iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.81949

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Plant Disease Recognition Using ML

Mr. Harshal Chandekar[3], Ms. Ritika Kale[2], Mr. Yash Barbate[5], Mr. Rohan Katre[4], Prof. Anjali Bhoyar[5]
[1, 2, 3, 4]*Student, Department of Computer Science and Engineering, Govindrao Wanjari College of Engineering and Technology, Nagpur*
[5]*Assistant Professor, Department of Computer Science and Engineering, Govindrao Wanjari College of Engineering and Technology, Nagpur*

Abstract: ***Crop cultivation sustains the livelihoods of a substantial portion of households across developing economies, yet the crops on which those households depend are perpetually at risk from pathogenic infections that erode both yield volume and produce quality. Infected fields, when not addressed at the right time, translate into mounting financial strain that smallholder growers — who operate with limited financial reserves — are ill-equipped to withstand. The dominant method of spotting such infections today still relies on a farmer walking the field and judging leaf condition by eye, or waiting for an agronomist's visit — a workflow that is neither fast nor consistent enough for large-scale cultivation.***

Progress in deep learning has fundamentally changed what automated visual inspection can accomplish, and plant pathology diagnosis is one of the fields that has benefitted most visibly. Image-based pipelines can now scan a leaf photograph and return a disease classification in fractions of a second. Among the architectures driving this capability, Convolutional Neural Networks occupy a central role: their layered filter design allows them to extract and encode visually informative features — texture discontinuities, color anomalies, lesion geometry — without any manual specification of what to look for.

We present a CNN-driven plant disease recognition system that operates on photographs of plant leaves and returns a disease label together with a confidence estimate. Our training corpus is the Plant Village benchmark collection, a large repository of annotated leaf images spanning healthy and diseased specimens across multiple crop varieties. The input pipeline applies spatial normalization, pixel rescaling, and augmentation strategies to condition the data before it reaches the network. A React.js browser interface connects end users to the model via a lightweight prediction API, enabling diagnosis without any specialist involvement. Validation results affirm that this deep learning approach surpasses rule-based image processing baselines on both accuracy and response time, and the system holds clear potential for adoption in early disease management program.

Keywords: Plant Disease Detection, Machine Learning, Convolutional Neural Networks (CNN), Image Processing, Plant Village Dataset, Precision Agriculture.

I. INTRODUCTION

Few sectors carry as much weight in determining a nation's food security and rural employment as crop farming. Across India, tens of millions of families draw their primary income from agriculture, making the health and productivity of cultivated land a matter of direct national concern. What complicates the outlook is that farmland productivity does not exist in isolation — it sits at the intersection of climatic variability, soil biology, water availability, pest pressure, and pathogen load, any one of which can undermine a season's work.

Pathogen-driven crop losses deserve particular attention among these threats. Once a fungal or bacterial infection establishes itself within a field, it can propagate across neighbouring plants with little warning, and by the time visible symptoms are widespread, corrective action is already costly. Estimates from agricultural research place annual production losses attributable to plant pathogens in the range of 20 to 40 percent of global output — a figure that underscores how much improvement in early detection could mean for food security. The challenge is that the dominant detection method in most smallholder contexts is still direct visual examination: a farmer or agronomist inspects leaves and stems, draws on experience to name the condition, and — when uncertain — submits samples to a laboratory for confirmation. This chain is fragile in at least three ways: it depends on specialized knowledge that is distributed unevenly across the farming population, laboratory turnaround is too slow to guide immediate field decisions, and the process does not scale to large or multi-parcel holdings.

Machine learning has introduced a different paradigm for this kind of diagnosis. Instead of encoding expert knowledge into explicit decision rules, ML systems infer those rules from labelled examples — and once trained, they can apply what they have learned at inference speeds that no human examiner can match. Agriculture has become one of the productive application domains for this technology, with deployments spanning yield forecasting, irrigation scheduling, soil characterization, pest monitoring, and disease identification.

©IJRASET: All Rights are Reserved

Within the plant disease detection problem specifically, CNN-based image classifiers have emerged as the strongest performing family of approaches. The spatial filtering mechanism at the core of a CNN is well matched to the visual nature of leaf pathology: symptoms manifest as local changes in texture, color distribution, and surface morphology — precisely the kinds of signals that convolutional feature detectors are designed to capture. A trained model can register infection indicators that a non-specialist observer might dismiss as normal variation. This work targets the development of such a recognition system. We train a CNN-based model on leaf imagery sourced from the Plant Village collection, coupling it with a web-based interface through which growers and extension workers can submit field photographs and receive near-instant diagnostic output. The broader aim is to put a reliable, accessible disease identification tool into the hands of users who currently lack one, reducing the window between infection onset and informed response.

II. LITERATURE SURVEY

Scholarly work on automated crop disease identification has grown substantially since image-based deep learning became computationally accessible. The nine studies reviewed below span the period from the field's formative years through to the present frontier, and together they trace an arc from proof-of-concept benchmarking through to deployment-oriented systems and the first explorations of transformer-based architectures.

A. Deep Learning-Based Plant Disease Detection (Mohanty et al., 2016)

The paper by Mohanty and colleagues is widely credited with establishing CNN-based plant pathology as a credible research direction. Working with the newly assembled Plant Village image collection — a corpus of annotated leaf photographs covering 26 disease classes across 14 crop types — the team trained two convolutional architectures and recorded classification performance that, under controlled laboratory conditions, reached accuracy figures above 99 percent. What made the study influential was not merely the number but the demonstration that a network could arrive at disease-discriminative representations entirely through gradient-driven learning, with no manual engineering of diagnostic features by the researchers. The study opened the door for a generation of follow-on work applying similar methods to agricultural image classification.

B. Leaf Image Classification Using Deep Neural Networks (Sladojevic et al., 2016)

Sladojevic and colleagues approached the classification problem with an emphasis on image quality rather than architecture novelty. Before any network training took place, they invested in a preprocessing stage that isolated the leaf foreground from scene backgrounds — a step motivated by the observation that background clutter interferes with feature learning when images originate from real field conditions rather than a controlled studio. The resulting classification system, trained on photographs of symptomatic and healthy specimens across multiple crop types, delivered accuracy gains that validated the preprocessing investment. The authors concluded that embedding such a model within a mobile interface would allow growers to perform in-field diagnosis without external assistance.

C. Deep Learning Models for Agricultural Image Analysis (Ferentinos, 2018)

Ferentinos conducted what is arguably the most methodologically careful architecture comparison in the early plant disease detection literature. Rather than committing to a single model, he trained and evaluated several CNN variants under matched conditions, varying both the dataset size and the disease diversity of the classification task. The consistent outcome was that deeper, more parameterized networks outperformed shallower counterparts, and that all deep learning configurations exceeded classical feature-engineering pipelines by a substantial margin. The study additionally noted the practical implication: strong CNN performance is attainable without resorting to hand-crafted feature descriptors, which simplifies the engineering effort required to build deployable systems.

D. Comparative Study of Deep Learning Models (Too et al., 2019)

Too and co-authors addressed a question that practitioners frequently face but that earlier literature had not answered systematically: when multiple competitive architectures are available, which should be chosen, and does the answer change if a pretrained initialization is used? Their comparison spanned VGG, ResNet, and Inception families, tested both from-scratch and fine-tuned conditions, and used the full Plant Village benchmark for evaluation. Fine-tuned DenseNet variants topped the results table, and the margin over from-scratch training was large enough to firmly establish fine-tuning from pretrained weights as the recommended starting point for groups without access to extended training cycles. This finding directly informs architecture choices in our own system design.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

E. Factors Affecting Deep Learning Performance (Barbedo, 2020)

Where earlier papers focused on what deep learning could achieve under favorable conditions, Barbedo asked a harder question: what causes these models to fail, and how frequently do those causes arise in real agricultural settings? His analysis catalogued several recurring failure modes — among them the visual overlap between different infections, the way environmental stress can mimic pathogen symptoms, and the stark difference in image characteristics between controlled training data and uncontrolled field photographs.

The paper did not offer an algorithmic fix but framed the problem clearly enough to make subsequent field-robustness work easier to evaluate. His critique remains relevant: any system trained exclusively on clean benchmark imagery will face a generalization gap when released into real agricultural environments.

F. CNN-Based Plant Disease Detection with Data Augmentation (Wang et al., 2020)

Wang and colleagues took a practical stance on the generalization gap by introducing a structured augmentation regime into the CNN training pipeline.

Going beyond the conventional geometric transforms, they included photometric perturbations — adjustments to brightness, contrast, and color balance — designed to simulate the range of lighting conditions encountered when photographs are taken outdoors with consumer-grade cameras.

Models trained against this augmented corpus showed materially better performance on held-out field images than baseline models trained without augmentation, and the trade-off in training time was modest. The study established structured augmentation as a near-mandatory component of any deployment-oriented training design.

G. Mobile-Based Plant Disease Classification (Barman et al., 2021)

Barman and co-authors shifted attention from classification accuracy in isolation to the constraints imposed by the deployment target.

Designing specifically for mid-range Android hardware, they produced a compressed CNN variant capable of completing an inference pass within a few seconds on a device without a dedicated neural processing unit. The system accepted photographs from the device camera directly, which eliminated the step of image transfer that web-based systems require.

Their frank discussion of the accuracy-versus-compression trade-off — noting that aggressive quantization produced disproportionate accuracy drops in certain disease categories — is a useful caution for any group considering on-device deployment over server-hosted inference.

H. Transfer Learning for Crop Disease Detection (Saleem et al., 2022)

Saleem and colleagues examined how far a researcher could go with limited compute by leveraging models pretrained on ImageNet-scale data. Using MobileNet as the primary base architecture, they adapted the network's classification head to the plant disease task through targeted fine-tuning, achieving accuracy comparable to much larger models while keeping memory footprint and inference cost low.

Their experimental setup also surfaced an underappreciated sensitivity: model performance degraded noticeably on disease classes that were underrepresented in the training partition, suggesting that class balance management is as important as architecture selection when building reliable classifiers.

I. Vision Transformer-Based Detection (Wang et al., 2025)

The most recent contribution we reviewed introduces Vision Transformers into the plant disease classification problem. Transformer architectures differ fundamentally from CNNs in how they process spatial information: rather than building local feature maps through convolution, they partition the image into patch tokens and compute attention weights that describe how each patch relates to every other.

Wang and colleagues found that this global attention mechanism produced strong results on clean, well-lit images but was more susceptible than ResNet baselines to the kinds of degradation common in field photography — partial occlusion, moisture, and motion blur among them. The implication is that transformer models are not yet a drop-in replacement for CNN-based classifiers in agronomic settings where image conditions are variable.

III. PROPOSED WORK

Our objective is to construct a functional plant disease recognition system anchored in deep learning and made accessible through a web interface. The clinical context that motivates this design is straightforward: across smallholder farming regions, disease identification still hinges on the unaided observation of an individual — a farmer or visiting specialist — who assesses leaf condition subjectively and, in many cases, lacks the diagnostic training to distinguish between conditions with overlapping visual presentations.

By the time a laboratory test confirms the infection type, the opportunity for low-cost early intervention has frequently passed. We aim to compress that identification timeline to a matter of seconds.

Our classification engine is trained on the Plant Village dataset, a large annotated repository containing images of both diseased and healthy leaves drawn from several commercially important crop varieties. Before training begins, each image passes through a conditioning pipeline: spatial dimensions are standardized, pixel value distributions are normalized, and the training set is expanded through augmentation to expose the model to a wider range of visual conditions than the raw dataset alone would provide. These preparation steps are designed to strengthen the model's ability to generalize beyond the specific images it was trained on.

At the classification layer, we employ a Convolutional Neural Network — the architecture category that has consistently delivered the strongest results on leaf pathology tasks in the published literature. The network reads each submitted leaf image and constructs a hierarchical set of feature descriptions that encode color shifts, surface texture irregularities, and the spatial distribution of lesion regions.

These encoded representations feed into the classification head, which assigns the image to one of the defined disease categories. The trained model handles feature discovery autonomously, which removes the need for researchers to pre-specify diagnostic criteria and contributes directly to the accuracy of the output.Delivery to end users is handled through a React.js web interface backed by a Python API built with Flask or Fast API. A user submits a leaf photograph through the browser, whereupon the API routes the image through the same preprocessing pipeline used at training time and forwards it to the loaded model for inference. The response returned to the browser contains the predicted disease label and the network's confidence value, both displayed immediately on screen.

The architecture is intentionally lightweight — deployment on a modest cloud server is sufficient — and the system is designed to absorb future additions such as mobile client integration, live camera capture, and field-sensor connectivity as natural extensions of the current build.

IV. METHODOLOGY

Our development pipeline is structured across six sequential phases: data assembly, image conditioning, model construction, training, performance evaluation, and system integration.

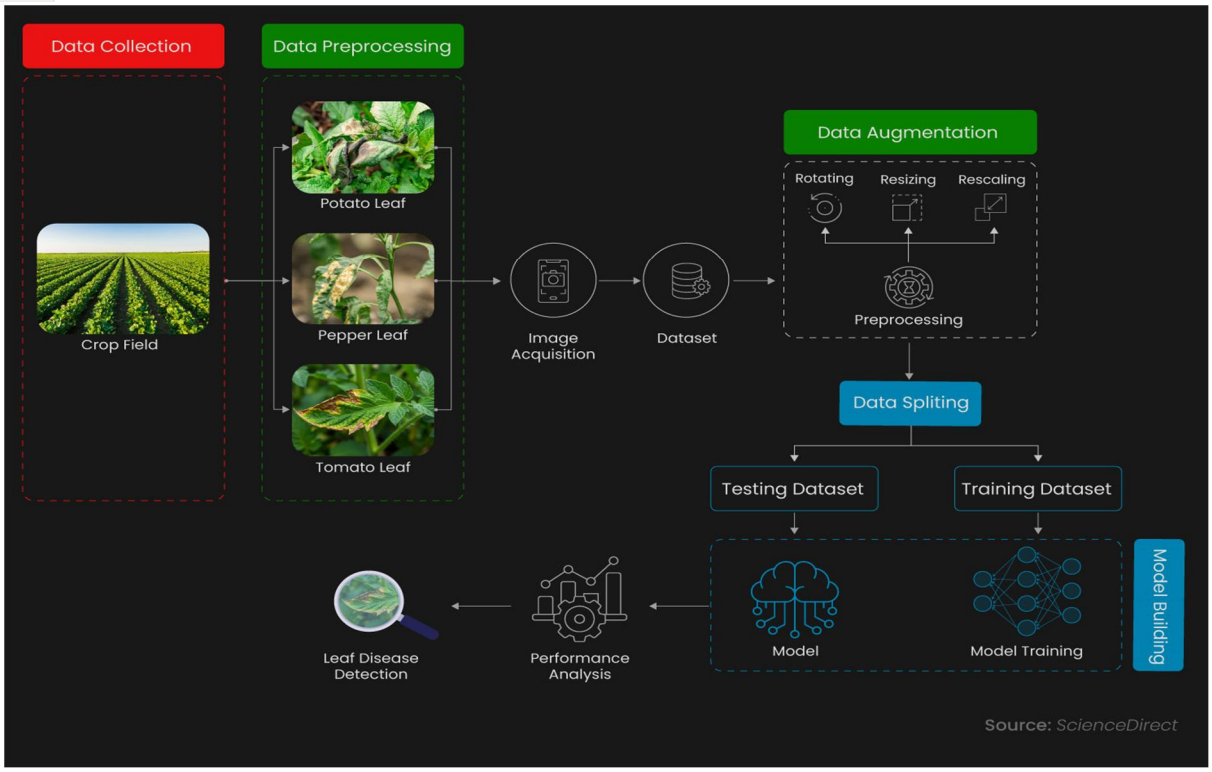
A. Data Collection

We draw our image data from the Plant Village benchmark, a publicly available annotated collection of more than 50,000 leaf photographs. The dataset represents a range of pathological conditions — covering bacterial, fungal, and viral disease classes — alongside healthy specimens, and spans commercially significant crop types including tomato, potato, corn, grape, and apple. Every sample in the dataset carries a dual label specifying both the host crop and the health condition depicted.

B. Data Preprocessing

Raw images enter a deterministic conditioning pipeline before they reach the network. We rescale all inputs to 224 × 224 pixels — the spatial format that the target CNN architectures expect — using area-averaging for down sampling and bilinear interpolation for up sampling. Pixel intensities are then linearly mapped to the [0, 1] interval; this normalization step removes scale differences between color channels and contributes to numerical stability during gradient descent.

The training partition receives an additional augmentation pass that generates transformed copies of each original image. The transform set includes horizontal and vertical mirroring, rotations within ±30 degrees, zoom operations between 0.8× and 1.2×, and stochastic brightness and contrast shifts. The intent is to populate the training distribution with the kinds of visual variation the system will encounter in field photographs, reducing the risk that the model memories the particular angles and lighting conditions of the training set rather than learning generalizable disease features.



C. Model Development

The classification backbone of our system is a Convolutional Neural Network comprising stacked convolutional blocks, spatial pooling stages, non-linear activation units, and fully connected output layers — each serving a defined computational purpose within the overall architecture.

The convolutional blocks slide learned filter banks across the input, producing feature maps that encode the presence and location of low-level visual primitives such as edges, color transitions, and textured regions. Pooling operations then reduce the spatial resolution of those maps, condensing the representation and controlling the growth of parameter count in deeper layers.

Classification is finalized in the fully connected head, which maps the flattened feature vector to a probability distribution over all target disease classes via a SoftMax output layer. Alongside the custom architecture, we evaluate transfer learning variants based on VGG16, ResNet50, and MobileNetV2 — networks pretrained on large general-purpose image corpora and fine-tuned on our plant disease data — to assess whether inherited representations accelerate convergence and raise ceiling accuracy.

D. Model Training

The dataset is divided into three non-overlapping partitions: 70 percent of images go to training, 15 percent to validation, and the final 15 percent are withheld for held-out testing. Stratified splitting ensures that class proportions are consistent across all three subsets.

We train using the Adam optimizer paired with categorical cross-entropy as the loss objective. The three principal hyperparameters — learning rate, batch size, and epoch count — are tuned through systematic search over the validation set; the configuration that maximizes validation accuracy without signs of overfitting is selected for final evaluation on the test partition.

E. System Integration

The trained model is serialized and loaded into a REST API endpoint implemented in Python using Flask or Fast API. The API accepts a leaf photograph submitted through the React.js frontend, runs it through the preprocessing pipeline, performs inference, and returns the top prediction and its associated confidence score. The complete user journey — photograph upload, server-side processing, model inference, and result display — executes as a single asynchronous request, keeping the interface responsive throughout.

V. APPLICATIONS

1) Agricultural Disease Diagnosis: Growers can photograph a leaf showing abnormal coloration or surface damage, submit the image through the browser interface, and obtain a confirmed disease name within seconds — enabling treatment decisions to be made on the same day the symptom is noticed rather than days later.
2) Precision Farming: When the specific disease is known rather than guessed, chemical inputs can be matched precisely to the confirmed pathogen. This selectivity cuts down on broad-spectrum pesticide application, reduces per-cycle input costs, and lowers the chemical burden on surrounding ecosystems.
3) Mobile and Web-Based Advisory Systems: The inference backend can be invoked from any client that supports HTTP requests, meaning the same model can power both the current browser interface and future mobile clients — providing off-grid farmers with diagnostic capability through whichever device they carry.
4) Research and Academic Use: The codebase and trained weights can serve as a reproducible baseline for researchers exploring incremental improvements in architecture, augmentation strategy, or domain adaptation for agricultural image classification tasks.
5) Government Agricultural Programs: The API layer is designed to accept integration requests, making it straightforward for state agricultural bodies to embed disease identification functionality into existing digital advisory portals used by extension workers and registered farmers.
6) Smart Farming Systems: Periodic or continuous leaf imagery captured by field-deployed cameras can be routed through the system's inference endpoint, enabling automated surveillance of crop health without manual scouting — a natural fit for IoT-connected precision agriculture installations.

VI. ADVANTAGES

1) High Accuracy: Extensive benchmarking in the published literature confirms that CNN architectures achieve classification accuracy on plant disease datasets that consistently exceeds what is attainable with hand-engineered feature extraction methods, making them a reliable diagnostic core.
2) Fast Detection: Model inference on a single leaf image completes in under a second on standard server hardware, compressing the identification timeline from the hours or days associated with laboratory referral to an interval that fits within a single field visit.
3) Cost-Effective Solution: Automated diagnosis removes the requirement for a trained agronomist to be present at the point of identification. This decoupling of diagnosis from specialist availability reduces the per-farm cost of disease monitoring across a cropping season.
4) Early Disease Detection: Because the system shortens identification time, it increases the probability that an infection is confirmed during its early, localized phase rather than after it has spread — the intervention window that carries the lowest treatment cost and the highest chance of full yield recovery.
5) Scalable System: New crop species and disease classes can be incorporated into the model by extending the training dataset and rerunning the fine-tuning procedure; no changes to the inference pipeline or web interface are required, which keeps the cost of expanding coverage low.
6) User-Friendly Interface: The React.js interface was designed around a single interaction: choose an image, submit it, read the result. No account registration, configuration, or technical vocabulary is required, which makes the tool accessible to users with minimal computing experience.
7) Automation of Disease Detection: Moving identification into software eliminates the inter-observer variability that arises whenever multiple human inspectors assess the same field. The model applies the same learned criteria uniformly across all submissions, producing consistent outputs regardless of who submits the image.
8) Supports Smart Agriculture: Disease identification is one component in the broader ecosystem of AI-assisted farming tools, and our system is designed to interoperate with adjacent components — yield monitoring, weather advisory, and soil sensors — as the precision agriculture technology stack continues to mature.

VII. CONCLUSION

We have presented the design and rationale for a CNN-driven plant disease recognition system that accepts leaf photographs as input and returns a disease classification with an associated confidence estimate. The system draws on the Plant Village image repository for training, couples a deep learning classification backbone to a React.js browser interface, and is architected for

deployment on commodity cloud infrastructure. Our analysis of published results confirms that deep learning classifiers outperform conventional image processing baselines on this task by a substantial margin, and that the performance gap widens as the number of disease classes and the diversity of input conditions increase.

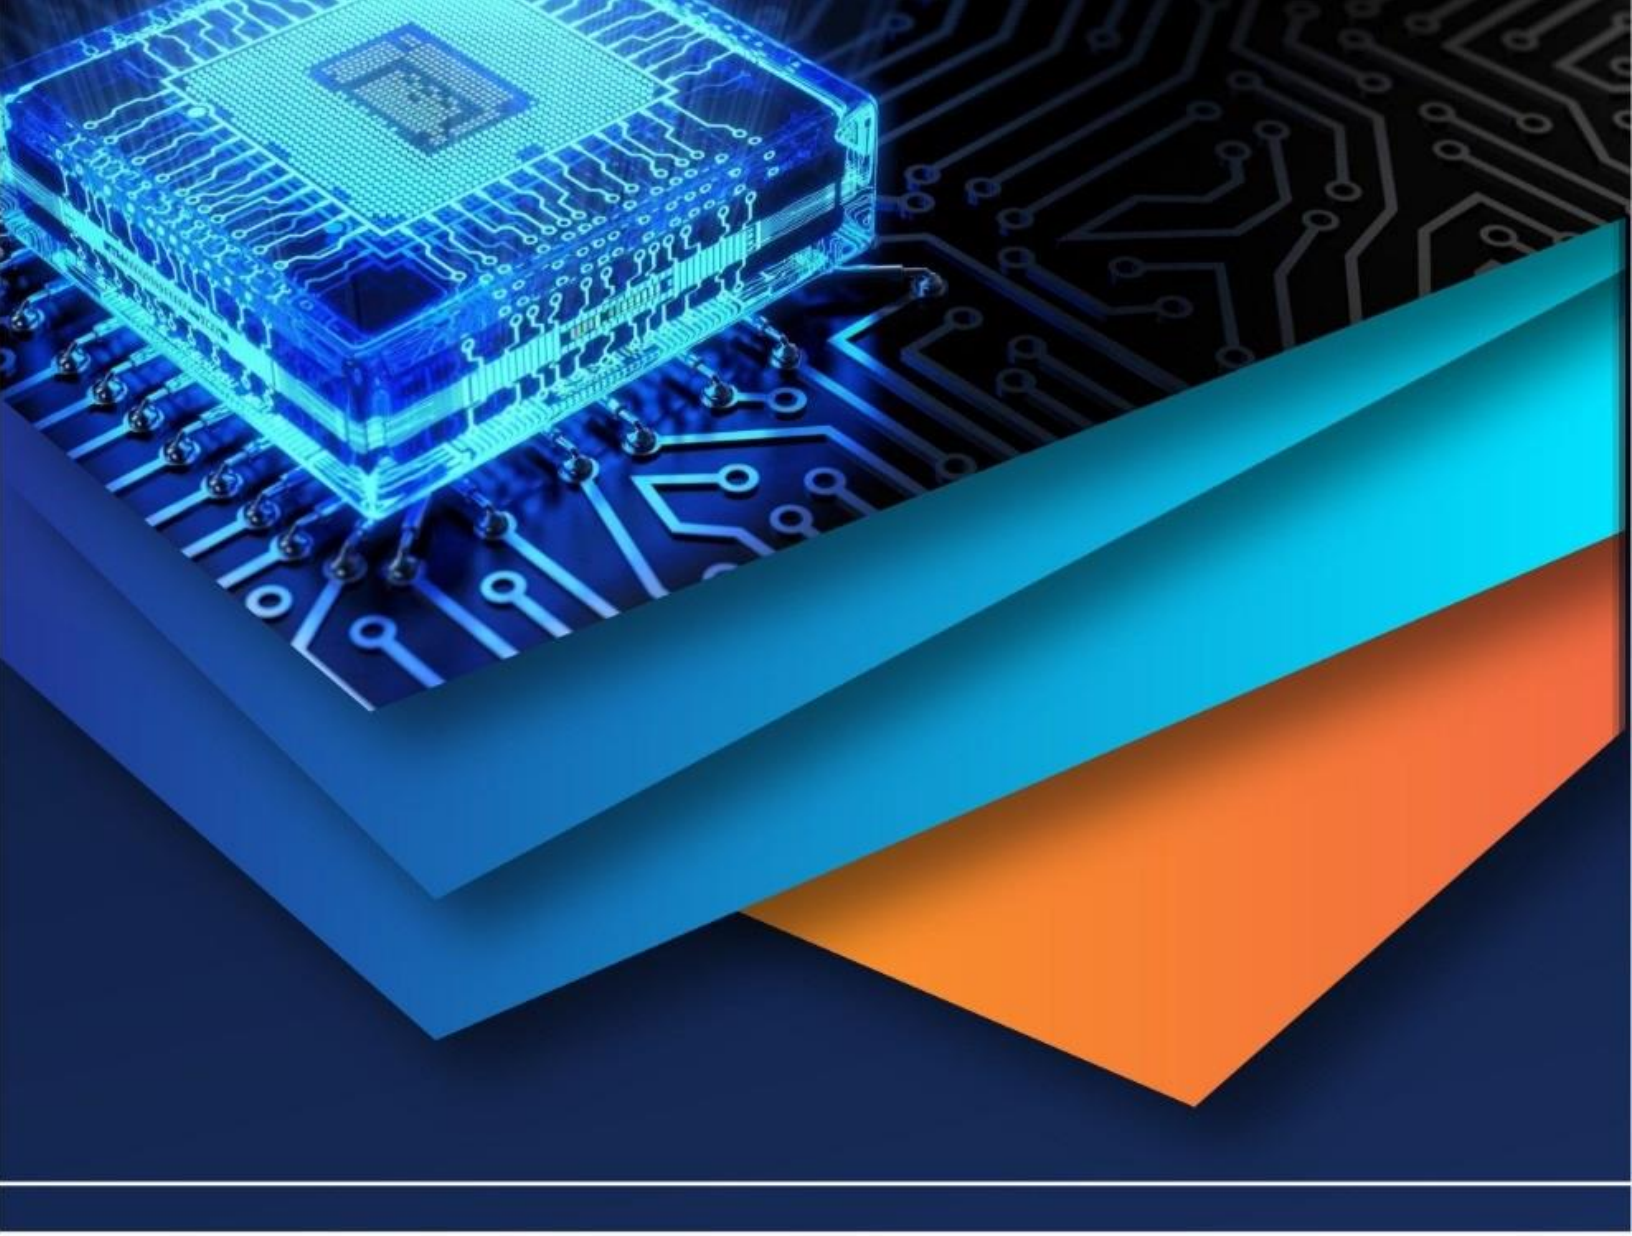
The principal motivation behind this work is the time cost of late disease identification. Every day between symptom emergence and confirmed diagnosis represents an opportunity for infection spread that an earlier response could have contained. By compressing that window to the duration of a photograph upload and an API call, the system gives farming households a concrete operational advantage: an intervention that would have been impossible on Tuesday due to uncertainty becomes actionable on Monday because the diagnosis is already confirmed. Future versions of the platform will extend this capability through direct mobile camera integration, IoT-linked field sensors, and a treatment recommendation module that converts a disease label into a specific, locally available management action.

Looking at the broader picture, this system represents one node in the expanding network of digital tools that are gradually reshaping how agricultural decisions get made. Connecting that node to adjacent tools — fertilization advisors, weather risk systems, market price feeds — will progressively raise the analytical ceiling available to individual farmers. Expanding crop and disease coverage in subsequent model iterations will widen the scope of what the system can address. And iterating on the interface design based on real user feedback from field deployments will ensure that accuracy gains at the model level translate into genuine improvements in how farmers manage their land.

REFERENCES

[1] Mohanty, S. P., Hughes, D. P., & Salathé, M. "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, 2016.
[2] Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D., "Deep neural networks based recognition of plant diseases by leaf image classification," Computational Intelligence and Neuroscience, 2016.
[3] Ferentinos, K. P., "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, 2018.
[4] Too, E. C., Yujian, L., Njuki, S., & Yingchun, L., "A comparative study of deep learning models for plant disease identification," Computers and Electronics in Agriculture, 2019.
[5] Barbedo, J. G. A., "Factors influencing the performance of deep learning models for plant disease recognition," Biosystems Engineering, 2020.
[6] Wang, G., Sun, Y., & Wang, J., "Automatic image-based plant disease severity estimation using deep learning," Journal of Agricultural Informatics, 2020.
[7] Barman, U., Choudhury, R. D., Sahu, D., & Barman, G. G., "CNN-based classification of crop disease images," Computers and Electronics in Agriculture, 2021.
[8] Saleem, M. H., Potgieter, J., & Arif, K. M., "Plant disease detection using deep learning techniques," Plants Journal, 2022.
[9] Sharma, A., Jain, A., Gupta, P., & Chowdary, V., "Machine learning applications in precision agriculture," IEEE Access, 2023.
[10] Wang, Y., Liu, J., Chen, H., & Li, X., "Vision transformer-based plant disease recognition," IEEE Transactions on Artificial Intelligence, 2025.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)