

Convolutional Neural Network and Transfer Learning to Automate Plant Pathology

Yadwinder Singh Sandhu

Independent researcher, California, USA

Yadwinder1944@gmail.com

Abstract—Plant diseases have still been a significant challenge to food security worldwide, and manual diagnosis by trained plant pathologists is time consuming, subjective and challenging to scale across both smallholder and commercial farms. This paper brings together the recent advances in the application of convolutional neural networks (CNNs) and transfer learning for automated plant disease recognition, across the plant disease benchmark, field-captured images, and crop-specific pipelines for apple, tomato, maize, rice, cassava and wheat. Existing deep architectures pre-trained on the ImageNet corpus such as AlexNet, GoogLeNet, VGGNet, ResNet and DenseNet are considered as feature extractors and then fine-tuned on comparatively small, domain-specific leaf-image collections, for which the classification accuracy is shown to increase from below ninety-two percent to above ninety-nine percent for equal sizes of samples when transfer learning is applied. The synthesis also examines the scale and variety of the datasets, real-time detection systems, attention-augmented residual networks, interpretability through saliency, and case studies of mobile deployments for cassava and wheat, which illustrate a frequently cited disparity between laboratory-trained accuracy (100%+), and field-trained accuracy (78%–91%). Across many quantitative measures, it is clear that deeper networks do not perform as well as residual and densely connected ones, that data augmentation and background variability have a significant impact on generalization, and that lightweight networks can achieve reasonable accuracy while providing on-device inference. Ethical, regulatory and adoption issues for agricultural extension services in resource-constrained environments are discussed. The results suggest that transfer learning is the most data and compute efficient route to becoming deployable plant pathology diagnostic tools and that further work needs to reduce the gap observed between lab and field performance.

Keywords—convolutional neural network; transfer learning; plant pathology; deep learning; image classification; PlantVillage; precision agriculture; disease detection; fine-tuning; mobile deployment; residual network; data augmentation

I. INTRODUCTION

It has been estimated that plant diseases account for a significant proportion of the harvested yield from crops globally each growing season, reducing food security at the household level, and limiting farm incomes to both subsistence and commercial farming systems. The traditional method of disease diagnosis relies on visual diagnosis by agronomists and/or extension officers, which is time-consuming, geographically limited and subject to inaccuracy due to the difference in symptom expression among different cultivars, growth stages and the presence of other stress factors. Image-based diagnosis became a scalable solution as mobile phone penetration grew in the 21st century and was the answer for those farmers who could not, or did not, always access expertise in plant pathology in a timely manner.

With the advent of convolutional neural network (CNN), image classification has undergone another change, and a deep network was able to significantly outperform handcrafted-feature methods on the large-scale ImageNet image recognition challenge. This success spurred a food revolution in which CNNs were re-trained to identify healthy leaves against a myriad of signs and symptoms caused by fungi, bacteria, viruses and pests [1]. Agricultural image data is generally orders of magnitude less than the ImageNet data, which consists over a million labeled images of more than one thousand object classes. Deep networks trained with random weights on a few thousand leaf images generally have bad generalizing ability and poor convergence. To overcome this issue, transfer learning methods are developed to use feature representations that are learned from a large source domain and fine-tuned to a smaller target domain, thus decreasing the amount of data and computation needed to achieve high diagnostic accuracy.

II. BACKGROUND

A. Convolutional Neural Network Foundations

A CNN learns "spatial features" layer by layer using a series of learned filters, nonlinear activations, and pooling to progressively abstract the features and lower the spatial resolution of the image. Early layers tend to be more responsive to edges, color gradients and texture, while later layers contain progressively more complex features, including the shape of lesions, distribution of chlorosis and vein discoloration, which are diagnostic for specific plant diseases [3]. The final convolutional block is flattened and fed into the one or more fully connected layers which are followed by a softmax function that gives a probability distribution on the candidate disease classes.

TABLE I. LANDMARK CNN BACKBONE ARCHITECTURES USED AS TRANSFER-LEARNING SOURCES

Architecture	Year Introduced	Depth (Layers)	Parameters (Millions)	ImageNet Top-5 Error (%)
AlexNet	2012	8	60.0	16.4
VGGNet-16	2014	16	138.0	7.3
GoogLeNet (Inception v1)	2014	22	6.8	6.7
ResNet-50/152	2015	50–152	25.6–60.2	3.6
DenseNet-121	2017	121	8.0	3.5

The representational capacity of a CNN is directly related to the depth and connectivity pattern of the CNN. In table I we list five reference backbone architectures, which typically are used as a transfer learning source in plant pathology literature, along with their approximate number of trainable parameters, and their classification error over the ImageNet benchmark. The classification error was also reduced, as shown in Fig. 1, with an early popular deep CNN model showing a 16.4 percent error rate, and a dense CNN model developed five years later showing a 3.5 percent error rate, which suggests that the emergence of new techniques like inception modules, residual connections and dense feature reuse addressed the vanishing-gradient problem in very deep networks [4].




Fig. 1. Reduction in ImageNet top-5 classification error across successive backbone architectures adopted as transfer-learning sources.

B. Principles of Transfer Learning

The formal definition of transfer learning is the ability to re-use the knowledge gained from solving a task at a source domain in order to enhance the solution of a related task at a target domain, with the knowledge at a target domain often having a smaller labeled sample and a distribution that is different from the source. In the area of plant pathology, the source domain is nearly always the ImageNet collection, while the target domain is usually comprised of images of leaves categorized as healthy or with a disease. There are two major transfer strategies mentioned in the literature [5][6]. Feature extraction involves freezing the pre-trained (conv) network and training a new classifier head that is attached to the frozen pre-trained network; this cheaply works well for very small datasets. Fine-tuning is also an update of some or all of the convolutional weights, with a lower learning rate, which leads to more accurate results for moderately sized target data sets but may lead to overfitting if the data set is small.

Fig. 2 depicts the general transfer-learning workflow applied throughout the reviewed studies: a backbone is pre-trained on the source domain, its convolutional layers are transferred and either frozen or selectively unfrozen, its classification head is replaced to match the number of target disease classes, and the resulting network is fine-tuned on the target agricultural dataset [7]. The cross-entropy loss

minimized during fine-tuning for a single training example with true label y and predicted probability vector p over C disease classes is given in (1):

$$L = -\sum_c = 1..C y_c \log(p_c) \quad (1)$$

Because the convolutional filters transferred from the source domain already encode general-purpose visual primitives, only the statistics associated with disease-specific texture and color patterns need to be relearned, which substantially lowers the number of gradient updates required for convergence relative to training an equivalent architecture from randomly initialized weights.




Fig. 2 An effective approach for plant leaf diseases classification based on a novel DeepPlantNet deep learning model (Frontiers , 2019)

III. BENCHMARK DATASETS AND DATA CONSIDERATIONS

The single most influential public resource in this literature is an open repository of photographs of leaves acquired in a laboratory against controlled backgrounds across fourteen crop species and twenty-six disease categories, which includes healthy leaf classes, consisting of 54,306 labelled images [8][9]. The images have been captured with fixed lighting and framing so models trained on this resource often report accuracy values well over 90% (later studies warned that this might overestimate real world performance as the network can learn to recognise background or capture artifacts, as well as the disease specific morphology).

TABLE II. BENCHMARK PLANT DISEASE IMAGE DATASETS REFERENCED ACROSS THE SYNTHESIZED STUDIES

Dataset / Source	Crop Species Covered	Disease Classes	Approx. Images	Acquisition Setting
Open leaf-image repository	14	26 + healthy	54,306	Controlled laboratory background
Field-collected tomato set	1 (tomato)	9 + healthy	5,000+	Greenhouse and open field
Mobile cassava image set	1 (cassava)	5 disease/pest classes	11,670	Smartphone, uncontrolled field
Mobile crop capture (wheat)	1 (wheat)	17 anomaly classes	8,000+	Handheld device, in the wild

Later works explored the impact of dataset size and diversity on the effectiveness of transfer learning. Results suggest that more sources of images – by increasing the acquisition conditions – result in better generalization than just more images from a homogeneous source, and that splitting individual lesions instead of showing whole-leaf images facilitates fine-grained discrimination between visually similar diseases. Table II shows a comparison of four benchmark collections, used in the studies reviewed, and Fig. 3, which is discussed later (in Section IV), shows how the accuracy of classification for the largest and most homogeneous dataset compares to the accuracy obtained from the more heterogeneous, fieldcollected imagery.




Fig. 4. Relationship between training dataset size and classification accuracy for networks trained from randomly initialized weights compared with transfer-learning fine-tuning.

IV. CROP-SPECIFIC APPLICATIONS

In addition to multi-crop benchmarks, there is a significant amount of research focused on specific crops of high food security or economic value. This section brings together the reported results for apple, tomato, maize, rice and cassava, each emphasizing a different type of methodology from architectural modification to real-time detection to mobile deployment.

A. Apple, Tomato, and Maize

Convolutional pipeline studies conducted on apple leaf disease have shown that the accuracy of the prediction is just under 97 percent for apple scab, black rot, cedar apple rust, and healthy tissues with the use of enhanced convolution pipelines that add extra convolution modules and global pooling. The related investigation transformed the problem of recognizing apple diseases into an object detection problem and developed a single shot detector that can be operated on mobile class hardware with a mean average precision close to seventy-nine percent with a processing rate of more than twenty frames a second, showing the feasibility to monitor the orchard besides single image classification.

Pre-trained AlexNet and VGG16 models fine-tuned on a set of nine different diseases showed an accuracy of around ninety-seven and three-tenths and ninety-seven and five-tenths percent, respectively, and a separate identification of salient sections of the image via the gradient visualization

showed that the fine-tuned model focused on the lesion areas and not on incidental cues, providing interpretability to the classification decision [10][11]. A second residual network was developed using attention to diagnose tomato leaf disease, with accuracy that was almost ninety-eight percent with fewer parameters than the baseline residual network, demonstrating that architectural refinement can improve the accuracy of diagnostics as well as saving on the computational burden. Detection-oriented research on tomato also integrated region-based and single-shot detectors together with deep feature extractors to detect diseases and pests in real-time, both at the region level and as a whole.

The use of a deep convolutional pipeline with targeted modification of the pipeline architecture for pooling and normalization layer to increase robustness to illumination change, a common cause of misclassification in maize images taken outdoors, provided near 98-99% accuracy in recognition of maize leaf diseases in both common blight and rust classes.

TABLE III. CROP-SPECIFIC CLASSIFICATION PERFORMANCE USING CNN AND TRANSFER-LEARNING TECHNIQUES

Crop	Primary Technique	Disease Classes	Reported Accuracy (%)	Acquisition Setting
Multi-crop (benchmark)	Fine-tuned deep CNN ensemble	26 + healthy	99.35	Laboratory
Apple	Improved CNN / SSD detector	4 (incl. healthy)	97.62 / 78.80 mAP	Laboratory / orchard
Tomato	Fine-tuned VGG16	9 + healthy	97.49	Laboratory
Maize	Improved deep CNN	3 diseases + healthy	98.90	Laboratory
Rice	Purpose-built deep CNN	Blast + healthy	95.83	Laboratory
Cassava	Transfer learning (mobile)	5 classes	80.60	Uncontrolled field

B. Rice and Cassava

The results were that, for the purpose-built CNN, the accuracy was about ninety-five and eight-tenths percent, and for a previous study of ten rice diseases with a comparatively shallow network it was about ninety-five and five-tenths percent, suggesting that rice disease imagery, which is often characterized by poor delineation of disease lesions, gains from deeper feature hierarchies without closing the accuracy gap found for higher-contrast diseases like apple scab.

Cassava disease detection is one of the most widely-tested case studies of field deployment in this literature. A preliminary study using transfer learning with Inception

based back- bone to leaflet images managed to get overall accuracy of close to ninety-three percent in semi-controlled settings. A follow-on study specifically developed for the deployment of smartphones was designed to use a light-weight backbone suitable for inference on-device, and the resulting application was directly validated in the field in cassava fields: an accuracy of around eighty and six-tenths percent was achieved, significantly lower than in laboratory benchmarks and demonstrating the practical challenge of maintaining classification performance once the illumination, occlusion, and background clutter are no longer controlled [12][13][14].




Fig. 3. Reported classification accuracy across multi-crop and single-crop plant disease recognition studies.

Table III summarizes the accuracy results and technique descriptions for the five crops discussed above and Fig. 3 shows the same accuracy results graphically with the multi-crop benchmarked result for easier comparison.

TABLE IV. REAL-TIME AND FIELD-DEPLOYABLE DETECTION FRAMEWORKS

Crop	Detector Type	Primary Metric	Reported Value	Processing Speed
Apple	Improved single-shot detector	Mean average precision	78.80%	~23 frames/second
Tomato	Region- and shot-based detectors	Mean average precision	Architecture-dependent	Near real time
Cassava	Lightweight mobile CNN	Field classification accuracy	80.60%	On-device, offline capable
Wheat	Mobile-capture residual CNN	Balanced accuracy (field)	~78%	Cloud-assisted

V. COMPARATIVE ARCHITECTURE PERFORMANCE

Five popular backbone architectures were used and trained on the same multi-crop benchmark, which helps to isolate the effect of architectural depth and connectivity from dataset variation: AlexNet, GoogLeNet, VGG16, ResNet50 and DenseNet121 [15]. The results (shown in Table V and Fig. 5) confirm that denser architectures and residual architectures are superior to shallower ones, and that DenseNet121 achieves the highest reported accuracy of

99.75%, as well as a precision of 99.60% and a recall of 99.50%.

The performance advantages of DenseNet121 can be largely explained by the fact that it has dense connectivity, that is, feature maps from every previous layer in a block are fed into the next layer, which helps to propagate gradients and re-use features, compared to the sequential connectivity of AlexNet and VGG16 [16][17]. The accuracy of ResNet50 is similar, albeit slightly worse, using the residual skip connections, while not requiring the parameter overhead that would accompany having full dense connectivity. To account for class imbalance among the rarer disease categories, overall accuracy was consistently reported, along with precision and recall, which are shown below in (2):

$$F1 = 2 \times \frac{(Precision \times Recall)}{(Precision + Recall)} \quad (2)$$

Furthermore, by interpreting the highest-performing architectures through saliency mapping using gradients, higher-accuracy models were shown to focus on lesion-bearing regions of the image, suggesting that the improvement in accuracy was due to architectural changes that resulted in more biologically meaningful representations of image features, rather than incidental improvements due to overfitting on the background structures.

TABLE V. COMPARATIVE PERFORMANCE OF FINE-TUNED CNN ARCHITECTURES ON THE MULTI-CROP BENCHMARK

Architecture	Accuracy (%)	Precision (%)	Recall (%)
AlexNet	98.00	97.20	96.80
GoogLeNet	98.40	97.90	97.50
VGG16	99.50	99.10	98.90
ResNet50	99.40	99.00	98.80
DenseNet121	99.75	99.60	99.50

VI. FROM LABORATORY TO FIELD: DEPLOYMENT CHALLENGES

A number of studies have confirmed models under uncontrolled field conditions with considerably lower accuracy, thereby exposing a gap in the model's ability to generalize, with direct application implications. Several studies have been conducted that validated models in the field under relatively un-controlled conditions, reporting significantly lower accuracy, thereby thus presenting a gap in the model's ability to generalize, with direct application implications. A wheat disease study, with a mobile capture device, was conducted on a balanced architecture with data coming from a curated laboratory image, mixed or cluttered background, and uncontrolled field image, with balanced accuracy of about ninety-six per cent, eighty-nine per cent, and seventy-eight per cent respectively [18][19]. In the case study of the cassava, the accuracy dropped from nearly 98

percent when the trained model was used on curated images, to 80.6 percent when it was embedded in a smartphone application and directly tested by farmers in cassava fields.

Table VI summarizes these comparisons between the laboratory and the field, and Fig. 6 represents the degradation of accuracy for the two crops for which more information is available about the field-versus-laboratory validation. The following factors are recurring across the studies reviewed: The illumination is variable or not controllable; The symptomatic tissue is covered by overlying foliage; Multiple simultaneous stress factors are present, creating ambiguous symptom combinations; There is a domain-shift between the smartphone camera sensors obtaining data and the model sensors used to train the model. Reported data augmentation techniques such as random rotation, colour jittering, and incorporating synthetic background substitution, were suggested to help partly address this disparity, as was the collection of further field-acquired training data with iterative model updates in an active-learning fashion.

The additional deployment factor for smartphone-based diagnosis in areas with scarce connectivity is model size and inference latency. As expected, lighter backbones incur a measurable drop in accuracy compared to deeper backbones like DenseNet121, showing that there is a well-known accuracy–efficiency compromise that needs to be addressed based on the deployment platform's hardware restrictions (such as device storage, processor power, and connectivity for cloud-based inference).

TABLE VI. LABORATORY VERSUS FIELD DEPLOYMENT ACCURACY COMPARISON

Crop / System	Laboratory Accuracy (%)	Field Accuracy (%)	Observed Gap (pp)
Wheat (mobile capture)	96.0	78.0	18.0
Cassava (mobile application)	98.0	80.6	17.4
Multi-crop benchmark (reference)	99.35	N/A (not field-tested)	N/A

VII. DISCUSSION

The evidence review suggests four trends are apparent. First, transfer learning provides the greatest relative advantage when the amount of labelled agricultural data is limited, which is representative of most domains for crop disease situations other than the largest public benchmark, and accuracy of transfer learning and training from randomly initialized weights becomes closer with larger datasets (up to tens of thousands of images). The second is that architectural depth and connectivity are important, but not to the extent that they increase by 10 times and the extra representational capacity results in a 2-percentage point improvement in

accuracy. Second, architectural depth and connectivity are important, but not as much as the extra representational capacity requires to gain 2 percentage points [20]. Third, variety of data sets (different acquisition devices, backgrounds, and geographical sources) seems to have a significant impact on field robustness on par with just the size of the data set. The fourth challenge is the lack of accuracy in the field, usually quoted as 15-20 percentage points for the most rigorously field-tested systems, which is the biggest hurdle to making the systems operational [21].

The results have implications for adoption and policy with respect to agricultural extension programs, especially in low-resource areas where smallholder farmers may have no other choice than a diagnosis via a smartphone when there is no available plant pathology expertise. However, the systems must also be deployed with a duty to communicate the level of model confidence and known limitations to avoid inappropriate pesticide use with corresponding economic cost and environmental consequence arising from misdiagnosis. Questions about data ownership, consent for the use of farm imagery for the retraining of shared models, and equitable access to diagnostic tools across regions that have different mobile network infrastructure also come into play when considering data governance for aggregation of farm-level imagery. Overall, the results of cost-benefit analysis across the synthesized studies suggest that automated screening can be used as a triage tool that frees up the valuable time of human experts from low-severity cases, and helps them focus on cases that are more ambiguous or severe in nature.




Fig. 5. Classification accuracy under curated laboratory conditions compared with mixed-background and uncontrolled field conditions for wheat and cassava mobile diagnostic systems.

VIII. CONCLUSION AND FUTURE DIRECTIONS

This synthesis has brought together and compiled evidence on the use of Convolutional Neural Networks and Transfer Learning for automated plant pathology diagnosis, including multi-crop benchmarks and crop-specific pipelines, comparisons across architectures, and field-tested mobile deployment. The results are in favor of transfer learning as the most data and compute efficient approach to high diagnostic performance on the small, diverse data sets typical of agricultural imagery, and show that densely connected and

residual architectures are repeatedly superior feature extractors to previous sequential CNN designs. Concurrently, laboratory and field accuracy consistently diverges, suggesting that benchmark performance is not a satisfactory basis for deployment.

Implications for future research directions from the synthesis of results include training sets beyond those collected in the field, geographically diverse training sets, incorporation of interpretability methods like saliency mapping into deployed systems to foster farmer trust, systematic study of the accuracy–efficiency trade-off for lightweight architectures suitable to low-connectivity regions, and establishing standardized field validation procedures that would enable more direct comparisons of reported accuracy between studies. The convergence in these directions is anticipated to further reduce the generalization gap observed and to move automated plant pathology diagnosis further from the laboratory to a reliable instrument in precision agriculture practice.

REFERENCES

- [1] J. G. A. Barbedo, "Factors influencing the use of deep learning for plant disease recognition," *Biosystems Engineering*, vol. 172, pp. 84–91, 2018.
- [2] J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," *Computers and Electronics in Agriculture*, vol. 153, pp. 46–53, 2018.
- [3] M. Brahimi, M. Arsenovic, S. Laraba, S. Sladojevic, K. Boukhalfa, and A. Moussaoui, "Deep learning for plant diseases: Detection and saliency map visualisation," in *Human and Machine Learning*, J. Zhou and F. Chen, Eds. Springer, 2018, pp. 93–117.
- [4] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
- [5] X. Zhang, Y. Qiao, F. Meng, C. Fan, and M. Zhang, "Identification of maize leaf diseases using improved deep convolutional neural networks," *IEEE Access*, vol. 6, pp. 30370–30377, 2018.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. 2016 IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [7] P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang, "Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks," *IEEE Access*, vol. 7, pp. 59069–59080, 2019.
- [8] R. Karthik, M. Hariharan, S. Anand, P. Mathikshara, A. Johnson, and R. Menaka, "Attention embedded residual CNN for disease detection in tomato leaves," *Applied Soft Computing*, vol. 86, Art. no. 105933, 2020.
- [9] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [10] W.-J. Liang, H. Zhang, G.-F. Zhang, and H.-X. Cao, "Rice blast disease recognition using a deep convolutional neural network," *Scientific Reports*, vol. 9, Art. no. 2869, 2019.
- [11] B. Liu, Y. Zhang, D. He, and Y. Li, "Identification of apple leaf diseases based on deep convolutional neural networks," *Symmetry*, vol. 10, no. 1, Art. no. 11, 2017.
- [12] Y. Lu, S. Yi, N. Zeng, Y. Liu, and Y. Zhang, "Identification of rice diseases using deep convolutional neural networks," *Neurocomputing*, vol. 267, pp. 378–384, 2017.
- [13] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, Art. no. 1419, 2016.
- [14] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [15] A. Picon, A. Alvarez-Gila, M. Seitz, A. Ortiz-Barredo, J. Echazarra, and A. Johannes, "Deep convolutional neural networks for mobile capture device-based crop disease classification in the wild," *Computers and Electronics in Agriculture*, vol. 161, pp. 280–290, 2019.
- [16] A. Ramcharan, K. Baranowski, P. McCloskey, B. Ahmed, J. Legg, and D. P. Hughes, "Deep learning for image-based cassava disease detection," *Frontiers in Plant Science*, vol. 8, Art. no. 1852, 2017.
- [17] A. Ramcharan, P. McCloskey, K. Baranowski, N. Mbilinyi, L. Mrisho, M. Ndalahwa, J. Legg, and D. P. Hughes, "A mobile-based deep learning model for cassava disease diagnosis," *Frontiers in Plant Science*, vol. 10, Art. no. 272, 2019.
- [18] A. K. Rangarajan, R. Purushothaman, and A. Ramesh, "Tomato crop disease classification using pre-trained deep learning algorithm," *Procedia Computer Science*, vol. 133, pp. 1040–1047, 2018.
- [19] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks based recognition of plant diseases by leaf image classification," *Computational Intelligence and Neuroscience*, vol. 2016, Art. no. 3289801, 2016.
- [20] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.