Comparative Analysis of Regression Models for Tesla Closing-Price Prediction

Ze Ni

Department of Economics, The Ohio State University, Columbus, OH 43210, USA
`ni.447@buckeyemail.osu.edu`

Abstract. This study addresses the challenge of short-term stock-price forecasting by comparing six regression techniques for predicting the same-day closing price of Tesla, Inc. (TSLA). A ten-year dataset (September 2014–September 2024) of daily open, high, low, close, and volume data was enriched with technical indicators—simple and exponential moving averages, relative strength index, and on-balance volume—and split chronologically into 80 % training and 20 % testing sets. Models evaluated include ordinary least squares, ridge regression (L2 regularization), lasso regression (L1 regularization), k-nearest neighbors, random forest, and gradient boosting. Hyperparameters were selected via nested five-fold, time-series cross-validation, and out-of-sample performance was measured by root mean squared error, mean absolute error, mean absolute percentage error, and coefficient of determination. Results indicate that ridge regression with a tuned penalty coefficient ($\alpha = 0.1$) achieved the lowest test RMSE of \$2.60, closely followed by ordinary least squares with RMSE of \$2.53, MAE near \$2.00, MAPE under 1 %, and R^2 above 0.99. In contrast, k-nearest neighbors and ensemble methods exhibited significant overfitting. These findings demonstrate that carefully engineered technical features combined with regularized linear models yield robust forecasts for highly volatile equities.

Keywords: Technical indicators; Stock-price prediction; Regression analysis; Cross-validation

1 Introduction

Stock-price forecasting remains one of the most challenging problems in quantitative finance owing to high volatility, noisy market dynamics, and non-stationary time series [1]. Accurate short-term predictions support algorithmic trading, risk management, and portfolio optimization, since modest improvements in forecast accuracy can translate into substantial financial gains [2]. Tesla, Inc. (TSLA) exemplifies a highly volatile growth equity driven by rapid technological innovation and shifting market sentiment; precise estimates of its daily closing price can inform both institutional and retail trading strategies.

Several comparative studies have evaluated statistical and machine-learning methods for equity forecasting but have typically relied on a single error metric. Patel

© The Author(s) 2026

A. J. Moshayedi (ed.), *Proceedings of the 2025 International Conference on Hybrid Commerce, Human Capital, and Economic Dynamics (ICHCH 2025)*, Advances in Economics, Business and Management Research 374,
https://doi.org/10.2991/978-2-38476-585-0_34

et al. compared ordinary least squares (OLS), support vector machines, and random forests on multiple stock indices, finding that ensemble regressors reduced root-mean-squared error by up to 15 percent relative to OLS [3]. Bao et al. demonstrated that a deep-learning framework combining stacked autoencoders with long short-term memory networks outperformed classical autoregressive models by capturing nonlinear dependencies [4]. Fischer and Krauss applied LSTM, multilayer perceptrons, and support vector regression to S&P 500 constituents and observed a 12 percent improvement in directional accuracy [5]. More recently, hybrid architectures combining attention-based Transformer networks with gradient-boosting ensembles have yielded further gains in forecast precision [6].

To address this gap, six regression techniques—ordinary least squares, Ridge regression, Lasso regression, k-nearest neighbors, Random Forest, and Gradient Boosting—are applied to ten years of Tesla daily OHLCV data enriched with technical indicators (moving averages, relative strength index, and on-balance volume). The dataset is split chronologically into an 80 % training set and a 20 % hold-out test set. Hyperparameters are selected by nested five-fold, time-series cross-validation on the training data. Out-of-sample performance is evaluated on the test set using root-mean-square error, mean absolute error, mean absolute percentage error, and coefficient of determination. The objective of this study is to determine which regression method offers the best tradeoff between predictive accuracy and robustness for forecasting Tesla's high volatility closing prices.

2 Theoretical Basis

2.1 Role of Technical Indicators

Technical indicators transform raw price and volume series into features that may capture market sentiment or momentum [7]. Simple moving averages (SMA) smooth short-term fluctuations to reveal enduring trends, while exponential moving averages (EMA) weight recent observations more heavily to enhance responsiveness to new information [8]. Momentum oscillators such as the 14-day Relative Strength Index (RSI) compare average gains and losses to flag overbought or oversold conditions [8]. On-Balance Volume (OBV) aggregates volume flows based on daily price direction, providing volume-pressure signals that often lead price movements [7]. Although widely used by practitioners, empirical evidence on the predictive power of individual indicators is mixed, suggesting that combinations of complementary features may yield better forecasts.

2.2 Linear and Regularized Regression

Regression methods offer interpretability and computational efficiency but can suffer from high variance when predictors are noisy or multicollinear [9]. Ordinary least squares (OLS) provides the best linear unbiased estimates under ideal assumptions and is fast to compute, but it lacks any mechanism to penalize large or irrelevant coefficients, which can lead to severe overfitting in high-dimensional or noisy settings.

Ridge regression introduces an L_2 penalty to shrink coefficients toward zero, improving stability and handling multicollinearity effectively, though at the cost of some bias [10]. Lasso regression adds an L_1 penalty, encouraging sparsity by driving some coefficients exactly to zero, which enhances interpretability and generalization when only a few predictors carry signal; however, it may arbitrarily select among highly correlated features [10]. Prior studies in asset pricing have shown that careful regularization can substantially improve out-of-sample performance when applied to large factor sets [9]. Tuning the penalty parameter via nested cross-validation is therefore critical, particularly in time-series contexts where autocorrelation can distort simple validation schemes [9].

2.3 Nonparametric and Ensemble Methods

Nonparametric approaches such as k-nearest neighbors (k-NN) estimate the response by averaging outcomes of the k most similar past observations, capturing nonlinear relationships without a predetermined functional form; this flexibility makes k-NN robust to model misspecification, but its performance degrades sharply in high dimensions and it incurs high prediction latency. Ensemble methods leverage multiple decision trees to improve predictive power and reduce overfitting. Random Forest builds many decorrelated trees on bootstrap samples and averages their predictions to lower variance; it handles mixed data types and outliers well, but offers limited extrapolation beyond the training range. Gradient Boosting sequentially fits trees to residual errors, providing fine-grained control over the bias–variance trade-off through learning rate, tree depth, and number of iterations; its high accuracy comes at the cost of greater sensitivity to hyperparameter settings and a propensity to overfit without shrinkage or early stopping. Effective hyperparameter selection for tree depth, number of estimators, and learning rate via nested cross-validation is therefore essential to achieve robust performance on noisy financial data [11].

3 Methodology

3.1 Data Acquisition and Cleaning

Daily open, high, low, close, and volume (OHLCV) data for Tesla, Inc. (TSLA) from September 15, 2014, through September 13, 2024 (2,517 trading days) were obtained from the “Tesla Stocks Dataset” on Kaggle [12]. Non-trading days and rows with missing values were removed; small gaps were forward-filled. Price and volume outliers exceeding five standard deviations from the rolling mean were winsorized to mitigate data errors or one-off events.

3.2 Feature Engineering

For each day, technical indicators were computed: simple moving averages over 5, 10, and 20 days; exponential moving averages over 5 and 10 days; 14-day RSI; and OBV.

These indicators augment the raw OHLCV features, yielding a total of 11 predictors per observation. All features were standardized (zero mean, unit variance) based on training-set statistics to ensure comparability and to improve convergence of regularized models.

3.3 Train/Test Split and Cross-Validation

Observations were ordered chronologically and split into an 80 percent training set (September 2014–December 2021; 2,010 days) and a 20 percent test set (January 2022–September 2024; 502 days). Model hyperparameters were tuned via nested five-fold, time-series cross-validation on the training set: the outer loop held out each of five chronological blocks for validation, and the inner loop conducted a grid search over each model's parameter grid to minimize validation RMSE.

3.4 Evaluation Metrics

Out-of-sample performance was measured on the held-out test set using four metrics: RMSE, MAE, MAPE, and R^2. Lower RMSE, MAE, and MAPE indicate more accurate forecasts; higher R^2 indicates better explained variance. Nested CV RMSE values were recorded to assess in-sample generalization and to aid hyperparameter selection.

4 Result

4.1 Baseline Fit under Default Settings

Table 1 shows each model's training, nested-CV, and out-of-sample RMSE before any feature standardization or hyperparameter tuning.

Table 1. Baseline RMSE (USD) on train, CV, and test sets

Model	Train RMSE	CV RMSE	Test RMSE
Linear	2.24	2.00	2.53
Ridge	2.59	2.21	3.22
Lasso	3.14	2.26	3.81
KNN	3.10	17.07	17.39
Forest	1.09	12.47	7.86
Gradient Boosting	0.02	12.18	7.53

4.2 Tuned-Model Performance

After standardizing inputs and selecting optimal hyperparameters by nested CV, Table 2 reports the final out-of-sample metrics (RMSE, MAE, MAPE, R^2).

Table 2. Test-set performance of tuned models (USD; MAPE in %)

Model	RMSE	MAE	MAPE(%)	R^2
Linear	2.527	1.988	0.98	0.996
Ridge	2.598	2.054	1.01	0.995
Lasso	4.144	3.194	1.57	0.988
KNN	17.386	14.455	7.73	0.790
Forest	7.858	5.845	3.17	0.957
Gradient Boosting	7.524	5.735	3.06	0.961

4.3 Discussion

The benchmark results highlight that hybrid architectures combining attention-augmented CNN-LSTM modules with gradient-boosting learners can capture both local temporal patterns and global nonlinearities, yielding superior predictive performance compared to standalone methods [13]. However, reliance on daily OHLCV data and a limited set of technical indicators may overlook fast-moving intraday dynamics and sentiment shifts, which have been shown to materially enhance short-term forecast accuracy when incorporated into explainable deep-learning frameworks [14]. Building on these insights, future research should explore minute-level price data, integrate news-media and social-media sentiment features, and develop hybrid models that blend linear regression with attention-based neural modules or gradient-boosting trees to further improve robustness in high-volatility equity forecasting [11, 12].

5 Conclusion

This study has demonstrated that, among six regression techniques applied to ten years of Tesla daily data enriched with moving averages, RSI, and on-balance volume, simple linear models—particularly ordinary least squares and Ridge regression with tuned regularization—offer the strongest balance of predictive accuracy and robustness. While nonparametric and ensemble methods can achieve lower training error, they are prone to overfitting and yield unstable out-of-sample performance.

The findings underscore three key insights. First, model simplicity with careful feature engineering can outperform more complex methods, especially when data are noisy and volatility is high. Second, hybrid architectures that integrate linear components with deep-learning modules—such as attention-augmented CNN-LSTM or Transformer blocks feeding into gradient-boosted trees—hold promise to capture both short- and long-term dependencies. Third, expanding beyond daily OHLCV to include intraday price movements, volatility-driven feature engineering, and sentiment indicators from news and social media may substantially enhance forecast precision.

Future research should therefore explore minute-level datasets, implement explainable hybrid models that blend linear regression with state-of-the-art sequence learners, and rigorously compare their performance under real-time trading conditions.

Such extensions will help develop more reliable, high-frequency forecasting tools for volatile equities.

References

1. Makridakis, S., Spiliotis, E., Assimakopoulos, V.: Statistical and machine learning forecasting methods: Concerns and ways forward. PLoS ONE 13(3), e0194889 (2018)
2. Patel, J., Shah, S., Thakkar, P., Kotecha, K.: Predicting stock market index using fusion of machine learning techniques. Expert Systems with Applications 42(4), 2162–2172 (2015)
3. Bao, W., Yue, J., Rao, Y.: A deep learning framework for financial time series using stacked autoencoders and long-short term memory. PLoS ONE 12(7), e0180944 (2017)
4. Fischer, T., Krauss, C.: Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research 270(2), 654–669 (2018)
5. Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., Sun, L.: Transformers in time series: A survey. In: IJCAI 2023 Survey Track, pp. 6778–6786 (2023)
6. Nie, Y., Nguyen, N.H., Sinthong, P., Kalagnanam, J.: A time series is worth 64 words: Long-term forecasting with transformers. arXiv:2211.14730 (2022)
7. Brock, W., Lakonishok, J., LeBaron, B.: Simple technical trading rules and the stochastic properties of stock returns. Journal of Finance 47(5), 1731–1764 (1992)
8. Achelis, S.B.: Technical Analysis from A to Z. 2nd edn. McGraw-Hill, New York (2001)
9. Gu, C., Kelly, B., Xiu, D.: Empirical asset pricing via machine learning. Review of Financial Studies 33(5), 2223–2273 (2020)
10. Tibshirani, R.: Regression shrinkage and selection via the Lasso. Journal of the Royal Statistical Society: Series B (Methodological) 58(1), 267–288 (1996)
11. Sadorsky, P.: A Random Forests approach to predicting clean energy stock prices. Journal of Risk and Financial Management 14(2), Article 48 (2021)
12. iamtanmayshukla: Tesla Stocks Dataset. https://www.kaggle.com/datasets/iamtanmayshukla/tesla-stocks-dataset, last accessed 2025/8/5
13. Shi, Z., Hu, Y., Mo, G., Wu, J.: Attention-based CNN-LSTM and XGBoost hybrid model for stock prediction. arXiv:2204.02623 (2022)
14. Su, J., Zhang, C., Huang, H.: Intraday and post-market investor sentiment for stock price prediction: A deep learning framework with explainability and quantitative trading strategy. Systems 13(5), Article 390 (2024)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.