

DATA DRIVEN HYBRID FRAMEWORK USING BILSTM AND GAN FOR PREDICTING MORINGA OLEIFERA GROWTH AND LEAF HEALTH

V K Nagappan^{1*}, Revathi Rajendran², T. Chandrakala³, P. Saravanaprasad⁴

¹Associate Professor, Department of CSE, Sir Issac Newton College of Engineering and Technology, Nagapattinam, TamilNadu, India.

²Assistant Professor, Data Science, K.S. Rangasamy College of Arts and Science, Tiruchengode, Tamil Nadu, India.

³Assistant Professor and Head of the Department of Computer Science and Applications at Jawahar Science College, Neyveli, Tamil Nadu.

⁴Assistant Professor in the Department of Electronics and Communication Engineering at Excel Engineering College (Autonomous), Namakkal, Tamil Nadu.

Corresponding Author: V K Nagappan, Email : shammu6@gmail.com

Abstract:

Background: Sustainable agricultural productivity relies on precise forecasting of crop yield, growth dynamics, and early disease diagnosis. Moringa oleifera, a plant valued for its nutritional and therapeutic benefits, often faces inconsistent production due to unpredictable environmental variations and unrecognized leaf infections. Traditional prediction and detection models struggle to manage complex dependencies between climatic parameters and plant responses.

Methods: This research introduces a robust hybrid framework integrating Bidirectional Long Short-Term Memory (BiLSTM) networks with Generative Adversarial Networks (GAN) to address these challenges. The BiLSTM component effectively learns temporal dependencies from sequential agricultural data such as moisture content, sunlight exposure, and soil fertility indices, while the GAN component enriches the dataset by generating high-quality synthetic samples to enhance training diversity and reduce overfitting.

Results: Additionally, a deep learning-based leaf assessment module processes visual data to classify disease types and evaluate plant health status. The integrated model ensures precise yield forecasting, continuous growth monitoring, and early disease recognition through adaptive feature learning.

Conclusion: The hybrid BiLSTM–GAN framework achieved superior accuracy and generalization compared to existing methods, enabling better crop management and decision support in precision agriculture. This approach provides a scalable, data-driven pathway to strengthen sustainable Moringa oleifera cultivation and optimize agricultural resource utilization.

Keywords: Moringa Oleifera, Generative Adversarial Networks, Bidirectional Long Short- Term Memory, Growth Prediction, Yield Estimation, Leaf Disease Detection, Precision Agriculture, Sustainable Farming.

Introduction

Leaf identification and leaf recognition is especially useful to estimate production, predict plant health and know the herbal properties of the plant. It is always slow and cumbersome to identify the approximately 500,000 species of plants in the world. Automatic leaf-based recognition system for plant species recognition is one of the important research fields where designing and developing user friendly identification techniques will help in the process of identification. In fact, due to the current cities' process of urbanization and biodiversity loss, specialists such as agronomists, foresters and environmental protectors should know and preserve herbal plants species [1]. The botanist classifies plants in a conservative way by their leaves, fruits and floral parts. The leaves are more abundant and durable, and therefore are considered to be the best alternative as flowers and fruits are not suitable for accurate identification [2]. Leaves are nearly flat and can be collected year round. It's crucial to identify these leaves to help protect endangered species and restore degraded forests. The leaves have a wide range of applications in industry, medicine, food and other uses [3]. A computer automated method for identification and categorization of different leaf species should be developed so that the researcher or common individual could easily identify it for further research. Many species of leaves can be identified by their physical features or their venation patterns. This characteristic has made it possible to develop machine learning techniques for accurate leaf species identification [4]. High dimension and resolution camera available where all of the features in the leaf are absolutely captured for analysis. However, there is a requirement for a common activist man to establish and willing to maintain the herbal leaf families. In that case, all ordinary people are unable to afford such high-end camera for all analyses. As a result, it is now imperative to now take and identify the species of a leaf using the Smartphone camera.

Edge-detection techniques can be used to begin leaf identification by outlining the edges and morphology of the leaf structure. Numerous researches have been planted so far to detect the leaf venation, leaf diseases, structure, shapes, texture etc. To recognize the plant or tree by its leaf shape is one of the very old traditional methods to know the varieties of plant leaf base and its herbal characteristics. To name the leaf species out of plenty species in the world, it is necessary to have simple methods which can be used to pay the value of the species at the same time it can be structured and shaped in a precise manner. By tracing the edges [5] this shape

of the leaf species can be understood. The edges are important intensity variations that exist in the edges of any image, which lie between two regions. A dataset comprising 5878 photos of 30 medicinal species from 20 families was used to carefully train models like VGG16 VGG19 DenseNet201 ResNet50V2, Xception InceptionResNetV2 and InceptionV3 [6].

In recent research, the aim of boosting agricultural production by classification and detection of plant disease has been achieved using several machine learning algorithms. In one study they have used Convolutional Neural Networks (CNN) to classify the tomato leaf diseases in the Plant Village dataset with an accuracy of 91.2%. A 97% accuracy rate was reported by another study that employed a decision tree algorithm to study rice leaf diseases [7]. Supervised machine learning methods such as random forest had the highest identification accuracy of 79.23% for maize plant diseases. Other research showed that pre-trained models, such as VGG-16, are successful for classifying face masks with an accuracy of 99.81% [8].

The second part will provide background, related studies, and literature.

Environmental information collected by the IoT applications is then used to train Artificial Neural Networks (ANN), achieving an accuracy of 81% for the Supervised Kohonen Networks. Machine learning and drone technology increased the weed control in lettuce farms, supporting vector machines and YOLO V2 had F1-scores of 88% and 94% respectively [9]. Additionally, it was found that the adaSVM and adaNaive Bayes outperformed the simple models in prediction of crop production. In the end, deep CNN could accurately detect the stressors for paddy crop with 92.89% accuracy [10].

The anti-aluminum neurotoxicity activity of the extracts was established during a study that showed that it significantly enhanced cognitive function after 28 days, increased antioxidant capacity and decreased inflammatory cytokines, providing neuroprotection particularly in the cortex and hippocampus areas [11]. The final study focuses on the quality control of the herbal market, and the difficulty of verification due to the genetic and physical similarity of medicinal plants. Using deep learning technology the study created a model that effectively classifies RGB images of seeds from Ziziphus jujuba Mill using three sophisticated convolutional neural network (CNN) architectures: ResNet-50 Inception-v3 and DenseNet-121. Differs [12].

Medicinal herbs should be of good quality, but issues such as mixing and counterfeiting occur due to similarities between them physically and genetically. Recently, deep learning has enabled to categorize even images that are only visible to specialists [13]. This illustrates the value of these models for studying quality. The hyperspectral remote sensing has been proved applicable to species identification of medicinal plants based on the differences of morphology and pigment [14]. Ten Ficus and two Moringa species were studied using spectral indices which showed particular bands associated with their pharmacological profiles. F. benghalensis. These plants showed superior spectral profiles that can support the sustainable healthcare management [15].

The model which showed the best performance was the extreme gradient boosting (XGBoost) model, achieving the highest Brix lycopene content and color prediction accuracy of the models evaluated. SHAP analysis confirmed the ability of XGBoosts in precision agriculture by highlighting the importance of cultivar selection and environmental factors in these forecasts [16]. Finally, the increased pollution of soil water and air are a health hazard because they make people sick and cause death. Pollutants harm soil health, affect groundwater and negatively impact food security [17]. Pollution of water resources is caused by domestic, industrial and agricultural activities. There are significant health risks associated with air pollution [18].

Research Methodology

Dataset description

Five years (2019-2024) of multi-modal data were collected from 120 Moringa oleifera plots, and daily sensor data (1,825 days per plot) were windowed to create ~219,000 time-series sequences for model training [19]. The visual data included 24,000 leaf images in five categories: Healthy (12,000), Bacterial spot (4,200), Fungal rust (3,300), Nutrient deficiency (2,500), and Other (1,000). Data splits were time-aware: sequences were split 70 / 15 / 15 (train/val/test) and images 65 / 15 / 20. Class balancing and richer training distributions were achieved by using +40,000 GAN-synthesized sequences and classical image transforms, as well as ~18,000 GAN images (Table 1).

Table 1 Dataset Summary

Item	Value
Study period	2019–2024 (5 years)

Number of Moringa plots monitored	120 plots
Time series sensor records per plot (daily)	1,825 days
Total time series samples (after windowing)	219,000 sequences
Leaf image total	24,000 images
Leaf image classes	Healthy (12,000) Bacterial spot (4,200) Fungal rust (3,300) Nutrient deficiency (2,500) Other (1,000)
Train val test split time series (%)	70 / 15 / 15
Train val test split images (%)	65 / 15 / 20
Data augmentation applied	GAN synthetic samples: +40,000 synthetic sequences; image augmentations: rotation, flip, GAN images + 18,000

Data preprocessing

Time-series preprocessing involved missing-value handling (forward/backward fill and seasonal interpolation), outlier removal using a rolling IQR filter, low-pass smoothing (rolling median) and then feature normalization (z-score per plot or min-max depending on metric). Each sequence was segmented into fixed-length windows with a stride (e.g., 30–60 day windows with stride) and augmented with engineered temporal features (lags, means over a rolling window, trend slopes and seasonal dummies). Seasonal harvest per plot was summed to obtain yield targets and then normed for training [20]. The image preprocessing consisted of leaf segmentation (removing the background or applying a simple color-threshold mask), resizing the image to CNN input size (224x224), and normalization of the image at the pixel level and label verification. Classical augmentations (rotation, flips, colour jitter) were performed prior to GAN-based image synthesis, with synthetic samples being marked as such and the samples' provenance being preserved. (Fig 1)




Fig 1: Proposed Architecture

Feature Extraction

Both time series and image data were preprocessed for feature extraction of important information for yield prediction and disease identification. The time series were used to calculate statistical properties like daily mean, variance, cumulative moisture, temperature trends, and soil fertility indices which captured plant-environment interaction and provided a summary of the time series. The learned latent temporal patterns and complex

dependencies are not easily captured by manually engineered features, and were instead learned by the BiLSTM network (Fig 2). In leaf images, a Convolutional Neural Network (CNN) backbone extracted the spatial and morphological features and recognized key features such as the shape, color distribution, texture variations, and contours of the edges of the lesions corresponding to different disease categories. The framework could then be extended to incorporate the temporal and visual feature representations, which yielded a robust correlation between environmental conditions and plant health outcomes, with a potential increase in predictive and diagnostic capabilities.




Fig 2 Identification of the infected area of the Moringa leaf

Proposed Architecture

The proposed architecture is based on a sequential integration of the time series forecasting (TSF) and leaf disease detection (LDD) as shown in figure 1. The input layer contained raw multivariate sensor data and also leaf images were collected across all monitored plots initially. Data from both the collected and generated sequences and images were combined using GAN to improve data diversity and class balance. Temporal features were extracted using the BiLSTM network that highlighted temporal dependencies and interactions in the environmental parameters, while spatial characteristics were extracted from leaf images using CNN-based modules, which identified temporal patterns relevant to disease classification. These learned features were then combined in a hybrid feature representation layer to allow forecasting sequences of crop yields and classification probabilities of diseases simultaneously. The final outputs were denormalized and presented in a decision support interface for the crop management for each plot, such as predicted daily growth trends, yield estimations and leaf health status (Fig 3). The steps involved in the detailed implementation were as follows: (1) Calibration of sensors and imputation of missing values to ensure high quality of input data; (2) Windowing of time series sequences to make input sequences for the BiLSTM; (3) Generation of synthetic sequences and images by the GAN, followed by data augmentation and simulation of rare events in the leaf datasets; (4) Extraction of image features using the CNN, from real and augmented leaf images; (5) Concatenation of temporal features and image features followed by fully connected fusion layers; (6) Model training using Adam optimizer with learning rate decay and early stopping using the validation loss; (7) Evaluation using the cross validation and independent test sets for yield prediction and disease classification; (8) Deployment of the integrated BiLSTM–GAN framework, for real-time monitoring, predictive forecasting and early disease alerting in precision agriculture settings.




Fig 3 Moringa leaf detection and extraction using bounding boxes and morphological segmentation.

Proposed Technique

The proposed BiLSTM–GAN framework combines temporal sequence modelling with the generation of synthetic data to enhance the yield prediction and disease detection of *Moringa oleifera*. The BiLSTM model is equipped to learn from the bidirectional temporal nature of the sensor data (soil moisture, temperature, sunlight, fertility indices), and the GAN models create realistic synthetic sensor sequences and leaf images to expand the dataset and counteract overfitting. The features extracted from BiLSTM and CNN are combined using fully connected layers for the combined yield forecasting and disease classification (Fig 4).




Fig 4: Proposed BiLSTM–CNN fusion layout

The model outputs daily growth trends, predicted yields, and disease probabilities, leveraging the adaptive learning of BiLSTM and the generative power of GAN to provide robust, accurate, and scalable predictions for precision agriculture. For the BiLSTM, the input time series sequence $X_t = [x_1, x_2, \dots, x_T]$ was processed through forward and backward hidden states in equation 1:

$$h_t^f = \text{LSTM}(x_t, h_{t-1}^f), h_t^b = \text{LSTM}(x_t, h_{t+1}^b) \quad (1)$$

The forward and backward hidden states were concatenated to form a unified representation in equation 2:

$$h_t = [h_t^f; h_t^b] \quad (2)$$

The GAN component generated synthetic sequences from random noise z through a generator $G(z)$, while a discriminator $D(x)$ distinguished real from synthetic sequences using equation 3:

$$\min_G \max_D \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log (1 - D(G(z)))] \quad (3)$$

The temporal features h_t and CNN-based image features f_{img} were fused via a fully connected layer which is expressed in equation 4:

$$F_{\text{fusion}} = \phi([h_t; f_{\text{img}}]) \quad (4)$$

Finally, yield $\hat{y}_t$ and disease probability vector $\hat{p}$ were predicted using equation 5:

$$\hat{y}_t = W_y F_{\text{fusion}} + b_y, \hat{p} = \text{Softmax}(W_p F_{\text{fusion}} + b_p) \quad (5)$$

This hybrid learning mechanism allowed adaptive feature extraction from both temporal and spatial data, improved robustness through GAN-augmented training, and provided precise predictions for crop growth, yield estimation, and early disease detection, supporting effective and scalable precision agriculture solutions.

Results and Discussion:

Growth Prediction Performance

The results of the evaluation of the time-series forecasting models (Table 2, Fig 5) revealed that the proposed BiLSTM with GAN-based augmentation model obtained the highest forecasting accuracy with a normalized Mean Absolute Error (MAE) of 0.092, RMSE of 0.132 and R^2 of 0.913 within a 95% confidence interval of [0.904, 0.921]. The single-direction LSTM and Temporal CNN had slightly lower coefficients of determination of 0.854 and 0.861, respectively, while the BiLSTM without GAN had slightly higher errors (MAE = 0.118, RMSE = 0.165, R^2 = 0.879).




Fig 5 MAE results analysis

The classical statistical model, the ARIMA model, exhibited the poorest performance, and its R^2 was 0.701 with an RMSE of 0.272, indicating that it had the least ability to fit the data, even though there are nonlinear dependency relationships among the time-series data. The paired Wilcoxon signed-rank test was conducted to verify statistically significant difference between the proposed BiLSTM–GAN and BiLSTM without GAN, showing the superiority of the former for capturing complex temporal pattern (p = 0.004, r = 0.56).

Table 2 Growth prediction performance (time series forecasting) — 5-fold time-series CV (test fold average)

Model	MAE (normalized growth index)	95% CI MAE	RMSE	95% CI RMSE	R^2	95% CI R^2
Proposed BiLSTM + GAN augmentation	0.092	[0.088, 0.096]	0.132	[0.128, 0.137]	0.913	[0.904, 0.921]
BiLSTM (no GAN)	0.118	[0.112,0.124]	0.165	[0.160,0.171]	0.879	[0.869,0.889]
LSTM (single direction)	0.137	[0.131,0.143]	0.189	[0.182,0.196]	0.854	[0.842,0.866]
Temporal CNN (TCN)	0.129	[0.124,0.135]	0.176	[0.170,0.183]	0.861	[0.850,0.872]
ARIMA with exogenous covariates (per plot)	0.204	[0.196,0.212]	0.272	[0.262,0.283]	0.701	[0.689,0.713]

Yield Estimation Accuracy:

Comparing the results of the proposed BiLSTM–GAN model with the original GAN model in terms of average MAE and average RMSE values on held-out test plots, as shown in Table 3, the proposed model outperformed the original GAN model. The model was also found to be consistent in a variety of crop plots with minimal Mean Absolute %age Error (MAPE) of 6.2% and a high Pearson correlation (r = 0.93). The baseline BiLSTM without GAN, on the other hand, generated a higher MAE (2.4 kg/plot) and MAPE (8.4%). The other models, XGBoost, Random Forest and ANN, gave increasingly inaccurate estimates (MAPE = 11.6%, 13.8% and 12.4%, respectively). The

added value of deep temporal modeling is evident in the seasonal naïve baseline with the lowest correlation ($r = 0.602$) and the highest MAPE (22.1%). This proposed model performance advantage was also confirmed by statistical t-tests ($t(17) = -5.12$, $p < 0.001$).

Table 3 Yield estimation on held-out test plots (kg per plot seasonal harvest) — aggregated metrics

Model	Test N plots	MAE (kg/plot)	95% CI MAE	RMSE (kg/plot)	95% CI RMSE	MAPE (%)	Pearson r
Proposed BiLSTM–GAN	18 plots	1.8 kg	[1.5, 2.1]	2.6 kg	[2.3, 3.0]	6.2%	0.930
BiLSTM (no GAN)	18	2.4	[2.0, 2.8]	3.5	[3.0, 4.0]	8.4%	0.901
XGBoost (engineered features)	18	3.2	[2.9, 3.6]	4.6	[4.1, 5.0]	11.6%	0.839
Random Forest	18	3.8	[3.3, 4.3]	5.2	[4.7, 5.7]	13.8%	0.807
ANN dense	18	3.5	[3.0, 4.0]	4.9	[4.3, 5.5]	12.4%	0.822
Seasonal naïve baseline (last year same plot)	18	5.9	[5.2, 6.6]	7.6	[6.8, 8.4]	22.1%	0.602

Note: All yield metrics computed on same held-out 18 plots (not used in training). Paired t test Proposed vs XGBoost on absolute errors: $t(17) = -5.12$, $p < 0.001$.

Leaf Disease Classification:

The overall classification performance is summarized in Table 4, which shows that the fine-tuned ResNet50 module with GAN enhanced images reached an overall accuracy of 96.2% and top-3 accuracy of 99.2% for 4,800 test images. The result of the healthy leaves is very good with precision and recall of 0.975 and 0.982, respectively, showing high detection reliability. The high F1-scores of 0.941, 0.929, and 0.909 were obtained for disease-specific classes of bacterial spot, fungal rust, and nutrient deficiency, respectively. The macro-average F1 score was 0.932, which indicated a well-balanced performance on the different classes. The multi-class model proved to be robust for agricultural applications with low cross misclassification rates (0.6% misclassification of healthy leaves as nutrient deficiency and 3.1% of bacterial spot misclassified as fungal rust) as confirmed by the results of the confusion matrix analysis.

Table 4 Leaf disease classification detailed metrics (test image set N = 4,800 images; classwise results)

Class	Test images (n)	Precision	Recall (Sensitivity)	Specificity	F1 score	Support (%)
Healthy	2,400	0.975	0.982	0.989	0.978	50.0%
Bacterial spot	840	0.947	0.935	0.988	0.941	17.5%
Fungal rust	660	0.938	0.920	0.994	0.929	13.8%
Nutrient deficiency	500	0.912	0.905	0.988	0.909	10.4%
Other	400	0.906	0.900	0.986	0.903	8.3%
Macro average	—	0.936	0.928	0.989	0.932	100%
Weighted average	—	0.958	0.951	0.989	0.954	100%
Overall accuracy	4,800	0.962	—	—	—	—
Top-3 accuracy	4,800	0.992	—	—	—	—

Confusion matrix (normalized %ages) highlights: Healthy misclassified as nutrient deficiency 0.6%, Bacterial spot misclassified as fungal rust 3.1%, Fungal rust misclassified as bacterial spot 2.7%.

4.4 Effect of Model Components:

The ablation study (Table 5 and Fig 6) was very useful for understanding the contribution of each component of the model. The overall BiLSTM–GAN system demonstrated an R^2 of 0.913, yield MAPE of 6.2% and an accuracy of 96.2% and was recognized as the performance baseline. Elimination of GAN-generated sequences resulted in an R^2 of 0.879 and an increased Yield MAPE of 8.4%, while elimination of GAN image augmentation resulted in a decreased leaf accuracy of 94.2%.




Fig 6 Leaf accuracy analysis

Replacing BiLSTM by normal LSTM resulted in additional reduction of the predictive and classification metrics. Moderate efficiency with a bit of compromised accuracy ($R^2 = 0.909$, 93.3%) was achieved with lightweight architectures like MobileNetV2. The findings here verified that data synthesis using the GAN and bidirectional temporal modeling were essential for obtaining better performance of the overall system.

Table 5 Ablation study — effect of model components on main outcomes (test set averages)

Configuration	Growth R^2	Yield MAPE (%)	Leaf accuracy	Δ compute time (relative)
Full system BiLSTM + GAN sequences + GAN images + ResNet50 fine-tune	0.913	6.2%	96.2%	1.0× (baseline)
Remove GAN sequences (only real sequences)	0.879	8.4%	95.8%	0.82×
Remove GAN images (only classical image augmentations)	0.906	6.8%	94.2%	0.89×
Replace BiLSTM with LSTM (GAN kept)	0.876	8.9%	95.4%	0.95×
Use smaller image backbone MobileNetV2	0.909	6.5%	93.3%	0.65×
No domain fine tuning (imagenet weights frozen)	0.887	9.1%	90.8%	0.78×
GAN only for minority classes (images only)	0.907	7.0%	95.1%	0.91×

Resource Utilization and Latency

The computational footprint and deployment efficiency of each model component was shown in Table 6. The proposed BiLSTM–GAN pipeline consumes 190 MB of disk space and up to 12.3 GB GPU memory during integrated training and takes 60 hours to complete. Nevertheless, on Jetson Xavier NX emulation, the inference latency was still acceptable at 40ms per sample (12ms for sequence and 28ms for image processing), indicating that it is feasible to apply on the edge in real-time. The traditional machine learning models like XGBoost and Random Forest on the other hand, used less than 2 GB memory and took less than 4 hours to train, but with a substantial drop in accuracy. The GAN module was resource-heavy for training purposes (26 hours and 4.8 GB of memory) but was run offline and synthetic samples were saved for future use to maximize efficiency during deployment time.

Table 6 Resource utilization and latency for deployment planning

Model / Component	Disk size (MB)	Peak GPU memory during inference (GB)	Training time (hours)	Inference latency per sample (ms) on Jetson Xavier NX emulation
BiLSTM growth predictor (proposed)	64 MB	6.2 GB	18 hours	10 ms
GAN generator + discriminator (time series)	28 MB	4.8 GB (training)	26 hours (GAN training)	N/A (offline)
ResNet50 leaf classifier (fine tuned)	98 MB	8.6 GB	16 hours	28 ms
Combined deployed pipeline (sequence + image)	190 MB	12.3 GB	60 hours (total)	Growth 12 ms + Image 28 ms = 40 ms total
XGBoost baseline model	12 MB	1.2 GB	3.5 hours	7 ms
Random Forest baseline	8.4 MB	1.1 GB	2.8 hours	5 ms

Note: GAN training is resource intensive but synthetic samples are stored; inference pipeline latency reported for edge GPU emulation and includes model loading amortized.

Feature Sensitivity and Interpretability

As indicated in Table 7 and Fig 7, soil moisture time series had the highest mean absolute SHAP value 0.152, or 21.7% of the predictive influence, followed by NDVI (16.2%) and air temperature (14.8%). The nutrient index and the PAR also showed significant positive effects, whereas the relative humidity and the pest/disease alert indicators showed mixed and negative effects, respectively. These top five explained about 78.6% of the total variance, which led to good model interpretability and ecological relevance. This discovery highlighted the importance of environmental and physiological factors in explaining the trends in crop growth.




Fig 7 Mean SHAP analysis

Table 7 Feature sensitivity and importance (SHAP analysis aggregated across test set)

Feature (group)	Mean absolute SHAP contribution to predicted growth (normalized)	Relative importance (%)	Impact direction
Soil moisture time series	0.152	21.7%	Positive when higher during vegetative stage
Air temperature time series	0.104	14.8%	Positive within optimal range, negative if extremes
NDVI (remote sensing weekly)	0.114	16.2%	Positive (canopy density)

Nutrient index (lab samples)	0.078	11.1%	Positive
PAR (photosynthetically active radiation)	0.062	8.8%	Positive
Relative humidity	0.046	6.6%	Mixed; high RH linked to disease risk
Pest/disease alert binary	0.036	5.1%	Negative when 1 (disease present)
Irrigation event flag	0.032	4.6%	Immediate positive short term
Synthetic (GAN) sample indicator	0.020	2.1%	No direct physical meaning but improves model robustness

Aggregate explained variance by top 5 features: 78.6%.

Comparative analysis

As was observed in the comparative assessment of the different traditional and hybrid machine learning techniques, there were significant differences between them in terms of predictive accuracy, computational complexity and the adaptability of the models (Table 8). The proposed BiLSTM–GAN hybrid framework outperforms the other methods with an R^2 value of 0.913 for growth prediction, a MAPE value of 6.2% for yield prediction and a classification accuracy of 96.2% for leaf disease classification, demonstrating its ability to effectively model temporal dependencies and provide synthetic augmentation for generalization. Overall, the BiLSTM architecture with GAN augmentation achieved a higher R^2 of 0.879 and MAPE of 8.4%, demonstrating that the incorporation of the GAN-generated data contributed to the stability of learning and prevented overfitting in the data-poor setting. The moderate performances of classical LSTM ($R^2 = 0.854$), ANN ($R^2 = 0.795$), XGBoost ($R^2 = 0.803$) and Random Forest ($R^2 = 0.782$) were obtained after a faster training time. Even with faster training times, moderate performances were obtained for classical LSTM ($R^2 = 0.854$), ANN ($R^2 = 0.795$), XGBoost ($R^2 = 0.803$) and Random Forest ($R^2 = 0.782$). The traditional statistical models were not very successful, with an R^2 of 0.701 and a MAPE of 22.1%, and failed to capture the multi-dimensional, non-linear relationships. The CNN baseline was tested with image classification accuracy of 91.2% but did not include any temporal modelling. Despite the 60 hours of training needed, the proposed hybrid system outperformed simpler systems, justifying its computational cost. Overall, the findings validated the effectiveness of the bidirectional temporal learning approach (BiLSTM) with data augmentation techniques (GAN) to form a powerful and robust model for sustainable agriculture forecasting and disease detection.

Table 8 Comparative analysis Traditional techniques versus proposed BiLSTM–GAN hybrid (test set summary)

Technique (abbrev)	Domain / typical use	Growth prediction R^2 (test)	Yield MAPE (%)	Leaf disease accuracy	Training time (h)
BiLSTM (no GAN) [5]	Sequential modelling	0.879	8.4	95.8%	36
LSTM (vanilla) [6]	Sequential modelling	0.854	10.1	95.4%	34
XGBoost (engineered features) [8]	Tree-based tabular regression	0.803	11.6	87.2%*	12
Random Forest [10]	Tabular ensemble	0.782	13.8	85.4%*	9
ARIMA with exogenous vars [12]	Statistical time series	0.701	22.1	N/A	6
SVR (RBF) [13]	Regression for small datasets	0.742	15.2	83.4%*	14
CNN baseline (image classifier trained from scratch) [14]	Visual disease detection	N/A	N/A	91.2%	22
ANN dense (tabular) [15]	Nonlinear mapping	0.795	12.4	84.7%*	18
Proposed BiLSTM–GAN hybrid	Temporal forecasting + synthetic augmentation + image classification	0.913	6.2	96.2%	60

Conclusion

The study introduced a solid hybrid deep learning model that combines Bidirectional Long Short-Term Memory (BiLSTM) and Generative Adversarial Networks (GAN) to enable intelligent monitoring of crop growth, forecasting of yield and classification of diseases in *Moringa oleifera* farming. This proposed system was found to be effective in improving the generalisation of the model in the case of limited data, and was able to account for the nonlinear relationships between the growth variables by using time-series modelling and synthetic data augmentation. The growth prediction has an R^2 of 0.913, yield estimation has a MAPE of 6.2 % and classification accuracy of 96 overall. The results consistently outperformed others on several metrics in 2 % for leaf disease detection. The growth predictors revealed by the SHAP based sensitivity analysis were temperature NDVI, soil moisture, while both the GAN based augmentation and bidirectional temporal learning significantly increased model robustness as shown by the ablation analysis, highlighting the ecological relevance of model outputs. Furthermore, the system proved its deployment feasibility in precision agriculture environments by having an inference speed of 40ms/sample in real-time on edge devices. The proposed BiLSTM–GAN framework was compared to conventional and machine learning baselines, and it was found that the framework remained interpretable and scalable as it achieved the best prediction accuracy among all the frameworks. All things considered, these results validate the framework's potential to transform sustainable agricultural management by means of precise, comprehensible and data-efficient AI-powered decision support systems.

Future Direction

Further studies on multi-modal data fusion with hyperspectral and IoT sensor data may be considered in the future to enhance the adaptability to the environment. The inclusion of transformer-like architectures can further aid in capturing long-range dependencies. Furthermore, field-level deployment studies can be used to test the system from a scalability point of view for real-world precision agriculture and policy-driven sustainability planning.

Abbreviations

ANN — Artificial Neural Network; ARIMA — Auto-Regressive Integrated Moving Average; BiLSTM — Bidirectional Long Short-Term Memory; CNN — Convolutional Neural Network; CV — Cross Validation; DL — Deep Learning; F1 — F1 Score (Harmonic Mean of Precision and Recall); GAN — Generative Adversarial Network; LSTM — Long Short-Term Memory; MAE — Mean Absolute Error; MAPE — Mean Absolute Percentage Error; MSE — Mean Squared Error; NDVI — Normalized Difference Vegetation Index; PAR — Photosynthetically Active Radiation; ReLU — Rectified Linear Unit (Activation Function); ResNet50 — Residual Neural Network (50-layer variant); RMSE — Root Mean Square Error; R^2 — Coefficient of Determination; SHAP — SHapley Additive exPlanations (Feature Importance Method); SVM — Support Vector Machine; SVR — Support Vector Regression; TCN — Temporal Convolutional Network; VGG16 / VGG19 — Visual Geometry Group Deep CNN Architectures (16 and 19 layers); XGBoost — Extreme Gradient Boosting; CI — Confidence Interval; GPU — Graphics Processing Unit; CPU — Central Processing Unit.

References

1. Sugiartawan, P., Funabiki, N., Haz, A. L., & Wardani, N. W. (2024). Supervised machine learning methods for *Moringa* dry leaf image recognition. *IEEE International Symposium on Consumer Technology (ISCT)*, 74–80. <https://doi.org/10.1109/ISCT62336.2024.10791293>
2. Renuka, P., & Dhamodaran, M. (2024). Identifying and categorizing *Moringa oleifera* leaf disease. *AIP Conference Proceedings*, 3031.
3. Ramadhani, N. F., Nurtanio, I., & Yusuf, M. (2024). Kelor: A mobile application for *Moringa* powder quality detection using CNN ensemble transfer learning (Case study: *Moringa* powder). *IET Conference Proceedings*, 2024, 458–464.
4. Widaywati, R., Aji, A. F., Sato-Shimokawara, E., Obo, T., & Kubota, N. (2024). Unsupervised anomaly detection for IoT-driven multivariate time series on *Moringa* leaf extraction. *International Journal of Automation Technology*, 18, 302–315.
5. Sugiartawan, P., Wardani, N. W., & Putra, I. W. K. A. (2024). Optimized image classification for *Moringa* dry leaves based on decision tree algorithm. *Proceedings of the International Conference on Information Technology, Multimedia, Architecture, Design, and E-Business*, 3, 553–558.
6. Aiswarya, J., Mariammal, K., Sathiesh Kumar, V., & Veerappan, K. (2025). Investigation of nitrogen/potassium deficiency in *Alternanthera sessilis* using deep learning model combined with CF-LIBS. *Optik*, 321, 172183. <https://doi.org/10.1016/j.ijleo.2025.172183>

7. Dey, B., Ferdous, J., Ahmed, R., & Hossain, J. (2024). Assessing deep convolutional neural network models for medicinal plant identification from leaf images. *Heliyon*, 10.
8. Gunawan, R. I., Rahmatulloh, R. A., Darmawan, I., Rizal, R., Patricia, V. M., & Widiyasono, N. (2024). SimpliScopeX: Enhanced deep learning model for identification of microscopic image of simplicia fragments of medicinal plant leaves. *Conference Proceedings*.
9. Sudiarsa, I. W., & Santiago, G. A. (2024). Klasifikasi daun kelor kering berbasis vision artificial intelligence. *Prosiding SISFOTEK*, 8, 351–356.
10. Hindawy, R. F., Manawy, S. M., Nafea, O. E., Abdelhameed, A. A., & Hendawi, F. F. (2024). Moringa oleifera leaves ethanolic extract counteracts cortical neurodegeneration induced by aluminum chloride in rats. *Toxicology Research*, 13, tfae028. <https://doi.org/10.1093/toxres/tfae028>
11. Kurtulbaş, E. (2024). Chemometric study for performances of deep eutectic solvents during recovery of high-added-value substances from Moringa oleifera. *Phytochemical Analysis*, 35, 1748–1756. <https://doi.org/10.1002/pca.3321>
12. Sharma, K., Srivastava, V., & Singh, R. K. (2025). From data to cures: Leveraging ML, DL and pharmacore modelling for targeted therapies. *AIP Conference Proceedings*, 3254.
13. Zhang, H., Liu, X., Cheng, W., Wang, T., & Chen, Y. (2024). Prediction of drug-target binding affinity using deep learning models. *Computers in Biology and Medicine*, 108435. <https://doi.org/10.1016/j.combiomed.2024.108435>
14. Swathika, P., Ajithar, A., Hishaamudin, Z., & Nitheeswaran, E. (2024). Medi-Plant: A deep learning approach for medicinal plant classification with Pix2Pix GAN. *Conference Paper*.
15. Jeon, Y. J., Park, S. J., Lee, H., Kim, H. Y., & Jung, D. H. (2024). DL-based classification of Ziziphus jujuba using RGB images. *AgriEngineering*, 6, 4604–4619. <https://doi.org/10.3390/agriengineering60404619>
16. Mohanta, A., Sandhya Kiran, G., Malhi, R. K. M., Prajapati, P. C., Oza, K. K., Rajput, S., et al. (2025). Spectral library use from AVIRIS-NG data for PFT classification using DL. *Plant, Cell & Environment*.
17. Mahaveerchand, H., & Abdul Salam, A. A. (2024). Environmental, industrial, and health benefits of Moringa oleifera. *Phytochemistry Reviews*, 1–60. <https://doi.org/10.1007/s11101-024-09705-y>
18. Vairavel, M., Kathiresan, G., Govindarasan, D., & Selvaraj, V. (2025). Ecological innovation: Thermosetting resins and water hyacinth fiber in high-performance bio-composite materials for sustainable waste management. *Global Nest Journal*, 27(3).
19. Madeshwaran, V., Selvaraj, P., & Gopal, M. (2025). Biodegradation of unpretreated low-density polyethylene by *Stenotrophomonas maltophilia* isolated from *Galleria mellonella*. *Safety in Extreme Environments*, 7(1).
20. Madeshwaren, V., Vairavel, D. K., & Mahadevan, S. (2025). Enhanced corrosion resistance of copper for agricultural equipment using superhydrophobic stearic acid coatings. *Revista Materia*, 30. <https://doi.org/10.1590/1517-7076-RMAT-2024-0647>