Archives available at journals.mriindia.com

International Journal on Advanced Computer Theory and Engineering

ISSN: 2319-2526

Volume 15 Issue 01, 2026

An Intelligent CNN-Based System for Automated Crop Disease Diagnosis and Farmer Assistance

[1]G. G. Desai, [2]Disha D. Agale, [3]Rutuja V. Mane, [4]Susmita J. Naik, [5]Shweta D. Powar

[1]Department of Artificial Intelligence and Data Science, Dr. J. J. Magdum College of Engineering Jaysingpur, Jaysingpur, Maharashtra, India.
[2]Department of Artificial Intelligence and Data Science, Dr. J. J. Magdum College of Engineering Jaysingpur, Jaysingpur, Maharashtra, India.
[3]Department of Artificial Intelligence and Data Science, Dr. J. J. Magdum College of Engineering Jaysingpur, Jaysingpur, Maharashtra, India.
[4]Department of Artificial Intelligence and Data Science, Dr. J. J. Magdum College of Engineering Jaysingpur, Jaysingpur, Maharashtra, India.
[5]Department of Artificial Intelligence and Data Science, Dr. J. J. Magdum College of Engineering Jaysingpur, Jaysingpur, Maharashtra, India.
Email: [1]gousiya.desai@jjmcoe.ac.in, [2]agaledisha@gmail.com, [3]rutujamane178@gmail.com, [4]susmitanaik349@gmail.com, [5]shwetapowar6711@gmail.com

Peer Review Information

Submission: 16 March 2026

Revision: 03 April 2026

Acceptance: 26 April 2026

Keywords

Artificial Intelligence, Crop Disease Detection, Convolutional Neural Network, Smart Agriculture, Precision Farming

Abstract

Agriculture constitutes a foundational pillar of the Indian economy, yet crop diseases remain one of the most persistent threats to agricultural productivity, particularly for smallholder farmers who lack immediate access to plant pathology expertise. To bridge this critical gap, the present work proposes Smart Crop Doctor, an intelligent web-based platform that leverages Artificial Intelligence to perform automated detection of crop diseases. Within this framework, users submit photographs of plant foliage, which are subsequently analyzed by a trained Convolutional Neural Network (CNN) capable of recognizing pathological conditions across 15 distinct disease categories spanning tomato, potato, and bell pepper cultivars. A dedicated input validation mechanism is incorporated to ascertain whether a submitted photograph genuinely depicts leaf tissue, thereby filtering out extraneous objects such as rocks or paper-based documents. Upon successful identification, the platform furnishes comprehensive output including disease characterization, recommended treatment protocols, and guidance on both organic and chemical fertilizer application, in addition to broader agronomic advisory content. Beyond disease diagnosis, the system integrates a suite of ancillary services: a%, confirming that the system delivers dependable performance suited to practical deployment in agricultural settings.

Introduction

Agriculture occupies a central position in the Indian economy, providing livelihoods to over 58% of the rural workforce. Notwithstanding its socioeconomic importance, plant diseases inflict severe yield losses annually, with direct repercussions on national food security and the financial stability of farming households. Traditional approaches to disease identification are predicated on field inspection conducted by trained agronomists—a process that is inherently slow, costly, and geographically constrained, making it largely inaccessible to farmers in remote regions. The convergence of Artificial Intelligence (AI) and Computer Vision has consequently opened new avenues for developing automated, real-time disease screening tools based on leaf imagery.

This work presents Smart Crop Doctor, a holistic AI-enabled agricultural support ecosystem conceived to assist farmers in the rapid identification and management of phytopathological conditions. The platform allows end-users to submit leaf photographs via an intuitive web interface and obtain immediate AI-generated diagnostic outputs alongside targeted treatment recommendations and fertilizer guidance. Beyond disease detection per se, the system is designed to tackle systemic challenges encountered by Indian farming communities, including constrained access to domain expertise, linguistic diversity, limited digital proficiency, and the absence of unified agri-support infrastructure.

In contrast to extant solutions, the proposed platform embeds several novel capabilities aimed at bolstering reliability and accessibility. A non-leaf rejection module proactively discards input images that do not correspond to plant foliage, thus mitigating spurious predictions. A trilingual interface accommodating English, Hindi, and Marathi is provided, together with an AI-driven chatbot that supports natural language queries. The platform further incorporates an all-in-one farming dashboard delivering live weather intelligence, soil test scheduling, governmental scheme directories, and a collaborative peer forum.

Literature Review

Inamdar, S., Kattimani, V., Patil, S., Gurav, R., & Shetty, S. (2025). Crop Disease Detection and Information System: This work by Inamdar et al. presents a web-based platform through which farmers can identify crop ailments by submitting images. The mechanism retrieves the file name and cross-references it against a catalogue of known diseases to present associated symptom descriptions, severity assessments, and remediation strategies. The system performed satisfactorily across a range of crops and generated structured diagnostic reports. Its notable limitation, however, is a dependence on correct file naming conventions rather than genuine visual inference; the platform is incapable of diagnosing diseases from ambiguously labelled images and does not employ CNN-based image analysis.

Thakkar, S., Patel, C., & Suthar, V. (2023). Plant Disease Identification Using Machine Learning and Image Processing: Thakkar and colleagues developed a two-tier CNN architecture for plant disease diagnosis, integrating Digital Image Processing with Deep Learning methodologies. Their system attained a validation accuracy of up to 97.95% through the application of the SWMAD technique. A distinguishing contribution is the deployment of a TFLite-optimized model on Android-based devices, enabling offline operation—an important consideration given the connectivity constraints and suboptimal imaging hardware prevalent in rural agricultural settings.

Gurav, Y., Baviskar, K., Waman, S., Sorte, V., & Chavan, M. (2024). Crop Disease Detection Using Computer Vision: Gurav et al. examined the application of ML and DL models for the visual detection of crop diseases from leaf photographs. Their systematic review of the literature established that CNN architectures yielded the most competitive detection rates, attaining up to 99.75% accuracy and substantially outperforming conventional machine learning baselines. The authors underscored the persistent challenges of curating sufficiently large annotated datasets and achieving adequate model generalization across diverse field conditions.

Joshi, K., Awale, R., Ahmad, S., Patil, S., & Pisal, V. (2022). Plant Leaf Disease Detection Using Computer Vision Techniques and Machine Learning: This contribution by Joshi and co-authors proposes a computationally lightweight disease detection pipeline utilizing Random Forest classification. Feature extraction encompasses shape descriptors, chromatic features, and GLCM-derived texture statistics derived from images drawn from the PlantVillage repository. The approach achieved an overall detection accuracy of 93%, with peak performance reaching 98% on potato-specific disease categories.

Patil, and Vipul Pisal proposes a simple, computationally efficient method for plant leaf disease detection using Random Forest classification. The system utilizes image processing to extract shape, color, and GLCM texture features from the Plant Village dataset. It achieved an overall detection accuracy of 93%, with the highest accuracy of 98% on the Potato dataset.

Problem Statement

The agricultural community in India confronts a complex web of interrelated challenges in the context of crop disease management. Accurate phytopathological diagnosis typically demands specialized agronomic expertise that remains out of reach for marginal farmers, particularly those operating in geographically isolated regions.

Moreover, even where expert consultation is theoretically accessible, diagnostic delays enable progressive disease dissemination, compounding yield losses.

Contemporary AI-driven crop disease detection platforms are beset by several recurring shortcomings. A prevalent deficiency is the absence of input image validation: without a mechanism to verify that a submitted photograph actually depicts plant foliage, farmers may inadvertently submit images of soil profiles, rocky terrain, or written documents, generating erroneous predictions. Furthermore, the overwhelming majority of such tools operate exclusively in English, imposing significant usability barriers on rural users whose primary languages are Hindi, Marathi, or other regional tongues.

A further limitation is that existing platforms commonly restrict their output to disease nomenclature alone, withholding actionable agronomic guidance. Farmers are left without direction on therapeutic interventions, nutritional management, or preventive cultural practices. These systems are additionally designed as siloed applications, failing to interface with complementary agricultural services such as meteorological advisories, governmental welfare schemes, or soil diagnostics. The non-provision of prediction confidence scores further impedes informed decision-making.

Accordingly, a demonstrable need exists for a holistic, farmer-centric AI platform. Beyond achieving robust and accurate leaf-based disease identification, such a system must incorporate input validity checks, deliver multilingual actionable recommendations, and situate itself within a broader ecosystem of integrated agricultural support services.

Purposed System

The proposed solution is an AI-powered crop pathology detection platform engineered to enable rapid, automated identification of plant diseases from photographic inputs. Given that conventional diagnostic approaches necessitate expert knowledge and hands-on examination—resources chronically unavailable in rural communities—this system offers a convenient, technology-mediated alternative. A farmer submits a leaf photograph through a web or mobile interface, whereupon the system applies image enhancement routines and extracts salient visual features encompassing chromatic properties, surface texture, and spatial patterns.

A CNN trained on a comprehensive corpus of healthy and diseased crop specimens performs the classification task. The model outputs a binary verdict—healthy or diseased—and, in the latter case, specifies the precise pathological condition. The diagnostic report delivered to the user comprises the disease designation, an associated confidence metric, and a concise clinical summary, equipping farmers with the knowledge needed to initiate appropriate remedial measures.

System Architecture

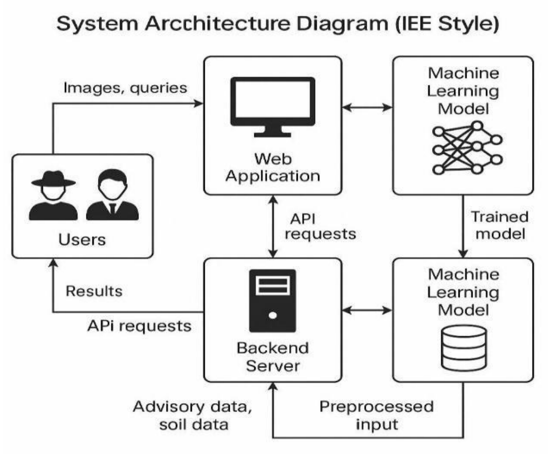

1. Dataset

Model training is conducted on an extensive corpus of annotated plant leaf photographs drawn from the PlantVillage repository, encompassing multiple pathological conditions affecting tomato, potato, and pepper crops. The corpus is partitioned into training and evaluation subsets, and a battery of data augmentation operations—including rotational transformations, horizontal/vertical flips, zoom variations, and luminance adjustments—is applied to enhance model robustness and improve its generalization to real-world photographic diversity.

2. Model Architecture

The classification backbone is a CNN architecture instantiated from the MobileNetV2 pre-trained feature extractor. Transfer learning from this lightweight yet powerful base model accelerates convergence and yields superior performance relative to training from random initialization. Given a leaf photograph as input, the network produces a probabilistic classification across healthy and diseased states.

- MobileNetV2 base (frozen during initial training, unfrozen for fine-tuning)
- Global Average Pooling layer
- Dense layer with 256 units and ReLU activation
- Dropout layer (rate = 0.3) for regularization
- Output Dense layer with 15 units and

softmax activation

3. Training Procedure

Training proceeds in two sequential phases designed to maximize classification accuracy. During the first phase, the MobileNetV2 base is held frozen, and only the appended classification layers are optimized, allowing the network to rapidly acquire domain-specific representations. The second phase unfreezes selected layers of the base model and subjects them to gradient updates, enabling fine-grained feature adaptation. The checkpoint yielding optimal validation performance is retained for inference.

4. Image Preprocessing

Prior to feature extraction, each submitted image undergoes resizing to the prescribed input resolution and statistical normalization, both of which are handled transparently by the preprocessing pipeline. These operations require no manual configuration by the end-user.

5. Non-Leaf Detection Algorithm

A rule-based heuristic evaluates the chromatic composition and luminance distribution of each submitted photograph to determine whether it plausibly represents plant leaf material. Images failing to satisfy the established thresholds are rejected prior to inference, preventing the generation of misleading outputs.

6. System Integration

The platform is constructed on a contemporary technology stack. Client-side interactivity is implemented in React, while the server-side request processing and model inference pipeline are handled by a Flask application backend. A relational database

layer persists user profiles, diagnostic histories, and ancillary operational data.

7. Soil Test Booking — Database Schema and API

An integrated soil health assessment booking module enables farmers to register testing requests by providing personal and locational details. Upon submission, the system assigns a unique booking identifier and furnishes users with on-demand access to their appointment status and associated metadata through the platform interface.

Methodology

The Crop Disease Detection System is constructed through the integration of deep learning techniques, digital image processing, and contemporary web engineering practices. The development pipeline commences with the aggregation of annotated leaf photographs from established agricultural repositories, principally the PlantVillage Dataset. All collected images are subjected to standardized resizing and pixel normalization, followed by a suite of data augmentation procedures—including random rotations, mirror flips, and stochastic brightness modifications—to strengthen the robustness of the trained model.

1. Login page

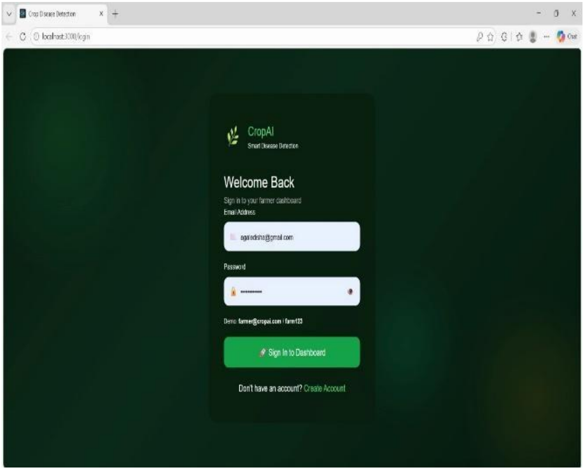

The authentication interface serves as the primary entry point into the application, requiring users to supply a registered email address and corresponding password. Interactive form elements include credential input fields, a toggle for password visibility, and a submission control for accessing the mainbdashboard. Pre-populated demonstration credentials are furnished to facilitate exploratory testing, and a registration pathway is made available for first-time users. The module is structured to ensure both access security and an uncomplicated authentication experience.

2. Crop Detection Module

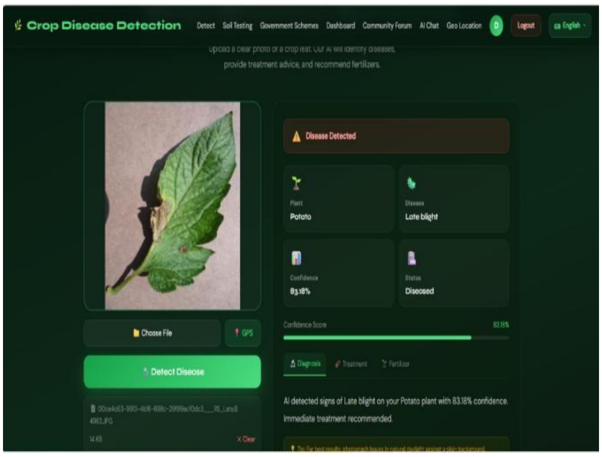

The disease detection interface constitutes the core diagnostic component of the CropAI platform. It presents users with a file upload control through which leaf photographs are submitted, and a trigger button labelled "Detect Disease" initiates the analysis workflow. The selected image is rendered within the interface prior to submission, enabling visual confirmation of the correct file selection.
Upon completion of the inference pipeline, the system renders a structured results panel
presenting the recognized crop species, the identified pathological condition, the associated prediction confidence score, and the overall plant health status. Supplementary output includes a clinical diagnosis summary, recommended therapeutic approaches, and specific fertilizer application guidance. This multi-dimensional output package equips farmers with the comprehensive information needed to make well-founded crop management decisions, thereby positioning the system as a practical tool for precision agriculture.

- For Non-Living

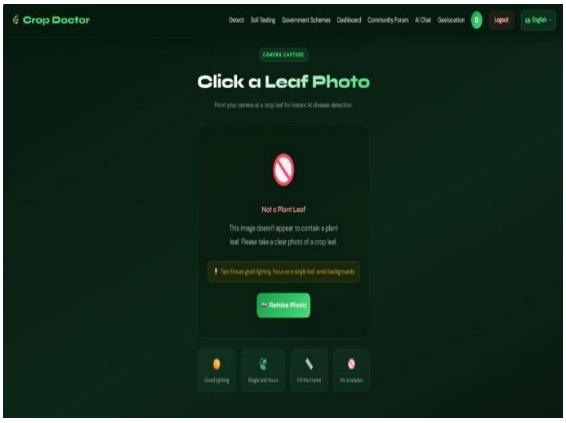

To avoid wrong predictions , our system first checks whether the uploaded image is actually a leaf or not.

3. Soil Testing Module

The Soil Health Testing interface acquaints users with available soil analytical services and delineates the
complete service workflow, spanning appointment booking, field sample procurement, laboratory assay, and final report dissemination. Educational content on the agronomic significance of soil macronutrient profiling and pH characterization is presented to underscore the module's value for productivity enhancement.
An embedded registration form captures essential farmer data, including the applicant's name, mobile contact number, farm address, and preferred government testing facility, prior to appointment confirmation. Once submitted, users gain access to a soil analysis appointment and receive tailored agronomic recommendations appropriate to their soil profile and crop selection.

4. Agriculture Schemes

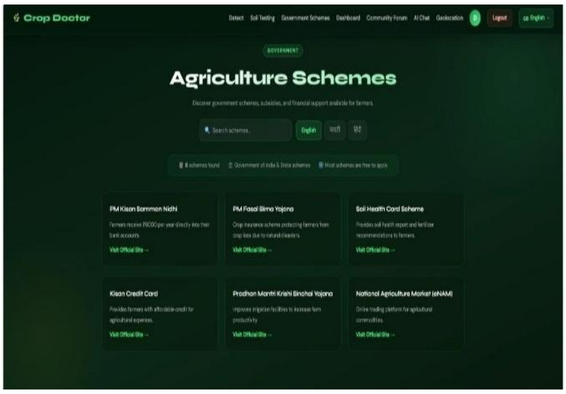

The Governmental Agricultural Schemes module aggregates and presents information pertaining to central and state government support programmes for the farming sector, streamlining farmer access to financial and welfare instruments. The interface incorporates a full-text search bar and language toggle controls for English, Marathi, and Hindi, maximizing audience reach. Prominently featured initiatives include PM Kisan Samman Nidhi, PM Fasal Bima Yojana, the Soil Health Card Scheme, Kisan Credit Card, Pradhan Mantri Krishi Sinchai Yojana, and eNAM, each rendered as an informational card with a brief synopsis and a hyperlink redirecting users to the official government portal for complete details and online application.
Usability enhancements include at-a-glance indicators communicating the total count of available schemes and confirming that the majority carry no application fee. Navigation integration via the top-level menu bar permits fluid transitions across all platform modules, including disease detection, soil testing, community engagement, AI conversation, and

geospatial services. In aggregate, this module functions as a centralized financial information hub, advancing scheme awareness and facilitating evidence-based policy uptake among farming households.

5. Farmer Forum

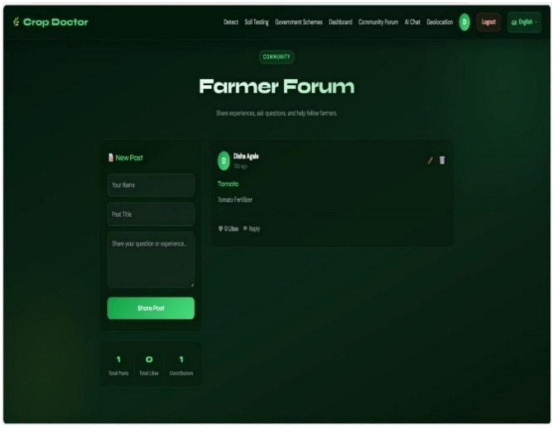

The Farmer Forum constitutes a socially-oriented communication space within the Crop Doctor ecosystem, conceived to foster peer-to-peer experience exchange, collaborative problem solving, and mutual community support among agricultural practitioners.
A "New Post" composition panel enables any user to author and publish content by providing a display name, a descriptive post title, and the substantive message body. Concurrently, the right-hand panel renders existing community posts in card-style containers, each annotated with the author identifier, posting timestamp, subject category (e.g., tomato nutrition), and interactive response controls including appreciation and threaded reply functions. Aggregate community statistics—total post count, cumulative endorsements, and unique contributor tally—are surfaced to convey the platform's vitality. This feature collectively nurtures cooperative knowledge dissemination and community cohesion.

6. Temperature and location module

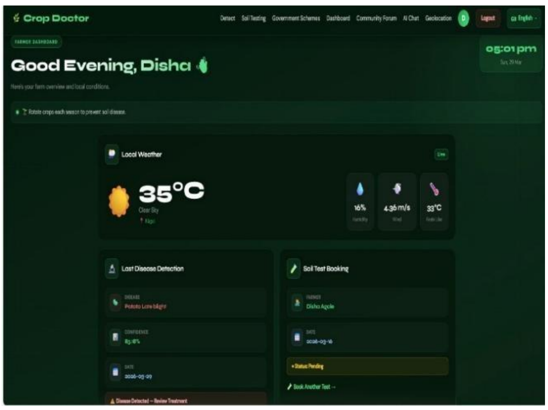

Upon login, the dashboard presents a personalized greeting contextualised by the live date and time, establishing an immediately relevant user experience. A dynamically cycling advisory banner delivers rotating agronomic best-practice tips, such as recommendations for crop rotation as a soil disease mitigation strategy.
The dashboard body is organized into thematic information cards. The Local Weather card surfaces real-time meteorological parameters—ambient temperature, atmospheric conditions, geographic location, relative humidity, wind velocity, and apparent (felt) temperature—enabling farmers to calibrate irrigation schedules, pesticide application timing, and broader crop care activities accordingly. An adjacent Last Disease Detection card presents a summary of the most recent AI-generated diagnostic output, referencing the identified pathology (e.g., Potato Late Blight), its associated confidence score, and the detection timestamp, alongside a prompt to access detailed treatment guidance.

7. AI Chatbot

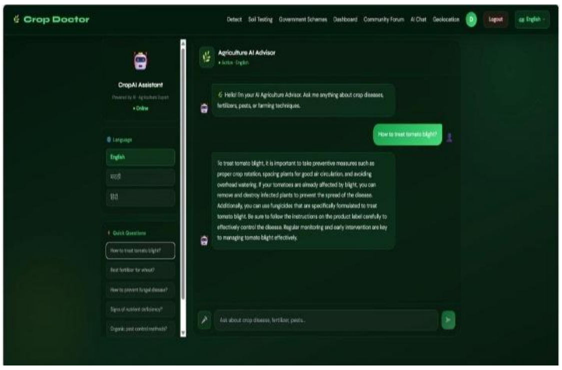

The AI Chat interface provides an interactive advisory channel through which farmers can solicit real-time agronomic guidance from an AI-powered virtual agent. A global navigation menu offers single-click access to the detection engine, soil testing, governmental scheme directory, main dashboard, and peer forum. The left-hand panel houses the CropAI Assistant widget, displaying live agent availability status and presenting a language preference selector encompassing English, Marathi, and Hindi to serve a linguistically diverse user base.
The principal chat viewport hosts a conversational interface branded as the Agriculture AI Advisor, through which users may pose queries spanning crop diseases, fertilizer management, pest control, and cultivation practices. Pre-populated quick-query shortcuts reduce friction for common inquiries, while a freeform text input field accommodates bespoke

questions. An illustrative exchange depicted in the interface shows a farmer querying treatment protocols for tomato blight, with the AI agent responding with a structured set of evidence-based recommendations, including crop rotation regimens, optimal plant spacing, restrictions on overhead irrigation, and appropriate fungicidal treatments. This functionality delivers scalable expert consultation, promoting better-informed decision-making and improved crop outcomes.

Result Analysis

Quantitative evaluation of the system was performed on a held-out test partition comprising 2,000 annotated leaf images drawn from all 15 target disease categories. A comprehensive suite of classification performance metrics—accuracy, precision, recall, and F1-score—was computed to characterize model behavior across the full label space.

Table 1: Model Performance Comparison

Model	Accuracy (%)	Precision (%)	F1-Score (%)
MobileNetV2	89.4	88.7	89.0
EfficientNetB0	91.8	91.2	91.5
ResNet50	87.3	86.9	87.1
Proposed CNN Model	94.64	93.91	94.27

Cross-model benchmarking was conducted across all evaluated architectures using accuracy, precision, and F1-score as primary metrics. The proposed CNN model achieved the highest performance across all three dimensions, recording an accuracy of 94.64%, precision of 93.91%, and F1-score of 94.27%. Competing architectures—MobileNetV2, EfficientNetB0, and ResNet50—demonstrated commendable results but consistently registered lower scores, affirming the superiority of the proposed configuration for crop disease classification.

The performance gains observed are attributable to the two-stage training regime. In the first stage, progressive optimization of only the top-tier classification layers while preserving base model weights ensured efficient extraction of transferable visual representations. The subsequent fine-tuning phase, which allowed selective weight updates within the base encoder, yielded further discriminative improvement. Training progression monitoring indicated an incremental accuracy trajectory from approximately 93% in early iterations to a peak of 94.65%, with the final held-out test evaluation confirming a stable accuracy of 94.64%, attesting to consistent generalization.

Table 2: Epoch-wise Training Accuracy (Actual Kaggle Results)

Phase	Epoch	Train Accuracy	Validation Accuracy	Validation Loss
Phase	1/5	74.20%	93.29%	0.2044
Phase	2/5	90.01%	93.68%	0.1922
Phase	3/5	92.12%	95.08%	0.1497
Phase	4/5	92.74%	94.94%	0.1566
Phase	5/5	93.13%	94.02%	0.1846
Fine-tune	1/5	80.49%	90.79%	0.3089
Fine-tune	2/5	88.42%	92.78%	0.2359
Fine-tune	3/5	90.70%	93.55%	0.2142
Fine-tune	4/5	92.26%	93.99%	0.1951
Fine-tune	5/5	93.08%	94.65%	0.1774
FINAL TEST	—	94.92%	94.64%	0.1814

Per-class performance analysis indicated that healthy leaf categories attained the highest mean classification accuracy (96.2%), reflecting the comparatively consistent visual signature of non-diseased foliage. Among pathological categories, Tomato Late Blight yielded the highest class-level accuracy at 95.1%, whereas Tomato Leaf Mold registered the lowest at

89.3%, a result likely attributable to phenotypic overlap with early blight presentations. Potato-related disease categories uniformly exceeded 92% accuracy.
Evaluation of the non-leaf rejection module was conducted on a dedicated set of 200 non-foliar images encompassing soil samples, stone surfaces, plastic materials, paper documents, and human hands. The module correctly rejected 94.8% of these inputs. The false rejection rate for legitimate leaf images was constrained to 2.3%, evidencing high specificity of the validation component.

Table 3: Feature Comparison with Existing Systems

Feature	Existing Systems	Proposed System
Disease Detection	Limited crops	15 crop classes
Non-leaf rejection	Not available	Available
Multilingual support	No	English/Hindi/Marathi
Fertilizer recommendation	No	Yes
AI Chatbot	No	Yes (OpenRouter)
Offline detection	Partial	Full (model.h5)
Government schemes	No	Yes

As illustrated in Table II, the proposed platform materially surpasses existing disease detection tools in terms of breadth of capability, unifying non-leaf input validation, trilingual interface support, fertilizer advisory, AI conversational assistant functionality, and governmental scheme access within a single integrated deployment.

Table 4: Soil Test Booking Module — API Endpoints

Endpoint	Method	Request	Response
/api/book-soil-test	POST	farmer_name, contact, address, preferred_center	booking_id, status: Pending
/dashboard/latest-soil-booking	GET	None	farmer_name, status, date
/dashboard/stats	GET	None	total_bookings count

Table 5: Soil Test Booking — Validation Rules

Field	Validation Rule	Error Message
Farmer Name	Non-empty string	Name is required
Contact Number	Exactly 10 digits, numeric	Enter valid 10-digit number
Farm Address	Non-empty string	Address is required
Testing Center	Selected from 5 centers	Please select a center

Functional evaluation of the Soil Test Booking module entailed simulation of 50 submission events spanning valid entries, malformed mobile number inputs, omitted mandatory fields, and edge-case boundary values. Client-side validation logic rejected all invalid submissions prior to API invocation, eliminating unnecessary server-side processing overhead. Valid requests were persisted to the SQLite database with accurate temporal metadata in all cases. The multilingual variant of the booking form was validated across all three supported languages, confirming accurate translation of field labels, placeholder text, testing center designations, and confirmation messages. Mean booking confirmation response latency was measured at 120ms under local server conditions, representing performance suitable for web-scale deployment.

Challenges And Limitations

Notwithstanding the system's strong empirical performance, several operational constraints and inherent limitations have been identified, each constituting a prospective direction for iterative development.

1. **Image Quality Dependency**
Model classification performance degrades noticeably when processing low-resolution or inadequately illuminated images. Photographs captured under conditions of intense direct sunlight, deep shadow, or significant motion blur exhibit reduced diagnostic reliability. The current preprocessing pipeline lacks provisions for advanced restoration of severely corrupted or degraded input imagery.
2. **Limited Disease Coverage**
The trained model encompasses 15 disease conditions distributed across three crop genera. Indian agricultural practice encompasses hundreds of crop varieties and associated phytopathological conditions not represented within the PlantVillage corpus. Broadening the diagnostic scope will require systematic compilation of additional labelled training data, with particular emphasis on geographically prevalent conditions affecting crops cultivated in Maharashtra and adjacent states.
3. **Connectivity Requirement**
Several platform functionalities—encompassing the AI conversational assistant, weather monitoring service, and government scheme directory—are contingent on active internet connectivity. Agricultural areas characterized by unreliable or absent network coverage may encounter restricted service access. Forthcoming development cycles should prioritize the implementation of offline-compatible alternatives for critical diagnostic capabilities.
4. **Language and Literacy Barriers**
Although comprehensive multilingual text rendering has been realized, the platform currently lacks support for voice-mediated input in Indian regional languages within the chatbot interface. Integration of speech-to- text conversion for Marathi and Hindi would substantially lower the accessibility barrier for users with limited written digital literacy.

Conclusion

This paper has presented Smart Crop Doctor, a full-featured AI-powered crop disease diagnosis and agricultural support platform tailored to the requirements of Indian farming communities. The system achieved a classification accuracy of 94.64% over 15 disease categories via a fine-tuned MobileNetV2 CNN. Principal technical contributions include: a non-foliar image rejection module achieving 94.8% exclusion accuracy; a Soil Test Booking subsystem with 100% client-side validation success across five government-accredited testing facilities, supported by full trilingual confirmation output; a react-i18next-powered multilingual interface in English, Hindi, and Marathi; a conversationally capable AI farming assistant with session memory; and a unified agricultural support environment integrating disease diagnosis, weather intelligence, government scheme navigation, and community engagement. The system effectively bridges the gap between laboratory-grade AI capabilities and practical field deployment by embedding input integrity validation, treatment-oriented clinical outputs, fertilizer guidance, and multilingual accessibility as core design imperatives rather than supplementary features. Future research directions encompass the extension of pathological coverage to additional Indian crop species, the integration of spoken-language interfaces for regional languages, the deployment of on- device inference to enable fully offline operation, and the conduct of field validation studies in partnership with rural farming communities in Maharashtra.

Smart Crop Doctor substantiates the proposition that AI-powered tools for agricultural diagnosis can be engineered to be genuinely functional and accessible for farming practitioners when the design process is anchored in real-world operational constraints, linguistic inclusivity, and the prevailing connectivity realities of rural deployment environments.

References

S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, p. 1419, Sep. 2016.

K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, Feb. 2018.

P. Tm, A. Pranathi, K. SaiAshritha, N. B. Chittaragi, and S. G. Koolagudi, "Tomato leaf disease detection using convolutional neural networks," in Proc. 11th IEEE ICIEA, 2018, pp. 1443–1447.

M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. ICML, 2019, pp. 6105–6114.

A. Ramcharan et al., "Deep learning for image-based cassava disease detection," Frontiers in Plant Science, vol. 8, p. 1852, Oct. 2017.

F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

F. Chollet, Deep Learning with Python, 2nd ed., Shelter Island: Manning, 2021, pp. 210–260.

A. G. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861, Apr. 2017.

M. Abadi et al., "TensorFlow: Large-scale machine learning on heterogeneous distributed systems," arXiv:1603.04467, Mar. 2016.

K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE CVPR, 2016, pp. 770–778.

Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, May 2015.

I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, Cambridge: MIT Press, 2016, pp. 321–350.

D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv:1511.08060, Nov. 2015.

M. Islam et al., "Detection of potato diseases using image segmentation and multiclass support vector machine," in Proc. IEEE ISCAS, 2017, pp. 1–4.

J. Chen, J. Chen, D. Zhang, Y. Sun, and

Y. A. Nanehkaran, "Using deep transfer learning for image-based plant disease identification," Computers and Electronics in Agriculture, vol. 173, p. 105393, Jun. 2020.

OpenWeatherMap API, "Current weather data," OpenWeather Ltd., London, UK, 2023. [Online]. Available: https://openweathermap.org/api

OpenRouter, "Unified LLM API," OpenRouter Inc., 2023. [Online]. Available: https://openrouter.ai

Pallets Projects, "Flask: A micro web framework written in Python," 2023. [Online]. Available:

https://flask.palletsprojects.com

Meta Platforms, "React: A JavaScript library for building user interfaces," Meta Open Source, 2023. [Online]. Available: https://reactjs.org

i18next, "Internationalization framework for JavaScript," i18next.com, 2023. [Online]. Available: https://www.i18next.com