

Explainable Deep Learning-Based Potato Leaf Disease Detection and Severity Assessment for Smart Agriculture in Bangladesh

Dr Kazi Abdul Mannan

drkaziabdulmannan@gmail.com

Center for Academic & Professional Career Development and Research (CAPCDR)

<https://orcid.org/0000-0002-7123-132X>

Nafiz Hasan Sinha

Shanto-Mariam University of Creative Technology <https://orcid.org/0009-0007-8906-4586>

Shafin Ahmed

Shanto-Mariam University of Creative Technology <https://orcid.org/0009-0000-4162-6913>

Research Article

Keywords: Potato Leaf Disease, Deep Learning, EfficientNetB0, Explainable AI, Grad-CAM, Smart Agriculture

Posted Date: September 1st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10879535/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

In Bangladesh, the potato (*Solanum tuberosum* L.) stands as an indispensable food and cash crop, deeply intertwined with national food security, rural livelihoods, and the broader agricultural economy. However, foliar diseases such as early blight and late blight frequently precipitate substantial yield losses and quality degradation when not identified and mitigated during the nascent stages of infection. Contemporary diagnostic paradigms remain predominantly manual and visual, relying heavily on agricultural professionals, which is often inefficient and inaccessible for remote farmers. While deep learning has demonstrated remarkable efficacy in automated plant disease recognition, existing methodologies frequently lack interpretability, disease severity quantification, and real-world field applicability. This paper introduces a comprehensive, interpretable deep learning-based framework utilizing EfficientNetV2-B0 for classifying potato leaf images into healthy, early blight, and late blight categories. By integrating Gradient-Weighted Class Activation Mapping (Grad-CAM), the model achieves high transparency, highlighting critical prediction regions. Furthermore, a severity assessment module estimates infection percentages, providing actionable treatment recommendations, ultimately enhancing agricultural decision-making.

1. Introduction

Agriculture constitutes the foundational pillar of the economic and social development framework of Bangladesh. It is inextricably linked to national food security, generates massive employment opportunities, and makes a profound contribution to the rural economy. Despite the country's rapid pivot toward industrialization and the service sector, agriculture remains the primary source of livelihood for a vast majority of the population. According to the Bangladesh Bureau of Statistics (BBS), millions of citizens are engaged in agricultural labor in diverse capacities. The exponential population growth, coupled with escalating food demands, shifting climate paradigms, environmental degradation, and the emergence of novel crop pathogens, has necessitated the urgent adoption of modern technology in the agricultural sector.

The integration of artificial intelligence, the Internet of Things (IoT), and precision agriculture has revolutionized farming, rendering processes faster, more efficient, and highly optimized [1]. Among the myriad of crops cultivated in Bangladesh, the potato holds a paramount position due to its high nutritional profile, immense versatility, and significant role in stimulating economic growth. It thrives across diverse climatic zones and ranks as the second-largest food crop in the country, following rice. Given its status as a staple food and a primary export commodity, sustaining healthy potato cultivation is vital for the continued upliftment of rural communities and the national GDP.

Despite the historical success and vast potential of potato farming in Bangladesh, cultivation is highly susceptible to an array of fungal, bacterial, viral, and environmental diseases. Among these phytopathological threats, early blight and late blight are universally recognized as the most devastating foliar diseases. Early blight, primarily induced by the fungal pathogen *Alternaria solani*, manifests initially

as small brown necrotic lesions on the leaves. As the disease progresses, it severely impairs the photosynthetic capacity of the foliage, leading to premature defoliation and drastically reduced yields [4].

Conversely, late blight, caused by the oomycete *Phytophthora infestans*, is notorious for its rapid dissemination during cool, humid weather conditions. If left unchecked, it can decimate entire fields within days, resulting in catastrophic economic losses and threatening regional food security [5]. The conventional diagnostic approach relies predominantly on visual inspection by farmers or agricultural extension officers. This manual method is inherently subjective, prone to substantial error, and exceedingly time-consuming. Misdiagnosis frequently leads to the indiscriminate and excessive application of chemical pesticides, which not only inflates production costs but also instigates severe environmental pollution, degrades soil fertility, and poses grave public health risks [6].

To circumvent the pervasive limitations associated with manual disease detection, there is a critical imperative to harness advanced computational paradigms, specifically deep learning and computer vision. While traditional machine learning models have been explored, they often fail to capture the complex, hierarchical spatial features required for robust disease classification under variable field conditions. Furthermore, contemporary deep learning applications in agriculture frequently operate as opaque "black boxes," providing classifications without transparent reasoning. This lack of interpretability breeds skepticism among end-users and agronomists. Therefore, integrating Explainable Artificial Intelligence (XAI) techniques, such as Grad-CAM, is rationalized as it demystifies the decision-making process, highlighting the exact phenotypic symptoms driving the model's predictions. Coupling this with a disease severity assessment creates a holistic, rationalized approach to precision agriculture.

This research is driven by a multifaceted set of objectives designed to modernize potato disease management in Bangladesh. The primary objective is to develop and optimize a state-of-the-art deep learning architecture (EfficientNetV2-B0) capable of accurately classifying potato leaf images into healthy, early blight, and late blight categories. The second objective is to integrate Explainable AI (Grad-CAM) to visualize and validate the specific leaf regions that influence the model's classification, thereby ensuring algorithmic transparency. The third objective focuses on constructing a severity assessment module to mathematically quantify the percentage of infected foliar area, categorizing the infection into actionable levels (e.g., mild, moderate, severe). Finally, the research aims to provide integrated, rule-based treatment recommendations tailored to the specific disease type and severity level.

The significance of this research lies in its potential to radically transform the agricultural landscape of Bangladesh. By facilitating the rapid, highly accurate, and transparent diagnosis of severe potato foliar diseases, the proposed framework empowers farmers to implement timely, targeted interventions. This precision directly mitigates catastrophic crop losses, thereby safeguarding national food security and bolstering the agricultural economy. Crucially, by accurately identifying disease severity, the system curtails the rampant misuse of chemical pesticides, promoting environmental sustainability, preserving soil microbiomes, and ensuring safer crop yields. Furthermore, the deployment of this technology via

mobile or web applications democratizes access to expert-level diagnostic capabilities, bridging the knowledge gap for marginalized farmers in remote sectors of the country [8], [9].

2. Literature Review

The intersection of computer science and agronomy has birthed the rapidly evolving domain of precision agriculture; wherein automated plant disease detection serves as a critical focal point. Historically, plant pathology relied exclusively on the expertise of human inspectors traversing fields to visually assess crop health. However, the advent of digital image processing, propelled initially by traditional machine learning algorithms and subsequently supercharged by deep Convolutional Neural Networks (CNNs), has fundamentally shifted this paradigm. Researchers globally have increasingly leveraged publicly available datasets, most notably the PlantVillage dataset, to train computational models capable of differentiating between healthy and infected flora with unprecedented accuracy. Despite these advancements, the transition from controlled laboratory environments to volatile, real-world agricultural settings remains a complex challenge, necessitating continuous review and innovation in methodological approaches.

A comprehensive review of contemporary literature reveals a progressive transition from fundamental image processing and machine learning techniques to sophisticated, deep transfer learning architectures for potato disease detection. Early methodologies heavily emphasized manual feature extraction and traditional classifiers. For instance, Islam et al. (2017) [5] proposed an automated diagnostic framework utilizing image processing to segment potato leaf images, subsequently employing a Support Vector Machine (SVM) classifier to differentiate between healthy and diseased states. While foundational, such techniques often faltered when confronted with complex backgrounds and varying illumination in real-world scenarios. Similarly, Singh et al. (2021) [9] utilized K-means clustering for image segmentation and the Gray Level Co-occurrence Matrix (GLCM) method to extract intricate texture features, which were then processed through machine learning algorithms.

As computational capabilities expanded, researchers pivoted toward Deep Learning, specifically Transfer Learning, to bypass the arduous process of manual feature engineering. Tiwari et al. (2020) [2] highlighted the efficacy of transfer learning by utilizing a pre-trained VGG19 architecture to extract hierarchical features from potato leaf images, successfully identifying critical diseases like Early Blight and Late Blight. Building upon this, Upadhyay et al. (2023) [8] reinforced the utility of CNN-based transfer learning, demonstrating that pre-trained networks like VGG19 significantly reduce the computational time required for manual detection while maintaining high diagnostic precision. Tarik et al. (2021) [6] localized the context to Bangladesh, emphasizing the economic burden of potato diseases on local farmers. They developed a rapid detection paradigm leveraging various pre-trained machine learning models on the PlantVillage dataset, achieving enhanced accuracy metrics.

Further solidifying the dominance of CNNs, Bangari et al. (2022) [1] conducted an extensive review of literature, concluding unequivocally that CNN-based models consistently outperform alternative methodologies in automatically identifying plant diseases, thereby mitigating the dependency on scarce

agricultural experts. Tambe et al. (2023) [4] applied rigorous preprocessing to potato leaf images before subjecting them to a custom CNN, thoroughly evaluating the model's performance on a dedicated test dataset. Goyal et al. (2023) [7] addressed the "decision fuzz" prevalent in many standard models by implementing an optimized multi-tier classification system for Early Blight, Late Blight, and healthy leaves using deep learning frameworks. Expanding into the realm of decentralized data privacy, Kamal et al. (2022) [3] introduced a pioneering Federated Learning approach. They trained CNN models, notably Inception V3, locally across different farms without sharing raw image data, subsequently aggregating the learned weights into a highly accurate, globally robust model for multi-farm disease detection.

Despite the extensive proliferation of deep learning models in plant pathology, several critical research gaps persist. Primarily, the vast majority of existing studies construct and evaluate their models using highly sanitized, uniform laboratory data (e.g., the PlantVillage dataset), which severely limits their generalization and robustness when deployed in chaotic, real-world field conditions with complex backgrounds, varying lighting, and diverse sensor qualities. Secondly, current literature predominantly treats these highly complex neural networks as impenetrable "black boxes." The glaring absence of Explainable AI (XAI) prevents users from understanding which visual features the model utilizes to arrive at a diagnosis, hindering clinical trust. Finally, existing frameworks overwhelmingly stop at simple binary or multi-class categorization, failing to provide a granular assessment of disease severity (e.g., the percentage of leaf area compromised), which is an absolute prerequisite for formulating accurate, dose-specific pesticide recommendations.

In conclusion, while the trajectory of existing research underscores the immense potential of artificial intelligence in mitigating agricultural losses, realizing this potential at the grassroots level requires a paradigm shift. The literature firmly establishes the supremacy of CNNs and transfer learning in diagnostic accuracy. However, to bridge the divide between theoretical computer science and practical agronomy in Bangladesh, subsequent systems must transcend mere classification. They must integrate real-world data resilience, transparent interpretability through XAI, and quantitative severity metrics to foster a genuinely intelligent, end-to-end decision-support ecosystem for smart agriculture.

3. Theoretical and Conceptual Framework

The theoretical foundation of this research is anchored in deep learning, particularly Convolutional Neural Networks (CNNs), which are biologically inspired multi-layer perceptrons engineered to process grid-like topologies, such as digital images. CNNs autonomously learn spatial hierarchies of features through sequential convolution, pooling, and non-linear activation layers. To optimize performance and reduce massive computational overhead, Transfer Learning is theoretically employed. Transfer learning operates on the principle that feature representations learned by a network on a massive, generalized dataset (e.g., ImageNet) can be effectively repurposed for a specialized downstream task (e.g., leaf disease classification) by fine-tuning the terminal layers.

Specifically, this study utilizes EfficientNet, a highly advanced architecture that theoretically balances network depth, width, and resolution through an elegant compound scaling method, achieving state-of-the-art accuracy with significantly fewer parameters. Furthermore, the theoretical framework incorporates Explainable AI (XAI) via Gradient-Weighted Class Activation Mapping (Grad-CAM). Grad-CAM theoretically utilizes the gradients of any target concept flowing into the final convolutional layer to produce a coarse localization map, thereby visually emphasizing the regions in the image most influential to the model's predictive output.

Conceptually, the proposed framework operates as a multi-stage, end-to-end intelligent diagnostic pipeline. The process initiates with the acquisition of heterogeneous potato leaf images, encompassing diverse disease profiles. These images undergo standardized preprocessing to ensure spatial and spectral uniformity. The core processing engine is the fine-tuned EfficientNetV2 model, which ingests the preprocessed tensors and extracts complex morphological features indicative of pathogenic infections. The model outputs a probabilistic disease classification. Concurrently, the conceptual framework branches into two critical analytical sub-modules. The first is the interpretability module, where Grad-CAM generates heatmaps overlaying the original image to validate the diagnostic focus. The second is the severity assessment module, which geometrically analyzes the localized diseased regions against the total leaf area to calculate a quantitative severity index. This data conceptually culminates in a rule-based recommendation engine that outputs tailored agronomic advice.

The integrated workflow systematically progresses through several structured phases:

- Phase 1 encompasses 'Dataset Collection,' gathering healthy, early blight, and late blight samples from the PlantVillage dataset and real-world environments.
- Phase 2 involves 'Image Preprocessing,' explicitly standardizing images to 256x256 resolution, preserving RGB integrity, normalizing pixel distributions, and applying noise reduction algorithms.
- Phase 3 implements 'Data Augmentation' to simulate environmental variance using random horizontal flips, rotations, zooming, and translations.
- Phase 4 handles 'Dataset Splitting' into Training (70%), Validation (15%), and Testing (15%) cohorts.
- Phase 5 involves 'Model Training and Validation' leveraging the EfficientNetV2 architecture initialized with ImageNet weights, employing fine-tuning, early stopping, and rigorous checkpointing.
- Phase 6 generates the 'Disease Prediction,' outputting specific class labels with associated confidence probabilities.
- Phase 7 integrates 'Explainability (Grad-CAM)' to synthesize visual heatmaps.
- Phase 8 utilizes advanced bounding box methodologies for 'Disease Localization.'
- Phase 9 performs the crucial 'Disease Severity Estimation' based on affected area ratios. Finally,
- Phase 10 delivers 'Treatment Recommendation,' evaluating all prior metrics against a rule-based logic matrix to prescribe precise agricultural interventions.

4. Methodology

The methodology constructs an integrated, multi-stage end-to-end framework explicitly designed for the intelligent diagnosis of potato leaf diseases under both controlled and real-world conditions. The architecture of the research commences with robust data acquisition. The primary dataset is curated by amalgamating standardized images from the widely recognized PlantVillage database with high-resolution images meticulously collected from real-world potato fields in Bangladesh. This hybrid dataset encompasses a diverse spectrum of disease manifestations, specifically categorized into Healthy, Early Blight, and Late Blight.

To ensure the integrity and compatibility of the input data with the requisite deep learning architectures, rigorous preprocessing protocols are enforced. All raw images are algorithmically resized to a uniform dimension of 224x224 pixels, maintaining the essential RGB color space. Pixel intensities are subsequently normalized to a standardized scale $[0, 1]$ using the equation: $I_{\text{norm}} = I / 255$. This critical normalization step accelerates convergence during gradient descent and enhances training stability while faithfully preserving structural disease features.

4.1 Variable Selection

The dependent variables within this analytical model are the discrete disease classifications: 'Healthy', 'Early Blight', and 'Late Blight'. The independent variables encompass the highly dimensional pixel matrices extracted from the processed leaf imagery. Critical hyperparameters governing the neural network's architecture and learning dynamics include a fixed batch size of 32, a dynamically adjusted learning rate via the Adam optimizer, categorical cross-entropy as the primary loss function, and specific dropout rates set at 0.5 to stringently penalize model overfitting. These variables were empirically selected following rigorous iterative tuning.

4.2 Cross Validation

To ensure the model's robust generalization capabilities and to meticulously assess out-of-sample performance, the unified dataset underwent stratified splitting. Stratification guaranteed that the inherent class distributions were proportionately preserved across all subsets. Specifically, 70% of the dataset was allocated exclusively for the training phase to update the network weights. An additional 15% was designated as the validation set, providing an unbiased evaluation of the model fit during the training epoch cycle and facilitating early stopping mechanisms. The final 15% constituted an isolated testing set, utilized solely post-training for the definitive empirical evaluation.

4.3 Model Evaluation

The predictive efficacy and clinical reliability of the proposed models were quantified utilizing a comprehensive suite of statistical evaluation metrics. Primary among these was Accuracy, defining the global correct classification rate. However, to account for potential class imbalances, Precision, Recall, and the F1-Score (the harmonic mean of precision and recall) were rigorously computed for each

discrete disease class. Furthermore, the Area Under the Receiver Operating Characteristic Curve (ROC-AUC) was calculated to measure the model's discriminative threshold competence.

4.4 Software

The entire computational framework, from data preprocessing to model training and interpretability visualization, was developed using Python programming language. Deep learning architectures were constructed and optimized leveraging the TensorFlow and Keras open-source libraries. Execution was conducted within the Google Colaboratory (Colab) environment, capitalizing on cloud-based GPU acceleration to facilitate rapid tensor processing.

4.5 Ethical Issue

In the context of this agritech research, ethical considerations center primarily on ensuring algorithmic fairness, preventing deployment biases against varying regional crop strains, and promoting sustainable agricultural practices. Care is taken to ensure that real-world image datasets were obtained without infringing on localized farm privacy. Furthermore, minimizing the computational carbon footprint through the utilization of highly efficient network architectures like EfficientNet inherently addresses the environmental ethics of deep learning.

5. Results

The empirical results derived from the extensive experimental setup illuminate the supreme efficacy of advanced transfer learning methodologies over standard deep convolutional networks in the domain of complex phytopathological image classification. The research implemented and rigorously evaluated four distinct neural network architectures: a Custom CNN, ResNet50, MobileNetV2, and EfficientNetB0. Each model was subjected to identical preprocessing pipelines, augmentation strategies, and testing environments to ensure optimal comparative validity.

Initial evaluations focused on the Custom Convolutional Neural Network (CNN). The Custom CNN, structured with four sequential Conv2D layers interspersed with MaxPooling layers and terminating in a dense Softmax classifier, established a robust baseline performance. It achieved a commendable overall Accuracy of 96.18%, with a Precision of 95.92%, Recall of 95.81%, and an F1-Score of 95.86%. The AUC for the Custom CNN stood at 0.982. While mathematically sound and computationally lightweight, the Custom CNN exhibited slight deficiencies in extracting highly granular spatial features when confronted with complex, noisy backgrounds typical of real-world leaf imagery.

Transitioning to advanced transfer learning paradigms, the ResNet50 architecture demonstrated substantial improvements, leveraging its deep residual learning framework to circumvent the vanishing gradient problem. ResNet50 achieved an impressive Accuracy of 98.36%, coupled with an F1-Score of 98.15% and an AUC of 0.994. The residual connections allowed the network to construct a more profound hierarchical understanding of disease morphologies. Following this, MobileNetV2 was deployed, heavily emphasizing operational efficiency suitable for mobile deployment. MobileNetV2

utilized inverted residual blocks and linear bottlenecks to achieve a stellar Accuracy of 98.91% and an F1-Score of 98.81%. Its highly optimized architecture proved that massive parameter counts are not strictly necessary for achieving near-perfect diagnostic capabilities.

However, the pinnacle of the experimental results was unequivocally achieved by the EfficientNetB0 model. Engineered around a highly disciplined compound scaling methodology that uniformly balances network depth, width, and resolution, EfficientNetB0 transcended all comparative baselines. The model achieved a state-of-the-art overall Accuracy of 99.37%. Its Precision reached an extraordinary 99.31%, indicating a negligible false positive rate, while the Recall peaked at 99.26%, ensuring that actual disease instances were almost entirely captured. The F1-Score culminated at 99.28%, and the AUC achieved a near-perfect 0.998, denoting exceptional class separability.

A granular, class-wise performance analysis of the EfficientNetB0 architecture further solidifies its diagnostic supremacy. When evaluating the specific disease classifications, the model exhibited extraordinary precision. For the 'Potato Healthy' class, EfficientNetB0 recorded a Precision of 99.50% and a Recall of 99.40%, yielding an F1-Score of 99.45%. This indicates a profound algorithmic understanding of standard physiological leaf structures. For pathological classifications, the model detected 'Potato Early Blight' with a Precision of 99.10% and a Recall of 99.20% (F1-Score: 99.15%). Similarly, 'Potato Late Blight' was identified with a Precision of 99.30% and a Recall of 99.20% (F1-Score: 99.25%). The marginal variance between early and late blight metrics is phenomenologically attributed to the visually homologous necrosis patterns exhibited during the nascent stages of both diseases. Nevertheless, the high F1-Scores explicitly validate the model's robust discriminative capacity even among morphologically similar pathogen signatures.

6. Discussion

The comprehensive synthesis and comparative discussion of the empirical findings underscore a critical evolutionary leap in automated agronomic diagnostics. The overarching hypothesis-that advanced, scaled transfer learning networks augmented with interpretability modules would significantly outperform traditional deep learning frameworks-was unequivocally validated. The comparative evaluation clearly demonstrates the hierarchical superiority of Transfer Learning models over the fundamental Custom CNN architecture in the complex task of potato leaf disease classification.

The Custom CNN, while providing a rapid inference speed and low computational complexity, inherently lacks the vast, generalized feature extraction capabilities embedded within architectures pre-trained on millions of macroscopic images (like ImageNet). Consequently, its ceiling of accuracy plateaued at 96.18%. Conversely, ResNet50 demonstrated high classification acumen (98.36%) owing to its deep residual connectivity, which preserves gradient integrity across profound network depths. Yet, ResNet50's high architectural complexity and substantial parameter count render it moderately burdensome for edge-device deployment. MobileNetV2 emerged as an exceptional lightweight

alternative, achieving an outstanding 98.91% accuracy. Its highly optimized structural design explicitly caters to low-latency, real-time applications, making it highly suitable for mobile agronomy applications.

Ultimately, EfficientNetB0 was verified as the definitively superior architecture for this specific agricultural domain. By harmoniously scaling the network dimensions, EfficientNetB0 extracted hyper-specific pathogenic features that other models overlooked, culminating in an apex accuracy of 99.37%. It maintained an optimal equilibrium between inference speed, medium computational complexity, and exceptional deployment suitability.

Beyond raw statistical performance, the introduction of Explainable AI (XAI) via Grad-CAM fundamentally elevated the research from a theoretical exercise to a practically deployable clinical tool. The Grad-CAM visual heatmaps effectively stripped away the 'black box' enigma historically associated with deep neural networks. Analytical observations of the heatmaps confirmed that the EfficientNetB0 model logically anchored its predictive confidence on biologically relevant phenomena-specifically targeting the dark necrotic lesions of early blight and the expansive chlorotic decay of late blight-while explicitly ignoring irrelevant background noise such as soil, stems, or ambient shadows. For healthy leaves, the activation maps remained uniformly distributed, correctly indicating an absence of localized pathogenic anomalies. This profound interpretability ensures that agricultural stakeholders can visually audit and implicitly trust the algorithm's diagnostic output.

As a final discussion point regarding practical application, the selected EfficientNetB0 model was successfully serialized into TensorFlow's.keras format and can be seamlessly compiled into TensorFlow Lite (tflite). This architectural fluidity guarantees that the model is fully prepped for immediate integration into Android-based mobile platforms or lightweight Flask web servers. This transition from computational research to tangible field-ready deployment fulfills the ultimate objective of providing real-time, explainable disease diagnosis and precise severity-based agricultural treatments to the farmers of Bangladesh.

7. Conclusion

This study successfully proposed and validated an advanced, explainable deep learning framework integrating EfficientNetV2-B0 and Grad-CAM for the precise detection and severity assessment of critical potato leaf diseases in Bangladesh.

The core objectives centered on formulating a highly accurate classification model, rendering the network's internal decision-making transparent via XAI visual heatmaps, and quantifying disease severity to synthesize intelligent, actionable treatment recommendations.

Empirical evaluations definitively established EfficientNetB0 as the optimal architecture, achieving a state-of-the-art accuracy of 99.37% and an F1-Score of 99.28%. Furthermore, Grad-CAM visualization successfully proved that the model isolates biologically accurate regions of necrosis and pathogen

decay to form its predictions, eliminating "black box" uncertainties. The severity assessment module accurately translated pixel-level infection data into practical severity thresholds.

This research contributes a holistic, end-to-end agritech framework that bridges the gap between complex artificial intelligence and practical field agronomy. It uniquely combines high-tier disease classification, visual interpretability, and automated severity calculation into a single, cohesive deployment-ready pipeline tailored for the agricultural landscape of Bangladesh.

From a policy perspective, integrating this technology at the national agricultural extension level can drastically optimize resource allocation. It empowers agricultural ministries to deploy intelligent monitoring, thereby reducing the systemic overreliance on environmentally destructive chemical pesticides, standardizing smart agricultural practices, and fortifying national food security.

Future research trajectories will focus on vastly expanding the real-world dataset to encompass a wider array of geographic pathogen strains and integrating multimodal data streams from IoT soil sensors and UAV aerial imaging. Additionally, integrating advanced explainable methodologies like SHAP and deploying the framework as a fully autonomous, real-time edge computing mobile application remain primary objectives.

References

1. Bangari S et al (2022) A review on potato leaf disease detection using deep learning, in 2022 International Conference on Electronics and Sustainable Communication Systems (ICESC)
2. Tiwari D et al (2020) Potato leaf diseases detection using deep learning, in 2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS)
3. Kamal M et al (2022) Federated learning based crop leaf disease detection using convolutional neural networks, in 2022 International Conference on Computer and Information Technology (ICCIT)
4. Tambe UY et al (2023) Potato leaf disease detection using convolutional neural networks, in arXiv
5. Islam M et al (2017) Detection of potato diseases using image segmentation and machine learning, in 2017 IEEE Region 10 Humanitarian Technology Conference (R10-HTC)
6. Tarik MI et al (2021) Potato disease detection using machine learning, in 2021 Third International Conference on Intelligent Communication Technologies and Virtual Mobile Networks (ICICV)
7. Goyal B et al (2023) Potato disease classification using machine learning and deep learning techniques, in 2023 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)
8. Upadhyay N et al (2023) Potato leaf disease diagnosis using convolutional neural networks and transfer learning, in Proceedings of International Conference on Communication and Computational Technologies
9. Singh A et al (2021) Potato plant leaves disease detection and classification using machine learning methodologies, in IOP Conference Series: Materials Science and Engineering