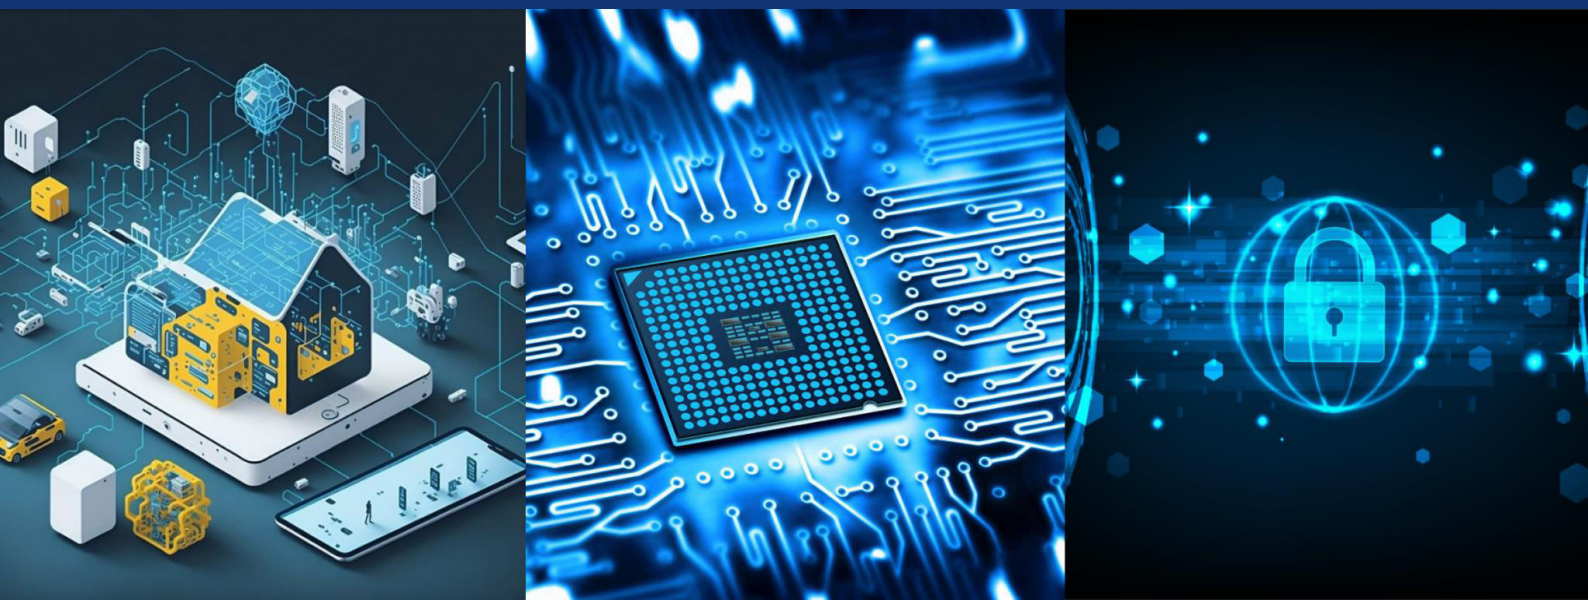
IJIRCCE

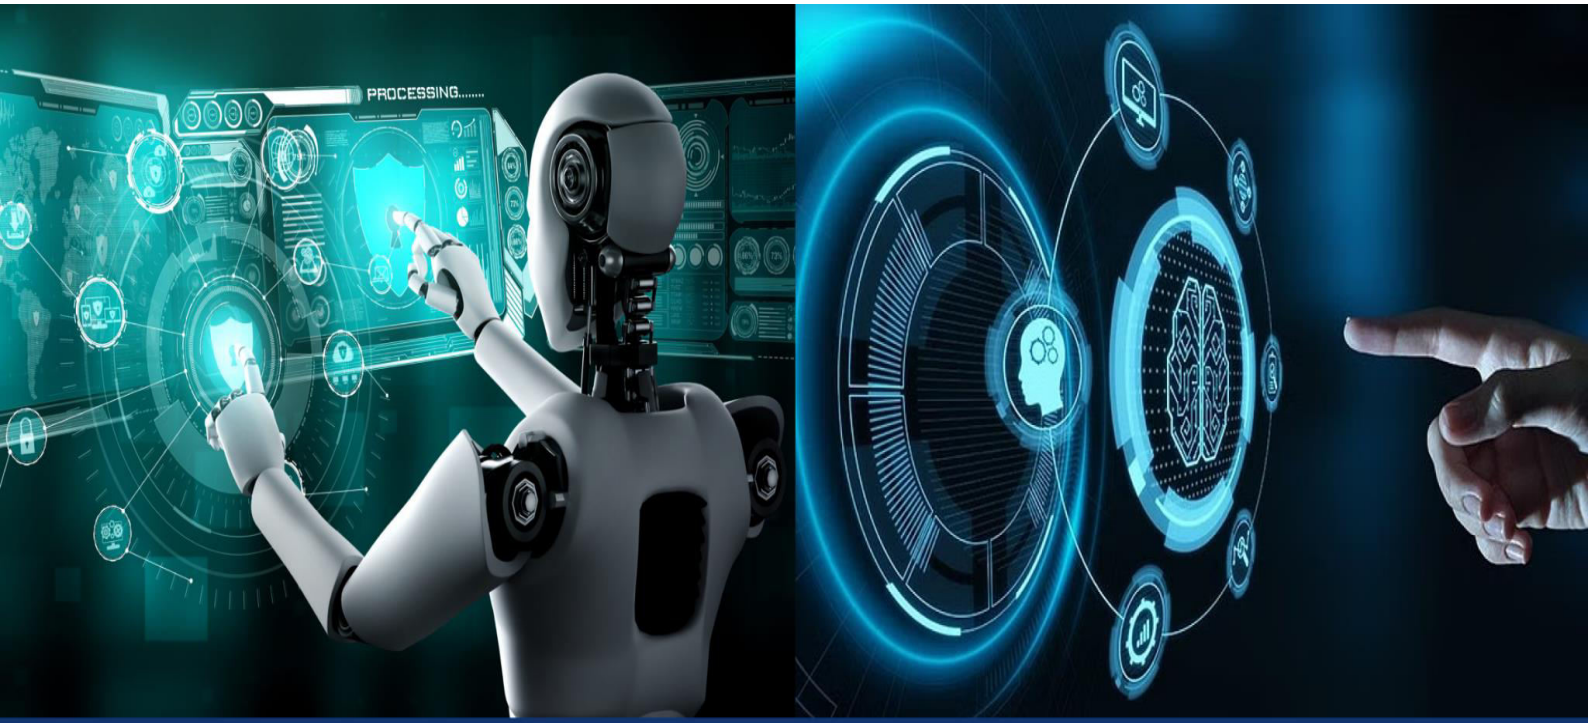
ISSN(O): 2320-9801

ISSN(P): 2320-9798



International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 9.927

Volume 14, Issue 11, November 2026

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

PranaDhanya: An AI-Powered Portal for Crop Disease Diagnosis, Livestock Health Screening and Farm Decision Support

Dinesh Kumar M, Dhanush M, Shashank M Aradhya, Adithya S, Kushal R

Professor, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India

Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India

Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India

Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India

Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India

ABSTRACT: Small and marginal farmers usually have to decide about crop disease, livestock illness, fertilizer and selling time without quick access to an expert. This paper presents PranaDhanya, a web portal that puts six machine-learning models behind one role-based interface for farmers, veterinary officers and administrators. Two ResNet9 convolutional networks screen leaf images (38 plant classes) and animal images (11 cattle, dog and goat classes); an XGBoost classifier recommends crops from soil and weather values; Random Forest models estimate crop yield, wholesale price and a suitable fertilizer. The models are served by a FastAPI microservice and consumed by a Spring Boot 3 backend with JWT security and a PostgreSQL database. On our held-out data the plant and animal image models reach 96.00% test accuracy, crop recommendation reaches 98.64–99.86% accuracy (F1 = 0.99), fertilizer prediction reaches 98.99–100% accuracy (F1 = 0.99), and the yield and price regressors reach R^2 of 0.74–0.85 (MAE 1.12 t/ha) and 0.77–0.85 (MAE Rs. 71 per quintal) respectively. We also discuss where the system still falls short, especially for field photographs taken against cluttered backgrounds.

KEYWORDS: Precision agriculture, plant disease detection, livestock diagnosis, ResNet9, XGBoost, Random Forest, crop recommendation, price forecasting, FastAPI, Spring Boot

I. INTRODUCTION

Farming in India still runs on a mix of experience, hearsay and whatever advice happens to be nearby. A farmer who spots brown patches on a maize leaf, or a lesion on a calf's foot, often has to wait days for a visit from an extension worker or a veterinarian, and by then the problem may have spread. The same is true of quieter decisions: which crop suits this season's soil, how much fertilizer is really needed, and whether to sell this week or wait for the mandi price to improve.

Machine learning can shorten that waiting time, and a good deal of work has shown that deep networks can recognise crop diseases from photographs [2, 4, 5]. What is much rarer is a single, usable system that joins these models together with the everyday workflow of a farm: logging a diagnosis, booking a vet, reading a treatment note in one's own language. Most published systems are either a notebook with a model or a mobile app wrapped around one classifier.

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

In this paper we describe PranaDhanya, a full-stack portal that tries to close that gap. Our contributions are:

- an integrated platform in which six trained models (two image classifiers and four tabular models) are exposed through one REST service and one web interface;
- two ResNet9 image classifiers, one for 38 plant-leaf classes and one for 11 livestock classes covering cattle, dogs and goats;
- a role-based workflow for farmers, veterinary officers and administrators, including consultation booking, record keeping, a multilingual interface and a voice assistant; and
- an honest evaluation that reports both the headline accuracies and the failure cases we saw when the system was tried on real, unsegmented photographs.

The rest of the paper is organised as follows. Section II reviews related work, Section III presents the proposed methodology (architecture, data and models), Section IV reports the experimental results together with limitations, and Section V concludes.

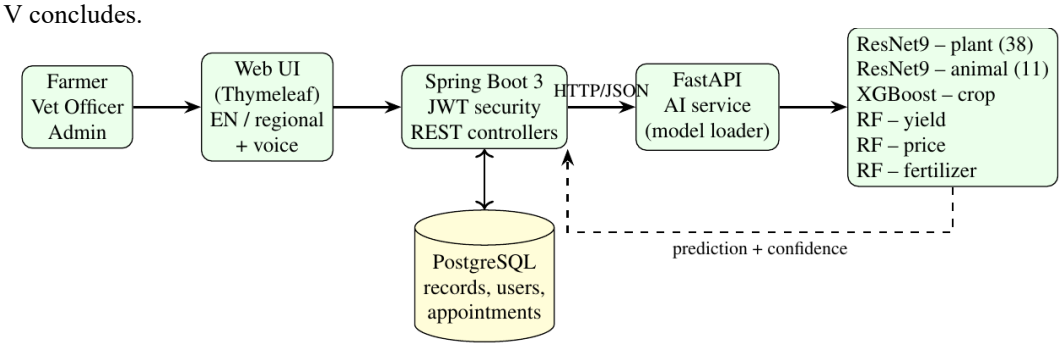


Fig. 1: Overall architecture of PranaDhanya. The Java backend never loads a model itself; it forwards images or feature values to the FastAPI service and stores the returned prediction.

II. RELATED WORK

Image-based plant disease detection became practical with the PlantVillage collection, on which Mohanty et al. trained deep networks that reached very high accuracy under laboratory conditions [2, 3]. Ferentinos later compared several architectures on a larger set of leaf images, and Kamilaris and Prenafeta-Boldú surveyed the wider use of deep learning in agriculture [4, 5]. A recurring observation in these studies is that accuracy drops noticeably when test images are taken in the field rather than against a plain background, which is a point we return to in Section VI.

Residual networks made it possible to train deeper models without the degradation problem [1], and compact variants such as ResNet9 offer a good compromise between accuracy and training time on modest hardware. For tabular agricultural data, gradient-boosted trees and Random Forests remain strong and easy-to-deploy baselines for crop recommendation, yield estimation and price forecasting [6, 7].

What sets PranaDhanya apart from these studies is less the individual models than the way they are packaged: crop, soil, market and livestock services live in one portal, they share one authentication layer, and a diagnosis can lead directly to a booked consultation with a veterinary officer.

III. PROPOSED METHODOLOGY AND DISCUSSION

A. System Architecture

PranaDhanya follows a three-tier design (Fig. 1). The presentation tier is a server-rendered web interface built with Thymeleaf. The application tier is a Spring Boot 3 backend that handles authentication with JSON Web Tokens (JWT), role-based access for the Farmer, Vet Officer and Admin roles, appointment booking and record keeping, with PostgreSQL as the persistent store. All machine-learning inference is isolated in a separate Python FastAPI microservice that loads the trained models once at start-up and answers REST calls from the backend. Keeping the models in their own service means they can be retrained and redeployed without touching the Java code, and the whole stack can be started with a single Docker Compose command.

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

1) User Roles and Modules

Once logged in, a farmer sees a console with the modules shown in Fig. 2: Plant Diagnosis, Livestock Diagnosis, Crop Analytics (crop recommendation and fertilizer advice), Price Forecast, Book Consultation, My Records, Govt Schemes and Profile Settings. Veterinary officers review the cases booked with them and can export clinical summaries, while administrators manage users and monitor the platform. The interface can be switched between languages, has a high-contrast toggle, and includes a floating microphone button for voice input, which matters for users who are more comfortable speaking than typing.

B. Datasets

Plant disease. The leaf classifier is trained on a PlantVillage-style collection that covers 14 crop species and 38 classes, ranging from Apple scab and Tomato late blight to healthy leaves of each species [3]. The images are background-removed leaves on a black canvas.

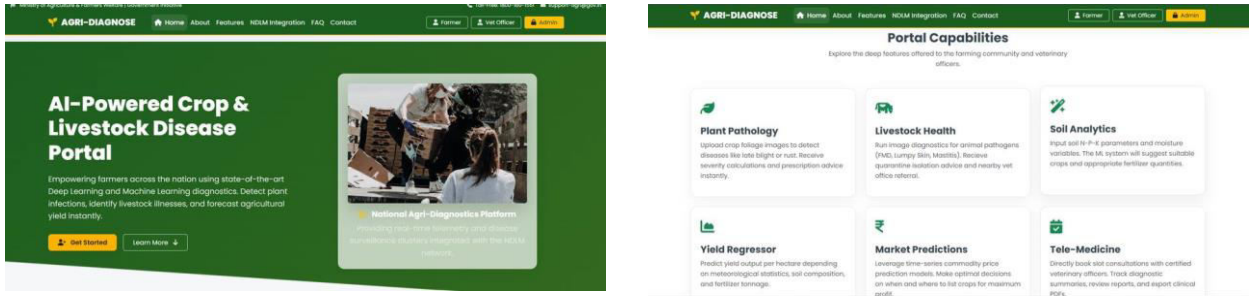


Fig. 2: The PranaDhanya landing page (left) and the portal capabilities section (right), showing the six functional areas offered to farmers and veterinary officers.

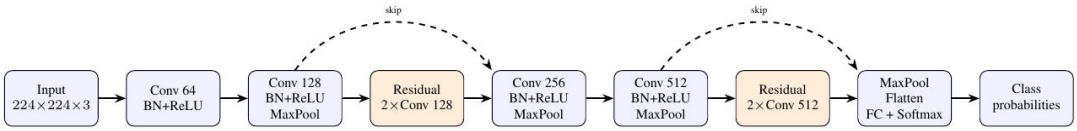


Fig. 3: ResNet9 used for plant and livestock classification. Orange blocks are residual blocks whose input is added back to their output through the dashed skip connections.

Livestock disease. The animal classifier covers 11 classes: cattle (foot-and-mouth disease, lumpy skin disease, healthy), dogs (demodicosis, dermatitis, fungal infections, ringworm, hypersensitivity, healthy) and goats (healthy, unhealthy). We assembled these from publicly available image collections; the dog skin set provided the bulk of the dermatological classes.

Tabular data. Crop recommendation uses the standard 2 200-sample table with 22 crop labels and seven inputs (N, P, K, temperature, humidity, pH, rainfall). Yield prediction uses area, rainfall, fertilizer, soil type, crop, temperature and humidity; price prediction uses crop, state, market and month; fertilizer prediction uses N, P, K, temperature, humidity, moisture, soil type and crop type.

C. Pre-processing

Both image models use the same input pipeline: images are converted to RGB, resized to 224 × 224 with bilinear interpolation, converted to tensors and normalised with the ImageNet statistics (mean = [0.485, 0.456, 0.406], standard deviation = [0.229, 0.224, 0.225]). Keeping the exact same transform at training and inference time turned out to matter in practice; a mismatch between the two is one of the easiest ways to lose several points of accuracy without any obvious error message. For the tabular models, categorical fields such as soil type, crop and state are label-encoded, and the feature order is saved with the model so that the API can rebuild the same vector at prediction time.

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

D. ResNet9 Image Classifier

For both image tasks we use a ResNet9 network: a convolutional stem, two residual blocks, three max-pooling stages and a final fully connected layer (Fig. 3) [1]. Each convolution is followed by batch normalisation and ReLU. The skip connections let the network learn corrections on top of the previous features, which helps training stay stable even with a small number of epochs. The same architecture is trained twice, once with a 38-way and once with an 11-way output layer.

Training used the Adam optimiser [8] with a learning rate of 10^{-3}, cross-entropy loss, a batch size of 32 and 10 epochs, implemented in PyTorch [9]. Table 1 summarises the configuration.

E. Tabular Models

Crop recommendation uses an XGBoost classifier with 100 trees and a multi-class soft-probability objective (multi:softprob, log-loss), which returns a ranked list of crops with probabilities [6]. **Yield** and **price prediction** use Random Forest regressors with 100 trees and a fixed random seed, built with scikit-learn [7, 10]. **Fertilizer prediction** also relies on a Random Forest model trained on soil nutrient, moisture and crop information. The portal then maps the predicted fertilizer to a dosage per acre and an organic alternative through a lookup table, so the farmer gets an actionable note rather than a bare label.

Table 1: Training configuration of the image models

Setting	Plant model	Animal model
Architecture	ResNet9	ResNet9
Output classes	38	11
Input size	224 × 224 RGB	224 × 224 RGB
Normalisation	ImageNet	ImageNet
Optimiser / LR	Adam / 10^{-3}	Adam / 10^{-3}
Loss	Cross-entropy	Cross-entropy
Batch size / epochs	32 / 10	32 / 10
Framework	PyTorch	PyTorch

Table 2: Performance of the six models in PranaDhanya

Task	Model	Result
Crop recommendation	XGBoost classifier	Acc. 98.64–99.86%, F1 = 0.99
Plant disease	ResNet9 CNN	Test acc. 96.00% (val. 95.88%)
Yield prediction	Random Forest regressor	R^2 = 0.74–0.85, MAE 1.12 t
Price prediction	Random Forest regressor	R^2 = 0.77–0.85, MAE Rs. 71/quintal
Animal disease	ResNet9 CNN	Test acc. 96.00% (val. 95.88%)
Fertilizer prediction	Random Forest	Acc. 98.99–100%, F1 = 0.99

F. Inference and Advisory Workflow

When a farmer uploads a photograph, the backend forwards it to the FastAPI service, which applies the pre-processing above, runs the ResNet9 model and returns the top class with its softmax confidence. The backend then attaches a knowledge-base entry for that class: description, symptoms, severity, organic and chemical treatment, and, for livestock, quarantine guidance (Fig. 4). The result is stored in the user's record history so that it can be shared with a veterinary officer later.

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

G. Evaluation Metrics

Classification models are evaluated with accuracy and the macro F1-score,

$$F1 = 2PR / (P + R), \quad (1)$$

where P and R are precision and recall. Regression models are evaluated with the coefficient of determination and the mean absolute error,

$$R^2 = 1 - \Sigma_i (y_i - \hat{y}_i)^2 / \Sigma_i (y_i - \bar{y})^2, \quad MAE = (1/n) \Sigma_{i=1}^{n} |y_i - \hat{y}_i|. \quad (2)$$

IV. EXPERIMENTAL RESULTS

A. Overall Performance

Table 2 lists the results of all six models. Where a range is given, the lower and upper values come from different evaluation runs of the same model.

The classification tasks that rely on structured soil and weather values are the easiest: both the crop and fertilizer models sit close to 99%, which is consistent with how cleanly separable these standard tables are. The regression tasks are harder and more honest about the uncertainty in farming. A yield R^2 between 0.74 and 0.85 with a mean error of about 1.1 tonnes means the model captures the main drivers (area, rainfall, fertilizer, crop) but cannot see local pests, irrigation or management practice. The price model's mean error of about Rs. 71 per quintal is small compared with typical mandi swings, which makes it useful for choosing a rough selling window rather than for predicting the exact day-to-day rate.

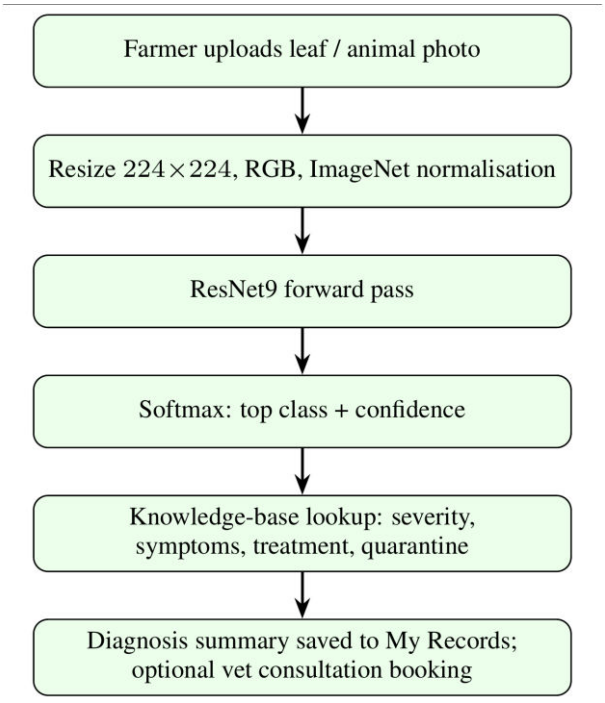


Fig. 4: Image diagnosis workflow from upload to stored report.

B. Training Behaviour of the Image Models

Fig. 5 shows the validation accuracy and the training/validation loss over the 10 epochs. Validation accuracy starts at roughly 69% after the first epoch, jumps to about 87% in the second, then dips around the fourth epoch before climbing steadily. The validation loss follows the same shape: it peaks at epoch four, which we attribute to the relatively high learning rate in the early epochs, and then falls quickly. From epoch seven onward, the training and validation losses almost coincide and both settle at a very low value, which suggests the network is not simply memorising the training set.

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

The livestock classifier uses the same architecture and training recipe and reaches the same 96.00% test accuracy. With only 11 classes and visually distinctive conditions such as lumpy skin nodules, this is plausible, but class imbalance in the animal data (healthy images are far easier to collect) should be kept in mind when reading a single accuracy number.

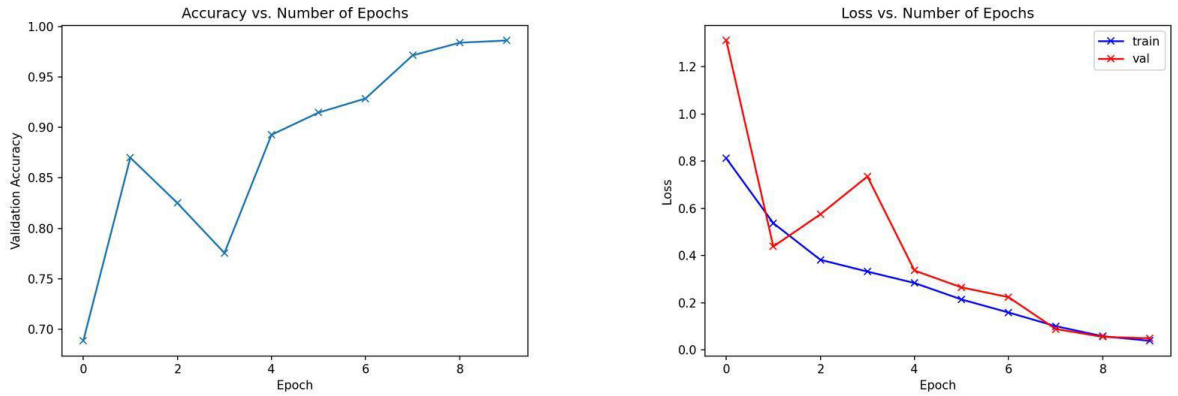


Fig. 5: Training history of the ResNet9 image classifier over 10 epochs: validation accuracy (left) and training versus validation loss (right).

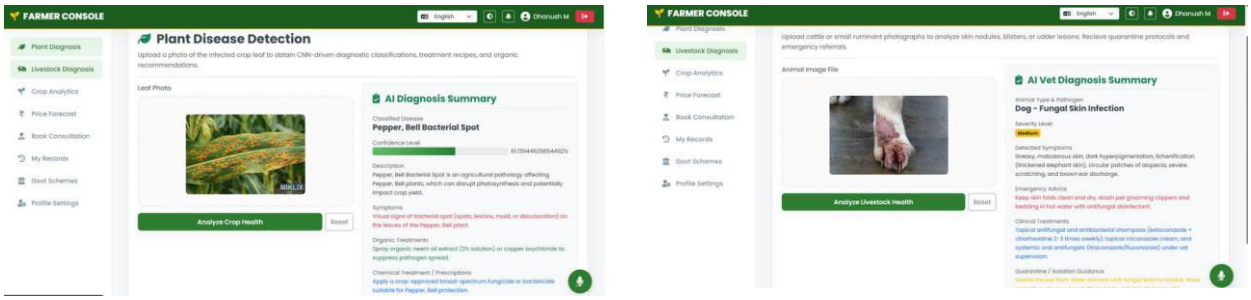


Fig. 6: Diagnosis modules in the farmer console. Left: plant disease detection on a field photograph of a maize leaf; the model returned Pepper, Bell – Bacterial Spot with 61.1% confidence, which is a misclassification. Right: livestock diagnosis of a dog paw returning Dog – Fungal Skin Infection with symptoms, treatment and quarantine advice.

C. Sample Outputs of the Deployed System

Fig. 6 shows two real outputs from the deployed portal. In the livestock module (right), a photograph of a dog's paw with red, inflamed and hairless skin was classified as Dog – Fungal Skin Infection with a Medium severity tag. The summary also lists the typical symptoms, immediate care instructions, clinical treatments to discuss with a vet and isolation guidance, which is exactly the type of information a farmer or pet owner needs before the consultation.

The plant module on the left of Fig. 6 is a more instructive case. We uploaded a field photograph of a maize leaf covered in orange rust pustules, with maize ears visible behind it. The portal returned Pepper, Bell – Bacterial Spot at about 61% confidence, although Corn – Common rust is one of the 38 classes. The lesion pattern is visually similar to bacterial spot on a segmented leaf, and the cluttered background and lighting are far from the black-background training images. We include this example deliberately: it shows that a 96% accuracy on the held-out split does not automatically carry over to photographs taken in the field, and that the confidence value displayed to the user should be treated as a warning sign when it is moderate.

Fig. 7 shows the tabular modules. For the Mandya (Karnataka) example, with Kharif season, clayey soil and N-P-K values of 45-50-40, the crop recommender returned jute with an expected yield of 4.5 t/ha; the fertilizer advisor suggested DAP at 80 kg per acre for rice with a rock phosphate and phosphate-solubilising bacteria alternative. The price module returns an expected wholesale price per quintal and plots it against recent values.

© 2026 IJIRCCE | Volume 14, Issue 11, November 2026 DOI: 10.15680/IJIRCCE.2026.1411006

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

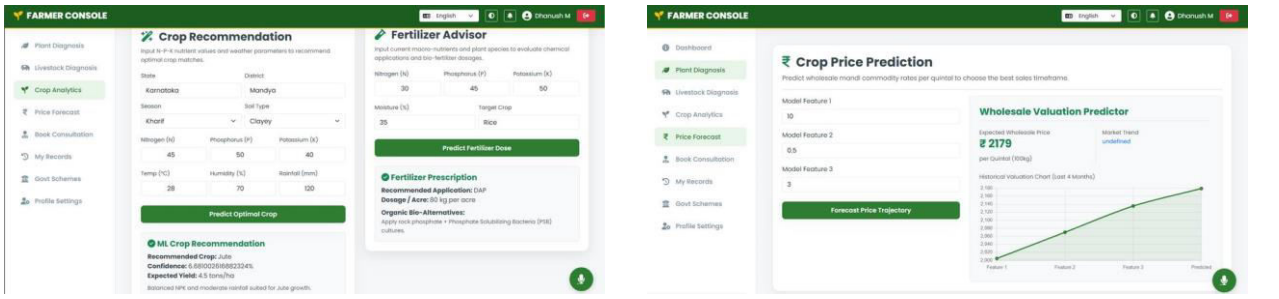


Fig. 7: Left: Crop Analytics module with crop recommendation and fertilizer advisor; the fertilizer advisor returns a DAP prescription of 80 kg per acre together with an organic bio-alternative. Right: Price Forecast module showing the predicted wholesale price per quintal and a four-month valuation chart.

D. Limitations

We want to be clear about what the current system does and does not do.

1. **Domain shift.** The plant model was trained on background-removed leaves. Field photographs with soil, hands, other leaves or strong shadows can reduce accuracy, as the maize example above shows. Retraining with field images and strong augmentation is the most important next step.
2. **Closed label set.** Both image models can only predict classes they have seen. A disease that is not in the 38 or 11 classes will be mapped to the nearest known one. An "unknown" or out-of-distribution option is needed.
3. **Limited livestock coverage.** Dog skin diseases are well represented, but cattle and goat classes are fewer and less diverse, and the system is not a substitute for a veterinary examination.
4. **Confidence calibration.** Softmax scores are not calibrated probabilities. In the crop-recommendation example above the top-ranked crop carried a low probability because inputs were spread across several plausible crops, so the interface should present the ranked list rather than a single winner.
5. **Regression uncertainty.** Yield and price predictions come without prediction intervals, and price forecasts do not capture sudden policy or weather shocks.
6. **Advisory content.** Treatment text comes from a fixed knowledge base. Doses, especially of chemical products, must be confirmed by an agronomist or veterinarian.

V. CONCLUSIONS AND FUTURE WORK

We presented PranaDhanya, an integrated portal in which two ResNet9 image classifiers and four tabular models support disease diagnosis, crop and fertilizer choice, and price awareness for farmers, with a veterinary consultation workflow built around them. The individual models perform well on their evaluation data, and the architecture keeps the machine-learning service independent of the web application so that each can evolve separately. Our field example also showed where the work is not finished.

In the future we plan to (i) collect and label field images and add domain-specific augmentation, (ii) add Grad-CAM style visual explanations so that users and veterinarians can see what the network looked at [11], (iii) calibrate the confidence scores and add an "uncertain" response, (iv) extend livestock coverage, including mastitis, and (v) deploy the interface on low-end phones with offline support and more regional languages, in line with national digital livestock initiatives.

VII. ACKNOWLEDGEMENT

The authors thank the Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, for the facilities and guidance provided during this work.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770–778.
[2] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Front. Plant Sci., vol. 7, Art. no. 1419, 2016.
[3] 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, 2015.
[4] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Comput. Electron. Agric., vol. 145, pp. 311–318, 2018.
[5] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," Comput. Electron. Agric., vol. 147, pp. 70–90, 2018.
[6] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, 2016, pp. 785–794.
[7] L. Breiman, "Random forests," Mach. Learn., vol. 45, no. 1, pp. 5–32, 2001.
[8] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learn. Represent. (ICLR), 2015.
[9] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in Adv. Neural Inf. Process. Syst. (NeurIPS), 2019, pp. 8024–8035.
[10] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.
[11] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), 2017, pp. 618–626.



INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

www.ijircce.com

Scan to save the contact details