

Farmers Disease Diagnostic and Reporting Portal: An AI-Based Mobile Platform for Crop Disease Screening

Miss. Rutika Sudhir Pradhan¹, Prof. Shubham Lotwala²

Department Of Computer Application SSBT'S COET, Jalgaon, Maharashtra^{1,2}

Abstract: A farmer who notices a damaged leaf needs both an initial assessment and a way to obtain further advice. This paper describes a mobile portal that connects these tasks through AI-assisted crop screening and structured case reporting. Farmers would record photographs, crop details, and the date of observation. A compact model would suggest possible disease categories for supported crops, while unclear images and uncertain cases would be referred for review. Each case would retain the original evidence, the model version, and any later assessment by a crop specialist. Offline storage would allow observations to be captured without a connection and submitted once connectivity returns. The study presents a system design rather than an implemented service. It defines requirements, processing steps, diagrams, and a plan for evaluating classification quality, synchronization, expert review, and farmer usability. Its contribution is a traceable route from the first observation to a reviewed response. Field trials are still needed to determine whether the design improves reporting or access to useful advice. A photograph-based suggestion would remain preliminary and would not replace field examination or laboratory confirmation when these are required.

I. INTRODUCTION

1.1 Background

A farmer who notices spots or curling on a leaf needs to know what evidence to collect and whom to consult. Appearance alone may not resolve the problem: pathogens, pests, nutrient imbalance, and environmental conditions can produce overlapping symptoms. Crop history, field examination, or laboratory analysis may be needed [7]. The portal described here would keep these observations together so that the farmer does not have to start the explanation again at each stage of consultation.

1.2 Importance

Earlier studies show that image recognition can assist crop screening, but performance measured on a prepared dataset may change when photographs are taken in a field [1], [2]. This distinction shapes the project. A quick answer is useful only if the farmer understands its limits and can obtain further help. Recording the evidence and the response is therefore as important to this design as displaying a predicted label.

1.3 The Farmers Disease Diagnostic and Reporting Portal

The portal concerns crop disease. Its basic unit is a case: photographs linked to crop details, an observation date, and subsequent decisions. Farmers submit cases and view responses; specialists review evidence and enter advice; administrators manage assignment and access. The interface identifies who produced each result and whether more information is needed.

1.4 Types of Diseases and Similar Crop Problems

The initial classifier should cover a declared set of crop-disease combinations supported by training and validation data. The following categories describe reporting needs, not a claim that every condition can be identified from a leaf image. Assessment should consider living and non-living causes [7].

1.5 Research Problems

The design must handle unclear images, unfamiliar conditions, failed uploads, and corrections to model predictions. Field-image research makes the first two issues particularly important [1]-[3]. Reporting introduces further requirements: retries should not create duplicate cases, and expert corrections should preserve the original prediction so that the

reasoning behind advice can be reviewed.

Category	Illustrative presentation	Portal requirement
Fungal and oomycete diseases	Leaf spots, powdery growth, rust-like lesions, or blight symptoms	Record the crop and affected organ; identify supported classes.
Bacterial diseases	Spots, blight, or wilt that can resemble other disorders	Flag the uncertainty and seek additional evidence.
Viral diseases	Mosaic patterns, curling, or stunted growth	Request plant-level images and route to expert assessment.
Other crop problems	Pest damage, nutrient imbalance, chemical injury, or water stress	Allow for non-disease and unknown outcomes.

1.6 Scope

The first version would cover selected crops, image screening, case tracking, and specialist feedback. Crop selection depends on available labelled data and reviewers. Marathi, Hindi, and English are candidate languages for a Maharashtra pilot. Laboratory confirmation, universal crop coverage, and automatic pesticide prescribing are excluded. Exact location sharing remains optional.

1.7 Objectives

The project aims to connect a crop observation to a reviewed response. It should identify uncertain inputs, retain reports during connection failures, prevent duplicates, and protect personal information. Tests will examine whether farmers can complete reports and whether reviewers receive sufficient evidence. Classification and reporting quality will be assessed separately.

II. LITERATURE REVIEW AND RELATED WORK

2.1 Image-Based Plant Disease Classification

Mohanty, Hughes, and Salathe [1] showed that deep learning could classify crop and disease combinations from PlantVillage leaf images. Performance was substantially weaker when the images came from different conditions. The study therefore supports image-based screening while showing the limits of evaluation within a prepared dataset. For the present project, an independent set of field images should be reserved for testing after data collection. No such test has yet been carried out.

2.2 Mobile Diagnosis in Agricultural Fields

Ramcharan et al. [2] took disease recognition onto a mobile device and examined cassava symptoms in field images and video. Their evaluation considered symptom severity and differences between testing conditions. This gives the present design a practical reference: the camera workflow and the visible state of the plant belong in the evaluation, not just the trained network. Their work also shows that offline crop screening already has research precedent.

2.3 Diverse Images and Dataset Limitations

PlantDoc, introduced by Singh et al. [3], provides internet-sourced plant images with conditions that differ from controlled leaf photography. It is relevant when considering cluttered backgrounds and variation in appearance. Nevertheless, using it would not demonstrate that a model works for local varieties or local phones. Before inclusion, its images would need checks for label suitability, duplicate content, provenance, and permitted use.

2.4 Models Suitable for Mobile Devices

MobileNetV3 was developed by Howard et al. [4] for efficient vision processing on mobile hardware. It is a reasonable candidate to test, rather than a predetermined winner. A crop-specific experiment must compare its recognition performance with its storage, memory, and timing requirements. The choice of model should follow those measurements on the intended phones.

2.5 Confidence and Visual Explanation

Two supporting methods are relevant to the review process. Guo et al. [5] investigated why a network's confidence may not match how often it is correct and examined calibration methods. Selvaraju et al. [6] developed Grad-CAM to show regions influencing a prediction. Calibration could inform referral thresholds, while visual explanations could help reviewers inspect suspicious predictions. Neither method establishes the biological cause of a symptom.

2.6 Research Gap and Proposed Contribution

Existing studies provide evidence for crop image classification, mobile inference, confidence calibration, and visual explanation [1], [2], [4]-[6]. This paper brings those ideas into a reporting workflow. It asks how an observation can retain its context as it moves from an automated suggestion to expert assessment and back to the farmer. The contribution lies in this system design; it does not establish a new classification algorithm or claim that similar agricultural services are absent.

The research questions are: How well does the selected model screen local field images? How does rejecting uncertain cases affect errors and expert workload? Can farmers complete reports despite connectivity interruptions? Does the workflow improve completeness and time to reviewed advice relative to the pilot service's existing process? These questions require empirical evaluation after implementation.

III. METHODOLOGY

This is a design study. Its outputs are requirements, workflows, diagrams, and an evaluation plan. They have been developed from the cited literature and an analysis of the reporting task. Interviews with farmers, image collection, model training, and pilot testing remain work for the implementation stage. Consequently, the paper assesses the logic of the design without claiming that it has already been effective in practice.

3.1 Feasibility Study

Technical feasibility: A mobile client, compact model, backend, and structured database can form the prototype. Key uncertainties are representative data, unknown-condition handling, and performance on affordable phones. **Operational feasibility:** Farmers need a short capture process and visible status; experts need adequate evidence and manageable queues. Agricultural partnerships are needed before promising response times.

Economic feasibility: Open-source components can reduce licensing expenses, but labelling, hosting, storage, maintenance, and expert review remain costs. A pilot should estimate cost per submitted and reviewed case. **Ethical feasibility:** Participation should be voluntary, data use understandable, and guidance reviewed for local conditions. These considerations define a feasibility assessment to be completed during implementation.

3.1 Risk Analysis

Risk	Consequence	Proposed control
Wrong or missed label	Inappropriate action or delayed assessment	Limit output to supported classes, apply uncertainty checks, and refer to experts.
Dataset bias or leakage	Inflated accuracy and uneven real-world performance	Split data by farm or session; test location and season independently.
Synchronization failure	Lost cases or duplicate submissions	Use a local queue, stable identifiers, acknowledgements, and retries.
Personal data exposure	Disclosure of identity or farm location	Collect minimally, restrict access, and aggregate summaries.
Expert backlog	Delayed or unreviewed reports	Show queue status and apply capacity-based escalation.

3.3 Requirement Collection and Identification

Requirements should be refined by asking farmers to complete a short reporting task and asking specialists to review the resulting cases. This would reveal where instructions are unclear or essential evidence is missing. Administrators should also examine assignment and access controls. Each requirement needs a concrete acceptance check: for example, retrying an interrupted upload should preserve one case rather than create a second one.

3.4 Software and Hardware Requirements Specification

The working software design consists of an Android client, a Python API, a relational database for case records, separate protected image storage, and an on-device inference runtime. A training workstation may use a GPU, but the farmer's phone should be evaluated on the hardware actually available. Required memory, supported operating-system versions, and server capacity will be determined by profiling. They are not established minimum specifications at this stage.

3.5 Assumptions

The service depends on more than a trained model. A participating organization must provide reviewers, define supported crops, and explain how reports will be used. Farmers need occasional connectivity to receive reviewed advice. Local screening is possible only after the model is installed. The design also assumes that some cases will remain unresolved: a clear photograph may still lack the evidence needed to identify a cause.

3.6 Proposed Data Processing

Capture and validation: The farmer selects the crop and records when symptoms were noticed. Camera guidance asks for a close-up of the affected part and, where useful, a whole-plant view. The app checks that the image can be opened and is usable before analysis. It can request another photograph without discarding the case. Symptom duration and crop stage provide context; a locality may be entered manually when precise location is unnecessary.

Preprocessing and inference: Images are resized and normalized consistently with model development. The model runs for supported inputs and records the candidate label, score, timestamp, and model version. Poor quality, an unsupported crop, or uncertain evidence results in a request for further assessment. A high confidence score alone cannot establish that an unfamiliar image is within the model's training distribution.

Reporting and synchronization: The phone assigns a case identifier before upload. If the connection fails, the same identifier is used on retry, allowing the server to recognize an existing submission. The screen shows a report as submitted only after acknowledgement. Observation time and upload time remain separate. When an expert disagrees with the model, the reviewed assessment is added to the history rather than replacing the original prediction.

Review and aggregation: Experts can request evidence, record an assessment, or recommend field or laboratory examination. Advice records include the reviewer, source, and revision date. Deduplicated reports are summarized by crop, time, and permitted locality, with reviewed and unreviewed counts separated. A cluster warrants investigation rather than an automatic outbreak declaration.

3.7 Possible Machine Learning Approach

Transfer learning with MobileNetV3 is a candidate approach [4]. Public datasets may support development, supplemented by consented local images labelled by crop experts. Partitions should separate farms, plants, and capture sessions, with near-duplicate removal before splitting. Augmentation belongs only in training. Validation data support model selection, calibration, and rejection thresholds; an untouched field test set supports final assessment.

A simple baseline and the compact model should use identical splits. Experiments should report macro F1, per-class recall and precision, confusion matrices, and calibration alongside memory, size, and latency. Quantized and full-precision versions should be compared before deployment. Thresholds should reflect validation evidence and an agreed error tolerance; no confidence cutoff is assumed universally safe [5].

Unknown-condition testing should include unsupported crops, non-disease damage, unrelated objects, and difficult field images. Expert corrections require checking before becoming training labels. Model updates should be versioned and evaluated against the previous release, with rollback available if field performance declines.

IV. DIAGRAM DESIGN

4.1 Data Flow Diagram




Figure 1. Proposed data flow for mobile crop screening and reporting.

4.2 State Chart




Figure 2. Report states from capture to review and closure.

4.3 Sequence Diagram




Figure 3. Interaction between the farmer, portal, model, server, and expert.

V. EXPECTED RESULTS AND DISCUSSION

5.1 Expected Functional Outcomes

The expected output is a case that a farmer can submit and a specialist can inspect without losing the original context. The farmer should be able to answer three questions from the screen: Is my report still on the phone or has it reached the server? Is this result from the model or a reviewer? What information or action is needed next? These are intended interface outcomes and have not yet been tested with users.

A useful case record should make it possible to request another image or clarify a symptom without starting a new report. This could improve continuity between the farmer and the reviewer, but it remains an intended benefit. A pilot should compare completeness and review time with the participating service's existing process. The number of reports alone cannot establish diagnostic quality, user satisfaction, or an improvement in crop yield.

5.2 Evaluation Framework

Evaluation should include offline model testing, end-to-end device and synchronization tests, and a supervised pilot. Specialists should establish reference labels independently of the model. Disagreements need adjudication, unresolved cases should be reported separately, and laboratory-supported labels should be identified distinctly. Recruitment, crop coverage, sample-size rationale, devices, and test conditions must be documented.

Area	Measures	Purpose
Classification	Macro F1, per-class precision and recall, confusion matrix	Reveal minority-class errors and missed cases.
Uncertainty	Coverage, error among accepted cases, unknown-case rejection	Measure automated screening against expert workload.
Confidence	Reliability plot and calibration error	Compare scores against observed correctness.
Mobile operation	Median and 95th-percentile latency, memory, model size	Assess performance on the intended devices.
Reporting	Upload success rate, duplicate rate, interruption recovery	Verify that retries preserve cases reliably.
Service and usability	Completion rate, task time, review turnaround, feedback	Assess ease of access and review capacity.

Precision measures how often predictions for a class are correct; recall measures how many reference cases of that class are found. F1 combines precision and recall, and macro F1 weights each class equally. Coverage is the fraction of eligible cases receiving an automated suggestion. Denominators and exclusions must be stated. Confidence intervals should account for related photographs from the same plant or farm.

5.3 Interpretation and Comparison

Model variants should be compared on identical held-out cases. The pilot should compare report completeness and turnaround with the existing process under comparable crops, periods, and review capacity. Recruitment bias and reviewer availability must be considered when interpreting differences. Seasonal and location-specific results should be reported separately so aggregate scores do not hide weak performance.

Referral policy creates a practical trade-off. Asking experts to review more cases may reduce incorrect automatic suggestions, but it can also lengthen the queue. Keeping larger images may help inspection while slowing uploads. Both decisions should be tested against the resources of the pilot service. The current design provides no measured accuracy, time saving, or reduction in crop loss.

VI. LIMITATIONS

The largest limitation is the absence of implementation evidence. The paper cannot establish how well a trained model or the reporting service would perform. Early symptoms, mixed infections, occluded leaves, and damage to roots or internal tissues would remain difficult for image-based screening. Even a plausible label may not identify the actual cause. These limitations make specialist referral part of the service rather than an optional extra [2], [7].

Public datasets may not reflect local crop varieties, seasons, phones, or farming practices. Farmers who can access the portal may also differ from those unable to participate. Report counts consequently reflect both crop problems and patterns of service use; their relationship to actual disease prevalence cannot be assumed. Estimating prevalence would require an appropriate sampling design and denominator. Offline capture helps preserve observations, but expert communication still needs connectivity.

VII. FUTURE SCOPE

Development should start with a small, declared set of crops and an advisory collection approved by reviewers. A supervised pilot can then identify failures in capture, classification, upload, and follow-up. Crop coverage should expand only when suitable data and reviewers are available. Audio instructions and assisted reporting through extension workers are possible additions, but their usefulness should be checked with the people expected to use them.

Later work could assess model-assisted requests for follow-up images, expert prioritization of uncertain cases, and monitoring of seasonal performance. Visual explanations could support expert inspection [6]. Weather context and regional trends may inform further research, but forecasting would need independent validation. Privacy-preserving learning could also be studied; it would not automatically resolve consent, security, or biased data.

VIII. CONCLUSION

A useful crop-health portal needs to do more than attach a disease name to a photograph. It must preserve the farmer's observation, explain uncertainty, deliver the case to a reviewer, and return a response that can be understood and followed up. This paper has described how those steps could be connected through a mobile interface, local storage, and a shared case record.

Existing research provides a basis for testing image-assisted screening, while also showing why performance outside prepared datasets cannot be assumed. The next stage is to implement the bounded workflow and test it with representative images and users. The value of the project will depend on the quality of reviewed cases and the usefulness of the response, as well as the classifier's scores.

REFERENCES

- [1] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, Vol. 7, Article 1419, 2016.
- [2] A. Ramcharan, P. McCloskey, K. Baranowski, N. Mbilinyi, L. Mrisho, M. Ndalawha, J. Legg, and D.P. Hughes, "A Mobile-Based Deep Learning Model for Cassava Disease Diagnosis," *Frontiers in Plant Science*, Vol. 10, Article 272, 2019.
- [3] D. Singh, N. Jain, P. Jain, P. Kayal, S. Kumawat, and N. Batra, "PlantDoc: A Dataset for Visual PlantDisease Detection," *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*, 2020.
- [4] A. Howard et al., "Searching for MobileNetV3," *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1314–1324, 2019.
- [5] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On Calibration of Modern Neural Networks," *Proceedings of Machine Learning Research*, Vol. 70, pp. 1321–1330, 2017. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," *Proceedings of the IEEE International Conference on Computer Vision*, pp. 618–626, 2017.
- [6] University of California Statewide Integrated Pest Management Program, "Diagnosing Plant Problems," *UC IPM Pest Management Guidelines: Floriculture and Ornamental Nurseries*, UC ANR Publication No. 3392, 2026.