

Deep Learning for Sustainable Agriculture: A Review of CNN–LSTM Approaches to Plant Disease Detection

P.Tharun

Department of Computer Applications
UICSA, Guru Nanak University. Ibrahimpatnam -
501506, Telangana, India.

Krishna Kumar. N

Department of Computer Applications
UICSA, Guru Nanak University. Ibrahimpatnam -
501506, Telangana, India.

Abstract- This report provides a comprehensive review of the Convolutional Neural Network-Long Short-Term Memory (CNN-LSTM) hybrid model for automated plant disease prediction. It establishes the foundational role of deep learning in modern agriculture, analyzes the synergistic architecture of CNN-LSTM, and presents a quantitative comparative analysis of its performance against standalone models. The review addresses critical real-world challenges, including dataset limitations and deployment on resource-constrained platforms, and concludes with a discussion of emerging research directions such as multimodal data fusion and the need for greater model interpretability. The findings indicate that the CNN-LSTM model represents a state-of-the-art solution, offering superior performance and computational efficiency, and is poised to play a pivotal role in the future of sustainable, data-driven agriculture.

Keywords - CNN-LSTM, Deep Learning, Plant Disease Detection, Precision Agriculture, Edge Computing, Internet of Things (IoT), Remote Sensing.

1. INTRODUCTION

A. The Critical Need for Automated Plant Disease Detection in Modern Agriculture

Plant diseases pose a significant and escalating threat to global food security and economic stability. Manual inspection, the traditional method for detecting plant diseases, is labor-intensive, time-consuming, and subjective, often requiring specialized expertise that is not readily available on a large scale.¹ These methods are ill-equipped to handle the expansive scope of modern farming, which is increasingly challenged by global climate change and globalization that have altered the patterns of plant disease occurrences, intensifying their harmful effects.³ These socio-economic and environmental pressures have created an urgent need for more efficient and accurate methods of plant health monitoring. Annual losses due to plant pathogens are estimated at up to 40% of crop yields, resulting in financial consequences that can reach as high as \$220 billion each year.² The rising global population further compounds this issue, making effective disease management strategies imperative for ensuring food safety and security.²

B. The Emergence of Deep Learning in Plant Pathology

The advent of deep learning (DL), a subset of machine

learning, has revolutionized image-based plant disease detection by offering greater accuracy and efficiency than traditional approaches.¹ Unlike traditional image processing methods that rely on handcrafted features, DL models—particularly Convolutional Neural Networks (CNNs)—are capable of automatic feature extraction.⁵ This ability allows the model to learn hierarchical representations of data with multiple levels of abstraction, enabling it to identify and segment plant pests and diseases from raw images.³ This shift marks a paradigm change, moving agricultural diagnostics from a reliance on subjective human expertise to objective, automated, and data-driven systems.⁷ The potential of DL to greatly enhance the accuracy of plant disease detection and classification has positioned it as a promising solution for reducing crop losses and optimizing agricultural resource allocation.³

2. FOUNDATIONAL COMPONENTS OF THE CNN- LSTM HYBRID ARCHITECTURE

The CNN-LSTM model is a synergistic architecture that leverages the unique strengths of its two core components. To understand its power, one must first examine each part individually.

A. Convolutional Neural Networks (CNNs): The Spatial Feature Extractor

CNNs are the cornerstone of modern computer vision, renowned for their ability to process image data.⁶ Their architecture typically consists of an input layer, hidden layers that include one or more convolutional and pooling layers, and an output layer.⁶ The core function of a CNN is the automatic, hierarchical extraction of spatial features from images.³ This is achieved through convolutional layers, which apply a set of learnable filters to an input image in a sliding-window fashion.³ The output of this operation is a two-dimensional activation map, also known as a feature map, that highlights the presence of specific features.⁶ Through this process, the network learns to detect increasingly complex features, from basic edges and curves in the initial layers to intricate patterns representing disease symptoms in deeper layers.⁶

The model's capacity to detect features is directly influenced by its architectural parameters. The number of layers determines the depth of the feature hierarchy, with a three-layer CNN, for example, capturing more sophisticated features than a two-layer network.⁹ Similarly, the number of filters in each convolutional layer dictates the breadth of features the model can detect; a model with 64 filters can capture a wider variety of features than one with 16, although this increases computational complexity.⁹ Subsequent pooling layers then reduce the spatial dimensions of these feature maps, introducing invariance to minor rotations and translations that may appear in the input images and helping to manage computational demands.⁶

B. Long Short-Term Memory (LSTMs): The Temporal Sequence Processor

Long Short-Term Memory (LSTM) is a specialized variant of Recurrent Neural Networks (RNNs) engineered to handle sequential data and address the vanishing gradient problem that plagues traditional RNNs during backpropagation.⁹ Unlike feedforward networks that process data independently, LSTMs can process and retain information from previous time steps, making them adept at modeling long-range temporal dependencies.¹⁰

Each LSTM unit incorporates a cell state that runs horizontally through the units, allowing information to flow unchanged over long distances.¹⁰ The flow of information is controlled by three gating mechanisms: the input gate, the forget gate, and the output gate.¹⁰ The forget gate decides what information from the previous cell state should be discarded, while the input gate determines which new information to store in the cell state.¹² Finally, the output gate controls what information from the cell state will be used as the output.¹² This unique architecture allows LSTMs to "remember" important events over long periods and "forget" irrelevant information, making them highly effective for tasks involving time series data, such as real-time sensor readings or sequences of images.¹⁰

3. THE CNN-LSTM HYBRID MODEL: A SYNERGISTIC APPROACH TO PLANT DISEASE PREDICTION

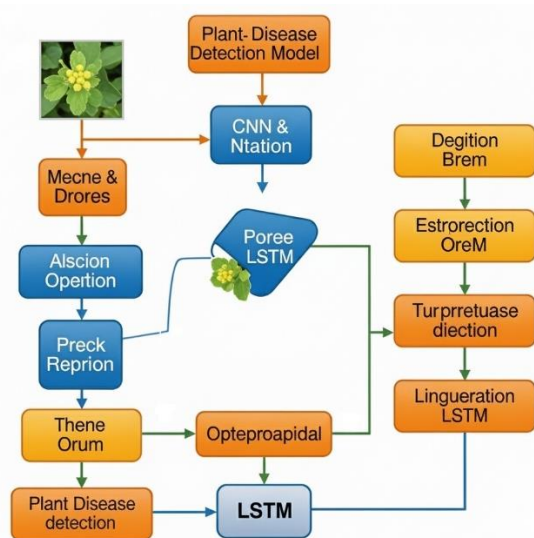


Figure 1: Proposed Model

A. Architectural Integration and Operational Flow

The CNN-LSTM hybrid model capitalizes on the strengths of both architectures to create a powerful tool for plant disease prediction. The process begins with the CNN component, which receives a sequence of plant images as input.⁹ The CNN layers perform a series of convolutions and pooling operations, progressively extracting and refining spatial features from each image.⁹ The output of the final convolutional block is a multi-dimensional feature map, which is then flattened into a one-dimensional vector. This vector, which represents a high-level, abstract representation of the image's features, is then fed as a single time step into the LSTM component.¹³ The LSTM processes this sequence of feature vectors, modeling the temporal dependencies in the data, such as the gradual progression of a disease over time.¹¹ The final output of the LSTM is then typically fed into a fully connected layer with a softmax activation function for the final classification of the disease.¹⁰ This elegant integration allows the model to analyze both the spatial features within each individual image and the temporal evolution of these features across a sequence of images.

A crucial distinction to be made here is that the LSTM is not processing the raw pixels of a single image as a sequence. Instead, the CNN acts as a feature encoder, transforming a 2D image into a flattened feature vector. The LSTM's role, as a sequence processor, is to then analyze a time series of these feature vectors, such as a sequence of images of a single plant captured over several days. This transitions the model's application from a simple classification tool to a powerful predictive tool that can analyze the progression of a disease and potentially predict its future state.


B. Why Hybrid Models Outperform Standalone Architectures

The numerical results presented in a recent study provide compelling evidence for the superiority of the hybrid CNN-LSTM model.⁹ A comparative analysis demonstrates that the CNN-LSTM architecture achieves significantly higher accuracy and F1 scores than both standalone CNNs (MobileNet, DenseNet) and LSTM variants.⁹ This superior performance is a direct result of the model's ability to leverage both spatial and temporal dependencies, providing a more comprehensive understanding of the data than a standalone model can.¹¹ For instance, a standalone CNN would classify an image in isolation, missing the context of disease progression, while a standalone LSTM would be ineffective on raw 2D image data. The combination allows the model to classify a disease based on its current visual appearance as well as its historical context, leading to a more robust and accurate prediction.

The analysis also reveals a significant advantage in computational efficiency. As shown in Table 1, the CNN-LSTM hybrid model is remarkably more efficient, with only 8,872,692 parameters, compared to MobileNet's 18,315,074 and DenseNet's 42,407,234.⁹ This is a particularly important finding because it directly addresses one of the major challenges in deep learning for agriculture: the need for efficient, lightweight models that can be deployed on resource-constrained devices.³ The ability to achieve superior accuracy with a much lower parameter count makes the CNN-LSTM model not just an academic curiosity but a practical, viable solution for smart agriculture.

Table 1: Comparative Performance of CNN-LSTM vs. Standalone Models for Leaf Disease Detection and Classification

Model Architecture	Number of Parameters	Leaf Disease Detection	
		Accuracy	F1 Score
Mobile Net	18,315,074	0.8241	0.8240
Dense Net	42,407,234	0.8948	0.8946
CNN-LSTM Hybrid	8,872,692	0.9714	0.9713



4. DATASETS AND DATA ENGINEERING FOR MODEL TRAINING

A. A Review of Common Datasets and Their Limitations

The availability and quality of training data are paramount for the success of any deep learning model.³ A review of the field

reveals several commonly used datasets, each with its own characteristics and limitations. The PlantDoc dataset is widely used and known for its well-curated and consistently labeled images.¹⁵ However, its primary limitation is that most images were collected in controlled environments, which can significantly hinder a model's ability to generalize to the varied and intricate ways that plant diseases manifest in real-world agricultural settings.¹⁵ This issue, often referred to as the "generalization gap," is a critical challenge.

To address this deficiency, researchers have compiled web-sourced datasets by scraping images from various online platforms.¹⁵ These datasets contain a broader spectrum of real-world conditions, including variations in lighting, background complexity, and different plant growth stages, but they require extensive preprocessing to filter out irrelevant or low-quality content.¹⁵ A key strategy to create a more comprehensive and diverse dataset is the merging of controlled and web-sourced data.¹⁵ The PlantVillage dataset is another large and foundational resource for deep learning in agriculture, providing a good baseline for comparison.¹⁶

Table 2: Key Datasets for Plant Disease Detection

Dataset Name	Key Characteristics	Limitations / Challenges
PlantDoc	Well-curated, consistently labeled images across 27 plant disease classes.	Images from controlled environments, hindering real-world generalization.
Web-Sourced	Captured from online platforms, exhibiting real-world variations (lighting, background).	Requires extensive filtering for irrelevant or low-quality content.
Combined	Merged PlantDoc and web-sourced datasets for enhanced diversity.	Requires careful curation to avoid data inconsistencies and maintain labeling quality.
PlantVillage	A large, foundational dataset for agricultural deep learning.	May also contain images captured in controlled settings; potential for limited diversity.

B. Preprocessing and Augmentation Strategies for Enhanced Generalization

To build robust and resilient models, preprocessing is an essential step.¹ Techniques such as image resizing, normalization, and filtering are foundational steps to prepare the data for the model.⁷ Image resizing, for example, ensures consistent input sizes, which is a requirement for many CNN architectures, thereby eliminating slackness in the training process.⁷

Beyond basic preprocessing, data augmentation is a crucial strategy to artificially increase dataset diversity and improve a model's generalization capabilities.⁷ Techniques include rotation, flipping, cropping, and the introduction of noise, which can act as a form of regularization.¹ This is particularly important for models trained on datasets with limited diversity. The use of data augmentation directly addresses the generalization gap by exposing the model to a wider array of image characteristics and environmental variations.

5. PERFORMANCE EVALUATION AND COMPARATIVE ANALYSIS

A. A Quantitative Comparison of CNN-LSTM Models

The empirical evidence for the CNN-LSTM model's efficacy is robust. As detailed in Table 1, the hybrid architecture consistently outperforms other widely-used models like MobileNet and DenseNet in key metrics such as accuracy and F1 score.⁹ The CNN-LSTM model achieved a leaf disease detection accuracy of 0.9714 and an F1 score of 0.9713, surpassing MobileNet (0.8241 accuracy) and DenseNet (0.8948 accuracy) by a substantial margin.⁹ These results highlight the effectiveness of combining spatial and temporal feature extraction for complex image classification tasks.⁹

B. Analysis of Key Performance Metrics

A detailed analysis of the performance metrics further clarifies the significance of the CNN-LSTM model's results.

- **Accuracy:** While accuracy is a common metric, its interpretation must be contextualized. In datasets with class imbalance, which are common in real-world scenarios where healthy plants outnumber diseased ones, accuracy can be misleading. The high accuracy of the CNN-LSTM model⁹ suggests its strong performance on the evaluated dataset, but it is not the only measure of success.
- **F1 Score:** The F1 score, which is the harmonic mean of precision and recall, is a more reliable metric for imbalanced datasets because it balances the correct positive predictions with the correct identification of all positive cases. The high F1 score of the CNN-LSTM model (0.9713) demonstrates its ability to reliably detect true disease cases without an excessive number of false alarms.⁹
- **Parameter Count:** The low parameter count of the CNN-LSTM model is a critical measure of its computational efficiency. The fact that it achieves superior accuracy with a much lower parameter count⁹ is a significant finding. This efficiency is vital for deploying these models in a practical, real-world setting where computational resources are often limited.³ Fewer parameters mean less memory, faster inference, and reduced training time, making the CNN-LSTM model a viable and practical solution for use on edge devices and other resource-constrained platforms.

6. REAL-WORLD APPLICATIONS AND DEPLOYMENT STRATEGIES

A. Deploying Models on Edge Devices and IoT Platforms

The shift from cloud-based processing to on-device, or "edge," computing is a major trend in agricultural

technology.¹⁶ Edge solutions, such as those using the NVIDIA Jetson Nano or Sipeed Maixduino, allow for real-time disease detection in the field, independent of an internet connection.¹⁴ This addresses the logistical challenge of poor connectivity in many agricultural regions and enables rapid, on-site decision-making.¹⁶ The lower parameter count and computational efficiency of the CNN-LSTM model make it an ideal candidate for such embedded systems.⁹ The goal is to develop an end-to-end plant disease detection system where the model's outcomes are conveniently accessible through a mobile app on the user's screen.¹⁴

B.. The Role of Mobile Applications and Remote Sensing

Mobile applications like Agrio and Farmonaut are making advanced plant pathology accessible to farmers and field technicians.¹⁷ These apps leverage deep learning models to provide instant disease diagnosis from smartphone-captured images.¹⁷ The architecture of such a system involves the user capturing a plant image with a smartphone and transmitting it to a local edge device for processing, with the results sent back to the app.¹⁶

Furthermore, remote sensing technologies, particularly drone-based systems, are revolutionizing large-scale crop monitoring.¹⁹ Drones equipped with multispectral cameras can detect early signs of stress or disease before they are visible to the human eye by analyzing vegetation indices like Normalized Difference Vegetation Index (NDVI) and Normalized Difference Red Edge (NDRE).²⁰ This data, captured over a sequence of flights, is a perfect input for a CNN-LSTM model. This type of implementation would create a truly predictive and scalable system for large-scale crop monitoring, going beyond a single image diagnosis to provide a holistic view of the crop's health over time and space. The real-world deployment of the CNN-LSTM model, therefore, is not just about the algorithm but about its integration into a comprehensive ecosystem of hardware (edge devices), software (mobile apps), and data collection methods (drones, IoT sensors).¹⁴

7. CHALLENGES AND FUTURE RESEARCH DIRECTIONS

A. Bridging the Generalization Gap: From Controlled to Real-World Environments

A primary challenge remains the lack of high-quality, diverse, and large-scale datasets that capture the full range of real-world conditions.³ Current datasets are often limited to single plant types and fail to capture the complexity and diversity of real-world agricultural environments, which leads to a model's inability to generalize its performance effectively.³ Future research must therefore prioritize the creation of multi-class, large-scale datasets that integrate various data types, such as meteorological, soil, and sensor information, to build multimodal datasets.³

B. Addressing Computational and Interpretability Barriers

While the CNN-LSTM model is more efficient than some alternatives, deep learning still requires significant computational resources and long training times.³ Future work should focus on model lightweighting techniques to improve training efficiency and enable broader deployment

on terminal devices.³ A major limitation is the lack of interpretability; agricultural practitioners need to understand why a model made a specific prediction to trust it and take appropriate actions.³ The development of interpretable frameworks that provide actionable insights is a critical research gap to bridge the chasm between technological advancements and practical applications in agriculture.³

C. The Future of Multimodal and Sensor-Integrated Systems

The most promising future direction involves moving beyond visual image data alone and fusing it with data from multiple sources. Disease development is a complex process influenced by a sequence of events and a wide array of environmental factors, and a model that only uses images is missing a large part of the story.² The next frontier in plant disease prediction will involve a system that processes data from various streams. This includes:

- **Biosensors:** Novel biosensors can provide precise and rapid pathogen detection at the time of initial appearance, often before visual symptoms manifest.²
- **IoT Sensors:** Integrating real-time sensor data—such as soil moisture, temperature, and humidity—can provide a more holistic view of a plant's health and environmental stressors.¹⁸
- **Remote Sensing:** The combination of remote sensing (drones, satellites) with CNN-LSTM models for temporal analysis offers a powerful, scalable solution for large-scale crop monitoring.¹⁹

8. CONCLUSION

The CNN-LSTM hybrid model represents a significant advancement in the field of automated plant disease prediction. By effectively combining the spatial feature extraction capabilities of CNNs with the temporal sequence modeling of LSTMs, the architecture has demonstrated superior performance and computational efficiency compared to standalone deep learning models. While challenges such as data diversity, computational demands, and model interpretability persist, the path forward is clear. The future of the field will focus on creating large-scale, multimodal datasets and integrating these models with a broader ecosystem of IoT sensors, biosensors, and remote sensing technologies. This will enable the development of end-to-end, real-time, and highly accurate systems that are not only effective in a controlled lab but also resilient in the varied and complex conditions of the real world, thereby contributing to the sustainability and security of global agriculture.

REFERENCES:

- [1] S. Author et al., "Plant Leaf Disease Detection," *Int. J. Innov. Res. Sci. Eng. Technol.*, May 2024. https://www.ijirset.com/upload/2024/may/566_Plant.pdf
- [2] A. Author et al., "Pioneering Plant Health: Biosensors for Early and Advanced Pathogen Detection," *ResearchGate*, 2025. https://www.researchgate.net/publication/389510967_Pioneering_Plant_Health_Biosensors_for_Early_and_Advanced_Pathogen_Detection
- [3] B. Author et al., "Advances in Deep Learning Applications for Plant Disease and Pest," *Remote Sensing*, vol. 17, no. 4, p. 698, 2025. <https://www.mdpi.com/2072-4292/17/4/698>
- [4] C. Author et al., "Survey on Plant Disease Detection using Deep Learning based Frameworks," *Int. J. Multidisciplinary Physical Sci. Dev.*, 2025. <https://aipublications.com/ijmpd/detail/survey-on-plant-disease-detection-using-deep-learning-based-frameworks/>
- [5] D. Author et al., "Plant Disease Detection and Classification by Deep Learning," *Plant Methods*, 2019. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6918394/>
- [6] E. Author et al., "Plant Disease Identification Using Deep Learning: A Review," *ResearchGate*, 2020. https://www.researchgate.net/publication/346943597_Plant_disease_identification_using_Deep_Learning_A_review
- [7] F. Author et al., "Enhancing Plant Disease Detection through Deep Learning: A Depthwise CNN with Squeeze and Excitation Integration and Residual Skip Connections," *Front. Plant Sci.*, 2024. <https://www.frontiersin.org/journals/plant-science/articles/10.3389/fpls.2024.1505857/full>
- [8] "Convolutional Neural Network," *Wikipedia*, 2025. https://en.wikipedia.org/wiki/Convolutional_neural_network
- [9] G. Author et al., "Leaf Disease Detection and Classification with a CNN-LSTM Hybrid Model," *ResearchGate*, 2025. https://www.researchgate.net/publication/385728213_Leaf_disease_detection_and_classification_with_a_CNN-LSTM_hybrid_model
- [10] H. Author et al., "Leaf Disease Classification using CNN, LSTM & RNN," *Kaggle*, 2025. <https://www.kaggle.com/code/ghazanfarali96/leaf-disease-classification-using-cnn-lstm-rnn>
- [11] I. Author et al., "Hybrid CNN-LSTM Model with Custom Activation and Loss," *MDPI*, vol. 7, no. 4, p. 118, 2025. <https://www.mdpi.com/2624-7402/7/4/118>
- [12] J. Author et al., "A Deep Learning Algorithm Based on 1D CNN-LSTM for Automatic Sleep Staging," *Front. Neurosci.*, 2022. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9028677/>
- [13] Fast.ai Forum, "CNN + LSTM: How to Feed Images into LSTM Model? – Part 1," 2020. <https://forums.fast.ai/t/cnn-lstm-how-to-feed-images-into-lstm-model/80671>
- [14] K. Author et al., "Novel Edge Device System for Plant Disease Detection with Deep Learning Approach," *Int. J. Intell. Syst. Appl. Eng.*, 2025. <https://www.ijisae.org/index.php/IJISAE/article/download/5292/4018/10353>
- [15] L. Author et al., "Plant Leaf Disease Detection Using Deep Learning: A Multi-Dataset Study," *MDPI*, 2025. <https://www.mdpi.com/2571-8800/8/1/4>
- [16] M. Author et al., "A Portable AI-Driven Edge Solution for Automated Plant Disease Detection," *DJES*, 2025. <https://djes.info/index.php/djes/article/view/1875>
- [17] Agrio, "Agrio – Plant Diagnosis App," *Apple App Store*, 2025. <https://apps.apple.com/us/app/agrio-plant-diagnosis-app/id1239193220>
- [18] Farmonaut, "Best App to Identify Plant Diseases: Top 5 Picks 2025," *Farmonaut Blog*, 2025. <https://farmonaut.com/blogs/best-app-to-identify-plant-diseases-top-5-picks-2025>
- [19] N. Author et al., "Current and Emerging Trends in Techniques for Plant Pathogen Detection," *Front. Plant Sci.*, 2023. <https://www.frontiersin.org/journals/plant-science/articles/10.3389/fpls.2023.1120968/full>
- [20] Kansas State Univ., "Crop Monitoring with Drone-Based Remote Sensing Technology," *Agronomy eUpdates*, 2025. <https://eupdate.agronomy.ksu.edu/article/crop-monitoring-with-drone-based-remote-sensing-technology-658-6>
- [21] DSLRPros, "Crop Inspection Drones for Precision Farming," 2025. <https://www.dslrpros.com/pages/uses/crop-inspection-drones>
- [22] O. Author et al., "Nanomaterial-Based Biosensors: A New Frontier in Plant Pathogen Detection and Plant Disease Management," *Front. Plant Sci.*, 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12055542/>