

Article

Utilizing Ensemble Learning Techniques to Enhance Corn Price Prediction: A Case Study on South Dakota

Jihene Kaabi , Youssef Harrath *  and Ethan Price

Beacom College of Computer and Cyber Sciences, Dakota State University, Madison, SD 57042, USA;
jihene.kaabi@dsu.edu (J.K.); ethan.price@dsu.edu (E.P.)

* Correspondence: youssef.harrath@dsu.edu

Abstract

Predicting crop prices is a complex challenge that farmers must navigate each year, but machine learning algorithms can provide valuable insights to support more informed decision making. In recent years, agricultural price prediction models have made significant advances, with architectures achieving varying degrees of success. However, ensuring the accuracy and reliability of these models remains an ongoing challenge. This research explores the use of stacking, an ensemble learning technique, to enhance the performance of base models in predicting corn prices in many regions of the state of South Dakota in the USA. We propose a hybrid architecture that combines Long Short-Term Memory Networks and Transformer Neural Networks, allowing us to leverage the strengths of both models. Using real corn price data from the last 11 years, our findings show that this stacked architecture not only outperforms its individual base models but also industry standard approaches. The ensemble model can predict within a one-month window, reaching a maximum mean average error of 6%. More research is needed to use multiple features, such as weather, crude oil price, and market demand, to provide an efficient decision support system.

Keywords: corn price prediction; ensemble learning; LSTM, transformer; ridge regression; stacking



Academic Editors: Xi Qiao, Shuo Zhang and Cong Wang

Received: 26 September 2025

Revised: 19 October 2025

Accepted: 20 October 2025

Published: 22 October 2025

Citation: Kaabi, J.; Harrath, Y.; Price, E. Utilizing Ensemble Learning Techniques to Enhance Corn Price Prediction: A Case Study on South Dakota. *Agronomy* **2025**, *15*, 2447. <https://doi.org/10.3390/agronomy15112447>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Agriculture has long been a cornerstone of developing civilizations and remains a key factor in the economic development of developing nations. Technological advances in agriculture have provided extraordinary opportunities for farmers to increase productivity, reduce risks, and optimize resources. Among these advancements, machine learning and deep learning algorithms significantly help farmers make informed decisions every day. From pest monitoring and plant disease detection to crop yield prediction, these technologies are transforming agricultural practice [1].

Despite advances in agricultural technology and market analytics, a significant challenge remains: identifying the optimal time to sell crops. The high volatility of crop prices presents substantial difficulties for farmers trying to navigate market conditions, often leading to severe consequences for their financial stability and overall well-being. Price fluctuations in agricultural commodities have long been a central concern within the economic domain. However, accurate prediction of crop prices remains an unresolved issue, largely due to the complex interplay of factors influencing prices, such as crop yield, weather variability, government policies, and global market dynamics [2]. This

persistent uncertainty carries not only economic ramifications but also psychological impacts. Research has shown that financial stress and the unpredictability associated with agricultural income contribute significantly to mental health challenges, including an increased risk of suicide among farmers [3,4].

Although agriculture is essential for any nation, it is a particularly vital component of the United States' economy, contributing 5.6% to the U.S. gross domestic product and positioning the country as a major player in global agricultural exports. Rural and urban communities across the country rely heavily on the stability of crop markets, and Midwest states such as South Dakota play an especially prominent role. In South Dakota, agriculture serves as the backbone of the economy, providing a substantial portion of the population with jobs and financial stability.

South Dakota's agricultural landscape is uniquely positioned, emphasizing commodities such as corn, soybeans, and cattle. The state is sixth nationally in corn production, with more than 42 million acres of farmland, more than 28,000 farms, and USD 12.9 billion in market value of agricultural products sold. Notably, 81% of these farms are family-owned or operated by individuals, which strengthens the personal connection to both financial risks and mental health challenges that farmers face. This makes South Dakota an ideal case study for addressing crop price prediction, with a focus on developing models specifically tailored to corn prices.

As mentioned above, corn markets are highly volatile and prices frequently fluctuate. This volatility is illustrated in Figure 1, which shows the historical price of corn in the United States over the past 59 years. The markets for any agricultural commodity can vary widely. Although national trends provide a general overview, local markets, such as those of South Dakota, can exhibit significantly different behaviors.




Figure 1. Historical corn prices in the United States.

Addressing this problem through machine learning-based commodity price prediction offers a potential solution to financial and mental health problems. Using data-driven models, we can provide farmers with the insights needed to make better and more educated decisions. By implementing this, we can alleviate the financial risk and agrarian stress that come with agriculture. Therefore, in this paper, we propose a hybrid machine learning model that uses an ensemble learning technique known as stacking and combines several models to predict South Dakota corn prices.

The rest of the paper proceeds as follows. Section 2 provides a comprehensive review of the current literature on the application of machine learning techniques to corn price

prediction. Section 3 details the methodological framework, including data sources, model architecture, and the evaluation metrics used. Section 4 presents the experimental findings in two parts: the first analyzes the performance of the model on historical corn price data, while the second evaluates the applicability and effectiveness of the model in real-world market scenarios. Finally, Section 5 summarizes the key contributions of the study and outlines directions for future research in this domain.

2. Literature Review

The prediction of corn prices has gained increasing attention due to its economic importance and the inherent volatility of agricultural markets. This section explores the progression of computational approaches applied to this problem, beginning with traditional machine learning models that have laid the groundwork for data-driven price forecasting. It then delves into the use of deep learning techniques, particularly Long Short-Term Memory (LSTM) networks and their variants, which have shown promise in capturing temporal dependencies in price data. The review also examines ensemble learning strategies that combine multiple predictive models to improve accuracy and robustness. Finally, it highlights the key gaps and limitations in existing research, identifying areas where current methodologies fall short and opportunities for future innovation remain.

2.1. Traditional Machine Learning Approaches

Early efforts in corn price forecasting frequently employed linear models and machine learning techniques such as ARIMA, Support Vector Machines (SVM), and Random Forests (RF). Yuan and Ling [5] evaluated the performance of the ARIMA, LSTM, Prophet, and XGBoost models in agricultural commodities, reporting reasonable accuracy but limited generalization capabilities. Similarly, Chen et al. [6] developed a web-based automated system to predict agricultural commodity prices using historical data from Malaysia. Five machine learning models: ARIMA, SVR, Prophet, XGBoost, and LSTM were compared through two experimental phases. The LSTM model achieved the best performance and was selected as the system's prediction engine. The paper did not explore hybridization or dynamic integration strategies. Gaur et al. [7] applied SVM and RF to the prediction of the price of corn, incorporating engineering characteristics such as weather and yield. Despite improved input quality, the RF model plateaued at a performance ceiling ($R^2 \approx 0.08$), suggesting limitations in its ability to capture non-linear temporal dependencies. Wibowo et al. [8] investigated the use of Ridge Regression with damping factors, demonstrating reduced error metrics compared to basic linear models. Although this work focused on general food commodities, its regularization strategy inspired the use of Ridge Regression as a meta-learner in our proposed ensemble model. Though these models have practical appeal due to their interpretability, they are often insufficient for modeling complex, dynamic market behaviors, especially in scenarios where feature interdependencies evolve over time.

2.2. Deep Learning Models: LSTM and Variations

Long Short-Term Memory (LSTM) networks have emerged as the dominant architecture for modeling time series data due to their ability to capture both short- and long-term dependencies. Quan and Shi [9] demonstrated improved forecasting accuracy using a hybrid CEEMDAN LSTM model for financial time series. Sabu and Kumar [10] showed that LSTM outperformed traditional models such as SARIMA and Holt-Winters in the context of predicting the price of areca nuts. Further extensions, such as Long- and Short-Term Time-series Network (LSTNet) proposed in [11], integrated convolutional

layers and skip connections to enhance temporal modeling across multivariate data. However, standard LSTM models often assume that historical price patterns persist over time, an assumption that can be problematic in rapidly shifting markets. To address this, Cai et al. [12], Zhou et al. [13] proposed adaptive LSTM variants that update predictions based on recent residuals or external signals, achieving improved results on oil and agricultural datasets. Despite their improvements, these models tend to operate as standalone learners. Their inability to dynamically interact with other model types (e.g., attention-based networks) limits their performance, particularly in heterogeneous forecasting environments.

2.3. Ensemble Learning in Price Prediction

Ensemble learning has gained traction as a strategy to enhance the robustness of the model by aggregating diverse learners. Techniques such as stacking, boosting, and bagging are increasingly employed in time-series domains. Adhikari et al. [14] developed an ensemble that integrates ARIMA and multivariate regression for cardamom price prediction, although their model suffered a high *RMSE* due to poor interaction between the base model. Sankareswari et al. [15] showed that the ensemble models outperformed the single learners when the base models were optimally configured. Shahhosseini et al. [16] confirmed similar benefits in corn yield prediction, noting that the ensemble methods were more resistant to noise and seasonality.

In the specific domain of corn price, Ribeiro et al. [17] combined neural networks and K-Nearest Neighbors (KNN) through differential evolution to reduce bias and variance. Silva et al. [18] employed ensemble models with LSTM and Support Vector Regression (SVR) to forecast corn and sugar prices, reporting improved short-term accuracy. However, these models typically relied on simplistic averaging techniques or heuristic weightings, with limited exploration of meta-learning strategies.

Transformer-based models represent a newer frontier in time series forecasting. Unlike LSTM, which processes input sequentially, transformers use self-attention mechanisms to model global dependencies in parallel. Xiao et al. [19] compared Transformer, GRU, and LSTM for stock price prediction, finding that LSTM achieved higher accuracy in data-rich local contexts, while Transformers underperformed without sufficient feature diversity. Similarly, Gade et al. [20] applied transformer-based models to crop yield and price prediction, demonstrating their potential for long-term prediction but highlighting their limitations in short sequences. Although promising, this architecture has not been widely adopted in ensemble systems focused specifically on corn pricing. Moreover, studies that leverage transformer-based architectures in hybrid frameworks remain scarce.

Advanced ensemble strategies such as stacking have proven effective in other domains but are rarely applied in agriculture. Sagi and Rokach [21] outlined how stacking can dynamically weight the output of the base learner through a meta-model, increasing robustness across volatile datasets. However, few studies explore stacking architectures that integrate statistical learners, LSTM, and transformers within a single framework, especially for agricultural forecasting.

A comparative summary of forecasting models, their strengths and limitations, and how our proposed model addresses these gaps is presented in Table 1.

Table 1. Comparative summary of corn price forecasting models.

Model	Strengths	Weaknesses	Role in This Study
ARIMA	Ideal for short-term forecasting of stationary and linear time series	Lacks the flexibility to handle non-linear dynamics, external factors, or long-range patterns	Used for baseline comparison in the literature; not implemented in our model
LSTM	Capable of modeling non-linear relationships and short-term temporal dependencies.	Limited in capturing long-range dependencies; may overfit to recent trends	Combined with Transformer to improve long-term sequence modeling
XGBoost	Powerful for non-linear regression and classification tasks	Lacks built-in mechanisms for modeling temporal dependencies in time series data.	Included in the literature review for contrast; not part of the final architecture
Transformers	Highly effective at modeling long-range dependencies through self-attention	Require large datasets and may underperform on short, noisy sequences	Used as a base learner to model extended temporal patterns
Ridge Regression (Meta-Learner)	Effectively manages multi-collinearity, limits overfitting, and enables efficient model ensemble integration.	Lacks the flexibility to capture non-linear or sequential patterns	Serves as the meta-learner in the stacking ensemble to combine base model outputs.
Traditional Ensembles	Improve robustness through model combination; simple to implement	Use static weights or averaging; lack a learning mechanism	Addressed by stacking ensemble with trainable meta-learner
Proposed Hybrid (Ours)	Integrates LSTM and Transformer outputs using Ridge Regression; adaptive and robust	Limited to univariate input; future work needed for multivariate extension	Core contribution of this study; improves accuracy over individual models

2.4. Gaps and Limitations in Existing Research

Although significant progress has been made, several gaps persist in the literature. Firstly, many models such as those used in [7,14] do not perform well due to the lack of adaptation of the dynamic model and the effective integration of the model output. Secondly, ensemble methods remain underexplored in the context of corn, with most work either applying basic averaging or integrating only two models.

Furthermore, research tends to focus on short-term price prediction ([17,18]), leaving the long-term prediction underdeveloped. There is also a scarcity of studies that combine deep learning models (e.g., LSTM, Transformer) with advanced ensemble techniques. Xiao et al. [19] suggest LSTM outperforms other architectures, but ensemble strategies that include Transformers and GRUs remain largely theoretical.

Furthermore, studies like Chen et al. [6] propose scalable solutions, but few address real-time processing or integration of various data sources such as satellite, weather, and market sentiment data. Finally, while multivariate features are occasionally included ([7,8]), they are not central to ensemble learning strategies. Our study addresses these challenges by introducing a novel stacking ensemble model that integrates LSTM, Transformer, and Ridge Regression. The LSTM captures short-term temporal dynamics, the Transformer models long-range dependencies, and Ridge Regression serves as a meta-learner to optimally combine outputs. This structure goes beyond basic averaging by learning to weight base predictions adaptively, improving robustness and generalization.

In addition, our model is trained on 11 years of South Dakota-specific corn price data, providing a highly localized and practically relevant solution. It also bridges short-term (1-month) and medium-term forecasting horizons, offering farmers better tools for operational and strategic decisions.

3. Methodology

This study employs a hybrid ensemble learning architecture that utilizes the stacking method to predict corn prices. Stacking, also known as stacking generalization, integrates multiple predictive models, known as base learners, whose outputs are then combined through a meta-learner to generate final predictions [22]. This methodology takes advantage of the diverse strengths of individual predictive models, effectively improving predictive accuracy and stability compared to standalone models.

3.1. Data Collection and Pre-Processing

This study uses 11 years of historical corn price data (2013–2024), consisting of 2807 daily records from the state of South Dakota. We used a 80:20 train-test split for the data. The dataset, sourced from Barchart, Chicago, IL, USA (<https://www.barchart.com/> (accessed on 9 December 2024)), includes corn prices recorded on business days, excluding weekends and holidays.

No normalization or feature scaling was applied, as the raw price values across the dataset fell within a consistent and comparable range. For modeling with LSTM and Transformer architectures, a rolling window of the last 10 time steps was used to predict the next price point, aligning with common practices in sequence-based forecasting.

Hyperparameter tuning was performed separately for both the LSTM and the Transformer models to identify the optimal configuration for each. This tuning process focused on improving accuracy by adjusting parameters such as the learning rate, number of units, number of layers, and batch size, ensuring that both models operated under conditions that maximized predictive performance. More details about hyperparameter tuning are given in Section 3.4.

3.2. Base Models

The stacking ensemble consists of two primary base models: Long Short-Term Memory (LSTM) neural networks and Transformer neural networks. These models were selected on the basis of their complementary capabilities in processing time series data.

- **Long Short-Term Memory (LSTM):** LSTM is a variant of Recurrent Neural Networks (RNN) specifically developed to address the problem of vanishing gradients encountered in traditional RNNs. LSTM networks incorporate special gating mechanisms—input, forget, and output gates, that allow the network to retain relevant historical information and eliminate unnecessary data. This enables effective modeling of long-term dependencies present in sequential data, which makes LSTM particularly suitable for the prediction of corn prices, where historical patterns significantly influence future values [23].
- **Transformer Neural Networks:** Introduced in [24], transformer models are based on self-attention mechanisms, which allow the model to dynamically weigh the importance of each data point in the input sequence. Unlike recurrent architectures, Transformers process sequences concurrently, enabling the capture of intricate temporal relationships without the sequential processing bottleneck. This characteristic makes Transformers highly effective at capturing complex dependencies and improving prediction accuracy for time-series data, such as corn prices.

3.3. Stacking and Meta-Learner

The stacking ensemble combines the predictions generated independently by the LSTM and Transformer base models. The outputs from these base learners serve as inputs to the meta-learner, Ridge Regression.

- **Stacking Mechanism:** Stacking works by training the base models independently on historical corn price data. The base models each produce predictions on a validation set.

These predictions become inputs to the meta-learner by aggregating the predictions of the base models and feeding them into the meta-learner. The meta-learner then learns to optimally combine these predictions, effectively harnessing the strengths and mitigating individual weaknesses of the base models [22].

- **Meta-Learner (Ridge Regression):** Ridge Regression, a linear regression model with $L2$ regularization, acts as the meta-learner. This method penalizes large coefficients through regularization, which helps reduce overfitting and improve the generalizability of the ensemble model [25]. When evaluating different meta-learners, penalizing the weights was a main factor so no coefficient became too powerful. $L2$ regularization adds a penalty term to the loss function that is equal to the sum of the squares of its coefficients. The decision to go with Ridge Regression as the meta-learner was both for its penalization of coefficients, balancing each baseline model's results, and success in other studies like [26]. Wang and Lu [26] utilized ridge regression as the meta-learner in a stacked ensemble model, discovering that the fewer base models combined, the stronger ridge regression worked as a meta-learner. By combining predictions from LSTM and Transformer models, Ridge Regression optimally weights each base model's output, leading to a more robust and accurate predictive model for corn prices.

3.4. Hyperparameter Tuning

Hyperparameter tuning is a process where a model's hyperparameters are selected to return the strongest result. Each base model underwent an independent version of this process before being aggregated into the meta-learner. The meta-learner then also completed this as well. Hyperparameter tuning is a tedious process, having to explore a large search space. It is extremely computationally expensive and can take time to complete. However, it is vital for the success of machine learning models as it ensures both the base learners and the meta-learner were optimized for best performance. As summarized in Table 2, the tuning process explored a defined search space for each parameter and identified the optimal value. Traditionally, there are two primary ways to complete hyperparameter tuning: Exhaustive Grid Search (EGS) and Randomized Parameter Optimization (RPO). In this study, we have decided to implement EGS as the way to complete the hyperparameter tuning process. EGS evaluates all parameter combinations, determining the strongest set by comparing them with the model's evaluation metrics found in Section 4.1. While this search space is large and the process is expensive, we found the benefit of searching each combination critical in evaluating the models.

Table 2. Hyperparameter tuning results for each model.

Model	Parameter	Search Range	Optimal Value
LSTM Model	units	32, 64, 128	units = 128
	batch_size	8, 16, 32	batch_size = 8
	epochs	20, 40, 60	epochs = 40
	dropout	0.0, 0.2, 0.4	dropout = 0.0
	activation	relu/tanh	activation = tanh
Transformer Model	num_heads	2, 4, 8	num_heads = 4
	dense_dim	32, 64, 128	dense_dim = 32
	dropout	0.1, 0.2, 0.3	dropout = 0.1
	learning_rate	0.0005, 0.001	learning_rate = 0.001
	batch_size	16, 32	batch_size = 32
	epochs	20, 40	epochs = 20
Stacking Meta-Learner (Ridge)	alpha	0.001, 0.01, 0.1, 1, 10, 100	alpha = 0.001
	solver	auto, svd, cholesky, lsqr, sparse_cg, saga	solver = auto
	fit_intercept	True/False	fit_intercept = True

3.5. Combined Model Architecture

The overall hybrid ensemble architecture involves the following:

1. Independently train LSTM and Transformer base models on preprocessed historical corn price data.
2. Generation of intermediate predictions from each base model.
3. Combine these intermediate predictions using Ridge Regression as the meta-learner to produce the final forecast.

This structured ensemble approach, as shown in Figure 2, leverages both the sequential modeling capabilities of LSTM and the comprehensive self-attention features of the Transformer architecture. The final combination via Ridge Regression capitalizes on the diverse strengths of each base model, significantly enhancing the accuracy and reliability of corn price predictions.




Figure 2. Stacking hybrid model architecture.

4. Experimental Study

This section details the implementation of the methodology introduced in Section 3. We begin by describing the evaluation metrics used to measure error in the regression problem. Next, we divide the discussion into two parts: the testing portion and the implementation portion, analyzing the performance differences between the two. The program was developed in Python 3.13, using the TensorFlow 2.20 and Keras 3.11.3 libraries as the foundation for building and training the neural networks. The experiment was carried out on a Lenovo ThinkPad running Windows 10, equipped with an 11th Gen Intel® Core™ i7-1165G7 processor (2.80 GHz, 2803 MHz, four cores, eight logical processors). Visual Studio Code 1.105 served as the Integrated Development Environment (IDE).

4.1. Evaluation Metrics

The performance of the model was assessed using several standard evaluation metrics summarized in Equations (1)–(4). In each of these metrics, n represents the amount of data used in the evaluation, y_i is the true value, and $\hat{y}_i$ is the predicted value.

- Mean Squared Error (MSE): Measures the mean squared differences between predicted and actual values, highlighting the larger errors significantly.

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (1)$$

- Root Mean Squared Error (RMSE): The square root of MSE, providing error values in the same units as the original data, making interpretation more intuitive.

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

- Mean Absolute Percentage Error (MAPE): Represents average absolute error in percentage terms, useful for assessing relative prediction accuracy.

$$MAPE = \frac{100}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (3)$$

- Coefficient of Determination (R^2 Score): Indicates how well the predicted values align with actual values, with a value closer to 1 indicating better accuracy.

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (4)$$

4.2. Traditional Machine Learning Models

This section describes traditional machine learning models that have been included to provide a baseline of comparison. While this study focuses on the benefits of the stacking ensemble technique, we can continue to see its benefits when compared to other traditional methods for price forecasting, especially in the crop markets. Further details about specific parameters for each model can be found in Table 3.

- Linear Regression: Linear Regression is one of the most fundamental statistical learning methods used for forecasting continuous outcomes. It models the dependent variable y as a linear combination of independent variables X , represented at $y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_n x_n + \epsilon$, where β_i are the model coefficients and ϵ is the error term. The Ordinary Least Squares (OLS) method is used to minimize the sum of squared residuals, optimizing the coefficients. Linear regression is a standard baseline to compare with for forecasting [27].
- Support Vector Regression: Support Vector Regression (SVR) extends the principles of Support Vector Machines, originally introduced in [28], by employing a kernel-based approach to linear and nonlinear relations. SVR finds a central hyperplane $f(x) = \langle w, x \rangle + b$ that deviates from the actual target values by at most a specified margin, denoted by ϵ . SVR then uses a kernel function, such as Radial Basis Function (RBF), polynomial, or sigmoid, that enables the model to project data into higher-dimensional spaces, creating a linear fit. SVR is another traditional machine learning model that can be found in time-series forecasting, making it a strong baseline to compare against.

- Extreme Gradient Boosting (XGBoost): XGBoost, standing for Extreme Gradient Boosting, is an optimized ensemble method based on gradient-boosted decision trees introduced in [29]. XGBoost sequentially builds an ensemble of weak learners, typically shallow decision trees. These weak learners attempt to correct the residual errors of the previous generations by minimizing a differentiable loss function through gradient descent. XGBoost has become a staple ensemble method to test as a baseline through the years to check for improvements.

Table 3. Hyperparameters for traditional models.

Model	Parameter	Search Range	Optimal Value
Linear Regression	All parameters	Default (no tuning required)	Default configuration
Support Vector Regression (SVR)	C	0.1, 1, 10	C = 1
	epsilon	0.001, 0.01, 0.1	epsilon = 0.01
	kernel	linear, rbf, poly	kernel = linear
XGBoost Regressor	learning_rate	0.001, 0.01, 0.1	learning_rate = 0.01
	max_depth	3, 5, 7	max_depth = 3
	n_estimators	100, 200, 300	n_estimators = 200
	subsample	0.8, 0.9, 1.0	subsample = 1.0

4.3. Testing Results

This section presents a comprehensive evaluation of the predictive performance of each model (i.e., traditional machine learning models, base models, and the stacked ensemble model) on the testing dataset. Importantly, during the testing, each model was provided with the true sequence values up to the point of prediction. Later, further testing was completed to examine how the models behave with the generated values. We compare individual models, namely LSTM and Transformer, against a stacked ensemble architecture that integrates their outputs. The results, summarized in Table 4, are evaluated using four key metrics: Mean Squared Error (*MSE*), Root Mean Squared Error (*RMSE*), Mean Absolute Percentage Error (*MAPE*), and the Coefficient of Determination (R^2 Score). Through this analysis, our objective was to highlight the strengths, limitations, and relative effectiveness of each model in capturing and predicting corn price trends.

Table 4. Comparative analysis of model performance metrics.

Model	<i>MSE</i>	<i>RMSE</i>	<i>MAPE</i> (%)	R^2 Score
Linear Regression	0.007	0.081	1.364	0.987
XGBoost Regressor	0.020	0.140	2.107	0.962
SVR	0.007	0.082	1.383	0.987
LSTM	0.012	0.111	1.590	0.991
Transformer	0.048	0.219	3.555	0.968
Stacking Ensemble	0.003	0.055	0.940	0.998

The comparative results in Table 4 show the strengths and limitations of each type of model. Among the traditional machine learning models, we still see the competitive performance with Linear Regression and SVR demonstrating similar error rates $MSE = 0.007$ and $RMSE \approx 0.081$ – 0.082 and low *MAPE* 1.36–1.38%. We can attribute this similarity to the linear kernel-function used by the SVR. While generally strong in handling complex nonlinearities, XGBoost recorded a slightly high error with an *MSE* of 0.02 and *MAPE* of 2.107%. This indicates that the lack of features contributed to poor

performance and do not provide sufficient depth for gradient-boosted trees to fully utilize to their advantage.

The standalone LSTM model achieved notable performance, yielding an *MSE* of 0.012, an *RMSE* of 0.111, and a *MAPE* of 1.59% as shown in Table 4. Furthermore, the LSTM achieved a very high R^2 score of 0.991, indicating a robust capacity to capture and predict temporal dependencies in corn price data effectively. These results, reflected in Figure 3, show the strength of the LSTM model in handling sequential data, making it a reliable baseline for time series forecasting.

In comparison, the Transformer model exhibited comparatively weaker performance, with an *MSE* of 0.048 and *RMSE* of 0.219, representing more than four times the error magnitude observed in the LSTM results. The Transformer's *MAPE* was also significantly higher at 3.56%, indicating less accurate predictions relative to the actual corn prices. Although the R^2 score of 0.968 remains high, it is significantly lower than that of the LSTM and ensemble model, suggesting that although the Transformer can capture complex dependencies through its self-attention mechanism, additional fine-tuning or adaptation to the dataset could enhance its predictive precision. This poor performance can also be related to other factors. That is, the limited dataset size is a huge issue to overcome with the Transformer model. Transformers rely on extensive datasets to fully make use of their self-attention mechanism. Furthermore, the limited features in the dataset can also have constrained the model's performance, preventing the model from generalizing as strongly as other models. Figure 4 displays the inaccuracies and weaknesses of the model compared to the other two models.

The stacked ensemble model significantly outperformed both individual base models as well as the traditional models in all evaluation metrics, demonstrating remarkable improvements with an *MSE* of 0.003 and an *RMSE* of 0.055 as shown in Table 4. This reflects a substantial reduction in predictive error compared to the LSTM and Transformer models. Figure 5 supports this claim, showing a closer result to the predicted values and the true values. The *MAPE* of the ensemble of 0.94% further emphasizes its superior precision, indicating that the predictions deviate by less than 1% on average from actual corn prices, a level of precision practically meaningful for agricultural decision making. Additionally, the ensemble's R^2 score of 0.998 is exceptionally close to perfect, underscoring its reliability and efficacy in accurately modeling corn price variations.

While a 1% *MAPE* score may seem statistically minor, applying to a practical situation like crop prices can show its impact. In South Dakota, the average market price of corn has fluctuated around \$4.50–\$5.00 per bushel. A forecasting error of 1% corresponds then to an average deviation of \$0.045–\$0.05 per bushel. For a mid-size producer harvesting 200,000 bushels annually, this deviation equates to a potential prediction variance of \$9000–\$10,000 in projected revenue. Comparing this with traditional models like SVR and XGBoost, whose *MAPEs* are between 1.3% and 2.1%, respectively, the ensemble's improvement of nearly one percentage point could reduce forecasting uncertainty by \$2000–\$5000 annually. Therefore, we can determine that even a minute change in *MAPE* is the difference between thousands of dollars for local farmers.

The improvement observed with the ensemble model highlights the value of combining various predictive techniques. Specifically, integrating the strengths of the sequential memory capacity of the LSTM with the global dependency capture of the Transformer resulted in significantly higher predictive performance. Furthermore, using Ridge Regression as a meta-learner allowed effective weighting of individual model predictions, leveraging their complementary strengths and systematically minimizing prediction errors.

Overall, this analysis demonstrates that the stacking ensemble architecture provides significant practical advantages in accurately and reliably predicting corn prices. By mitigating individual model weaknesses and capitalizing on their respective strengths, the ensemble offers robust predictive power critical for strategic agricultural planning and decision making.




Figure 3. Comparison between the actual and predicted corn prices using the LSTM model.




Figure 4. Comparison between the actual and predicted corn prices using the Transformer model.




Figure 5. Comparison between the actual and predicted corn prices using the Ensemble model.

4.4. Implemented Results

After evaluating the performance of the model with sequences derived from the actual values of the corn price, we examined how these models performed when forecasting current corn prices. For each day, we tracked both the true price of corn and the corresponding predictions made by our models. We then calculated the absolute percent

error for each day to evaluate the performance of each model over the following months. Our results indicate an initial improvement in performance by the stacked ensemble model within the first month, which subsequently declines.

Table 5 presents the corn price data and the relative error results for the first month. By averaging the errors during this period, we found that the ensemble model demonstrated the lowest error, with a Mean Absolute Percentage Error (*MAPE*) of 5.25%. In comparison, the LSTM and Transformer base models showed slightly higher errors, with *MAPE* values of 5.49% and 5.61%, respectively. This indicates that, in the short-term prediction horizon of one month, our stacked ensemble architecture consistently outperforms the individual base models.

Table 5. One-month forecasting errors.

Date	Actual	LSTM	Transformer	Ensemble	LSTM Error	Transformer Error	Ensemble Error
15 January 2025	4.1307	4.08519	3.99710	4.02816	1.10	3.23	2.48
16 January 2025	4.0866	4.11898	4.00153	4.06386	0.79	2.08	0.56
17 January 2025	4.1834	4.15653	3.99990	4.10412	0.64	4.39	1.90
20 January 2025	4.2058	4.18362	3.99650	4.13337	0.53	4.98	1.72
21 January 2025	4.2282	4.21259	3.99138	4.16480	0.37	5.60	1.50
22 January 2025	4.1707	4.24393	3.99500	4.19794	1.76	4.21	0.65
23 January 2025	4.2223	4.27415	3.99488	4.23024	1.23	5.39	0.19
24 January 2025	4.1898	4.30499	3.99443	4.26323	2.75	4.66	1.75
27 January 2025	4.1448	4.33725	3.99810	4.29735	4.64	3.54	3.68
28 January 2025	4.1773	4.37031	4.00008	4.33249	4.62	4.24	3.72
29 January 2025	4.2935	4.40336	3.99007	4.36872	2.56	7.07	1.75
30 January 2025	4.2258	4.43742	3.98933	4.40517	5.01	5.60	4.24
31 January 2025	4.1422	4.47212	3.98805	4.44236	7.96	3.72	7.25
3 February 2025	4.2378	4.50736	3.98681	4.48012	6.36	5.92	5.72
4 February 2025	4.2945	4.54332	3.98579	4.51863	5.79	7.19	5.22
5 February 2025	4.2820	4.57999	3.98520	4.55787	6.96	6.93	6.44
6 February 2025	4.3263	4.61737	3.98418	4.59789	6.73	7.91	6.28
7 February 2025	4.2488	4.65549	3.98306	4.63872	9.57	6.25	9.18
10 February 2025	4.2880	4.69439	3.98186	4.68038	9.48	7.14	9.15
11 February 2025	4.2137	4.73406	3.98016	4.72293	12.35	5.54	12.09
12 February 2025	4.2762	4.77454	3.97807	4.76636	11.65	6.97	11.46
13 February 2025	4.3184	4.81596	3.97681	4.81072	11.52	7.91	11.40
14 February 2025	4.3459	4.85826	3.97549	4.85604	11.79	8.52	11.74

Table 6 extends the analysis by presenting results for approximately seven additional weeks. Over this extended period, a clear limitation of our models becomes evident as they struggle to accurately capture price relationships past the one-month mark. Using the same evaluation metric as in Table 5, we observed significantly higher error rates for the LSTM and ensemble models, with *MAPEs* of 41.73% and 43.23%, respectively. In contrast, the Transformer model maintained a substantially lower *MAPE* of 3.12%. However, despite the Transformer's lower numerical error, a closer examination reveals notable shortcomings. Specifically, as seen in Tables 5 and 6, the Transformer's predictions fluctuate minimally, typically within a narrow 2–4 cent range, failing to respond dynamically to actual market variations.

Several factors are likely to contribute to the performance issues. Firstly, unlike the testing methodologies described in Section 4.3, these forecasts are based on previously predicted prices as input for subsequent predictions. This recursive nature is a central issue in any field involved in price forecasting, as poor predictions feed back into the model. This results in compounding errors, where the accumulation of incorrect predictions builds

on each other. Table 6 demonstrates this idea. The error values for both the LSTM and Ensemble start low, 12.04% and 12.06%, respectively, growing linearly as the prediction date increases. Furthermore, this model suffers significantly with the limited features of the dataset. These models are trained purely on the corn price value, with no additional features to support it. While still a strong performance from the Ensemble model in short-term forecasting, all models lose valuable information concerning outside factors that affect corn prices (e.g., crude oil price, annual corn yield, weather, government influence). These two primary issues in the long-term forecasting lead to a model that struggles to capture the dynamic relationship, particularly when relying solely on the price of corn as the single-input feature.

In conclusion, our analysis demonstrates that, for short-term predictions within a one-month time frame, the stacked ensemble model consistently exhibits the lowest error rate, highlighting the potential of hybrid ensemble architectures for near-term forecasting of corn prices.

Table 6. Additional forecasting errors (over one-month).

Date	Actual	LSTM	Transformer	Ensemble	LSTM Error	Transformer Error	Ensemble Error
17 February 2025	4.3748	4.90149	3.97418	4.90234	12.04	9.16	12.06
18 February 2025	4.4037	4.94564	3.97285	4.94963	12.31	9.78	12.40
19 February 2025	4.3619	4.99072	3.97149	4.99792	14.42	8.95	14.58
20 February 2025	4.3711	5.03675	3.97005	5.04722	15.23	9.17	15.47
21 February 2025	4.3036	5.08372	3.96857	5.09754	18.13	7.78	18.45
24 February 2025	4.1969	5.13163	3.96705	5.14887	22.27	5.48	22.68
25 February 2025	4.1682	5.18058	3.96550	5.20129	24.29	4.86	24.79
26 February 2025	4.1537	5.23051	3.96396	5.25478	25.92	4.57	26.51
27 February 2025	4.0264	5.28143	3.96248	5.30931	31.17	1.59	31.86
28 February 2025	3.9598	5.33332	3.96098	5.36489	34.69	0.03	35.48
3 March 2025	3.8917	5.38633	3.95945	5.42166	38.41	1.74	39.31
4 March 2025	3.8206	5.44066	3.95791	5.47985	42.40	3.59	43.43
5 March 2025	3.8745	5.49618	3.95634	5.53930	41.86	2.11	42.97
6 March 2025	3.9574	5.55288	3.95475	5.60003	40.32	0.07	41.51
7 March 2025	4.0099	5.61076	3.95314	5.66200	39.92	1.42	41.20
10 March 2025	4.0475	5.66964	3.95153	5.72506	40.08	2.37	41.45
11 March 2025	4.0300	5.72938	3.94990	5.78903	42.17	1.99	43.65
12 March 2025	3.9395	5.79002	3.94826	5.85396	46.97	0.22	48.60
13 March 2025	3.9868	5.85149	3.94661	5.91978	46.77	1.01	48.48
14 March 2025	3.9482	5.91376	3.94495	5.98646	49.78	0.08	51.63
17 March 2025	4.0004	5.97686	3.94327	6.05403	49.41	1.43	51.34
18 March 2025	3.9779	6.04072	3.94157	6.12240	51.86	0.91	53.91
19 March 2025	4.0260	6.10526	3.93985	6.19151	51.65	2.14	53.79
20 March 2025	4.0988	6.17043	3.93812	6.26130	50.54	3.92	52.76
21 March 2025	4.0519	6.23616	3.93638	6.33167	53.91	2.85	56.26
24 March 2025	4.0544	6.30236	3.93462	6.40256	55.44	2.95	57.92
25 March 2025	3.9938	6.36895	3.93284	6.47387	59.47	1.53	62.10
26 March 2025	3.9292	6.43586	3.93106	6.54552	63.80	0.05	66.59
27 March 2025	3.9181	6.50301	3.92925	6.61742	65.97	0.28	68.89
28 March 2025	3.9540	6.57031	3.92743	6.68948	66.17	0.67	69.18
31 March 2025	4.0220	6.63767	3.92559	6.76162	65.03	2.40	68.12
1 April 2025	4.1127	6.70503	3.92373	6.83374	63.03	4.59	66.16

4.5. Discussion on Model Limitation and Potential Improvements

From Sections 4.3 and 4.4, we not only see the successes of the ensemble technique, but also the limitations the model faces. As previously mentioned, the current architecture has two primary issues that limit its ability to forecast corn prices past one month periods: recursive forecasting nature and data limitations. In Section 4.4, the stacking ensemble model requires the base model's predictions as input to make predictions about future prices. However, as inputs drift away from true values and only contain predicted ones, the errors of those predicted values begin to compound on one another. This leads to future prediction results being affected by current, high error predictions, and ultimately, the outcome of the new predictions is a greater error rate. This cycle is then repeated over and over, showing the limitation of the model to capture corn price relationships over one month. Furthermore, the relationships found for corn price are also affected by the limited dataset only containing one feature. Corn prices are affected by a myriad of factors, both local and global. Locally, corn price variations are affected by the supply–demand of corn in the area, the annual corn yield, and many weather metrics. Globally, corn price deals with many market factors, crude oil, ethanol, livestock feed, each raising and lowering the prices. It is this diverse portfolio that makes this market so volatile. Therefore, it is reasonable to assume that a single-feature input is not sufficient to capture these complex relationships.

However, many improvements can be implemented and tested to fix these limitations. Firstly, grouping multiple individual datasets to create a new dataset that has diverse features would be advantageous to this price forecasting problem. By utilizing these outside factors, we would expect that model performance past one month would improve as more complex relationships can be embedded not only in the base models but also the meta-learner. Furthermore, improvement towards the recursive nature of forecasting could be made by adjusting the sequences that serve as input to the base models. For example, extending the length of time steps for sequence creation could improve the forecasting ability by providing true values for longer periods of time. Additional enhancements with how the model handles sequences could be experimented on, trying different input–output configurations (i.e., what is being inputted into the model and how it is outputting results).

5. Conclusions and Future Work

This study presented an innovative stacking ensemble architecture that integrates Long Short-Term Memory (LSTM) networks and Transformer neural networks, combined through a Ridge Regression meta-learner, to accurately predict corn prices in South Dakota. The proposed model demonstrated substantial improvements over the individual base models and achieved remarkable precision, with a Mean Absolute Percentage Error (MAPE) of less than 1%. Our results underscore the effectiveness of ensemble learning techniques in capturing complex temporal patterns inherent in agricultural price data.

However, despite these advances, there remain opportunities to further enhance predictive accuracy and practical applicability. Future research should prioritize the incorporation of additional diverse features such as weather patterns, crude oil prices, market demand fluctuations, and global economic indicators. Integrating these multivariate datasets can enhance the model's ability to capture more intricate relationships influencing corn prices, providing more robust and accurate predictions.

Furthermore, future work should aim to develop a comprehensive Decision Support System (DSS) specifically tailored for farmers and agricultural stakeholders. This DSS could take advantage of real-time data streams, allowing dynamic decision-making support and timely adjustments based on changing market conditions.

Finally, focusing predictive models within a smaller geographic radius can further improve predictive accuracy by accounting for hyper-local variations in market conditions

and buyer-specific basis prices. This localized approach would empower farmers with precise insights relevant to their specific markets, significantly improving economic stability and reducing agrarian stress.

In conclusion, our proposed ensemble method provides a strong foundation for sophisticated agricultural price prediction. Continued research efforts, particularly around the incorporation of additional predictive features, the development of practical decision support tools, and the optimization of predictions at local scales, promise significant benefits for both economic outcomes and farmer well-being in agricultural communities.

Author Contributions: Conceptualization, J.K., Y.H. and E.P.; methodology, J.K., Y.H. and E.P.; software, E.P.; validation, J.K., Y.H. and E.P.; formal analysis, J.K., Y.H. and E.P.; investigation, E.P.; resources, E.P.; data curation, Y.H.; writing—original draft preparation, E.P.; writing—review and editing, J.K., Y.H. and E.P.; visualization, E.P.; supervision, J.K. and Y.H.; project administration, Y.H.; funding acquisition, Y.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Dakota State University through the funding program “Rising II for Faculty Retention” under grant number 81R203.

Data Availability Statement: The datasets presented in this article are not readily available because we signed an agreement with the company owning the data to not share it outside the University. Requests to access the datasets should be directed to Barchart, Chicago, IL, USA (<https://www.barchart.com/>).

Acknowledgments: We edited the paper using the overleaf with an embedded Writefull tool to check the grammar. We used ChatGPT 4o to help create an organizational flow to the paper, helping with the structure (i.e., sections, subsections, layout). Furthermore, ChatGPT 4o helped to rewrite sentences that felt similar to each other when comparing the results of the different models.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Pandey, C.; Sethy, P.K.; Behera, S.K.; Vishwakarma, J.; Tande, V. Smart agriculture: Technological advancements on agriculture—A systematical review. In *Deep Learning for Sustainable Agriculture*; Academic Press: Cambridge, MA, USA, 2022; pp. 1–56.
2. Westcott, P.C.; Hoffman, L.A. *Price Determination for Corn and Wheat: The Role of Market Factors and Government Programs*; Technical Bulletin No. 1878; AgEcon Search: St. Paul, MN, USA, 1999.
3. Guo, H.; Woodruff, A.; Yadav, A. Improving lives of indebted farmers using deep learning: Predicting agricultural produce prices using convolutional neural networks. In Proceedings of the 34th AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 13294–13299.
4. Scheyett, A.; Marburger, I.L.; Scarrow, A.; Hollifield, S.M.; Dunn, J.W. What Do Farmers Need for Suicide Prevention: Considerations for a Hard-to-Reach Population. *Neuropsychiatr. Dis. Treat.* **2024**, *20*, 341–352. [[CrossRef](#)] [[PubMed](#)]
5. Yuan, C.Z.; Ling, S.K. Long short-term memory model based agriculture commodity price prediction application. In Proceedings of the 2020 2nd International Conference on Information Technology and Computer Communications, Kuala Lumpur, Malaysia, 12–14 August 2020; pp. 43–49.
6. Chen, Z.; Goh, H.S.; Sin, K.L.; Lim, K.; Chung, N.K.H.; Liew, X.Y. Automated agriculture commodity price prediction system with machine learning techniques. *arXiv* **2021**, arXiv:2106.12747 [[CrossRef](#)]
7. Gaur, S.; Mahajan, J.; Sharma, M.; Hussain, D.; Kakani, R.; Saxena, A. Precision Corn Price Prediction with Advanced ML Techniques. In Proceedings of the 2024 International Conference on Trends in Quantum Computing and Emerging Business Technologies, Pune, India, 22–23 March 2024; pp. 1–5.
8. Wibowo, A.; Yasmina, I.; Wibowo, A. Food price prediction using time series linear ridge regression with the best damping factor. *Adv. Sci. Technol. Eng. Syst. J.* **2021**, *6*, 694–698. [[CrossRef](#)]
9. Quan, P.; Shi, W. Application of CEEMDAN and LSTM for Futures Price Forecasting. In Proceedings of the 2024 International Conference on Machine Intelligence and Digital Applications, Ningbo, China, 30–31 May 2024; pp. 249–255.
10. Sabu, K.M.; Kumar, T.M. Predictive analytics in Agriculture: Forecasting prices of Arecanuts in Kerala. *Procedia Comput. Sci.* **2020**, *171*, 699–708. [[CrossRef](#)]
11. Ouyang, H.; Wei, X.; Wu, Q. Agricultural commodity futures prices prediction via long-and short-term time series network. *J. Appl. Econ.* **2019**, *22*, 468–483. [[CrossRef](#)]

12. Cai, Y.; Zhang, N.; Zhang, S. GRU and LSTM Based Adaptive Prediction Model of Crude Oil Prices: Post-COVID-19 and Russian Ukraine War. In Proceedings of the 2023 6th International Conference on Computers in Management and Business, Macau, China, 13–15 January 2023; pp. 9–15.
13. Zhou, J.; Ye, J.; Ouyang, Y.; Tong, M.; Pan, X.; Gao, J. On Building Real Time Intelligent Agricultural Commodity Trading Models. In Proceedings of the 2022 IEEE Eighth International Conference on Big Data Computing Service and Applications (BigDataService), Newark, CA, USA, 15–18 August 2022; pp. 89–95.
14. Adhikari, B.; Sobin, C. Building Market Intelligence Systems for Agricultural Commodities: A Case Study based on Cardamom. In Proceedings of the 2021 2nd International Conference on Secure Cyber Computing and Communications (ICSCCC), Jalandhar, India, 21–23 May 2021; pp. 12–16.
15. Sankareswari, K.; Sujatha, G. Evaluation of an Ensemble Technique for Prediction of Crop Yield. In Proceedings of the ICIMMI 2023: 5th International Conference on Information Management & Machine Intelligence, Jaipur, India, 23–25 November 2023. [[CrossRef](#)]
16. Shahhosseini, M.; Hu, G.; Archontoulis, S.V. Forecasting corn yield with machine learning ensembles. *Front. Plant Sci.* **2020**, *11*, 1120. [[CrossRef](#)] [[PubMed](#)]
17. Ribeiro, M.H.D.M.; Ribeiro, V.H.A.; Reynoso-Meza, G.; dos Santos Coelho, L. Multi-objective ensemble model for short-term price forecasting in corn price time series. In Proceedings of the 2019 International Joint Conference on Neural Networks (IJCNN), Budapest, Hungary, 14–19 July 2019; pp. 1–8.
18. Silva, R.F.; Barreira, B.L.; Cugnasca, C.E. Prediction of corn and sugar prices using machine learning, econometrics, and ensemble models. *Eng. Proc.* **2021**, *9*, 31.
19. Xiao, J.; Deng, T.; Bi, S. Comparative analysis of LSTM, GRU, and transformer models for stock price prediction. In Proceedings of the International Conference on Digital Economy, Blockchain and Artificial Intelligence, Guangzhou, China, 23–25 August 2024; pp. 103–108.
20. Gade, S.; Singh, A.; Patil, S. Enriching Crop Yield and Price Prediction Using Transformers. *Int. J. Res. Eng. Sci. Manag.* **2024**, *7*, 87–93.
21. Sagi, O.; Rokach, L. Ensemble learning: A survey. *WIREs Data Min. Knowl. Discov.* **2018**, *8*, e1249. [[CrossRef](#)]
22. De Alwis, T.P.; Samadi, S.Y. Stacking-based neural network for nonlinear time series analysis. *Stat. Methods Appl.* **2024**, *33*, 901–924. [[CrossRef](#)]
23. Sherstinsky, A. Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. *Phys. D Nonlinear Phenom.* **2020**, *404*, 132306. [[CrossRef](#)]
24. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA, 4–9 December 2017; pp. 5998–6008.
25. Ahmed, A.M.; Sharma, E.; Jui, S.J.J.; Deo, R.C.; Nguyen-Huy, T.; Ali, M. Kernel ridge regression hybrid method for wheat yield prediction with satellite-derived predictors. *Remote Sens.* **2022**, *14*, 1136. [[CrossRef](#)]
26. Wang, Q.; Lu, H. A novel stacking ensemble learner for predicting residual strength of corroded pipelines. *Npj Mater. Degrad.* **2024**, *8*, 87. [[CrossRef](#)]
27. Montgomery, D.C.; Peck, E.A.; Vining, G.G. *Introduction to Linear Regression Analysis*, 5th ed.; John Wiley & Sons: Hoboken, NJ, USA, 2012.
28. Drucker, H.; Burges, C.J.C.; Kaufman, L.; Smola, A.; Vapnik, V. Support Vector Regression Machines. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 1997; Volume 9, pp. 155–161.
29. Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; ACM: New York, NY, USA, 2016; pp. 785–794.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.