that comprise it to provide an indication of the performance of the market segment that the index is intended to represent. These indices are often applied to compare the performance of stock portfolios. Typically, it may not be appropriate to compare the performance of a portfolio to a market index if the portfolio includes stocks that are not part of the index. However, since this research only covers stocks from the S&P 500 index, the comparison is fair and relevant.

The second benchmark comes from existing literature. For this study, we utilize the methodology presented by (Yang et al., 2019). The method from their study proposed a scheme in which stock price predictions were done using an ANN and the ELM. The forecasted prices are used as one of the criteria for choosing the stocks, together with other fundamental factors. Weights on these factors are optimized using an objective function similar to that in our paper. The resources, in their method, are equally distributed between the selected stocks.

Table 6. Assessment of the approaches' sensitivity.

Stock	MAPE			MSE			RMSE		
	GRU	RNN	LSTM	GRU	RNN	LSTM	GRU	RNN	LSTM
Baseline performance	3.17%	3.40%	4.72%	12.34	14.67	19.56	3.51	3.83	4.42
Updated performance	3.25% (+0.08%)	12.89 (+0.55)	3.59 (+0.08)	3.48% (+0.08%)	15.12 (+0.45)	3.89 (+0.06)	4.95% (+0.23%)	20.77 (+1.21)	4.56 (+0.14)

The second point of reference comes from the existing literature. For this study, we used the methodology presented by Yang et al. (2019). The method of their study proposed a scheme in which stock price predictions were made using an ANN and the ELM. The predicted prices are used as one of the criteria for choosing stocks, along with other fundamental factors. The weights of these factors are optimized using an objective function similar to that of our paper. The resources, in their method, are distributed equally among the selected stocks.

The scoring for each stock was carried out by using the described combined weighted factors in Equation (1). The ranked scores were then used to select the top 5% of stocks for portfolio inclusion. The scoring system will differentiate between high-potential and low-potential stocks on the basis of their predicted performance and technical and fundamental indicators.

The scoring of each stock was carried out using the combined weighted factors described in equation (1). The ranked scores were then used to select the top 5% of stocks for inclusion in the portfolio. The scoring system will differentiate between high and low potential stocks based on their predicted performance and technical and fundamental indicators.

The market's average and the S&P 500 index serve as baseline benchmarks. The market's average return fluctuates significantly, reflecting broader economic and market conditions. The S&P 500 index generally shows more stability but still mirrors overall market trends, including downturns in December 2018, May 2019, and the significant drops in February and March 2020, driven by the COVID-19 pandemic.

The work of Yang et al. from 2019 consistently returns values that are above the average market and S&P 500 index for much of the periods. These significant returns are 14.72% in January 2019, 26.42% in April 2020, and 19.23% in November 2020. This approach performed well during market declines, such as in February and March 2020, with returns of -2.42% and -3.32%, respectively. On the other hand, the proposed approach also constantly outperformed the market average and the S&P 500 index. It achieved very strong performances in February 2019 (10.53%), June 2019 (13.45%), and July 2020 (18.17%). During market downturns, the proposed approach acted robustly with positive returns in February 2020 at 1.79% and smaller negative returns in March 2020 at -4.12%.

Early in 2019, both the Yang et al. method and the proposed approach achieved a much higher performance than all benchmarks. Specifically, the returns of the proposed method were 7.98% and 10.53% in January and February 2019, respectively, against S&P 500 returns at 7.29% and 2.89%. This indicates that only advanced prediction and selection models provided good performance in the case of stable recoveries. From March to December 2019, the proposed method continued to register high performances, especially in months like June at 13.45%, July at 9.16%, and October at 9.05%. The Yang et al. method also strongly performed, with an 8.37% performance in April and 9.50% in November.

Table 7. Returns of the benchmarks and the proposal. Returns generated by algorithms are calculated as the average over twenty iterations.

Date	Market's average	S&P's 500	Yang et al. (2019)	Selected stocks
November 2018	3.09%	1.75%	6.75%	0.92%
December 2018	-9.77%	-10.11%	-6.03%	-5.80%
January 2019	8.92%	7.29%	14.72%	7.98%
February 2019	4.41%	2.89%	8.12%	10.53%
March 2019	0.72%	1.76%	4.87%	6.55%
April 2019	3.71%	3.78%	8.37%	7.38%
May 2019	-6.54%	-7.04%	-0.70%	2.78%
June 2019	7.60%	6.45%	11.04%	13.45%
July 2019	1.13%	1.30%	4.11%	9.16%
August 2019	-2.46%	-1.84%	2.69%	6.74%
September 2019	2.88%	1.69%	3.70%	3.65%
October 2019	1.28%	2.00%	7.87%	9.05%
November 2019	3.41%	3.29%	9.50%	12.41%
December 2019	2.58%	2.78%	6.09%	7.40%
January 2020	0.67%	-0.16%	5.80%	9.00%
February 2020	-10.58%	-9.18%	-2.42%	1.79%
March 2020	-17.68%	-14.30%	-3.32%	-4.12%
April 2020	14.42%	11.26%	26.42%	20.49%
May 2020	5.29%	4.33%	11.99%	16.18%
June 2020	1.15%	1.81%	8.39%	6.63%
July 2020	4.88%	5.22%	11.46%	18.17%
August 2020	4.62%	6.55%	13.96%	12.76%
September 2020	-2.57%	-4.08%	0.41%	4.02%
October 2020	-0.80%	-2.85%	2.83%	6.42%
November 2020	14.28%	9.71%	19.23%	14.79%
December 2020	4.01%	3.58%	9.14%	11.09%
January 2021	-0.67%	-1.13%	4.36%	1.33%
February 2021	6.11%	2.54%	14.46%	9.37%
March 2021	5.63%	4.07%	9.24%	7.26%
April 2021	4.63%	4.98%	8.21%	8.15%

The market was highly volatile during the pandemic period, and could therefore be seen in extreme downturns during February -9.18% and March -14.30% of 2020. Both the proposed method and that by Yang et al. proved resilient, with less pronounced losses than those found in the benchmarks. Then came the surprise recovery in April and May 2020, with returns of 20.49% and 16.18%, beating the market and that of Yang et al., respectively. The proposed approach has been doing decently at the beginning of 2021. In February 2021, for example, it did better than the S&P 500 index, returning 9.37% compared to 2.54%. The performance of the Yang et al. method was relatively good for February 2021, with a return of 14.46%. This is assured by the ANOVA on three independent samples and Tukey's HSD test on the differences among the three observed performances of the proposed method, the Yang et al. method, and benchmarks at 0.95 confidence. These statistics support the reliability and effectiveness of the approach and method proposed by Yang et al., particularly in generating higher returns and better managing risk than traditional market benchmarks.

Analyzing the results from Table 7, one can note that the proposed stock selection outperformed the market average and the S&P 500 index. It also compares favorably with the approach of Yang et al., even in times of turbulence and recovery in the market. The robustness of the results of the proposed approach is assured by the consistency of outperformance and statistical significance, hence proving the validity of the approach as a capable tool for optimization of the stock portfolio.

In contrast, the Yang et al. proposed method had a Sharpe ratio of 1.09 with a Sortino ratio of 5.34, while our proposed system provided a Sharpe ratio of 1.23 and a Sortino ratio of 4.04. The Tukey's HSD test based on ANOVA for three independent samples supported the statistical significance of these differences at the 95% level of confidence. These performances enlighten the risk-adjusted returns of the investment strategies. The Sharpe ratio, which reflects the average return earned over and above the risk-free rate per unit of volatility, serves as evidence that our proposed system indeed outperforms the method by Yang et al. in terms of risk-adjusted performance. Similarly, the Sortino ratio considers only downside volatility against the general volatility and serves as an indication of the strength of the proposed method in dealing with negative returns more effectively. The high Sharpe and Sortino ratios for our proposed system show that besides yielding strong returns, it also demonstrates effective risk management in outperforming the other strategies. This therefore supports the efficiency of our integrated model, which combines advanced prediction techniques with optimized stock selection to deliver superior investment performance.

Conclusions

This paper proposes a comprehensive model for stock price prediction and stock selection that uses advanced computational techniques to improve investment decision-making. The deep learning algorithm for portfolio optimization, combined with differential evolution, has shown a significant improvement in results compared to traditional methods and benchmarks.

By comparing the various prediction models, it has been observed that GRU, RNN, and LSTM provided better accuracy than the SARIMA model in predicting stock prices. These three neural network-based methods are statistically equivalent, while SARIMA is significantly worse. This indicates that the use of advanced deep learning techniques for financial time series predictions is efficient. Furthermore, the stock selection component of the model that combines return forecasts with technical and fundamental analysis indicators performed well. The results of the proposed method were consistently above the market average, S&P 500 index, and literature benchmarks over various time periods of market downturns. The resilience and superior returns of the proposed approach indicate its potential to become a good tool for stock portfolio optimization. ANOVA and Tukey's HSD test further validated these performance differences statistically. High Sharpe and Sortino ratios evidence that the proposed approach did a great job of keeping risks within controlled levels and producing higher risk-adjusted returns.

The combination of deep learning for stock price prediction and differential evolution for portfolio optimization provides a solid framework for improving investment strategies. This model not only improves forecast accuracy but also optimizes stock selection, leading to an increasingly better performing portfolio. Additional financial indicators could be incorporated in the future, and this model could be applied to other financial markets to further validate its generalizability and effectiveness.

Future research lines include further optimization that considers resource allocation to selected assets and further assessment of forecasting methods. Additionally, it is convenient to perform backtesting in more recent periods.

Conflict of interest

The authors have no conflict of interest to declare.

Funding

The authors received no specific funding for this work.

References

Abiodun, O. I., Jantan, A., Omolara, A. E., Dada, K. V., Mohamed, N. A., & Arshad, H. (2018). State-of-the-art in artificial neural network applications: A survey. *Heliyon*, *4*(11), e00938.

Achelis, S. B. (2001). *Technical Analysis from A to Z*. McGraw Hill New York.

Adewole, A. I. (2024). On the hybrid of ARIMA and GARCH model in modelling volatilities in Nigeria Stock Exchange. *Bima Journal of Science and Technology*, *8*(2A), 169-180.

Albawi, S., Mohammed, T. A., & Al-Zawi, S. (2017). Understanding of a convolutional neural network. *2017 International Conference on Engineering and Technology (ICET)*, 1–6. https://doi.org/10.1109/ICEngTechnol.2017.8308186

Aldridge, I. (2013). *High-frequency trading: a practical guide to algorithmic strategies and trading systems*. John Wiley & Sons.

Andriosopoulos, D., Doumpos, M., Pardalos, P. M., & Zopounidis, C. (2019). Computational approaches and data analytics in financial services: A literature review. *Journal of the Operational Research Society*, *70*(10), 1581-1599. https://doi.org/10.1080/01605682.2019.1595193

Armano, G., Marchesi, M., & Murru, A. (2005). A hybrid genetic-neural architecture for stock indexes forecasting. *Information sciences*, *170*(1), 3-33. https://doi.org/10.1016/j.ins.2003.03.023

Becker, Y. L., Fei, P., & Lester, A. M. (2007). Stock Selection: An Innovative Application of Genetic Programming Methodology. In *Genetic and evolutionary computation* (pp. 315–334). https://doi.org/10.1007/978-0-387-49650-4_19

Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *The review of economics and statistics*, 542-547. https://doi.org/10.2307/1925546

Bustos, O., & Pomares-Quimbaya, A. (2020). Stock market movement forecast: A systematic review. *Expert Systems with Applications*, *156*, 113464. https://doi.org/10.1016/j.eswa.2020.113464

Chong, E., Han, C., & Park, F. C. (2017). Deep learning networks for stock market analysis and prediction: Methodology, data representations, and case studies. *Expert Systems with Applications*, *83*, 187-205. https://doi.org/10.1016/j.eswa.2017.04.030

Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*. https://doi.org/10.48550/arXiv.1412.3555

Coyne, S., Madiraju, P., & Coelho, J. (2018). Forecasting stock prices using social media analysis. *Proceedings - 2017 IEEE 15th International Conference on Dependable, Autonomic and Secure Computing, 2017 IEEE 15th International Conference on Pervasive Intelligence and Computing, 2017 IEEE 3rd International Conference on Big Data Intelligence and Compu*, *2018-Janua*, 1031–1038. https://doi.org/10.1109/DASC-PICom-DataCom-CyberSciTec.2017.169

Das, S., & Suganthan, P. N. (2010). Differential evolution: A survey of the state-of-the-art. *IEEE transactions on evolutionary computation*, *15*(1), 4-31. https://doi.org/10.1109/TEVC.2010.2059031

Dash, R., & Dash, P. K. (2016). A comparative study of radial basis function network with different basis functions for stock trend prediction. *2015 IEEE Power, Communication and Information Technology Conference, PCITC 2015 - Proceedings*, 430–435. https://doi.org/10.1109/PCITC.2015.7438204

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the econometric society*, 987-1007. https://doi.org/10.2307/1912773

Nazário, R. T. F., e Silva, J. L., Sobreiro, V. A., & Kimura, H. (2017). A literature review of technical analysis on stock markets. *The Quarterly Review of Economics and Finance*, *66*, 115-126. https://doi.org/10.1016/j.qref.2017.01.014

Fernandez, E., Navarro, J., Efrain, S., & Coello Coello, C. (2020). Using evolutionary computation to infer the decision maker's preference model in presence of imperfect knowledge: A case study in portfolio optimization. *Swarm and Evolutionary Computation*, *54*, 100648.
https://doi.org/10.1016/j.swevo.2020.100648

Fernandez, E., Navarro, J., Solares, E., & Coello, C. C. (2019). A novel approach to select the best portfolio considering the preferences of the decision maker. *Swarm and Evolutionary Computation*, *46*(July 2018), 140–153.
https://doi.org/10.1016/j.swevo.2019.02.002

Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research*, *270*(2), 654-669.
https://doi.org/10.1016/j.ejor.2017.11.054

Gorgulho, A., Neves, R., & Horta, N. (2011a). Applying a GA kernel on optimizing technical analysis rules for stock picking and portfolio composition. *Expert Systems with Applications*, *38*(11), 14072–14085.
https://doi.org/10.1016/j.eswa.2011.04.216

Gorgulho, A., Neves, R., & Horta, N. (2011b). Applying a GA kernel on optimizing technical analysis rules for stock picking and portfolio composition. *Expert Systems with Applications*, *38*(11), 14072–14085.
https://doi.org/10.1016/j.eswa.2011.04.216

Guo, Z., Ye, W., Yang, J., & Zeng, Y. (2017). Financial index time series prediction based on bidirectional two dimensional locality preserving projection. In *2017 IEEE 2nd International Conference on Big Data Analysis (ICBDA)* (pp. 934-938). IEEE.
https://doi.org/10.1109/ICBDA.2017.8078775

Han, Y., Liu, Y., Zhou, G., & Zhu, Y. (2024). Technical analysis in the stock market: A review. *Handbook of Investment Analysis, Portfolio Management, and Financial Derivatives: In 4 Volumes*, 1893–1928.
https://doi.org/10.1142/9789811269943_0059

Hu, H., Tang, L., Zhang, S., & Wang, H. (2018). Predicting the direction of stock markets using optimized neural networks with Google Trends. *Neurocomputing*, *285*, 188–195.
https://doi.org/10.1016/j.neucom.2018.01.038

Huang, C. F., & Li, H. C. (2017). An Evolutionary Method for Financial Forecasting in Microscopic High-Speed Trading Environment. *Computational Intelligence and Neuroscience*, *2017*(1), 9580815.
https://doi.org/10.1155/2017/9580815

Huang, G.-B., & Siew, C.-K. (2004). Extreme learning machine: RBF network case. *2004 8th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, *2*(January), 1029–1036.
https://doi.org/10.1109/icarcv.2004.1468985

Huang, G.-B., Zhu, Q.-Y., & Siew, C.-K. (2006). Extreme learning machine: Theory and applications. *Neurocomputing*, *70*, 489–501.
https://doi.org/10.1016/j.neucom.2005.12.126

Kaczmarek, T., & Perez, K. (2022). Building portfolios based on machine learning predictions. *Economic research-Ekonomska istraživanja*, *35*(1), 19-37.
https://doi.org/10.1080/1331677X.2021.1875865

Kara, Y., Boyacioglu, M. A., & Baykan, Ö. K. (2011). Predicting direction of stock price index movement using artificial neural networks and support vector machines: The sample of the Istanbul Stock Exchange. *Expert systems with Applications*, *38*(5), 5311-5319.
https://doi.org/10.1016/j.eswa.2010.10.027

Khuat, T. T., & Le, M. H. (2017). An application of artificial neural networks and fuzzy logic on the stock price prediction problem. *JOIV: International Journal on Informatics Visualization*, *1*(2), 40-49.

Kyriakou, I., Mousavi, P., Nielsen, J. P., & Scholz, M. (2021). Short-term exuberance and long-term stability: A simultaneous optimization of stock return predictions for short and long horizons. *Mathematics*, *9*(6), 620.
https://doi.org/10.3390/math9060620

Ma, Y. L., Han, R. Z., & Wang, W. Z. (2020). Prediction-Based Portfolio Optimization Models Using Deep Neural Networks. *IEEE Access*, *8*(15), 11540–15393.
https://doi.org/10.1109/access.2020.3003819

Macedo, L. L., Godinho, P., & Alves, M. J. (2017). Mean-semivariance portfolio optimization with multiobjective evolutionary algorithms and technical analysis rules. *Expert Systems with Applications*, *79*, 33–43.
https://doi.org/10.1016/j.eswa.2017.02.033

Marasović, B., Poklepović, T., & Aljinović, Z. (2011). Markowitz'Model with Fundamental and Technical Analysis–Complementary Methods or Not. *Croatian operational research review*, *2*(1), 122–132.

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, *7*(1), 77–91.
https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Nassirtoussi, A. K., Aghabozorgi, S., Wah, T. Y., & Ngo, D. C. L. (2014). Text mining for market prediction: A systematic review. *Expert Systems with Applications*, *41*(16), 7653-7670. https://doi.org/10.1016/j.eswa.2014.06.009

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, 347-370. https://doi.org/10.2307/2938260

Nti, I. K., Adekoya, A. F., & Weyori, B. A. (2020). A systematic review of fundamental and technical analysis of stock market predictions. *Artificial Intelligence Review*, *53*(4), 3007–3057. https://doi.org/10.1007/s10462-019-09754-z

Patel, J., Shah, S., Thakkar, P., & Kotecha, K. (2015a). Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. *Expert systems with applications*, *42*(1), 259-268. https://doi.org/10.1016/j.eswa.2014.07.040

Patel, J., Shah, S., Thakkar, P., & Kotecha, K. (2015b). Predicting stock market index using fusion of machine learning techniques. *Expert systems with applications*, *42*(4), 2162-2172. https://doi.org/10.1016/j.eswa.2014.10.031

Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied soft computing*, *90*, 106181. https://doi.org/10.1016/j.asoc.2020.106181

Solares, E., De-Leon-Gomez, V., Salas, F. G., & Díaz, R. (2022). A comprehensive decision support system for stock investment decisions. *Expert Systems with Applications*, *210*, 118485. https://doi.org/10.1016/j.eswa.2022.118485

Solares, E., Salas, F. G., De-Leon-Gomez, V., & Díaz, R. (2022). A comprehensive soft computing-based approach to portfolio management by discarding undesirable stocks. *IEEE Access*, *10*, 40467–40481. https://doi.org/10.1109/ACCESS.2022.3167153

Taleb, N. N. (2010). *The Black Swan:: The Impact of the Highly Improbable: With a new section:" On Robustness and Fragility"* (Vol. 2). Random house trade paperbacks.

Tsay, R. S. (2005). *Analysis of financial time series*. John wiley & sons.

Wati, M., Haviluddin, A. M., Septiarini, A., & Hatta, H. R. (2023). Autoregressive Integrated Moving Average (ARIMA) Model for Forecasting Indonesian Crude Oil Price. *Jurnal Ilmiah Teknik Elektro Komputer dan Informatika (JITEKI)*, *9*(3), 720-730.

Wei, Y., Wang, F., Young, H., Ke, D., & Rodriguez, J. (2023). Autoregressive moving average model-free predictive current control for PMSM drives. *IEEE Journal of Emerging and Selected Topics in Power Electronics*, *11*(4), 3874–3884. https://doi.org/10.1109/JESTPE.2023.3275562

William, W., & Wei, S. (2006). Time series analysis: univariate and multivariate methods. *USA, Pearson Addison Wesley, Second Edition. Cap*, *10*, 212–235.

Wong, T., & Luo, Z. (2018). Recurrent auto-encoder model for large-scale industrial sensor signal analysis. In *International Conference on Engineering Applications of Neural Networks* (pp. 203-216). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-98204-5_17

Xidonas, P., Mavrotas, G., & Psarras, J. (2009). A multicriteria methodology for equity selection using financial analysis. *Computers & operations research*, *36*(12), 3187-3203. https://doi.org/10.1016/j.cor.2009.02.009

Yang, F., Chen, Z., Li, J., & Tang, L. (2019). A novel hybrid stock selection method with stock prediction. *Applied Soft Computing Journal*, *80*, 820–831. https://doi.org/10.1016/j.asoc.2019.03.028

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, *50*, 159–175. https://doi.org/10.1016/S0925-2312(01)00702-0

Zhong, X., & Enke, D. (2017). A comprehensive cluster and classification mining procedure for daily stock market return forecasting. *Neurocomputing*, *267*, 152–168. https://doi.org/10.1016/j.neucom.2017.06.010

Zolfaghari, M., & Gholami, S. (2021). A hybrid approach of adaptive wavelet transform, long short-term memory and ARIMA-GARCH family models for the stock index prediction. *Expert Systems with Applications*, *182*, 115149. https://doi.org/10.1016/j.eswa.2021.115149