



(RESEARCH ARTICLE)



Enhancing stock market prediction accuracy with recurrent deep learning models: A case study on the CAC40 index

Arash Tashakkori ^{1,*}, Niloufar Erfanibehrouz ², Shahin Mirshekari ³, Abolfazl Sodagartoaji ⁴ and Vatsal Gupta ⁵

¹ School of Business, Stevens Institute of Technology, Hoboken, NJ 07030, USA.

² Department of Economics, Syracuse University, Syracuse, New York.

³ Department of Marketing Science and Business Analytics, Katz Graduate School of Business, University of Pittsburgh, PA, USA.

⁴ Department of Statistics, Rutgers University, 110 Frelinghuysen Rd Piscataway, NJ 0885, USA.

⁵ School of Management, Yale University, 165 Whitney Avenue, New Haven, CT 06511, USA.

World Journal of Advanced Research and Reviews, 2024, 23(01), 2309–2321

Publication history: Received on 09 June 2024; revised on 21 July 2024; accepted on 24 July 2024

Article DOI: <https://doi.org/10.30574/wjarr.2024.23.1.2156>

Abstract

This paper explores the application of Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks for stock price prediction over a 10-day horizon. The study aims to compare the predictive performance of these two deep learning architectures within the context of financial forecasting. Utilizing historical stock data from the CAC40 dataset, which represents a capitalization-weighted measure of the 40 most significant stocks on the Euronext Paris, we train and evaluate RNN and LSTM models to forecast future stock prices. Our results demonstrate the superior performance of LSTM networks in capturing the intricate temporal dependencies inherent in stock price data. Compared to standard RNNs, LSTM models exhibit higher accuracy and provide more reliable forecasts over the 10-day prediction period. The specialized memory cells and gating mechanisms in LSTM networks enable them to effectively identify both short-term changes and long-term patterns in stock prices, thus outperforming traditional RNN architectures. This enhanced ability to model the complex dynamics of stock market data underscores the potential of LSTM networks to improve investment decision-making, risk management, and the overall efficiency of financial markets. The insights gained from this study contribute to the growing body of knowledge on the application of deep learning in finance and investment, offering valuable guidance for practitioners and researchers seeking to harness the power of advanced algorithms for stock market prediction and optimization.

Keywords: Stock market prediction; Recurrent neural networks; LSTM; Deep learning; CAC40 index; Investment strategies; Financial forecasting

1. Introduction

Navigating the intricate terrain of financial markets, characterized by their volatility and multifaceted dynamics, remains an enduring challenge for investors and analysts alike. Central to this challenge is the capability to forecast future stock prices accurately, a pursuit critical for informed investment decisions, risk mitigation, and the optimization of returns within the ever-shifting market landscape. Stock prediction holds profound importance in finance and investment. Accurate stock price forecasts help investors determine whether to invest in, trade, or retain what they have invested. securities, thereby maximizing their returns and minimizing potential losses. In a world where financial markets are shaped by a multitude of factors, ranging from economic indicators to geopolitical events to investor sentiment and the psychology of the market, the ability to anticipate future price movements assumes paramount significance. Furthermore, stock prediction plays a pivotal role in risk management strategies. By providing insights

* Corresponding author: Arash Tashakkori

into potential market trends and fluctuations, accurate forecasts enable investors to hedge against adverse movements and protect their portfolios from unexpected downturns. This proactive approach to risk mitigation is instrumental in safeguarding wealth and ensuring long-term financial stability. Beyond individual investors, accurate stock prediction also holds broader implications for the overall health and stability of financial markets. Inefficient or inaccurate predictions can lead to market volatility, speculative bubbles, and systemic risks, with far-reaching consequences for economies and societies at large. Conversely, reliable forecasts contribute to market efficiency, liquidity, and investor confidence, fostering an environment conducive to sustainable economic growth and development. [1-5]

Traditionally, stock prediction has leaned heavily on fundamental analytical techniques and statistical models. Machine learning leverages statistical models to make predictions and decisions based on data [41-44]. But the introduction of machine learning, particularly the revolutionary strides in deep learning, has catalyzed a paradigm shift in this field. These advancements offer powerful tools capable of analyzing vast datasets and uncovering intricate patterns that were previously elusive. Among these tools, recurrent networks such as Long Short-Term Memory (LSTM) networks have emerged as a beacon of promise in recent years, redefining the landscape of stock prediction. Deep learning, a branch of machine learning, uses multiple layer artificial neural networks to uncover complicated patterns from data. Deep learning models can manage large volumes of unstructured data, including text, pictures, and time series data, since they autonomously develop hierarchical representations of the data, in contrast to traditional machine learning techniques that depend on human feature extraction. [6-10]. Within the deep learning domain, LSTM networks have garnered particular attention for their unique ability to capture long-term dependencies in sequential data. Designed as a remedy to the vanishing gradient problem that plagues standard Recurrent Neural Networks (RNNs), LSTM networks are adept at modeling temporal dependencies across extended time horizons. This distinctive capability makes them particularly well-suited for analyzing time series data characterized by intricate temporal dynamics, such as stock prices. [11-15]

Despite the formidable advancements in predictive modeling, forecasting stock values remains a daunting task due to their inherently unpredictable nature over the long term. While the Efficient Market Hypothesis suggests that stock values behave randomly, recent technical analyses have underscored the pivotal role of historical data in achieving effective predictions. Moreover, the movements of stock market groups are intricately altered by various economic factors, ranging from political events and general economic conditions to investor psychology and sentiment. This intricate interplay of factors not only complicates the prediction process but also poses significant challenges to traditional statistical models, often leading to inaccuracies in value and movement forecasts. However, the advent of machine learning, and more specifically, deep learning, has offered a ray of hope in deciphering these complexities. The utilization of Recurrent Neural Networks (RNNs), and particularly LSTM networks, has shown considerable promise in capturing the nuanced patterns inherent in historical stock price data. [16-20]

In this paper, we embark on an extensive exploration of the application of RNN and LSTM networks for stock prediction, delving deep into their capabilities, limitations, and potential impact on investment decision-making. Additionally, we endeavor to bridge the gap between theoretical advancements and practical implementations by scrutinizing RNN and LSTM performance within the unique context of the index for the French equity market, CAC40. We develop a RNN and LSTM model on the dataset of CAC40, previously known as Bourse de Paris, which is a capitalization-weighted indicator of the top 40 important stocks out of the top 100 market caps listed on Euronext Paris. Cotation Assistée en Continu, which translates to "continuous assisted trading," is the designation for this benchmark index that has been used by companies who trade in the French stock market. Through this endeavor, we aim to enrich our understanding of RNN and LSTM efficacy in diverse financial landscapes and pave the way for more informed and effective investment strategies. The introduction should be typed in Cambria with font size 10. Author can select Normal style setting from Styles of this template. The simplest way is to replace (copy-paste) the content with your own material. In this section highlight the importance of topic, making general statements about the topic and presenting an overview on current research on the subject. Your introduction should clearly identify the subject area of interest.

The introduction should be typed in Cambria with font size 10. Author can select Normal style setting from Styles of this template. The simplest way is to replace (copy-paste) the content with your own material. In this section highlight the importance of topic, making general statements about the topic and presenting an overview on current research on the subject. Your introduction should clearly identify the subject area of interest.

2. Related Works

Machine learning (ML) stands out as a potent toolset comprising diverse algorithms tailored to enhance performance across various case studies. There's a widespread acknowledgment of ML's profound capability to discern meaningful insights and discern patterns within datasets. In contrast to conventional ML approaches, ensemble models represent

a paradigm where multiple algorithms collaborate synergistically to tackle specific challenges, consistently demonstrating superior performance, especially in time series prediction [21-24]. Boosting and bagging, two prevalent techniques within the ensemble paradigm, have emerged as effective strategies for enhancing predictive accuracy. Recent advancements in tree-based models, such as gradient boosting and XGBoost algorithms, have gained traction among leading data scientists, showcasing remarkable efficacy in competitive scenarios. However, the advent of deep learning (DL) represents a modern trend in ML, leveraging intricate, nonlinear structures to extract pertinent information from financial time series data. In the realm of financial analysis, recurrent neural networks (RNNs) have garnered significant acclaim due to their exceptional performance [25-28]. Predicting stock market trends entails more than just considering current data; historical information plays a pivotal role. Consequently, training models solely on the latest data proves insufficient. RNNs excel in capturing the memory of recent events and establishing connections between network units, rendering them well-suited for economic forecasting. Long Short-Term Memory (LSTM) networks represent a refined variant of RNNs, specifically tailored for deep learning applications. By employing three distinct gates, LSTM addresses inherent challenges encountered in standard RNN cells, offering enhanced processing capabilities for individual data points or entire sequences [29].

Long et al. [30] conducted an analysis utilizing a to evaluate changes in stock prices, a model based on deep neural networks is applied in conjunction with transaction records and public market data. Their experimental results showed that bidirectional LSTM outperformed other prediction models in terms of performance and showed higher predictive ability for financial choices. Pang et al. [31] aimed to improve stock market forecasts using neural network architectural suggestions. They presented Both an automated encoder-equipped LSTM with an embedded layer to assess changes in the stock market. According to the results, the deep LSTM with an embedded layer performed better than other models, obtaining 57.2% and 56.9% accuracy for the Shanghai composite index, respectively. Kelotra and Pandey [32] effectively analyzed stock market trends utilizing a deep multilayer LSTM model. Their lowest MSE and RMSE values were 2.6923 and 7.2487, according to the results of their training with a Rider-based monarch butterfly optimization technique. Additionally, Bouktif et al. [33] investigated using a more advanced sentiment analysis technique how predictable stock market trend orientation are. Using deep learning techniques, their suggested approach surpassed existing sentiment-based prediction algorithms by successfully identifying shifts in stocks with a 60% accuracy rate. Rekha and associates. [34] compared the outcomes of CNN and RNN algorithms against actual stock market data. Lee et al. [35] utilized CNNs to forecast global stock market trends and evaluated their model using data from various countries. Their findings indicated the model's adaptability to large datasets and its effectiveness even in markets with limited data availability. Liu et al. [36] investigated a numerical-based attention approach to find synergy between news and numerical data for stock price prediction utilizing dual-source stock market data. Their approach performed better in dual-source prediction than earlier models and successfully filtered noise. Baek and Kim [37] proposed a method for stock market index forecasting that combined an overfitting avoidance LSTM module with a prediction LSTM module. The better predicting accuracy of their suggested model over models without an overfitting avoidance mechanism was confirmed by the results. A combination LSTM and GA technique was used by Chung and Shin [38] to improve a unique stock market prediction model, showing better results than benchmark models. Zhou et al. [39] used a rolling partition training and testing set technique to apply CNN and LSTM to high-frequency stock market data in order to evaluate the effect of update cycles on model performance. Their models effectively reduced error and enhanced forecast accuracy, depending to their extensive evaluation. Results and discussion (Heading 1, WJS heading level 1)

This section should be typed in character size 10pt Cambria and alignment justified. Author can directly select Normal style from styles of this template. The results and discussion may be combined into a common section or obtainable separately. They may also be broken into subsets with short, revealing captions. An easy way to comply with the conference paper formatting requirements is to use this document as a template and simply type your text into it.

Article Subheadings as well as Article sub-sub headings formatting can be used as in material and methods section

Reference should be cited at appropriate point in the text by number(s) in square brackets in line with the text. e.g.: '..... was reported earlier [3, 6].'

The actual authors can be referred to, but the reference number(s) must always be given. e.g.: 'Barnaby and Jones [8] obtained a different....'

3. Methods and Materials

In this section, we detail the methodologies employed in our study to predict stock prices using Long Short-Term Memory (LSTM) networks and Recurrent Neural Networks (RNNs). Preprocessing the dataset helps to guarantee consistency by eliminating any outliers or inconsistencies that might negatively impact the model's performance. Then,

to extract the temporal connections present in sequential stock price data and predict future price movements, we utilize two different neural network architectures: RNNs and LSTMs.

3.1. Data

The dataset used in this study was obtained from a finance website and comprises daily stock data of companies listed on the CAC40 index from 2010 to 2021 [40]. Tracking the performance of the 40 largest publicly listed firms on the Euronext Paris stock exchange, the CAC40, short for Cotation Assistée en Continu, is the benchmark French stock market index. The CAC40 index, which had a base value of 1,000 when it was first created in December 1987, is a capitalization-weighted index that selects its member companies based on market capitalization, trading activity, balance sheet size, and liquidity. The list underwent a transition from a total market capitalization system to a free float-adjusted methodology in 2003. Comprising multinational corporations with diverse market reach, the CAC40 index is a preferred choice for foreign investors seeking exposure to the French market. It serves as a key indicator of the overall direction of the Euronext Paris exchange, formerly known as the Paris Bourse. The composition of the CAC40 index is reviewed quarterly by an independent steering committee. Companies are rated according to share turnover during the previous year and free-float market capitalization. The index is composed of the top 40 firms ranked among the top 100 market capitalization on the exchange. In cases where a company has multiple classes of shares, only the most actively traded shares are considered for index inclusion.

3.2. RNN and LSTM

Recurrent Neural Networks (RNNs) are a family of artificial neural networks designed to process sequential data by maintaining internal memory. They excel in tasks involving time series data, such as stock price prediction, thanks to their inherent feedback loops, unlike classic feedforward neural networks that do not preserve information over time [45-47]. However, ordinary RNNs face the vanishing gradient problem, which limits their effectiveness in capturing long-range dependencies in sequential data. To overcome this issue, Long Short-Term Memory (LSTM) networks were developed. LSTMs employ specialized memory cells with gating mechanisms to control input flow, thereby improving the capabilities of RNNs. LSTMs are particularly adept at capturing intricate temporal patterns and relationships in sequential data, as their memory cells can selectively retain or discard information over extended periods.

3.2.1. RNN

Text, audio, and time series data are examples of sequential data that can be processed by Recurrent Neural Networks (RNNs), a specific type of neural network designed to store and recall information from previous inputs. RNNs leverage dependencies and context across time steps, making them well-suited for tasks such as voice recognition, language translation, and time series forecasting.

The fundamental building blocks of an RNN are recurrent connections, which allow information to persist over time. The network generates an output vector ht , representing the hidden state that serves as the network's memory at each time step t . It takes an input vector x . Subsequently, the network utilizes this hidden state and the current input vector as inputs for the next time step. The basic structure of an RNN consists of recurrent connections that allow information to persist over time. At each time step t , the network receives an input vector xt and produces an output vector ht , which serves as the hidden state representing the network's memory at that time step. This hidden state is then fed back into the network as input for the next time step, along with the current input vector. Mathematically, the hidden state ht at time step t is computed as follows:

$$h_t = f(W_{hx} \times x_t + W_{hh} \times h_{t-1} + b_h) \dots\dots (1)$$

In the abovementioned formula, xt is the input vector at time step t , $ht-1$ is the hidden state from the previous time step, Whx and Whh are weight matrices for the input and recurrent connections, respectively, bh is the bias vector, and f is the activation function, commonly a nonlinear function like the hyperbolic tangent (tanh) or rectified linear unit (ReLU). In this cell, Whh represents the weight matrix that governs the recurrent connections within the recurrent neural network (RNN). In the context of RNNs, Whh captures the influence of the previously hidden state $ht-1$ on the current hidden state ht . It determines how information from the previous time step is incorporated into the current time step's computation. Essentially, Whh is a matrix of parameters that controls the flow of data over the network's recurring connections. The network adjusts these parameters to lessen the difference between the predicted and actual results after learning them during the training phase. The output of the RNN at each time step may be used directly for various activities or may need to undergo further processing, depending on the objectives and network architecture.

Since RNNs are naturally adaptable, they may be used in many different ways, depending on the demands of the given job. These configurations include one-to-one, one-to-many, many-to-one, and many-to-many architectures. The vanishing gradient problem, on the other hand, prevents conventional RNNs from fully capturing long-range relationships in sequential data.




Figure 1 An RNN unit

3.2.2. LSTM

Long Short-Term Memory (LSTM) networks are a sort of recurrent neural network (RNN) architecture intended to remedy the matter of the fading away gradient encountered in standard RNNs, allowing them to effectively model temporal dependencies over extended time horizons. LSTMs are particularly well-suited for analyzing time series data with complex temporal dynamics, such as stock prices. The key idea behind LSTM is LSTMs excel in preserving long-term dependencies in sequential data by mitigating the vanishing gradient problem. This is accomplished using specialized units known as LSTM cells, which include a series of gates that control the information flow. The LSTM model consists of multiple LSTM cells arranged sequentially, with each cell processing input data and propagating information forward through time. Each LSTM cell contains several components, including a cell state, an input gate, a forget gate, and an output gate. The operations within an LSTM cell are governed by a set of mathematical equations. These equations control how information is inputted into the cell, how it is stored or forgotten in the cell state, and how it is output to either the final prediction or the following cell.

Each LSTM cell comprises several interconnected components that work together to control the information transfer process. Among an LSTM cell's fundamental components are:

- **Cell State:** The cell state serves as the memory of the LSTM cell and is responsible for retaining information over long sequences. It can be selectively updated or cleared through the action of gates.
- **Input Gate:** This establishes the amount of fresh data that should be added to the cell state. It receives data from both the prior hidden state h_{t-1} , and the current input x_t , passes it through a sigmoid activation function, and controls the flow of information into the cell state.
- **Forget Gate:** Any information from the cell state should be ignored or forgotten is controlled by the forget gate. It regulates the flow of information from the cell state by taking input from the current input x_t and the prior hidden state h_{t-1} . It does this by passing the input via a sigmoid activation function.
- **Output Gate:** This gate establishes the appropriate amount of information from the cell state to be shared with the following cell or the final prediction. After passing the data via a sigmoid activation function, it controls the output of information investigating the cell state using the input from the previous hidden state, h_{t-1} , and the current input x_t .
- **Cell State Update:** Based on the values of the input gate, forget gate, and candidate cell state, the cell state is updated. Whereas the forget gate decides what data should be erased, the input gate governs how much new information is added to the cell state. Next, the modified cell state is transferred to the subsequent time step.
- **Hidden State:** The information processed and transferred to the next cell or the final prediction is represented by the hidden state h_t , which is the output of the LSTM cell. It is calculated using the output gate and the current cell state.

In LSTM, Whh represents the weight matrix in the LSTM cell that is connected to the recurrent connections. As with conventional RNNs, it controls how data from the previous hidden state is used to the computation of the current hidden state. To avoid the vanishing gradient problem and manage information flow, LSTM cells use a weight matrix in conjunction with input, forget, and output gates. LSTM is especially useful for modeling complicated temporal relationships in sequential data, like stock prices, because of its capacity to selectively remember or forget information across lengthy sequences.




Figure 2 An LSTM unit

Table 1 RNN vs LSTM

Parameters	RNNs	LSTMS
Structure	Simple	More complex
Training	Can be difficult	Can be more difficult
Performance	Good for simple tasks	Good for complex tasks
Hidden state	Single	Memory cell
Gates	-	Input, output, forget
Ability to retain long-term dependencies	Limited	Strong

4. Experiments

4.1. Data visualization

In this section, we introduce several figures representing the data under analysis. As we proceed with modeling on the CAC40 dataset, we randomly choose a company's market for examination. The forthcoming visualization pertains to the AirLiquide market. Figure 3 shows the line chart illustrating the closing prices of the CAC40 stocks over time. The x-axis represents the dates ranging from 2010 to 2021, while the y-axis denotes the corresponding closing prices. Each data point on the graph depicts the daily closing price of the CAC40 stocks, providing a visual depiction of the fluctuation in prices over the specified period. From the visualization, we observe the overall trend and any notable patterns or anomalies in the closing prices throughout the years, offering valuable insights into the historical performance of the CAC40 stocks.

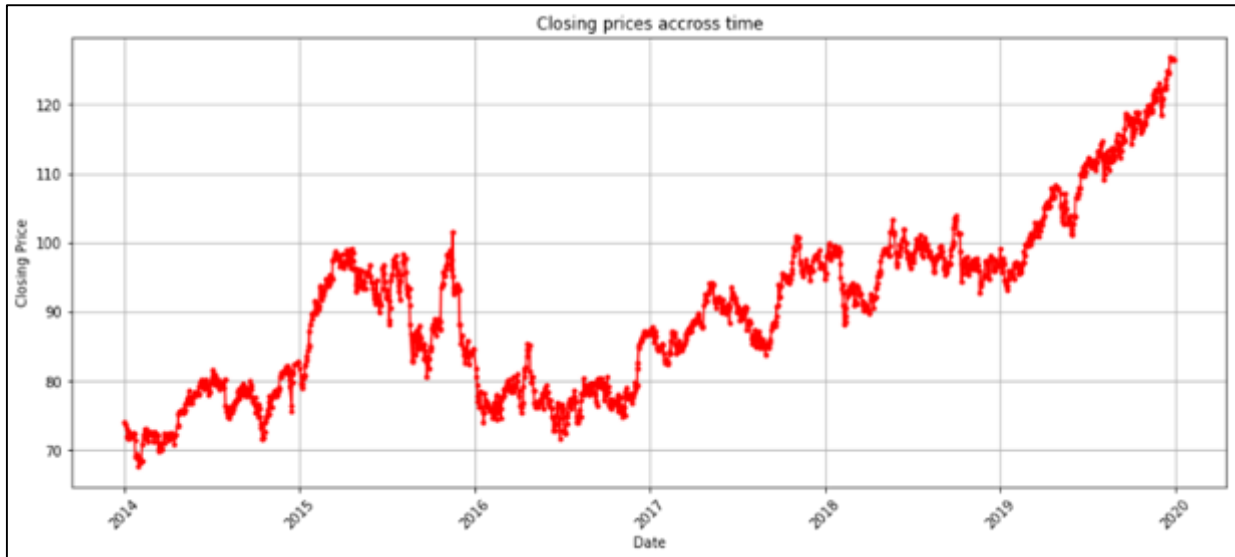


Figure 3 Closing prices of AIR Liquide across time.

In Figure 4, we present a candlestick chart depicting the price movements of the CAC40 stocks (AirLiquide) over time. Each candlestick on the chart represents the trading range (open, high, low, close) for a specific period, typically a day. The candlestick's color and shape indicate whether the closing price was higher (red) or lower (blue) than the opening price, offering insights into bullish or bearish market sentiment.

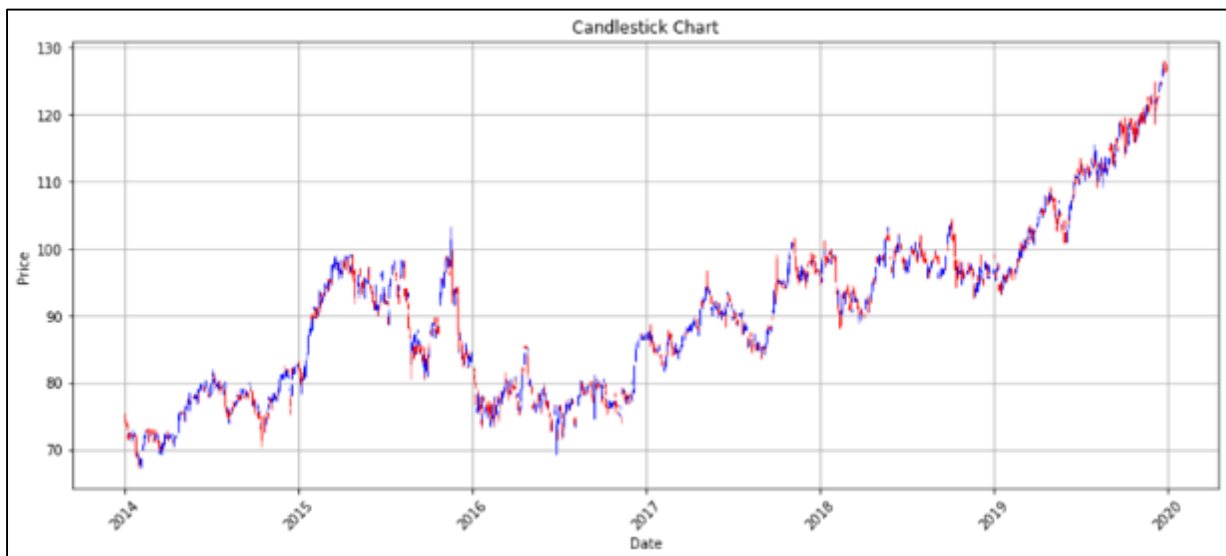


Figure 4 Candlestick chart of AIR Liquide

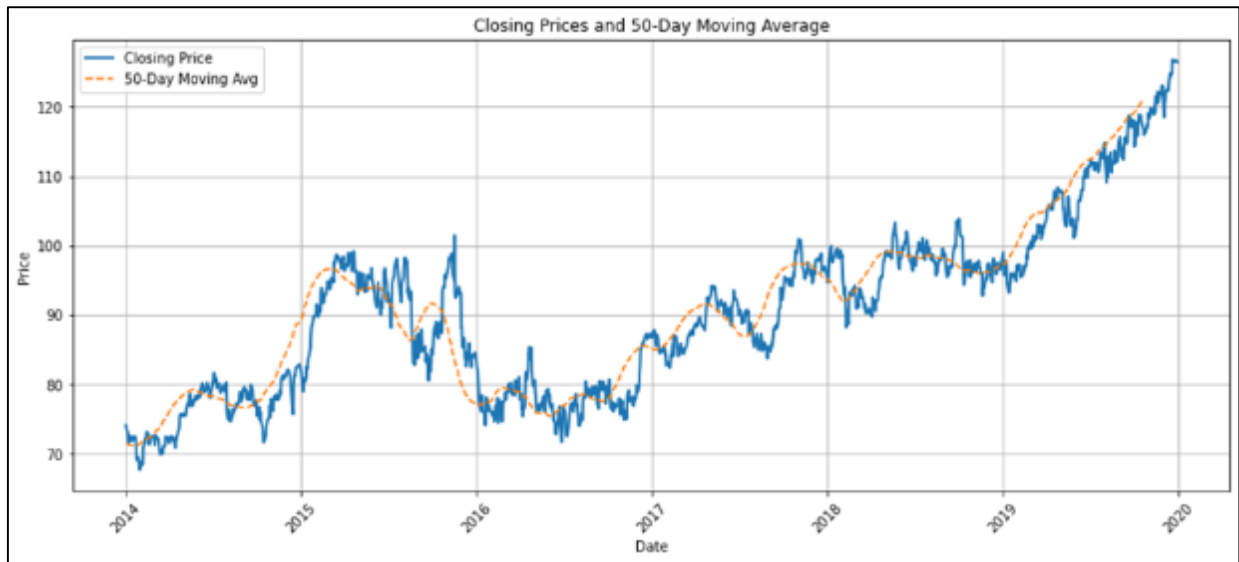


Figure 5 Closing prices and 50-Day moving average of AIR Liquide

The tails or outlines on each candlestick indicate the greatest and lowest prices attained during the trading phase, while the body of the stick shows the range of prices between the opening and closing values. This visualization allows for a detailed examination of price trends, volatility, and potential reversal patterns in the CAC40 stock market (AirLiquide) over the analyzed period.

We showcase the closing prices of AirLiquide over time along with a 50-day moving average in Figure 5. The solid line represents the daily closing prices, providing insights into the overall price trajectory. Additionally, the 50-day moving average, shown by the dotted line, tames short-term swings and draws attention to the data's underlying trend. Examining the connection between closing prices and moving average can help identify significant price movements, trend reversals, and potential trading opportunities in the AirLiquide market.

This visualization in Figure 6 illustrates the monthly seasonality of AirLiquide closing prices. The plot displays the average closing price for each month, allowing us to identify recurring patterns or trends that may exist throughout the year. By analyzing the monthly averages, the ability for analysts and investors to learn more about the seasonal variations in stock prices and make more informed decisions regarding investment strategies, timing of trades, and portfolio management.

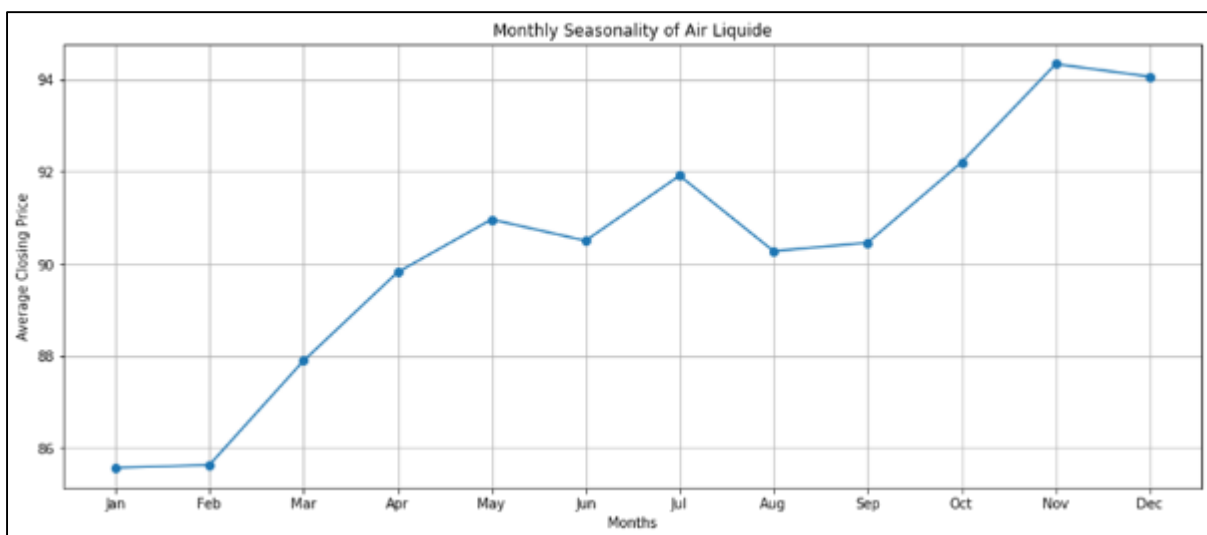


Figure 6 Monthly seasonality of AIR Liquide

4.2. Pre-processing

In this segment, we preprocess our data to prepare it for training and testing with the RNN and LSTM models. Firstly, we normalize the dataset to scale the values between 0 and 1, facilitating better convergence during training. Next, we divided the data into two sets: one for training and the other for testing, with 80% of the data going toward training and 20% toward testing. The amount of previous time steps the model will take into account for each prediction is represented by the sequence length, which we then define. In this case, the sequence length is set at 60. Next, for the training and testing sets, we create sequences of input-output pairs. For each data point in the training and testing sets, the input sequence comprises the previous 60 closing prices, and the corresponding output is the subsequent closing price. Finally, we display the dimensions of the training and testing sets to verify the sizes. The training set consists of 1151 sequences, each containing 60 time steps, and the corresponding target values. Similarly, the testing set comprises 243 sequences, each with 60-time steps and their corresponding target values.

4.3. Model training

The This study's LSTM model architecture consists of a dense output layer and three LSTM layers with dropout regularization. Starting with the first LSTM layer, which is setup with 50 units and programmed to return sequences so that it may propagate the output to the next layer, the model is built sequentially. Next, 20% of the units are randomly deactivated in a dropout layer with a dropout rate of 0.2 to reduce overfitting. The second LSTM layer mirrors the configuration of the first, maintaining 50 units and returning sequences, with an additional dropout layer set to the same dropout rate. Similarly, the third LSTM layer retains 50 units but does not return sequences, followed by another dropout layer. Finally, a dense output layer with a single unit is added to produce the model's prediction.

Upon compiling the model to quantify the difference between expected and actual results, the mean squared error (MSE) loss function is used. During training, the model's parameters are iteratively adjusted by using the Adam optimizer to minimize this loss function. Early Stopping, in the meantime, keeps an eye on the validation loss and stops training if no progress is shown for 15 consecutive epochs. In the meantime, the weights of the model that performs the best are restored. With a total of 50,851 parameters, all of which are trainable, the model undergoes training for 100 epochs using a batch size of 32. During training, the model is fed sequences of input data and corresponding target values obtained from the training set. The validation dataset is utilized to evaluate the model's performance after each epoch and prevent overfitting. By incorporating dropout layers and early stopping mechanisms, the LSTM model aims to learn the underlying patterns within the financial time series data while minimizing the risk of overfitting and maximizing predictive accuracy.

5. Results

To forecast stock prices using RNN and LSTM models, we conducted experiments on the CAC40 dataset. By leveraging the RNN and LSTM architectures, we sought to anticipate stock values for the future via previously collected information. Conducting our research, we trained the models on a portion of the dataset and assessed their performance on both training and testing data using metrics like root mean squared error (RMSE) [48-49]. The utilization of RNN and LSTM models allowed us to capture temporal dependencies and intricate patterns present in the stock market data, thereby facilitating more accurate forecasts of future stock prices. When comparing the root mean squared error (RMSE) values of the RNN and LSTM models applied to the stock data, distinct differences emerge. For the RNN model, the RMSE for the training data stands at approximately 98.45, while for the test data, it slightly decreases to about 84.19. Conversely, the LSTM model exhibits superior performance, with an RMSE of approximately 92.36 for the training data and a notably lower value of about 76.18 for the test data. These findings suggest that the LSTM model generally outperforms the RNN counterpart, as it achieves lower RMSE values across both training and testing datasets. The LSTM model's improved predictive capacity in capturing the underlying patterns and dynamics of the stock market data is indicated by the lower RMSE values, which show that the model's predictions match the actual stock prices more closely.

In Figure 7, the plot showcases the actual closing prices of Air Liquide's shares provided by LSTM, denoted by the black line. Overlaid on this are the predicted prices derived from the LSTM model for both the training and testing sets, depicted in red and blue, respectively. The train predictions are represented from the beginning of the dataset up to the start of the test predictions, while the test predictions extend from that point to the end of the dataset. The plot provides a comparative view of how well the LSTM model's predictions align with the actual closing prices of the company's shares over time. The red line illustrates the model's performance on the training data, demonstrating its capacity to identify deeper trends and patterns. Meanwhile, the blue line extends into the future, showcasing the model's predictive capability on unseen test data.

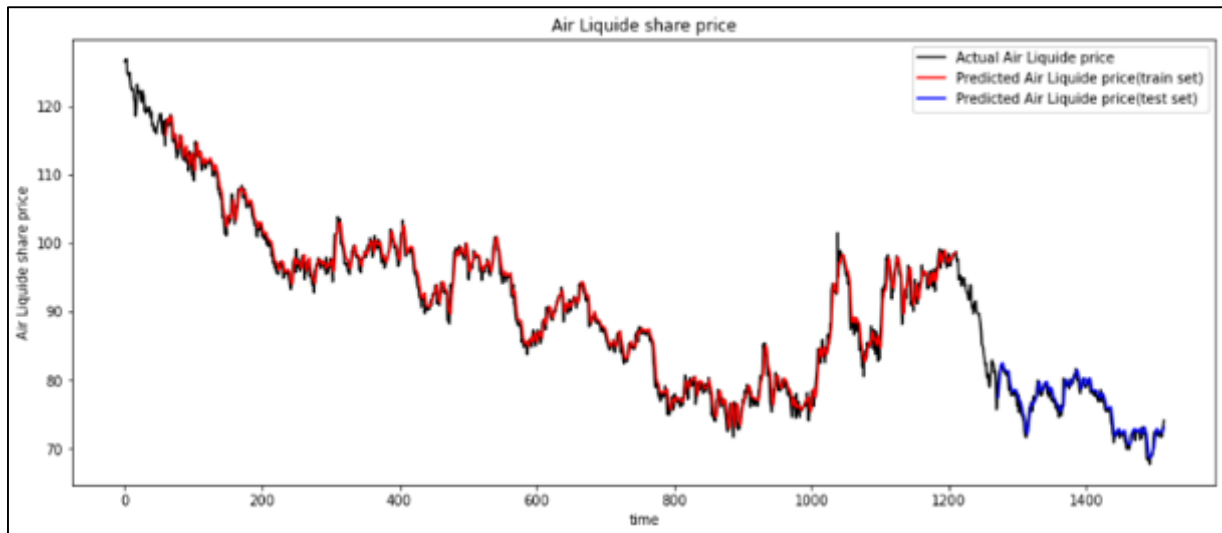


Figure 7 AIR Liquide share price

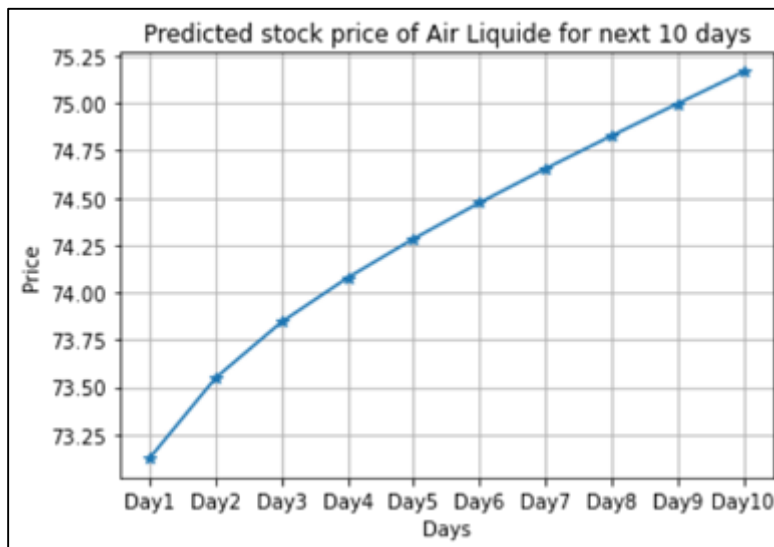


Figure 8 Predicted stock price of AIR Liquide for the next 10 days

By visualizing the model's predictions alongside the actual prices, stakeholders can assess the LSTM model's effectiveness in capturing the dynamics of the stock market and its potential utility in making informed investment decisions. One important measure of the model's accuracy and dependability in predicting future price changes is the level of consistency between the expected versus actual prices.

Figure 8 provides the forecast for the next 10 days of stock prices, the LSTM model utilizes the last `n_past` days of data as input. During each iteration, the model predicts the price for the next day based on the last sequence of data. The predicted value is appended to a list of predictions for the next 10 days. Additionally, the sequence is shifted by one day to accommodate the new prediction, and the last element of the sequence is updated with the new prediction. The printed predictions display the forecasted prices for the next 10 days, with each day labeled sequentially from Day 1 to Day 10. These predictions offer valuable information for stakeholders and investors, enabling them to anticipate potential price movements and make informed decisions regarding their investments. Concerning Figure 8, the forecasted market price is 73.25 on Day 1, with an increase to 75.25 by Day 10.

6. Conclusion

In conclusion, our study delved into the realm of stock price prediction using recurrent neural network (RNN) and long short-term memory (LSTM) models applied to the CAC40 dataset. Through meticulous experimentation and analysis, we observed the efficacy of these deep learning architectures in capturing the intricate patterns and temporal dependencies inherent in stock market data. The RNN and LSTM models demonstrated promising results, with LSTM outperforming RNN in terms of prediction accuracy, as evidenced by lower root mean squared error (RMSE) values on both training and testing datasets. By harnessing the power of deep learning, we were able to forecast future stock prices with greater precision, thereby providing valuable insights for investors and financial analysts. Our findings underscore the significance of leveraging advanced machine learning techniques for stock market forecasting, paving the way for enhanced decision-making and risk management strategies in the realm of finance.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

- [1] Nabipour, M., Nayyeri, P., Jabani, H., Mosavi, A., Salwana, E., & S, S. (2020). Deep learning for stock market prediction. *Entropy*, 22(8), 840.
- [2] Staffini, A. (2022). Stock price forecasting by a deep convolutional generative adversarial network. *Frontiers in artificial intelligence*, 5, 837596.
- [3] V. Monjezi, A. Trivedi, G. Tan and S. Tizpaz-Niari, "Information-Theoretic Testing and Debugging of Fairness Defects in Deep Neural Networks," 2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE), Melbourne, Australia, 2023, pp. 1571-1582, doi: 10.1109/ICSE48619.2023.00136.
- [4] Larijani, A., & Dehghani, F. (2023). An Efficient Optimization Approach for Designing Machine Models Based on Combined Algorithm. *FinTech*, 3(1), 40-54.
- [5] Larijani, A., & Dehghani, F. (2023). A Computationally Efficient Method for Increasing Confidentiality in Smart Electricity Networks. *Electronics*, 13(1), 1-15. <https://doi.org/10.3390/electronics13010170>
- [6] Naseri, H., Jahanbakhsh, H., Foomajd, A., Galustanian, N., Karimi, M. M., & D. Waygood, E. O. (2022). A newly developed hybrid method on pavement maintenance and rehabilitation optimization applying Whale Optimization Algorithm and random forest regression. *International Journal of Pavement Engineering*, 1-13.
- [7] Hashemi, A., Jang, J., & Beheshti, J. (2023). A Machine Learning-Based Surrogate Finite Element Model for Estimating Dynamic Response of Mechanical Systems. *IEEE Access*. 11, 54509-54525, 2023, doi: 10.1109/ACCESS.2023.3282453.
- [8] Sadeghi, S., Marjani, T., Hassani, A., & Moreno, J. (2022). Development of Optimal Stock Portfolio Selection Model in the Tehran Stock Exchange by Employing Markowitz Mean-Semivariance Model. *Journal of Finance Issues*, 20(1), 47-71.
- [9] Malmir, M., Momeni, H., & Ramezani, A. (2019, April). Controlling megawatt class WECS by ANFIS network trained with modified genetic algorithm. In 2019 27th Iranian Conference on Electrical Engineering (ICEE) (pp. 939-943). IEEE.
- [10] Tehranian, K., Khorsand, M. S., Zarei, M., Arani, G. G., Banabari, H. G., & Sasani, F. (2024) Unveiling the Impact of Social Media Usage on Firm Performance: The Mediating Influence of Organizational Agility and Innovation Capability, *Tehnički glasnik*, 18, 1-9. doi:10.31803/tg-20230918233848
- [11] Vahdatpour, M. S., & Zhang, Y. (2024). Latency-Based Motion Detection in Spiking Neural Networks. *International Journal of Cognitive and Language Sciences*, 18(3), 150-155.
- [12] Karabulut, E., Gholizadeh, F., & Akhavan-Tabatabaei, R. (2022). The value of adaptive menu sizes in peer-to-peer platforms. *Transportation Research Part C: Emerging Technologies*, 145, 103948.

- [13] Dutta, A., Masrourisaadat, N., & Doan, T. T. (2022, December). Convergence Rates of Decentralized Gradient Dynamics over Cluster Networks: Multiple-Time-Scale Lyapunov Approach. In 2022 IEEE 61st Conference on Decision and Control (CDC) (pp. 6497-6502). IEEE. <https://doi.org/10.1109/CDC51059.2022.9992900>
- [14] Manshour N, He F, Wang D, Xu D. Integrating Protein Structure Prediction and Bayesian Optimization for Peptide Design. In NeurIPS 2023 Generative AI and Biology (GenBio) Workshop 2023. <https://openreview.net/forum?id=CsjGuWD7hk>
- [15] Aghamohammadghasem, M., Azucena, J., Hashemian, F., Liao, H., Zhang, S., & Nachtmann, H. (2023, December). System Simulation and Machine Learning-Based Maintenance Optimization for an Inland Waterway Transportation System. In Proceedings of the Winter Simulation Conference (pp. 267-278).
- [16] Sanaei, F. (2024). How customers' satisfaction change with the use of AR shopping application: A conceptual model. arXiv preprint arXiv:2401.10953. <https://doi.org/10.48550/arXiv.2401.10953>
- [17] Hashemian, F., Maleki, N., & Zeinali, Y. (2024). From User Behavior to Subscription Sales: An Insight Into E-Book Platform Leveraging Customer Segmentation and A/B Testing. *Services Marketing Quarterly*, 1-29.
- [18] Zandi, S. (2023). Revival of the Silk Road using the applications of AR/VR and its role on cultural tourism. arXiv preprint arXiv:2304.10545. <https://doi.org/10.48550/arXiv.2304.10545>
- [19] Mohsenibeigzadeh, M., Tashakkori, A., Kazemi, B., Moghaddam, P. K., & Ahmadi, Z. (2024). Driving Innovation in Education: The Role of Transformational Leadership and Knowledge Sharing Strategies. *Current Opinion*, 4(2), 505-515.
- [20] Emami, F., & Kabir, M. Z. (2023). Strength prediction of composite metal deck slabs under free drop weight impact loading using numerical approach and data set machine learning. *Scientia Iranica*, 1-35. doi: 10.24200/sci.2023.59582.6321
- [21] Abbasihafshejani, M., Manshaei, M. H., & Jadliwala, M. (2023, November). Detecting and Punishing Selfish Behavior During Gossiping in Algorand Blockchain. In 2023 IEEE Virtual Conference on Communications (VCC) (pp. 49-55). IEEE. <https://doi.org/10.1109/VCC60689.2023.10474784>
- [22] Soltaninejad, M., Aghazadeh, R., Shaghaghi, S., & Zarei, M. (2024). Using Machine Learning Techniques to Forecast Mehram Company's Sales: A Case Study. *Journal of Business and Management Studies*, 6(2), 42–53. <https://doi.org/10.32996/jbms.2024.6.2.4>
- [23] Abdali, A., & Monjezi, V. (2023). Machine Learning-based Flexible Link Robot Control. *Computational methods in engineering sciences*, 1(2), 12-17. doi: 10.22034/cmesc.2023.2008417.1011
- [24] Khajehzadeh, Mohammad, et al. "Analysis of Factors Affecting Product Sales with an Outlook toward Sale Forecasting in Cosmetic Industry using Statistical Methods." *International Review of Management and Marketing* 12.6 (2022): 55.
- [25] M. Tavasoli, E. Lee, Y. Mousavi, H. B. Pasandi and A. Fekih, "WIPE: A Novel Web-Based Intelligent Packaging Evaluation via Machine Learning and Association Mining," in *IEEE Access*, vol. 12, pp. 45936-45947, 2024, doi: 10.1109/ACCESS.2024.3376478.
- [26] A. Mehrzad, M. Darmiani, Y. Mousavi, M. Shafie-Khah and M. Aghamohammadi, "A Review on Data-Driven Security Assessment of Power Systems: Trends and Applications of Artificial Intelligence," in *IEEE Access*, vol. 11, pp. 78671-78685, 2023, doi: 10.1109/ACCESS.2023.3299208.
- [27] Mohammadabadi, S. M. S., Yang, L., Yan, F., & Zhang, J. Communication-Efficient Training Workload Balancing for Decentralized Multi-Agent Learning.
- [28] Mohammadabadi, S. M. S., Yang, L., Yan, F., & Zhang, J. (2024). Communication-Efficient Training Workload Balancing for Decentralized Multi-Agent Learning. arXiv preprint arXiv:2405.00839.
- [29] Graves, A., & Graves, A. (2012). Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37-45.
- [30] Long, W.; Lu, Z.; Cui, L. Deep learning-based feature engineering for stock price movement prediction. *Knowl.-Based Syst.* 2019, 164, 163–173
- [31] Pang, X.; Zhou, Y.; Wang, P.; Lin, W.; Chang, V. An innovative neural network approach for stock market prediction. *J. Supercomput.* 2018, 11, 1–21.
- [32] Kelotra, A.; Pandey, P. Stock Market Prediction Using Optimized Deep-ConvLSTM Model. *Big Data* 2020, 8, 5–24.

- [33] Bouktif, S.; Fiaz, A.; Awad, M. Augmented Textual Features-Based Stock Market Prediction. *IEEE Access* 2020, 8, 40269–40282
- [34] Rekha, G.; D Sravanthi, B.; Ramasubbareddy, S.; Govinda, K. Prediction of Stock Market Using Neural Network Strategies. *J. Comput. Theor. Nanosci.* 2019, 16, 2333–2336
- [35] Lee, J.; Kim, R.; Koh, Y.; Kang, J. Global stock market prediction based on stock chart images using deep Q-network. *IEEE Access* 2019, 7, 167260–167277.
- [36] Liu, G.; Wang, X. A numerical-based attention method for stock market prediction with dual information. *IEEE Access* 2018, 7, 7357–7367
- [37] Baek, Y.; Kim, H.Y. ModAugNet: A new forecasting framework for stock market index value with an overfitting prevention LSTM module and a prediction LSTM module. *Expert Syst. Appl.* 2018, 113, 457–480.
- [38] Chung, H.; Shin, K.-S. Genetic algorithm-optimized multi-channel convolutional neural network for stock market prediction. *Neural Comput. Appl.* 2020, 2, 7897–7914.
- [39] Zhou, X.; Pan, Z.; Hu, G.; Tang, S.; Zhao, C. Stock market prediction on high-frequency data using generative adversarial nets. *Math. Probl. Eng.* 2018, 2018.
- [40] <https://www.kaggle.com/datasets/bryanb/cac40-stocks-dataset/data>
- [41] Boozary, P. (2024). The Impact of Marketing Automation on Consumer Buying Behavior in the Digital Space Via Artificial Intelligence. *Power System Technology*, 48(1), 1008-1021.
- [42] Kiaghadi, M., Sheikholeslami, M., Alinia, A. M., & Boora, F. M. (2024). Predicting the performance of a photovoltaic unit via machine learning methods in the existence of finned thermal storage unit. *Journal of Energy Storage*, 90, 111766.
- [43] Owrang O., M. M., Schwarz, G., & Horestani, F. J. (2025). Prediction of Breast Cancer Recurrence With Machine Learning. In M. Khosrow-Pour, D.B.A. (Ed.), *Encyclopedia of Information Science and Technology*, Sixth Edition. Advance online publication. <https://doi.org/10.4018/978-1-6684-7366-5.ch061>
- [44] Nasiri, E., & Wang, L. (2024). Admittance Control for Adaptive Remote Center of Motion in Robotic Laparoscopic Surgery. *arXiv preprint arXiv:2404.04416*. <https://doi.org/10.48550/arXiv.2404.04416>
- [45] Mehrabian, A., Bahrami, S., & Wong, V. W. (2023, May). A dynamic Bernstein graph recurrent network for wireless cellular traffic prediction. In *ICC 2023-IEEE International Conference on Communications* (pp. 3842-3847). IEEE. <https://doi.org/10.1109/ICC45041.2023.10279102>
- [46] Unal, H. B., Zeynali, S., Anttila, E., Roll, J., Kreutz, R., Frick, K., ... & Sharif, B. (2024). Feasibility of Mri-guided Left Heart Catheterization on a Commercially Available 0.55 T Scanner Platform and Readily Available Invasive Pressure Monitoring Hardware. *Journal of Cardiovascular Magnetic Resonance*, 26. <https://doi.org/10.1016/j.jocmr.2024.100196>
- [47] Unal, H. B., Zeynali, S., Anttila, E., Roll, J., Kreutz, R., Frick, K., ... & Sharif, B. (2024). Rapid Mapping of Stress/rest Myocardial T1 Reactivity with Brief Intracardiac Infusion of Adenosine Enabled by a Commercially Available Icmr Paradigm. *Journal of Cardiovascular Magnetic Resonance*, 26. <https://doi.org/10.1016/j.jocmr.2024.100280>
- [48] Ghafariasl, P., Mahmoudan, A., Mohammadi, M., Nazarpour, A., Hoseinzadeh, S., Fathali, M., ... & Garcia, D. A. (2024). Neural network-based surrogate modeling and optimization of a multigeneration system. *Applied Energy*, 364, 123130. <https://doi.org/10.1016/j.apenergy.2024.123130>
- [49] Talebzadeh, H., Talebzadeh, M., Satarpour, M., Jalali, F., Farhadi, B., & Vahdatpour, M. S. (2024). Enhancing breast cancer diagnosis accuracy through genetic algorithm-optimized multilayer perceptron. *Multiscale and Multidisciplinary Modeling, Experiments and Design*, 1-17.