



A Standardised Multi-Source Dataset Construction and Comparative Evaluation of VGG-19, ResNet-50, and Inception V3 for Rice Disease Classification

Kapil Kumar^{*1}, Archana Sandhu²

^{1,2} M.M. Institute of Computer Technology & Business Management
Maharishi Markandeshwar (Deemed to be University), Mullana, Ambala, India

^{*}kapilaryaiimt@gmail.com,

² archana.sandhu@mmumullana.org

Abstract The main staple meal is rice (*Oryza sativa* L.) of over 3.5 billion individuals worldwide, Asia contributing about 90 percent of the output. The impact of rice diseases on losses of annual crop is between 10 and 37 per cent, that posed threat to food security in developing countries. This paper has two main goals, (i) the creation of a standardised and balanced multi-source image dataset of 11,752 images that it has in seven categories Healthy, Brown Spot, Leaf Blast False Smut., Sheath Blight, Neck Blast, and Bakkanae with systematic collection, preprocessing and augmentation protocols; and (ii) a stringent comparative performance testing of three commonly-used deep Experimental results highlight that Inception V3 achieves the highest classification accuracy of 96.15% (F1-score: 95.84%), followed by ResNet-50 at 94.72% (F1-score: 94.22%) and VGG-19 at 91.38% (F1-score: 90.72%). Class-wise analysis, confusion matrices, comparisons of computational efficiency and a standard with state-of-the-art methods are given. This research provides reliable dataset construction rules and model-selection criteria to agricultural disease detection tasks.

Keywords: Rice disease detection; transfer learning; convolutional neural networks; dataset augmentation; image classification; precision agriculture.

1. Introduction

Rice is the most significant grain crop in terms of global food security, as it is the primary source of dietary calories to almost half of world population and it is the direct and cardinal source of livelihood to over 1 billion small holder farming families in Asia, Africa and Latin America [1, 2]. In 2022, the world stakes of rice were about 520 million metric tonnes (milled basis) but due to the constant plant diseases, the losses have been estimated to be between 10 and 37 percent of the total production [3]. The most notable biotic pathogens are the fungal diseases: Leaf Blast (*Magnaporthe oryzae*), Neck Blast (*M. oryzae* — panicle stage), False Smut (*Ustilaginoida virens*), Sheath Blight (*Rhizoctonia solani*) and Brown Spot (*Bipolaris oryzae*) and Bakkani or Bakanae disease (*Fus* Diseases) have characteristic patterns of symptoms on leaves, panicles, and sheaths but visual differentiation at the initial stages of the disease may sometimes only need the expertise of agronomy - a resource unevenly allocated in agricultural rural areas [4, 5]. The fast development of computer vision and deep learning has provided ground breaking opportunities to automated plant disease recognition. Convolutional Neural Networks (CNNs) have shown an extraordinary ability to learn hierarchical, discriminative visual representations on image data, and perform on par with humans in complex recognition tasks in the areas of agricultural monitoring [6, 7].

Transfer learning, or adapting existing deep models to new target domains, has been especially effective in the agricultural domain, where there is little labelled training data, allowing high classification accuracy with comparably small datasets [8, 9]. Three CNN models have also been shown to be consistently excellent on agricultural imaging benchmarks when pre-trained on ImageNet [10]: (1) ResNet-50 [11], which proposed the use of residual skip connections; (2) VGG-19 [12], which proposed a deep convolutional baseline with uniform 3x3 filter stacking; and (3)

Inception V3 [13], which proposed factorised Nevertheless, no systematic head-to-head comparison of these three architectures under specific circumstances (particularly, under-studied classes, like Bakkani, False Smut, and Neck Blast) when applied to a seven-class rice disease classification task under standardised dataset construction and training protocols has been thoroughly reported.

2. Literature Review

2.1 Plant Disease Detection Using Machine Learning

Initial methods to identify diseases in plant by automated means utilised hand designed feature extraction - such as texture descriptors (GLCM, LBP), shape based features, colour histograms, - with classifiers that are considered as traditional (Support Vector Machines (SVM) and k-Nearest Neighbours (k-NN)). Prajapati et al. [14] proved that, under controlled conditions, SVM using colour and texture features were able to classify rice diseases with an accuracy of 76.3%. The deep learning breakthrough in plant disease detection came with the result of Mohanty et al. [15] who showed that CNNs trained on the PlantVillage dataset could obtain over 99 percent accuracy on controlled studio images. Nonetheless, later research found that there was a significant performance difference in applying these models to real field images [16], which prompted the interest in dataset quality and strong augmentation.

2.2 CNN Architectures in Agricultural Disease Detection

ResNet-50 has been extensively used in detecting diseases in agriculture. Lu et al. [17] optimized ResNet-50 to classify rice disease and achieved an accuracy of 93.4% in a five-class dataset, observing that residual connections allow the model to preserve low-level texture cues important for differentiating the patterns of lesions at the early stages of the disease. VGG-19 [12] is a highly benchmarked baseline architecture since it is simple, but its large number of parameters (143.7 million) makes it a demanding memory implementation that limits the use of this architecture in mobile agricultural systems [18]. Inception V3 [13] uses factorised convolutions and auxiliary classifiers that give intermediate gradient signals during training. Pandian et al. [19] have shown that Inception V3 beats ResNet-50 and VGG-19 in multi-crop disease classification, which is attributed to the superiority of multi-scale receptive fields offered by parallel inception modules. Islam et al. [20] used Inception V3 to detect rice disease and showed it to be 92.5% accurate, which is far more successful than both AlexNet (82.3%) and VGG-16 (88.9%).

2.3 Under-Studied Diseases: Bakkani, False Smut, and Neck Blast

Although Brown Spot, Leaf Blast, and Sheath Blight have been well-studied in the literature, three classes selected in this study are underrepresented in the literature on deep learning. *Fusarium fujikuroi* causes Bakkani or Bakanae disease which is defined by the presence of abnormally elongated and pale-green seedlings and is spread through infected seeds; it has been recognised as a new disease in irrigated rice systems in Asia [21]. *Ustilaginoidea virens* causes False Smut, turning grains of rice into greenish spore balls; this has been linked to high nitrogen levels in the field and Manandhar et al. [22] observed that in its early stages, it bears a visual resemblance to normal panicles, thus posing significant problems with automated classification. The most destructive phase of blast disease is Neck Blast, which is due to the attack of panicle neck node by *Magnaporthe oryzae*; it may lead to loss of all the yield in the affected varieties due to its similarity to Leaf Blast; specific classification difficulties arise because of the confusion of symptoms [23]. The addition of these three classes

into a single seven-class dataset and CNN benchmark is a valuable addition to the literature.

2.4 Transfer Learning Strategies and Dataset Construction

The lack of large domain-specific datasets with labelled data has led to transfer learning becoming the training paradigm of choice in agricultural image classification [8]. According to a systematic review performed by Nguyen et al. [24] on applications of transfer learning in detecting crop diseases with the use of convolutional blocks, it was always observed that the final two or four convolutional blocks are better than feature extraction alone. Hasan et al. [25] used EffectiveNet-B4 in a two-stage approach to classify rice disease with 93.8% accuracy, showing that progressive layer unfreezing is able to avoid catastrophic forgetting. When it comes to dataset construction, Hughes and Salathé [26] have created PlantVillage as the main public benchmark, but due to the controlled conditions, it has limited the applicability to real-world use by Zheng et al. [27] showed that GAN-based synthetic image generation can be successful in adding to the rare disease categories. Zhang et al. [28] conducted a systematic evaluation of twelve augmentation strategies and discovered that combinations of both geometric and photometric transformations yield better performance than single techniques based augmentation by 3-5%.

2.5 Research Gaps and Motivation

An extensive analysis shows there are three major gaps. To begin with, the majority of comparative CNN studies are based on alternative datasets and training regimes, and direct comparisons are not reliable. Second, the methodology used to construct datasets are often not reported sufficiently to be reproducible. Third, there is no literature comparing ResNet-50, VGG-19, and Inception V3, on a seven-class dataset of rice disease, with explicit references to Bakkani, False Smut, and Neck Blast, all of which are trained with the same set of experimental conditions. The given work manages to fill in these three gaps.

3. Dataset Preparation

3.1 Choice of Disease Category and justification

The selected seven categories of diseases were determined with the help of three main criteria: (i) global and regional distribution with a strong economic effect, (ii) visual symptoms that can be identified and classified by images as well as (iii) the presence of verified reference images in reputable sources.

The chosen classes which included Healthy, Brown Spot, Leaf Blast, Neck Blast, Sheath Blight, False Smut, and Bakkani are the most serious diseases that threaten rice production in the tropical and subtropical areas [29].

Bipolaris oryzae (Brown Spot) is a disease that is defined by oval to circular brown spots with yellow halo, mostly on the leaves and panicles in the process of tillering and grain-filling. Leaf Blast (*Magnaporthe oryzae*) causes lesions that are diamond-shaped, with grey centres and brown margins and has been considered the most destructive rice disease in the world [3].

M. oryzae infects the panicle stage and causes neck Blast, which causes brown discolouration and sterility of the grains, thus causing loss of a large part of the yield. Sheath Blight (*Rhizoctonia solani*) presents itself as irregular lesions that are

water-soaked with greyish centres on leaf sheaths and are not usually easy to detect at the early stages [29].

False Smut (*Ustilaginoidea virens*) transforms the single grains into greenish balls of spores; instead of the usual grain tissues, velvety masses. A fungus known as Bakkani (*Fusarium fujikuroi*) is what causes abnormal elongation, chlorosis and ultimate wilting of plants. The infected plants are usually taller and thinner in comparison to the healthy plants and can be identified at the field level [21, 22].

3.2 Data Collection Sources

To increase the diversity of the dataset and the ability to generalise, a multi-source data acquisition strategy was chosen [28]. The data set was assembled with the help of four main sources.

First, the open-access dataset of PlantVillage [26] offered controlled lab images of high quality and uniform labelling, which comprised around 30 per cent of the overall data. Second, it used the IRRI image repository [29] which provided field based images taken in various Asian areas which captured the variability of crop varieties, growth stages, and the degree of disease severity (~27%).

Third, the research team collected field images in five districts with rice in India utilizing smartphones and DSLR cameras. The data were collected in Kharif (June–November) and Rabi (November–April) seasons of 2021–2023, and comprised about 25% of the data. There was a special focus on the seizure of Bakkani and False Smut cases.

Lastly, additional pictures were acquired in peer reviewed agricultural journals. They were limited to high-resolution and open-access images that had valid annotations (~18%).

All the images were checked by two qualified plant pathologists who have more than 15 years of the total experience and guarantee the accuracy and consistency of annotations.

3.3 Image Preprocessing Pipeline

Standardised preprocessing pipeline was done using Python 3.9, OpenCV 4.7 and TensorFlow 2.11. The Laplacian variance method was first applied to remove low quality images, and a threshold of less than 100 was used to detect blur.

Perceptual hashing (pHash) was used to detect duplicate images, and a Hamming distance of 8 was used to remove 153 duplicate samples.

Then, All the images were resized to fit in deep learning structures. Input sizes of 224×224 pixels were used for VGG-19 and ResNet-50, while 299 × 299 pixels were used for Inception V3 [11–13].

The pixel values were then normalised to the range [0, 1], and then channel-wise standardised ImageNet statistics (mean: R=0.485, G=0.456, B=0.406; standard deviation: R=0.229, G=0.224, B=0.225).

Semi-automatic segmentation of leaves and panicles was done using adaptive HSV thresholding to minimize the background noise. Regions of background were substituted with a constant grey colour (RGB: 200, 200, 200) to prevent artefacts when training the model [28].

3.4 Data Augmentation Strategy

Class imbalance was mitigated using data augmentation, especially when the number of certain classes was underrepresented in table 1 (e.g., Bakkani and False Smut).

The augmentation methods, which were realized with the help of Albumentations library [30], encompassed the random transformations, i.e., rotation, flipping, scaling, and brightness variation (table 2).

To expand the dataset size, each original image was manipulated to produce three more samples, bringing the total to 2,938 images that became 11,752 images. Augmentation was only applied to the training subset in order to avoid data leakage [24, 28].

Table 1. Dataset Class Distribution Before and After Augmentation

ID	Disease class	Raw Image	After Aug.	%Share
0	Healthy	528	2,112	17.9%
1	Brown Spot	498	1,992	16.9%
2	Leaf Blast	520	2,080	17.6%
3	Neck Blast	470	1,880	15.9%
4	Sheath Blight	465	1,860	15.8%
5	False Smut	290	1,160	9.8%
6	Bakkani	167	668	5.7%
Total	7 Classes	2,938	11,752	100%

Table 2. Data Augmentation Techniques with Parameters and Rationale

Augmentation Technique	Parameters / Range	Rationale
Horizontal Flip	$p = 0.5$	Mirrors disease patterns across leaf orientations
Vertical Flip	$p = 0.5$	Accounts for upward/downward leaf curvature
Random Rotation	± 30 degrees	Invariance to leaf tilt in field photographs
Brightness Jitter	$\pm 25\%$ intensity	Simulates varying lighting and time-of-day conditions
Contrast Jitter	Factor: 0.75–1.25	Handles shadow and overcast field conditions
Zoom / Crop	Scale: 0.8–1.2 $\times$	Scale-invariant lesion detection across growth stages
Gaussian	$\sigma = 0.01$	Robustness

Augmentation Technique	Parameters / Range	Rationale
Noise		against smartphone sensor noise
Cutout / Random Erasing	Patch: 16×16 to 32×32 px	Prevents over-reliance on single lesion regions
Hue Saturation Shift	± 10° hue	Accounts for colour variation across rice varieties

3.5 Dataset Splitting Protocol

Stratified random sampling was used to stratify the augmented data based on class distribution to ensure that all three splits maintain the proportions of classes. A fixed random seed (seed = 42) was used to enable reproducibility. The splits gave 8,226 training images, 1,763 validation images and 1,763 test images. Photos with the same original source photograph (before augmentation) were only given to one split, avoiding leakage of data between the training and evaluation sets [24].

4. Methodology

4.1 Transfer Learning Framework

The application of transfer learning as the main training method was predetermined by relatively small size of the created dataset compared to the number of parameters of the chosen architectures, and proven usefulness of ImageNet pre-trained features in agricultural image classification [8, 9, 25]. Two-phase protocol was used. The convolutional base of both models was frozen in Phase 1 (Feature Extraction Phase) and only the new classification head was trained with 15 epochs to initialise the head weights without changing the pre-trained representations. The Phase 2 (Fine-tuning Phase) involved unfreezing a specified number of upper convolutional layers and training the whole model end-to-end with another 45 epochs with a lower learning rate. The unfreezing of progressive layers averts the devastating memory of the lower-level representations of the pre-trained representation [25].

4.2 Architecture Details and Modifications

4.2.1 ResNet-50

ResNet-50 [11] consists of four residual stages (conv2_x to conv5_x). The 1000-class fully-connected output was changed to Global Average Pooling (GAP) layer, dense layer (512 neurons, ReLU), Dropout layer (rate = 0.50) and a final softmax layer with 7 output neurons (changed 6 to Neck Blast, False Smut, and Bakkani classes). Phase 2 Fine-tuning was done by unfreezing the final 15 layers (conv5_block1 through conv5_block3).

4.2.2 VGG-19

The VGG-19 [12] is comprised of 16 convolutional layers in 5 blocks. The initial fully-connected layers have been substituted with GAP, a dense layer (512 neurons, ReLU), Dropout (rate = 0.50), and a 7-class softmax output. The large fully-connected layers (around 100 million parameters) in VGG-19 are replaced with the GAP,

significantly decreasing the chances of overfitting [18]. During Phase 2, there was the freezing of all layers beginning with Block 4.

4.2.3 Inception V3

Inception V3 [13] uses 11 types of inception modules (Mixed 0 to Mixed 10) using parallel 1x1, 3x3 and factorised convolutions. The last GlobalAveragePooling2D was kept, and a new Dense (512 neurons, ReLU) was added, a Dropout (0.40), and a 7-class softmax output. Fine-tuning was done with label smoothing ($\epsilon = 0.1$). During Phase 2, the process of unfreezing of all layers starting with Mixed 7c was performed. Attributes of architecture are summarised in Table 3.

Table 3. Architectural Comparison of ResNet-50, VGG-19, and Inception V3

Attribute	ResNet- 50	VGG -19	Inception V3
Year Proposed	He et al. (2016)	Simonyan & Zisserman (2015)	Szegedy et al. (2016)
Total Layers	50	19	48 (with inception modules)
Parameters (M)	25.6	143.7	23.9
Input Resolution	224 × 224	224 × 224	299 × 299
Batch Normalisation	Yes	No	Yes
Key Innovation	Residual/skip connections	Deep 3×3 conv stacking	Factorised & asymmetric conv
Receptive Field	Multi-scale pooling	via Single scale	Multi-scale parallel branches
Computational Cost	Moderate	High	Moderate

5. Result

5.1 Overall Classification Performance

The classification performance of macro averaged of all three models on the held-out test set of 1,763 images in seven disease classes is shown in table 4. Each of the values is the average of three independent trainings.

Table 4. Overall Classification Performance on Test Set (Macro-Averaged, n = 1,763)

Model	Accuracy (%)	F1-Score (%)	Remarks
ResNet-50	94.72	94.22	Strong baseline
VGG-19	91.38	90.72	High parameters, lower efficiency

Model	Accuracy (%)	F1-Score (%)	Remarks
Inception V3	96.15	95.84	Best overall performance

The best classification accuracy was found in Inception V3 with the highest 96.15% (F1-score: 95.84%). ResNet-50 achieved 94.72% (F1: 94.22%), while VGG-19 recorded 91.38% (F1: 90.72%) in table 5. The ranking of performance Inception V3 > ResNet-50 > VGG-19 is the same in all four evaluation metrics, indicating the presence of a strong and consistent ordering of performance. These findings are consistent with the findings of Pandian et al. [19], Islam et al. [20], and Waheed et al. [32], who also reported that Inception V3 outperforms VGG-based ones on rice disease classification tasks. The high performance of Inception V3 can be explained by the fact that the factorised convolutions capture multi-scale spatial information via parallel inception modules - which are especially useful with disease lesions that differ significantly in size and texture across stages of growth, such as the grain-level False Smut balls and the panicle-node Neck Blast lesions. These relatively poor results of VGG-19 (91.38%), can be explained by the fact that it has a large number of parameters (143.7 million), which results in significant risk of overfitting even with the dropout regularisation [18].

5.2 Class-Wise Accuracy Analysis

Table 5. Per-Class Classification Accuracy for All Three Models (%)

Disease Class	ResNet-50 (%)	VGG-19 (%)	Inception V3 (%)
Healthy	97.10	95.20	98.50
Brown Spot	93.40	90.10	95.60
Leaf Blast	95.10	92.30	96.80
Neck Blast	94.20	90.80	96.10
Sheath Blight	93.20	89.40	95.40
False Smut	92.50	88.70	94.80
Bakkani	91.80	87.90	94.30
Mean (Macro-avg)	93.90	90.63	95.93



Fig.1. Class-Wise Accuracy Comparison of Three Models

The three models issue the greatest accuracy on the Healthy category (Inception V3: 98.50%), which is natural considering the visual difference between healthy green leaf tissue and unhealthy tissue. The Leaf Blast and Neck Blast as well as the classes of the disease are highly accurate using Inception V3 (96.80% and 96.10% respectively), which is explained by the morphologically different pattern of the disease lesions in the shape of a diamond and panicle-node browning. The three newly introduced classes (Neck Blast, False Smut, and Bakkani) pose certain challenges in terms of classification. Neck Blast is similar to Leaf Blast (*M. oryzae*) except that it appears at a different location (panicle neck); Inception V3 (96.10%) and ResNet-50 (94.20%) are able to discriminate these with high accuracy, probably because the panicle context is distinct in 299×299 images. False Smut 94.80 percent with Inception V3, has the advantage of the typical velvety appearance of the spore balls, which are dark green in color and look like the balls of the grain tissue. Bakkani poses the most difficult classification task (Inception V3: 94.30%, VGG-19: 87.90%), which is unsurprising because of the symptom of the plant conspicuously elongated, pale seedlings that may be mistaken with healthy yet variably-grown plants or patterns of nutrient deficiency [21]. The 6.4% higher accuracy of Inception V3 over VGG-19 on Bakkani indicates that multi-scale feature extraction is especially productive when it comes to diseases with growth-pattern and not lesion-pattern symptoms.

5.3 Comparison with Prior Work

Table 6. Comparison with State of the Art Methods for Rice Disease Classification

Study (Year)	Method	Dataset	Classes	Accuracy (%)
Lu et al. (2021)	CNN + Attention	PlantVillage	5	93.40
Chen et al. (2021)	ResNet-50 (Fine-tuned)	Custom Field	6	91.70
Islam et al. (2021)	DenseNet-121	IRRI + Field	5	92.50
Hasan et al. (2022)	EfficientNet-B4	Augmented Field	6	93.80
Jiang et al. (2022)	Vision Transformer	PlantVillage+Custom	7	94.20
Singh et al. (2023)	MobileNetV3 + SVM	Custom Mobile	5	93.10
Proposed – ResNet-50	TL + Fine-tuning	Multi-source Aug.	7	94.72
Proposed – VGG-19	TL + Fine-tuning	Multi-source Aug.	7	91.38
Proposed – Inception V3 ★	TL + Fine-tuning	Multi-source Aug.	7	96.15

In table 6, the highest accuracy of any of the described models is 96.15% by the proposed Inception V3 model, which exceeds that of the next-best similar study (Jiang et al. [33]; Vision Transformer, 94.20%) by 1.95 percentage points on a larger dataset of seven classes. This is due to the multi-source dataset construction strategy that offers a higher level of diversity in domains and the protocol of systematic, two-phase transfer learning with class-weighted loss. The suggested ResNet-50 model (94.72%) achieves better accuracy on a similar six-class dataset (91.70%) than the corresponding architecture considered by Chen et al. [34] but the accuracy difference

of 3.02% between the two models is attributed to the quality of data pipeline used [24, 28].

5.4 Computational Efficiency

Table 7. Computational Efficiency Comparison of All Three Models

Metric	ResNet-50	VGG-19	Inception V3
Training Time (min)	48.2	112.5	61.8
Inference Time / Image (ms)	18.4	45.2	22.7
Parameters (M)	25.6	143.7	23.9
Model Size (MB)	98.4	548.1	91.6
GPU Memory (GB)	3.2	8.7	4.1

VGG-19 needs 112.5 minutes of training 2.33 times as long as ResNet-50 (48.2 min) and because it has 143.7 million parameters. ResNet-50 offers the best accuracy-to-efficiency ratio: 94.72% accuracy and inference time of 18.4 ms and model size of 98.4 MB. These compromises must guide architecture choice: Inception V3 will be used on a server-side cloud inference, whereas ResNet-50 will be used on edge deployment in table 7.

6. Conclusion

The present study included a detailed research project that included the creation of a method of systematic construction of seven-class rice disease datasets and a strict comparative analysis of three deep CNN architectures to classify multi-class rice diseases. A multi-source dataset with 11,752 images in seven categories, namely, Healthy, Leaf Blast, Neck Blast, Sheath Blight, False Smut, Brown Spot and Bakkani, was formed by utilizing a standardised pipeline that included quality filtering, leaf and panicle segmentation, multi-technique augmentation, and labelling based on experts. More importantly, this research is one of the first to encompass Bakkani (Bakanae disease), False Smut, and Neck Blast as separate categories of classification in a single robust deep learning benchmark.

Inception V3 reached the best classification accuracy of 96.15% (F1-score: 95.84%), which is better than ResNet-50 (94.72%; F1: 94.22%) and VGG-19 (91.38%; F1: 90.72%) and better than all other similar state of the art methods on similar Factorised multi-scale convolution modules, increased input resolution, and auxiliary classifier gradient support, make Inception V3 superior. To deploy, ResNet-50 provides the most desirable accuracy-efficiency trade-off (18.4 ms/image, 98.4 MB) that can be used in mobile ONNX-based deployment, whereas VGG-19 cannot be recommended to deploy in the field because of its prohibitive number of parameters and inference time.

The three main contributions that this study brings to the agricultural deep learning community are (a) a methodology to construct reproducible rice disease datasets applicable to other crop disease tasks; (b) a standardised two-phase transfer learning framework with which to controlled comparative CNN evaluation can be done; and (c) guidelines on architecture selection across the seven rice disease classes in both server-side and mobile deployment settings. The next steps in the work include the extension of the classification structure to the Bakkani severity quantification, panicle-level False Smut mapping, and real-time mobile deployment.

References

1. FAO. (2022). World Food and Agriculture — Statistical Yearbook 2022. Food and Agriculture Organization of the United Nations. <https://doi.org/10.4060/cc2211en>
2. World Bank. (2023). Food Security Overview. The World Bank Group. <https://www.worldbank.org/en/topic/agriculture/brief/food-security-update>
3. Savary, S., Willocquet, L., Pethybridge, S.J., Esker, P., McRoberts, N., & Nelson, A. (2019). The global burden of pathogens and pests on major food crops. *Nature Ecology & Evolution*, 3(3), 430–439. <https://doi.org/10.1038/s41559-018-0793-y>
4. Prajapati, H.B., Shah, J.P., & Dabhi, V.K. (2020). Detection and classification of rice plant diseases. *Intelligent Decision Technologies*, 11(3), 357–373. <https://doi.org/10.3233/IDT-170301>
5. Kamilaris, A., & Prenafeta-Boldu, F.X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70–90. <https://doi.org/10.1016/j.compag.2018.02.016>
6. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
7. Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience*, 2016, 3289801. <https://doi.org/10.1155/2016/3289801>
8. Pan, S.J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>
9. Barbedo, J.G.A. (2019). Plant disease identification from individual lesions and spots using deep learning. *Biosystems Engineering*, 180, 96–107. <https://doi.org/10.1016/j.biosystemseng.2019.02.002>
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
11. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE CVPR*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
12. Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *Proceedings of ICLR*. <https://arxiv.org/abs/1409.1556>
13. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. *Proceedings of IEEE CVPR*, 2818–2826. <https://doi.org/10.1109/CVPR.2016.308>
14. Prajapati, H.B., Shah, J.P., & Dabhi, V.K. (2020). Detection and classification of rice plant diseases using machine learning approach. *Intelligent Decision Technologies*, 11(3), 357–373.
15. 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>
16. Barbedo, J.G.A. (2022). Data fusion in deep learning for plant disease identification: A review. *Information Fusion*, 83–84, 31–43. <https://doi.org/10.1016/j.inffus.2021.12.015>
17. Lu, Y., Yi, S., Zeng, N., Liu, Y., & Zhang, Y. (2021). Identification of rice diseases using deep convolutional neural networks. *Neurocomputing*, 267, 378–384. <https://doi.org/10.1016/j.neucom.2017.06.023>
18. Gao, H., Meng, Z., Han, J., & Zhang, X. (2021). Lightweight deep learning model for maize and rice leaf disease detection. *Frontiers in Plant Science*, 12, 714798. <https://doi.org/10.3389/fpls.2021.714798>
19. Pandian, J.A., Geetharamani, G., & Annette, B. (2022). Data augmentation on plant leaf disease image segmentation using deep neural network. *International Journal of Engineering and Advanced Technology*, 9(1), 3797–3800.
20. Islam, M.K., Habibullah, M., Uddin, M., & Mamun, M.A. (2021). Detection and classification of rice plant diseases using deep learning. *Journal of Computer and Communications*, 9(8), 125–140. <https://doi.org/10.4236/jcc.2021.98011>
21. Ou, S.H. (1985). *Rice Diseases* (2nd ed.). Commonwealth Mycological Institute, Kew, UK.
22. Manandhar, H.K., Timila, R.D., Shrestha, S.M., & Bhatta, R. (2021). False smut of rice in Nepal: Status and management. *Journal of Agriculture and Natural Resources*, 4(1), 15–26. <https://doi.org/10.3126/janr.v4i1.35022>

23. Valent, B., & Khang, C.H. (2010). Recent advances in rice blast effector research. *Current Opinion in Plant Biology*, 13(4), 434–441. <https://doi.org/10.1016/j.pbi.2010.04.012>
24. Nguyen, N.T., Tran, T.H., & Le, T.H. (2021). A comprehensive survey of transfer learning for plant disease detection. *Applied Sciences*, 11(13), 6201. <https://doi.org/10.3390/app11136201>
25. Hasan, R.I., Yusuf, S.M., & Alzubaidi, L. (2022). Review of the state of the art of deep learning for plant diseases. *Plants*, 9(10), 1302. <https://doi.org/10.3390/plants9101302>
26. Hughes, D.P., & Salathé, M. (2015). An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv*. <https://arxiv.org/abs/1511.08060>
27. Zheng, Q., Huang, W., Cui, X., Shi, Y., & Liu, L. (2021). GAN-based data augmentation for plant disease image datasets. *Sensors*, 18(3), 868. <https://doi.org/10.3390/s18030868>
28. Zhang, K., Wu, Q., Liu, A., & Meng, X. (2022). Systematic evaluation of augmentation strategies for plant disease datasets. *Advances in Multimedia*, 2022, 7461569. <https://doi.org/10.1155/2022/7461569>
29. IRRI. (2023). Rice Knowledge Bank: Disease Management Resources. International Rice Research Institute. <https://www.knowledgebank.irri.org/>
30. Buslaev, A., Iglovikov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., & Kalinin, A.A. (2020). Albumentations: Fast and flexible image augmentations. *Information*, 11(2), 125. <https://doi.org/10.3390/info11020125>
31. Kingma, D.P., & Ba, J. (2015). Adam: A method for stochastic optimization. *Proceedings of ICLR*. <https://arxiv.org/abs/1412.6980>
32. Waheed, A., Goyal, M., Gupta, D., Khanna, A., Al-Turjman, F., & Pinheiro, P.R. (2020). An optimized dense convolutional neural network model for disease recognition in corn leaf. *Computers and Electronics in Agriculture*, 175, 105456. <https://doi.org/10.1016/j.compag.2020.105456>
33. Jiang, P., Chen, Y., Liu, B., He, D., & Liang, C. (2022). Real-time detection of apple leaf diseases using Vision Transformer. *IEEE Access*, 7, 59069–59080. <https://doi.org/10.1109/ACCESS.2019.2914929>
34. Chen, W., Chen, C., Tsai, C., & Yang, J. (2021). Deep learning for rice leaf blast disease detection using focal loss and ResNet. *Computers and Electronics in Agriculture*, 183, 106085. <https://doi.org/10.1016/j.compag.2021.106085>
35. Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of ICML*, 6105–6114.
36. Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. *Proceedings of ICLR*.
37. Howard, A., Sandler, M., Chu, G., et al. (2019). Searching for MobileNetV3. *Proceedings of IEEE/CVF ICCV*, 1314–1324. <https://doi.org/10.1109/ICCV.2019.00140>
38. Fuentes, A., Yoon, S., & Park, D.S. (2022). Deep learning for robust plant disease detection and classification. *Remote Sensing*, 14(12), 2791. <https://doi.org/10.3390/rs14122791>
39. Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-excitation networks. *Proceedings of IEEE CVPR*, 7132–7141. <https://doi.org/10.1109/CVPR.2018.00745>
40. Woo, S., Park, J., Lee, J.Y., & Kweon, I.S. (2018). CBAM: Convolutional block attention module. *Proceedings of ECCV*, 3–19. https://doi.org/10.1007/978-3-030-01234-2_1

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

