Leveraging Supplier Stock Dynamics to Predict Tesla's Market Performance

1st Srihari Ananthan
M.Sc. Big Data and Artificial Intelligence
SRH University Heidelberg
Berlin, Germany
asrihari707@gmail.com

2nd Gagandeep Shivanna
M.Sc. Big Data and Artificial Intelligence
SRH University Heidelberg
Berlin, Germany
gagan017ad@gmail.com

3rd Rohan Rajendra Dhanawade
M.Sc. Big Data and Artificial Intelligence
SRH University Heidelberg
Berlin, Germany
rohan.dhanawade97@gmail.com

4th Kaustubh Salunke
M.Sc. Big Data and Artificial Intelligence
SRH University Heidelberg
Berlin, Germany
kaustubh.v.salunke@gmail.com

5th Klaus Schwarz
SRH University Heidelberg
Berlin, Germany
klaus.schwarz@srh.de

***Abstract*—In this evolving era of AI and automation, stock market analysis and prediction methods have grown increasingly diverse. While traditional approaches like basic statistics, linear regression, and LSTM networks (Long Short-Term Memory networks) dominate the landscape, they often overlook complex interdependencies with a company's supply chain. This study explores a novel angle by incorporating supply chain dynamics into stock price prediction. We focus on Tesla Inc., a leading giant in electric vehicles and automation, with a relatively compact but influential supplier network. From engine to chip manufacturers, Tesla's performance closely relates to that of its partners. Recognizing this, we hypothesize that trends in supplier stock prices can provide early signals for Tesla's stock movements due to operational linkages. Stock price data (2018–2024) was extracted via the Yahoo Finance API, and pre-processing included computing daily returns and one-day lagged prices. After refining an extensive list of Tesla suppliers through correlation heatmaps, a stacking ensemble model was implemented using LSTM and Random Forest as base learners and XGBoost as the meta-learner. Performance was evaluated with five-fold time-series cross-validation using RMSE and R^2. Our results show that while supplier data alone yields poor predictions (RMSE = 97.58, R^2 = –1.16), combining it with Tesla's recent 1-day lagged history improves accuracy significantly (RMSE = 13.95, R^2 = 0.96), achieving a 14% improvement over Tesla-only models. However, using longer-lagged features (e.g., 20-day lags) degraded performance drastically. This research offers a novel perspective in stock market prediction, revealing that supplier stock signals, although secondary, offer valuable fine-tuning potential when properly timed, contributing toward a more integrated financial modeling framework.**

***Index Terms*—Supply chain, Tesla, LSTM (Long Short-Term Memory network), Random forest regression, XGBoost regression, Stacking ensemble, AI in finance.**

I. INTRODUCTION

The accurate prediction of stock prices remains a significant challenge in financial markets, as it is influenced by a number of factors. Traditional models often rely on a company's historical price data and broader market indicators. It makes more sense to also consider the financial performance of key suppliers, as it could also exert an influence on a target company's stock valuation because of the increasing interconnectedness of global supply chains. This research investigates the extent to which the stock prices of supply chain of companies contribute to the predictability of a target firm's stock price, specifically focusing on Tesla (TSLA). Amiri et al. [1] demonstrated that inter-stock relationships can be leveraged to improve energy stock price prediction. Building upon this, our work investigates a specific type of inter-stock relationship, the supply chain, to determine its unique contribution to predictive accuracy. Tesla was chosen as the target company due to its complex, understandable, and extensive global supply chain, involving numerous critical component manufacturers and raw material providers, such as CATL, BYD, TSMC, Sony, Fuyao Glass, TE Connectivity, BHP, and Texas Instruments. It is crucial to understand the dynamics between a target company and its supply chain, as the target firm's operational capacity, production costs, its market perception, and financial performance can be directly impacted by the disruptions or successes within the supplier network. However, quantitative evidence regarding the direct predictive power of supplier stock prices on a primary firm's stock price is often limited, and this study aims to fill that gap.

As shown in Fig. 1, a multi-tiered modeling approach has been adopted to comprehensively assess the contribution of supplier stock performance. This involves a progression of models, each building upon the previous one with additional feature sets. To predict Tesla's stock price, the base model utilizes only supplier and general market indicator features, intentionally excluding Tesla's own historical data. The Level 1 Model focuses solely on Tesla's own lagged price history, specifically a 1-day lag, and it establishes a core baseline for self-predictability. The Level 2 Model combines Tesla's price history with a 20-day lag with supplier company features, aiming to evaluate if longer-term historical Tesla data, complemented by supplier information, enhances prediction

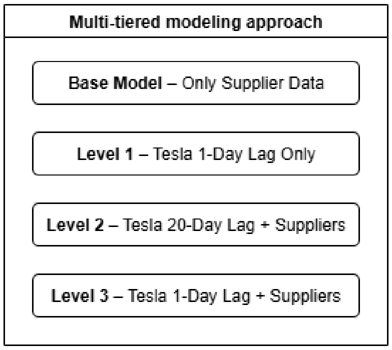

Fig. 1. Multi-tiered modeling approach

accuracy. Finally, the Level 3 Model integrates Tesla's recent (1-day) price history and supplier data, and it determines if this comprehensive feature set provides superior predictive power. Broader market indicators like VIX (market volatility), QQQ (Nasdaq100 tech index), XLI (industrial sector), ICLN (clean energy sector), and ARKQ (autonomous tech and robotics ETF) were selected to capture macroeconomic trends and sector-specific dynamics relevant to Tesla's operations. This setup is used to test whether combining internal and external signals yields the strongest predictive performance. Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), R-squared (R^2), and Directional Accuracy are some of the key metrics to evaluate each model's performance. For consistency and ease of analysis, a standardized "save model results" function was employed to systematically store predictions, true values, and metrics for each model level in both Pickle and JSON formats. The models were tested on approximately 261 samples over a period generally spanning from late October 2023 to late December 2024, ensuring a consistent evaluation window across different model configurations.

II. Related Work

Several foundational studies have validated the use of ensemble and deep learning models in financial forecasting. Gyamerah et al. [2] introduced stacking ensemble learning for stock market movement prediction using a blend of LSTM, Random Forest, and XGBoost, showing improved accuracy through model diversity. Similarly, Krauss et al. [3] demonstrated that combining deep learning models with tree-based algorithms enhances statistical arbitrage on the S&P 500. Jiang et al. [4] proposed an improved stacking framework that integrates deep and tree-based models for stock index prediction. Our work builds upon these studies by not only employing stacking as a meta-model framework but also introducing a multi-tiered ensemble architecture (Base, Level 1, 2, and 3), which allows us to incrementally test and isolate the predictive contribution of Tesla's internal data, supplier stock movements, and broader market indicators.

LSTM networks have gained prominence due to their effectiveness in time-series modeling. Fischer and Krauss [5] confirmed that LSTM outperforms traditional models in capturing sequential dependencies in financial data. Jiang et al. [6] further explored LSTM's power by integrating it with optimization techniques like the Crow Search Algorithm for better forecasting. While we also employ LSTM in all model levels, our novelty lies in integrating LSTM within a hybrid ensemble framework that includes Random Forest and XGBoost. Instead of optimizing a single LSTM model, we leverage the strengths of multiple learners through stacking, particularly focusing on how this combination can reveal short-term dependencies between supplier stock trends and Tesla's price movements.

Finally, Hendricks and Singhal [7] established a strong empirical link between supply chain disruptions and long-term stock performance, laying the groundwork for considering supply chain data in financial models. However, their focus remained on event-driven impacts rather than predictive modeling. Our research diverges by applying granular, real-time supplier stock data not just disruptions to forecast Tesla's stock. Unlike previous work, we quantify the incremental predictive power of suppliers, using correlation heatmaps and tiered modeling to validate their influence. By isolating supplier contributions and showing their modest yet statistically significant impact when paired with recent Tesla price history, we contribute a novel supply chain-aware framework to stock prediction literature.

III. Methodology

This study adopts a layered modeling strategy to explore how much influence the stock prices of Tesla's supply chain partners have on forecasting Tesla's own stock (TSLA). The aim is to assess whether movements in supplier stock prices can meaningfully enhance the accuracy of TSLA price predictions.

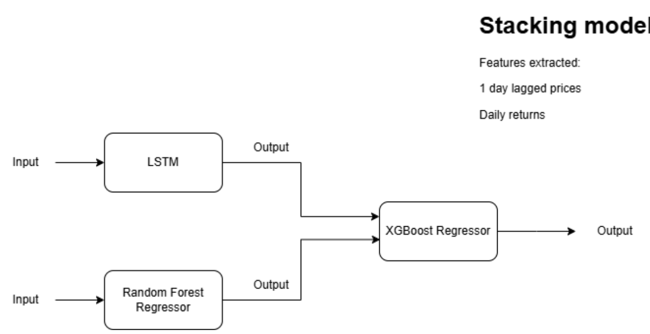

Fig. 2. Architecture of the stacking model used for Tesla stock prediction.

A. Rationale for Research Focus: Supply Chain and Tesla

In their work, Lei et al. [8] demonstrate that machine learning-based approaches can be used for supply chain financial risk prevention. This finding provides a motivation behind this research, which stems from the idea that a few main supplier companies, along with broader market factors, may show predictive value for Tesla's stock behavior. Tesla was selected for this case study because of its influential role in the electric vehicle and automation sectors, as well as its high visibility in financial markets. The company's relatively small yet strategically vital supplier chain system presents an ideal setting for focused analysis. These suppliers play a crucial role in Tesla's operations, covering essential components like chips and power systems. Due to this close-knit and traceable supplier network, the study suggests that fluctuations in their stock prices may offer early indicators of Tesla's market performance. These signals could reflect underlying operational dependencies, hinting at Tesla's stock shifts even before they occur.

B. Data Collection and Preparation

The integration of supplier stock data into the Tesla stock prediction framework was carried out through a structured, three-phase process:

1. Initial Supplier Identification:
An initial pool of 31 potential Tesla suppliers was compiled using a systematic and multi-source methodology. This involved a detailed analysis of publicly available resources, including Tesla's official website, investor reports, financial filings, and sustainability documentation. Additionally, industry news articles, press releases, and third-party analyses were consulted to validate supplier relevance, particularly in key domains such as battery production, raw materials, and semiconductor supply. This rigorous filtering ensured the inclusion of only those companies with direct and strategic operational ties to Tesla. To preserve the native financial trends within each market, all stock prices were retained in their original currencies, avoiding distortions caused by foreign exchange conversions.

2. Data Sourcing and Filtering:
Historical stock price data for Tesla and the shortlisted suppliers were collected via the Yahoo Finance API, covering the period from January 1, 2018, to December 31, 2024. Suppliers lacking publicly traded stock data, such as Bosch, ZF Friedrichshafen, and IDRA, or those that transitioned from public to private ownership (e.g., OmniVision) were excluded. The finalized supplier set included: CATL, BYD, TSMC, Sony, Fuyao Glass, Samsung Display, TE Connectivity, BHP, and Texas Instruments.

3. Data Pre-processing:
To ensure data consistency and compatibility with the modeling pipeline, the following pre-processing steps were applied:

- Data Cleaning: Removed missing values and non-trading days to ensure the integrity and continuity of the time-series data.
- Feature Scaling: Applied Min-Max scaling individually to all closing prices to normalize their widely different ranges, which prevents models from being biased toward features with larger numerical values.
- Correlation Analysis: Generated correlation matrices based on both daily returns and scaled prices. This was done to assess the potential contribution of each supplier's stock movement to Tesla's performance. It is crucial to note that while correlation can indicate a relationship, it does not imply causation.
- Visualization: Visualized correlations using a heatmap to identify suppliers with the strongest relationships to Tesla's stock movements. Appendix Fig. 7 shows the heatmap of the final list of suppliers.

C. Feature Engineering and Selection

To enhance the dataset's suitability for time-series modeling, multiple feature engineering techniques were implemented. Initially, raw closing prices were converted into daily returns, which better represent relative price fluctuations and help achieve stationarity, a key requirement for many predictive algorithms. In addition to returns, lagged features were introduced to incorporate temporal dependencies. Specifically, a one-day lag was used to capture short-term influences, reflecting the potential immediate impact of a supplier's prior-day performance on Tesla's stock movement. To explore longer-term behavioral trends and periodic patterns, a 20-day lag was also computed. This combination of short- and medium-term lag features aimed to enrich the model's ability to detect both immediate and delayed correlations across the supply chain network.

D. Modeling Approach - Multi-tiered Strategy

A hierarchical, multi-stage modeling framework was employed, as illustrated in Fig. 1. Each successive model layer incrementally incorporated additional feature sets to evaluate their respective contributions to prediction accuracy. For consistency and robustness, the evaluation primarily focused on the predictions and metrics derived from the stacking ensemble model, as depicted in Fig. 2. Model training and testing were conducted over a time window ranging from late October 2023 to late December 2024, covering approximately 258 to 261 daily samples.

The four distinct model levels assessed were:

Base Model (Suppliers Only): This foundational model utilized only supplier stock data and market indicators to predict Tesla's stock price. Tesla's own historical stock data was intentionally excluded to assess the standalone predictive value of external signals. Market indicators incorporated in this model included VIX, QQQ, XLI, ICLN, and ARKQ.

Level 1 Model (Tesla Only): Serving as a secondary baseline, this model relied solely on Tesla's one-day lagged closing price as its input feature. No supplier data or market-wide indicators were included, allowing for direct comparison with models utilizing broader contextual features.

Level 2 Model (Tesla 20d + Suppliers + Market Indicators): This intermediate model integrated Tesla's 20-day lagged price history alongside supplier data and market indicators. The inclusion of a longer lag aimed to capture extended temporal patterns and assess whether such historical context enhances the contribution of external variables.

Level 3 Model (Tesla 1d + Suppliers + Market Indicators): The final and most comprehensive model combined all three categories of input: Tesla's recent one-day lagged price, supplier stock data, and market indicators. This configuration was designed to capture both immediate and external influences on Tesla's stock movement. By comparing the predictive outcomes across Levels 1 through 3, the incremental benefit of incorporating supplier and market features can be quantitatively assessed.

E. Ensemble Model Architecture (Stacking)

To improve predictive accuracy, a stacked ensemble model was employed, as illustrated in Fig. 2. This framework aggregates the outputs of three progressively complex feature sets from Levels 1, 2, and 3. By integrating diverse learning algorithms, the ensemble capitalizes on the unique strengths of each. The base models (Level 0) consisted of a Long Short-Term Memory (LSTM)[10] network and a Random Forest Regressor. The LSTM, a variant of recurrent neural networks (RNNs), is well-suited for modeling sequential patterns and time-dependent relationships in stock price data. In contrast, the Random Forest Regressor, built on an ensemble of decision trees with bootstrapped sampling, offers robustness to overfitting and handles non-linearities effectively.

Both models were trained in parallel on the same input features to encourage diverse pattern learning. The individual predictions from each were horizontally concatenated to form a new feature matrix, which was then passed to a meta-model (Level 1). This meta-learner was implemented using an XGBoost Regressor, chosen for its efficiency and capability to handle complex relationships. The XGBoost model learned to intelligently weigh and combine the outputs of the LSTM and Random Forest models, thus producing a final refined prediction. This ensemble architecture leverages the complementary strengths of its base learners, resulting in more stable and accurate performance across various data scenarios.

F. Model Evaluation and Results Management

Model performance was assessed using a standardized set of regression metrics, including Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), R-squared (R^2), Directional Accuracy, and the total number of evaluation samples. RMSE quantifies the average magnitude of prediction errors, emphasizing larger deviations, while MAE provides an absolute measure of error regardless of direction. MAPE calculates the average percentage error with safeguards against division-by-zero errors. The coefficient of determination R^2 measures the proportion of variance in the target variable explained by the model. Directional Accuracy evaluates the percentage of correct predictions concerning the direction of price movement.

To ensure reproducibility and consistency, all metrics were computed programmatically and stored in both detailed and summary formats. Predictions, ground truth values, and evaluation summaries were systematically archived within a structured directory hierarchy to facilitate subsequent analysis and model comparison.

G. Comparison Methodology and Analysis

A dedicated evaluation pipeline was established to systematically load and analyze stored results for comparative assessment of different model configurations. The primary focus was to determine the impact of incorporating supplier data on Tesla stock price prediction. The Tesla-only model (Level 1) was compared against the most comprehensive model (Level 3), which integrates Tesla's one-day lagged price, supplier data, and broader market indicators.

An incremental comparison strategy was employed to isolate the contribution of each feature set. Initially, models utilizing solely supplier features were contrasted with those based exclusively on Tesla's historical prices. Subsequently, the performance difference between Level 1 and Level 2 models was examined to assess the effect of integrating supplier data and extending Tesla's price history to a 20-day lag. Finally, comparisons between Level 2 and Level 3 models evaluated the impact of incorporating recent Tesla price data (1-day lag) alongside supplier and market indicators.

A range of visualization techniques supported model interpretation, including bar charts for RMSE and R^2 metrics, actual versus predicted value plots, error distribution graphs, and rolling MAE trend lines.

IV. RESULTS

In this section, we look at the performance evaluation of four distinct model levels, which are designed to predict Tesla's stock price, focusing on supplier contribution and market indicator data. We use regression metrics for the evaluation of the stacking model's performance, primarily highlighted for analysis unless specified otherwise.

A. Model Performance Overview

The four model levels assessed are:

- **Base Model (Suppliers Only):** This model established a baseline for the prediction potential of external features alone by using just market and supplier variables.
- **Level 1 Model (Tesla Only):** Serving as a control, this model relied exclusively on Tesla's 1-day lagged price history to establish a benchmark for self-predictability.
- **Level 2 Model (Tesla 20d + Suppliers):** This model evaluated if a longer-term baseline may improve the predictive contribution of the suppliers by combining supplier data with a 20-day delayed price history for Tesla.
- **Level 3 Model (Tesla 1d + Suppliers):** This represented the most comprehensive modeling approach, integrating Tesla's 1-day lagged price history, supplier data, and market indicators.

A summary of the stacking model's performance across these levels is provided below, based on a test period from late 2023 to late 2024, involving 258 to 261 samples. Table 1 presents a comparison of the results for each model level, providing a comprehensive overview of their performance.

Note: RMSE measures the average magnitude of errors; MAE measures average absolute errors; R^2 indicates the

Model	RMSE	MAE	R^2	Dir. Acc.	Samples
Base (Suppliers Only)	97.58	75.29	-1.1641	21.5%	261
Level 1 (Tesla Only)	16.23	9.57	0.9402	39.6%	261
Level 2 (Tesla 20d + Sup.)	53.06	38.70	0.3660	33.9%	258
Level 3 (Tesla 1d + Sup.)	13.95	9.23	0.9558	37.3%	261

TABLE I
MODEL PERFORMANCE OVERVIEW

proportion of variance explained by the model; Directional Accuracy indicates the percentage of correctly predicted price movement directions.

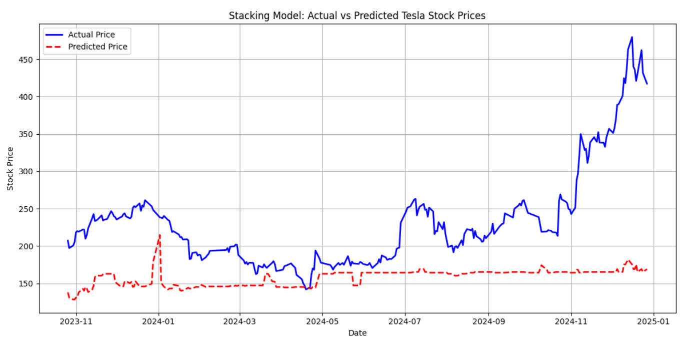

Fig. 3. Actual vs Predicted Tesla prices of base model

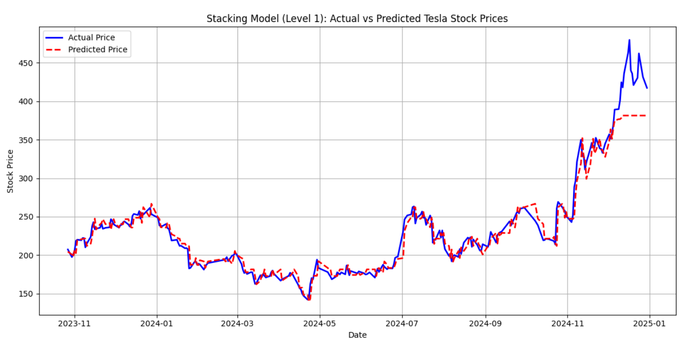

Fig. 4. Actual vs Predicted Tesla prices of Level 1 model

B. *Key Comparisons and Incremental Analysis*

The analysis addresses the primary question of whether suppliers improve Tesla predictions and examines the incremental contribution of different Levels of models.

1) Primary Question: Do Suppliers Improve Tesla Predictions?
   To isolate the pure supplier contribution, a comparison was made between Level 1 (Tesla Only) and Level 3 (Tesla + Suppliers) models, both using a 1-day lag for Tesla's historical price data.
   - Level 1 (Tesla Only) RMSE: 16.2267
   - Level 3 (Tesla + Suppliers) RMSE: 13.9526

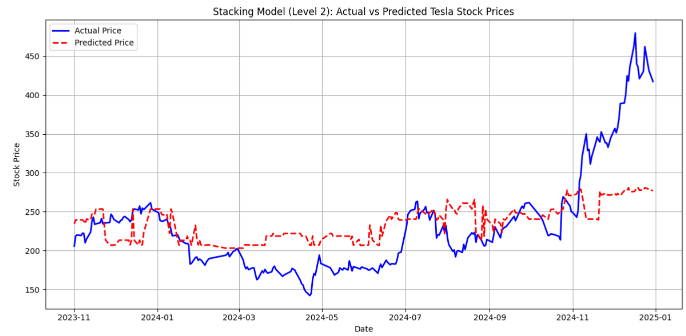

Fig. 5. Actual vs Predicted Tesla prices of Level 2 model

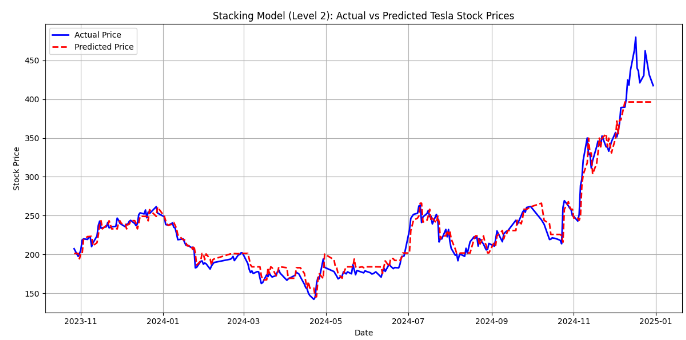

Fig. 6. Actual vs Predicted Tesla prices of Level 3 model

   This shows an improvement of +14.0% when adding suppliers and market indicators to the 1-day lagged Tesla history model.
2) Incremental Analysis
   This investigation looks at how adding various feature sets affects performance. The Appendix's Fig. 8 plots and displays the results.
   - **Base → Level 1:** By including Tesla's price history, the RMSE improved by +83.4% (from 97.5816 to 16.2267), demonstrating that it is a considerably better predictor than supplier and market data alone.
   - **Level 1 → Level 2:** Adding suppliers with a 20-day lag led to a -227.0% RMSE change (16.2267 to 53.0557), indicating a substantial degradation in performance. This suggests suppliers with a longer lag on Tesla's history hurt performance.
   - **Level 1 → Level 3:** Adding suppliers and market indicators with a 1-day lag improved RMSE by +14.0% (16.2267 to 13.9526), showing a modest contribution when combined with recent Tesla data.
   - **Level 2 → Level 3:** Switching from a 20-day lag to a 1-day lag, while including suppliers, resulted in a +73.7% RMSE improvement (53.0557 to 13.9526), highlighting the critical importance of using recent Tesla price history.

C. Key Insights and Conclusion on Supplier Contribution

Based on the performance metrics and incremental analysis, several key insights emerge:

Suppliers by themselves are not very good predictors: Only supplier and market indicators were used in the Base Model, which had a very high RMSE (97.58) and a negative R^2 (-1.16), demonstrating that these characteristics alone did not include enough predictive information. These forecasts were worse than the mean of the observed data. Fig. 3 shows the comparison between the model's actual and anticipated values. It can be seen here that the model failed to follow the trends in Tesla's prices, performing below a random baseline and suggesting that the features were not informative on their own.

The primary factor is Tesla's own pricing history: 94% of the variation in Tesla's price can be explained by the Level 1 model (R^2 = 0.94), demonstrating a high degree of predictive capacity based on Tesla's own momentum. The actual and anticipated values for this model are shown in Fig. 4. Here, trained on Tesla's own lagged price history, it was able to effectively follow the stock's price movements, demonstrating the predictive value of a company's internal data.

When paired with recent Tesla data, adding suppliers provides a slight improvement: The Level 3 model increases RMSE by roughly 14%, but only when a 1-day lag for Tesla's history is used. Fig. 6 shows the Level 3 model's performance by contrasting real and anticipated values.

Timing is crucial: using more recent data (a 1-day lag in Level 3) improves performance, whereas using longer lags (a 20-day lag in Level 2, for example) decreases performance. The Level 2 model's actual and anticipated values are shown in Fig. 5.

Directional accuracy is still difficult to obtain: Even the best models only manage to reach 37–39% directional accuracy, which suggests that it is difficult to accurately predict the direction of price movement.

Alongside Tesla's own recent price history, which continues to be the main driver of predictive accuracy, supplier data makes a minor but incremental contribution to Tesla stock prediction.

V. Future Enhancement

Even though this study clearly shows how useful stacking models and one-day lag features are for enhancing stock price prediction, there are still a number of areas that might use improvement. The low directional accuracy (37–39%) is a major issue that restricts the usefulness of trading in practice. This could be improved in future studies by redefining the job as a classification problem (predicting upward or downward movement) and adding attitude cues from social media sites like Reddit and Twitter, as well as financial news and earnings reports. Furthermore, trading decisions based on anticipated movements rather than price levels could be optimized by investigating reinforcement learning techniques. A promising area for future work is to incorporate alternative error metrics, such as the Mean Arctangent Absolute Percentage Error (MAAPE), as proposed by Kim and Kim [10], to address the limitations of MAPE in a financial forecasting context.

Using Granger causality tests to formally establish a causal relationship between supplier stock fluctuations and the target firm is another improvement. Despite being supported by the existing system, this was not used; using it in further research could more thoroughly guide feature selection and validate directional influence. A more comprehensive use of SHAP (SHapley Additive Explanations) to the interpretation of model outputs can further enhance model transparency. SHAP values can provide insights into supply chain relevance and assist in determining which relationships promote forecast accuracy by measuring the impact of each supplier or market indication. Analyzing the effect of changing all foreign currency stock data to USD on forecasting accuracy is one possible topic for future research. Another choice is to increase the size of the supplier dataset. Future research might include a larger group of supply chain partners, such as logistics companies or second-tier suppliers, rather than just filtering based on correlation. It may also be possible to capture delayed effects that daily stock prices might overlook by incorporating event-level supply chain data, such as disruptions, manufacturing milestones, or partnership announcements.

Lastly, testing the generalizability of this supply chain-aware modeling technique across different companies and industries would uncover characteristics unique to each sector. Future models might potentially be expanded to incorporate equity risk analysis and long-term forecasting, taking use of the proven correlation between supply chain structure and a company's long-term market performance.

VI. Conclusion and Outlook

The Level 1 model's R2 of 0.9402 indicates that Tesla's price history is by far the most important factor in predicting its stock price, accounting for over 94% of the variation alone. With an R2 of -1.1641 and an RMSE of 97.5816, models that just use supplier and market data (the Base Model) perform worse than chance, suggesting that they have virtually no independent predictive ability for changes in Tesla's stock. The prediction accuracy slightly improves when suppliers are included in Tesla's price history. When the Level 3 model (Tesla history + suppliers + market indicators) was compared to the Level 1 model (Tesla history only), the Root Mean Squared Error (RMSE) decreased from 16.2267 to 13.9526, or 14.0%. Additionally, this resulted in a marginal rise in R2 from 0.9402 to 0.9558, signifying an additional 2% of variance explained. It is said that this improvement is statistically significant. Crucially important, though, is the date of the historical data used for Tesla's features. Longer 20-day lags with suppliers (Level 2) dramatically worsen performance, with a -227.0% negative impact when compared to the Tesla-only model, even though 1-day lags with suppliers (Level 3) are helpful. Ultimately, suppliers are not the primary drivers of Tesla stock predictions, even though they do provide an incremental, fine-tuning value. Their contribution offers moderate rather than significant predictive value, and it is subordinated to the momentum of Tesla's stock and other market conditions.

References

[1] Amiri, B., Haddadi, A. & Farajpour Mojdehi, K. A Novel Hybrid GCN-LSTM Algorithm for Energy Stock Price Prediction: Leveraging Temporal Dynamics and Inter-Stock Relationships. *IEEE Access*. **13** pp. 24815-24832 (2025)

[2] Gyamerah, S., Ngare, P. & Ikpe, D. On Stock Market Movement Prediction Via Stacking Ensemble Learning Method. *2019 IEEE Conference On Computational Intelligence For Financial Engineering & Economics (CIFEr)*. pp. 1-8 (2019). https://doi.org/10.1109/CIFEr.2019.8759062

[3] Krauss, C., Do, X. & Huck, N. Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal Of Operational Research*. **259**, 689-702 (2017), https://doi.org/10.1016/j.ejor.2016.10.031

[4] Jiang, M., Liu, J., Zhang, L. & Liu, C. An improved Stacking framework for stock index prediction by leveraging tree-based ensemble models and deep learning algorithms. *Physica A: Statistical Mechanics And Its Applications*. **541** pp. 122272 (2020), https://doi.org/10.1016/j.physa.2019.122272

[5] Fischer, T. & Krauss, C. Deep learning with long short-term memory networks for financial market predictions. *European Journal Of Operational Research*. **270**, 654-669 (2018), https://doi.org/10.1016/j.ejor.2017.11.054

[6] Jiang, C., Tsai, Y., Shao, Z., Lee, S., Hsueh, C. & Huang, K. Hybrid Crow Search Algorithm–LSTM System for Enhanced Stock Price Forecasting. *Applied Sciences*. **14** (2024), https://doi.org/10.3390/app142311380

[7] Hendricks, K. B., & Singhal, V. R. (2005). An Empirical Analysis of the Effect of Supply Chain Disruptions on Long-Run Stock Price Performance and Equity Risk of the Firm. Production and Operations Management, 14(1), 35-52. https://doi.org/10.1111/j.1937-5956.2005.tb00008.x (Original work published 2005)

[8] Lei, Y., Qiaoming, H. & Tong, Z. Research on Supply Chain Financial Risk Prevention Based on Machine Learning. *Computational Intelligence And Neuroscience*. **2023**, 6531154 (2023), https://onlinelibrary.wiley.com/doi/abs/10.1155/2023/6531154

[9] Zhao, C., Hu, P., Liu, X., Lan, X. & Zhang, H. Stock Market Analysis Using Time Series Relational Models for Stock Price Prediction. *Mathematics*. **11** (2023), https://www.mdpi.com/2227-7390/11/5/1130

[10] Kim, S. & Kim, H. A new metric of absolute percentage error for intermittent demand forecasts. *International Journal Of Forecasting*. **32**, 669-679 (2016), https://www.sciencedirect.com/science/article/pii/S0169207016000121

VII. Appendix

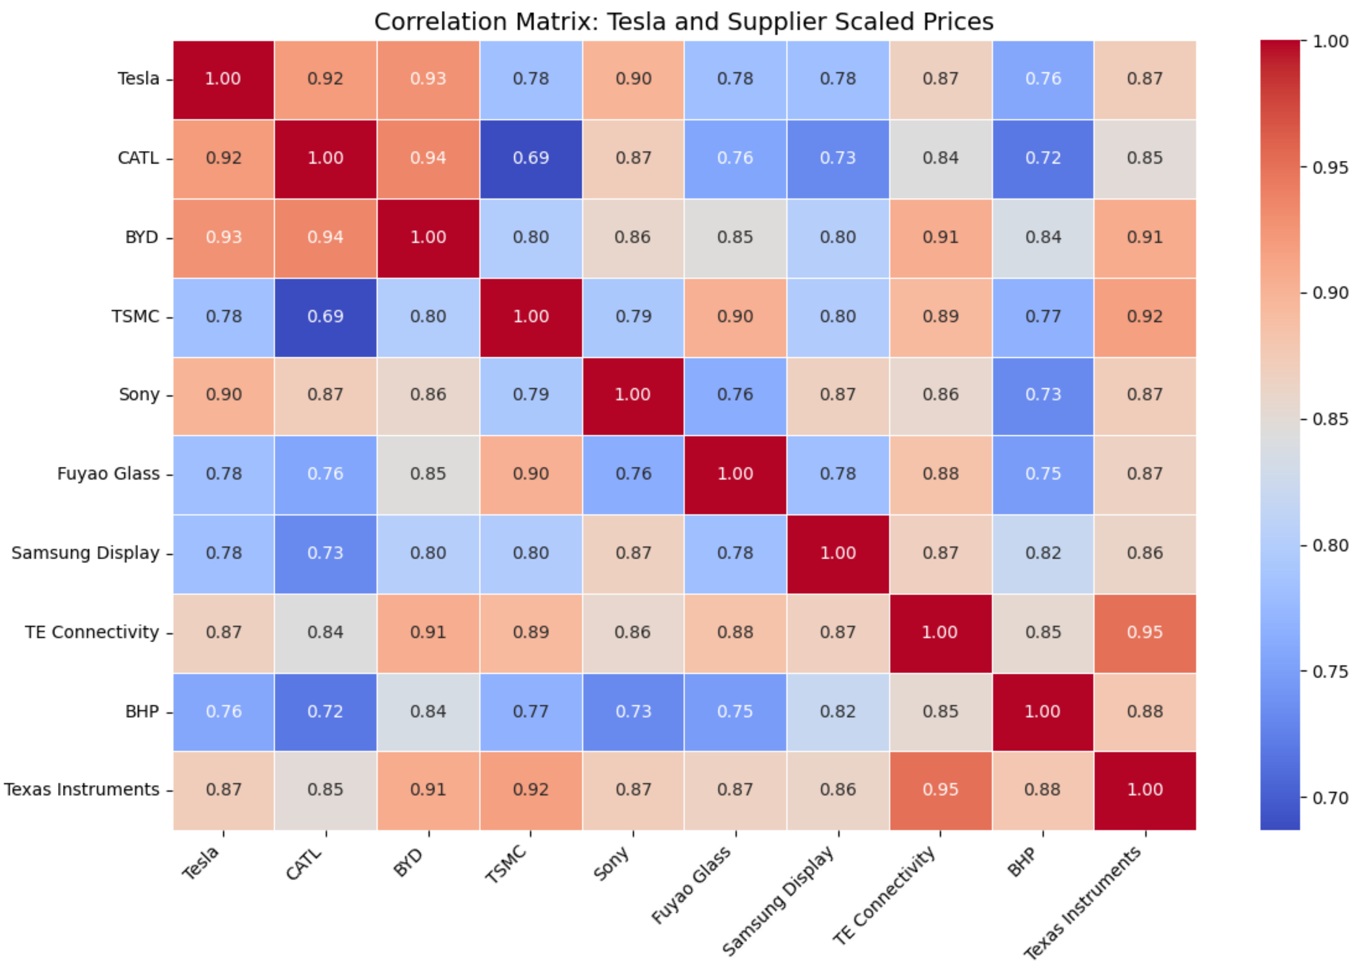

Fig. 7. Heatmap of scaled prices.

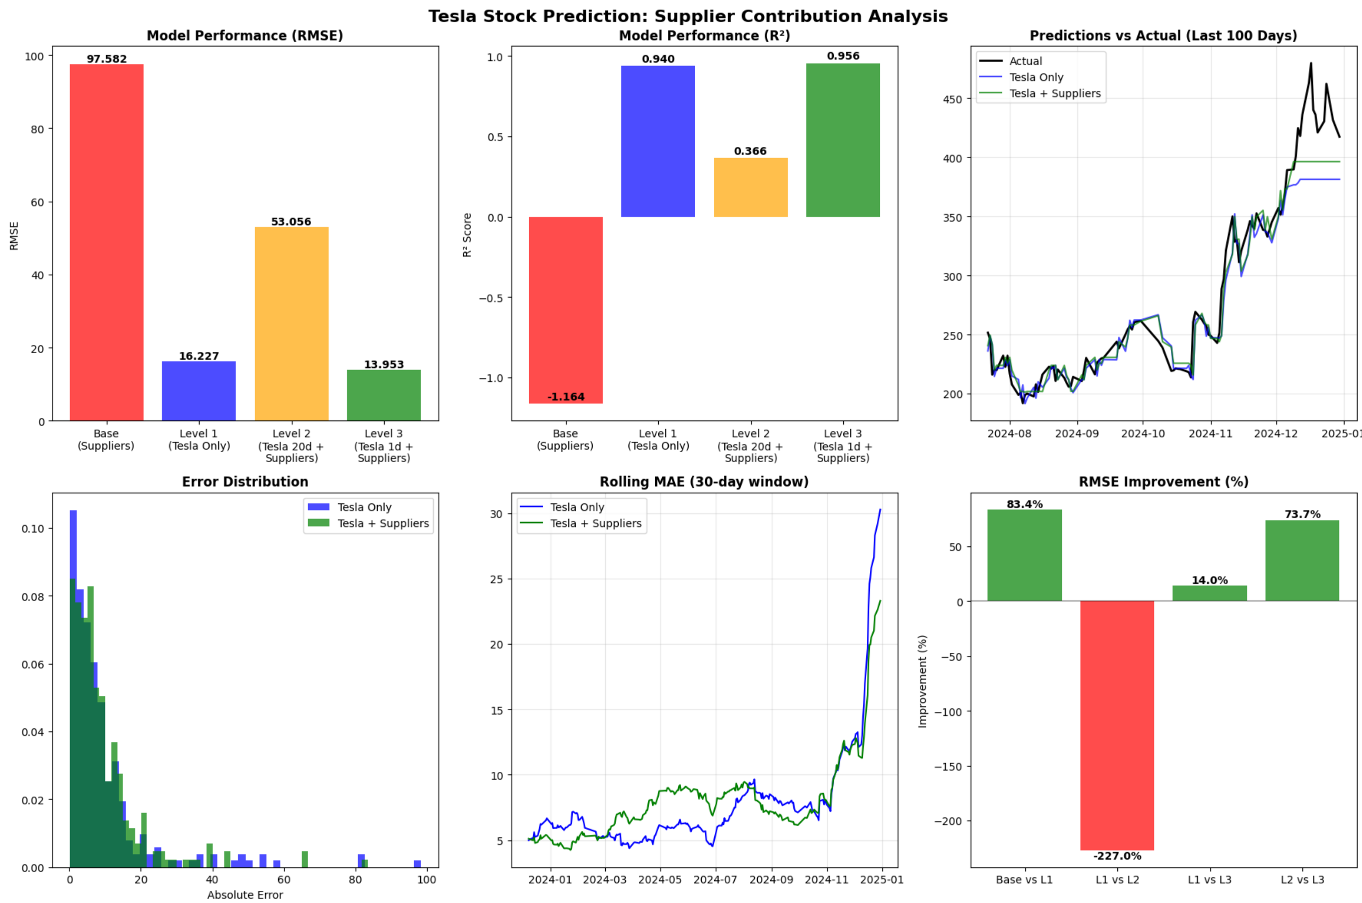

Fig. 8. Result comparisons of all Model levels