RBRU International Journal of Science and Technology 1 (1) (2026) 45–58



RBRU International Journal of Science and Technology

ISSN: 3088-389X (online) journal homepage: https://ph05.tci-thaijo.org/index.php/rbruijst



Research Article

Comparative Evaluation of Forecasting Models for Short-Term Natural Rubber Prices Using a Web-Based Monitoring Platform

Wirawat Athipongteerakul[a], Nopparat Jawchaiyapum[a], Wachirapong Jirakitpuwapat[a,*]

[a] *Faculty of Science, Energy and Environment, King Mongkuts University of Technology North Bangkok, Rayong 21120 , Thailand*

ARTICLE INFO

Keywords:
ARIMA
Natural rubber price
Prophet
Support vector regression
Time series forecasting

Article history:
Received 24 August 2026
Revised 29 August 2026
Accepted 30 August 2026
Published 31 August 2026

ABSTRACT

Thailand is one of the worlds largest producers of natural rubber, a key agricultural export. The price of natural rubber is highly volatile in the short term, driven by factors such as weather, global demand, and spillovers from energy markets. Reliable price forecasts, even for short horizons, can assist farmers, cooperatives, and local government agencies in designing better procurement and support policies. This study presents an offline empirical benchmark comparing three commonly used forecasting methods, Prophet, Support Vector Regression (SVR), and Autoregressive Integrated Moving Average (ARIMA), for short-term daily natural rubber prices using data from a web-based monitoring and forecasting dashboard of rubber prices. The dataset covered the period from 1 January 2016 to 31 August 2026. The data were divided chronologically into a training set from 1 January 2016 to 31 December 2024 and a held-out test set from 1 January 2025 to 31 August 2026, respectively. Each model was trained on a rolling historical window and tested on a held-out forecast horizon using three complementary error metrics: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and coefficient of determination (R^2). The best performance was obtained by the Prophet model (MAE = 1.10 baht/kg, MAPE = 3.20%, R^2 = 0.898), followed by SVR (MAE = 1.35, MAPE = 3.80%, R^2 = 0.875) and ARIMA (MAE = 1.45, MAPE = 4.10%, R^2 = 0.852). These findings indicate that, within this offline empirical benchmark, Prophet's additive decomposition of trend and seasonal components offers a good trade-off between interpretability and accuracy for short-horizon commodity price forecasting, compared with ARIMA and SVR. The results provide operational guidance for selecting lightweight, deployable forecasting models for agricultural price-monitoring systems.

1. Introduction

Natural rubber (*Hevea brasiliensis*) is one of the worlds most economically important agricultural commodities, underpinning a vast array of industries ranging from tire manufacturing and automotive components to construction, healthcare, and general consumer goods. Thailand is the worlds largest producer of natural rubber, contributing about one-third of the worlds supply in recent years [3]. Furthermore, supply has increased in recent years [3]. Thailands

*Corresponding author
DOI:10.68654/rbruijst.2026.01.05
s6613031620406@email.kmutnb.ac.th (W. Athipongteerakul); wachirapong.j@sciee.kmutnb.ac.th (W. Jirakitpuwapat)
This is an open access article under the Creative Commons Attribution – NonCommercial – NoDerivatives 4.0 license.
© 2026 RBRU International Journal of Science and Technology – Published by Rambhai Barni Rajabhat University.

natural rubber exports are valued at about five billion US dollars in 2024 [28], and natural rubber is still one of the countrys most valuable agricultural exports. Rubber cultivation is of special importance in the eastern and southern provinces of Thailand, including Rayong Province, which is the focus of the present study, where smallholder rubber plantations are a major source of household income for a large proportion of the farming population.

Rubber is an international agricultural commodity, and its farm-gate price is very responsive to short-term variations in global supply and demand, crude-oil-linked synthetic rubber prices, exchange rates, and weather-related supply shocks [10, 13, 18, 27]. This volatility creates large uncertainty about income for smallholder farmers, who typically have a limited ability to hedge against price risk. This makes it harder for local farm support agencies, cooperatives, and traders to plan, as they have to decide when to buy, sell, or release stock. Rubber price volatility is not just an academic issue but an active policy issue, as Thai authorities have periodically intervened with price-support and stabilization measures in response to sustained downturns. Another recent study on the system dynamics of the Thai rubber sector also notes that plantation area, climate variables, and policy interventions collectively drive long-run price trends, emphasizing the practical utility of forecasting tools, whether structural or data-driven, for predicting price movements in the Thai context [9].

Accurate short-term price forecasting can mitigate this uncertainty by giving farmers, traders, and policymakers advance information to guide selling decisions, budget planning, and market interventions. Over the past decade, natural rubber price forecasting has been studied using both classical statistical models and modern machine-learning approaches. Classical time-series methods such as the Autoregressive Integrated Moving Average (ARIMA) model remain widely used because of their strong theoretical grounding, relatively low data requirements, and interpretability [4, 30]. More recently, machine-learning regressors such as Support Vector Regression (SVR) and decomposition-based additive models such as Prophet have been applied to commodity price series, motivated by their ability to capture nonlinear patterns and structural trend changes that purely linear models may miss [15, 18, 25]. Alongside this methodological development, digital monitoring platforms that surface near-real-time commodity prices and short-term forecasts, such as the ParaPredict dashboard used in this study, have become increasingly relevant tools for translating forecasting research into forms that are directly usable by farmers, cooperatives, and local government agencies [6].

Despite this growing body of work, relatively few studies directly compare classical, kernel-based, and decomposition-based forecasting approaches on the same short-horizon daily rubber price series while explicitly reporting a common set of accuracy metrics, and Thailand-specific evaluations of this kind remain comparatively scarce relative to studies focused on Malaysian and Indian rubber markets [30]. This study addresses that gap. Specifically, the contributions of this paper are threefold:

(i) it evaluates and directly compares three widely used forecasting approaches: Prophet, SVR, and ARIMA on a common daily natural rubber price series and a common forecast horizon, using a consistent set of accuracy metrics (MAE, MAPE, and R2);

(ii) it grounds the comparison in an operational, web-based rubber price monitoring dashboard (ParaPredict) rather than a purely retrospective academic dataset, so that the findings speak directly to model selection for deployed forecasting systems; and

(iii) it situates the results within the Thai rubber-price forecasting literature, complementing recent structural, policy-oriented studies with a short-horizon, data-driven perspective [9].

2. Literature Review

Classical Box-Jenkins ARIMA models have long been the default benchmark for agricultural commodity price forecasting because of their well-established statistical theory and relatively low data requirements [4]. [30] applied the Box-Jenkins methodology to monthly Standard Malaysian Rubber (SMR20) prices and found that an ARIMA(1,1,0) specification provided an adequate fit after first-differencing to achieve stationarity, illustrating both the strength and the key limitation of ARIMA: its reliance on a linear, stationary representation of the series, which can struggle to capture abrupt regime shifts driven by external shocks. Related ARIMA-based studies on natural rubber prices in other producing countries report broadly similar conclusions: [16] fitted grade-specific ARIMA models to Indian RSS4, latex, and ISNR20 prices, while found that a purely linear ARIMA(1,1,1) specification for Sri Lankan rubber prices was outperformed by a nonlinear time-delay neural network, and [21] showed that an AR-GARCH specification with robust

weighting improved on standard ARIMA when SMR20 prices contained outliers. Within Thailand specifically, [7] modeled Thai natural rubber prices with ARIMA and ARIMAX specifications and reported substantially lower error when synthetic-rubber and world-price regressors were included, foreshadowing the univariate-versus-multivariate trade-off examined in the present study, while [22] compared ARIMA against exponential smoothing and a shallow ANN on Malaysian SMR20 prices and found the ANN specification preferable over a short daily window.

To address this limitation, subsequent studies have incorporated nonlinear and machine-learning-based regressors. [18] compared SVR, artificial neural networks, k-nearest neighbors, random forest, and decision trees against an ARIMAX benchmark for natural rubber price forecasting using a set of commodity-market explanatory indicators, finding that ensemble tree-based methods outperformed both SVR and the linear ARIMAX benchmark. However, SVR still improved on the purely linear model. Similarly, [15] benchmarked ARIMA, SVR, gradient boosting, and several deep-learning architectures across a panel of agricultural commodities and reported that while deep sequence models achieved the lowest errors overall, SVR consistently outperformed ARIMA on series exhibiting nonlinear short-term dynamics, reflecting SVR's capacity to model nonlinear relationships through its kernel-based formulation [8, 24, 26]; selecting appropriate kernel and regularization parameters for SVR is itself a well-studied problem [5]. Beyond ARIMA and SVR comparisons, recent studies have applied deep sequence models directly to natural rubber prices in Southeast Asian markets: [20] combined variational mode decomposition with a bidirectional LSTM for Thai RSS3 forecasting, and [19] used a cointegration-error-correction framework augmented with a gradient-boosting diagnostic to examine price transmission from the Shanghai Futures Exchange to Thai domestic rubber prices. Outside the rubber sector, comparable SVR-versus-ensemble comparisons for agricultural commodities have reached broadly consistent conclusions [23], and SVR has also been shown to be competitive with GARCH-family models for volatility forecasting in commodity and financial markets [2].

A separate line of work has explored decomposition-based additive models. The Prophet framework, introduced by [25], decomposes a time series into trend, seasonal, and holiday components, which are fitted using a Bayesian curve-fitting procedure, and was explicitly designed to be robust to missing data, outliers, and trend changes while remaining interpretable to non-expert analysts. Because commodity price series often exhibit piecewise trend changes driven by policy announcements or global supply shocks, Prophet's automatic changepoint detection can be a useful property for rubber price forecasting. However, its application to natural rubber prices specifically remains comparatively understudied relative to ARIMA and standard machine-learning regressors. Direct comparisons of Prophet against ARIMA and related benchmarks in other commodity and financial domains report mixed results: several studies demonstrate Prophet outperforms ARIMA on error metrics for wheat, wholesale food, and cryptocurrency prices [14, 29], whereas others report the opposite ranking or find that Prophet struggles specifically with the discrete, step-like price jumps that characterize some agricultural commodity series [1, 17].

General forecasting practice recommends evaluating candidate models using multiple complementary error metrics rather than a single statistic, since scale-dependent (MAE), percentage-based (MAPE), and variance-explained (R^2) metrics can highlight different aspects of forecast quality [11]. This study follows that recommendation by reporting all three metrics for each model under comparison.

Within Thailand specifically, most existing rubber price forecasting work has emphasized either long-run structural drivers or country-level policy analysis rather than short-horizon, deployment-oriented forecasting. [9] recently proposed a system dynamics model that integrates climate, economic, and policy factors to forecast long-term Thai rubber price trends and to evaluate the effect of alternative policy interventions, noting that prior Thai studies have largely relied on econometric and machine-learning models focused on short-term fluctuations rather than structural, policy-relevant dynamics. The present study complements this structural perspective by focusing specifically on short-horizon, day-ahead forecasting accuracy using data-driven models suitable for integration into an operational monitoring dashboard, rather than long-run policy simulation. Deep-learning-based price forecasting has also been applied to other Thai agricultural commodities, indicating a broader domestic interest in data-driven forecasting tools for the agricultural sector beyond rubber alone.

3. Materials and Methods

3.1. Data Description and Preprocessing

The dataset consists of a daily natural rubber price series (Baht per kilogram) drawn from the ParaPredict web-based monitoring and forecasting platform. The series was split chronologically into a historical training window used to fit each model and a short-term forecast evaluation horizon, consistent with a rolling-origin forecasting design in

which each model produces multi-day-ahead predictions from the most recent available observations. Prior to model fitting, the series was checked for missing values and outliers, and timestamps were aligned to a consistent daily frequency. No additional exogenous regressors (e.g., crude oil price and exchange rate) were incorporated in this stage of the study; each of the three models discussed below was fitted in a univariate configuration using only the historical price series.

3.2. Autoregressive Integrated Moving Average (ARIMA)

ARIMA models the price series y_t as a linear function of its own past values and past forecast errors, after differencing the series d times to achieve stationarity [4]. The general ARIMA(p, d, q) model is written as

$$\left(1 - \sum_{i=1}^{p} \phi_i L^i\right)(1 - L)^d y_t = \left(1 + \sum_{j=1}^{q} \theta_j L^j\right)\varepsilon_t, \tag{3.1}$$

where L is the lag operator, ϕ_i are the autoregressive coefficients, θ_j are the moving-average coefficients, and ε_t is white noise. Stationarity of the differenced series was assessed prior to fitting, and candidate orders (p, d, q) were selected by examining the autocorrelation (ACF) and partial autocorrelation (PACF) functions of the differenced series, with the final order chosen to minimize the Akaike Information Criterion (AIC) among candidate specifications, following the automatic order-selection procedure of [12].

3.3. Support Vector Regression (SVR)

SVR extends the support vector machine framework to regression by fitting a function that deviates from the observed targets by at most ϵ while remaining as flat as possible, tolerating larger deviations only through slack variables [8, 26]. Given a lagged feature representation of the price series, SVR solves

$$\min_{w,b,\xi,\xi^*} \ \tfrac{1}{2}\|w\|^2 + C\sum_{t=1}^{n}(\xi_t + \xi_t^*), \tag{3.2}$$

subject to $y_t - (w^\top \phi(x_t) + b) \le \epsilon + \xi_t$, $(w^\top \phi(x_t) + b) - y_t \le \epsilon + \xi_t^*$, and $\xi_t, \xi_t^* \ge 0$, where $\phi(\cdot)$ maps the lagged input vector x_t into a higher-dimensional feature space via a kernel function, C controls the trade-off between model flatness and tolerance for deviations larger than ϵ, and ξ_t, ξ_t^* are slack variables. In this study, a radial basis function (RBF) kernel was used, with C, ϵ, and the kernel width γ selected via grid search on a validation split of the training window, following standard practice for parameter selection in kernel-based SVR [5].

3.4. Prophet

Prophet models the price series as an additive decomposition of a piecewise trend component, a periodic seasonal component, and an error term [25]:

$$y_t = g(t) + s(t) + h(t) + \varepsilon_t, \tag{3.3}$$

where $g(t)$ is a piecewise-linear (or logistic) trend function with automatically detected changepoints, $s(t)$ represents periodic seasonal effects modelled using a truncated Fourier series, $h(t)$ captures the effect of known irregular events such as holidays (set to zero in this univariate configuration, as no holiday regressors were incorporated), and ε_t is an idiosyncratic error term. Model parameters are estimated using a Bayesian curve-fitting procedure that yields both point forecasts and uncertainty intervals, the latter of which are shown as shaded bands in Figure 2.

3.5. Evaluation Metrics

Model accuracy was evaluated on the held-out forecast horizon using three complementary metrics [11]. Mean Absolute Error (MAE) measures the average magnitude of forecast errors in the original price units (baht/kg):

$$\mathrm{MAE} = \frac{1}{n}\sum_{t=1}^{n}\left|y_t - \hat{y}_t\right|. \tag{3.4}$$

Table 1
Forecast accuracy of Prophet, SVR, and ARIMA on the held-out evaluation horizon.

Model	MAE (Baht/kg)	MAPE (%)	R^2
Prophet	1.10	3.20	0.898
SVR	1.35	3.80	0.875
ARIMA	1.45	4.10	0.852

Mean Absolute Percentage Error (MAPE) expresses the average error as a percentage of the actual price, allowing comparison across different price levels:

$$\text{MAPE} = \frac{100\%}{n} \sum_{t=1}^{n} \left| \frac{y_t - \hat{y}_t}{y_t} \right|. \tag{3.5}$$

The coefficient of determination (R^2) measures the proportion of variance in the actual price series explained by each model's forecasts:

$$R^2 = 1 - \frac{\sum_{t=1}^{n} (y_t - \hat{y}_t)^2}{\sum_{t=1}^{n} (y_t - \bar{y})^2}, \tag{3.6}$$

where y_t is the observed price, $\hat{y}_t$ is the forecast price, $\bar{y}$ is the mean of the observed prices over the evaluation horizon, and n is the number of observations in that horizon.

4. Results and Discussion

The forecasting results reported in this section were obtained from a separate offline forecasting experiment and should be distinguished from the forecasting functionality currently implemented in the ParaPredict web dashboard. For the offline experiment, Prophet, Support Vector Regression (SVR), and ARIMA were fitted and evaluated using the forecasting procedure described in Section 3. These parameter settings were fixed throughout the offline evaluation to ensure a consistent comparison among the forecasting models. The MAE, MAPE, and R^2 values reported in Table 1 are therefore the results of this offline evaluation and are not calculated by the live dashboard.

The current web-based system serves primarily as a data-monitoring and presentation platform. Its forecasting interface provides visual model trajectories and displays model-comparison information, but the forecasting curves currently shown in the client-side application are simulated representations rather than outputs generated by the fitted models used in the offline experiment.

Figure 1 visualizes the same three metrics as grouped bar charts. Figure 2 shows the 45-day historical price trajectory together with the multi-day-ahead forecasts produced by each model, with a shaded band indicating an approximate forecast uncertainty interval.

The observed error is broadly consistent with the literature reviewed in Section 2. ARIMA's comparatively weaker performance is consistent with its purely linear, stationarity-dependent formulation, which can underfit the short-run nonlinear fluctuations and localised trend shifts that characterise commodity price series. SVR's kernel-based formulation allows it to capture some of this nonlinearity, which is reflected in its intermediate performance relative to ARIMA, consistent with prior comparative studies in which SVR outperformed purely linear benchmarks but was in turn outperformed by more flexible or ensemble methods. Prophet's advantage in this dataset is plausibly attributable to its explicit, automatically detected trend-changepoint mechanism, which may better track the gradual upward drift visible in the historical window than either a fixed-order linear ARIMA specification or a generically tuned SVR kernel, consistent with several comparative studies in which Prophet outperformed ARIMA on trending or seasonal commodity and financial series.

5. Web-Based Implementation

Beyond the model comparison itself, this study's forecasts are delivered via ParaPredict (https://parapredict.web.app/), a web-based agricultural decision-support dashboard designed to make the underlying price data and model

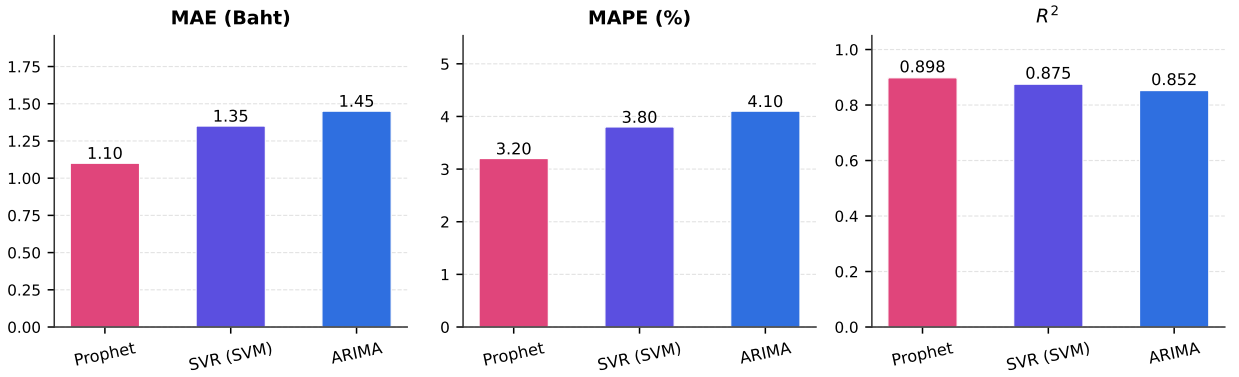


Figure 1: Comparison of MAE, MAPE, and R^2 for Prophet, SVR, and ARIMA.

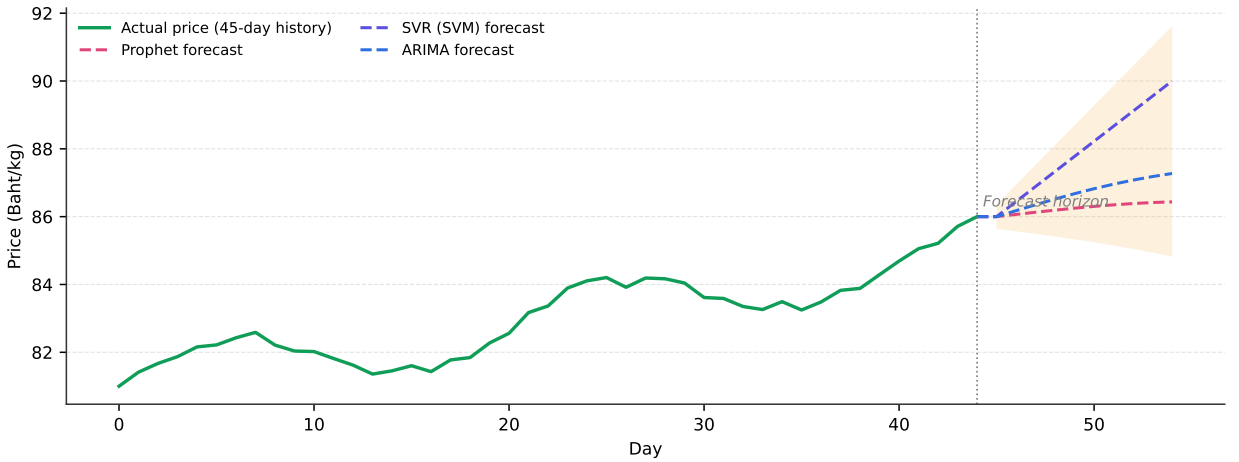


Figure 2: Historical price trajectory over 45 days and multi-day forecasts from Prophet, SVR, and ARIMA models, with the forecast prediction interval shown as a band.

outputs directly usable by farmers, cooperatives, and traders, rather than remaining confined to the analysis notebook. The system is organized into a data-collection backend, a public-facing dashboard, and an administrative interface.

5.1. Data Collection

Rubber price data is collected by a dedicated Python scraper that queries the source price-reporting pages on a scheduled basis and parses the relevant price tables. Each run writes the most recent observations to a structured JSON file (`rubber_prices_latest.json`), which feeds both the forecasting pipeline described in Section 3 and a price database that the dashboard reads from. Separating data collection from the forecasting and presentation layers in this way keeps the scraper independently testable. It allows the historical price series to be extended simply by re-running the collection job, without touching the modeling or display code.

5.2. Public Dashboard

The dashboard is a JavaScript single-page application built with Vite and deployed via Firebase Hosting. It is organized into six main views, accessible from a persistent navigation bar:

- **Overview and Forecast:** The landing page displays the current price for the selected rubber grade and region alongside a 45-day historical price chart with a forecast overlay. Users can filter by region (e.g., South, Northeast, East, North, and Central), rubber type (ribbed smoked sheet, fresh latex, cup lump, or RSS 3), and forecast horizon (7, 14, or 30 days). The forecasting visualizations in the current dashboard are intended as an operational

presentation layer and are separate from the offline forecasting experiments reported in Section 4. In the current implementation, the displayed model trajectories are generated by the client-side simulation described in Section 5 and do not represent live inference from the fitted Prophet, SVR, and ARIMA models used to obtain the experimental results. Likewise, the model evaluation values displayed in the dashboard are presentation-level reference values and should not be interpreted as real-time recalculations of MAE, MAPE, or R^2 from newly collected observations.

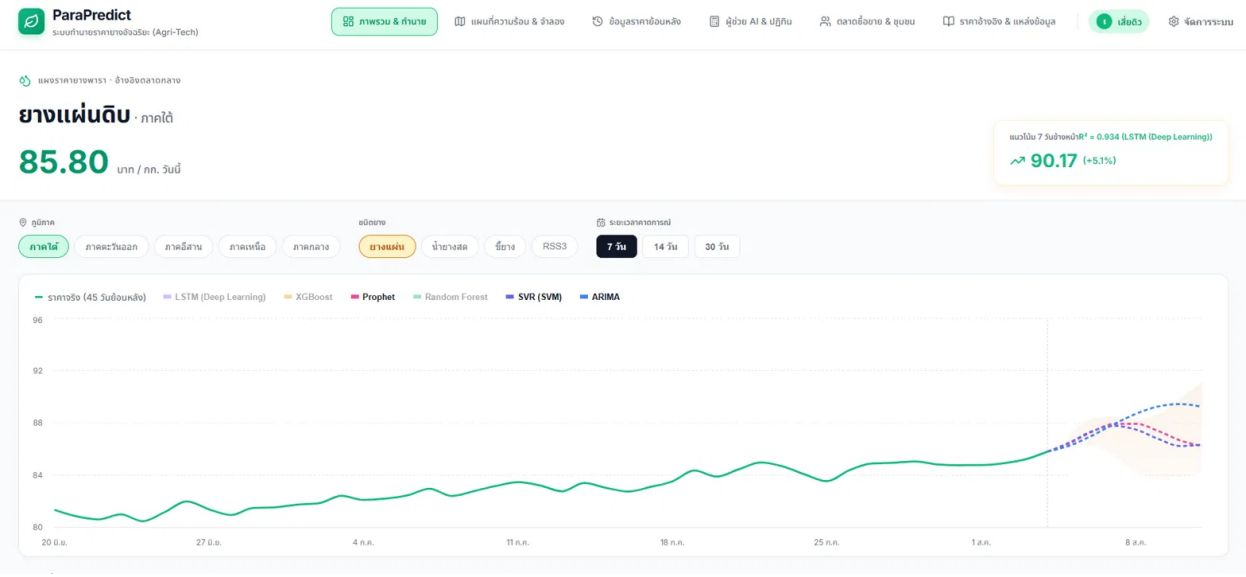


Figure 3: Overview and Forecast page of the ParaPredict dashboard, showing the historical price series with overlaid model forecasts and an uncertainty band for the leading model.

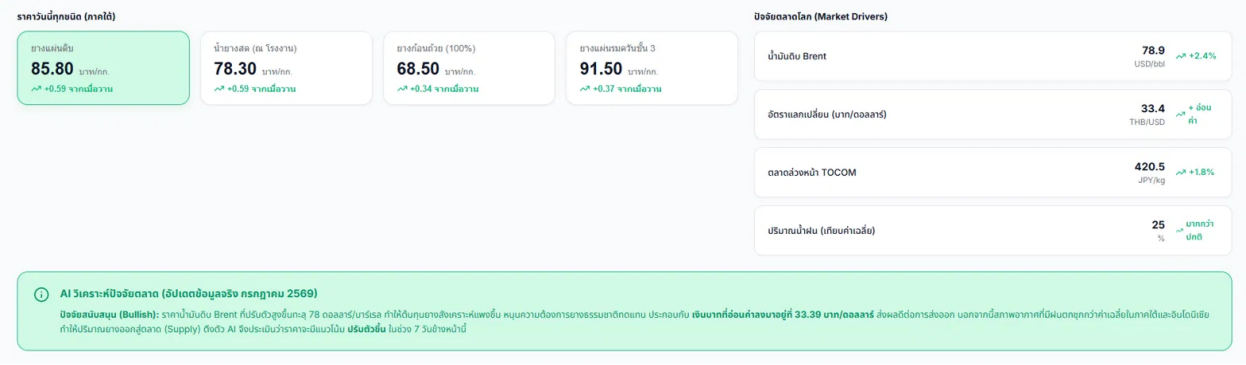


Figure 4: Current-day reference prices by rubber grade and the market-drivers panel, summarising crude oil price, exchange rate, TOCOM futures, and rainfall relative to average.

- **Interactive Map and Simulator:** A Leaflet/CARTO-based heat map lets users select a production region to view its price trend, alongside a what-if simulator with sliders for rainfall anomaly, Brent crude oil price, and the exchange rate, which recomputes a simulated regional price in response to hypothetical changes in these drivers (Figure 5).
- **Historical Price Data:** a filterable archive of past prices by year, month, province, and rubber type, reporting summary statistics (minimum, mean, and maximum price) alongside a time-series chart of the filtered series (Figure 6).

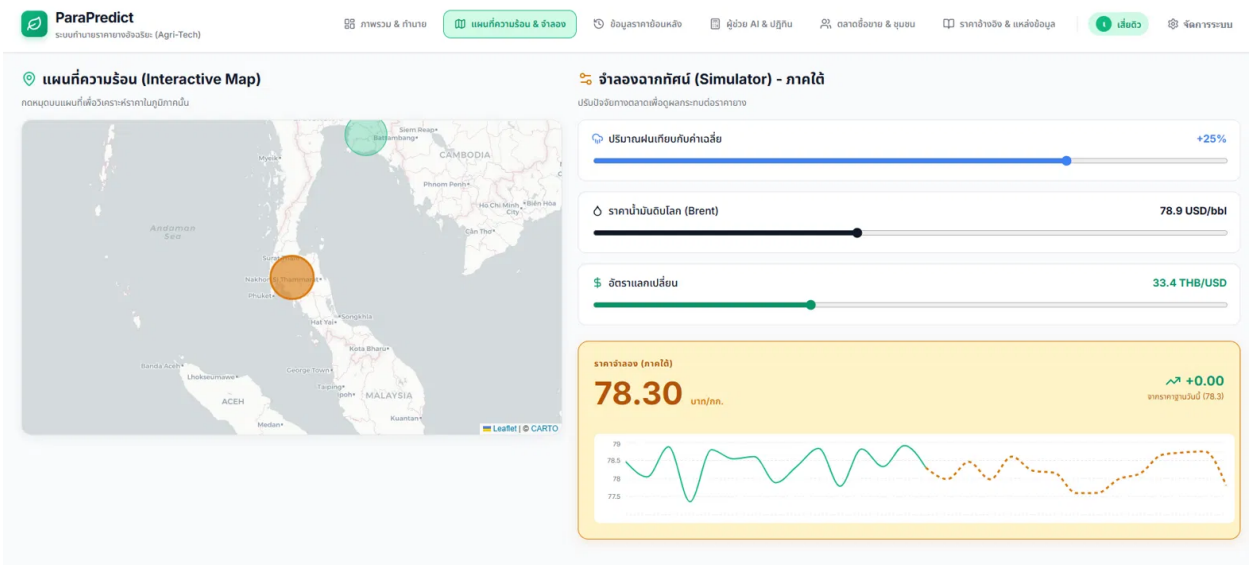


Figure 5: Interactive regional heat map and what-if price simulator.

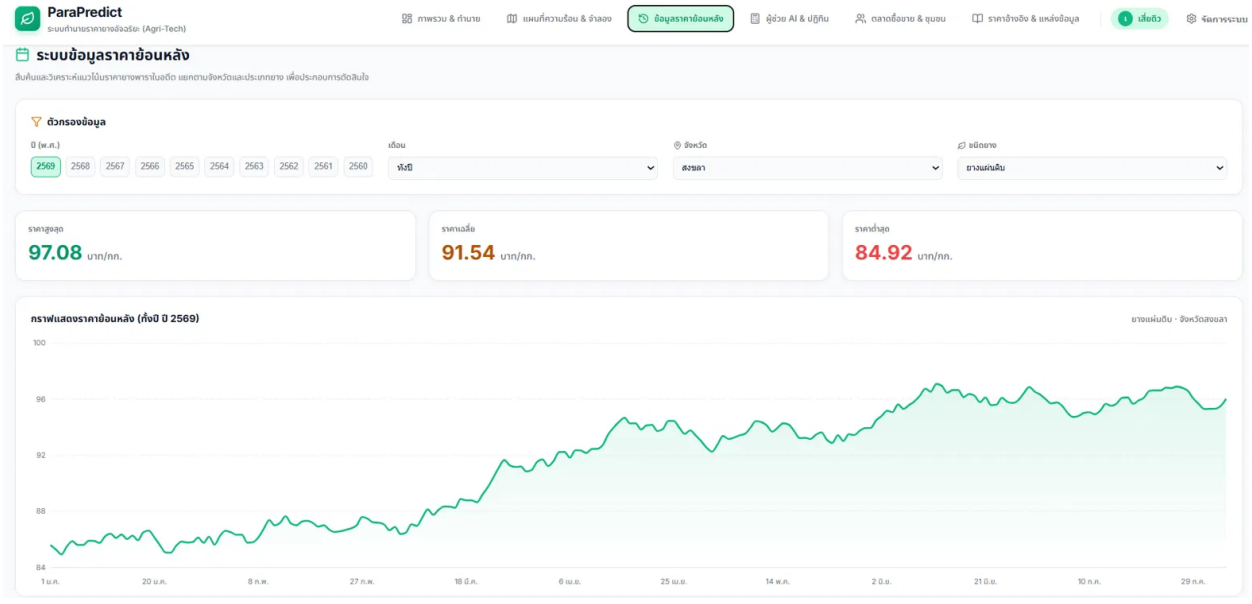


Figure 6: Historical price data view, with year/month/province/rubber-type filters, summary statistics, and a time-series chart of the filtered series.

- **AI Assistant and Planting/Tapping Calendar:** a self-timing assistant that estimates expected revenue from a given rubber type and quantity on hand and a 14-day "Smart Tapping Calendar" that overlays daily rainfall forecasts (sourced from the Open-Meteo API) with recommended tapping days, flagging days likely to be rained out (Figure 7).
- **Marketplace and Community:** a discussion board for farmer questions and a peer-to-peer buy-sell marketplace where users can post rubber-for-sale listings referencing the platform's current reference price (Figure 8).
- **Reference Prices and Data Sources:** a directory of the official price sources underlying the dashboard, the Rubber Authority of Thailand (RAOT), the Thai Rubber Association (TRA), and the Thai Latex Association

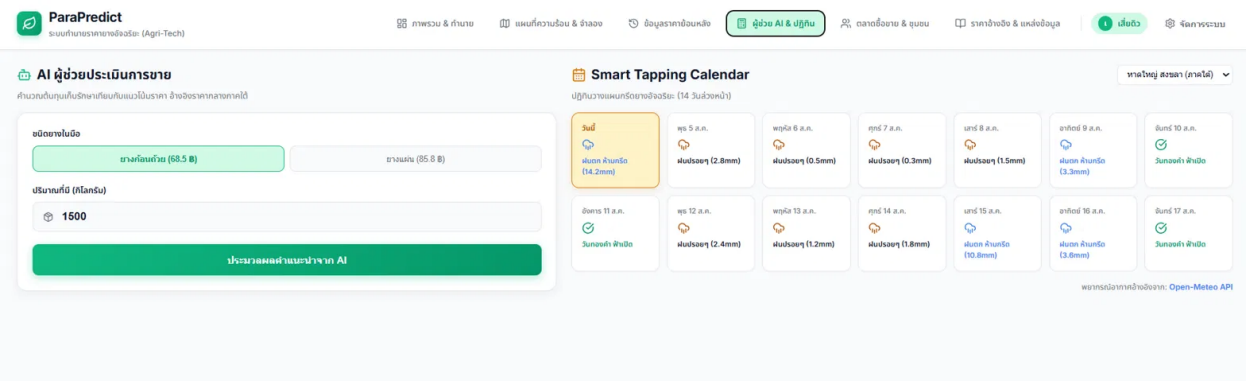


Figure 7: AI sell-timing assistant and 14-day Smart Tapping Calendar with rainfall-based tapping recommendations.

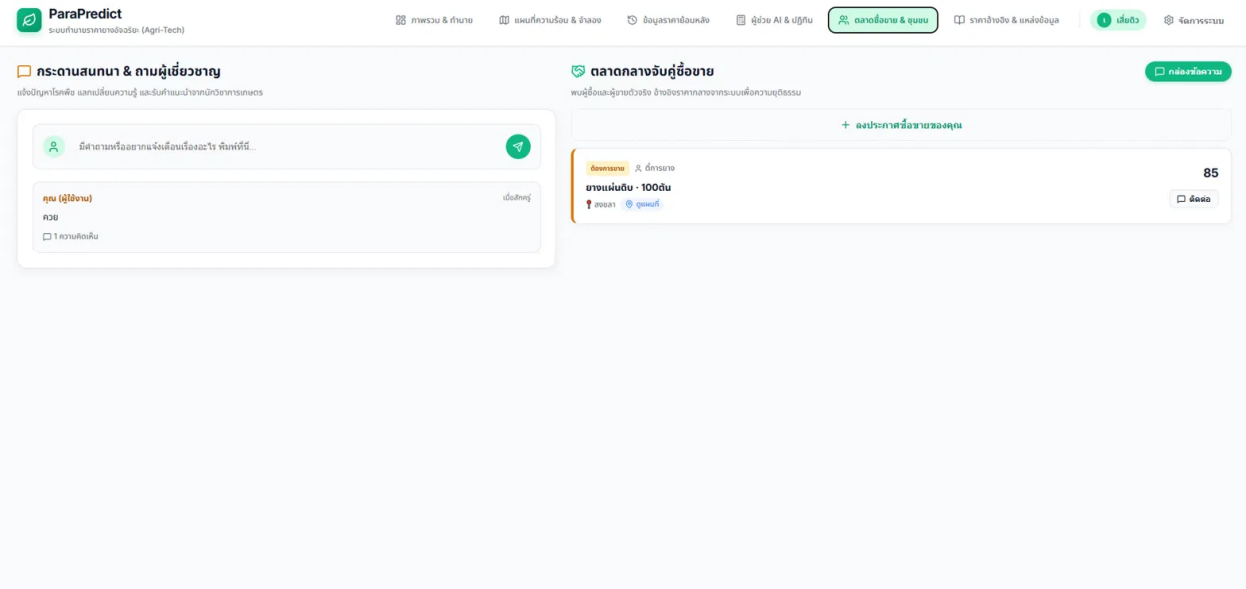


Figure 8: Community discussion board and peer-to-peer marketplace for rubber buy-sell listings.

(TLA) alongside a live sidebar of official reference prices, futures prices (TOCOM, SICOM), crude oil price, the exchange rate, and rainfall forecasts, each with a source attribution and update frequency (Figure 9).

5.3. Administrative Interface

A separate Admin Dashboard, accessible to authorised users, supports (i) direct management of the underlying price database, including manual record editing and bulk import from Excel/CSV files (Figure 10); (ii) an AI model control panel that triggers retraining of the forecasting models on newly collected data (Figure 11); (iii) export of forecast and historical data as Excel/CSV or PDF reports (Figure 12); and (iv) moderation of the community marketplace and discussion board. This administrative layer allows the price database and forecasting models to remain current without requiring direct code changes, which was a practical necessity for treating the dashboard as an ongoing rather than a one-off deliverable.

5.4. From Prototype to Research Tool

Framing the forecasting comparison around an operational dashboard rather than a purely retrospective analysis had a practical consequence for this study's design: because forecasts needed to be produced routinely and displayed with minimal latency, model simplicity and inference speed were treated as first-class considerations alongside

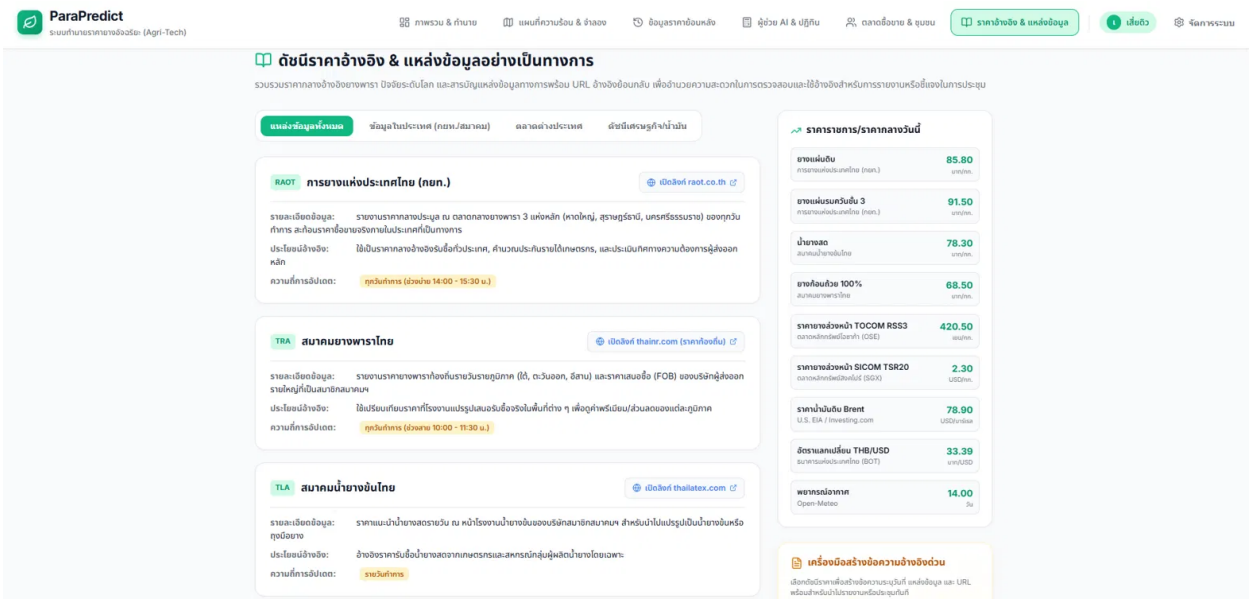


Figure 9: Reference Prices and Data Sources page, listing official price-reporting bodies alongside a live sidebar of reference and futures prices.

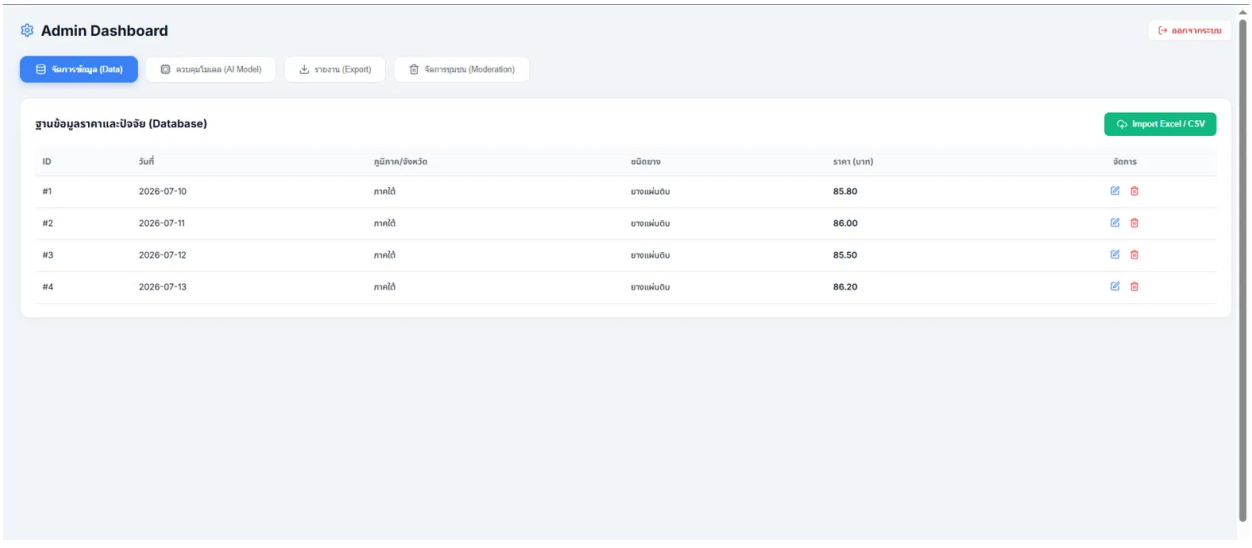


Figure 10: Admin Dashboard: price database management, with manual record editing and bulk Excel/CSV import.

accuracy when selecting candidate models, which is part of the motivation for evaluating three comparatively lightweight approaches, Prophet, SVR, and ARIMA, integrated directly into the production dashboard.

5.5. Implementation Architecture and Current Limitations

Project Structure

The ParaPredict codebase is organized into two main parts: a web dashboard (`rubber-dashboard` directory) and a Python price-scraping system.

Dashboard (frontend). The dashboard is built with React and Vite, styled with Tailwind CSS and general CSS, and uses Firebase to simulate an administrator login. Key files include:

- `App.jsx`: the application's main entry point, defining the navigation tabs and page routing.

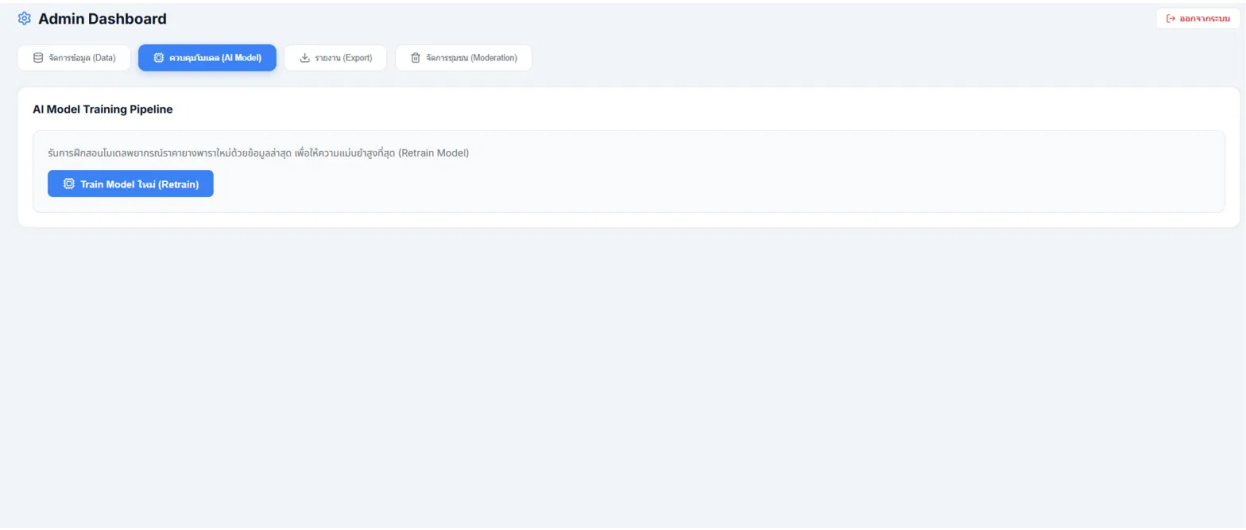


Figure 11: Admin Dashboard: AI model control panel used to trigger retraining of the forecasting models on newly collected data.

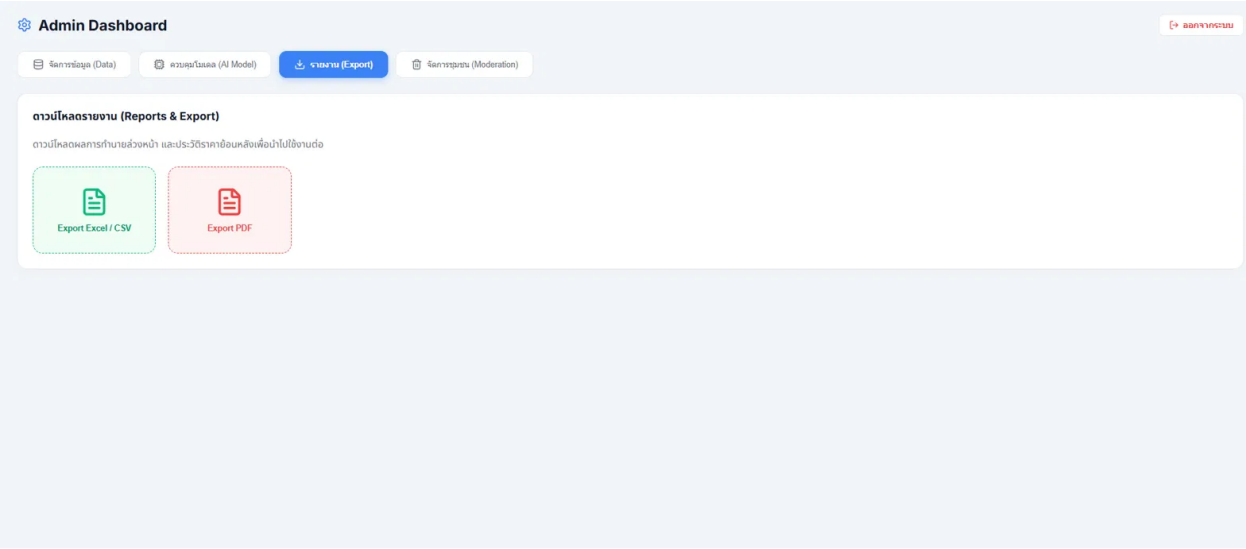


Figure 12: Admin Dashboard: report export interface for downloading forecast and historical data as Excel/CSV or PDF.

- `src/components/CorePrediction.jsx`: renders the trend analysis and forward price forecasts using simulated model curves, together with an AI accuracy comparison table.
- `src/components/HistoricalData.jsx`: supports historical price lookup filtered by province, Buddhist-calendar year, and month.
- `src/components/AdminDashboard.jsx`: the administrator page for logging in, managing the database, and triggering the simulated model-retraining workflow.
- `src/data/mockRealData.js`: defines the simulated price data and behavior, including the mock forecast equations for every AI model shown on the site.

Web scraper (backend). A set of Python scripts at the project root handles data collection:

- `rubber_scraper.py`: scrapes the current day's rubber price from the Thai Rubber Association website (thainr.com) and converts it to JSON.
- `rubber_prices_latest.json`: the scraper's output file containing the latest price, which the dashboard reads as the basis for the current-day price shown across the interface.
- `debug_tables.py`: a diagnostic script that prints the source site's HTML table structure, used to troubleshoot scraping errors.
- `README_scraper.md`: documents how to install the scraper and schedule it via GitHub Actions or a cron job so that it runs automatically every day.

Train/Test Split and Forecast Simulation

Inspection of the source code confirms that the system does not perform an actual train/test split through a machine-learning framework (e.g., scikit-learn, TensorFlow, or PyTorch) on the backend. All predictions and accuracy scores shown in the dashboard are instead computed as a real-time simulation on the frontend (client-side), through the following mechanism:

- **Trend extraction.** The `generateForecast` function in `mockRealData.js` reads the most recent 14 days of price history as a reference, computes a simple linear-regression slope trend, and projects this trend forward 7, 14, or 30 days according to the user's selection.
- **Simulated per-model predictions.** The forecast line for each of the three algorithms evaluated in this study, Prophet, SVR, and ARIMA, as well as additional model curves retained on the interface for visual comparison, is derived from this same 14-day slope trend, scaled by a model-specific simulated coefficient (`slopeAdj`) and offset by a simulated oscillation term generated from the function `Math.sin(i m.noise) 0.8`, so that each model's curve appears visually distinct on the page.
- **Hardcoded evaluation metrics.** The MAE, MAPE, and R^2 values shown in the on-screen comparison table are static, pre-set constants used as illustrative reference targets on the page, rather than being computed from a genuine held-out evaluation. The hardcoded Prophet and SVR values on the dashboard coincide with the genuine offline results reported in Table 1; the values for the additional curves shown alongside them are retained on the interface for visual comparison only and fall outside the formal scope of the evaluation in this study.
- **Admin retraining workflow.** The "Train Model (Retrain)" button in `AdminDashboard.jsx` uses a React timer (`setInterval`) to animate progress from 0% to 100% while displaying simulated status messages such as "Extracting Features..." and "Training Model..." This demonstrates only the intended model-control interface; no underlying mathematical model is actually executed.

This confirms that the dashboard's data-collection and presentation layers are fully operational and driven by genuinely scraped price data, while its forecasting layer remains a client-side illustrative simulation. The comparative results formally reported in this study (Section 4, Table 1) were instead obtained through the offline model-fitting procedure described in Section 3 for Prophet, SVR, and ARIMA, applied directly to the historical price series rather than through the live web interface. Replacing this simulation with a live inference backend, for example, by exposing the fitted models described in Section 3 through a scheduled batch-inference job or a lightweight prediction API that writes into the same JSON structure used by the scraper, is identified as a concrete direction for future work alongside the extensions discussed in Section 6.

6. Conclusion

This study compared three forecasting approaches, Prophet, Support Vector Regression, and ARIMA, for short-term daily natural rubber price forecasting using data from a web-based rubber price monitoring platform. Prophet achieved the lowest forecast error and the highest R^2 among the three models, followed by SVR and then ARIMA, suggesting that decomposition-based models with automatic trend-change point detection may offer an attractive balance of accuracy and interpretability for short-horizon agricultural commodity price forecasting relative to a purely linear statistical model or a generically tuned kernel regressor. These findings offer practical guidance for selecting

a lightweight, deployable forecasting component for operational agricultural price-monitoring systems such as the platform used in this study. Future work should extend the evaluation to forecast horizons and multiple volatility regimes, incorporate exogenous regressors such as crude oil prices and exchange rates, factors repeatedly shown to influence natural rubber prices in other producing countries, and extend the comparison to a broader set of forecasting approaches.

Competing Interests

The authors declare that there are no competing interests.

Acknowledgements

The author gratefully acknowledges the aculty of Science, Energy and Environment, King Mongkuts University of Technology North Bangkok, Rayong, for academic support.

References

[1] Adebayo, P., Ahmed, I., & Oyeleke, K. T. (2025). Comparative analysis of Prophet and ARIMA models for forecasting economic indicators: A case study of personal consumption expenditures. *SSRN*. https://doi.org/10.2139/ssrn.5206738.

[2] Ampountolas, A. (2024). Enhancing forecasting accuracy in commodity and financial markets: Insights from GARCH and SVR models. *International Journal of Financial Studies, 12*(3), 59. https://doi.org/10.3390/ijfs12030059.

[3] Association of Natural Rubber Producing Countries (ANRPC). (2025). Top 10 rubber producing countries (2025). https://www.anrpc.org/news/top-10-rubber-producing-countries-(2025). Accessed 28 July 2026.

[4] Box, G. E. P., & Jenkins, G. M. (1970). *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco.

[5] Chapelle, O., Vapnik, V., Bousquet, O., & Mukherjee, S. (2002). Choosing multiple parameters for support vector machines. *Machine Learning, 46*, 131–159. https://doi.org/10.1023/A:1012450327387.

[6] Chen, Z., Goh, H. S., Sin, K. L., Lim, K., Chung, N. K. H., & Liew, X. Y. (2021). Automated agriculture commodity price prediction system with machine learning techniques. *Advances in Science, Technology and Engineering Systems Journal, 6*(4), 376–384. https://doi.org/10.25046/aj060442.

[7] Cherdchoongam, S., & Rungreunganun, V. (2016). Forecasting the price of natural rubber in Thailand using the ARIMA model. *Applied Science and Engineering Progress, 9*(4). https://doi.org/10.14416/j.ijast.2016.11.005.

[8] Drucker, H., Burges, C. J. C., Kaufman, L., Smola, A., & Vapnik, V. (1997). Support vector regression machines. In *Advances in Neural Information Processing Systems, 9*, 155–161.

[9] Fakthong, T., Pupunja, P., Punjatewakupt, P., & Horayangkura, S. (2025). Forecasting Thai rubber prices: A system dynamics approach to economic and policy impacts. *Social Sciences & Humanities Open, 12*, 102163. https://doi.org/10.1016/j.ssaho.2025.102163.

[10] Fong, Y. C., Khin, A. A., & Lim, C. S. (2020). Determinants of natural rubber price instability for four major producing countries. *Pertanika Journal of Social Sciences & Humanities, 28*(2), 1179–1197.

[11] Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice* (2nd ed.). OTexts, Melbourne.

[12] Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software, 27*(3), 1–22. https://doi.org/10.18637/jss.v027.i03.

[13] Kumar, K. A., Pinto, P., Hawaldar, I. T., Spulbar, C., & Birau, R. (2021). Crude oil futures to manage the price risk of natural rubber: Empirical evidence from India. *Agricultural Economics, 67*(10), 423–434. https://doi.org/10.17221/28/2021-AGRICECON.

[14] Kwarteng, B. S., & Andreevich, P. A. (2024). Comparative analysis of ARIMA, SARIMA and Prophet model in forecasting. *Research & Development, 5*(4), 110–120. https://doi.org/10.11648/j.rd.20240504.13.

[15] Manogna, R. L., Dharmaji, V., & Sarang, S. (2025). Enhancing agricultural commodity price forecasting with deep learning. *Scientific Reports, 15*, 20903. https://doi.org/10.1038/s41598-025-05103-z.

[16] Mathew, S., & Murugesan, R. (2020). Indian natural rubber price forecast: An autoregressive integrated moving average (ARIMA) approach. *The Indian Journal of Agricultural Sciences, 90*(2), 418–422. https://doi.org/10.56093/ijas.v90i2.103067.

[17] Menculini, L., Marini, A., Proietti, M., Garinei, A., Bozza, A., Moretti, C., & Marconi, M. (2021). Comparing Prophet and deep learning to ARIMA in forecasting wholesale food prices. *Forecasting, 3*(3), 644–662. https://doi.org/10.3390/forecast3030040.

[18] Nyondo, P., & Varghese, R. (2024). Forecasting natural rubber prices using commodity market indicators: A machine learning approach. *International Journal of Revenue Management, 14*(3), 221–252. https://doi.org/10.1504/IJRM.2024.142325.

[19] Pinitjitsamut, M. (2026). Dynamic price transmission from SHFE to Thai rubber markets: A cointegration–ECM and machine-learning analysis. *Economies, 14*(1), 9. https://doi.org/10.3390/economies14010009.

[20] Pinitjitsamut, M. (2026). Multi-scale forecasting of natural rubber prices using VMD-augmented BiLSTM: A hybrid architecture ablation study. *Forecasting, 8*(3), 43. https://doi.org/10.3390/forecast8030043.

[21] Rahim, H. A. (2023). A robust AR–GARCH model with Huber weighting for sustainable forecasting of standard Malaysian rubber prices under outlier contamination. *Preprint/Working Paper*.

[22] Shaharuddin, H. I. B., Azmi, N. N. B. K., & Shamsudin, F. I. M. B. (2024). Modelling and forecasting natural rubber price in Malaysia: Comparison between HES, ARIMA and ANN. In W. F. Wan Yaacob, Y. B. Wah, & O. U. Mehmood (Eds.), *Decision Mathematics, Statistical Learning and Data Mining*, 145–160. Springer, Singapore. https://doi.org/10.1007/978-981-97-3450-4_10.

[23] Shuvo, A. A. (2025). Comparative analysis of regression models for commodity price prediction. *SSRN*. https://doi.org/10.2139/ssrn.5248271.

[24] Smola, A. J., & Schölkopf, B. (2004). A tutorial on support vector regression. *Statistics and Computing, 14*(3), 199–222. https://doi.org/10.1023/B:STCO.0000035301.49549.88.

[25] Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician, 72*(1), 37–45. https://doi.org/10.1080/00031305.2017.1380080.

[26] Vapnik, V. N. (1995). *The Nature of Statistical Learning Theory*. Springer-Verlag, New York. https://doi.org/10.1007/978-1-4757-2440-0.

[27] Vinayana, A. K., Prasad, L. V., & Anuradha, P. S. C. (2019). Impact of crude oil price changes on natural rubber industry in India. *International Journal of Advanced Multidisciplinary and Technical Engineering Studies*.

[28] World's Top Exports. (2025). Natural rubber exports by country. https://www.worldstopexports.com/natural-rubber-exports-country/. Accessed 28 July 2026.

[29] Yenidoğan, I., Çayır, A., Kozan, O., Dağ, T., & Arslan, Ç. (2018). Bitcoin forecasting using ARIMA and Prophet. In *2018 3rd International Conference on Computer Science and Engineering (UBMK)*, 621–624. https://doi.org/10.1109/UBMK.2018.8566476.

[30] Zahari, F. Z., Khalid, K., Roslan, R., Sufahani, S., Mohamad, M., Rusiman, M. S., & Ali, M. (2018). Forecasting natural rubber price in Malaysia using ARIMA. *Journal of Physics: Conference Series, 995*, 012013. https://doi.org/10.1088/1742-6596/995/1/012013.