2020

Towards Deep Learning for Weed Detection: Deep Convolutional Neural Network Architectures for Plant Seedling Classification

Martinson Ofori
Dakota State University

Omar F. El-Gayar
Dakota State University

Towards Deep Learning for Weed Detection: Deep Convolutional Neural Network Architectures for Plant Seedling Classification

Conference Paper · July 2020

2 authors:

Martinson Ofori
Dakota State University

Omar F. El-Gayar
Dakota State University

Towards Deep Learning for Weed Detection: Deep Convolutional Neural Network Architectures for Plant Seedling Classification

Martinson Ofori
Dakota State University, martinson.ofori@trojans.dsu.edu

Omar El-Gayar
Dakota State University, omar.el-gayar@dsu.edu

Towards Deep Learning for Weed Detection: Deep Convolutional Neural Network Architectures for Plant Seedling Classification

Completed Research Full Paper

Martinson Ofori
Dakota State University
Martinson.Ofori@trojans.dsu.edu

Omar El-Gayar
Dakota State University
Omar.El-Gayar@dsu.edu

Abstract

Traditional means of on-farm weed control mostly relies on manual labor. This process is time consuming, costly and contributes to major yield losses. The conventional application of chemical weed control, however, goes against the strive for sustainability. To solve this using computer vision, precision agriculture researchers have used remote sensing weed maps, but this has been largely ineffective for early season weed control due to problems such as solar reflectance and cloud cover in satellite imagery. With the current advances in artificial intelligence, this study leverages the automatic feature extraction capabilities of deep convolutional neural networks (DCNN) to classify plant seedlings. In a comparative study, we demonstrate that DCNNs can successfully classify crops and weeds in various phenological growth stages. Our results indicate that while training DCNNs from scratch can achieve state-of-the-art performance for weed classification tasks, model performance can be improved by fine-tuning a pre-trained model.

Keywords

Artificial intelligence, deep learning, transfer learning, weed detection, plant seedling, classification, sustainable production, precision agriculture, smart farming

Introduction

Traditionally, weed control has been done using manual labor, a process which has proven to be time consuming, costly and a major contributor to yield losses (Gianessi 2009). However in the last 50 years, the go-to method for managing weeds has become the use of chemical weed killers such as herbicides and pesticides (Bastiaans et al. 2008). However, the continuous employment of chemical weed killers undermines the United Nations' Sustainable Development Goals (SDG). Specifically Goal 12 which targets sound management of chemical release to the environment and harnessing technology for more sustainable production (UN 2015). For example, the intensive use of chemicals like atrazine and alachlor, the main active ingredients in many commercial herbicide formulations, frequently leads to detection in water due to their persistence in the environment and their low biodegradability (Furtado et al. 2019; Tian et al. 1999).

In the last few decades, there has been a substantial acceptance of agricultural information technology (AIT) as part of everyday agricultural practices (Wang et al. 2019). It is this uptake in AITs that has resulted in the management practice known as Precision Agriculture (PA)[1]. PA allows the use of AITs to monitor inherent field conditions (soil, weather, etc.) and aim for better decision making that allows optimum profitability, increased sustainability, and environmental protection (Bongiovanni and Lowenberg-Deboer 2004). However, uniform applicators often used for chemical weed control does not minimize environmental effects as required in PA. For PA researchers meeting the SDG needs implies the practice of site-specific weed management (SSWM). SSWM adjusts herbicide application according to weed density and/or weed species. SSWM methods include remote and proximal (ground-based sensors and cameras)

[1] http://www.grap.udl.cat/en/presentation/pa_definitions.html

sensing (López-Granados 2011). López-Granados (2011) indicate the necessity of weed control at early stages of plant growth in their study. However, the use of remote sensing weed maps raises three issues: 1) the similar reflectance characteristics in early plant growth makes it difficult to detect small variations in reflectivity, 2) the need for high resolution images when weeds are distributed in small patches and 3) the reflectance of soil background interferes with detection (López-Granados 2011; Thorp and Tian 2004).

Proximal sensing has been posited as the best method for site-specific weed control especially ones that rely on ground-based computer vision (Gerhards 2010). The problem for such systems is that at early stages of plant growth, crops and weeds are almost indistinguishable to traditional pixel-based classification and as such, rule-based methods that rely on edge detection for leaf shape and texture recognition are used (Golzarian and Frick 2011). Feature extraction in such systems are also not robust enough to be generalized in different farming scenarios.

In recent times, the resurgence of artificial intelligence (AI) in the from Machine Learning (ML), which allows computers to learn without being explicitly programmed, has achieved phenomenal results in various problem domains. One such example is computer vision which has benefited from the application of deep Convolutional Neural Networks (DCNN) for image classification (Simonyan and Zisserman 2014). The availability of large hand-annotated datasets, and the ever-increasing speed and decreasing cost of infrastructure, has allowed the use of various DCNN architectures.

In instances where training data is limited, DCNNs can quickly overfit to training data. This means that the model learns the expected output of every data point rather than the general distribution of the data. In such instances, the model can predict from the training data with high accuracy but does poorly on any other data it has not seen before, i.e. test data. One method that has been proposed in various studies has been the use of pre-trained models applying a technique called Transfer learning (TL) (Yosinski et al. 2014). TL is useful in such instances as it allows the use of representations learned from previous data to solve problems in new datasets. Other methods of combatting overfitting include the use of image augmentation, optimizers and non-linearities

The overarching aim of this research paper, therefore, is to leverage the automatic feature extraction capabilities of DCNNs to classify plant seedlings. Theoretically, we intend to demonstrate through a performance comparison of various off-the-shelf DCNNs that they can successfully classify crops and weeds in various phenological growth stages. We also intend to identify limitations with these techniques that can further guide future research. Given the growing use of information systems (IS) to enhance sustainability across various sectors (Watson et al. 2008), we contribute to the growing body of literature on nascent technology use for sustainable practices in agriculture by demonstrating another area where artificial intelligence application can be beneficial. In effect, this paper shows that DCNNs are a viable alternative to current methods used for weed control in plant seedlings. This research will be relevant to researcher and producers of computer vision equipment, especially low-cost variable-rate technology (VRT) for ground-based site-specific weed control.

The focus of this study, therefore, is to evaluate and benchmark different off-the-shelf DCNN architectures using the Plant Seedlings dataset (Giselsson et al. 2017). Accordingly, the objectives of this paper are to perform a systematic evaluation using:

a) *Train from scratch with random weights initialization of the network.*
b) *Explore the use of transfer learning (TL) to improve generalizability under two conditions:*
   - *Model as a fixed feature extraction*
   - *Model fine-tuning*

Background

From a labor-intensive manual practice, agriculture has undergone a paradigm shift to become a technology-friendly practice that employs satellite technology, drones, robotics, big data, artificial intelligence, and other modern technology to provide selective rather than homogenous treatment for farm management (Aubert et al. 2012; El-Gayar and Ofori 2020). This practice – known as PA – has been discussed at length in past research as inextricably linked to sustainability. Bongiovanni and Lowenberg-Deboer (2004) in their review discuss this relationship at length and show that spatial management employed in PA can reduce environmental impact on sensitive areas while still maintaining profitability.

In past research, this idea of profitability has been touted as a major pull factor for farmers to adopt the PA (Wolfert et al. 2017). Even though some AITs in PA has been adopted just as fast as any other technology in history, the use of variable rate technology (VRT) – which provides the best sustainability gains – has rarely exceeded 20% of farms (Lowenberg-DeBoer and Erickson 2019). The criticism is that farmers may be convinced of the idea of VRTs but not their value. However, PA, associated technologies, and their role in sustainability and climate change remain a recurrent theme in news and social media (Lakshmi and Corbett 2020; Ofori and El-Gayar 2019). Therefore, it important to continuously improve these technologies and enhance their value.

One such technology that needs to be improved upon is machine vision technology for detecting and spot-spraying of weeds as an economic alternative to manual spraying and environmental-friendly alternative to uniform spraying. This is where the idea of artificial intelligence (or machine learning) for developing VRTs can be most beneficial. Specifically, the use of Deep Learning - a subset of ML that refer to *deep* algorithms designed to exploit unknown structures that exist in data. DL algorithms uncover representations at multiple levels, with higher level learned features expressed in terms of lower level features (Bengio 2013). In essence, by abstracting higher level representations, DL learns features in some input data by combining several lower level features that allows the mapping of input unto some output classification. While experts have not agreed on where shallow learning ends and deep learning begins (Schmidhuber 2015), it is generally accepted that an ML model is considered *deep* if there is at least one hidden layer between the input and output layer.

Neural Networks (NN) refer to the ML algorithms modeled after the human brain with the aim of recognizing patterns in data. There are several types of NNs – Feed-forward Neural Networks, Recurrent Neural Networks, Convolutional Neural Networks (CNN), to name a few. This study focuses on the application of CNNs, which refer to the DL technique aimed at processing data in the form of multiple arrays, such as images (often composed of three 2D arrays that represent the pixel intensities of the color channels present in the RGB color image) (LeCun et al. 2015). This means DL, or more specifically DCNNs, align nicely with the objective of machine vision by allowing computers to *see* and segregate the contents of digital color images. DCNN models have shown great promise in image classification. The authors of (Krizhevsky et al. 2012) were very influential in the argument for the application of DCNNs to images after they achieved a top-5 test error rate of 15.3% in the ImageNet LSVRC-2010 contest as compared to 26.2% achieved by the second-best entry.

In Precision Agriculture, weed detection using computer vision is not novelty. It is an active research area which has evaded researchers over the years. Methods proposed by earlier studies used spectral imaging techniques available on drones (Goel et al. 2002; Vioix et al. 2002). These methods are still most effective when spraying of an entire field is required rather than the site-specific, or crop-specific, needs of precision and smart agriculture. In their survey paper on deep learning applications in agriculture, the authors of (Kamilaris and Prenafeta-Boldú 2018) uncovered 8 important papers that used deep learning for weed detection. However, majority of the available studies (5 papers) still relied on imagery from satellite-based remote sensing, while the remaining used unmanned aerial vehicles (3 papers). Although the first 6 to 8 weeks of plant growth is the time when seedlings compete with weeds for most water and nutrients, most work in this area do not explore the use of plant seedlings for weed detection except for (Ashqar et al. 2019). They applied deep learning to the segmented images in the Plant Seedling Dataset proposed by (Giselsson et al. 2017) the VGG16 architecture but did not provide evidence of using the benchmark suggested by the original authors of the dataset. In this paper, we complement prior research by comparing the performance of popular DCNN architectures, that have demonstrated state-of-the-art performance on popular general image datasets such as the ImageNet and the CIFAR, on a dataset specific to plant seedlings.

Methods

Dataset

Giselsson et al. (2017) introduced the public image database for benchmarking plant seedling classification aimed at ground-based weed or specie spotting[2]. This dataset consists of 5,539 images of approximately

[2] https://vision.eng.au.dk/plant-seedlings-dataset/

960 unique plants belonging to 12 species at several growth stages. The plants were grown indoors in Styrofoam boxes and images were captured over a 20-day period. As overlapping plant leaves are minimal at the onset of plant growth, where most weed control such as broadcast spraying is undertaken, the images were captured in non-overlapping mode. Also, to avoid errors that may occur in pixel-based segmentation algorithms, plants were grown in soil which is covered in small stones. Figure 1 shows the samples, scientific names, counts and images from each class.

Figure 1. Plant Seedlings

This dataset is intended for researchers to perform object analysis, species recognition or plant appearance analysis without the hard task of image acquisition, segmentation, and annotation. The creators of the dataset also suggest a performance benchmark to permit easy replication of result and easy comparison of algorithm performance.

Analysis

As has been done in other supervised machine learning applications, the data is divided into 70% training, 15% validation and 15% test sets. All the neural network models are developed using the Keras library with TensorFlow backend (Abadi et al. 2016; Chollet and others 2015).

Experimental Setup

The experiments were run Google Colab which employs a Tesla K80 GPU having 2496 CUDA core, compute of 3.7, and 12GB GDDR5 VRAM. In all cases, we trained the models for 20 epochs with the mini-batch sizes of 32 image instances. Our initial learning was rate 0.0001 decreased by a factor of 0.5 after every 3 epochs when validation accuracy does not improve. We also applied the following preprocessing techniques:

- **Image resizing.** All images were resized to 128x128 pixels to ensure same aspect ratio.
- **Normalization of pixel values.** This was done to ensure that all the pixels have similar data distribution. Pixel normalization aids the convergence of neural networks.

Model Architectures

We choose models based on their availability in the TensorFlow Keras library and architectural properties as depicted in (Khan et al. 2019). A summary of the model architectures and performance validation on ImageNet is presented in Table 2:

- **Spatial Exploitation Based** – These kinds of networks take advantage of spatial filters to improve performance of the network. We use the **VGG** (Simonyan and Zisserman 2014), a popular DCNN network that replaced previous large filters with a smaller set of 3x3 filters and pushing depth to 16 and 19 layers. The VGG won second place in the ImageNet Challenge 2014 classification track. Specifically, we use the 16-layered network: VGG16.

- **Multi-Path Based** – To reduce the problem of performance degradation, gradient vanishing or explosion problems, these networks connect one layer to another by skipping some intermediate layers while still allowing flow of information across the layers through multiple paths or shortcut connections. Here we use the **DenseNet** (Huang et al. 2018). DenseNets connects each layer to every other layer in a feed-forward fashion such that feature-maps of all preceding layers are used as input to subsequent ones. The 121-layered DenseNet is employed in our tests.
- **Width Based Multi-Connection** – Instead of the traditional focus on the depth of a network, these models increase width to improve learning. The **Xception** network introduced by (Chollet 2017) does this by using depth-wise separable convolutions. It is an extreme version of the Inception network that maps the spatial correlations for each output channel separately, and then performs a pointwise convolution (1x1) to capture cross-channel correlation. The Xception is known to perform better than the Inception on ImageNet.
- **Depth and Multi-Path Based** – The **ResNet** (He et al. 2015) which won the of ImageNet 2015 challenge in image classification, detection, and localization, as well as Winner of MS COCO 2015 detection, and segmentation uses both depth and multiple connections. It is a very deep network that learns the residual representation functions instead of learning the signal representations directly. In this case we use the ResNet152V2, the 152-layered version of the network.
- **Depth Based** – The basic assumption for these networks is that the deeper the network, the better it performs as it improves feature representations. We employ the **Inception** (Szegedy, Vanhoucke, et al. 2015; Szegedy, Wei Liu, et al. 2015), a model introduced to increase the depth and width of a network while ensuring computational cost remains low. The InceptionV3 was the 1st runner up of the 2015 ImageNet image classification challenge. We choose the InceptionV3 as a representative of this category.

Model Architecture	Model Name	Depth	Number of Parameters	Top-1 Accuracy	Top-5 Accuracy
Spatial Exploitation Based	VGG16	23	138,357,544	0.713	0.901
Depth + Width Based	InceptionV3	159	23,851,784	0.779	0.937
Multi-Path Based	DenseNet121.	121	8,062,504	0.750	0.923
Width Based Multi-Connection	Xception	126	22,910,480	0.790	0.945
Depth + Multi-Path Based	ResNet152V2	152	60,380,648	0.780	0.942

Table 2. Architecture and performance of the pre-trained models

Model Evaluation

Two transfer learning scenarios are considered in this study:

- **Fixed feature extractor.** In this scenario, instead of retraining the entire network by initializing the network with random weights, we employ the pre-trained weights of each model and consider it as a fixed feature extractor. The network is then used as a contributor of feature vectors consisting of generic properties applicable to other datasets different from the one which it was originally trained for. We remove the classifier head (often the 1000-class ImageNet dataset) and replace with our own 12-class classifier representing the 12 species available in our dataset.
- **Fine-tuning.** This strategy allows retraining of higher-level portions of the network while keeping the lower levels which usually contains generic features useful for many tasks. The higher levels are chosen for retraining as they are known to contain feature more specific to the original dataset for which they were trained (Yosinski et al. 2014).

Model Evaluation

Model performance evaluation uses the proposed benchmarks suggested by the authors of the dataset (Giselsson et al. 2017), namely, Precision (P_C), Recall (R_C), and Mean Weighted Average f_1-scores.

Results

This section discusses the results of our experiments using the five pre-trained models (VGG16, InceptionV3, DenseNet121, Xception, and ResNet152V2).

Table 3 presents a detailed comparison of the models under different circumstances using the macro averages of the proposed evaluation metrics. Figure 2 shows the corresponding confusion matrices generated from the predictions on the test set by models trained from scratch and fine-tuned models. Figures 3 – 5 shows the training regimes.

Model	**Accuracy**			**Precision**			**Recall**			**f1-Score**		
	TS - RWI	**TL - FFE**	**TL - FT**	**TS - RWI**	**TL - FFE**	**TL - FT**	**TS - RWI**	**TL - FFE**	**TL - FT**	**TS - RWI**	**TL - FFE**	**TL - FT**
VGG16	0.9149	0.6703	0.9197	0.9082	0.6827	0.9101	0.9023	0.6147	0.9163	0.9044	0.6282	0.9126
InceptionV3	0.8981	0.6763	0.8981	0.8936	0.6726	0.8879	0.8731	0.6543	0.8844	0.8811	0.6618	0.8853
DenseNet121	0.9029	0.7350	0.9221	0.8980	0.7189	0.9156	0.8836	0.7065	0.9156	0.8881	0.7080	0.9149
Xception	0.9053	0.6811	0.8957	0.9067	0.6858	0.8979	0.9053	0.6436	0.8816	0.9042	0.6535	0.8873
ResNet152V2	0.8897	0.7422	0.9293	0.8870	0.7304	0.9182	0.8641	0.7422	0.9192	0.8709	0.7346	0.9173

Table 3. Performance evaluation on Plants Seedling Dataset

Training from scratch with random weights initialization

Per the objectives of this research, our first experiment involved training the network from scratch by initializing the network weights with random Gaussian distributions. The VGG16 achieves the best result on the test set in this round of training with an accuracy of 91.49%, followed by the Xception (90.53%), DenseNet121 (90.29%), InceptionV3 (89.81%) and the ResNet152V2 (88.97). Generally, this same result is echoed in the figures for precision, recall and f1-socres.

Transfer learning with model as a fixed feature extractor

Our second experiment used the ImageNet weights of pre-trained models as a fixed feature extractor. Training was generally smoother in this regime (as shown in Figure 4) but the models quickly overfitted to the training data and subsequently did poorly on the validation set (during training) and the test set (after training). In this case, the ResNet152V2 and the DenseNet121 achieves 74.22% and 73.5% respectively on the test set. The Xception network reaches 68.11%, the DenseNet121 reaches 67.63%, while the VGG16 which has achieved the best results when trained from scratch does worse than all four models with an accuracy of 67.03. While the remaining results followed a similar pattern, the VGG16 did better in terms of precision than the InceptionV3.

Transfer learning with fine-tuning

The third experiment involved fine tuning through gradual unfreezing of layers in the model. Although both the training accuracy increased in each case with more stability, the validation accuracy curve displayed a more erratic behavior. This behavior can however be attributed to the image augmentation used for training in this round. Without unfreezing more than half of the network, we found the following layers to be most optimal in each model:

- **VGG16**: block4_conv2
- **InceptionV3**: conv2d_51
- **DenseNet121**: conv4 _block 12 1 _conv
- **Xception**: block10_sepconv 1 _act
- **ResNet152V2**: conv4 _block 14 2 _conv

The ResNet152V2 achieved the best accuracy when fine-tuned (92.93%), an improvement of up to 4% over training from scratch. The DenseNet121 achieves the next best result with 92.21% also increasing accuracy by 2% over training from scratch. and 73.5% respectively on the test set. Similarly, the VGG16 increases accuracy by 1%. On the other hand, the InceptionV3 achieves same accuracy while the Xception dropped accuracy by 0.5%.

Discussion

The current study was motivated by the need to investigate the use of *new* and robust applications of AI to support sustainable development in agriculture. Undoubtedly, the results above demonstrate the potential suitability of using DCNN models for Plant Seedling Classification. Although, we have found that models performed poorly when using the ImageNet weights as a fixed-feature extractor, this could be attributed to the fundamental differences between the feature representations of the two image datasets – ImageNet consists of general images while the Plant Seedling dataset consists of only plant images. It can be observed that training models from scratch achieved decent results but performing fine-tuning by freezing part of the network improved the classification results. Surprisingly, the ResNet152V2 model performs best overall compared to the other four DCNNs. This is interesting given that it achieved the lowest accuracy in our initial experiment when the models were trained from scratch.

Overall, we find that our result is better than the average performance (Accuracy < 87% or f_1-scores < 0.8) exhibited by weed detection and crop type classification tasks (Kamilaris and Prenafeta-Boldú 2018). Given the observation that the two highest misclassifications occurred between Black-grass and Loose Silky-bent (both weeds). And the objective to distinguish food crops from weeds, then in a food crop-or-weed scenario, most of the networks would have easily achieved above 99% accuracy. So, while we wait for better chemical agents with higher biodegradability and lower environmental persistence, our result paves the way for employing further digital transformation which leverages new technology capabilities in a bid to ensure sustainable development in PA. Partnerships such as the John Deere and NVIDIA has for commercializing AI-enabled precision spraying technology could be the key to unlocking a new age for PA where chemicals are only applied in a *see-and-spray* manner.

Additionally, even though this research has been framed in the context of chemical weed killers in agriculture and environmental sustainability, the success of fine-tuning a pre-trained model using the transfer learning approach is a huge milestone for Green IS. Often the drawback for deep learning and other machine learning tasks is their requirement for huge amounts of data for training which has a direct impact on both energy consumption and computing power of the infrastructure involved. The fact that transfer learning in this case outperforms training from scratch means very little time will be spent training an accurate model for VRTs such as precision spraying equipment. Similarly, with just 8M trainable parameters, the performance of the DenseNet architecture (second best in the fine-tuned transfer learning approach) means even less computing power will be required if this model is adopted compared to the ResNet model (60M) or even the VGG model (138M).

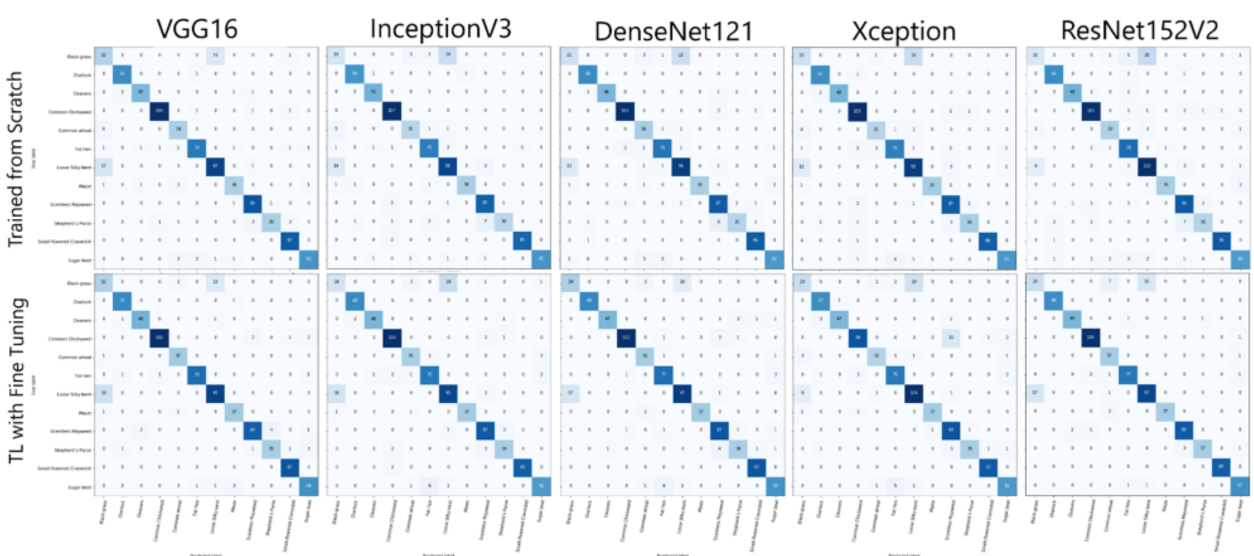

Figure 2. Confusion matrices

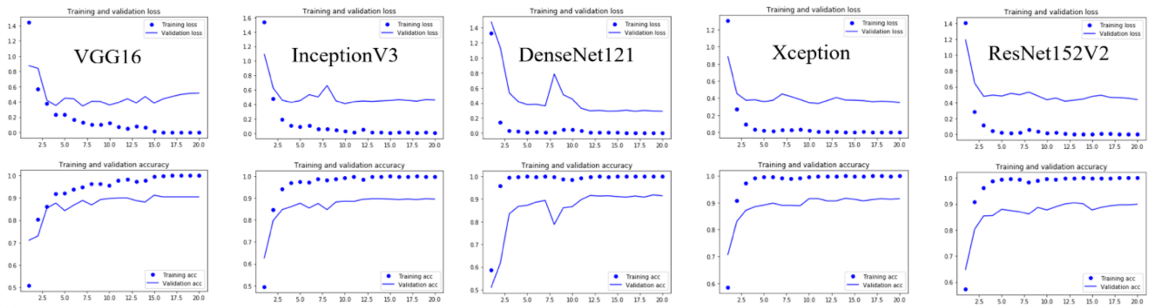

Figure 3. Training from scratch with random weights initialization

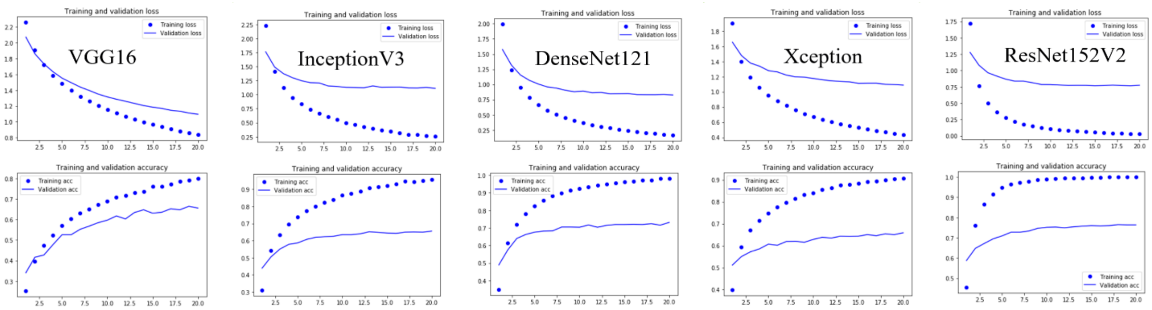

Figure 4. Transfer learning with model as a fixed feature extractor

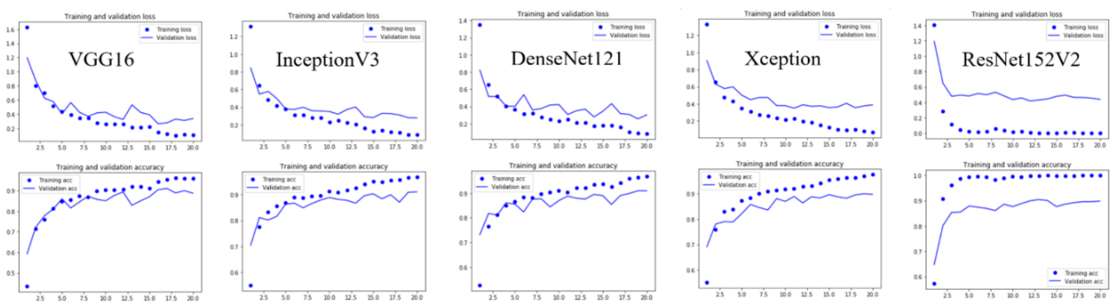

Figure 5. Transfer learning with fine-tuning

Conclusion

In this report we investigated how different deep learning training schemes influence classification of weeds in non-segmented plant seedlings. We conducted three different experiments: training a model from scratch with random weight initialization, using pre-trained models as fixed-feature extractors and sequentially fine-tuning the model by unfreezing parts of the pre-trained model We conclude that DCNN models that have achieved state-of-the-art performance in image classification can be applicable to Precision Agriculture for classifying plant seedlings and, by extension, distinguish weeds from food crops. While using the transfer learning scenario which involved models trained on ImageNet as a fixed feature extractor did not achieve great results, we found that fine-tuning the higher levels of the networks lead to the best results. The relative performance of the models revealed that fine-tuning ResNet152V2, DenseNet121, and VGG16 resulted in higher performance compared to training from scratch while the InceptionV3 and Xception does better when initialized with random weights. Our result indicates the readiness of employing AI-enabled computer vision systems for precision treatment. Such systems will go a long way to address food production and sustainability, and potentially battle the effects of climate change caused by interactions between agricultural chemical agents, soil, water, energy, and climate.

We suggest that future researchers investigate how the results presented in this study may be improved using multi-stage transfer learning which transfers weights from an intermediary dataset much closer to the target dataset. Also, novel models inspired by the DenseNet or ResNet architectures, as well as lightweight architectures like the MobileNets, may achieve improved the performance and should be investigated (Ashqar et al. 2019). For the IS – and Green IS – communities, the proposition that farmers may not be convinced of the value of such systems, and that often their experience differs from the intended use by manufactures and researchers (Lowenberg-DeBoer et al. 2019; Lowenberg-DeBoer and Erickson 2019), suggests that there may be other underlying issues beside the accuracy and cost of the systems that should be investigated. We recommend that further research further explore issues related to AI adoption in PA such as the deployment and use of AI applications from a farmer and technology-provider perspective to support sustainability and food production.

REFERENCES

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mane, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viegas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. 2016. “TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems,” *ArXiv:1603.04467 [Cs]*. (http://arxiv.org/abs/1603.04467).

Ashqar, B. A. M., Abu-Nasser, B. S., and Abu-Naser, S. S. 2019. *Plant Seedlings Classification Using Deep Learning*, (3:1), p. 8.

Aubert, B. A., Schroeder, A., and Grimaudo, J. 2012. “IT as Enabler of Sustainable Farming: An Empirical Analysis of Farmers’ Adoption Decision of Precision Agriculture Technology,” *Decision Support Systems* (54:1), pp. 510–520. (https://doi.org/10.1016/j.dss.2012.07.002).

Bastiaans, L., Paolini, R., and Baumann, D. T. 2008. “Focus on Ecological Weed Management: What Is Hindering Adoption?,” *Weed Research* (48:6), pp. 481–491. (https://doi.org/10.1111/j.1365-3180.2008.00662.x).

Bengio, Y. 2013. “Deep Learning of Representations: Looking Forward,” in *Statistical Language and Speech Processing* (Vol. 7978), A.-H. Dediu, C. Martín-Vide, R. Mitkov, and B. Truthe (eds.), Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 1–37. (https://doi.org/10.1007/978-3-642-39593-2_1).

Bongiovanni, R., and Lowenberg-Deboer, J. 2004. “Precision Agriculture and Sustainability,” *Precision Agriculture* (5:4), pp. 359–387. (https://doi.org/10.1023/B:PRAG.0000040806.39604.aa).

Chollet, F. 2017. “Xception: Deep Learning with Depthwise Separable Convolutions,” *ArXiv:1610.02357 [Cs]*. (http://arxiv.org/abs/1610.02357).

Chollet, F., and others. 2015. *Keras*. (https://keras.io).

El-Gayar, O., and Ofori, M. 2020. “Disrupting Agriculture: The Status and Prospects for AI and Big Data in Smart Agriculture,” in *AI and Big Data’s Potential for Disruptive Innovation:*, Advances in Computational Intelligence and Robotics, M. Strydom and S. Buckley (eds.), IGI Global. (https://doi.org/10.4018/978-1-5225-9687-5).

Furtado, R. X. de S., Azevedo, E. B., and Motheo, A. D. J. 2019. “Electrochemical Degradation of Aqueous Alachlor and Atrazine: Products Identification, Lipophilicity, and Ecotoxicity,” *Eclética Química Journal* (44:1SI), p. 12. (https://doi.org/10.26850/1678-4618eqj.v44.1SI.2019.p12-25).

Gerhards, R. 2010. “Spatial and Temporal Dynamics of Weed Populations,” in *Precision Crop Protection - the Challenge and Use of Heterogeneity*, E.-C. Oerke, R. Gerhards, G. Menz, and R. A. Sikora (eds.), Dordrecht: Springer Netherlands, pp. 17–25. (https://doi.org/10.1007/978-90-481-9277-9_2).

Gianessi, L. P. 2009. “Solving Africa’s Weed Problem: Increasing Crop Production & Improving the Lives of Women.,” *Aspects of Applied Biology* (No.96), pp. 9–23.

Giselsson, T. M., Jørgensen, R. N., Jensen, P. K., Dyrmann, M., and Midtiby, H. S. 2017. “A Public Image Database for Benchmark of Plant Seedling Classification Algorithms,” *ArXiv:1711.05458 [Cs]*. (http://arxiv.org/abs/1711.05458).

Goel, P. K., Prasher, S. O., Patel, R. M., Smith, D. L., and DiTommaso, A. 2002. “Use of Airborne Multi– Spectral Imagery for Weed Detection in Field Crops,” *TRANSACTIONS OF THE ASAE* (45), p. 9.

Golzarian, M. R., and Frick, R. A. 2011. “Classification of Images of Wheat, Ryegrass and Brome Grass Species at Early Growth Stages Using Principal Component Analysis,” *Plant Methods* (7:1), p. 28. (https://doi.org/10.1186/1746-4811-7-28).

He, K., Zhang, X., Ren, S., and Sun, J. 2015. “Deep Residual Learning for Image Recognition,” *ArXiv:1512.03385 [Cs]*. (http://arxiv.org/abs/1512.03385).

Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K. Q. 2018. "Densely Connected Convolutional Networks," *ArXiv:1608.06993 [Cs]*. (http://arxiv.org/abs/1608.06993).

Kamilaris, A., and Prenafeta-Boldú, F. X. 2018. "Deep Learning in Agriculture: A Survey," *Computers and Electronics in Agriculture* (147), pp. 70–90. (https://doi.org/10.1016/j.compag.2018.02.016).

Khan, A., Sohail, A., Umme Zahoora, and Qureshi, A. S. 2019. *A Survey of the Recent Architectures of Deep Convolutional Neural Networks*. (https://doi.org/arXiv:1901.06032).

Krizhevsky, A., Sutskever, I., and Hinton, G. E. 2012. "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (eds.), Curran Associates, Inc., pp. 1097–1105. (http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf).

Lakshmi, V., and Corbett, J. 2020. *How Artificial Intelligence Improves Agricultural Productivity and Sustainability: A Global Thematic Analysis*, , January 7. (https://doi.org/10.24251/HICSS.2020.639).

LeCun, Y., Bengio, Y., and Hinton, G. 2015. "Deep Learning," *Nature* (521:7553), pp. 436–444. (https://doi.org/10.1038/nature14539).

López-Granados, F. 2011. "Weed Detection for Site-Specific Weed Management: Mapping and Real-Time Approaches: Weed Detection for Site-Specific Weed Management," *Weed Research* (51:1), pp. 1–11. (https://doi.org/10.1111/j.1365-3180.2010.00829.x).

Lowenberg-DeBoer, J., and Erickson, B. 2019. "Setting the Record Straight on Precision Agriculture Adoption," *Agronomy Journal* (111:4), p. 1552. (https://doi.org/10.2134/agronj2018.12.0779).

Lowenberg-DeBoer, J., Huang, I. Y., Grigoriadis, V., and Blackmore, S. 2019. "Economics of Robots and Automation in Field Crop Production," *Precision Agriculture*. (https://doi.org/10.1007/s11119-019-09667-5).

Ofori, M., and El-Gayar, O. 2019. "The State and Future of Smart Agriculture: Insights from Mining Social Media," in *2019 IEEE International Conference on Big Data (Big Data)*, Los Angeles, CA, USA: IEEE, December, pp. 5152–5161. (https://doi.org/10.1109/BigData47090.2019.9006587).

Schmidhuber, J. 2015. "Deep Learning in Neural Networks: An Overview," *Neural Networks* (61), pp. 85–117. (https://doi.org/10.1016/j.neunet.2014.09.003).

Simonyan, K., and Zisserman, A. 2014. "Very Deep Convolutional Networks for Large-Scale Image Recognition," *ArXiv:1409.1556 [Cs]*. (http://arxiv.org/abs/1409.1556).

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. 2015. "Rethinking the Inception Architecture for Computer Vision," *ArXiv:1512.00567 [Cs]*. (http://arxiv.org/abs/1512.00567).

Szegedy, C., Wei Liu, Yangqing Jia, Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. 2015. "Going Deeper with Convolutions," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, USA: IEEE, June, pp. 1–9. (https://doi.org/10.1109/CVPR.2015.7298594).

Thorp, K. R., and Tian, L. F. 2004. "A Review on Remote Sensing of Weeds in Agriculture," *Precision Agriculture* (5:5), pp. 477–508. (https://doi.org/10.1007/s11119-004-5321-1).

Tian, L., Reid, J. F., and Hummel, J. W. 1999. "Development of a Precision Sprayer for Site-Specific Weed Management," *Transactions of the ASAE* (42:4), pp. 893–900. (https://doi.org/10.13031/2013.13269).

UN. 2015. "Sustainable Consumption and Production," , September 10. (https://www.un.org/sustainabledevelopment/sustainable-consumption-production/, accessed February 24, 2020).

Vioix, J.-B., Douzals, J.-P., Truchetet, F., Assémat, L., and Guillemin, J.-P. 2002. "Spatial and Spectral Methods for Weed Detection and Localization," *EURASIP Journal on Advances in Signal Processing* (2002:7), p. 793080. (https://doi.org/10.1155/S1110865702204072).

Wang, Y., Jin, L., and Mao, H. 2019. "Farmer Cooperatives' Intention to Adopt Agricultural Information Technology—Mediating Effects of Attitude," *Information Systems Frontiers* (21:3), pp. 565–580. (https://doi.org/10.1007/s10796-019-09909-x).

Watson, R., Boudreau, M.-C., Chen, A., and Huber, M. 2008. "Green IS: Building Sustainable Business Practices," *Information Systems: A Global Text*, Athens, GA: Global Text Project, pp. 1–17.

Wolfert, S., Ge, L., Verdouw, C., and Bogaardt, M.-J. 2017. "Big Data in Smart Farming – A Review," *Agricultural Systems* (153), pp. 69–80. (https://doi.org/10.1016/j.agsy.2017.01.023).

Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. 2014. "How Transferable Are Features in Deep Neural Networks?," *ArXiv:1411.1792 [Cs]*. (http://arxiv.org/abs/1411.1792).