

Option Pricing Using Machine Learning Models: A Comparative Study with the Black-Scholes Model

1. Ali. Abed Lazim^{ID}: Ph.D Student in Finance – Financial Engineering, Aras International Campus, University of Tehran, Jolfa, Iran
2. Ali. Sour^{ID}: Associate Professor of Theoretical Economics, Faculty of Economics, University of Tehran, Tehran, Iran
3. Ezatollah. Abbasian^{ID*}: Professor of Financial Economics, Department of Financial Engineering, Faculty of Management, University of Tehran, Tehran, Iran
4. Eslam. Fakher^{ID}: Assistant Professor of Shahid Chamran University of Ahvaz, Ahvaz, Iran

*corresponding author's email: e.abbasian@ut.ac.ir

ABSTRACT

The capital market has always been accompanied by significant fluctuations, influenced by a variety of factors. These conditions expose market participants to different types of risks, thereby increasing the necessity of employing risk management methods. One of the widely used instruments in this context is the option contract. However, one of the main challenges related to options is their pricing. Among the most well-known models for option pricing is the Black-Scholes model, which is extensively used in practice. Nevertheless, due to its simplifying assumptions, its outcomes do not always align with real market conditions. Machine learning models, owing to their high flexibility and ability to analyze complex datasets, hold strong potential for improving the accuracy of option pricing. This study focuses on pricing options using machine learning algorithms. In this research, 181 option contracts traded on the Tehran Stock Exchange from April 2018 to October 2024 were examined. After data collection and standardization, the option prices were predicted using both machine learning algorithms and the Black-Scholes model. The predicted prices were then compared with actual market prices using Mean Absolute Error and Root Mean Squared Error as evaluation metrics. Additionally, the models' performance was analyzed based on contract maturity and moneyness. The results indicate that machine learning algorithms outperformed the Black-Scholes model both overall and within the maturity and moneyness subcategories.

Keywords: Machine learning models – Stock option – Black-Scholes model – LSTM

Introduction

Financial markets operate under conditions of continuous uncertainty in which asset prices respond to changes in macroeconomic variables, firm-specific information, investor expectations, liquidity conditions, geopolitical events, and shifts in market sentiment. These fluctuations expose investors, portfolio managers, financial institutions, and other market participants to substantial price risk and create an enduring need for instruments capable of transferring, managing, or hedging such risk. Derivative securities have therefore become an important component of modern financial markets, with option contracts providing investors with the right, but not the obligation, to buy or sell an underlying asset at a predetermined price within or at the end of a specified period. In addition to their role in hedging, options facilitate speculation, portfolio construction, volatility trading, and the



Article history:
Received 01 May 2026
Revised 16 September 2026
Accepted 21 September 2026
Initial Publish 24 October 2026
Published online 01 July 2027

How to cite this article:

Abed Lazim, A., Sour, A., Abbasian, E., & Fakher, E. (2027). Option Pricing Using Machine Learning Models: A Comparative Study with the Black-Scholes Model. *Journal of Management and Business Solutions*, 5(7), 1-18.
<https://doi.org/10.61838/jmbs.465>



incorporation of market expectations into financial decisions. The usefulness of these instruments, however, depends fundamentally on determining an appropriate theoretical price. A pricing model that systematically overvalues or undervalues contracts may result in inappropriate investment decisions, inaccurate risk measurement, inefficient hedging strategies, and distorted assessments of market opportunities. Consequently, option pricing remains one of the central problems of quantitative finance. Contemporary research continues to investigate the performance and applicability of classical option-pricing approaches, particularly the Black-Scholes framework, because of its continuing importance as a theoretical and practical benchmark (1). At the same time, modified lattice approaches continue to be developed for contracts whose characteristics cannot be adequately represented by elementary analytical solutions; for example, modified bino-trinomial tree structures have recently been investigated for barrier-option valuation and loss-minimization strategies (2). These developments illustrate that although option-pricing theory is mature, the search for more accurate and adaptable valuation techniques remains highly relevant.

The Black-Scholes model represents one of the most influential developments in financial economics because it provides a mathematically tractable framework for deriving the theoretical value of European options from a relatively small number of observable or estimable inputs. Under its conventional specification, option value is determined principally by the current price of the underlying asset, strike price, time remaining to maturity, volatility, and the risk-free interest rate. Mathematical derivations of the model demonstrate how assumptions concerning stochastic price behavior, continuous trading, and risk-neutral valuation lead to a closed-form solution for option prices (3). The model has subsequently become a benchmark against which both conventional extensions and data-driven valuation techniques can be evaluated. Empirical investigations have nevertheless shown that its pricing efficiency can vary across markets and contract characteristics (4). Iranian evidence similarly demonstrates the continuing importance of comparing Black-Scholes with alternative models. Amiri examined option pricing under Black-Scholes, Boness, and binomial-tree specifications using gold coin option contracts in the Iran Mercantile Exchange, emphasizing that valuation outcomes depend on both modeling assumptions and the characteristics of the market being studied (5). Abvali and colleagues likewise proposed a new analytical method in relation to the Black-Scholes equation, demonstrating continued interest in improving or reformulating classical option-pricing procedures (6). More recently, Mahdavi directly investigated the effectiveness of machine-learning models relative to Black-Scholes for option pricing, reflecting a broader transition from exclusively theory-driven valuation toward empirical and algorithmic approaches (7).

Despite its foundational importance, the Black-Scholes framework relies on assumptions that may depart considerably from actual market behavior. Standard implementations generally presume frictionless trading, continuous asset-price dynamics, stable or constant volatility, an adequately specified risk-free rate, and distributional properties that simplify the behavior of underlying returns. Actual financial markets, in contrast, may exhibit volatility clustering, nonlinear price adjustments, jumps, changing liquidity, asymmetric responses to information, heterogeneous investor behavior, transaction costs, and structural changes over time. These departures can produce persistent differences between theoretical and observed market prices. Even relatively technical aspects of Black-Scholes implementation, such as discretization and hedging error, remain active areas of investigation; Rakshit and colleagues examined hedge errors through an asymptotic finite-difference perspective, illustrating that practical implementation raises issues beyond the conceptual elegance of the continuous-time model (8). The challenge is especially consequential when the pricing environment differs from highly liquid

developed markets. Market microstructure, depth, trading frequency, available maturities, investor composition, and institutional conditions can all influence observed option prices. Therefore, the empirical adequacy of a pricing framework cannot simply be assumed to transfer unchanged across national markets. In the Iranian context, where exchange-traded options are comparatively recent, these considerations are particularly relevant. The study underlying the present article examines 181 option contracts traded on the Tehran Stock Exchange between April 2018 and October 2024 and evaluates pricing accuracy not only overall but also according to maturity and moneyness, thereby permitting the performance of competing models to be assessed under different contract conditions.

The growing availability of financial data and advances in computational capacity have created an alternative approach to this problem through artificial intelligence and machine learning. Artificial intelligence is increasingly regarded not merely as a collection of application-specific tools but as a powerful methodological paradigm capable of transforming scientific discovery, modeling, prediction, and optimization (9). Bianchini and colleagues similarly described artificial intelligence as an emerging general method of invention, highlighting its capacity to reshape the production and application of knowledge across multiple scientific domains (10). Within finance, this transformation is particularly important because financial datasets frequently contain nonlinear dependencies, high-dimensional interactions, temporal structures, and patterns that are difficult to specify *ex ante* in conventional parametric equations. Machine-learning models differ from traditional analytical pricing models in an important conceptual respect: rather than requiring the researcher to fully specify the functional relationship among explanatory variables and option value in advance, they can estimate such relationships from observed data. This characteristic creates the possibility of capturing complex market behavior without imposing all of the restrictive assumptions associated with classical models. However, greater flexibility also introduces challenges related to overfitting, model selection, transparency, and interpretation. Miller's work on explanation in artificial intelligence emphasizes that predictive performance alone is not always sufficient and that understanding how algorithmic systems produce their outputs can be important for their practical acceptance and meaningful use (11). This concern is relevant to financial applications in which pricing outputs may influence investment, hedging, and risk-management decisions.

The suitability of machine learning for financial applications is supported by a growing body of work on neural-network and deep-learning methods. Alipour and Parsinejad designed neural-network models to improve stock portfolio prediction, demonstrating the usefulness of artificial neural structures for extracting relationships from financial data that may not be adequately represented using conventional linear techniques (12). Karimi Dastgerdi and Zamani Boroujeni reviewed deep-learning approaches to financial-market prediction and identified recurrent neural networks, long short-term memory networks, convolutional neural networks, and hybrid architectures as important methodological developments in financial forecasting (13). One reason for the appeal of these techniques is their ability to model nonlinear functions and, in recurrent architectures, dependencies extending across sequential observations. Financial time series frequently contain temporal relationships in which current prices reflect not only contemporary explanatory variables but also patterns embedded in previous observations. Long Short-Term Memory (LSTM) networks are specifically designed to address this feature by retaining or discarding information through gating mechanisms and thereby reducing the vanishing-gradient problems that constrain conventional recurrent neural networks. Li and colleagues proposed an evolutionary attention-based LSTM architecture for time-series prediction, demonstrating how attention and evolutionary mechanisms may further strengthen the capacity of LSTM-based systems to identify informative temporal patterns (14). The applicability of

LSTM is not limited to finance; bidirectional LSTM architectures have also demonstrated their capacity to capture complicated temporal relationships in environmental forecasting applications (15). Similarly, LSTM-based models have been explored in classification settings, although the article by Singh and colleagues listed in the present reference set has subsequently been marked as retracted and should therefore not be interpreted as reliable substantive evidence for model effectiveness (16). Taken together, the broader methodological literature nonetheless provides a strong rationale for examining recurrent deep-learning architectures where prediction depends on nonlinear and sequential information.

The quality of financial information is another critical consideration because sophisticated algorithms cannot completely compensate for data that are noisy, sparse, nonrepresentative, or structurally unstable. Financial prices are generated through continuous interactions among informed and uninformed trading, liquidity shocks, behavioral reactions, market-wide events, and measurement conventions. Consequently, observed financial series often contain substantial noise. Mohammadi and Mansourfar demonstrated that financial data noise can alter assessments of long-term stock-market co-movement, showing that the informational structure of financial series can materially affect empirical conclusions (17). This issue is particularly important in machine-learning applications because algorithms may learn both economically meaningful patterns and idiosyncratic noise. Appropriate preprocessing, filtering, standardization, training procedures, and out-of-sample validation are therefore essential to determining whether lower estimation error reflects genuine predictive capacity rather than sample-specific fitting. At the same time, different algorithms manage the bias-variance tradeoff in different ways. Decision-tree models, for example, are comparatively interpretable and can represent nonlinear threshold relationships, but unrestricted trees are vulnerable to overfitting. Ibrahim discusses the evolution from individual decision trees toward bagging and Random Forest approaches specifically from a bias-variance perspective, illustrating why ensemble strategies can improve generalization (18). Kappelhof and colleagues, although working in a medical prediction setting, further demonstrate the practical use of decision trees and evolutionary algorithms for predictive classification and provide methodological evidence concerning the use of tree-based structures in complex empirical data environments (19). Such evidence is relevant methodologically because option pricing likewise requires algorithms capable of learning nonlinear interactions while maintaining acceptable out-of-sample performance.

Support-vector methods provide another important machine-learning framework for nonlinear financial prediction. Support Vector Machines are grounded in statistical learning theory and attempt to determine a decision boundary with favorable generalization properties by maximizing the margin defined by influential observations. For continuous prediction problems, Support Vector Regression extends the same principles to function estimation. With kernel functions, support-vector models can map input variables into higher-dimensional feature spaces and represent nonlinear relationships without requiring those relationships to be explicitly parameterized in their original form. Piccialli and Sciandrone examined the relationship between nonlinear optimization and support-vector machines, highlighting the mathematical optimization foundations underlying SVM estimation (20). Yeh and colleagues developed a support-vector-machine formulation based on a convolution product, further demonstrating how the architecture of support-vector approaches can be modified to address complex patterns (21). These properties make SVR relevant to option valuation, where interactions among underlying price, strike price, maturity, volatility, and interest rates may generate nonlinear pricing surfaces. Compared with deep neural networks, support-vector approaches may also offer advantages when available datasets are more limited, although their effectiveness

depends substantially on kernel choice, hyperparameters, scaling, and the distributional characteristics of the training sample.

Direct evidence from derivative markets increasingly indicates that machine-learning models can provide competitive or superior pricing performance relative to classical analytical methods. Ivaşcu investigated option pricing using machine-learning techniques and found that data-driven approaches could improve pricing accuracy relative to Black-Scholes under the empirical conditions examined (22). Jayaraman and colleagues evaluated several data-driven approaches to option pricing, including Decision Trees, neural networks, Random Forests, and Gradient Boosting, demonstrating the potential for algorithmic methods to reproduce observed option prices more accurately than traditional benchmark specifications in some settings (23). Liu and Zhang focused specifically on LSTM-based option pricing from the perspective of realized skewness and compared recurrent and other machine-learning approaches, providing evidence that LSTM architectures can be particularly effective in learning option-pricing relationships (24). Fan and Sirignano more recently investigated machine-learning methods for pricing financial derivatives and further demonstrated the capacity of neural and hybrid approaches to model derivative prices using market information (25). Collectively, these studies suggest that machine learning should not simply be understood as an auxiliary forecasting technology; rather, it represents a fundamentally different way of estimating the pricing map between market state variables and derivative values. Nevertheless, superiority cannot be presumed universally because algorithmic performance can vary with sample size, market structure, liquidity, training design, contract characteristics, and evaluation metrics.

An important unresolved question is therefore whether machine-learning models maintain their relative advantages across different characteristics of option contracts. Aggregate error measures can conceal considerable heterogeneity in pricing performance. Moneyness, for example, changes the relationship between the underlying asset price and strike price and influences the composition of intrinsic and time value. In-the-money and out-of-the-money options can therefore exhibit different pricing behavior and different sensitivities to model misspecification. Maturity creates another important source of heterogeneity because the remaining life of a contract affects uncertainty concerning the terminal asset price, sensitivity to volatility, and the contribution of time value. A model that performs well for longer-term contracts may not necessarily retain the same accuracy near expiration, where price behavior can become particularly sensitive to movements in the underlying asset. These considerations provide a rationale for complementing overall accuracy comparisons with subgroup analyses. They are also important from a practical perspective because investors and risk managers require valuation models that are reliable for the particular contracts they trade rather than merely accurate on average. The continuing development of Black-Scholes-based research (1), lattice-based modifications (2), investigations of numerical and hedging errors (8), and machine-learning comparisons (7) indicates that contemporary option-pricing research is increasingly concerned not with a single universally optimal model but with determining how different frameworks perform under varying empirical circumstances.

This issue is especially significant for the Iranian capital market. Much of the international machine-learning option-pricing literature has been constructed using comparatively mature markets characterized by longer derivative-market histories, deeper liquidity, and substantially larger datasets. Results obtained under those institutional conditions may not automatically generalize to a younger option market. Evidence from Iranian financial studies has demonstrated the applicability of neural networks to market prediction (12), examined alternative option-pricing formulations (6), compared traditional valuation models in domestic derivative markets (5), and more recently

directly considered machine learning relative to Black-Scholes (7). Nevertheless, there remains value in a multi-model empirical comparison that simultaneously evaluates conceptually distinct approaches—tree-based learning, support-vector regression, recurrent deep learning, and Black-Scholes—using the same option-market observations and identical performance criteria. Evaluating these models with Mean Absolute Error and Root Mean Squared Error provides complementary information because MAE summarizes the magnitude of absolute prediction deviations, whereas RMSE places greater weight on larger errors. Moreover, examining performance by both maturity and moneyness can reveal whether aggregate differences between classical and machine-learning approaches persist across economically important option categories. Such evidence can contribute to the literature on computational finance while providing market-specific information relevant to investors, analysts, portfolio managers, and institutions operating in the Tehran Stock Exchange.

Accordingly, the aim of this study was to evaluate and compare the accuracy of Decision Tree, Support Vector Regression, Long Short-Term Memory, and Black-Scholes models in pricing options traded on the Tehran Stock Exchange and to determine their relative predictive performance overall and across different levels of contract maturity and option moneyness.

Methods and Materials

The dataset used in this study consists of option contracts traded on the Tehran Stock Exchange between April 2018 and October 2024. Given the importance of data quality, a two-stage filtering process was applied to ensure that only contracts with sufficient trading activity and reliability were selected for training and evaluating the models.

In the first stage, the following two filters were applied:

1. Each option contract must have at least 50 trading days.
2. The number of trading days for each underlying symbol must be at least 70% of the total trading days in the full period.

To further enhance the robustness of the model, two additional filters were implemented:

1. Removal of records with a market price below 10 IRR.
2. Removal of contracts with a remaining maturity of fewer than 5 days, as option prices tend to exhibit higher volatility in the final days before expiration.

After applying the above filters, a total of 13,077 records were obtained from 181 option contracts across 14 different underlying symbols.

Finally, model evaluation was carried out using an out-of-sample validation approach, where 80% of the dataset was used for training and the remaining 20% was reserved for performance evaluation.

In 1973, Fischer Black and Myron Scholes sparked a revolution in the world of finance by introducing their renowned option pricing model. Built on sophisticated mathematical and economic concepts, the Black-Scholes model offered an efficient and practical method for valuing European-style options. Although the model was developed under idealized conditions and specific assumptions, it continues to serve as the foundation for many advanced models and remains a fundamental tool in financial education and research (Kaushik, 2017).

The Black-Scholes model determines the fair value of an option based on five key variables: the current price of the underlying asset, the strike price, time to maturity, price volatility, and the risk-free interest rate (Rakshit et al.,

2024). The following equations demonstrate how, by combining these elements and applying the mathematical structure of the model, the theoretical price of an option can be computed (Majumdar et al., 2018):

$$C = S_t \cdot N(d_1) - K \cdot e^{(-rT)} \cdot N(d_2) \quad (1)$$

$$P = K \cdot e^{(-rT)} \cdot N(-d_2) - S_0 \cdot N(-d_1) \quad (2)$$

The model uses the following standard notations:

C: Call option price

P: Put option price

S_t : Current price of the underlying

K: Strike price of the option

asset

T: Time remaining until expiration

r: Risk-free interest rate

σ : Historical volatility of the asset's returns

Two crucial parameters in the Black-Scholes model, d_1 and d_2 , play a significant role in the option valuation process. These parameters represent the likelihood of two key events:

d_1 : Indicates the probability that the asset price at maturity will exceed the strike price. In other words, it reflects the likelihood of the call option being profitable.

d_2 : Represents the present value of the strike price at maturity, essentially showing how much the exercise price is worth in the future.

Together with the standard normal distribution function, these two parameters are embedded in the model's complex formulas to estimate the precise value of an option. In essence, d_1 and d_2 assist in estimating both the probability of profitability and the expected future value of the option (Majumdar et al., 2018).

Decision Tree (DT). The Decision Tree algorithm was first introduced in the 1960s. Primarily used for classification and prediction tasks, this algorithm builds tree-like models to make decisions. It utilizes concepts such as information gain and entropy to select the best feature for splitting the data. Due to its simplicity and high interpretability, the decision tree has been widely applied in various fields, including machine learning and data mining (Ibrahim, 2022). Despite its simplicity, decision trees are prone to overfitting, especially when they become too complex and tailored to the training data. To mitigate this issue, techniques such as tree pruning or restricting the maximum depth of the tree can be employed (Kappelhof et al., 2021).

Support Vector Machine (SVM) is a machine learning algorithm initially introduced in the 1960s and later developed in the 1990s by Vladimir Vapnik and colleagues. Rooted in statistical learning theory, SVM quickly became one of the most popular methods for solving classification problems, particularly in high-dimensional and non-linear domains (Piccialli & Sciandrone, 2022). The primary objective of SVM is to separate data points using an optimal hyperplane. This hyperplane is selected in a way that maximizes the margin between the closest points of each class, known as support vectors. In cases where data are not linearly separable, the algorithm employs the kernel trick—a technique that maps data into a higher-dimensional space where a linear separation becomes feasible. This ability to capture complex, non-linear relationships has made SVM a powerful tool for a wide range of classification tasks across various domains (Yeh et al., 2021).

Long Short-Term Memory (LSTM) was introduced in the 1990s. The core idea behind its development emerged from the limitation of traditional neural networks in capturing long-term dependencies within sequential data. LSTM is a type of Recurrent Neural Network (RNN) specifically designed to learn long-range temporal relationships. It overcomes the vanishing gradient problem by incorporating a gating mechanism that regulates the flow of information through the network (Li et al., 2019). Initially, LSTM received limited attention. However, with advancements in computing hardware and the emergence of new applications—particularly in fields such as natural language processing and speech recognition—its popularity grew significantly. By the 2010s, with the rise of modern deep learning frameworks, LSTM had become one of the most essential tools for analyzing sequential and time-series data (Singh et al., 2022).

This section outlines the methodology used to evaluate and compare the predictive performance of the models. The analysis includes the definition of key grouping criteria—namely option moneyness and contract maturity—as well as the metrics used to measure model accuracy.

Option Moneyness: Options are classified based on their moneyness into three categories: in-the-money (ITM), at-the-money (ATM), and out-of-the-money (OTM). However, due to the limited number of ATM options in the dataset, this group was excluded from the analysis. As a result, the evaluation was carried out on two categories only: ITM and OTM options.

Contract Maturity: Based on the structure of the TSE, option contracts are typically issued with maturities of either 6 months or 12 months. In this study, contracts are divided into two groups accordingly: short-term contracts (6 months) and long-term contracts (12 months).

For the Black-Scholes model, theoretical prices are first calculated using the model's formula and relevant input variables. These theoretical prices are then compared to actual market prices. In machine learning models, the input features are fed into each algorithm to predict option prices, which are similarly compared against actual market prices.

To evaluate and compare the accuracy of the pricing models, common error metrics are used, including Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). These metrics allow for a precise assessment of each model's predictive performance across different categories of moneyness and maturity.

Findings and Results

The collected parameters include the underlying asset price, time to maturity, risk-free interest rate, strike price, and the market price of the options contracts. For these parameters, the following statistical indicators are measured: mean, standard deviation, minimum, first quartile, median, third quartile, and maximum.

Table 1. Descriptive Statistics

Parameters	Mean	SD	Minimum	First quartile	Median	Third quartile	Maximum
Underlying asset price	3718.2	3158.3	799	1915.75	2482	3541.5	15610
Time to maturity	0.457	0.255	0.07	0.24	0.45	0.66	0.95
Risk-free interest rate	0.284	0.035	0.21	0.27	0.27	0.30	0.34
Strike price	3021.1	2410.8	1000	1600	2200	3000	10800
options market price	987.02	1348.4	10	175	450	1058	8977

As shown in the table, the minimum and maximum underlying asset prices for the options contracts are 799 Rials and 15,610 Rials, respectively. Additionally, the strike prices for these options range from 1,000 Rials to 10,800 Rials. The remaining time to maturity varies between 0.07 and 0.95 years. The risk-free interest rate fluctuates

between 0.21 and 0.34, and the market prices of the options contracts fall within the range of 10 Rials to 8,977 Rials.

In this section, the predicted prices from each model are compared with the actual market prices of the options. Following this comparison, the prediction errors of each model are examined to assess their accuracy in estimating option prices. Error metrics such as RMSE and MAE are used to evaluate the performance of the models. The following chart visually presents the prediction error of each model in forecasting option prices.




Figure 1. Comparison of the prediction accuracy of each model based on statistical metrics: MAE and RMSE

Based on the presented data, the results indicate that machine learning algorithms outperformed the Black-Scholes model in predicting option prices. Among them, the Long Short-Term Memory (LSTM) algorithm demonstrated the best performance, achieving a Mean Absolute Error (MAE) of 0.1084 and a Root Mean Squared Error (RMSE) of 0.0502. To better illustrate the prediction errors of each model, the following scatter plots compare the predicted option prices generated by the LSTM and Black-Scholes models against the actual market prices.




Figure 2. Scatter Plot of Actual vs. Predicted Option Prices for LSTM and Black-Scholes Models



As observed in the scatter plots, the distribution of predicted versus actual prices in the LSTM model aligns more closely with the diagonal line, indicating a stronger linear relationship. This suggests that the LSTM algorithm provides more accurate predictions compared to the Black-Scholes model.

This section analyzes the predictive performance of machine learning algorithms and the Black-Scholes model in estimating the prices of options contracts with different maturities. Options are categorized into two groups based on their contract duration: short-term and long-term. In the context of the Iranian financial market, these correspond to 6-month and 12-month contracts, respectively. The predictive accuracy of both machine learning models and the Black-Scholes model has been evaluated separately for each group to examine how contract maturity affects model performance.

Table 2. Short-term vs. Long-term Contract Analysis

Maturity	Model	Mae	Rmse
Short-term	LSTM	0.0486	0.0992
	Decision Tree	0.0660	0.1263
	SVR	0.0793	0.1330
	Black-Scholes	0.1475	0.2418
Long-term	LSTM	0.0500	0.1047
	Decision Tree	0.0602	0.1289
	SVR	0.0633	0.1242
	Black-Scholes	0.1042	0.1908

In Figures 3 and 4, the prediction accuracy of machine learning algorithms and the Black-Scholes model in estimating option prices is compared based on different contract durations. This comparison illustrates how each method performs in predicting prices for short-term and long-term contracts and highlights the differences in their performance across contract periods.

**Figure 3. Model Performance on Short-Term Options****Figure 4. Model Performance on Long-Term Options**

Based on Table 2 and Figures 3 & 4, the performance of machine learning algorithms in predicting option prices has been superior to that of the Black-Scholes model. Additionally, another insight gained from the data indicates that both models predict long-term options more accurately than short-term ones.

This section analyzes the predictive performance of machine learning algorithms and the Black-Scholes model in estimating the prices of options contracts based on their moneyness. In financial terms, options are typically categorized into three groups based on moneyness: in-the-money, at-the-money, and out-of-the-money. However, due to the very limited number of at-the-money contracts in the dataset, this category has been excluded from the analysis. As a result, the study focuses only on in-the-money and out-of-the-money options. The accuracy of both the machine learning algorithms and the Black-Scholes model in predicting the prices of these two categories is evaluated and compared in the following analysis.

Table 3. In-the-money vs. Out-of-the-money options Analysis

Moneyiness	Model	Mae	Rmse
In-the-money	LSTM	0.0603	0.1049
	Decision Tree	0.0749	0.1368
	SVR	0.0755	0.1289
	Black-Scholes	0.1427	0.2297
Out-of-the-money	LSTM	0.1292	0.2530
	Decision Tree	0.1738	0.3527
	SVR	0.2605	0.4243
	Black-Scholes	0.4554	0.6964

In Figures 5 and 6, the prediction accuracy of machine learning algorithms and the Black-Scholes model is compared based on different levels of option moneyness. This comparison illustrates how each method performs in estimating the prices of in-the-money and out-of-the-money options, and it highlights the differences in their predictive performance across these categories.




Figure 5. Model Performance on In-the-money Options




Figure 6. Model Performance on Out-of-the-money Options

Based on Table 3 and Figures 5 and 6, machine learning algorithms have demonstrated better performance than the Black-Scholes model in predicting the prices of both in-the-money and out-of-the-money options. Furthermore, the results indicate that both approaches achieve higher prediction accuracy for in-the-money options compared to out-of-the-money options.

Discussion and Conclusion

The present study compared the predictive performance of three machine-learning approaches—Long Short-Term Memory (LSTM), Decision Tree (DT), and Support Vector Regression (SVR)—with the conventional Black-Scholes model for pricing option contracts traded on the Tehran Stock Exchange. The findings provide consistent evidence that the data-driven algorithms produced lower pricing errors than the Black-Scholes benchmark, both in the aggregate evaluation and when contracts were classified according to maturity and moneyness. Across the overall comparison, LSTM exhibited the strongest predictive performance, followed by Decision Tree and SVR, whereas Black-Scholes generated the largest prediction errors. The scatter-plot comparison also showed that LSTM-predicted prices were more closely concentrated around the diagonal representing equality between predicted and observed market prices, whereas Black-Scholes predictions displayed substantially greater dispersion. This result suggests that the flexible nonlinear structure of machine-learning models, and particularly the capacity of LSTM to learn complex sequential dependencies, is better suited to the empirical pricing structure observed in the Iranian option market than a model based on fixed theoretical assumptions.

The overall superiority of machine-learning approaches is consistent with the growing literature demonstrating that option prices can be estimated effectively through data-driven methods. Ivaşcu reported that machine-learning approaches yielded lower pricing errors than the Black-Scholes model when applied to crude-oil options, indicating that market-derived nonlinear relationships can be captured more effectively by flexible learning algorithms than by a predetermined analytical specification (22). Similarly, Jayaraman et al. compared several data-driven algorithms, including Decision Trees, neural networks, Random Forests, and Gradient Boosting, with conventional option-pricing approaches and found that machine-learning techniques could provide substantial improvements in pricing accuracy (23). Fan and Sirignano likewise demonstrated the usefulness of machine-learning methods for financial-derivative valuation, particularly where nonlinear pricing relationships cannot be adequately represented through a

single traditional analytical structure (25). More recent evidence reported by Mahdavi also supports the effectiveness of machine-learning models relative to the Black-Scholes benchmark in option pricing (7). Therefore, the present findings extend an increasingly consistent empirical pattern to the Tehran Stock Exchange and suggest that the advantages of machine learning are not limited to highly developed derivative markets.

The weaker performance of Black-Scholes can be interpreted in light of the assumptions on which the model is constructed. Although Black-Scholes remains one of the most important theoretical foundations of derivative valuation, it assumes a relatively simplified financial environment in which essential parameters such as volatility and the risk-free rate are treated within a restrictive mathematical structure. The model continues to be useful as a benchmark and provides an elegant theoretical relation among the underlying asset price, strike price, volatility, interest rate, and maturity (3). However, real markets may exhibit changing volatility, asymmetric price dynamics, non-normal returns, transaction frictions, discontinuities, and heterogeneous investor behavior. Empirical evaluations in different markets have therefore documented variation in the practical efficiency of Black-Scholes (4). In Iran, Amiri's comparison of Black-Scholes with Boness and binomial-tree models also illustrates that the choice of valuation framework can materially affect option-price estimates (5). Abvali et al. similarly attempted to improve option pricing through an alternative analytical treatment of the Black-Scholes equation, further reflecting the limitations involved in applying its conventional form without modification (6). The hedge-error analysis of Rakshit et al. additionally indicates that departures between theoretical and practical pricing outcomes can arise even when the Black-Scholes framework is considered from a more technical numerical perspective (8). Accordingly, the relatively large Black-Scholes errors obtained in the present study may reflect a mismatch between restrictive theoretical assumptions and the empirical dynamics of Tehran Stock Exchange option prices.

Among the machine-learning algorithms, LSTM showed the strongest performance. This finding is theoretically plausible because option-market observations are generated within a dynamic financial environment in which current values may contain information associated with preceding market conditions. LSTM was specifically developed to retain relevant information across sequences while controlling information flow through its gating mechanisms, making it particularly appropriate for time-dependent data. Li et al. demonstrated that attention-enhanced LSTM architectures can effectively model complex time-series relationships (14), while Karimi Dastgerdi and Zamani Boroujeni emphasized the usefulness of LSTM and related deep-learning architectures for financial-market forecasting because of their ability to capture nonlinear and temporal dependencies (13). Liu and Zhang provide even more directly comparable evidence: in their investigation of option pricing using realized skewness and several competing machine-learning methods, LSTM achieved highly favorable predictive performance relative to alternative models and Black-Scholes (24). Although Ghayoumi Zadeh et al. applied bidirectional LSTM outside finance, their findings similarly demonstrated the ability of recurrent deep-learning architectures to learn nonlinear temporal relationships from sequential data (15). Thus, the present finding that LSTM produced the most accurate option-price predictions is compatible with the broader methodological evidence concerning its capacity to retain temporally relevant information while modeling nonlinear interactions.

Decision Tree and SVR also outperformed Black-Scholes, although neither achieved the accuracy of LSTM. Decision trees are capable of generating nonlinear decision structures by recursively dividing the input space according to informative predictor thresholds. Their performance in the present study indicates that option pricing may involve conditional and nonlinear relationships that can be approximated effectively without imposing a prespecified functional form. Ibrahim's discussion of decision trees and the subsequent evolution toward Random

Forest methods highlights both the predictive utility and the bias-variance challenges of tree-based learning (18). Kappelhof et al. similarly demonstrated that decision-tree techniques can generate useful predictions in complex multivariable datasets, although careful control of model complexity is necessary to reduce overfitting (19). SVR, by contrast, relies on statistical-learning principles and kernel-based transformations to represent nonlinear relationships. Its favorable performance relative to Black-Scholes is therefore also theoretically reasonable. Piccialli and Sciandrone emphasized the close relationship between support-vector methods and nonlinear optimization (20), while Yeh et al. demonstrated how alternative SVM structures can further expand the algorithm's capacity to model complex patterns (21). Nevertheless, the stronger performance of LSTM suggests that explicitly learning temporal dependencies may provide an additional advantage beyond nonlinear approximation alone.

A second important result concerned contract maturity. For short-term options, LSTM recorded an MAE of 0.0486 and RMSE of 0.0992, followed by Decision Tree with 0.0660 and 0.1263 and SVR with 0.0793 and 0.1330, whereas Black-Scholes generated considerably higher errors of 0.1475 and 0.2418. For long-term contracts, the corresponding LSTM errors were 0.0500 and 0.1047; Decision Tree produced 0.0602 and 0.1289, SVR 0.0633 and 0.1242, and Black-Scholes 0.1042 and 0.1908. Thus, machine-learning algorithms retained their advantage within both maturity classes rather than achieving superior performance only through the composition of the overall sample. The article also concludes that model predictions were generally more accurate for long-term than short-term contracts. This pattern may indicate that shorter-duration contracts are more strongly influenced by rapidly changing market conditions as expiration approaches, whereas longer-maturity contracts provide a comparatively smoother pricing structure from which models can learn. Near maturity, relatively small movements in the underlying price can produce pronounced changes in option value, particularly around economically important strike regions, which can make accurate prediction more difficult.

The maturity findings can also be interpreted through the contrast between fixed-form pricing equations and adaptive learning systems. Classical option-pricing models explicitly include time to maturity, but the manner in which maturity interacts with other inputs is imposed by the theoretical model. Machine-learning algorithms instead estimate these interactions from empirical observations, allowing the effect of time to vary according to patterns in the training data. Recent studies continue to modify tree-based and analytical pricing structures precisely because different option characteristics can challenge standard valuation procedures. Rahayu et al., for example, investigated a modified bino-trinomial tree approach for barrier-option pricing as a strategy for reducing valuation losses (2), while Xiang's recent analysis of Black-Scholes reinforces the continued importance of evaluating how well its theoretical pricing structure corresponds to real option-market conditions (1). The present maturity-specific findings therefore suggest that performance comparisons should not rely solely on aggregate errors, since the relative accuracy of alternative approaches may depend on contractual characteristics.

A third major finding relates to moneyness. Machine-learning approaches outperformed Black-Scholes in both in-the-money (ITM) and out-of-the-money (OTM) categories. For ITM options, LSTM achieved MAE and RMSE values of 0.0603 and 0.1049, respectively; Decision Tree generated 0.0749 and 0.1368; SVR produced 0.0755 and 0.1289; and Black-Scholes produced substantially larger errors of 0.1427 and 0.2297. For OTM contracts, prediction became more difficult for every model: LSTM errors increased to 0.1292 and 0.2530, Decision Tree to 0.1738 and 0.3527, SVR to 0.2605 and 0.4243, and Black-Scholes to 0.4554 and 0.6964. Accordingly, two conclusions emerge. First, the superiority of machine learning was robust across moneyness categories. Second, all models estimated

ITM option prices more accurately than OTM prices. The latter result is especially important because it demonstrates that pricing difficulty is not uniform across the option surface.

The greater error observed for OTM options may reflect the comparatively complex valuation of contracts whose value depends strongly on the probability of future favorable price movement. Such options can have relatively low observed prices and may be particularly sensitive to volatility expectations, remaining maturity, market sentiment, and small changes in the underlying asset. These conditions may amplify relative pricing instability and make the underlying price relationship harder to estimate. By contrast, ITM options contain a more substantial intrinsic-value component, potentially providing a more stable relationship between the underlying asset and option price. The capacity of machine-learning models to accommodate nonlinear interactions appears to mitigate, but not eliminate, this difficulty. This interpretation is consistent with evidence that financial prediction involves substantial nonlinear structure and time variation (13). It is also compatible with Mohammadi and Mansourfar's observation that noise in financial data can materially influence estimated relationships among market variables (17). Because low-priced or less actively traded contracts can be especially susceptible to market microstructure noise, the markedly larger errors for OTM options underscore the importance of considering data quality when interpreting model performance.

More broadly, the results support the conceptual argument that artificial intelligence can serve as a general-purpose empirical methodology for identifying relationships that are difficult to specify analytically. Xu et al. described artificial intelligence as an increasingly powerful paradigm for scientific research (9), while Bianchini et al. characterized AI as an emerging general method of invention capable of reshaping scientific problem solving (10). In financial valuation, this advantage is evident in the ability of learning algorithms to derive mappings between market variables and observed option prices without requiring all relationships to follow predetermined analytical assumptions. Nevertheless, predictive superiority should not obscure the importance of model transparency. Miller emphasized that explanation and interpretability are central concerns in artificial intelligence, particularly where algorithmic decisions have consequential practical applications (11). In option markets, this implies that financial institutions should not select a model solely because it minimizes error; stability, interpretability, robustness, computational requirements, and sensitivity to changing market regimes should also be considered. The present study therefore supports machine learning as a strong complement or alternative to conventional pricing methods while retaining Black-Scholes as a useful theoretical benchmark rather than rendering classical financial theory irrelevant.

Overall, the convergence of the findings across the aggregate analysis, maturity groups, and moneyness categories strengthens the conclusion that machine-learning methods are well suited to option pricing in the Tehran Stock Exchange. LSTM was particularly effective, suggesting that combining nonlinear approximation with sequential learning provides an important advantage in modeling option-market behavior. The results also reveal, however, that pricing accuracy remains conditional on the characteristics of the option contract: OTM contracts were more difficult to value accurately than ITM contracts, and performance differed according to maturity. Thus, the principal contribution of the study is not merely the demonstration that machine learning can outperform Black-Scholes, but also the evidence that this advantage persists across economically meaningful contract classifications while the absolute level of prediction difficulty continues to vary across those classifications.

The study has several limitations that should be considered when interpreting the findings. First, the analysis was restricted to 181 option contracts traded on the Tehran Stock Exchange between April 2018 and October 2024, and the characteristics of this comparatively young derivatives market may limit generalization to other financial

markets. Second, the dataset was subject to filtering requirements concerning trading history, underlying-symbol activity, minimum market price, and remaining maturity; although these procedures improved data quality, they also excluded some contracts and may have reduced representation of illiquid or near-expiration options. Third, the analysis examined only three machine-learning approaches and did not include ensemble methods, boosting algorithms, transformer architectures, or hybrid financial-machine-learning models. Fourth, the study assessed predictive performance mainly through MAE and RMSE; additional measures could provide further insight into economically significant pricing errors. Fifth, at-the-money contracts were excluded because their number was insufficient for meaningful analysis, preventing a complete three-category comparison of moneyness. Finally, the study focused primarily on predictive accuracy and did not explicitly evaluate model interpretability, computational cost, hedging effectiveness, transaction costs, or performance under distinct volatility and market-regime conditions.

Future studies should expand the dataset across longer periods and, where possible, incorporate data from additional derivative markets to determine whether the observed performance patterns remain stable across institutional and market environments. Researchers could compare LSTM, Decision Tree, and SVR with Random Forest, XGBoost, LightGBM, recurrent-convolutional hybrid networks, attention-based neural networks, transformers, and ensemble architectures. Hybrid approaches that integrate theoretically derived option-pricing information with machine-learning predictions may also be particularly valuable. Future analyses should separately investigate high- and low-volatility periods, bullish and bearish regimes, liquidity conditions, and crisis episodes to determine the stability of model rankings under structural changes. Larger samples should be used to permit detailed analysis of at-the-money contracts and finer subdivisions of maturity and moneyness. Researchers should also examine alternative volatility measures, implied volatility, trading volume, bid-ask spreads, liquidity, underlying-return characteristics, and macroeconomic variables as predictive inputs. In addition to statistical accuracy, future studies should evaluate hedging performance, economic profitability, computational efficiency, model stability, explainability, and robustness through rolling-window and genuinely forward-looking validation designs.

For practical option valuation in the Tehran Stock Exchange, investors, analysts, brokerage firms, portfolio managers, and risk-management units may benefit from using machine-learning models alongside conventional analytical pricing methods rather than relying exclusively on Black-Scholes. LSTM can be considered especially useful where sufficient sequential market data are available, while Decision Tree and SVR can provide complementary estimates with different levels of computational complexity and interpretability. Pricing systems should be recalibrated regularly as new observations become available because relationships learned from historical data may change when volatility, liquidity, investor behavior, or market structure shifts. Practitioners should also evaluate predictions separately by maturity and moneyness instead of assuming uniform accuracy across all contracts, with particular caution applied to out-of-the-money options because they displayed substantially larger errors. A multi-model framework in which algorithmic prices are compared with Black-Scholes values and actual market conditions may provide a more informative basis for identifying unusual pricing, assessing trading opportunities, and supporting hedging decisions than reliance on a single valuation method.

Acknowledgments

We would like to express our appreciation and gratitude to all those who helped us carrying out this study.

Authors' Contributions

All authors equally contributed to this study.

Declaration of Interest

The authors of this article declared no conflict of interest.

Ethical Considerations

All ethical principles were adhered in conducting and writing this article.

Transparency of Data

In accordance with the principles of transparency and open research, we declare that all data and materials used in this study are available upon request.

Funding

This research was carried out independently with personal funding and without the financial support of any governmental or private institution or organization.

References

1. Xiang S. Research on Option Pricing Method Based on the Black-Scholes Model. *Economics and Management Innovation*. 2025;2(1):69-77. doi: 10.71222/k8mkc798.
2. Rahayu RA, Agustina F, Sidarto KA. Pricing Modified Barrier Options Using the Bino-Trinomial Tree Model: A Strategy for Loss Minimization. *CAUCHY: Jurnal Matematika Murni dan Aplikasi*. 2025;10(2):556-70. doi: 10.18860/cauchy.v10i2.33239.
3. Majumdar R, Mariano P, Peng L, Sisti A. A Derivation of the Black-Scholes Option Pricing Model Using a Central Limit Theorem Argument. *arXiv*; 2018.
4. Kaushik V. Efficiency of Black and Scholes Model for Pricing Options at Indian Stock Market. *Management Dynamics*. 2017;17(2):52-9. doi: 10.57198/2583-4932.1053.
5. Amiri M. Option Pricing under Black-Scholes, Boness and Binomial Tree Models: Evidence from the Gold Coin Option Contracts in Iran Mercantile Exchange. *Journal of Securities Exchange*. 2020;13(50):141-70.
6. Abvali M, Khaliliaraghi M, Hasanabadi H, Yaghoobnezhad A. Optional Trading Pricing with a New Analytic Method for the Black-Scholes Equation. *Financial Management Strategy*. 2019;7(3):135-55.
7. Mahdavi R. Investigating the Effectiveness of Option Pricing Using Machine Learning Models Compared to the Black-Scholes Model. *Financial Management Strategy*. 2025;13(49):51-71. doi: 10.22051/jfm.2025.48366.2974.
8. Rakshit A, Bandyopadhyay G, Chakraborty T. Hedge Error Analysis in Black Scholes Option Pricing Model: An Asymptotic Approach towards Finite Difference. *arXiv*; 2024.
9. Xu Y, Liu X, Cao X, Huang C, Liu E, Qian S, et al. Artificial Intelligence: A Powerful Paradigm for Scientific Research. *The Innovation*. 2021;2(4). doi: 10.1016/j.xinn.2021.100179.
10. Bianchini S, Müller M, Pelletier P. Artificial Intelligence in Science: An Emerging General Method of Invention. *Research Policy*. 2022;51(10):104604. doi: 10.1016/j.respol.2022.104604.
11. Miller T. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence*. 2019;267:1-38. doi: 10.1016/j.artint.2018.07.007.

12. Alipour MA, Parsinejad S. Designing an Appropriate Neural Network to Assist in More Accurate Stock Portfolio Predictions. *Transactions on Machine Intelligence*. 2018;1(4):228-42. doi: 10.47176/TMI.2018.228.
13. Karimi Dastgerdi A, Zamani Boroujeni F. A Review of Deep Learning Methods for Financial Market Prediction. *Transactions on Data Analysis in Social Science*. 2020;2(3):164-72. doi: 10.47176/TDASS.2020.164.
14. Li Y, Zhu Z, Kong D, Han H, Zhao Y. EA-LSTM: Evolutionary Attention-Based LSTM for Time Series Prediction. *Knowledge-Based Systems*. 2019;181:104785. doi: 10.1016/j.knosys.2019.05.028.
15. Ghayoumi Zadeh H, Fayazi A, Rahmani Seryasat O, Rabiee H. A Bidirectional Long Short-Term Neural Network Model to Predict Air Pollutant Concentrations: A Case Study of Tehran, Iran. *Transactions on Machine Intelligence*. 2022;5(2):63-76. doi: 10.47176/TMI.2022.63.
16. Singh A, Dargar SK, Gupta A, Kumar A, Srivastava AK, Srivastava M, et al. Evolving Long Short-Term Memory Network-Based Text Classification. *Computational Intelligence and Neuroscience*. 2022;2022(1):4725639. doi: 10.1155/2022/4725639.
17. Mohammadi S, Mansourfar G. The Effect of Financial Data Noise on the Long-Term Co-Movement of Stock Markets. *Transactions on Data Analysis in Social Science*. 2022;4(1):9-21. doi: 10.47176/TDASS/2022.9.
18. Ibrahim M. Evolution of Random Forest from Decision Tree and Bagging: A Bias-Variance Perspective. *Dhaka University Journal of Applied Science and Engineering*. 2022;7(1):66-71. doi: 10.3329/dujase.v7i1.62888.
19. Kappelhof N, Ramos LA, Kappelhof M, van Os HJ, Chalos V, van Kranendonk KR, et al. Evolutionary Algorithms and Decision Trees for Predicting Poor Outcome after Endovascular Treatment for Acute Ischemic Stroke. *Computers in Biology and Medicine*. 2021;133:104414. doi: 10.1016/j.combiomed.2021.104414.
20. Piccialli V, Sciandrone M. Nonlinear Optimization and Support Vector Machines. *Annals of Operations Research*. 2022;314(1):15-47. doi: 10.1007/s10479-022-04655-x.
21. Yeh WC, Jiang Y, Tan SY, Yeh CY. A New Support Vector Machine Based on Convolution Product. *Complexity*. 2021;2021(1):9932292. doi: 10.1155/2021/9932292.
22. Ivaşcu CF. Option Pricing Using Machine Learning. *Expert Systems with Applications*. 2021;163:113799. doi: 10.1016/j.eswa.2020.113799.
23. Jayaraman JD, Zhu X, Rabaa'i AA, editors. An Evaluation of Data Driven Machine Learning Approaches to Option Pricing. *Proceedings of the Northeast Business & Economics Association*; 2022.
24. Liu Y, Zhang X. Option Pricing Using LSTM: A Perspective of Realized Skewness. *Mathematics*. 2023;11(2):314. doi: 10.3390/math11020314.
25. Fan L, Sirignano J. Machine Learning Methods for Pricing Financial Derivatives. *arXiv*; 2024.