Unearthing Crop Insights: Plant Seedling Classification Using Convolutional Neural Networks

Madison Hall & Jackie Ennis
{mhall, jqennis}@stanford.edu

Link to Video: https://www.youtube.com/watch?v=NKfKXmBQMG4

Introduction

Agriculture faces unprecedented sustainability and productivity challenges. Precision agriculture - automatically managing weeds among crops - could reduce the cost and environmental impact of herbicidal treatments. Our approach proposes a convolutional neural network with weights pre-trained on ImageNet to classify 12 species of crops and weeds at critical early growth stages. The model takes in a 64x64 image of a crop seedling and predicts its species among 12 classes, achieving

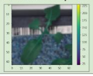

[0: "Charlock"] → a weed!

Dataset

Our dataset contains 5,539 labelled color PNG images of 960 individual plants at various growth stages, constituting 12 different species. This dataset was recorded at Aarhus University Flakkebjerg Research station and provided by a collaboration between University of Southern Denmark and Aarhus University (Giselsson et al., 2017). We resize our images to 64x64, normalize the data using feature-scaling, and performing data augmentation by applying various image transformations.

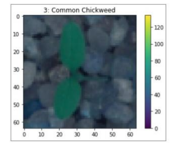

Figure 1: example image labeled "Common Chickweed"

Model

Baseline

For our baseline, we trained a model composed of two 2D convolutional layers, accompanied by a 2D max pooling to reduce dimensionality, and 10% drop out, to prevent overfitting. We compiled the model using the ADAM Optimization algorithm. The model achieved 70% accuracy on the training set, 73% on the testing set.

Activation Functions

- ReLU: $A = RELU(Z) = max(0, Z)$
- Softmax: $\sigma(\mathbf{z})_j = \frac{e^{z_j}}{\sum_{k=1}^{K} e^{z_k}}$

Loss Function

- Categorical cross-entropy loss (commonly used in multiclass classification:)

$$J = -\frac{1}{m}\sum_{i=0}^{m}\left(y^{(i)}\log\left(a^{[2](i)}\right) + (1-y^{(i)})\log\left(1-a^{[2](i)}\right)\right)$$

Pre-Trained Networks on ImageNet

We decide to leverage pre-trained architectures with weights trained on ImageNet to improve our model's performance despite the small size of our dataset. We use transfer learning to use each pre-trained network as a base for our model. We vary the number of frozen layers, finally settling on freezing all but the bottom layer for the best performance. We add a dropout rate of 0.4 to reduce overfitting.

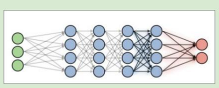

Visualization of freezing all but last layers of a pre-trained network (Amidi & Amidi, 2018)

DenseNet121

DenseNet121 is a 121-layer deep Densely Connected Convolutional Networks with feed-forward layers.

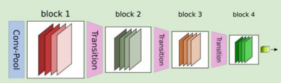

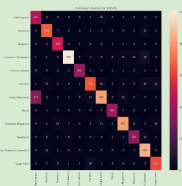

VGG16

VGG16 is a 16-layer deep convolutional network designed for image detection that uses small, 3×3 convolutional layers and 2x2 max-pooling layers.

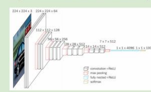

ResNet50

ResNet is a 50-layer Deep Residual Learning designed for image recognition (He et al, 2016). It is well-known for its robustness against the accuracy degradation problem when deepermodels converge.

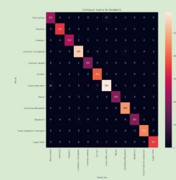

Results & Discussion

	Train Accuracy	Test Accuracy	Error	F1 Score
Baseline	70.05%	73.62%	0.0083	74.85
DenseNet	62.46%	56.94%	0.6152	56.94
VGG	63.24%	59.44%	0.8731	59.21
ResNet	90.82%	56.91%	0.6711	56.48

- Number of images in train set: 4,432
- Number of images in test set: 1,107

Analysis of Results

The baseline achieved the highest F1 score on the test data. We are surprised that our model from scratch outperformed the pre-trained networks, but posit that there was too little data for the larger networks to distinguish between almost identical looking leaves, and that our dataset is not similar enough to ImageNet to garner the best results.

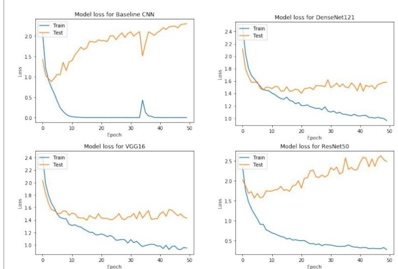

Future Work

We would perform more extensive hyperparameter tuning to improve the performance of the pre-trained networks. We are also interested in using texture analysis of the leaves as feature extraction.

References

Giselsson, T. M., Jørgensen, R. N., Jensen, P. K., Dyrmann, M., \& Midtiby, H. S. (2017). A public image database for benchmark of plant seedling classification algorithms. arXiv preprint arXiv:1711.05458

Huang, G., Liu, Z., Van Der Maaten, L., \& Weinberger, K. Q. (2017). Densely connected convolutional networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 4700-4708).