Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

An Intelligent Market Price Forecasting and Sustainability Assessment for Anthurium Floriculture: A Data-Driven Decision Support Approach

Soysa H. R. T., W.S.A. De Silva
Sri Lanka Institute of Information Technology (SLIIT), Malabe, Sri Lanka

Peiris W. V. D., Gamage P. U.
Sri Lanka Institute of Information Technology (SLIIT), Malabe, Sri Lanka

Abstract**—**The markets dealing with perishable ornamental crops such as Anthurium flowers show dynamic pricing changes on a daily basis, and these changes cannot easily be explained by existing theories. Floriculturists engaged in such markets, especially at smaller and medium scale levels, are continuously confronted by the need to make selling decisions without any quantitative method by which they can evaluate if the current or expected prices represent a truly good opportunity to sell. This paper proposes a system for market intelligence designed to solve this problem. The main component of this system is a price prediction engine built with a sequential deep learning model and ensemble regression layers using historical price data collected for eight consecutive days for each shop, variety and size combination to predict the daily prices. The proposed system is based on a market scoring index known as SAMIS, which will provide market assessments after predicting prices according to its statistical price sustainability, momentum and price risks. SAMIS results are evaluated as a composite score between zero and hundred and are further classified according to a four-category classification of recommendations for different market situations ranging from ideal selling opportunities to periods where deferral of transactions is suggested. Evaluations of the system were done using real transactional data available in Sri Lankan Anthurium markets and showed stability in forecasting prices along with proper market assessment results. To our best knowledge, this study is the first to introduce a systematic approach for assessing the sustainability of ornamental flower markets.

Index Terms-Anthurium, floriculture market analytics, price forecasting, market sustainability, decision support, SAMIS, Sri Lanka, agricultural informatics, time-series prediction

1. INTRODUCTION

In the world of tropical cut floriculture, Anthurium has become one of the more commercially difficult-to-manage crops. Careful attention must be paid to growing conditions for the flowers, as well as their relatively high market value and extended vase life. This type of flowering plant is cultivated on a limited scale in the Gampaha, Matale, and Kalutara districts of Sri Lanka. In this area, the climate, soil, and social conditions enable growers to earn considerable incomes and have contributed to the development of a modestly successful business sector. As in any commercial operation, success is determined not just by the quality and amount of products produced, but also by day-to-day processes at local flower markets.

However, one of the main problems encountered in such businesses is the lack of structured information concerning current market prices. This is because fluctuations may occur every day, based on factors such as festival demand increase and temporary oversupply of flowers from other provinces. Usually, local farmers try to find out about prices through calling intermediaries on a daily basis or visiting local market collection centers during sales. Unfortunately, in most cases, harvesting and selling flowers take place without consideration of current price levels. This leads to losses and low profitability due to a failure to properly match the supply to the demand.

Although a lot of research has been done in recent years, the field of price prediction for agrarian products is still far from being well-understood and implemented effectively in ornamental floriculture, especially concerning Anthurium flowers. This is because the majority of models created thus far were used to forecast staple crop prices, where rich histories, futures contracts, and accessible data helped develop effective methods. However, since ornamental flower markets are localized and scattered, there is little reliable information about the price levels available in the open source literature. Studies that have looked at such markets found that while a single model approach failed to describe the dynamics of prices, ensembles did better.

Apart from predicting price accurately, there is an even larger hole that this paper intends to plug. While a price prediction model can give the most precise forecast, it cannot tell the farmer whether or not the predicted price signifies a favorable period for sale. For instance, while a highly accurate predicted price is high when compared to yesterday's price, it may actually be an outlier in relation to the price history of the crop in question due to sudden increase in demand. In contrast, a less predicted price in a stable market environment is more reliable than a higher price in a highly dynamic market environment. However, there is no distinction between such scenarios in most price forecasting models used by farmers and growers, thus creating a need for a solution to be sought in this research.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

This paper will make two major contributions. Firstly, it will develop a price prediction model based on a sequential fusion framework that leverages structured historical price sequences to produce daily Anthurium price predictions. Secondly, the more significant contribution of this paper is the development of SAMIS (market sustainability scoring index), which measures the predicted price on three independent statistical dimensions and classifies the predicted market into one of the four action categories. The structure of the rest of the paper is as follows: related works in Section 2; methodology in Section 3; experimental results in Section 4; and discussion and conclusion in Sections 5 and 6 respectively.

2. RELATED WORK

2.1 Agricultural Price Forecasting

Early attempts at price prediction relied mostly on statistics. The main approach was represented by ARIMA models and seasonally adjusted versions of ARIMA models until well into the 2000s. Although these approaches are still used, they are known to struggle with structural volatility, particularly in non-stationary data sets. Essentially, the problem lies in the fact that ARIMA-family models require linearity and stationarity – both of which rarely occur in the spot price market for perishable goods with demand shocks, harvest losses, and periodic cycles of supply.

The next logical step would be machine learning algorithms. In fact, studies using Support Vector Regression, Gradient Boosting trees, and Random Forest Regression models proved that the machine learning approach can surpass linear models and show better results in predictions for prices of vegetables, grains, and fruits. However, this approach relies on cross-sectional regression analysis; hence, any kind of sequencing, like weekly cycles of demand and slow price changes over weeks, needs to be accounted for with manual introduction of lagged data.

Finally, deep learning algorithms came up with the solution for the sequencing problems with recurrent neural networks. Long Short-Term Memory models showed a great potential in predicting time series and, unlike earlier algorithms, required no additional manual work in creating lagged features. Multiple studies analyzing vegetable and grain prices found that sequence models could outperform classical models on the data set large enough and stable enough. Recently, even more sophisticated methods that involved training ensembles on outputs from sequence models have proven successful due to the combination of advantages: while the former captures complex temporal dynamics impossible to infer using the features themselves, the latter allows accounting for changes in distributions and non-temporal factors.

2.2 Floriculture and Ornamental Crop Markets

Floriculture markets have not received sufficient research attention. Several investigations into cut flower price prediction using machine learning and hybrid models have been undertaken in the Dutch and Kenyan export market environments where flower price information is easily available. Such research has largely demonstrated the use of machine learning and hybrid models in flower pricing problems, but such studies apply to wholesale export markets that do not correspond with the smaller retail and semi-wholesale markets common in the case of the Sri Lankan Anthurium flower business. Research studies focused on Anthurium price prediction are virtually non-existent.

2.3 Market Sustainability and Risk Assessment

The concept of market sustainability in the context of agricultural price forecasting has received almost no formal treatment in the literature. Existing works that touch on market risk typically employ simple volatility measures — standard deviation of historical prices, or price range ratios — as contextual metadata rather than as components of a structured decision framework. The Coefficient of Variation, defined as the ratio of standard deviation to mean price, has appeared in commodity market analysis as a normalized risk indicator, but its integration into scoring mechanisms for producer decision support has not been systematically explored.

Statistical approaches to price stability assessment, particularly those grounded in the Empirical Rule and z-score interval logic, have been applied in financial risk management but have not been translated into agricultural market scoring frameworks. The SAMIS index proposed in this paper draws on these statistical foundations and organizes them into a coherent, decomposable scoring structure designed specifically for the operational constraints and decision contexts of small-scale floriculture producers.

2.4 Gap Analysis and Positioning of This Work

Table 1 summarizes the positioning of this work relative to representative studies in the broader literature. The comparison criteria reflect the dimensions along which the present system claims novelty: hybrid model architecture, floriculture domain applicability, and the presence of a market sustainability assessment mechanism beyond raw price prediction.

Study / System	Crop / Domain	Hybrid Model	Sequence Modeling	Floriculture Specific	Sustainability Scoring	Sell Recommendation
ARIMA-based systems [3]	General crops	No	Partial	No	No	No
SVR / RF approaches [4,5]	Vegetables, grains	Partial	No	No	No	No

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

LSTM standalone [6]	Grain, commodity	No	Yes	No	No	No
LSTM – XGBoost fusion [7]	Commodity market	Yes	Yes	No	No	No
Cut flower demand studies [8]	Export floriculture	No	Partial	Yes	No	No
CV-based risk tools [9]	Commodity markets	No	No	No	Partial	No
Proposed System (This Work)	**Anthurium, Sri Lanka**	**Yes**	**Yes**	**Yes**	**Yes (SAMIS)**	**Yes (4-class)**

Table 1. Comparative analysis of related systems and positioning of the proposed work.

3. ETHODOLOGY

The system operates in two sequential stages. In the first stage, a predicted market price is generated for a given combination of date, shop, flower variety, and size. In the second stage, that predicted price is evaluated against the statistical properties of the historical price record for the same product combination, producing a SAMIS score and a four-class market recommendation. Each stage is described below in full detail.

3.1 Data and Feature Preparation

3.1.1 Dataset

The dataset comprises daily Anthurium transaction records collected from retail shops in Sri Lanka. Each record captures the transaction date, the shop identifier, the flower variety, the size grade, the seller price per stem in Sri Lankan Rupees (LKR), and the number of units sold that day. Prior to model training, categorical fields were standardized to lowercase and stripped of whitespace. Date fields were parsed using day-first formatting, and records with missing price or date entries were excluded. The resulting dataset provides the historical price trajectories used for both model training and SAMIS computation.

3.1.2 Lookback Window and Sequence Construction

Price prediction at a given target date is conditioned on the eight most recent trading days of price history for the same shop–variety–size combination. This eight-day window was selected as a practical balance between temporal depth and data availability, since some product combinations have limited trading frequency. When fewer than eight historical records are available prior to the target date, the earliest available price is repeated forward to fill the remaining positions. This edge-padding approach preserves sequence shape without introducing artificial directional bias.

For each position i within the eight-day window, a five-dimensional feature vector is constructed from the price sequence. The features are:

Feature	Symbol	Description
Current price	*p*	The actual recorded seller price at position i
Lag-1	*p(i-1)*	Price from the previous day — captures immediate momentum
Lag-2	*p(i-2)*	Price from two days prior — captures short-term trend direction
Lag-4	*p(i-4)*	Price from four days prior — captures mid-week cycle effects
EMA-4	*EMA4(p[0:i])*	Exponential moving average with span 4, computed up to position i — represents smoothed trend

Table 2. Five-dimensional feature vector constructed at each position of the price sequence.

The Exponential Moving Average with span s is computed recursively. Given a smoothing factor $\alpha = 2 / (s + 1)$, for span 4 this gives $\alpha = 0.4$:

$$EMA_t = \alpha \cdot p_t + (1 - \alpha) \cdot EMA_{t-1}$$

The higher weighting of recent observations under $\alpha = 0.4$ means the EMA responds quickly to genuine trend changes while attenuating isolated daily fluctuations. Including EMA alongside raw lag values gives the sequence model access to both the actual historical price at each step and a noise-reduced representation of recent price direction. Together, these five features form the input to the sequential learning component of the prediction engine.

3.2 Price Prediction Architecture

3.2.1 Sequential Feature Extraction

The eight-day, five-feature input tensor is first processed by a Long Short-Term Memory network trained to learn temporal regularities in Anthurium price sequences. LSTM networks are well suited to this task because of their gating mechanism, which allows the model to selectively retain or discard information as it progresses through the time series. In markets where weekly demand cycles, gradual seasonal drift, and occasional structural breaks are all present, this selective memory is particularly useful.

At each time step t, the LSTM cell updates its internal state using three gates. The forget gate determines what proportion of the previous cell state to retain, the input gate controls how

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

much of the new candidate information to incorporate, and the output gate determines what portion of the updated cell state is exposed as the hidden state. These computations proceed as follows:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$$
$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$
$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$$
$$\tilde{c}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c)$$
$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t$$
$$h_t = o_t \odot \tanh(c_t)$$

The final hidden state h_8, produced after processing all eight time steps, is passed through a linear output layer to generate a preliminary price estimate referred to throughout this paper as lstm_pred. This value represents the network's learned expectation of the next price given the observed historical sequence.

3.2.2 Ensemble Fusion for Final Prediction

While the sequential model captures temporal price dynamics effectively, it operates on a fixed feature structure and cannot easily incorporate the kind of categorical context — shop identity, flower variety, size grade — that meaningfully influences Anthurium price formation in local markets. Different shops serve different buyer demographics and have different pricing power. Different varieties command structurally different price levels. A model that ignores this context will exhibit systematic bias across product combinations.

To address this, the lstm_pred value is assembled into a richer feature vector that includes contextual market attributes alongside statistical features derived from the same price window:

$$x_{RF} = [shop,\ variety,\ size,\ lstm\_pred,\ units\_sold,\ lag\_1,\ lag\_2,\ lag\_4,\ EMA\_4]$$

This vector is passed through a pre-trained preprocessing pipeline that applies ordinal encoding to categorical fields and scaling to numerical features, then fed into a Random Forest Regression model. The Random Forest generates the final price prediction by averaging across an ensemble of B independently grown decision trees:

$$\hat{p}_{final} = (1/B) \cdot \Sigma_{b=1}^{B}\ T_b(x_{RF})$$

The ensemble averaging suppresses variance in the final prediction — a property that matters considerably in markets where individual transaction outliers can distort a single tree's output. The final predicted price is rounded to the nearest 0.5 LKR increment, reflecting the half-rupee pricing conventions observed in Sri Lankan Anthurium retail markets. This rounding step is implemented as:

$$\hat{p}_{rounded} = round(\hat{p}_{final} \times 2) / 2$$

The overall prediction pipeline can be summarized as a sequential operation in which the temporal model extracts a price representation from the historical sequence, and the ensemble model refines that representation using contextual market information. Neither stage is redundant: removing the LSTM stage leaves the Random Forest without access to temporal pattern information, while removing the Random Forest stage leaves the prediction without contextual grounding.

3.3 SAMIS: A Market Sustainability Scoring Framework

The predicted price $\hat{p}$ is a point estimate. On its own, it tells a grower what the market is likely to do but says nothing about whether that expectation should change their behavior. A high predicted price in a chronically unstable market carries different implications than the same predicted price in a historically reliable one. A rising price against a backdrop of unusual volatility warrants more caution than the same price increase in a market with consistent trading patterns. These distinctions are the motivation for SAMIS.

SAMIS — Sustainability, Analysis, and Market Intelligence Score — is a composite index defined on the interval [0, 100]. It decomposes market assessment into three components, each capturing a distinct and orthogonal dimension of market quality: S, measuring how statistically consistent the predicted price is with historical behavior; T, measuring the direction and magnitude of recent price momentum; and R, measuring the underlying volatility of the market as reflected in its historical price dispersion. The composite score is:

$$SAMIS = S(\hat{p}, \mu, \sigma) + T(\hat{p}, p_{recent}) + R(\sigma, \mu)$$

Each component is computed independently, so that a weakness on one dimension — for example, a falling price trend — does not artificially inflate or deflate the other components. This independence is an intentional design property: it allows the recommendation engine to respond to the full pattern of component values, not just the composite total.

3.3.1 Statistical Foundations

Before defining the components, it is useful to establish the statistical measures on which they depend. Given a historical price series $\{p_1, p_2, ..., p_n\}$ for a particular shop–variety–size combination, three summary statistics are computed.

The mean price μ represents the central tendency of the historical market:

$$\mu = (1/n) \cdot \Sigma_{i=1}^{n} p_i$$

The standard deviation σ measures the average magnitude of price deviations from the mean:

$$\sigma = \sqrt{(1/n) \cdot \Sigma_{i=1}^{n} (p_i - \mu)^2}$$

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

The Coefficient of Variation CV expresses standard deviation as a proportion of the mean, producing a scale-independent measure of relative price dispersion:

$$CV = \sigma / \mu$$

CV is preferred over raw standard deviation for cross-product comparisons because it normalizes volatility by the price level. A market trading at a mean of 200 LKR with σ = 16 LKR has identical relative volatility to one trading at 100 LKR with σ = 8 LKR, and CV correctly assigns them the same risk score. Using σ directly would make the higher-priced variety appear inherently riskier, which would be misleading.

3.3.2 Sustainability Component (S)

The Sustainability Component measures the degree to which the predicted price is consistent with the historical distribution of prices for the same product combination. A predicted price that falls well within the normal range of historical variation is statistically plausible — it represents a continuation of established market behavior. A predicted price that lies far outside that range, whether unusually high or unusually low, suggests an atypical market condition that may not persist.

This assessment is formalized using z-score interval logic, which is grounded in the Empirical Rule for normally distributed data: approximately 68% of historical prices lie within one standard deviation of the mean, roughly 95% within two standard deviations. Predicted prices falling within successively wider intervals therefore represent successively less typical market conditions. The scoring structure is as follows:

Condition	Score	Classification	Interpretation
$\mu - \sigma \leq \hat{p} \leq \mu + \sigma$	40	Highly Sustainable	Predicted price within the ce tral 68% of historical distribution
$\mu - 1.5\sigma \leq \hat{p} \leq \mu + 1.5\sigma$	30	Sustainable	Moderate deviation from mean; still within expected range
$\mu - 2\sigma \leq \hat{p} \leq \mu + 2\sigma$	20	Moderately Sustainable	Price approaching statistical boundary; unusual but not extreme
$\hat{p} \notin [\mu - 2\sigma, \mu + 2\sigma]$	10	Not Sustainable	Price is a statistical outlier; likely reflects transient market anomaly

Table 3. Sustainability Component scoring based on z-score interval analysis.

An important property of this formulation is that it responds symmetrically to outliers in both directions. A predicted price that is unusually high — perhaps reflecting a temporary demand surge — receives the same low sustainability score as one that is unusually low. From the perspective of market reliability, both extremes represent unstable conditions that may not support a dependable selling decision. The grower is informed not only about the expected price, but about how representative that price is of the market's normal operating range.

The score ceiling of 40 for this component, relative to 30 for each of the other two, reflects the interpretive emphasis placed on statistical price normality in the overall framework. A market whose predicted price falls squarely within historical norms is fundamentally more reliable as a selling environment than one characterized by high momentum or low volatility alone.

3.3.3 Trend Component (T)

The Trend Component captures the direction and strength of recent price movement. Rather than evaluating the predicted price against the full historical distribution — as the Sustainability Component does — this component focuses on the short-term trajectory: where has the price been going over the past week, and does the predicted price continue that direction?

Let p_recent denote the first price in the trailing seven-day window. The percentage change between this reference price and the predicted price is:

$$\Delta\% = (\hat{p} - p\_recent) / p\_recent \times 100$$

Condition	Trend Label	Score	Practical Meaning
$\Delta\% > 3\%$	Rising	30	Prices are appreciating; selling now or soon captures the upward movement
$-3\% \leq \Delta\% \leq 3\%$	Stable	15	Price is in equilibrium; no strong directional signal either way
$\Delta\% < -3\%$	Falling	5	Market is contracting; deferring where possible may recover value

Table 4. Trend Component scoring based on percentage price change over the trailing seven-day window.

The ±3% band around the stable zone is an important design choice. In day-to-day flower markets, price variation of up to

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

3% is common and does not reliably signal a directional trend — it is more accurately characterized as market noise. Classifying every sub-3% movement as a meaningful trend would produce unstable, frequently changing recommendations that erode grower trust and lead to decision paralysis. The band suppresses this noise and ensures that only genuine, sustained price movements trigger Rising or Falling classifications.

The score asymmetry — 30 for Rising versus 5 for Falling — is intentional and reflects the agricultural market context. For a perishable commodity, a rising price is an active signal that rewards immediate action, while a falling price, though unfavorable, still leaves the grower with limited options if the product cannot be stored. The score differential encodes this asymmetry without eliminating the falling market signal entirely.

3.3.4 Risk Component (R)

While the Sustainability Component evaluates where the predicted price sits within the historical distribution, and the Trend Component evaluates short-term direction, the Risk Component evaluates the reliability of the market as a whole. Specifically, it asks: regardless of where today's price falls, how predictable and stable has this market historically been?

A market with a long history of tightly clustered prices around the mean is one in which predictions can be trusted and selling decisions can be made with reasonable confidence. A market with large swings in both directions is one in which even an accurate predicted price may be followed immediately by a sharp reversal. The grower's effective risk is higher in the latter case, even if today's predicted price happens to be favorable.

The Coefficient of Variation captures this distinction precisely. By expressing price dispersion as a proportion of the mean, it provides a normalized measure of relative volatility that is comparable across product combinations with different price scales:

CV Range	Risk Level	Score	Market Interpretation
$CV < 0.05$	Low	30	Highly stable market; price fluctuations are minor and predictions are reliable
$0.05 \leq CV \leq 0.10$	Medium	20	Moderate volatility; reasonable confidence in forecast with some caution warranted
$CV > 0.10$	High	10	Volatile market history; predicted price should be treated ith reduced confidence

Table 5. Risk Component scoring based on the Coefficient of Variation.

The CV thresholds of 0.05 and 0.10 represent meaningful boundaries in the context of retail flower markets. A CV below 0.05 implies that prices rarely deviate by more than 5% of the mean — a market condition consistent with stable consumer demand and predictable supply. A CV above 0.10 implies that price swings of more than 10% of the mean are common, indicating a market where external shocks — seasonal demand surges, supply disruptions, competitor pricing — are frequent enough to make any single-day price forecast inherently uncertain.

3.3.5 Composite SAMIS Score and Score Structure

The three components are summed directly to produce the SAMIS score:

$$SAMIS = S + T + R, \quad 0 \leq SAMIS \leq 100$$

Component	Min Score	Max Score	Weight of Total	Primary Measure
S — Sustainability	10	40	40%	z-score interval logic
T — Trend	5	30	30%	7-day % price change
R — Risk	10	30	30%	Coefficient f Variation
SAMIS Total	**25**	**100**	**100%**	**Composite index**

Table 6. SAMIS component structure, score ranges, and relative contribution to the composite index.

The minimum achievable score of 25 arises from the floor values of each component (S=10, T=5, R=10) and is assigned to markets where the predicted price is a statistical outlier, the recent trend is falling, and historical volatility is high simultaneously. This floor is intentionally non-zero, reflecting the fact that even in severely unfavorable conditions, some residual market information remains available. The maximum of 100 represents a market where the predicted price is fully within the historical normal range, a rising trend is confirmed, and historical volatility is low — conditions that together constitute the most reliable possible selling environment.

3.4 Four-Class Market Recommendation

The SAMIS score is not simply thresholded at a single cutoff. A four-class recommendation framework evaluates the composite score alongside the individual component classifications, ensuring that specific combinations of conditions receive appropriate handling rather than being averaged away.

Class 1 — Optimal Selling Conditions

The highest recommendation tier requires all three components to reach their maximum scores. Sustainability must be 40 (predicted price within one standard deviation of historical mean), Trend must be 30 (price rising by more than 3%), and Risk must be 30 (CV below 0.05). This triple conjunction ensures that the Optimal classification is reserved for markets that are simultaneously normal, appreciating, and

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

stable — not merely one or two of those things. The typical SAMIS range for Class 1 conditions is 85 to 100.

Class 2 — Favorable Selling Conditions

Favorable conditions are identified when at least two components are at or near their maximum values and the third is not at its minimum. This captures markets where the overall picture is positive but one dimension — typically either a flat trend or moderate volatility — introduces some uncertainty. Growers in Class 2 markets are advised to proceed with selling, with awareness that the market is not fully optimal. SAMIS scores in the range of approximately 65 to 84 characterize this class.

Class 3 — Moderate Conditions

The Moderate classification covers a broad middle range in which the market is functional but no strong directional case exists for or against immediate selling. This class also captures a specific edge case: when the Sustainability score is exactly 20 — meaning the predicted price lies between the 1.5σ and 2σ boundaries — the market receives a Class 3 classification regardless of trend and risk scores, because the marginal statistical plausibility of the predicted price introduces a structural level of uncertainty. SAMIS scores from 40 to 64 typically characterize this class.

Class 4 — Unfavorable Conditions

The Unfavorable classification is triggered when the composite SAMIS score falls below 40, or when the Sustainability score reaches its minimum of 10 (predicted price beyond 2σ), or when both Trend and Risk are simultaneously at their minimum scores. The composite score threshold and the individual component triggers operate as alternative conditions, meaning any one of them is sufficient to generate a Class 4 recommendation. Where post-harvest storage is feasible, the system advises deferring the selling decision until market conditions improve.

Class	SAMIS Range	S Score	T Score	R Score	Recommendation
1 — Optimal	85 – 100	40	30	30	Sell immediately
2 — Favorable	65 – 84	≥ 30	≥ 15	≥ 20	Recommended to sell
3 — Moderate	40 – 64	20	Any	Any	Monitor; partial sale possible
4 — Unfavorable	< 40	10	5	10	Defer if storage allows

Table 7. Four-class market recommendation framework derived from SAMIS component scores.

4. RESULTS AND DISCUSSION

4.1 Price Prediction Performance

The prediction pipeline was evaluated on a held-out test set of Anthurium transaction records not used during training. Comparison was made against an ARIMA baseline, a standalone ensemble regressor trained on the same feature set without the sequential pre-processing layer, and the sequential model operating without ensemble fusion. The fusion architecture produced the most stable predictions across the evaluation period, with notably lower error variance during weeks characterized by demand volatility. The reduction in prediction variance — rather than just mean error — is of particular practical significance in this context, since growers are harmed not just by average forecast inaccuracy but by unpredictable forecast quality.

Model Configuration	MAE (LKR)	RMSE (LKR)	Error Std. Dev.
ARIMA Baseline	18.42	24.17	15.83
Ensemble Regressor Only	12.65	16.94	11.27
Sequential Model Only	10.38	14.22	9.61
Sequential–Ensemble Fusion (Proposed)	**6.73**	**9.48**	**6.12**

Table 8. Predictive accuracy comparison across model configurations.

4.2 SAMIS Assessment Evaluation

The performance of the SAMIS classification scheme was assessed through analyzing whether SAMIS-classified periods matched up with the market behavior that happened after those times based on previously known information regarding the results. Times classified under Class 1 or Class 2 were compared to the weeks where the subsequent market values did not fall below the forecasted price, indicating that these periods represented a good time for selling. The correlation between SAMIS classification and the previously observed market conditions proved to be significantly greater than what would be expected based solely on the price forecast.

A particularly instructive set of observations involved market periods where the predicted price was high but the Sustainability Component received a score of 10, indicating that the predicted price lay beyond the two standard deviation boundary of historical distribution. In each such case, subsequent market prices returned toward the historical mean within several days, confirming that these episodes represented transient price spikes. The SAMIS framework correctly classified these as unfavorable or moderate

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

conditions despite the superficially attractive predicted prices - a distinction that a price-only recommendation system would have missed.

The Risk Component also demonstrated discriminative value at the product-combination level. Variety–shop combinations with high historical CV received consistent Class 3 or Class 4 classifications even during periods of rising trend, reflecting genuine market unpredictability that was subsequently confirmed in price trajectory data. Low-CV combinations, by contrast, received stable Class 1 and Class 2 classifications during normal market periods, providing growers with reliable positive signals when the underlying market conditions genuinely warranted them.

4.3 Recommendation Engine Behavior

The manner in which the four-class model reacts in extreme situations is something worth looking into, especially because it is during the more ambiguous market situations when the added value of using the recommendation algorithm becomes evident over merely using a score cutoff. For instance, when there is a significant increase in the price level (T = 30), while at the same time there is a statistical outlier for the forecasted price level (S = 10) with high volatility (R = 10), the resultant SAMIS score of 50, if applied to a single-threshold system, will categorize the market situation as moderate. However, the component-based decision-making process will override the result by classifying it under the Class 4 category based solely on its Not Sustainable classification, thus providing valuable advice not to invest in such a market.

5. CONCLUSION

In this study, a market intelligence solution for predicting daily prices of Anthurium and assessing their sustainability was introduced, which was developed based on the realistic information-related challenges experienced by small floriculture growers in Sri Lanka. The prediction model utilizes a temporal sequence model in combination with an ensemble regression layer within a sequential fusion framework to identify both the time-dependent features and contextual aspects of the market, generating predictions of daily prices that are more robust against market fluctuations compared to each model used alone.

The truly valuable component of this study is the introduction of SAMIS, a method for transforming the simple predicted price value into a detailed market analysis along three separate statistical axes. By analyzing the predicted price relative to past values for the exact combination of goods, relative to the current trend in the market, and relative to the volatility of those past prices in the market, SAMIS offers a much richer context for decision making when compared to a simple single-axis system. The classification algorithm applied to the outputs of SAMIS accounts for special cases where increasing prices exist in volatile markets and normal prices exist in steady markets in a way consistent with agricultural economics.

It is important to recognize the limitations of this study. The parameters that were used to define the thresholds of SAMIS - the standard deviations to determine sustainability, the ±3% window to determine the nature of the trends, and the CV levels as the indicators of risk – were chosen based on considerations of both statistics and the market dynamics, rather than being empirically determined using a validation sample. A more accurate calibration of these parameters using empirical methods can be helpful to achieve better sensitivity of the system in case of abnormal distributions of prices in markets under investigation.

Future research will concentrate on the calibration of thresholds, integration of external signals from the market environment, and applying this methodology to new types of ornamental flowers and market settings. This paper has presented architecture for two different kinds of models, which are applicable not only to the flower business but also any other perishable agricultural commodity markets; the authors believe that this paper will be able to provide the basis for future works.

REFERENCES

[1] ochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

[2] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.

[3] Box, G. E. P., & Jenkins, G. M. (1976). Time Series Analysis: Forecasting and Control. Holden-Day, San Francisco.

[4] Vapnik, V. N. (1995). The Nature of Statistical Learning Theory. Springer, New York.

[5] iaw, A., & Wiener, M. (2002). Classification and regression by randomForest. R News, 2(3), 18–22.

[6] ers, F. A., Schmidhuber, J., & Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. Neural Computation, 12(10), 2451–2471.

[7] C en, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794.

[8] uguma, H. T., & Abdi, D. (2019). Analysis of cut flower export market concentration and trend in Ethiopia. Journal of Agricultural Economics, Extension and Rural Development, 7(1), 781–788.

[9] endoza, M. G., & Rosegrant, M. W. (1995). Pricing conduct of spatially differentiated markets. International Agricultural Trade Research Consortium Working Paper.

[10] Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189–1232.

[11] Shumway, R. H., & Stoffer, D. S. (2017). Time Series Analysis and Its Applications (4th ed.). Springer.

[12] Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688.

[13] Abadi, M., et al. (2016). TensorFlow: A system for large-scale machine learning. 12th USENIX Symposium on Operating Systems Design and Implementation, 265–283.

[14] McKinney, W. (2010). Data structures for statistical computing in Python. Proceedings of the 9th Python in Science Conference, 51–56.

[15] Kouadio, L., et al. (2018). Artificial intelligence approach for the prediction of Robusta coffee yield using soil fertility properties. Computers and Electronics in Agriculture, 155, 324–330.