Multivariate integration of time series with ML for corn price forecasting in Colombia.

Adelaida Ojeda Beltran[a,*], Mario E. Suaza-Medina[b], F. Javier Zarazaga-Soria[b], Emiro De-La-Hoz-Franco[c] and José Escorcia-Gutierrez[c]

[a]*Faculty of Economic Sciences, Universidad del Atlántico, Puerto Colombia - Atlántico, Colombia*
[b]*Aragon Institute of Engineering Research (I3A), Universidad de Zaragoza, Zaragoza, Spain.*
[c]*Department of Computational Science and Electronics, Universidad de la Costa, CUC, Barranquilla, Colombia*

ARTICLE INFO

Keywords:
Multivariate integration
Time series
Machine Learning
Forecasting
Agriculture

ABSTRACT

The volatility of corn prices poses a significant challenge for both producers and policymakers. This study proposes a hybrid model that combines Extreme Gradient Boosting (XGBoost) and Light Gradient Boosting Machine (LightGBM), optimized through Particle Swarm Optimization with Cuckoo Search (PSO-CS), for accurate corn price forecasting. The approach integrates multivariate time series data, including local prices from the Atlántico market and international futures prices from the Chicago Board of Trade (CBOT). Empirical Mode Decomposition (EMD) is applied to enhance signal clarity and improve model performance. Model performance is assessed through sensitivity analysis and statistical comparison using the Diebold-Mariano (DM) test. The results demonstrate that the proposed ensemble outperforms both individual models and neural network combinations, achieving a Mean Absolute Percentage Error (MAPE) of 2.06

1. Introduction

Agriculture plays a fundamental role in the economies of developing countries due to its contribution to domestic production and employment, serving as a source of livelihood for approximately 2.57 billion people and acting as a cornerstone of food security, according to data from the Food and Agriculture Organization of the United Nations (FAO, 2020). However, despite technological advancements, the sector still faces challenges, particularly for small scale producers. In this context, the Organisation for Economic Cooperation and Development (OECD) and the FAO, in their *Agricultural Outlook 2021–2030* report, highlight the rising global demand for agricultural products especially cereals such as maize driven by shifting consumption patterns and the growing share of animal-based products in global diets (OECD and FAO,2021). Maize and protein meals are projected to account for more than 66% of total forage consumption by 2030, while the share of maize used for biofuels is expected to decline from 15.8% to 13.7%. However, the anticipated increase in global maize production and the release of strategic reserves are expected to meet the increasing demand for food, feed, and biofuels, potentially leading to a downward trend in the international reference price of maize.

Yellow maize is a strategic crop in Colombia, as it serves as the foundation for the production of animal feed, an essential input for the livestock sector, particularly the poultry industry, which supplies chicken meat and eggs, both of which are fundamental to the Colombian diet and household consumption (Arbeláez et al., 2024). Between 2010 and 2023, national mechanized yellow maize production increased by 49. 3% (from 0.5 to 0.8 million tons), while imports increased by 71.6% (from 3.4 to 5.9 million tons) (WFP, 2024). This growth has been driven by increasing demand for poultry products, alongside improvements in production processes and mechanization (CIMMYT and CIAT, 2019). However, domestic production remains insufficient, forcing the country to meet approximately 80% of its demand through imports, mainly from the United States, Brazil, and Argentina countries with higher yields and more competitive production costs (UPRA, 2022). As a result, Colombia faces a deficit of nearly 4.6 million tons of yellow maize, increasing its vulnerability to international price fluctuations and directly impacting the cost of poultry and other staple food products (BMC, 2023).

To address the volatility of agricultural product prices and support strategic decision-making, predictive models based on advanced time series analysis techniques have been developed. These models are capable of handling the non-stationary and nonlinear nature of agricultural data (Zhao, 2021). Among traditional statistical approaches, models such as ARIMA, ARCH, and GARCH have been widely used, although they present notable limitations particularly the need for large volumes of historical data that meet strict statistical assumptions (Nafkha and Suchodolska, 2024). In contrast, machine learning (ML) techniques, particularly those based on artificial neural networks (ANNs), have shown greater reliability in modeling nonlinear relationships (Hiyam et al., 2024), outperforming statistical methods by capturing complex patterns and uncovering hidden interactions within market

*Corresponding author
Email address: adelaidaojeda@mail.uniatlantico.edu.co (A.O. Beltran)
ORCID(s): 0000-0003-4530-5644 (A.O. Beltran); 0000-0002-2112-3478 (M.E. Suaza-Medina); 0000-0002-6557-2494 (F.J. Zarazaga-Soria); 0000-0002-4926-7414 (E. De-La-Hoz-Franco); 0000-0003-0518-3187 (J. Escorcia-Gutierrez)

dynamics (Adebiyi et al., 2014). Within this category, recurrent neural networks (RNNs) are especially effective for time series analysis, as their feedback connections allow them to retain past information—unlike conventional ANNs, which fail to adequately capture the temporal effects of significant historical events.

Various ML techniques have been successfully applied to agricultural time series forecasting. Multilayer Perceptron (MLP) networks are valued for their ability to approximate nonlinear functions and process multivariate inputs, though they require the definition of meaningful input output mappings (Lee and Xia, 2024). Convolutional Neural Networks (CNNs) complement these advantages by offering pattern recognition capabilities, but they fall short in capturing temporal dependencies. Meanwhile, Support Vector Regression (SVR), based on the principle of structural risk minimization and kernel-based nonlinear mapping, has been successfully applied to complex agricultural contexts (Zheng et al., 2020). Likewise, the Extreme Gradient Boosting (XGBoost) algorithm incorporates bagging and boosting techniques to optimize differentiable loss functions (Wu et al., 2024). Finally, Long Short-Term Memory (LSTM) networks, specifically designed to capture temporal dependencies and address vanishing/exploding gradient problems, have demonstrated strong predictive performance. However, their training process demands significant computational resources and careful parameter tuning (Chung and Shin, 2018).

Recently, a growing trend has emerged in combining ML algorithms with time series decomposition methods and optimization techniques to enhance model accuracy and robustness in agricultural forecasting. Decomposition methods such as STL (Seasonal and Trend decomposition using Loess) and EMD (Empirical Mode Decomposition) allow time series to be broken down into simpler components trend, seasonality, and residuals facilitating the modeling of complex patterns when used in conjunction with ML techniques. Simultaneously, optimization algorithms such as Particle Swarm Optimization (PSO), Genetic Algorithms (GA), and Cuckoo Search (CS) have been employed to fine-tune model hyperparameters and architectures, improving performance in nonlinear environments. These hybrid combinations have led to the emergence of highly efficient models, some of which are classified as state of the art (SOTA) algorithms, representing the most advanced solutions available. Notable examples include LSTM-EMD-PSO and XGBoost-STL-GA, which integrate temporal modeling, adaptive decomposition, and optimized configuration to achieve higher accuracy in maize price forecasting.

Nevertheless, despite these advancements, a significant gap remains in the integration of multivariate data combining both national and international sourcesparticularly within the context of developing countries. In the specific case of Colombia's maize market, most existing studies focus exclusively on local price series, overlooking the predictive potential of incorporating variables such as international futures prices (e.g., CBOT), climate indicators, or seasonal trends.

This paper is organized as follows: Section 2 details the contributions derived from the development of the article. Section 3 presents a literature review to contextualize the proposed work. Section 4 describes the methodology used. Section 5 outlines the preliminary concepts relevant to the development of the study, while Section 6 presents the results and discussion. Finally, Section 8 closes the document with the conclusions.

2. Contributions

The primary objective of this research is to predict the price of technified yellow corn in Colombia using a time series of historical prices from a market in the Colombian Caribbean region. Additionally, the study aims to assess how the inclusion of corn futures contract prices from the Chicago Board of Trade (CBOT) can enhance the performance of the predictive model. Specifically, it seeks to determine whether incorporating these futures prices improves the model's ability to forecast price fluctuations. This research contributes to strengthening strategic decision-making processes for market participants by providing them with a more accurate and reliable predictive model. The key contributions of this article are as follows:

- The application of the Empirical Mode Decomposition (EMD) method to the time series of corn prices in Colombia, enabling the identification of the optimal intrinsic mode functions (IMFs) that best capture price fluctuations.

- The development of predictive models was carried out using various algorithms to forecast corn prices in Colombia, based on historical price data. The selected models include: Fully Connected Network (FCN), Recurrent Neural Network (RNN), eXtreme Gradient Boosting (XGBoost), and Light Gradient Boosting Machine (LightGBM). The optimal model was identified based on its performance according to standard evaluation metrics.

- Another contribution of this study is the systematic use of evaluation metrics to validate the predictive performance of the models. The metrics applied were Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and the Coefficient of Determination (R^2). These measures enabled a rigorous comparison of both individual and ensemble models under univariate and multivariate settings. .

- Development of a novel multivariate forecasting architecture that employs an ensemble strategy optimized through Particle Swarm Optimization (PSO) to integrate heterogeneous predictive models. This approach dynamically assigns optimal weights to Fully Connected Networks (FCN), Recurrent

Neural Networks (RNN), XGBoost, and LightGBM based on their individual performance, with the aim of minimizing overall forecasting error. By incorporating both local and external variables into a unified multivariate framework and enhancing it through PSO-based optimization, the model achieves higher predictive accuracy and robustness.

- The validation of the top-performing predictive model using independent databases from two different markets in Colombia, distinct from the original training set. This validation is conducted by applying the model without adjustments to its parameters, aiming to assess its generalization capability and robustness in contexts that differ from the original training environment.

3. Related Works

The analysis and forecasting of agricultural commodity prices has attracted a great deal of interest due to the complexity and volatility of the futures market. In this context, various model combination strategies have been proposed to improve forecast performance compared to individual models.

Several studies have explored the use of hybrid approaches and decomposition techniques to improve the forecasting accuracy of non-stationary and non-linear time series. (Ji et al., 2022) propose a mixed model combining ARIMA with PLS regression to predict agricultural commodity prices, achieving more accurate weekly forecasts by integrating spatial and temporal factors across markets. Similarly, (Bahri and Vahidnia, 2022) employs a combination of empirical decomposition techniques (SEEMD) and neural networks (CNN and LSTM) to predict complex signals with high accuracy, highlighting the ability to handle fluctuations in nonstationary data. In the energy domain, (Bedi and Toshniwal, 2020) use Variational Mode Decomposition (VMD) together with LSTM networks to forecast energy demand, outperforming traditional models by capturing seasonal and historical patterns in the data.

On the other (Zhang et al., 2020), a new model selection framework incorporating time series features and forecast horizons is proposed. Twenty-nine features are used to represent agricultural product prices, and three intelligent models are specified as candidate forecasting models: namely, Artificial Neural Network (ANN), Support Vector Regression (SVR), and Extreme Learning Machine (ELM).s. For its part (Suaza et al., 2023), it analyzes how data acquisition delays impact price prediction systems when different forecasting algorithms are used. The study (Chaowalit et al., 2025) employs multivariate time series and deep learning techniques, such as CNN and LSTM, to predict the import price of soybean meal in Thailand, a key input in the animal feed industry. The Bidirectional LSTM (Bi-LSTM) model proved to be the most efficient, helping importers plan costs by considering shipping time. In addition, future improvements are suggested by incorporating variables such as weather and geopolitical events to optimize predictions.

Time series forecasting has been studied through various decomposition approaches and hybrid models. (Bahri and Vahidnia, 2022) analyzed the effectiveness of smoothing ensemble empirical mode decomposition (SEEMD) combined with LSTM and CNN neural networks, achieving outstanding results in noisy and nonlinear time series. Similarly, (Li et al., 2023) evaluated the adaptive Fourier decomposition (AFD) method compared to EMD and FDM, concluding that AFD offers greater sensitivity to peaks in financial time series, enhancing trend detectiona and structural changeidentification. (Liu et al., 2020) highlighted the performance of VMD in analyzing mill vibration signals, demonstrating its ability to mitigate aliasing and boundary effect issues present in EMD. (Yang and Yang, 2020) proposed hybrid approaches based on EEMD and neural networks, showing superior performance by integrating multiple learning algorithms.

Finally, (Bedi and Toshniwal, 2020) designed a hybrid model that combines VMD with LSTM networks for energy demand forecasting, achieving significantly higher accuracy than traditional models and AI-based techniques. (Wang et al., 2022) used a hybrid ap approach based on the artificial bee colony (ABC) algorithm to predict agricultural futures prices, combining decomposition techniques (SSA, EMD, VMD) with models such as ARIMA and neural networks, achieving significant improvements in accuracy, (Zhang and Tang, 2024) proposes a novel VMD-SGMD-LSTM model that combines enhanced quadratic decomposition techniques with an artificial intelligence framework. Ultimately, artificial intelligence plays a crucial role in modern agriculture.(Kundu et al., 2022) propose the MaizeNet model to detect maize diseases and estimate yield losses with high accuracy. Meanwhile, Sachithra and (Sachithra and Subhashini, 2023) highlight how AI promotes agricultural sustainability by optimizing resources and reducing environmental impacts. In the data preprocessing stage, the VMD is applied to decompose the original futures price data, while the Second Generation Mode Decomposition (SGMD) is used for further refinement of the remaining components.

Studies conducted by (Safari et al., 2020) analyze the performance of the EMD method regarding methods for decomposing time series into simpler components to improve prediction accuracy.

In summary, although state-of-the-art models such as Transformers and hybrid architectures with attention mechanisms have demonstrated important advances in agricultural price forecasting in international contexts, their effectiveness has been validated primarily under conditions involving long, clean, and multisource time series, conditions that differ significantly from the realities of the Colombian agri-food system. Studies such (Zeng et al.,

2023) and (Guindani et al., 2024b) highlight not only the potential predictive performance of these models but also their high computational requirements and the complexity of fine-tuning, which limits their applicability in data and resourceconstrained environments. In contrast, the present study adopts a more robust, interpretable, and adaptable architecture based on EMD decomposition, decision tree models, and conventional neural networks aligned with the structure of the dataset used: weekly time series with structural noise and strong seasonality, typical of the Colombian maize market. This methodological choice not only ensures greater technical efficiency but also provides a better fit with the operational conditions of the country, without ruling out the future integration of more complex models as data ecosystems and technical capacities evolve.

To build upon these findings, this research conducted a comprehensive review of existing studies on agricultural price prediction. Table 1 provides a detailed summary, including crop types, data sources, models employed, algorithms applied, and evaluation metrics used to assess model performance.

4. Methodology

The proposed methodological architecture comprising Empirical Mode Decomposition (EMD), supervised learning models, and optimization through Particle Swarm Optimization Cuckoo Search (PSO-CS)was designed to address the specific challenges of agricultural price forecasting systems in Colombia. In particular, the maize market faces structural constraints such as high seasonality, the influence of exogenous variables, and limitations in both the availability and quality of historical data (Arbeláez et al., 2024) EMD was selected for its ability to decompose nonlinear time series into interpretable components, thereby reducing noise and enhancing the predictive signal (Safari et al., 2020). Additionally, models such as LightGBM, XGBoost, FCN, and RNN were chosen for their ability to capture complex patterns in multivariate datasets, while PSO-CS ensures an efficient integration of these models by dynamically adjusting their weights based on relative performance. This framework is not merely a technical reproduction of existing literature but rather a configuration tailored to the real-world constraints and needs of the Colombian agricultural sector.

Figure 1 presents the proposed methodological framework for corn price forecasting. The process is structured into three main phases: (1) prediction using a univariate time series, (2) integration of multivariate time series data, and (3) model validation and Figure 2 presents the architecture developed for forecasting corn prices in Colombia, based on the integration of multivariate time series and Figure 2 shows the forecasting architecture for corn prices in Colombia, structured in three phases: (1) univariate prediction using EMD and ML models, (2) multivariate modeling with CBOT futures as exogenous input, and (3) optimization of ensemble forecasts using the PSO-CS metaheuristic to improve accuracy and robustness.

Phase 1, prediction from a univariate series. This first phase is detailed in the following steps.

Step 1: Data preparation.

In this process, data cleaning is performed to remove outliers, missing values, and redundant variables. Then, the transformed data is normalized and structured for use in the prediction model.

Step 2: Data decomposition.

The data is decomposed into simple and independent components using EMD decomposition techniques. For this phase, the database of historical corn prices in Colombia, generated by the Gran Central de Abastos del Caribe (GRANABASTOS) detailed in the table 2

Step 3: Individual forecasting.

The components are fed into the individual forecast models XGBoost, LightGBM, FCN, and RNN, and individual forecasts are generated.

Step 4: Combination of forecasts.

The PSO-CS weight allocation method is used to determine the optimal weights of the individual models with different combination strategies, and then the individual predictions are weighted to obtain the combined results for each of the

Step 5: Model Combination

At this stage, neural networks (FCN and RNN) are integrated with each other, and tree-based models (XGBoost and LightGBM) are combined among them.

Phase 2: Multivariate Data Integration: Data1-corn and Data2-corn

Step 1:Verification of correlation and cointegration between series

The methodology presented in Step is extended by integrating data on corn futures contract prices from the CBO. This dataset is structured as a sequence that follows the same timeline as the original historical corn price series in Colombia, provided by GRANABASTOS, transforming it into a multivariate time series model.

Step 2: : Individual Multivariate Predictive Models.

In this step, individual predictions were generated using the RNN, FCN, XGBoost, and LightGBM models on the multivariate time series constructed from the integration of Data1 and Data2. Each model was independently trained and validated on the multivariate dataset in order to assess its individual performance before being considered for combination in later stages of the study.

Step 3:Statistical Significance of the Diebold-Mariano Test (DM).

In this step, the Diebold-Mariano (DM) test was applied to evaluate whether the differences in predictive accuracy between individual models were statistically significant. The test compared the forecast errors pairwise to determine if one model consistently outperformed another

Step 4: Combination of Models for Data1-Corn and Data2-Corn.

Table 1
Summary of recent studies on crop price forecasting using time series and ML. Comparison of crops, input variables, data sources, forecasting algorithms, best-performing models, and evaluation metrics across referenced publications.

Ref	Year	Crop	Input Variables	Type of Source	Algorithms Used	Best Performing Algorithm	Evaluation Measure
(Sediyono et al., 2025)	2025	Multi-commodity (maize, rice, wheat)	Price history, weather, seasonality, volatility indices	Public + satellite data	Transformer, LSTM-VAE + Attention	Attention-boosted Transformer + LSTM-VAE ensemble	RMSE, MAPE, Directional Accuracy
(Guindani et al., 2024a)	2025	Multiple (maize included)	Time series data + economic indicators	Public databases (global)	CNN, LSTM, TCN, XGBoost, hybrid DL models	Hybrid CNN + TCN + LSTM	RMSE, MAE, MAPE, R^2
(Zhang et al., 2025)	2025	Avocado	Time series + exogenous variables (climate, region)	Public datasets (unspecified)	TCN, MLP, Attention Mechanism	TCN + MLP + Attention (hybrid)	MSE, RMSE
(Yoon et al., 2025)	2025	Tomato and Apple	24 environmental Climatological Variables	Public data (KREI OASIS)	LSTM (with time delay), SHAP for explainability	LSTM with time delay	Nash-Sutcliffe Efficiency (NSE), MAE, RMSE, SHAP values R^2
(Rana et al., 2024)	2024	Spinach	Time series Historical	Historical-Government Data	Regressive Moving Average (ARIMA), Random Forest (RF), Long-Short-Term Memory (LSTM)	Long-Short-Term Memory (LSTM)	MAE, MSE, RMSE, Squared correlation coefficient R^2
(Sari et al., 2024).	2024	Corn, Sugar, Soybean, Rice, Oat, Cotton, Coffee, Cocoa, Soybean oil, Lumber	Time series Historical	Historical-Government Data	Extreme Learning + Genetic Algorithm, Long Short-Term, Memory + Genetic Algorithm + autoregressive integrated moving average model	Extreme Learning + Genetic Algorithm	RMSE, MAPE, MAE
(Jin and Xu, 2024)	2024	Corn	Time series Historical	Wholesale market in China	ARIMA and AR-GARCH, Artificial Neural Networks (ANN), Support Vector Regression (SVR), and Gaussian Process Regression (GPR)	Gaussian Process Regression (GPR)	RMSE, MAE
(Garai et al., 2023).	2023	Onion	Time series Historical	Historical-Government Data	ARIMA-GARCH, ANN-SVR	ARMA, RF, SVR	RMSE, MAE, MAPE
(Mohanty et al., 2023)	2023	Corn	Time series Historical	Department of Agriculture, United States	ARIMA, Statistical Regression (SR), Decision Tree Regressor (DTR), Random Forest (RF), K-nearest, neighbor regressor (KNN) and kernel ridge(KR)	Decision Tree Regressor (DTR)	RMSE, MAE
(Purohit et al., 2021)	2023	Cotton	Time series: Historical Prices	Historical-Government Data	Linear Regression, Bayesian Linear Regression, Decision Trees, Random Forest(RF)	Decision Trees (DT)	MSE, RMSE, MAE, MAPE
(Banerjee et al., 2022)	2022	Soybean, Corn	Time series: Historical Prices	Daily Price Data—Chicago Board of Trade (CBOT)	Artificial bee colony algorithm (ABC), ARIMA, SVR, (RNN)-Gated Recurrent	Artificial Bee Colony Algorithm (ABC)	RMSE, MAE, MAPE
(Wang et al., 2022)	2021	Cardamom	Time series Historical	Historical-Government Data	Multiple linear regression (MLR), ARIMA	Multiple linear regression (MLR)	RMSE, R^2
(Adhikari et al., 2021)	2021	Tomato, Onion, Potato	Time series: Historical Prices	Historical Information-Government-Report of Horticultural Statistics Division	(Additive-ETSSVM, Additive-ETS-LSTM)-(Multiplicative-ETS-ANN, Multiplicative-ETS SVM, Multiplicative-ETS-LSTM	Additive-ARIMA-ANN-SVM Multiplicativo-ETS	RMSE, MAE, MAPE
(Deepa et al., 2023)	2021	Corn	Time series: Historical Prices	Historical data-Governmental-Time orn futures prices, traded on CBOT.	K-means, k-nearest neighbor (kNN) algorithm	K-means, k-nearest neighbor (kNN) algorithm	RMSE, MAE, MAPE
(Li et al., 2021)	2020	Grape	Time series Historical	Historical-Government Data	Least Squares Support Vector Machine (LSSVM), Extreme Learning Machine (ELM)	Least Squares Support Vector Machine (LSSVM), Extreme Learning Machine (ELM)	RMSE, MAE, Improvement Rate (IR)

In this step, model combination was carried out using the PSO-CS algorithm to create hybrid ensembles. Specifically, the models XGBoost and LGBM were combined into one ensemble, and FCN and RNN into another. The goal was to leverage the complementary strengths of each model, with PSO-CS optimizing the weight assignment in each ensemble to improve overall forecasting accuracy.

Step 5: Sensitivity and Stability Analysis of the Predictive Ensemble.

In this step, a sensitivity and stability analysis of the predictive ensemble was conducted by evaluating the model's behavior under different hyperparameter configurations of the PSO-CS algorithm. Various settings were tested to assess the consistency of performance and the stability of the weight distribution assigned to each base model. This analysis helps validate the robustness of the ensemble under diverse optimization scenarios.

Step 6: Uncertainty Quantification and Model Reliability.

This step assesses the robustness of the model by estimating the prediction uncertainty through block bootstrap. It provides a 95% confidence interval around the ensemble mean, allowing evaluation of the statistical reliability of the forecasts.

Phase 3: Model Validation: Multivariate Integration

Step 1: Out of Time Validation.

This step evaluates the model's performance using entirely new and different data not involved in the training or validation stages.

Step 2: Model Validation Using Data3-Corn and Data4-Cor.

In this phase, the best-performing models, both individual and ensemble, were validated using two new datasets from different markets: Data3-Corn and Data4-Corn. No additional parameter tuning was applied, in order to evaluate the models' generalization capacity. The objective was to assess their performance in previously unseen real-world scenarios.

Step 3 Comparative Analysis with Previous Studies Using the Same Time Series.

This step consisted of comparing the proposed ensemble model with previous state-of-the-art approaches applied to the same CBOT time series. The results reported by Wang et al. (2022) and Zeng et al. (2023) were used as reference benchmarks.

4.1. Hypothesis Analysis for Model Selection

The selection of predictive models was based on two fundamental hypotheses. Hypothesis H1 states that decision tree algorithms (XGBoost and LightGBM) offer greater accuracy and stability in agricultural contexts such as Colombia, where data tend to be noisy, nonlinear, and temporally limited (Bedi and Toshniwal, 2020). Hypothesis H2 suggests that although neural networks (FCN and RNN) are capable of capturing sequential patterns, their performance is negatively affected in short and highly seasonal time series when adequate regularization mechanisms are not applied (Chung and Shin, 2018).

5. Model descriptions

This section describes the models compared in the experiments.

5.1. Empirical modal decomposition

In this study, Empirical Mode Decomposition (EMD) was employed as the primary preprocessing technique prior to predictive modeling. Its selection is supported by a prior comparative analysis in which various decomposition methods were evaluated using the same time series of corn prices, including EMD, Variational Mode Decomposition (VMD), Singular Spectrum Analysis (SSA), Ensemble Empirical Mode Decomposition (EEMD), Wavelet Transform (WT), and a baseline case without decomposition. The study concluded that EMD and SSA achieved the best predictive performance when combined with machine learning models such as LightGBM, XGBoost, and neural networks. The findings were published in the article titled Hybridization of Variational Mode Decomposition and Singular Spectrum Analysis for Corn Price Forecasting, included in the volume *Advances in Computing and Data Sciences* (Springer CCIS series) (Ojeda-Beltran et al., 2024). Therefore, EMD was selected for its proven ability to extract nonlinear and nonstationary components from complex agricultural time series, which are often affected by structural noise and seasonal variability.

Nevertheless, EMD also presents important limitations: high sensitivity to noise, dependence on interpolation techniques, and the absence of a formal mathematical foundation, which can compromise the consistency of the extracted modes. In scenarios with a low signal-to-noise ratio or when more stable spectral separation is required, methods such as VMD or WT may offer better performance. However, these approaches often demand more complex calibration and tend to be less flexible, limiting their applicability in settings such as the Colombian maize market, which is characterized by irregular signal structures and noisy data.

The EMD first introduced by Huang et al. (1998), is an adaptive time-frequency signal processing method designed to decompose nonlinear and nonstationary signals into a finite set of Intrinsic Mode Functions (IMFs) and a residual component. Each IMF represents a distinct frequency characteristic of the original signal and has a compact Fourier spectrum. The decomposition process, known as sifting, involves iteratively identifying local extrema, computing upper and lower envelopes via cubic spline interpolation, averaging these envelopes, and subtracting the mean to extract each IMF (Yang and Fan, 2022). Once all IMFs are extracted, the residual captures the long-term trend:

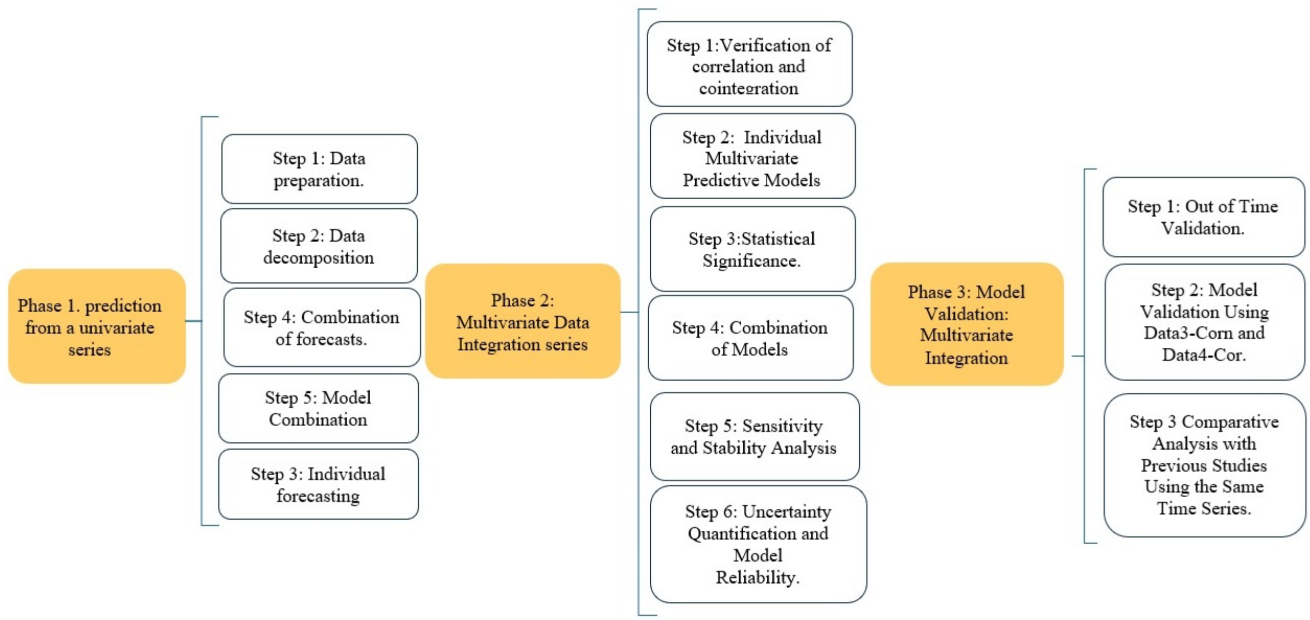

Figure 1: Proposed methodological framework for corn price forecasting The process is divided into three phases: (1) prediction based on a univariate time series, (2) multivariate data integration, and (3) model validation.

$$X(t) = \sum_{k=1}^{N} c_k(t) + r_N(t) \tag{1}$$

where $c_k(t)$ denotes the k-th IMF and $r_N(t)$ the final residual.

EMD is widely used for forecasting due to its capacity to separate stationary and nonstationary components (Lahmiri, 2017). However, it presents limitations such as sensitivity to noise, reliance on interpolation methods, and lack of a formal mathematical foundation (Yang and Chen, 2019). Some studies recommend discarding the first IMF (IMF_1) due to its potential to introduce forecast errors (Yu et al., 2008), while others suggest training it separately using deep learning approaches. Alternative formulations with improved robustness, such as the one proposed by Park et al. (2011), address these limitations by reducing sensitivity to stopping criteria and local fluctuations.

5.2. Fully Connected Networks

FCN are neural networks in which each neuron in one layer is connected to every neuron in the next layer. They are classic deep learning models primarily used for classification and regression tasks.

In these networks, each neuron in one layer is connected to every neuron in the next layer. They are highly flexible and commonly used for tasks such as classification, regression, and pattern recognition. The key advantage of FCN is their ability to capture complex relationships between input features. However, their main limitation is the lack of explicit handling of temporal or spatial relationships within the data. Mathematically, a FCN operates by calculating a linear combination of input values, followed by the application of a nonlinear activation function:

$$z = Wx + b, \quad a = f(z) \tag{2}$$

Here, W is the weight matrix, b is the bias vector, x is the input vector, and f is the activation function. The output a is obtained by applying f to the linear transformation z.

5.3. Recurrent Neural Networks

They are designed to process sequential data, such as time series or text. Unlike FCNs, RNN feature recurrent connections that enable them to retain information from previous states, making them particularly well-suited for tasks where temporal context is essential.

The operation of an RNN is based on iterating through time, updating a hidden state h_t at each time step:

$$h_t = f(Wx_t + Uh_{t-1} + b) \tag{3}$$

Here, x_t is the input at time step t, h_{t-1} is the hidden state from the previous time step, W and U are weight matrices, b is the bias vector, and f is a nonlinear activation function (e.g., *tanh* or *ReLU*).

Variants of RNN, such as *Long Short-Term Memory* (LSTM) networks and *Gated Recurrent Units* (GRU), include mechanisms to mitigate issues such as vanishing or exploding gradients, thereby improving the capture of long-term dependencies.

5.4. XGBoost and LightGBM

XGBoost (*Extreme Gradient Boosting*) and LightGBM (*Light Gradient Boosting Machine*) are two of the most

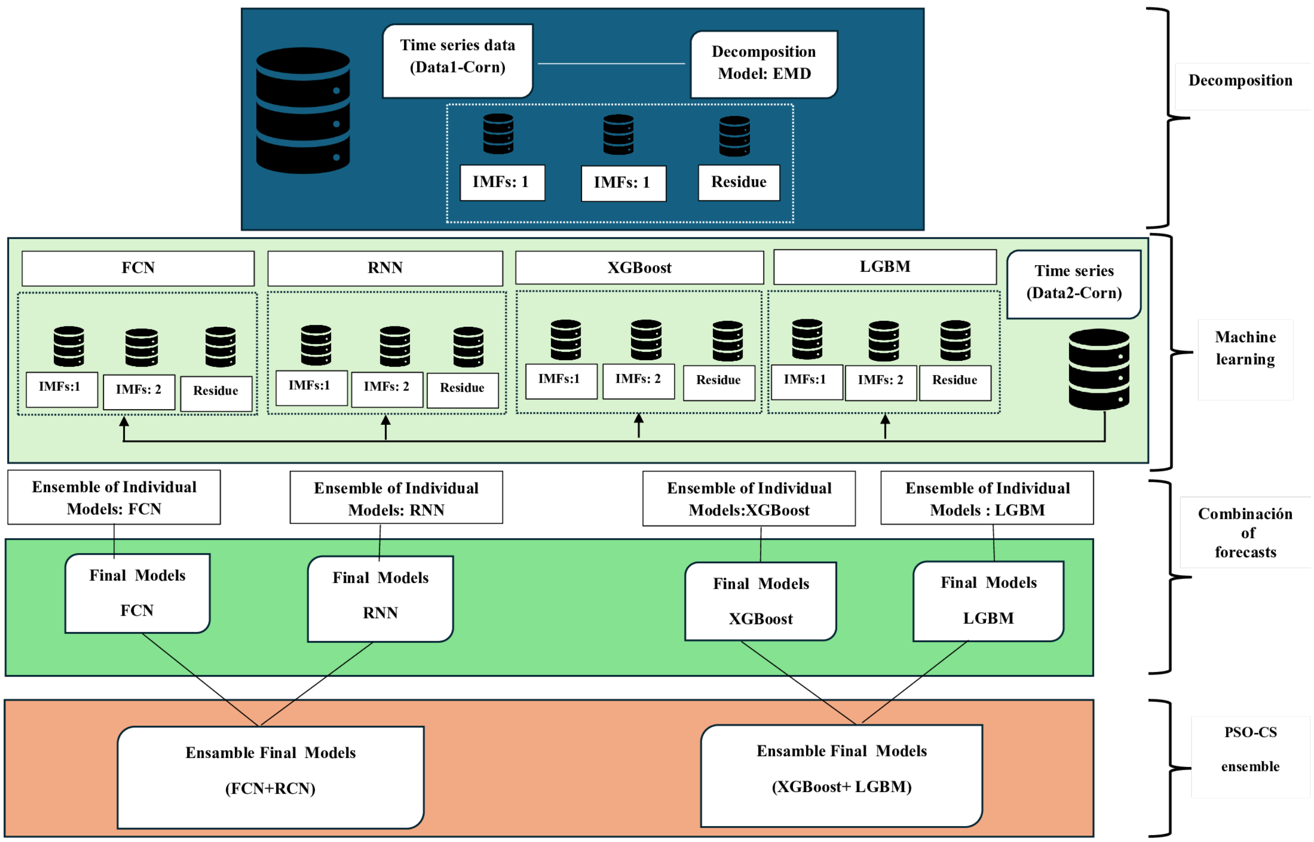

Figure 2: Architecture for forecasting corn prices in Colombia using multivariate time series and ensemble learning models. The figure illustrates the three-phase methodology: (1) univariate prediction using EMD and ML models, (2) integration of CBOT corn futures data to create a multivariate model, and (3) Ensemble predictions are optimized through a PSO-CS metaheuristic to enhance accuracy and robustness.

popular and powerful gradient boosting algorithms used in ML. Both algorithms are designed to optimize predictive performance while maintaining computational efficiency. XGBoost employs advanced regularization techniques, such as L1 and L2, to prevent overfitting and handles missing data efficiently. It is particularly well-suited for structured/tabular data and often achieve state-of-the-art results in competitions. On the other hand, LightGBM is optimized for speed and scalability, using techniques such as histogram-based learning and leaf-wise tree growth to reduce memory usage and computational time. LightGBM is especially effective for large datasets and high-dimensional feature spaces. Both algorithms support parallel and distributed computing, making them highly efficient for large-scale ML tasks. While XGBoost focuses on robustness and accuracy, LightGBM prioritizes speed and scalability, making the choice between them dependent on the specific requirements of the problem.

The objective function for XGBoost is defined as:

$$\text{Obj} = \sum_{i=1}^{n} L\left(y_i, \hat{y}_i\right) + \sum_{k=1}^{K} \Omega(f_k) \tag{4}$$

Here, $L(y_i, \hat{y}_i)$ is the loss function that measures the difference between the predicted value $\hat{y}_i$ and the true value y_i. The term $\Omega(f_k)$ represents the regularization function applied to each of the K trees f_k, which controls model complexity.

The regularization term $\Omega(f_k)$ is expressed as:

$$\Omega(f_k) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2 \tag{5}$$

Here, $\Omega(f_k)$ is the regularization term used in the XGBoost objective function. The first term γT penalizes the number of leaves T in the tree, while the second term $\frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2$ penalizes large leaf weights w_j using the regularization parameter λ.

The objective function for LightGBM is similarly defined as:

$$\text{Obj} = \sum_{i=1}^{n} L(y_i, \hat{y}_i) + \Omega(f_k) \tag{6}$$

Here, $L(y_i, \hat{y}_i)$ is the loss function that measures the prediction error for each observation, and $\Omega(f_k)$ is the regularization term that controls the complexity of the model.

LightGBM uses a leaf-wise growth strategy, selecting the leaf that maximizes the reduction in loss. The loss reduction is calculated as:

$$\Delta L = \frac{G^2}{H + \lambda} \tag{7}$$

Here, G is the sum of the first-order gradients (residuals) of the loss function for the leaf, H is the sum of the second-order gradients (Hessians), and λ is the regularization parameter. The expression ΔL quantifies the gain from splitting a node.

5.5. Dataset and material

This section presents the data collection used to train the prediction models, analyzes the relationships between the variables and describes the subset selected for the experiments. Additionally, a descriptive analysis of the input data is provided. The selected wholesale corn price data set in Colombia corresponds to the GRANABASTOS market located in the department of Atlántico (data1-corn) data was obtained from the Agricultural Sector Supply and Price Information System (SIPSA) of the National Administrative Department of Statistics in Colombia (DANE) DANE - SIPSA API The prices of shelled corn are reported in COP per kilogram (COP/kg).

data1-corn was created on Friday, September 30, 2024, and contains weekly yellow corn price data from January 1, 2014, through September 25, 2024. Table 2 contains the description of the data fields, and Table 3 shows the summary statistics.

The set Data2-corn corresponds to the prices of corn futures contracts on the Chicago Board of Trade (CBOT) FUTURES-CONTRACTS - API and includes the fields Date, Last Price, Opening Price, High, Low Price, and Volume % change. On the Investing.com website, corn prices in the United States are quoted in US dollars (USD) per bushel. A bushel of corn is approximately equivalent to 25.4 kilograms. Futures markets are financial platforms where standardized contracts are traded to buy or sell an underlying asset, such as agricultural commodities, oil, metals, etc., at a future date at a pre-agreed price. Agricultural futures contracts, including corn, have become a key financial asset to mitigate price risk in volatile markets, offering opportunities for portfolio diversification and hedging against extreme price fluctuations (Mensi et al., 2025). Table 4 shows the description of the Data2-corn fields, and Table 5 shows the summary statistics.

The date field is common to both data sets, Data1-corn and Data2-corn, which allows integration to the models for their operation from multivariate time series. Data1-corn and Data2-corn were created on Friday, September 30, 2024, and contain data from January 1, 2014, through September 25, 2024.

The datasets Data3-corn and Data4-corn are explicitly used for model validation. Data3-corn corresponds to the wholesale prices of corn in the market of the department of

Table 2
Variables in Data1-corn (Granabastos market). Daily records of yellow maize prices used in univariate modeling.

Variable	Description
Date	Observation day
Source	Observed market (Granabastos)
Item	Yellow cracked maize
Average Price	Weighted average of observed prices, considering traded volumes per price range.
Minimum Price	Lowest price recorded for the product during the day.
Maximum Price	Highest price recorded for the product during the day.

Table 3
Summary statistics of average prices in Data1-corn. Descriptive values in COP/kg from Granabastos market (2014–2024).

Parameter	Value
Count	554
Mean	1746.01
Std	484.70
Min	848
25%	1448
50%	1543
75%	2162.5
Max	2850

Table 4
Variables in Data2-corn (CBOT futures). Daily corn futures data used for multivariate modeling.

Variable	Description
Date	Day of the record (DD/MM/YYYY)
Last	Closing price of the day
Open	Opening price
High	Highest price of the day
Low	Lowest price of the day
Volume	Total traded volume
% Change	Price variation from previous day

Table 5
Summary statistics of Data2-corn (CBOT futures). Descriptive values in USD/bushel from 2014 to 2024.

Parameter	Value
Count	3787
Mean	484.88
Std	138.65
Min	301.5
25%	369.25
50%	427.75
75%	608
Max	838.75

Monteria, a Colombian region. Data4-corn corresponds to the wholesale prices of yellow corn in the Barranquillita market, which belongs to the department of Atlántico Colombian Caribbean region. It contains weekly data on the price of yellow corn from January 1, 2014, to September 25, 2024.

5.6. Input and Output Structure of the Models

For the Data1-corn dataset, which corresponds to local maize prices in Colombia, the input structure was defined using lagged variables generated from the univariate time series. Specifically, seven past values of the maize price (labeled as “retraso-1” to “retraso-7”) were used, representing weeks t-1 to t-7, respectively. These values form the feature vector for each training instance. The model output corresponds to the "PROMEDIO" column, which represents the maize price at week t. This transformation of the univariate series into a supervised learning format enables the model to capture relevant temporal dependencies for weekly price forecasting.

For the Data1-corn and Data2-corn, which integrates both local and international information, a multivariate input structure was employed. This structure consisted of seven lagged values of the Colombian maize price (labeled as “retraso-1” to “retraso-7”) and an additional variable named “Futuro,” which represents an external component associated with the international maize price, sourced from the Chicago Board of Trade (CBOT) futures market. This combination allows the model to capture not only internal temporal dynamics but also the influence of external market factors. The model output remains the “PROMEDIO” variable, which corresponds to the maize price in week t. This configuration enhances the model’s predictive capacity by integrating multiple relevant sources of information for the local market.

5.7. Evaluation measures

The performance metrics used for model evaluation are MAE, RMSE, MAPE, and R². Each metric provides a unique perspective on model behaviour, allowing analysis of absolute and percentage errors.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} \left| y_i - \hat{y}_i \right| \tag{8}$$

Here, y_i denotes the actual value, $\hat{y}_i$ is the predicted value, and n is the total number of observations. The mean absolute error (MAE) measures the average magnitude of the errors in a set of predictions, without considering their direction.

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{9}$$

Here, y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations. The root mean squared error (RMSE) measures the square root of the average squared differences between predicted and actual values, giving higher weight to larger errors.

$$\text{MAPE} = \frac{1}{n}\sum_{i=1}^{n} \frac{|y_i - \hat{y}_i|}{y_i} \tag{10}$$

Here, y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations. The mean absolute percentage error (MAPE) expresses the prediction error as a percentage of the actual value, making it useful for comparing errors across different scales.

$$R^2 = 1 - \frac{\text{SSE}}{\text{SST}} \tag{11}$$

Here, SSE is the sum of squared errors, and SST is the total sum of squares. The coefficient of determination R^2 measures the proportion of the variance in the dependent variable that is predictable from the independent variables.

6. Results and discussion

This section presents the results obtained from the application of the proposed methodology.

6.1. Forecast from Data1-corn

The first subsection of the results addresses Phase 1 prediction from a univariate time series called Data1-Corn. The analysis has four main steps.

Step 1. Data preparation

Before model training, IMF1, IMF2, and the residual were normalized using the MinMaxScaler method to scale values between -1 and 1, improving convergence in neural networks. Lag variables (up to 7 time steps) were created to capture temporal dependencies. The dataset was split chronologically into training (80%), validation (10%), and testing (10%) sets. For XGBoost and LightGBM models, a time series cross validation strategy using TimeSeriesSplit

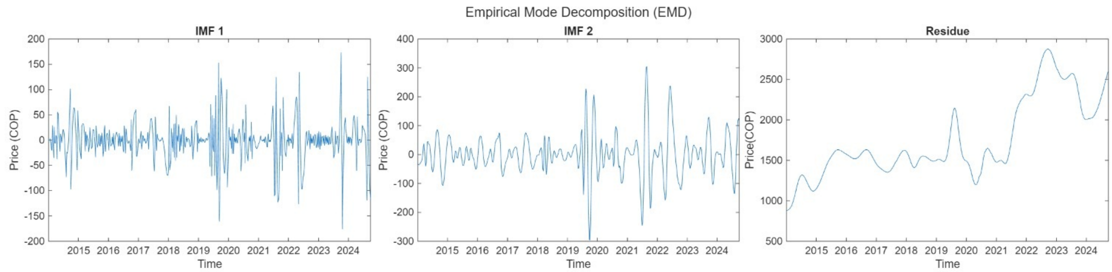

Figure 3: EMD decomposition of historical corn prices from Granabastos. The series is split into two intrinsic mode functions (IMF1, IMF2) and a residual, representing short-, medium-, and long-term components.

with three folds was applied to ensure robust evaluation and prevent data leakage.

Step 2: Data decomposition

Decomposition experiments are performed using the EMD technique, applied to Data1-corn. Figure 3 shows the original corn price series from Data1-corn, decomposed using the Empirical Mode Decomposition (EMD) technique into two intrinsic mode functions (IMFs) and a residual. These components are stored under the fields IMF1, IMF2, and Residual, respectively. IMF1 represents high-frequency fluctuations associated with short-term noise or random events, characterized by high volatility and no clear pattern. IMF2 captures intermediate-frequency variations, possibly linked to recurring seasonal cycles, and displays a more ordered structure. Finally, the Residual field reflects a long-term trend marked by sustained growth, potentially indicating structural evolution in the underlying data.

Step 3: Individual Forecasting

After applying the Empirical Mode Decomposition (EMD) technique to the original time series, three components were obtained: two intrinsic mode functions (IMF1 and IMF2) and a residual. Each of these components was modeled independently, considering their specific behavior and level of complexity. Prior to training, the data were normalized to a range between -1 and 1 using the MinMaxScaler. Lag variables were incorporated up to 7 steps, and the dataset was split into training (80%), validation (10%), and test (10%) subsets.

For the neural network models, FCN and RNN were implemented. FCN included hidden layers with tanh activation, a linear output layer, and was optimized using the RMSprop algorithm. RNN was constructed with LSTM layers, whose architecture and hyperparameters were optimized using the Optuna library. Both models were trained separately for each IMF and the residual. Meanwhile, XGBoost and LightGBM were tuned through hyperparameter search, adjusting values such as learning rate, maximum depth, number of leaves (num_leaves), and number of estimators, with time series cross-validation.

The results showed that tree-based models (XGBoost and LightGBM) consistently outperformed the neural networks across all components. In IMF2, LightGBM achieved an MSE of 0.0763 and XGBoost 0.0891, while FCN had a considerably higher MSE (0.2850). For the residual component, the tree models again demonstrated superior ability to capture long-term trends, with XGBoost reaching an MSE of 0.0355 and LightGBM 0.0479. Table 6 reports the results of the metrics applied to the MSE values for each EMD component across individual models, including the evaluation of IMF1, IMF2, and the residual.

Step 4: Combination of forecasts

For the four models analyzed (XGBoost, FCN, RNN, and LightGBM), the assembly was performed by combining the predictions of the components derived from the EMD: IMFs 1, IMFs 2 and residual. Subsequently, the predictions generated by the trained models for each mode were combined by an additive ensemble, in which the individual predictions of IMFs 1, IMFs 2 and residual were summed to reconstruct the original time series.

The analysis of the metrics shows that in the validation set, the XGBoost model shows that the root mean square error (RMSE) is 69.22 COP and a mean absolute error (MAE) of 54.70 COP, indicating considerable accuracy in the predictions. The mean absolute percentage error (MAPE) is 2.52%, and the correlation coefficient (R) reaches 0.87, suggesting a strong relationship between predictions and actual values.

On the other hand, the LightGBM (LGBM) model also exhibits a solid performance, with an RMSE of 66.23 weights and an MAE of 51.92 weights. Its MAPE is 2.33%, and the R is 0.91, indicating an even stronger correlation than XGBoost.

In contrast, FCN and RNN underperform in this set, with RMSE of 213.65 and 188.11 weights, respectively, and MAE of 172.04 and 134.29 COP. Both models have higher MAPEs, 7.92% for FCN and 6.05% for RNN, and negative correlation coefficients, suggesting low accuracy and an inverse relationship between predictions and actual values.

On the test set, XGBoost significantly improves its performance, recording a root mean square error (RMSE) of 26.84 (COP) and a mean absolute error (MAE) of 20.51. The mean absolute percentage error (MAPE) decreases to 0.78%, and the correlation coefficient (R) increases to 0.94,

Table 6
MSE values for each EMD component across individual models. Evaluation of IMF1, IMF2, and the residual using XGBoost, FCN, RNN, and LGBM.

IMFs	Metric	XGBoost	FCN	RNN	LGBM
IMF 1	MSE	0.2827	0.1016	0.0832	0.2776
IMF 2	MSE	0.0891	0.2850	0.3087	0.0763
Residual	MSE	0.0355	0.1448	0.2163	0.0479

Table 7
Evaluation metrics of individual and ensemble models on Data1-corn. Results for validation and test sets.

Model	Metric	Validation	Test
XGBoost	RMSE	69,2236	26,8398
	MAE	54,7005	20,5127
	MAPE	2,52%	0,78%
	R^2	0,8700	0,94
FCN	RMSE	213,6510	156,2484
	MAE	172,0380	135,8828
	MAPE	7,92%	5,44%
	R^2	-0,2720	-3,1400
RNN	RMSE	180,1100	234,6564
	MAE	134,2850	190,2517
	MAPE	6,05%	7,89%
	R^2	-0,6700	-12,75
LGBM	RMSE	66,2290	36,9358
	MAE	51,9157	29,0004
	MAPE	2,33%	1,09%
	R^2	0,9100	0,8900
(XGB+LGBM)	MSE	718,29	3,876.51
	RMSE	26,8	62,26
	MAE	20,33	45,96
	MAPE	0,78%	2,10%
(FCN+RNN)	MSE	11,285.97	4,374.96
	RMSE	106,24	66,14
	MAE	90,9	51,43
	MAPE	3,54%	2,35%

indicating excellent accuracy and a robust correlation between predictions and actual values.

LGBM also shows good performance in this set, with an RMSE of 36.94 COP and an MAE of 29.00 COP. Its MAPE is 1.09%, and the R is 0.89, maintaining a strong correlation.

FCN and RNN continue to perform poorly, with RMSE of 156.25 and 234.66 COP, respectively, and MAE of 135.88 COP and 190.25 COP. The MAPEs are 5.44% for FCN and 7.89% for RNN, with negative correlation coefficients, indicating a persistent low accuracy and inverse relationship in predictions.

Table 7 details the results of the evaluation metrics of the XGBoost, LightGBM , FCN AND RNN after assembly. These results allow us to compare their performance and efficiency, highlighting their ability to make accurate predictions after the process of assembling the modes and the residual. Figure 4 illustrate the test window and predictions for the RNN, LightGBM, and FCN models on the Data1-corn dataset. The red line represents the actual values, while the green line shows the model predictions.

Step 5: Model Combination

Model assembly was performed using the PSO-CS algorithm (Particle Swarm Optimization Cuckoo Search), combining the predictions of the RNN and FCN models by optimizing weights to minimize the mean square error (MSE) in the validation data. The process began with loading the predictions generated by each model for the validation and test sets. Subsequently, an objective function was defined based on the MSE, where the combined predictions were calculated as a weighted sum of the individual model predictions, with weights adjusted by the PSO-CS algorithm.

The PSO-CS algorithm began with the initialization of a population of particles, each representing combinations of normalized weights. During the process, the particles updated their positions and velocities using the best-identified local and global solutions. In addition, the global search component based on Lévy flights (proprietary to Cuckoo Search) introduced random variations to prevent the algorithm from becoming trapped in local minima.

After 20,000 iterations, the optimal weights were determined: 53.88% for the RNN and 46.12% for the FCN, reflecting their relative contribution to the assembly.

The final predictions were calculated as a weighted combination of the model outputs; the results of the performance metrics in the validation set are as follows: the mean squared error (MSE) is 11,285.97, indicating a considerable mean squared difference between the predictions and the actual values. The root mean squared error (RMSE) has a value of 106.24, which shows that the average error in the same units as the original data is significant. The mean absolute error (MAE) is 90.9 weights, showing that, on average, the model predictions deviate from 90.9 weights from the true values. The mean absolute percentage error (MAPE) is calculated as 3.54%, indicating that the model, on average, has a low relative error, which is acceptable in percentage terms.

In the test set, the model improves markedly, showing a mean squared error (MSE) of 4,374.96, which represents a significant decrease in the mean squared differences. The root mean squared error (RMSE) is reduced to 66.14 COP, suggesting improved prediction accuracy. The mean absolute error (MAE) is also lower, with a value of 51.43 COP, indicating that the predictions are closer to the actual values. Finally, the mean absolute percentage error (MAPE) is 2.35%, reflecting a very solid and accurate performance in relative terms.

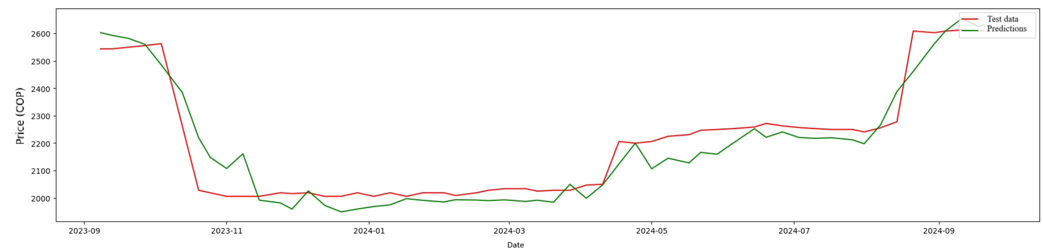

(a) XGBoost on Data1-corn

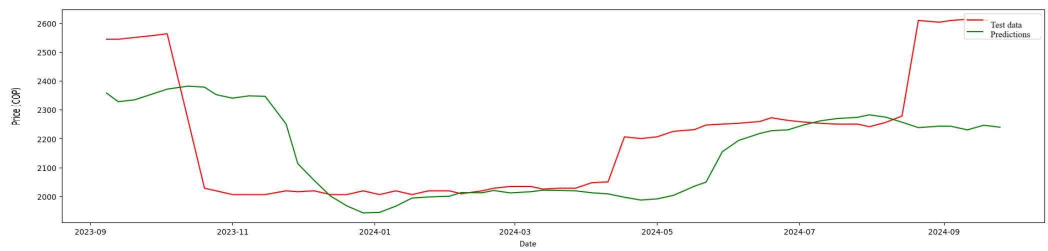

(b) RNN on Data1-corn

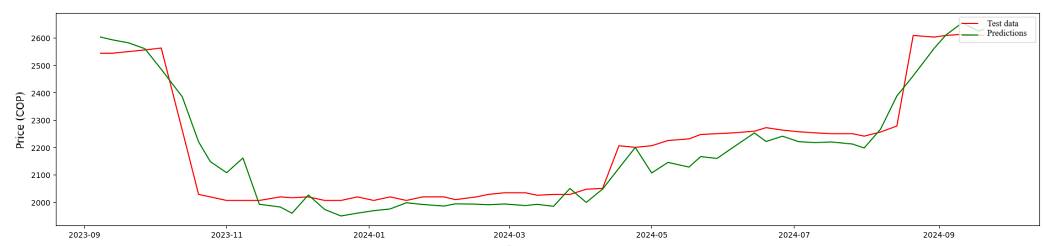

(c) LightGBM on Data1-corn

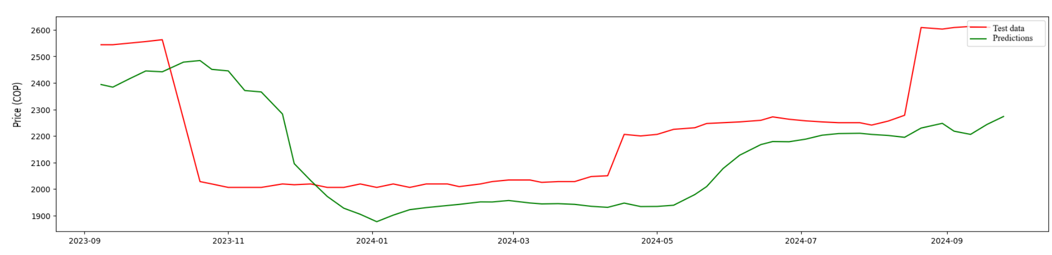

(d) FCN on Data1-corn

Figure 4: Individual model forecasts using Data1-corn. Individual forecasts using Data1-corn by XGBoost, LGBM, FCN, and RNN, based on the ‘PROMEDIO‘ series. Each model predicts the EMD-derived components ‘IMF1‘, ‘IMF2‘, and ‘Residual‘, and the final forecast is obtained by summing their predictions.

Model assembly using XGBoost and LightGBM, optimizing the weights assigned to each using the PSO-CS algorithm. In this case, the PSO-CS algorithm was responsible for optimizing the weights for the XGBoost and LGBM predictions, assigning them optimal values of 53.88% for XGBoost and 46.12% for LGBM. The result is a weighted combination of the predictions of both models.

The PSO-CS algorithm combines the global optimization features of Particle Swarm Optimization (PSO) with the exploratory capability of Cuckoo Search (CS). PSO initialized a population of particles, each representing combinations of normalized weights, and updated their positions and velocities using the best-identified local and global solutions. Additionally, CS introduced Lévy flights to generate random variations in the positions, preventing the algorithm from being trapped in local minima.

After 20,000 iterations, the optimal weights were identified: 53.88% for XGBoost and 46.12% for LightGBM, reflecting the higher contribution of XGBoost to the final performance.

In the validation set, the mean squared error (MSE) is 718.29, indicating that, on average, the predictions differ from the actual values by approximately 26.8 COP based

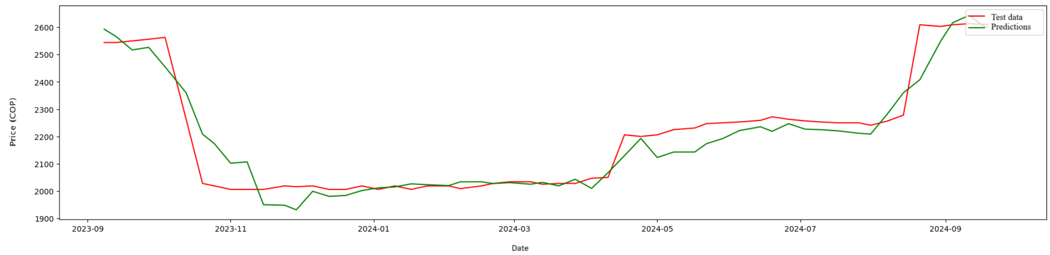

(a) XGBoost + LightGBM on Data1-corn

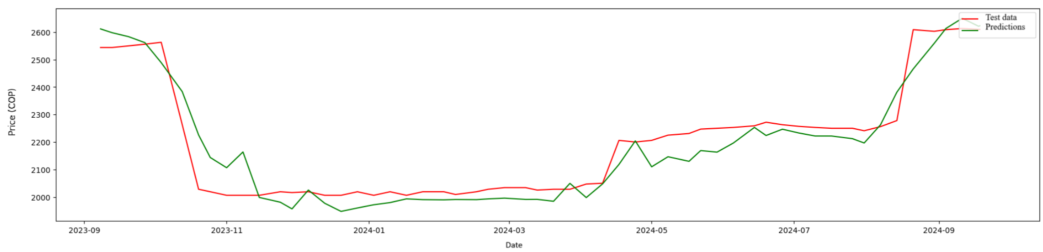

(b) FCN + RNN on Data1-corn

Figure 5: Ensemble forecasts using Data1-corn. Results from combining tree-based models and neural networks with PSO-CS optimization.

on the root mean squared error (RMSE). The mean absolute error (MAE) is 20.33 COP, reflecting a minimal average deviation between predictions and actual values. In addition, the mean absolute percentage error (MAPE) is 0.78%, meaning that the model predictions have a very low percentage deviation from the actual values. In the test set, the model maintains solid performance, with an MSE of 3,876.51 and an RMSE of 62.26 COP, indicating a somewhat higher but still acceptable average deviation. The MAE is 45.96 COP, and the MAPE stands at 2.1%, demonstrating that, although there is a slight increase in errors when applying the model to unseen data, accuracy remains high.

These results are presented in Table 7 and allow us to compare their performance and efficiency, highlighting their ability to make accurate predictions after the combination process.

The figure 5 shows the prediction results generated by the ensemble models applied to the Data1-corn dataset. The red lines represent the actual values from the test set, while the green lines indicate the predictions generated by the combined models.

6.2. Forecast from Data1-corn and Data2-corn

This section reports the results related to the integration of the Data1-corn and Data2-corn series.

Step 1: Verification of correlation and cointegration between series

An analysis of correlation and cointegration was conducted between the time series Data1-corn and Data2-corn. To ensure comparability, both series were converted to Colombian pesos (COP): Data2-corn, originally denominated in U.S. dollars, was converted using the historical exchange rate corresponding to each record's date. The series were then temporally aligned and standardized to the same unit of measurement (kilograms).

The Augmented Dickey-Fuller (ADF) stationarity test indicated that both series were non-stationary in their original form. Data1-corn showed a p-value of 0.4419, and Data2-corn a p-value of 0.5506. After first-order differencing, the p-values dropped to 0.0000 and 0.00000011 respectively, confirming stationarity and enabling their use in predictive analysis.

The Pearson correlation coefficient between Data1-corn and Data2-corn was 0.8338, with a p-value of 9.93×10^{-14}, indicating a strong and statistically significant positive relationship. Additionally, the Engle-Granger cointegration test revealed a long-term equilibrium relationship between the series, with an ADF statistic of -3.6281 and a p-value of 0.0052, leading to the rejection of the null hypothesis of no cointegration.

The Spearman rank correlation was 0.7335, with a p-value of 8.53×10^{-93}, also indicating a strong and statistically significant positive relationship in terms of order that is, both series tend to follow a similar ranking pattern regardless of their exact values.

These results support the use of futures contract prices as a valid predictive input for analyzing local corn prices in market contexts such as GRANABASTOS.

The figure 6 presents the results of the stationarity analysis for the Data2-corn and Data1-corn datasets, respectively. It displays the differenced time series along with their corresponding autocorrelation (ACF) and partial autocorrelation (PACF) plots. These visual tools are useful for identifying temporal dependencies and confirming stationarity before model implementation.

Step 2: Individual Multivariate Predictive Models

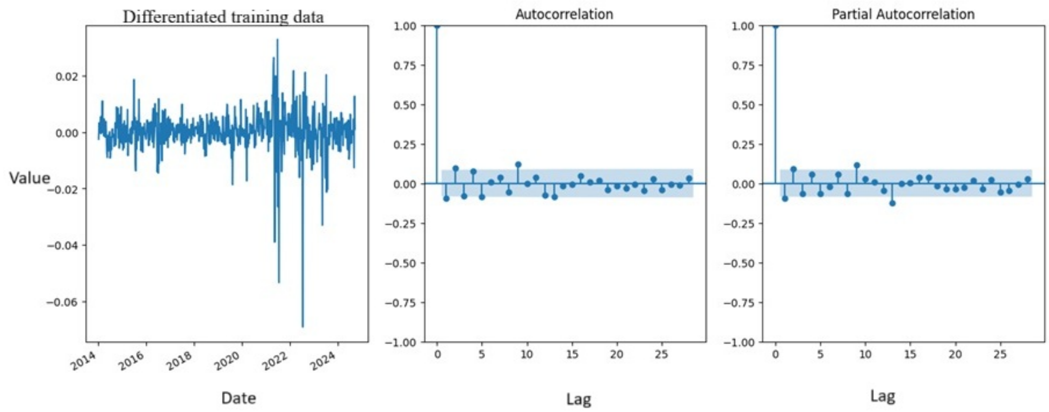

(a) ACF and differencing – Data2-Corn

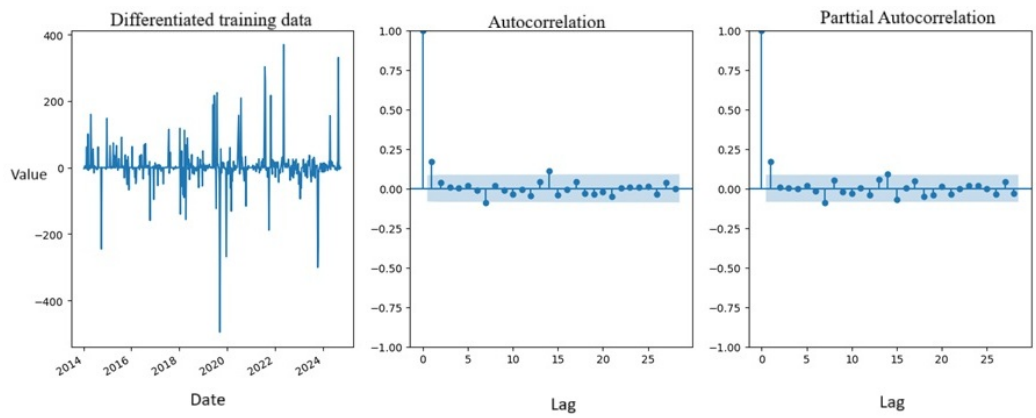

(b) ACF and differencing – Data1-corn

Figure 6: Stationarity analysis of time series using ACF and differencing. Visual assessment of autocorrelation and differentiation applied to Data1-Corn and Data2-Corn for the ADF test.

The results presented in Table 9 reveal clear differences in the performance of predictive models when incorporating the *Data1-corn* and *Data2-corn* series. Among the four evaluated models, XGBoost and LGBM—both based on decision trees—demonstrated significantly superior performance compared to the neural networks FCN and RNN.

The XGBoost model produced outstanding metrics: an RMSE of 22.19 COP, MAE of 16.76 COP, and MAPE of just 0.64% during the validation phase, with an R^2 of 0.96. In the testing phase, the model maintained strong generalization capabilities, achieving an RMSE of 60.82 COP, MAE of 47.11 COP, MAPE of 2.15%, and an R^2 of 0.90. Similarly, the LGBM model exhibited robust performance, with an RMSE of 28.43 COP and MAE of 23.25 COP during validation, and a MAPE of only 0.89%. In the testing phase, it continued to perform well with an RMSE of 60.96 COP, MAE of 45.68 COP, MAPE of 2.05%, and an R^2 of 0.90.

In contrast, the neural networks displayed limited predictive capability. The FCN model yielded an RMSE of 554.82 COP and MAE of 549.4 COP during validation, with a MAPE of 26.43% and a negative R^2 of -75.9, indicating a poor model fit. These metrics worsened during testing, with an RMSE of 743.66 COP, MAE of 698.29 COP, MAPE of 53.5%, and an R^2 of -62.4. Similarly, the RNN model showed an RMSE of 518.46 COP and MAE of 508.02 COP in validation, with a MAPE of 23.99% and an R^2 of -99.3. In testing, the RMSE increased to 598.27 COP, MAE to 531.34 COP, MAPE to 31.74%, and R^2 improved slightly but remained negative at -3.74.

The poor performance of the FCN and RNN models (MAPE = 53.5%, negative R^2) is attributed to structural limitations inherent to neural networks when applied to short, seasonal, and noisy time series. In the case of the RNN, the potential vanishing gradient problem hindered the model's ability to capture temporal dependencies, particularly with limited sequence lengths (Chung and Shin, 2018). Meanwhile, the FCN model, which does not explicitly model the temporal dimension, exhibited a strong tendency to overfit, exacerbated by the scarcity of weekly data and the inherent variability of the Colombian agricultural market.

This diagnosis is reinforced by visual evidence observed in the validation and test plots, where both models exhibit a clear overfitting pattern. The RNN curve, in particular, displays a more pronounced divergence from the validation curve, reflecting its difficulty in maintaining

Table 8
Evaluation metrics of models using combined Data1-corn and Data2-corn. Results on validation and test sets.

Model	Metric	Validation	Test
XGBoost	RMSE	22.19	60.82
	MAE	16.76	47.11
	MAPE	0.64%	2.15%
	R^2	0.96	0.9
FCN	RMSE	554.82	743.66
	MAE	549.4	698.29
	MAPE	26.43%	53.5%
	R^2	-75,9	-2,74
RNN	RMSE	518.46	598.27
	MAE	508.02	531.34
	MAPE	23.99%	31.74%
	R^2	-99,3	-3,74
LGBM	RMSE	28.43	60.96
	MAE	23.25	45.68
	MAPE	0.89%	2.05%
	R^2	0.94	0.9

Table 9
Statistical significance of the Diebold-Mariano (DM) test applied to individual predictive models using Data1-corn and Data2-corn

Comparison	DM Statistic	P-value	Statistical Decision ($\alpha = 0.05$)
XGBoost vs LGBM	0.44	0.659	Fail to reject H_0
XGBoost vs RNN	-5.38	0.0000016	Reject H_0
XGBoost vs FCN	-4.83	0.0000112	Reject H_0
LGBM vs RNN	-5.53	0.0000009	Reject H_0
LGBM vs FCN	-4.85	0.0000104	Reject H_0
FCN vs RNN	2.69	0.0096	Reject H_0

stability outside the training environment. These observations support the conclusion that, under these data conditions, the FCN and RNN architectures exhibit limited generalization capacity, particularly in highly seasonal agricultural contexts.

These results may be explained by the greater capacity of tree-based models to adapt to complex nonlinear structures, handle noise, and detect hierarchical relationships among variables. In contrast, neural networks require more stringent conditions for successful training, including carefully tuned architectures, larger amounts of data, and more exhaustive preprocessing. The findings of these results are summarized in Table 8 presents the results of the metrics applied to the individual forecasts and Figure 7 shows the prediction results generated by the XGBoost, RNN, FCN, and LightGBM models applied to the Data1-corn and Data2-corn datasets. In each subfigure, the red line represents the actual values from the test set, while the green line indicates the predictions made by each model. This visual comparison allows for an evaluation of each model's performance in capturing the trend of the real data.

Step 3:Statistical Significance of the Diebold-Mariano Test (DM)

To assess whether the differences in predictive accuracy among the individual models were statistically significant, the Diebold-Mariano (DM) test was employed. This procedure involved calculating the squared errors of each model and performing pairwise comparisons to evaluate the significance of the error differentials. The results indicated no significant difference between XGBoost and LGBM, suggesting that these tree-based models exhibit statistically comparable performance.

In contrast, significant differences emerged when comparing the tree-based models to the neural networks, with the fully connected network (FCN) achieving the best predictive results, followed by the recurrent neural network (RNN). These findings support the progression toward a model integration phase, utilizing optimization strategies such as the Particle Swarm Optimization (PSO) ensemble to combine the strengths of both model families and enhance forecasting accuracy. The table 9 presents the results obtained.

Step 4: Combination of Models for Data1-Corn and Data2-Corn

The results presented in Table 10 demonstrate that the PSO-CS ensemble (XGBoost + LGBM) achieved the best predictive performance, with a MAPE of 2.06%, outperforming both the individual models and the neural network ensemble (FCN + RNN, MAPE = 30.74%). This performance gap can be explained by several methodological factors and characteristics of the data.

First, XGBoost and LightGBM are gradient-boosted decision tree algorithms well-suited for modeling tabular, multivariate data. Their capacity to handle nonlinear interactions, noisy features, and collinearity makes them particularly robust in scenarios where data variability is high—such as in maize pricing, which is influenced by both national and international market dynamics (e.g., CBOT futures). These models offer enhanced generalization even in moderately sized datasets.

In contrast, the weaker performance of the neural network ensemble (FCN + RNN) is likely due to the limited weekly data volume, high seasonality in the agricultural market, and the sensitivity of neural networks to noise. Although RNNs are designed to capture temporal dependencies, their effectiveness diminishes when the sequence length is short and the signal-to-noise ratio is low. Moreover, neural networks tend to emphasize internal patterns within the target series, while tree-based models such as XGBoost and LGBM are better able to leverage exogenous variables like CBOT prices, which are essential in the proposed multivariate configuration.Figure 8. presents the prediction results generated by ensemble models applied to the Data1-corn and Data2-corn datasets. The red line represents the actual values from the test set,

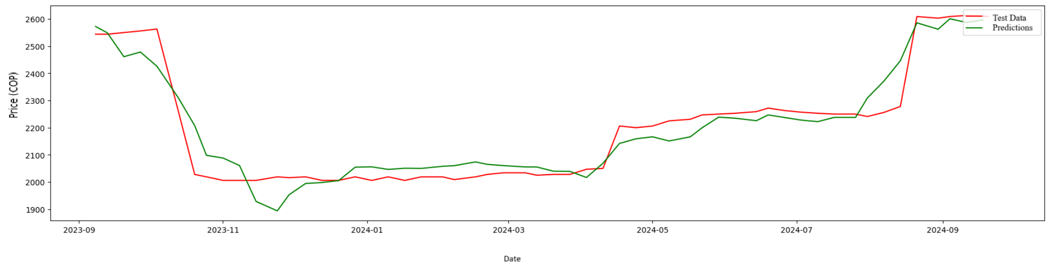

(a) XGBoost on Data1-corn and Data2-corn

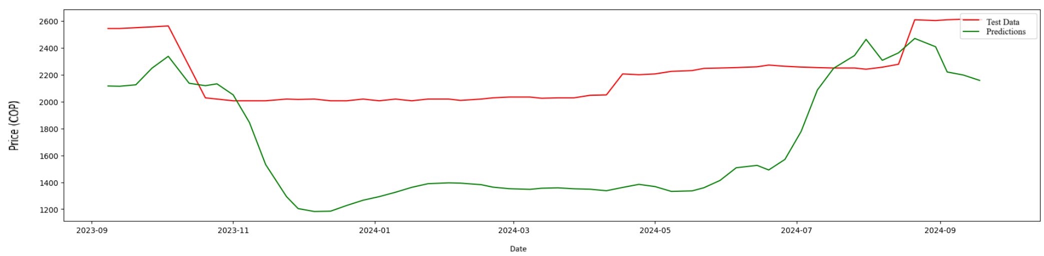

(b) RNN on Data1-corn and Data2-corn

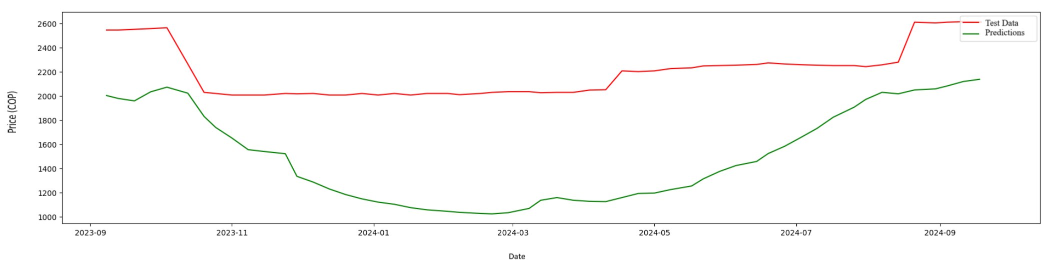

(c) FCN on Data1-corn and Data2-corn

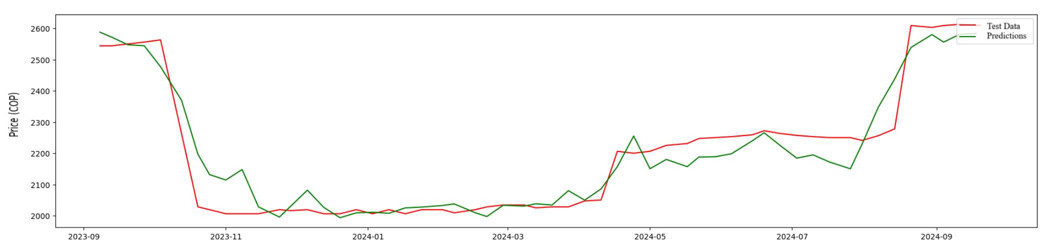

(d) LightGBM on Data1-corn and Data2-corn

Figure 7: Forecasts using multivariate models on corn data. Predictions from four ML models applied to the integrated Data1-corn and Data2-corn series.

while the green line indicates the predictions made by the combined models.

The results obtained in this phase indicate that the PSO-CS ensemble model (XGBoost + LGBM) achieve superior forecasting performance, reaching a MAPE of 2.19%. This model outperforms both the individual learners and the neural network-based ensemble (RNN + FCN). Such improvement is attributed to the synergy between methodological robustness and the intrinsic characteristics of the combined datasets used in the modeling process.

First, XGBoost and LightGBM are supervised learning algorithms based on gradient-boosted decision trees, known for their effectiveness in predicting tabular data and their ability to handle non-linear relationships, noisy variables, and multicollinearity. In the context of corn pricing, which involves highly volatile local and international factors (such as CBOT prices), these models offer greater robustness and generalization, even when working with multivariate data structures.

On the other hand, the lower performance of the neural network ensemble (RNN + FCN), which recorded a MAPE of 30.74%, may be explained by the high sensitivity of these architectures to noise, as well as their dependence on large volumes of training data to effectively learn complex

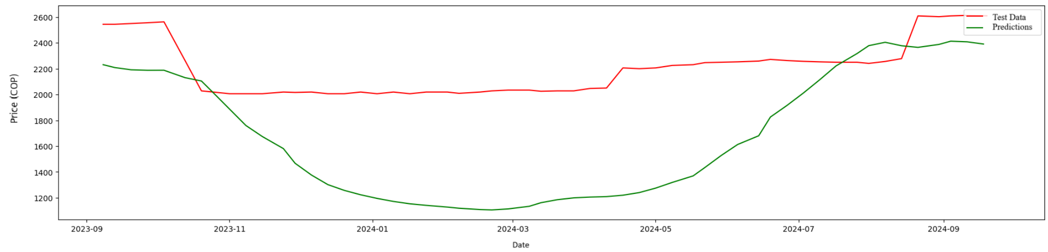

(a) FCN + RNN on Data1-corn and Data2-corn

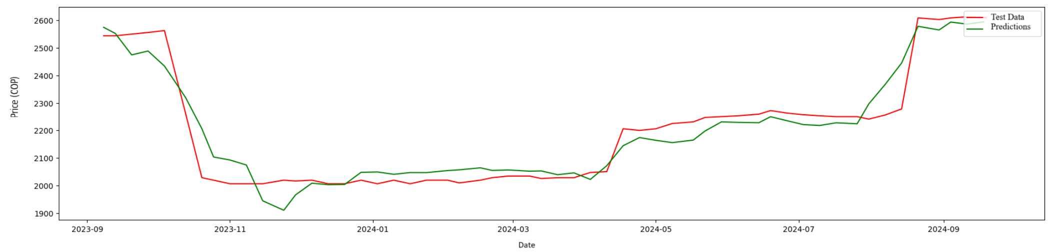

(b) XGBoost + LightGBM on Data1-corn and Data2-corn

Figure 8: Ensemble models using multivariate time series. Forecast results from the combination of neural and tree-based models on integrated corn price data.

Table 10
Performance of PSO-CS ensemble using Data1-corn and Data2-corn. Evaluation metrics for the combined model (XGBoost + LGBM) optimized via PSO-CS.

Model	Metric	Validation	Test
(XGBoost+LGBM)	RMSE	21.94	58.32
	MAE	16.71	45.44
	MAPE	0.64%	2.06%
	R^2	0.96	0.91
(FCN+RNN)	RMSE	339.69	628.61
	MAE	324.74	540.55
	MAPE	14.11%	30.74%
	R^2	-40,53	-0,7

temporal patterns. In this case, the limited length of the weekly time series, combined with the seasonal variability inherent to the agricultural market, constrained the RNN's ability to accurately capture sequential dependencies.

Additionally, neural networks tend to prioritize internal patterns within the target series, whereas tree-based models like XGBoost and LGBM more efficiently integrate exogenous signals from international markets. This results in a better exploitation of CBOT future prices in the multivariate model.

The superiority of the proposed ensemble is further supported by the sensitivity analysis, which showed a balanced weight distribution and consistent performance across different configurations of the PSO-CS hyperparameters, indicating high model stability and generalization capacity.

Furthermore, the results of the Diebold-Mariano statistical test confirmed that the differences in prediction errors between the XGBoost + LGBM ensemble and the neural network ensemble (recurrent and dense) are statistically significant ($p < 0.05$). This provides strong evidence supporting the superiority of the proposed approach in terms of predictive accuracy and statistical robustness.

Although the sensitivity to noise and boundary effects discussed in Subsection 5.1 are recognized as theoretical limitations of the EMD method, their practical impact in this study was minimal. This is supported by the error metrics obtained from the XGBoost and LightGBM models, which achieved MAPE values below 1% during validation (0.64% and 0.89%, respectively) and around 2% in the test phase (2.15% and 2.05%). Furthermore, the forecast curves remained stable even at the boundaries of the series, where edge effects typically manifest more strongly. These findings suggest that, although boundary effects and noise sensitivity are methodologically relevant, their influence did not compromise the predictive capacity or generalization ability of the model in the multivariate integration setting.

Step 5: Sensitivity and Stability Analysis of the Predictive Ensemble

The table 12 presents the results of the sensitivity analysis of the PSO-CS algorithm was conducted, applied to the ensemble of recurrent neural networks (RNN + FCN) and gradient boosting models (XGBoost + LGBM), both trained with the multivariate time series *Datos1_Corn* and *Datos2_Corn*. The RNN + FCN ensemble demonstrated superior accuracy in maize price forecasting, achieving significantly lower error metrics: MSE ranged from 70,662 to 99,884 COP^2, MAE between 189 and 234 COP, and MAPE between 7.97% and 9.93%. These results indicate a strong fitting capacity despite the complexity of the data. Although the coefficient of determination (R^2)

Table 11

Hyperparameter settings for models applied to Data1-corn and Data2-corn. Configuration used during training and optimization of individual learning algorithms.

Model	Type	Configuration
FCN	Fully Connected Neural Network	Sequential: Layer 1: Dense(8, activation='tanh') Layer 2: Dense(9, activation='tanh') Layer 3: Dense(52, 'tanh') Layer 4: Dense(25, 'tanh') Salida: Dense(1, linear) Optimizador: RMSprop Loss: MSE
RNN	Recurrent Neural Network (LSTM)	Sequential: Layer 1: LSTM(8) Layer 2: LSTM(25) Layer 3: LSTM(57) Output layer: Dense(1) Learning rate: 0.0002 Optimizador: Adam o RMSprop EarlyStopping used
XGBoost	Tree-based model	XGBRegressor con hiperparámetros: n_estimators=1000 early_stopping_rounds=50 learning_rate ∈ {0.01, 0.1, 0.2} max_depth ∈ {2–8} subsample ∈ {0.8, 0.9, 1.0}
LightGBM	Tree-based model	LGBMRegressor: n_estimators = 1000 num_leaves ∈ {600–1000} max_depth ∈ {8–1000} learning_rate = 0.01 Histogram-based tree learner used
PSO-CS (RNN+FCN)	Ensemble metaheuristic	N = 50, T = 20000 w_max = 0.99, w_min = 0.1 c1 = 2.0, c2 = 2.0 pa = 0.25, α = 0.01 assigned weights ∼ w1 = 0.46–0.60 / w2 = 0.40–0.54
PSO-CS (XGBoost+LGBM)	Ensemble metaheuristic	N = 50, T = 20000 w_max = 0.99 w_min = 0.1, c1 = 2.0, c2 = 2.0 pa = 0.25, α = 0.01 assigned weights ∼ w1 = 0.48–0.55 / w2 = 0.45–0.52

was negative (from -0.67 to -1.36), the model exhibited behavior that reflects alignment with the temporal dynamics of the series. The PSO-CS algorithm assigned balanced weights to the individual models, ranging from 0.46 to 0.60 for the RNN and from 0.40 to 0.54 for the FCN, suggesting complementary contributions. Conversely, the XGBoost + LGBM ensemble yielded lower predictive accuracy, with an average MSE of 249,935 COP^2, MAE of 443 COP, and MAPE of 21.19%, but demonstrated higher stability across different PSO-CS configurations. The weights assigned by PSO-CS also remained balanced (0.48–0.55 for XGBoost and 0.45–0.52 for LGBM), reinforcing the robustness of this combination. In conclusion, while the RNN + FCN ensemble is more suitable when precision is the primary objective, XGBoost + LGBM offers an advantage in terms of model stability under varying optimization settings.

In the sensitivity analysis of the PSO-CS algorithm, several hyperparameters were adjusted to evaluate their influence on the performance of the ensemble models. *N* represents the population size, that is, the number of candidate solutions evaluated per iteration. *T* indicates the number of iterations, which determines the duration of the optimization process. *c1* and *c2* are acceleration coefficients that balance individual learning and social influence. *w_max* is the inertia weight factor, which

Table 12
Sensitivity analysis of PSO-CS ensemble models (RNN+FCN and XGBoost+LGBM). Variation in model performance based on changes in PSO hyperparameters.

Model	N	T	c1	c2	w_max	MSE (COP^2)	MAE (COP)	MAPE (%)	R^2	w1	w2
XGBoost + LGBM	10	5000	1	1	0.6	249935.76	443.23	21.19	-4.93	0.48	0.51
	30	10000	2	2	0.8	249938.63	443.24	21.19	-4.93	0.49	0.50
	50	20000	3	3	0.99	249949.20	443.26	21.19	-4.93	0.53	0.46
	10	10000	2	1	0.6	249954.16	443.27	21.19	-4.93	0.55	0.44
	30	5000	1	3	0.99	249935.19	443.23	21.19	-4.93	0.48	0.51
	50	10000	3	1	0.8	249936.56	443.23	21.19	-4.93	0.48	0.51
RNN + FCN	10	5000	1	1	0.6	70662.65	189.25	7.97	-0.68	0.46	0.53
	30	10000	2	2	0.8	97092.70	229.31	9.69	-1.30	0.58	0.41
	50	20000	3	3	0.99	98416.45	231.79	9.81	-1.34	0.59	0.41
	10	10000	2	1	0.6	76602.42	197.07	8.29	-0.82	0.49	0.50
	30	5000	1	3	0.99	99884.48	234.50	9.93	-1.37	0.60	0.40
	50	10000	3	1	0.8	87545.82	212.77	8.96	-1.08	0.54	0.45

regulates the trade-off between exploration and exploitation in the search space. Together, these parameters define the algorithm's behavior and affect the accuracy and stability of the ensemble models.

For the model combination process using the PSO-CS algorithm, a total of 20,000 iterations was established. This value was determined through empirical testing with configurations of 5,000, 10,000, and 15,000 iterations, where the objective function (based on the mean squared error) was observed to stabilize after approximately 15,000 iterations, without significant performance improvements beyond that point. Therefore, 20,000 iterations were selected as a suitable compromise to ensure adequate convergence while maintaining a reasonable computational cost.

Regarding the hyperparameter tuning of the LSTM and FCN models, the Optuna library was used, employing the Tree-structured Parzen Estimator (TPE) algorithm across 100 optimization trials per model. For LSTM, the optimized parameters included the number of layers (1 to 3), units per layer (8 to 64), learning rate (0.0001 to 0.01), batch size, activation function, and dropout rate (0 to 0.5). For FCN, various dense layer configurations were explored, including different numbers of neurons, activation functions (tanh and ReLU), and learning rates. Time-series cross-validation via TimeSeriesSplit was applied to preserve temporal integrity, and the EarlyStopping mechanism was activated to prevent overfitting. The average training time per trial was approximately 8 minutes in a GPU-enabled environment.

This tuning strategy ensured a balance between predictive accuracy and computational efficiency, allowing each base model to operate under its optimal configuration within the ensemble process.

Table 11 presents the results obtained for the RNN + FCN and XGBoost + LGBM ensemble models under different PSO-CS configurations, including performance metrics (MSE, MAE, MAPE, R^2) and the weights assigned to each model (w1 and w2).

Step 6: Uncertainty Quantification and Model Reliability

Uncertainty quantification is a key component in validating the robustness and reliability of predictive models. In this study, a block bootstrap resampling approach was implemented, which is particularly suitable for time series data as it preserves the temporal dependence between consecutive observations. This technique enables the generation of an empirical distribution of predictions based on multiple subsampled datasets, from which a 95% confidence interval (CI) was constructed around the ensemble model's mean prediction. This interval is visually represented by a shaded area, offering an explicit estimate of the range within which the forecasts are expected to lie, accounting for the inherent variability of the prediction process.

The results demonstrate that the ensemble model successfully captures the overall trend in the price series, while also adequately representing the uncertainty associated with the market dynamics. Both the predicted and actual values mostly fall within the 95% CI, reinforcing the statistical reliability of the model. Nevertheless, slight underestimations are observed in some extreme values, particularly during periods of rapid price increase a common feature in aggregated models that tend to smooth out fluctuations. As a potential improvement, the integration of complementary models more sensitive to sudden changes, or adjusting the ensemble's internal weighting scheme, may enhance responsiveness to highly volatile scenarios without compromising the model's overall forecast stability. Figure 9 presents the results obtained.

6.3. Model Validation: Multivariate Integration

To validate the robustness, accuracy, and generalization capability of the multivariate model, this section presents a strategy structured in three stages. First, an out-of-time validation is conducted using unseen data beyond the training horizon to evaluate the model's forecasting ability in real future scenarios. Second, an external validation is

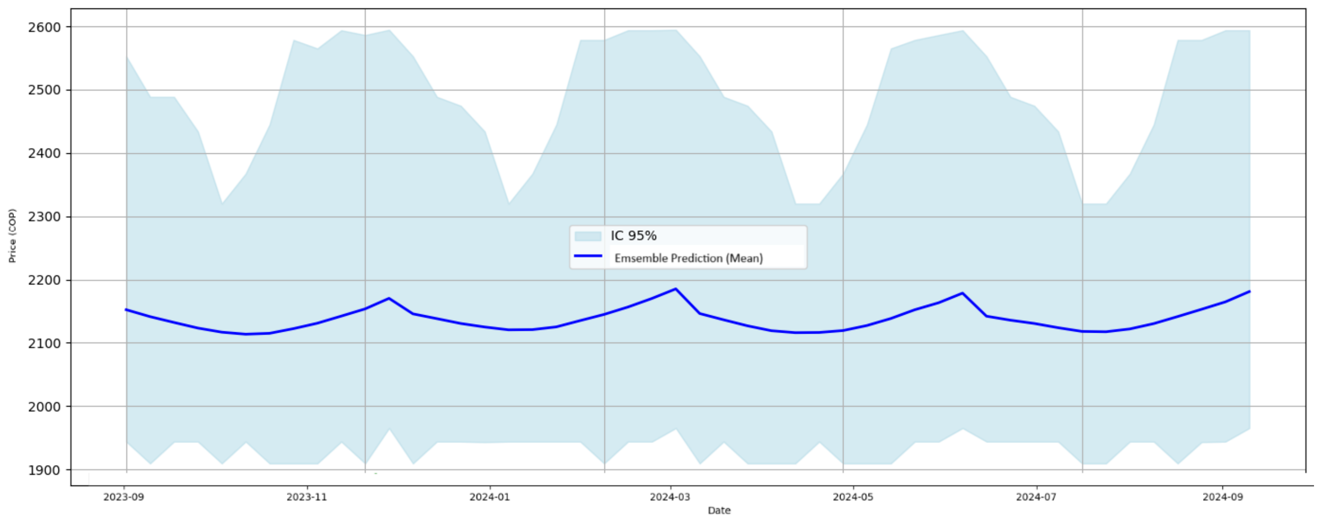

Figure 9: Uncertainty quantification and model reliability using ensemble prediction. The figure displays the mean prediction generated by an ensemble model (XGBoost+LGBM), accompanied by a 95% confidence interval estimated via block bootstrap. The shaded area reflects the uncertainty range across resampled predictions, highlighting the model's ability to provide statistically reliable forecasts.

performed using two different regional datasets (Data3-Corn and Data4-Corn) from various Colombian markets to assess the model's transferability across diverse geographical and commercial contexts, without the need to adjust its parameters. Third, a comparative analysis is carried out against benchmark models reported in recent literature that use the same time series, allowing the performance of the proposed approach to be contextualized within the current state of the art. Together, these three validation strategies provide a solid framework for evaluating both the predictive accuracy and the practical applicability of the proposed ensemble model.

Step 1: Out of Time Validation

To evaluate the temporal robustness of the proposed ensemble approach, an out-of-time validation was performed using unseen data from October 22, 2024, to March 1, 2025. The models were tested on their ability to generalize beyond the training period. Results show that XGBoost achieved an RMSE of 21.47, MAE of 17.44, MAPE of 0.62%, and R^2 of 0.49. In contrast, LGBM showed a significantly higher RMSE of 66.22 and MAPE of 2.37%, despite achieving the same R^2. Notably, the ensemble model combining XGBoost and LGBM via the PSO-CS algorithm slightly improved performance with an RMSE of 21.45, MAE of 17.63, and MAPE of 0.63%, maintaining the same R^2 of 0.49. These results suggest that while the ensemble model provided marginal error reductions, the overall explanatory power (R^2) was moderate across models in this extended forecast horizon. Nevertheless, the consistently low MAPE values, especially under 1%, highlight the models' precision in predicting weekly maize prices in out-of-sample conditions.

The detailed results of the out-of-time validation are presented in Table 13, which summarizes the predictive performance of the individual models (XGBoost and LGBM) and the ensemble model (XGBoost + LGBM optimized via PSO-CS). The results obtained from this validation show a slight decline in model accuracy compared to the metrics reported in Tables 8 and 10 for the standard validation and test sets. This behavior is expected, as increasing the temporal distance between the training and prediction data also increases the uncertainty associated with structural, seasonal, or exogenous changes not captured in the original dataset. Nevertheless, the ensemble model (XGBoost + LGBM optimized with PSO-CS) maintained acceptable performance in terms of both absolute and percentage error, demonstrating adequate

Table 13
Out-of-Time Validation Results (October 2024 – March 2025). Predictive performance of individual models (XGBoost and LightGBM) and their ensemble (XGBoost + LGBM) evaluated on unseen data beyond the training period. Metrics reported: RMSE, MAE, MAPE, and R^2.

Model	Metric	Out-of-Time Test Set
XGBoost	RMSE	21.47
	MAE	17.44
	MAPE	0.62%
	R^2	0.49
LGBM	RMSE	66.22
	MAE	51.91
	MAPE	2.37%
	R^2	0.90
XGBoost + LGBM (PSO-CS)	RMSE	21.45
	MAE	17.63
	MAPE	0.63%
	R^2	0.49

Table 14
Real vs. Predicted Prices – Out-of-Time Evaluation. Comparison between actual maize prices and predicted values by XGBoost, LGBM, and the ensemble model (PSO-CS) for the period December 2024 to February 2025.

Date	Real Price	Predicted Price (XGBoost)	Predicted Price (LGBM)	Predicted Price (PSO-CS)
2024-12-11	2813.00	2773.84	2779.29	2774.69
2024-12-18	2803.00	2797.36	2798.90	2797.60
2024-12-27	2797.00	2797.94	2787.74	2796.36
2025-01-03	2822.00	2797.84	2795.43	2797.47
2025-01-08	2791.00	2811.94	2805.68	2810.97
2025-01-17	2796.00	2795.48	2796.20	2796.73
2025-01-22	2778.00	2795.14	2788.26	2796.63
2025-01-29	2820.00	2806.37	2803.79	2806.31
2025-02-05	2796.00	2810.69	2807.54	2810.62
2025-02-12	2844.00	2808.35	2806.29	2807.46
2025-02-21	2869.00	2806.84	2804.77	2806.78
2025-02-26	2878.00	2807.78	2804.56	2808.26

generalization capability. Although the coefficient of determination (R^2) was lower than in the standard tests, this is consistent with the additional challenge of forecasting future events that may exhibit different characteristics. Overall, the results suggest that the proposed model can adapt reasonably well to real-world and evolving scenarios.

The table 14 presents a detailed comparison between the actual maize prices and the predictions generated by the XGBoost, LGBM, and the ensemble model optimized using PSO-CS. It shows the results corresponding to a specific period within the out-of-time validation, covering December 2024 to February 2025. All values are expressed in Colombian pesos (COP) and reflect the behavior of maize prices in the Granabastos market, located in the Atlántico department. This comparison allows for the evaluation of each model's predictive accuracy against real-world data not used during training. The table highlights the models' ability to follow short-term fluctuations in weekly prices and detect turning points in price trends. It also demonstrates the superior consistency of the ensemble model in approximating actual values, particularly during volatile weeks. Furthermore, the closeness between predicted and observed values supports the model's capacity to generalize temporal and market-specific dynamics. These findings reinforce the practical applicability of the proposed ensemble for operational decision-making in local agricultural markets.

Step 2: Model Validation Using Data3-Corn and Data4-Corn

To validate the model, two different market time series, Data3-Corn and Data4-Corn, were used without making specific adjustments to the model parameters. The model employed for validation corresponds to LGBM, XGBoost, and their combination, as these achieved the best performance in the forecasting phase using Data1-Corn and Data2-Corn. This approach allows for the evaluation of the model's generalization ability, that is, its capacity to make accurate predictions in different market contexts without the need for additional customization or calibration.

By not adjusting the parameters for each series, overfitting is avoided, and a more objective measure of the model's performance in real world scenarios is obtained. This method is useful for assessing the robustness and versatility of the model when faced with diverse datasets.

Figure 10 presents the external validation results of the individual models LightGBM and XGBoost applied to the Data3-corn (Montería) dataset, while Figure Figure 11 shows the predictions generated by the combined models FCN + RNN and XGBoost + LightGBM on the same dataset. In both figures, the red line represents the actual test values, and the green line corresponds to the model predictions. This comparison allows for the evaluation of the individual versus ensemble model performance in an external validation scenario.

The results of the validation of the individual LGBM and XGBoost decision tree models and their combination are presented in Table 15.

The validation results using *Data3-Corn* provide evidence that the model is a reliable tool for predicting corn prices in Colombia. Using the XGBoost model yields a coefficient of determination R^2 of 0.85, indicating that the model can explain 85% of the variability of the test data, which is a good level of fit. The RMSE of 91.78 COP, an MAE of 74.62 COP, and a MAPE of 3.49% show smaller deviations and predictions closer to the actual values.

For *Data4-Corn*, the model shows reasonable performance. The R^2 of 0.92 indicates that the model is able to explain a significant part of the variability of corn prices. The RMSE of 66.92 COP, an MAE of 50.23 COP, and a MAPE of 2.12% show smaller deviations and predictions closer to the real values.

The model product of the combination of (XGBoost+LGBM) shows that for *Data3-corn*, a coefficient of determination R^2 of 0.85 is obtained, indicating that the model can explain 85% of the variability of the test data, which is a good level of fit. The RMSE of 86.66 COP, a MAE of 68.55 COP, and a MAPE of 3.16% show smaller deviations and predictions closer to the true values, which outperform the individual models. For *Data4-corn*, this combination obtains a coefficient of determination R^2 of 0.78, indicating that the model can explain 78% of the variability in the test data, which is an acceptable level of fit. The RMSE of 128.16 COP, a MAE of 107.16 COP, and a MAPE of 4.21% show smaller deviations and predictions closer to the true values, without outperforming the single XGBoost model.

Figure 12 shows the external validation of the individual XGBoost and LightGBM models applied to the Data4-corn (Barranquillita) dataset, Figure 13 presents the results obtained from the ensemble model XGBoost + LightGBM on the same dataset. In both figures, the red line represents the actual values from the test set, while the green line corresponds to the predictions generated by the models. This comparison allows for the evaluation of

(a) LightGBM on Data3-corn

(b) XGBoost on Data3-corn

Figure 10: Validation of individual models on Data3-corn (Montería). Forecasting results of LightGBM and XGBoost applied to an external validation dataset.

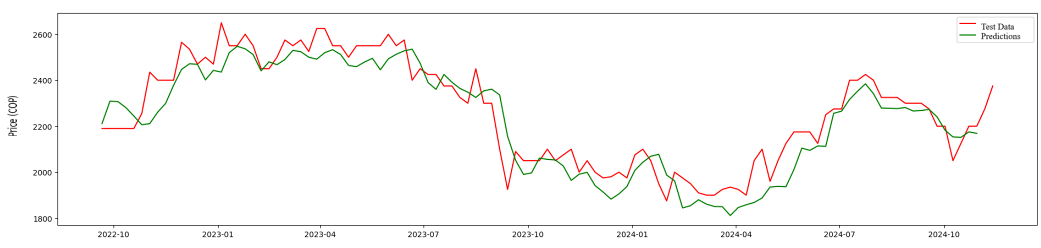

Figure 11: Ensemble model forecast on Data3-corn (Montería). Prediction results using the XGBoost + LightGBM combination optimized with PSO-CS.

individual model performance versus that of the combined model in an external validation scenario.

The results from the Data3-corn and Data4-corn datasets highlight the models' ability to generalize to unseen data. While individual models captured general trends, they struggled with abrupt changes. In contrast, the ensemble models provided more accurate and stable predictions, reinforcing the value of hybrid approaches for forecasting agricultural prices in variable real-world conditions.

Step 3 Comparative Analysis with Previous Studies Using the Same Time Series

When comparing the performance metrics of the proposed multivariate model with those reported by Wang et al. (2022) and Zeng et al. (2023) both based on the same CBOT time series it is evident that the PSO-CS optimized ensemble (XGBoost + LGBM) achieves a lower MAPE (2.06%) than Wang's (2.25%) and Zeng's (3.41%) models. This highlights the competitiveness of the proposed approach, which also outperforms the neural ensemble (RNN + FCN) developed in the same study. MAPE is particularly suitable for this comparison, as it provides a scale independent metric for consistent evaluation. These

Table 15
Performance of individual models on Data3-corn and Data4-corn. Evaluation based on RMSE, MAE, MAPE, and R^2 using external validation datasets.

Model	Metric	Data3-corn (Monteria) Test	Data4-corn (Barranquillita) Test
XGBoost	RMSE	91.78	66.2
	MAE	74.62	50.23
	MAPE	3.49%	2.12%
	R^2	0.85	0.92
LGBM	RMSE	99.15	241.29
	MAE	79.3	205.56
	MAPE	3.55%	7.62%
	R^2	0.8	0.43
(XGBoost+LGBM)	RMSE	86.66	128.16
	MAE	68.55	107.16
	MAPE	3.16%	4.21%
	R^2	0.85	0.78

(a) XGBoost on Data4-corn

(b) LightGBM on Data4-corn

Figure 12: Validation of individual models on Data4-corn (Barranquillita). Forecasting results of XGBoost and LightGBM applied to an external test set.

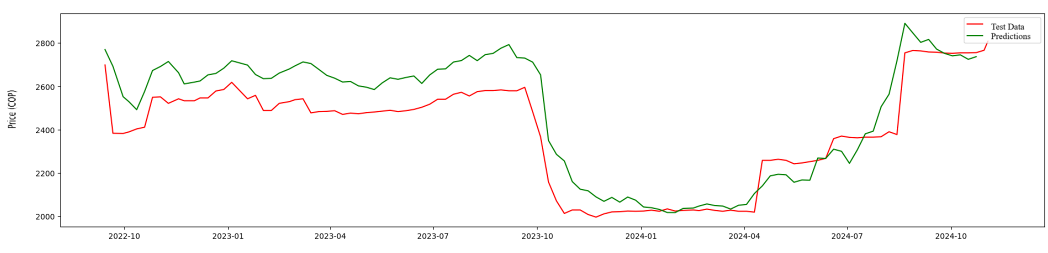

Figure 13: Ensemble model forecast on Data4-corn (Barranquillita). Prediction results using the XGBoost + LightGBM combination optimized with PSO-CS.

Table 16
Comparison of best-performing ensemble models from reviewed studies and this research. Highlights performance improvements achieved in this study relative to the literature.

Combined Model	MAPE (%)
PSO-CS (Wang et al., 2022)	2.25
PSO-CS (Zeng et al., 2023)	3.41
PSO-CS EMD Trees (This study)	2.06

results position the proposed model as a robust and accurate alternative within the current state of the art.Table 16 presents a comparison of the most accurate ensemble models reported in the literature with the one proposed in this study. The comparison highlights the improvement in MAPE achieved by the PSO-CS EMD Trees ensemble.

In summary, the comparative analysis confirms that the proposed multivariate ensemble model optimized with PSO-CS not only demonstrates superior predictive accuracy over individual and neural-based models but also outperforms benchmark approaches reported in recent literature using the same CBOT time series. The consistently lower MAPE values underscore the effectiveness of integrating heterogeneous models with multivariate data and advanced optimization techniques.

7. Implications and Scalability of the Proposed Model in the Colombian Context

The results obtained from the proposed model demonstrate not only methodological advancements in multivariable price forecasting but also significant practical potential. In the case of maize in Colombia, the model exhibited strong capabilities in anticipating price fluctuations, which is particularly critical given that approximately 80% of national demand is met through imports (UPRA, 2022), and that the prices of basic food products such as chicken and eggs are closely tied to the cost of this grain (Colombia Mercantile Exchange, 2023; Arbeláez et al., 2024). These findings support the feasibility of employing such tools to inform decisions regarding public procurement, subsidies, reference pricing, and early warning systems. Furthermore, the proposed application aligns with the guidelines of the National Development Plan 2022–2026 (Departamento Nacional de Planeación (DNP), 2022), which prioritizes food security, territorial productive planning, and the strengthening of agricultural information systems.

Given its flexible and scalable design, the proposed model can be adapted to other strategic agricultural commodities such as rice, soybeans, or coffee, as well as to different regions of the country with diverse productive and commercial dynamics. This would require adjusting the explanatory variables and incorporating additional data sources specific to each context, including factors such as logistics infrastructure availability or regional agricultural policies.

One of the main challenges identified for the practical implementation of this type of system in the Colombian agricultural sector is the effective integration of climatic and public policy variables elements not included in this initial stage of the study, but whose incorporation could significantly enhance the model's predictive capacity and broaden its applicability. Additionally, it is necessary to improve interoperability across data platforms, ensure the quality and frequency of collected information, and promote the adoption of advanced analytical tools among producers, associations, and public institutions. Addressing these barriers will help consolidate a more comprehensive and operational forecasting system, contributing to a more informed, resilient, and sustainable agricultural sector in Colombia.

8. Conclusions

This study highlights how combining multiple time series sources can represent a significant step forward in agricultural price forecasting, particularly for maize in Colombia. By combining local data, such as historical prices from the Granabastos market, with external information, such as maize futures from the Chicago Board of Trade (CBOT), the model was able to capture important interactions between domestic and global markets, ultimately leading to more accurate predictions.

The hybrid model developed in this work, which combines XGBoost and LightGBM and fine-tunes them using the PSO-CS algorithm, achieved a MAPE of 2. 06% in the multivariate test set (Data1-Corn and Data2-Corn). This performance significantly outperformed both the individual models and the neural network ensemble (FCN + RNN), which reached a much higher MAPE of 30.74%. These results underscore the importance of integrating information from diverse sources, particularly in environments where international market trends have a significant impact on local prices.

Even when tested in out of time scenarios from October 2024 to March 2025, the model maintained a high level of accuracy (MAPE = 0.63%), demonstrating its ability to adapt to new, unseen conditions. In external validations using datasets from other Colombian markets Data3-Corn (Montería) and Data4-Corn (Barranquillita) the model continued to perform well without requiring retraining, achieving MAPEs of 3.16% and 4.21%, respectively, along with R^2 values of 0.85 and 0.78. This consistency suggests strong generalization capacity across regions with distinct characteristics.

The strong performance of XGBoost and LightGBM can be attributed to their ability to model complex, nonlinear relationships, manage multicollinearity, and effectively incorporate external variables. In contrast, neural networks struggled due to their sensitivity to noise and the larger data volumes they typically require. Furthermore, applying Empirical Mode Decomposition (EMD) enabled the breakdown of the original time series into simpler components IMF1, IMF2, and the residual making it easier for the model to capture distinct temporal patterns.

Lastly, the use of PSO-CS for optimizing model weights played a key role in ensuring a balanced and stable ensemble. The sensitivity analysis confirmed the reliability of this optimization process across different configurations. Taken together, these findings demonstrate that integrating multivariate time series is not only a viable but also a powerful approach to enhancing prediction accuracy and model resilience in the face of market volatility. This work offers a valuable methodological contribution to the field of agricultural price forecasting, especially in contexts where local and international dynamics intersect.

Finally, the proposed model, based on the integration of multivariate time series, is adaptable to other agricultural commodities such as rice, potatoes, or beans, provided that relevant price data and exogenous variables are available. Its application in different regions will depend on the quality and availability of data, as well as institutional support for its management. The main challenges for its implementation in Colombia include technological gaps in rural areas, the need for technical training, and inter-institutional coordination to consolidate a robust agricultural information system. Nevertheless, the results show that it is possible to build highly accurate predictive tools that support strategic decision-making and contribute to the country's food security.

Limitations

Although the proposed approach has shown strength in terms of accuracy, stability, and multi-source integration, several limitations must be acknowledged. The model presents high computational complexity and depends on careful hyperparameter tuning, as well as the quality of the initial decomposition process (EMD). Additionally, while out-of-time validation and testing on external data sources not used during training supported its generalization capability, extending the model to other agricultural products or international markets still requires further validation and potential adjustments based on the specific characteristics of each time series. These considerations are essential to guide future implementations of the model in diverse contexts.

Credit authorship contribution statement

Adelaida Ojeda Beltrán: Conceptualization, methodology, Investigation, Data curation, Formal analysis, Writing - Original Draft. **Mario E. Suaza-Medina:** Validation, Visualization, Writing - Original Draft. **F. Javier Zarazaga-Soria:** Conceptualization, Formal analysis, Writing - Review & Editing **Emiro De La Hoz Franco:** Software, Supervision, Writing - Review & Editing **José Escorcia-Gutierrez:** Conceptualization, methodology, Validation, Formal analysis, Supervision, Writing - Review & Editing.

Acknowledgments

This article is part of the research project by the Ministry of Science, Technology, and Innovation of Colombia through the Bicentennial Doctoral Excellence Scholarship Program, as well as project T59_23R, funded by the Aragon Regional Government (Spain). The work of Mario Suaza has been partially supported by the Ministry of Science, Technology, and Innovation of Colombia (MINCIENCIAS 885/2020). This research was conducted by the Advanced Information Systems Group at Universidad de la Costa (Colombia) and the Sustainable Organizations Research Group at Universidad del Atlántico (Colombia).

Declaration of AI and AI-assisted technologies in the writing process

During the revision of this work, the authors used the writefull language check embedded in Overleaf to improve English grammar, proofread, and rephrase certain sentences in order to minimize the iThenticate similarity index. Following the use of these tools and services, the authors reviewed and edited the content as necessary, assuming full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this work.

References

Adebiyi, A. A., Adewumi, A. O., and Ayo, C. K. (2014). Comparison of arima and artificial neural networks models for stock price prediction. *Journal of Applied Mathematics*, 2014:Article ID 614342, 7 pages. DOI: 10.1109/CLOUD49709.2020.00009.

Adhikari, B., P, J. A., and C, S. C. (2021). Building market intelligence systems for agricultural commodities: A case study based on cardamom. In *2021 2nd International Conference on Secure Cyber Computing and Communications (ICSCCC)*, pages 12–16, Jalandhar, India. IEEE. DOI: 10.1109/INFOMAN.2017.7950359.

Arbeláez, M. A., Ramírez, S., Ayala, P., Montoya, N., Parra, J. J., and Silva, C. (2024). Trade policy of the yellow maize production chain in colombia. Fedesarrollo, Colombia. Document in Spanish. Retrieved from https://www.repository.fedesarrollo.org.co/bitstream/handle/11445/4669/Repor_Agosto_2024_Arbel%c3%a1ez_et%20al..pdf?sequence=1&isAllowed=y.

Bahri, M. Z. and Vahidnia, S. (2022). Time series forecasting using smoothing ensemble empirical mode decomposition and machine learning techniques. In *Proceedings of the International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*. IEEE. DOI: 10.1109/ICECCME55909.2022.9988336 .

Banerjee, T., Sinha, S., and Choudhury, P. (2022). Long term and short term forecasting of horticultural produce based on the lstm network model. *Applied Intelligence*, 52(8):9117–9147. DOI: 10.1007/s10489-021-02845-x.

Bedi, J. and Toshniwal, D. (2020). Energy load time-series forecast using decomposition and autoencoder integrated memory network. *Applied Soft Computing Journal*, 93:106390. DOI: 10.1016/j.asoc.2020.106390 .

Chaowalit, O., Selanon, W., Jaipanya, S., Nopparatmaitree, M., and Arampongsanuwat, S. (2025). Multivariate time series forecasting thailand soybean meal price with deep learning models. *ICIC Express Letters Part B: Applications*, 16(3):317–323. DOI: 10.24507/icicelb.16.03.317.

Chung, H. and Shin, K. S. (2018). Genetic algorithm-optimized long short-term memory network for stock market prediction. *Sustainability (Switzerland)*, 10(10). DOI: 10.3390/su10103765.

Colombia Mercantile Exchange (2023). Product analysis – maize market. Bolsa Mercantil de Colombia, Colombia. Document in Spanish. Published by the Corporate Management of Analytics and Business Intelligence. Retrieved from https://www.bolsamercantil.com.co/sites/default/files/2023-12/Analisis_de_producto_Maiz_2023.pdf.

Deepa, S., Alli, A., Sheetac, and Gokila, S. (2023). Machine learning regression model for material synthesis prices prediction in agriculture. *Materials Today: Proceedings*, 81:989–993. DOI: 10.1016/j.matpr.2021.04.327.

Departamento Nacional de Planeación (DNP) (2022). Plan nacional de desarrollo 2022–2026: "colombia, potencia mundial de la vida". Presidencia de la República de Colombia, Bogotá. Retrieved from https://www.dnp.gov.co/PND/Paginas/pnd-colombia-potencia-vida.aspx.

Food and Agriculture Organization of the United Nations (2020). The state of food and agriculture 2020: Overcoming water challenges in agriculture. FAO, Rome. Document in English. Retrieved from https://doi.org/10.4060/cb1447en.

Garai, S., Paul, R. K., Rakshit, D., Yeasin, M., Emam, W., Tashkandy, Y., and Chesneau, C. (2023). Wavelets in combination with stochastic and machine learning models to predict agricultural prices. *Mathematics*, 11(13):2896. DOI: 10.3390/math11132896.

Guindani, L. G., Oliveirai, G. A., Ribeiro, M. H. D. M., Gonzalez, G. V., and de Lima, J. D. (2024a). Exploring current trends in agricultural commodities forecasting methods through text mining: Developments in statistical and artificial intelligence methods. *Heliyon*, 10(23).

Guindani, T., Ortega, L., and Nishida, K. (2024b). Hybrid transformer with temporal encoding for long-term commodity price prediction. In *Proceedings of the 28th International Conference on Artificial Intelligence Applications*, Lecture Notes in Artificial Intelligence, pages 225–236. Springer. DOI: 10.1007/978-3-031-45210-3_18.

Hiyam, A., Alshibani, S. M., Ibrahim, O., and Elhag, A. A. (2024). Prediction opec oil price utilizing long short-term memory and multi-layer perceptron models. *Alexandria Engineering Journal*. DOI: 10.1016/j.aej.2024.10.057.

International Maize and Wheat Improvement Center (CIMMYT) and International Center for Tropical Agriculture (CIAT) (2019). *Maize for Colombia: Vision 2030*. International Maize and Wheat Improvement Center (CIMMYT) and International Center for Tropical Agriculture (CIAT), Colombia. Edited by Eduardo Figueroa. Retrieved from https://hdl.handle.net/10883/20382. Document in Spanish.

Ji, M., Liu, P., Deng, Z., and Wu, Q. (2022). Prediction of national agricultural products wholesale price index in china using deep learning. *Progress in Artificial Intelligence*, 11(1):121–129. DOI: 10.1007/s13748-021-00264-0.

Jin, B. and Xu, X. (2024). Forecasting wholesale prices of yellow corn through the gaussian process regression. *Neural Computing and Applications*, 36(15):8693–8710. DOI: 10.1007/s00521-024-09531-2.

Kundu, N., Rani, G., Dhaka, V. S., Gupta, K., Nayaka, S. C., Vocaturo, E., and Zumpano, E. (2022). Disease detection, severity prediction, and crop loss estimation in maize crop using deep learning. *Artificial Intelligence in Agriculture*, 6:276–291. DOI: 10.1016/j.aiia.2022.11.002.

Lee, J. and Xia, B. (2024). Analyzing the dynamics between crude oil spot prices and futures prices by maturity terms: Deep learning approaches to futures-based forecasting. *Results in Engineering*. DOI: 10.1016/j.rineng.2024.103086.

Li, B., Ding, J., Yin, Z., Li, K., Zhao, X., and Zhang, L. (2021). Optimized neural network combined model based on the induced ordered weighted averaging operator for vegetable price forecasting. *Expert Systems with Applications*, 168:114232. DOI: 10.1016/j.eswa.2020.1142327.

Li, J., Yang, X., Qian, T., and Xie, Q. (2023). The adaptive fourier decomposition for financial time series. *Engineering Analysis with Boundary Elements*, 150:139–153. DOI: 10.1016/j.enganabound.2023.01.037.

Liu, Z., Chai, T., Tang, J., and Yu, W. (2020). Signal analysis of mill shell vibration based on variational modal decomposition. In *Proceedings of the 39th Chinese Control Conference*, pages 1168–1173. DOI: 10.23919/CCC50068.2020.9189103.

Mensi, W., Rehman, M. U., Gök, R., Gemici, E., and Vo, X. V. (2025). Risk spillovers and diversification benefits between crude oil and agricultural commodity futures markets. *Research in International Business and Finance*, 73:102579.

Mohanty, M. K., Thakurta, P. K. G., and Kar, S. (2023). Agricultural commodity price prediction model: a machine learning framework. *Neural Computing and Applications*, 35(20):15109–15128. DOI: 10.1007/s00521-023-08528-7.

Nafkha, R. and Suchodolska, D. Ż. (2024). Modelling and estimating volatility using conditional variance-type models. *Procedia Computer Science*, 246:2655–2664.

Ojeda-Beltran, A., De-La-Hoz-Franco, E., and Escorcia-Gutierrez, J. (2024). Decomposition models for agricultural commodity price time series: A comparative research. In *Colombian Conference on Computing*, pages 234–248. Springer.

Organisation for Economic Co-operation and Development (OECD) and Food and Agriculture Organization of the United Nations (FAO) (2021). *OECD-FAO Agricultural Outlook 2021-2030*. OECD Publishing, Paris. Retrieved from https://doi.org/10.1787/19428846-en, DOI: 10.1787/19428846-en.

Purohit, S. K., Panigrahi, S., Sethy, P. K., and Behera, S. K. (2021). Time series forecasting of price of agricultural products using hybrid methods. *Applied Artificial Intelligence*, 35(15):1388–1406. DOI: 10.1080/08839514.2021.1981659.

Rana, H., Farooq, M. U., Kazi, A. K., Baig, M. A., and Akhtar, M. A. (2024). Prediction of agricultural commodity prices using big data framework. *Engineering Technology and Applied Science Research*, 14(1):12652–12658. DOI: 10.48084/etasr.6468.

Sachithra, V. and Subhashini, L. D. C. S. (2023). How artificial intelligence uses to achieve agriculture sustainability: Systematic review. *Artificial Intelligence in Agriculture*, 8:46–59. DOI: 10.1016/j.aiia.2023.04.002.

Safari, N., Price, G. C. D., and Chung, C. Y. (2020). Analysis of empirical mode decomposition-based load and renewable time series forecasting. In *2020 IEEE Electric Power and Energy Conference (EPEC)*, pages 1–6. IEEE. DOI: 10.1109/EPEC48502.2020.9320072.

Sari, M., Duran, S., Kutlu, H., Guloglu, B., and Atik, Z. (2024). Various optimized machine learning techniques to predict agricultural commodity prices. *Neural Computing and Applications*, 36(19):11439–11459. DOI: 10.1007/s00521-024-09679-x.

Sediyono, E., Hartomo, K. D., Arthur, C., Utami, I., Prabowo, R., and Chiong, R. (2025). An integrated framework for multi-commodity agricultural price forecasting and anomaly detection using attention-boosted models. *Journal of Agriculture and Food Research*, page 102021.

Suaza, M. E., Zarazaga-Soria, F. J., Pinilla-Lopez, J., Lopez-Pellicer, F. J., and Lacasta, J. (2023). Effects of data time lag in a decision-making system using machine learning for pork price prediction. *Neural Computing and Applications*, 35:19221–19233. DOI: 10.1007/s00521-023-08730-7.

Unidad de Planificación Rural Agropecuaria (UPRA) (2022). Productive planning plan for the maize chain in colombia. Ministry of Agriculture and Rural Development, Bogotá, Colombia. Document in Spanish. Retrieved from https://www.andi.com.co/Uploads/POP_Cadena_Maiz.pdf.

Wang, J., Wang, Z., Li, X., and Zhou, H. (2022). Artificial bee colony-based combination approach to forecasting agricultural commodity prices. *International Journal of Forecasting*, 38(1):21–34. DOI: 10.1016/j.ijforecast.2019.08.006.

World Food Programme (WFP) (2024). Food security assessment for the colombian population – 2024. World Food Programme (WFP), Rome. Document in English. Retrieved from https://es.wfp.org/publicaciones/evaluacion-de-la-seguridad-alimentaria-para-la-poblacion-colombiana-2024.

Wu, L., Chang, P., Chen, W., and Pei, T. (2024). Generalized load graphical forecasting method based on modal decomposition. *Global Energy Interconnection*, 7(2):166–178.

Yang, Y. and Yang, Y. (2020). Hybrid method for short-term time series forecasting based on eemd. *IEEE Access*, 8:61915–61929. DOI: 10.1109/ACCESS.2020.2983588.

Yoon, H. J., Kim, D., and Lee, S. J. (2025). Data-driven analysis of climate impact on tomato and apple prices using machine learning. *Heliyon*, 11(6). DOI: 10.1016/j.heliyon.2024.e20987.

Zeng, L., Zhang, X., and Liu, F. (2023). Informer-based transformer model for agricultural price forecasting in asian commodity markets. *Expert Systems with Applications*, 210:118657. DOI: 10.1016/j.eswa.2023.118657.

Zhang, D., Chen, S., Ling, L., and Xia, Q. (2020). Forecasting agricultural commodity prices using model selection framework with time series features and forecast horizons. *IEEE Access*, 8:28197–28209. DOI: 10.1109/ACCESS.2020.2971591 .

Zhang, L., Liang, R., et al. (2025). Avocado price prediction using a hybrid deep learning model: Tcn-mlp-attention architecture. *arXiv preprint arXiv:2505.09907*.

Zhang, T. and Tang, Z. (2024). Agricultural commodity futures prices prediction based on a new hybrid forecasting model combining quadratic decomposition technology and lstm model. *Frontiers in Sustainable Food Systems*, 8:1334098. DOI: 10.3389/fsufs.2024.1334098.

Zhao, H. (2021). Futures price prediction of agricultural products based on machine learning. *Neural Computing and Applications*, 33:837–850. DOI: 10.1007/s00521-020-05399-9.

Zheng, G., Zhang, H., Han, J., Zhuang, C., and Xi, L. (2020). The research on agricultural product price forecasting service based on combination model. In *2020 ieee 13th international conference on cloud computing (cloud)*, pages 4–9. IEEE.