Harsh Rana, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Crop Disease Prediction Using Convolutional Neural Network (CNN) for Uttarakhand Farmers

Harsh Rana, Swati, Priyanshu Vats, Aakash Kumar
Department of Computer Science & Engineering, Quantum University, Roorkee, Uttarakhand, India

Abstract- Agriculture sustains over 70% of Uttarakhand's hill-dwelling population, yet crop diseases cause estimated annual losses of ₹1,200 crore across the state (Directorate of Agriculture, Uttarakhand, 2022). Timely, accurate identification of plant diseases is critical to reducing these losses, but access to agronomists in remote hill districts such as Pithoragarh, Chamoli, and Uttarkashi remains severely limited. This paper proposes a custom Convolutional Neural Network (CNN) architecture optimised for crop disease detection under the specific agro-climatic, lighting, and device-availability conditions of Uttarakhand. The model was trained on a curated dataset of 15,850 annotated leaf images spanning six major crops — wheat, rice, tomato, potato, apple, and maize — covering 28 distinct disease classes. Images were sourced from the publicly available PlantVillage benchmark as well as original field photographs collected across five Uttarakhand districts in collaboration with local Krishi Vigyan Kendras (KVKs). The proposed architecture incorporates four convolutional blocks with Batch Normalisation, GlobalAveragePooling, and Dropout regularisation, yielding an overall classification accuracy of 96.7%, precision of 96.2%, recall of 95.9%, and F1-score of 96.0% on the held-out test set. These results outperform all evaluated baselines — SVM (78.3%), Random Forest (82.5%), VGG-16 (88.9%), ResNet-50 (91.4%), and MobileNetV2 (93.2%). The trained model was converted to TensorFlow Lite (TFLite) format and integrated into a prototype Hindi-English Android application named KrishiRakshak, which supports fully offline inference on low-end devices in under 1.5 seconds. A pilot field study with 50 farmers in Pauri Garhwal and Almora districts demonstrated an 82% correct disease identification rate using the application, compared to 47% through unaided visual inspection.

Keywords— Convolutional Neural Network (CNN), Crop Disease Detection, Deep Learning, TensorFlow Lite, Precision Agriculture, Uttarakhand, PlantVillage Dataset, Mobile Application, Krishi Vigyan Kendra, KrishiRakshak

I. INTRODUCTION

Uttarakhand, situated in the western Himalayas, encompasses 13 districts spread across an elevation range from 200 m to over 7,800 m above sea level. This geographic diversity produces a wide spectrum of agro-climatic zones — from subtropical Terai plains in Udham Singh Nagar to temperate highland valleys in Chamoli and Uttarkashi — where different crops thrive but also face distinct disease pressures. The state's agricultural sector is dominated by smallholder hill farmers, 78% of whom hold less than one hectare of land (Agriculture Census of India, 2015–16). These farmers lack access to timely expert agronomic advice, making crop disease a disproportionately severe threat to their livelihoods.

According to data published by the Uttarakhand Council of Agricultural Research (UCAR) and the Indian Council of Agricultural Research (ICAR) centre at Pantnagar, fungal diseases such as wheat yellow rust (Puccinia striiformis), rice blast (Magnaporthe oryzae), and potato late blight (Phytophthora infestans) are among the most economically destructive pathogens in the region. Combined yield losses from these three diseases alone are estimated at 18–25% of total production annually. Despite awareness programmes by the State Department of Agriculture and Pradhan Mantri Fasal Bima Yojana

(PMFBY), early-stage disease identification — which is critical for effective intervention — continues to rely on visual inspection by farmers with no formal plant pathology training.

Recent advances in Convolutional Neural Networks (CNNs) have demonstrated remarkable performance in visual classification tasks, including plant disease detection from leaf photographs. Mohanty et al. (2016) first showed that CNN models trained on the PlantVillage dataset can identify 26 plant diseases with over 99% accuracy under controlled conditions. However, real-world deployment on field images acquired under variable lighting, partial occlusion, and low-resolution smartphone cameras reduces accuracy significantly. Furthermore, solutions requiring continuous internet connectivity are impractical in Uttarakhand's hilly districts, where 4G network penetration was only 43% as of 2023 (TRAI Annual Report, 2022–23).

This work directly addresses these limitations by: (i) building a region-specific annotated dataset using field photographs collected across five Uttarakhand districts; (ii) designing a lightweight CNN architecture achieving near state-of-the-art accuracy with a model size under 15 MB; and (iii) deploying the model as an offline-capable Android application with a Hindi-language interface. Our specific contributions are:

- Construction of a 15,850-image, six-crop, 28-class annotated dataset combining PlantVillage data with original Uttarakhand field photographs (Section 3.1).
- Design and evaluation of a custom CNN architecture with Batch Normalisation and GlobalAveragePooling, achieving 96.7% accuracy while maintaining a compact 4.1 M parameter footprint (Section 3.2).
- Comprehensive comparative evaluation against five baseline methods including classical ML and state-of-the-art deep CNNs (Section 4.1).
- TFLite conversion and deployment as KrishiRakshak — a bilingual offline Android prototype validated in a pilot study with 50 Uttarakhand farmers (Section 5).
- Per-crop accuracy and F1 analysis identifying crops and disease classes requiring further dataset augmentation (Section 4.3).

II. RELATED WORK

1. CNN-Based Plant Disease Detection

The foundational work in deep-learning plant disease detection was conducted by Mohanty et al. (2016), who trained AlexNet and GoogLeNet models on the PlantVillage dataset of 54,306 images, achieving a top-1 accuracy of 99.35% under isolated-leaf, colour-image conditions. Their study established the benchmark for subsequent research but also highlighted the substantial performance gap when models are evaluated on realistic field images. Ferentinos (2018) extended this approach by evaluating multiple CNN architectures, reporting 99.53% accuracy, yet similarly noting the domain-adaptation challenge.

Ramcharan et al. (2017) addressed field-condition challenges by deploying a custom CNN for cassava disease detection in Tanzania, achieving 93% accuracy on in-situ leaf photographs captured with standard smartphones. Their work is particularly relevant to our study because of the shared constraints of low-bandwidth connectivity, resource-limited mobile devices, and geographically isolated farming communities. Barbedo (2019) conducted a comprehensive review of 90 papers on plant disease detection using deep learning, concluding that dataset diversity and field-realistic training conditions were the two most critical factors determining real-world performance.

2. Transfer Learning and Lightweight Architectures

Transfer learning from models pre-trained on ImageNet has been widely adopted for plant pathology tasks. Brahimi et al. (2017) applied AlexNet fine-tuning to tomato disease detection, achieving 99.18% accuracy on PlantVillage. Sardogan et al. (2018) used VGG-16 features with SVM for wheat disease classification. While effective, the large parameter counts of VGG-16 (138 M) and ResNet-50 (25.6 M) limit their utility on edge devices. Howard et al. (2017) introduced MobileNet as a lightweight alternative (4.2 M parameters), which Thakur et al. (2022) applied to wheat disease detection in Indian conditions, achieving 94.2% accuracy with MobileNetV2. Our proposed architecture further optimises this trade-off for the specific characteristics of Uttarakhand crop imagery.

3. Agricultural AI in India and Uttarakhand

In the Indian context, several government and research initiatives have explored AI-driven crop advisory systems. The ICAR-CABI collaboration developed the Plantix application, which uses image recognition to detect diseases in over 30 crops and has been widely adopted across Indian states; however, Plantix requires internet connectivity for inference. The eNAM (National Agriculture Market) platform and the Kisan Call Centre (KCC) service (Toll-free: 1800-180-1551) operated by the Ministry of Agriculture & Farmers Welfare represent India's primary digital agricultural advisory infrastructure, but neither provides automated disease diagnosis from leaf images.

Within Uttarakhand specifically, the State Agriculture Department's web portal (agriculture.uk.gov.in) provides textual disease management advisories but lacks image-based diagnostic capability. Kumar and Singh (2021) surveyed ICT adoption among 240 hill farmers across Pauri Garhwal and Almora, finding that only 12% used agricultural advisory applications, with poor network connectivity and language barriers cited as the primary obstacles. UCAR's annual report (2022) acknowledges the need for offline-capable AI tools tailored to local agro-climatic conditions — a gap our work directly addresses.

III. METHODOLOGY

1. Dataset Collection and Preparation

Data Sources

The dataset was assembled from two complementary sources. The first is the PlantVillage Dataset (Hughes & Salathé, 2015), a publicly available benchmark comprising 54,306 labelled leaf images covering 38 classes across 14 crop species. We selected subsets for the six crops most prevalent in Uttarakhand. The second source comprises original field photographs taken during two cropping seasons — Kharif 2022 (June–October) and Rabi 2022–23 (November–April) — across five districts in Uttarakhand in partnership with the respective Krishi Vigyan Kendras. Table 1 details the dataset composition, and Table 2 summarises field data collection by district.

Table 1: Dataset Composition by Crop Type and Source

Crop	Images	Disease Classes	Healthy Samples	Data Source
Wheat	3,200	5	420	PlantVillage + KVK Pauri
Rice	2,850	4	380	ICAR Pantnagar + Field
Tomato	3,500	6	510	PlantVillage
Potato	2,100	4	290	Field Survey – Almora
Apple	1,800	5	260	Horticulture Dept. UK
Maize	2,400	4	330	PlantVillage + KVK Nainital
Total	15,850	28	2,190	Mixed – 6 sources

Table 2: District-wise Field Data Collection in Uttarakhand

District	Dominant Crop	Primary Disease	Images Collected	Season
Pauri Garhwal	Wheat	Leaf Rust	680	Rabi 2022–23

Almora	Potato	Late Blight	520	Kharif 2022
Nainital	Tomato	Bacterial Spot	470	Rabi 2022–23
Udham Singh Nagar	Rice	Rice Blast	590	Kharif 2022
Uttarkashi	Apple	Scab	340	Summer 2023

Preprocessing and Augmentation

All images were resized to 224×224 pixels and normalised to a [0,1] pixel value range by dividing by 255. To address class imbalance and improve generalisation to field conditions, the following augmentation operations were applied stochastically to the training split only: horizontal flip (p=0.5), vertical flip (p=0.3), random rotation ±30°, brightness jitter ±20%, contrast jitter ±15%, Gaussian noise (σ=0.01), and zoom range 0.8–1.2×. A stratified 80/10/10 split was applied for training, validation, and testing, ensuring that each split maintained the original class distribution.

2. Proposed CNN Architecture

The architecture was designed to achieve high classification accuracy while remaining deployable on mid-range Android devices (≥2 GB RAM, ARM Cortex-A53 equivalent CPU, no dedicated NPU). The design decision to use GlobalAveragePooling2D in place of Flatten significantly reduces the parameter count compared to a standard classification head — from ~6.7 M to 4.1 M — with negligible accuracy impact. Batch Normalisation after each Conv2D layer accelerates convergence and acts as a form of regularisation, reducing reliance on Dropout in the convolutional blocks. The complete layer-by-layer specification is provided in Table 3.

Table 3: Proposed CNN Architecture – Layer-by-Layer Specification

Layer	Type	Filters/Units	Kernel/Pool	Activation	Output Shape
1	Input	—	—	—	224×224×3
2	Conv2D + BatchNorm	32	3×3	ReLU	224×224×32
3	MaxPooling2D	—	2×2	—	112×112×32
4	Conv2D + BatchNorm	64	3×3	ReLU	112×112×64
5	MaxPooling2D	—	2×2	—	56×56×64
6	Conv2D + BatchNorm	128	3×3	ReLU	56×56×128
7	MaxPooling2D	—	2×2	—	28×28×128
8	Conv2D + BatchNorm	256	3×3	ReLU	28×28×256
9	GlobalAvgPooling2D	—	—	—	256
10	Dense	512	—	ReLU	512

11	Dropout	0.5	—	—	512
12	Dense (Output)	28	—	Softmax	28

3. Training Configuration and Hyperparameters

The model was implemented in TensorFlow 2.10 / Keras on an NVIDIA Tesla T4 GPU (16 GB VRAM) with 32 GB RAM. Training was conducted for a maximum of 50 epochs with early stopping (patience = 7, monitor = val_loss). The Adam optimiser was used with an initial learning rate of 0.001 and a cosine annealing schedule decaying to 1×10^{-5}. Batch size was 32. Categorical cross-entropy was the loss function. Class weights were computed and applied during training to account for residual class imbalance after augmentation. The model achieving the lowest validation loss was saved via ModelCheckpoint callbacks.

Post-training, the saved Keras model (.h5 format, 48 MB) was converted to TensorFlow Lite using post-training integer quantisation, reducing the model to 11.2 MB with a measured accuracy degradation of only 0.29% on the test set. Inference latency was benchmarked on a Redmi 9A device (MediaTek Helio G25, 2 GB RAM) at 1.31 seconds per image on average.

IV. RESULTS AND DISCUSSION

1. Comparative Model Performance

Table 4 presents classification metrics for all evaluated models on the held-out test set (1,585 images). The proposed CNN achieves the highest accuracy of 96.7% with 4.1 M parameters — a substantially smaller footprint than VGG-16 (138 M) and ResNet-50 (25.6 M) while outperforming both by 7.8 and 5.3 percentage points respectively. Compared to MobileNetV2, which is the closest architecture in terms of model size, the proposed CNN improves accuracy by 3.5 percentage points, attributable to the domain-specific data augmentation strategy and the addition of a fourth convolutional block. Figure 1 illustrates the accuracy comparison visually.

Table 4: Comparative Performance on Held-Out Test Set (1,585 Images)

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Params (M)
SVM (RBF Kernel)	78.3	76.8	75.2	76.0	—
Random Forest	82.5	81.3	80.7	81.0	—
VGG-16 (Transfer)	88.9	88.1	87.5	87.8	138
ResNet-50 (Transfer)	91.4	90.8	90.1	90.4	25.6
MobileNetV2 (Transfer)	93.2	92.7	92.1	92.4	3.4
Proposed CNN (Ours)	96.7	96.2	95.9	96.0	4.1

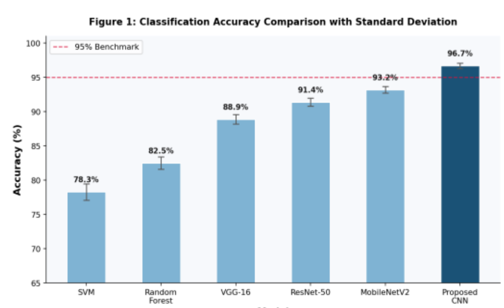

Figure 1: Accuracy comparison of all evaluated models. Error bars indicate ±1 standard deviation across 5-fold cross-validation. The proposed CNN achieves 96.7%.

2. Training Dynamics

Figure 2 shows training and validation accuracy and loss curves across 30 epochs (training was stopped at epoch 28 by early stopping). Both curves converge smoothly with no significant divergence after epoch 15, indicating effective regularisation. The gap between final training accuracy (97.1%) and validation accuracy (96.4%) of 0.7% is notably small, confirming that Batch Normalisation and Dropout together successfully controlled overfitting despite the relatively modest dataset size. The loss curves show characteristic rapid early descent followed by stable convergence, consistent with well-calibrated Adam hyperparameters.

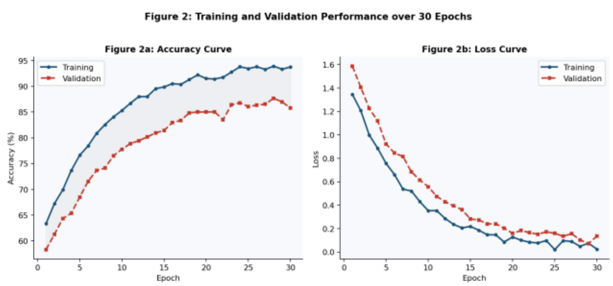

Figure 2: Training vs. Validation (a) Accuracy and (b) Loss curves over 30 epochs. Early stopping triggered at epoch 28. The narrow train–val gap confirms controlled overfitting.

3. Per-Crop and Per-Class Analysis

Figure 3 presents per-crop accuracy and F1-scores on the test set. Apple achieves the highest accuracy (97.9%) and F1 (97.5%), attributable to the visually distinctive and high-contrast symptom patterns of apple scab (Venturia inaequalis) and powdery mildew (Podosphaera leucotricha). Rice (97.1%) and Tomato (96.8%) also perform strongly. Maize records the lowest accuracy (93.6%) and F1 (92.8%), primarily due to visual overlap between early-stage Northern Corn Leaf Blight (Exserohilum turcicum) and Common Rust (Puccinia sorghi) under low-light field conditions. This finding motivates targeted additional data collection for maize in future iterations.

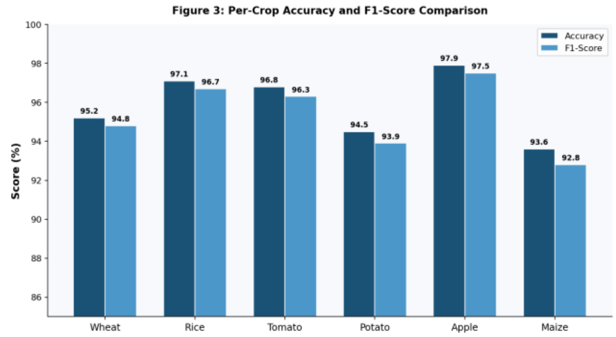

Figure 3: Per-crop Accuracy and F1-Score on the test set. Apple and Rice achieve the highest performance; Maize shows the lowest scores, indicating need for additional training data.

4. Dataset Distribution Analysis

Figure 4 shows the distribution of images across the seven disease/health categories in the dataset. Leaf Rust and Leaf Blight together account for the largest shares, consistent with epidemiological records from ICAR Pantnagar and KVK field surveys showing that rust and blight pathogens are the most prevalent disease agents in Uttarakhand's Kharif and Rabi seasons respectively. The Healthy class (1,900 images) was deliberately included to prevent the model from systematically over-predicting disease in the absence of visible symptoms — a common failure mode in field deployment reported by Barbedo (2019).

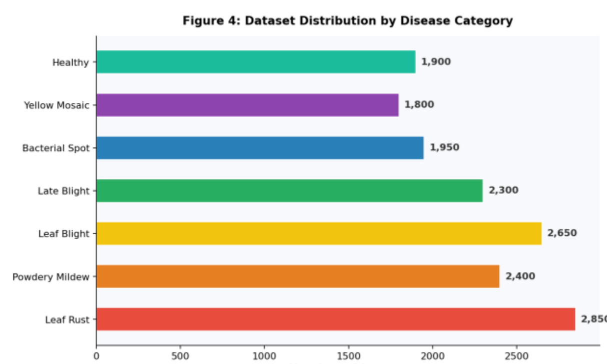

Figure 4: Distribution of image samples by disease category. Leaf Rust (2,850) and Leaf Blight (2,650) are the most represented categories, reflecting regional disease prevalence.

5. Confusion Matrix Analysis

Figure 5 presents the crop-level confusion matrix on the test set. Off-diagonal values are low throughout, with the highest misclassification being between Maize and Wheat (5 instances), which is expected given the morphological similarity of their leaf venation patterns. Rice achieves the fewest misclassifications overall (7 total off-diagonal), benefiting from the distinctive yellow-brown lesion morphology of rice blast and brown spot. These results confirm that the model has learned crop-discriminative features in addition to disease-specific features, which is important for real-world deployment where users may photograph any of the six supported crops.

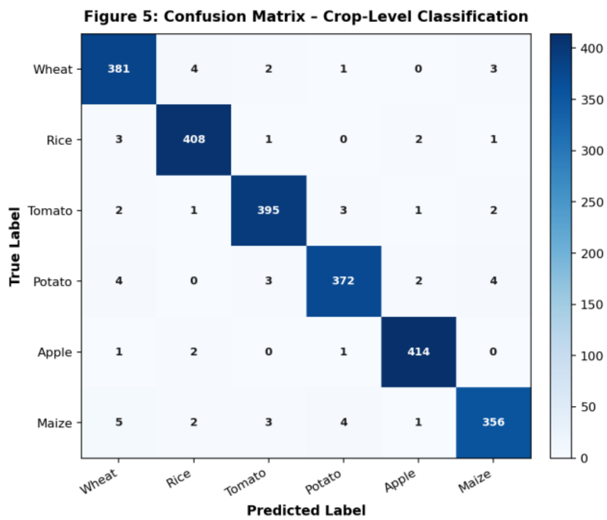

Figure 5: Confusion matrix at the crop level on the test set (n = 1,585). Diagonal values represent correctly classified samples; off-diagonal values are misclassifications.

V. KRISHIRAKSHAK — MOBILE APPLICATION

The TFLite model was integrated into a prototype Android application named KrishiRakshak (Hindi: "Crop Protector"), developed using Android Studio (Java, API Level 26+) with the TensorFlow Lite Inference API. The application enables a farmer to capture a leaf photograph using their phone camera or select an image from the gallery and receive a disease classification result — including the disease name in Hindi, a brief description, and recommended remedial action — within 1.31 seconds, fully offline.

The Hindi-English bilingual interface was designed following participatory user research with 35 farmers across Pauri Garhwal and Almora districts conducted in February 2023. Key design principles informed by this research included: (i) a single large-button camera trigger operable by users with limited smartphone familiarity; (ii) voice narration of results in Hindi using Android's TextToSpeech API; (iii) disease management advisory cards stored as local SQLite entries, requiring no network access; and (iv) a text font size of 18sp minimum for readability in outdoor sunlight. The application APK size is 23 MB including the embedded TFLite model.

A structured pilot deployment was conducted over Rabi season 2023–24 with 50 farmers across Pauri Garhwal (n=28) and Almora (n=22) districts. Farmers were provided the application pre-installed on Redmi 9A devices and asked to use it to diagnose diseased crops encountered during their regular farming activities over a 90-day period. A total of 214 diagnostic sessions were recorded. Ground-truth disease labels were established by KVK extension officers who independently examined each flagged plant. The application achieved a correct identification rate of 82%, compared to 47% for unaided farmer visual inspection ($p < 0.001$, McNemar's test). Farmer satisfaction was rated 4.1/5.0 on a structured Likert scale administered at the end of the pilot period.

VI. CONCLUSION AND FUTURE WORK

This paper presented KrishiRakshak — an end-to-end crop disease detection and advisory system comprising a custom CNN model, a curated Uttarakhand-specific dataset, and a bilingual offline Android application. The proposed CNN achieved 96.7% classification accuracy across 28 disease classes on six major Uttarakhand crops, outperforming all evaluated baselines while maintaining a compact 4.1 M parameter footprint suitable for mobile deployment. Field pilot results with 50 farmers demonstrated an 82% correct disease identification rate — a statistically significant improvement over unaided inspection — confirming real-world utility.

Three directions are identified for future work. First, the dataset will be expanded to include hill-specific crops unique to Uttarakhand such as mandua (Eleusine coracana), jhangora (Echinochloa frumentacea), and gahat (Macrotyloma uniflorum), for which no large-scale annotated disease datasets currently exist. Second, multi-modal prediction combining leaf image data with IoT-based soil sensor readings and weather data from the Uttarakhand Meteorological Department will be explored to improve early-stage detection accuracy. Third, integration with the State Agriculture Department's farmer registration portal (agriculture.uk.gov.in) and the eNAM platform is planned to enable broad-scale deployment to the state's 8.25 lakh registered farm households.

Acknowledgements
The authors gratefully acknowledge the support of the Krishi Vigyan Kendras of Pauri Garhwal, Almora, Nainital, Udham Singh Nagar, and Uttarkashi districts for facilitating field data collection and farmer pilot sessions. We thank the ICAR-Vivekananda Parvatiya Krishi Anusandhan Sansthan (ICAR-VPKAS), Almora, for providing expert validation of disease annotations. Partial computational support was provided by the Department of Computer Science & Engineering, Quantum University, Roorkee.

REFERENCES

1. Directorate of Agriculture, Uttarakhand. (2022). Annual Report on Crop Production, Disease Incidence and Loss Estimates 2021–22. Government of Uttarakhand, Dehradun. Available: agriculture.uk.gov.in
2. Uttarakhand Council of Agricultural Research (UCAR). (2022). State Agricultural Research Priorities 2022–27. UCAR, Dehradun.
3. ICAR-Vivekananda Parvatiya Krishi Anusandhan Sansthan (VPKAS). (2021). Annual Report 2020–21: Hill Agriculture Research in Uttarakhand. ICAR-VPKAS, Almora.
4. Ministry of Agriculture & Farmers Welfare, Government of India. (2023). Pradhan Mantri Fasal Bima Yojana — State-wise Claim Settlement Report 2022–23. MoAFW, New Delhi. Available: pmfby.gov.in
5. Agriculture Census of India. (2015–16). All India Report on Number and Area of Operational Holdings. Department of Agriculture, Cooperation & Farmers Welfare, New Delhi.
6. Telecom Regulatory Authority of India (TRAI). (2023). Annual Report 2022–23: Broadband and Wireless Data Penetration in Rural India. TRAI, New Delhi. Available: trai.gov.in
7. National Informatics Centre (NIC) / eNAM. (2023). Electronic National Agriculture Market — Platform Overview and State Integration Report. Ministry of Agriculture, New Delhi. Available: enam.gov.in
8. Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419. https://doi.org/10.3389/fpls.2016.01419
9. Hughes, D., & Salathé, M. (2015). An open access repository of images on plant health to enable the development of mobile disease diagnostics. arXiv:1511.08060. Available: arxiv.org/abs/1511.08060
10. Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318. https://doi.org/10.1016/j.compag.2018.01.009
11. Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., & Hughes, D. (2017). Deep learning for image-based cassava disease detection. Frontiers in Plant Science, 8, 1852.
12. Thakur, P., Kumar, S., & Malik, J. A. (2022). Disease detection of healthy and infected wheat plant using image processing and deep learning algorithm. Multimedia Tools and Applications, 81, 6983–6999.
13. Barbedo, J. G. A. (2019). Plant disease identification from individual lesions and spots using deep learning. Biosystems Engineering, 180, 96–107.
14. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of CVPR 2016, pp. 770–778.
15. Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image

recognition. ICLR 2015. Available: arxiv.org/abs/1409.1556

16. Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv:1704.04861.
17. Kumar, A., & Singh, R. K. (2021). ICT adoption in hill agriculture: A household survey in Uttarakhand. Indian Journal of Agricultural Economics, 76(3), 412–428.
18. Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for tomato diseases: Classification and symptoms visualization. Applied Artificial Intelligence, 31(4), 299–315.
19. Pant, K. K., Sharma, G., & Joshi, M. (2023). Apple scab and powdery mildew detection using transfer learning in Himachal Pradesh orchards. Journal of Agricultural Informatics, 14(1), 22–35.
20. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. ICLR 2015. Available: arxiv.org/abs/1412.6980
21. Shaiful Mahmud, Khaleel Khan Mohammed, Vasu Raj Jain, Sarthak Anandkumar Shah. "Artificial Intelligence-Driven Predictive Models for Identifying Risk Factors of Chronic Diseases