Towards Establishing a Predictive Machine Learning Model in Agriculture applying Convolutional Neural Networks

Bohdan Junior bohdanjr@stanford.edu
Joshua Adams adamsdsc@stanford.edu

Introduction

Machine learning is playing an increasingly important role in precision agriculture. We conducted a study using a collection of 3-channel images with 500 x 500 (910 total images) to predict the most possible classification of carbon levels in soil samples from its cellphone images. Building upon previous experiments, we conducted new experiments with deep learning CNN models (AlexNet, VGG, and Resnet). There was improvement of accuracy from 50% to 75% in more recent models.

Data

Soil samples were collected at five sites in south of Brazil and photos taken by cell phone in a controlled light environment. C-Levels converted gases were measured in lab by samples combustion. Number of samples = 910.

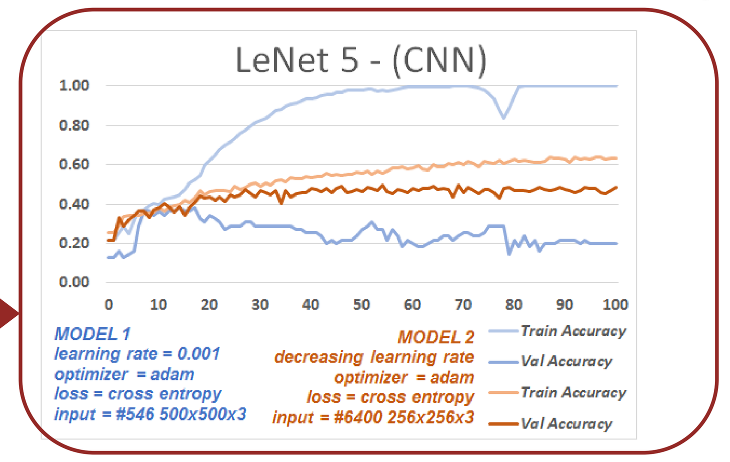

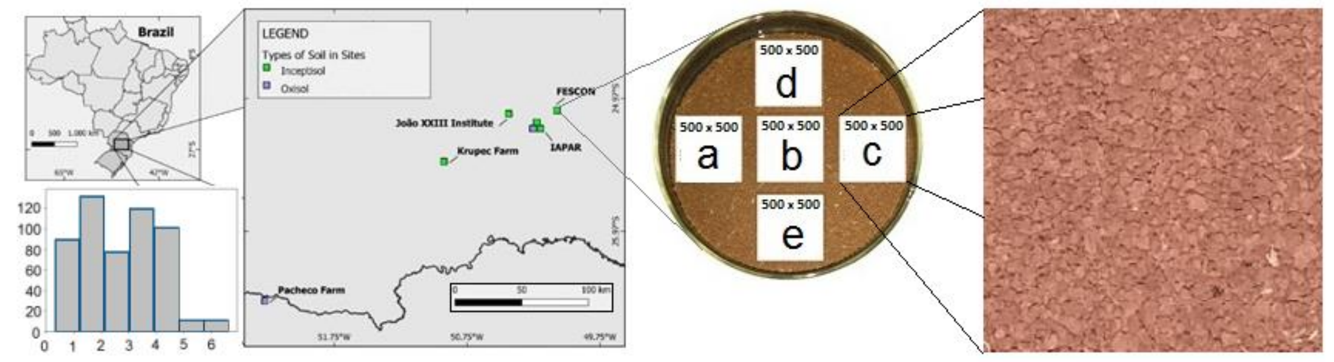

Features

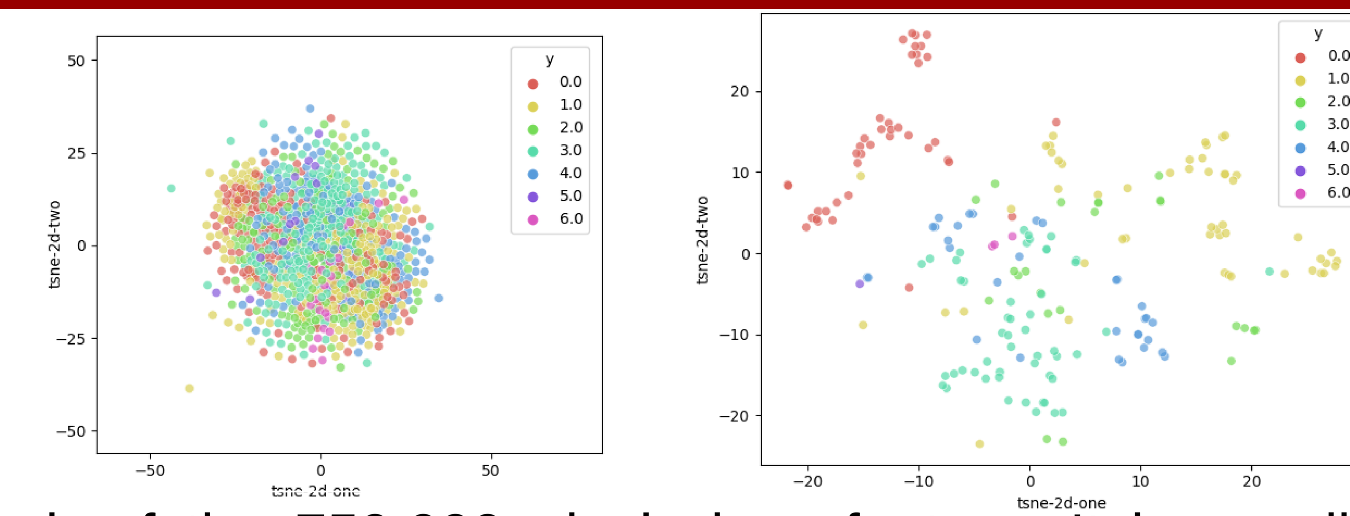

Each of the 750,000 pixels is a feature. It is possible to visualize this feature space in 2D using t-sne (left), without a clear separation of the classes. After training an AlexNet network, it is possible to visualize its last layer with 4,096 features, showing clusters for each of the classes.

Models and Results

Model	Architecture	Training curve	Confusion Matrix	Accuracy	Precision/Recall
LeNet-5	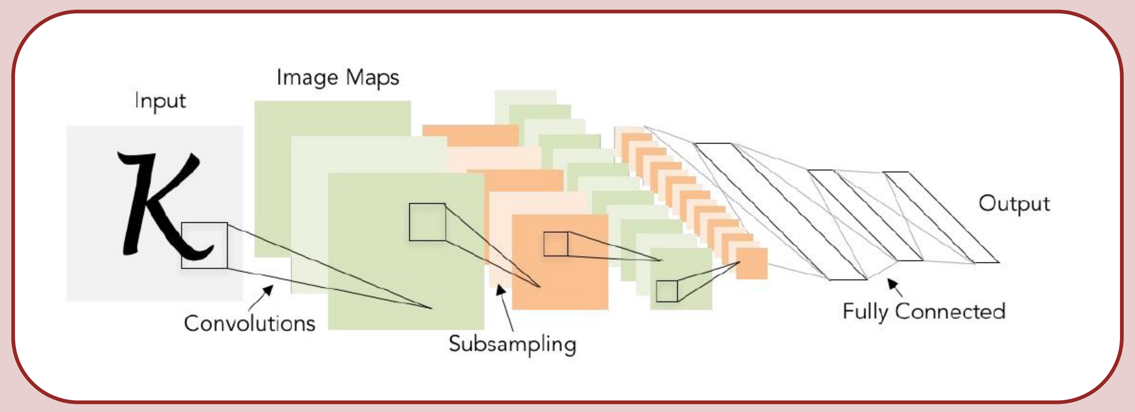	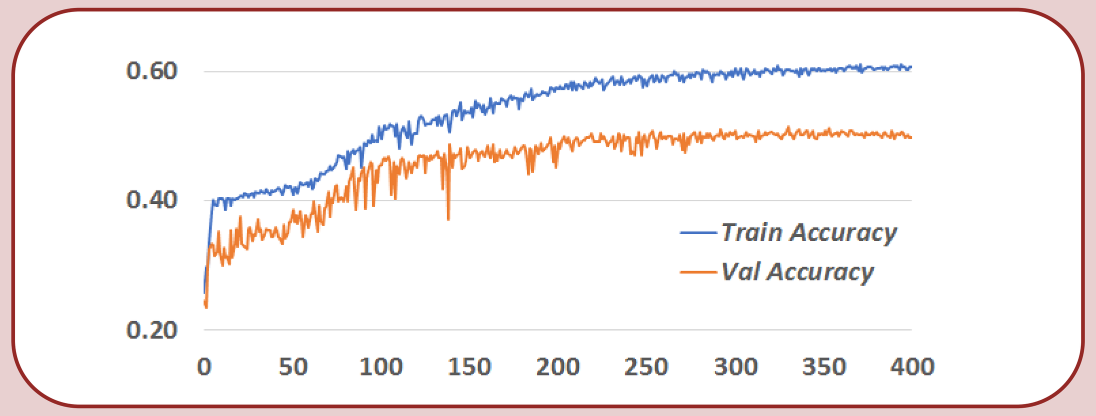	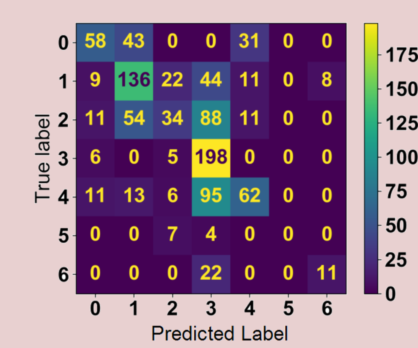	train: 0.6113 dev: 0.5136 test: 0.4990	Class 0: 0.61; 0.44 Class 1: 0.55; 0.59 Class 2: 0.46; 0.17 Class 3: 0.44; 0.95 Class 4: 0.54; 0.33 Class 5: 0.00; 0.00 Class 6: 0.58; 0.33
AlexNet	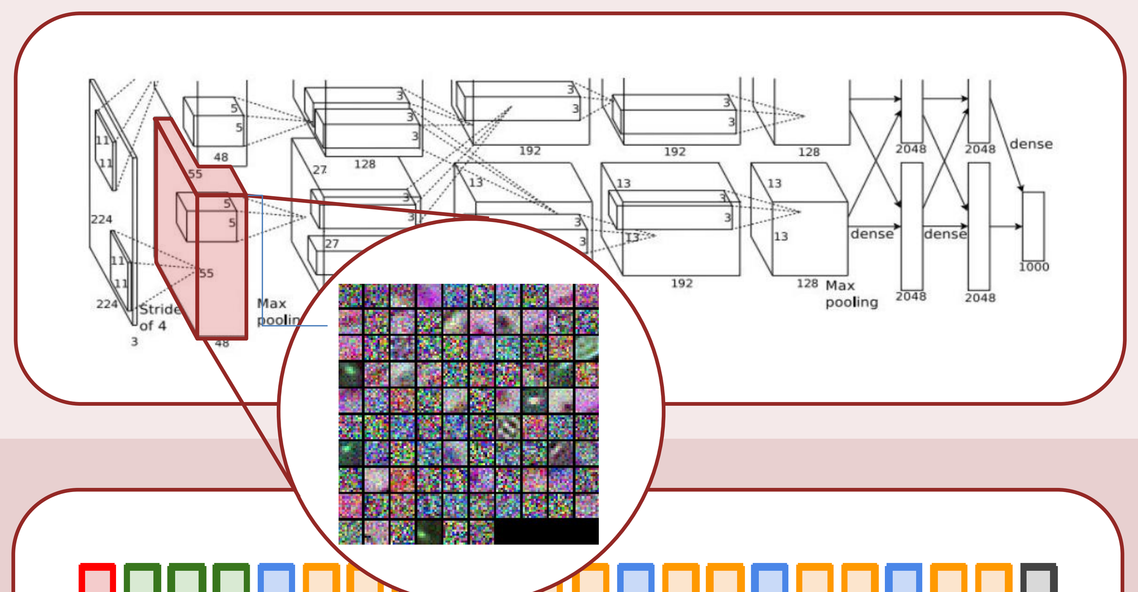	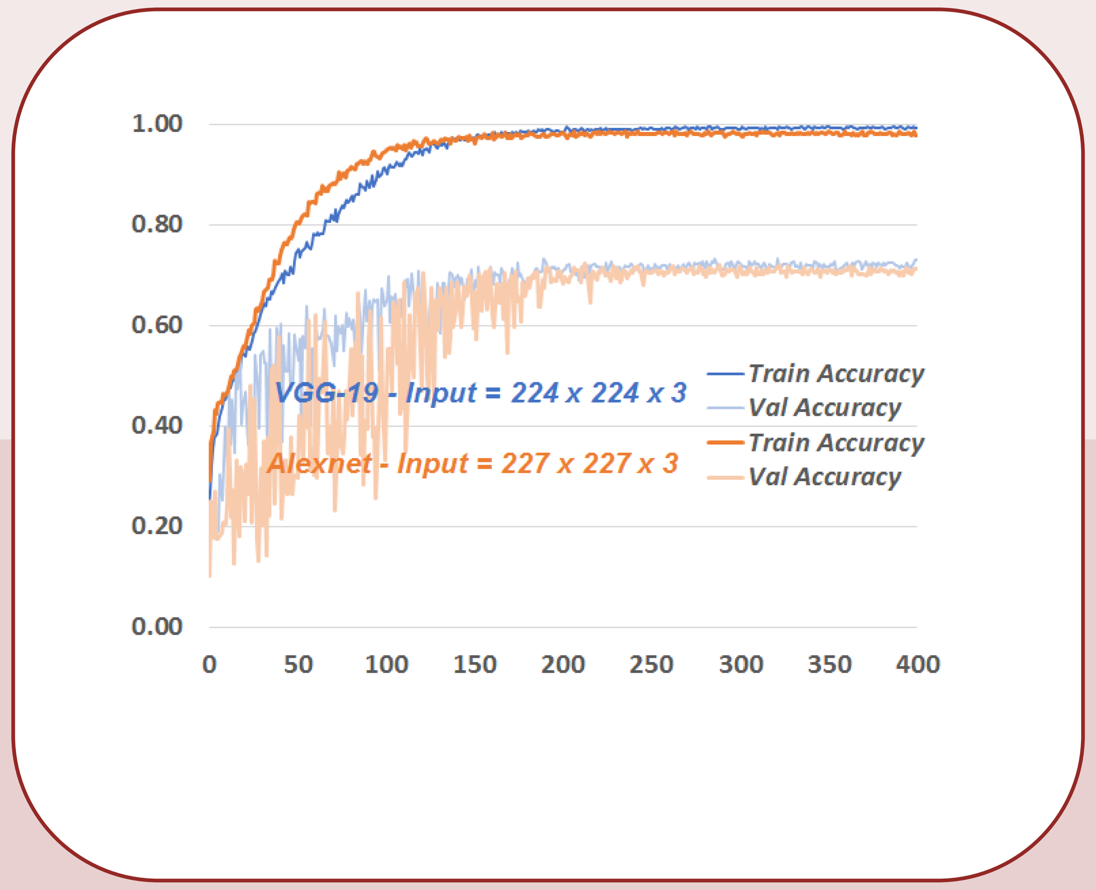	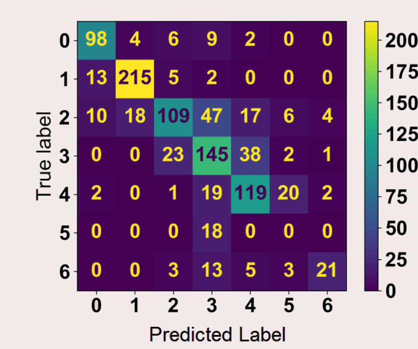	train: 0.9880 dev: 0.7247 test: 0.7070	Class 0: 0.80; 0.82 Class 1: 0.91; 0.91 Class 2: 0.74; 0.52 Class 3: 0.57; 0.69 Class 4: 0.73; 0.69 Class 5: 0.00; 0.00 Class 6: 0.47; 0.58
VGG-19	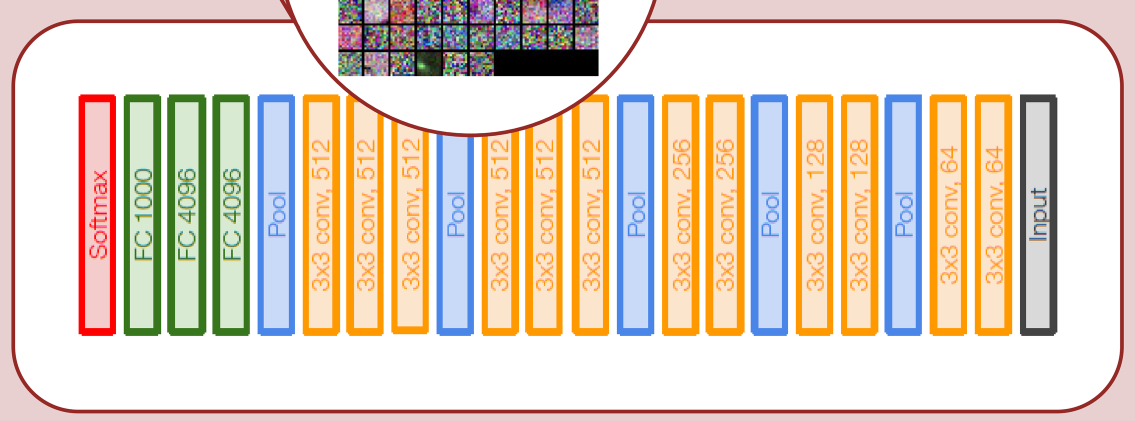		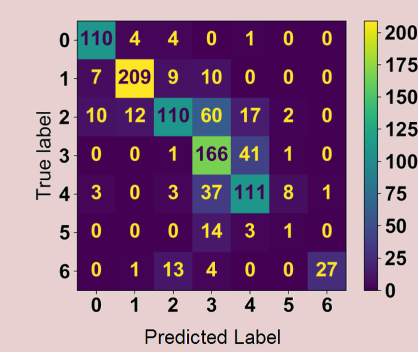	train: 0.9963 dev: 0.7333 test: 0.7340	Class 0: 0.85; 0.92 Class 1: 0.92; 0.89 Class 2: 0.79; 0.52 Class 3: 0.57; 0.79 Class 4: 0.64; 0.68 Class 5: 0.08; 0.06 Class 6: 0.96; 0.60
ResNet 50	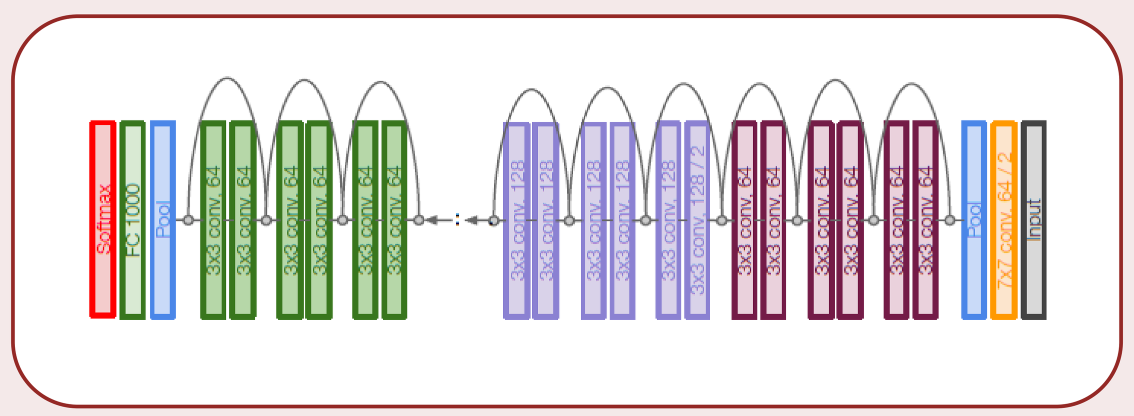	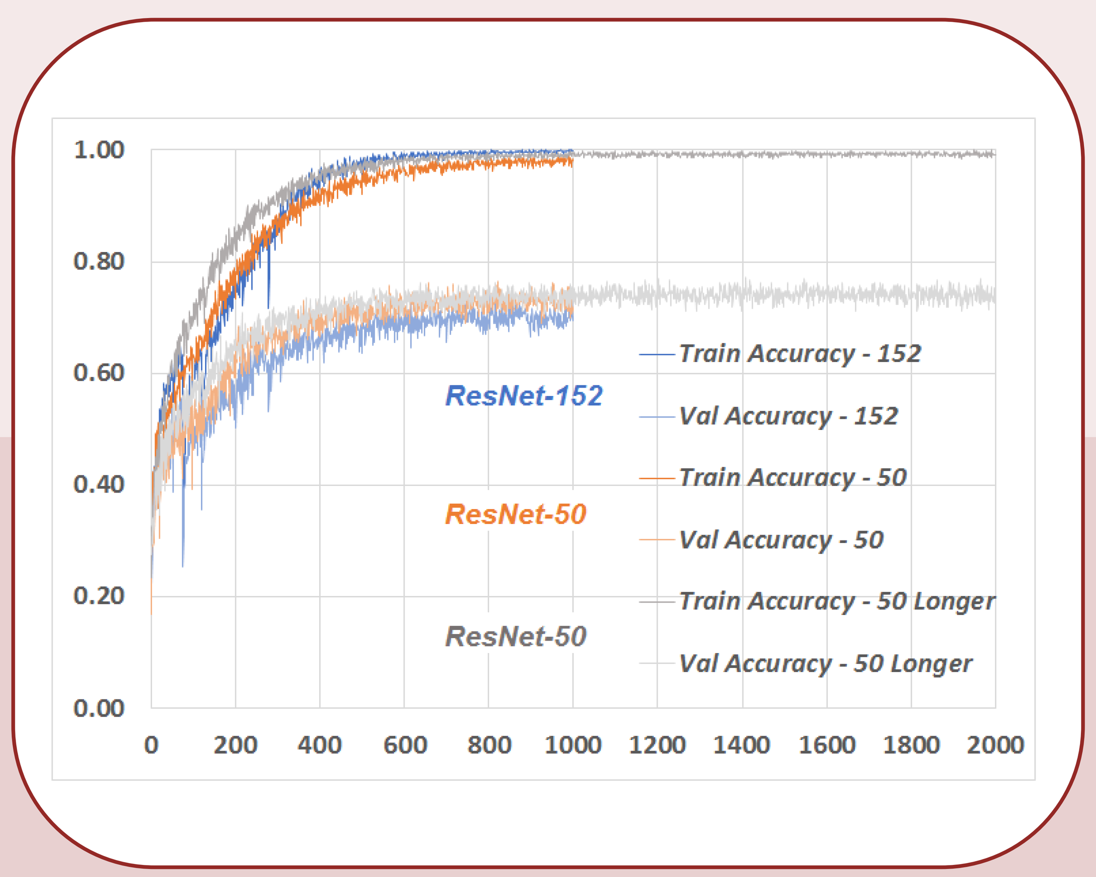	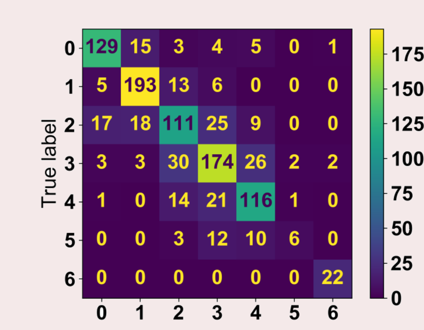	train: 1.0000 dev: 0.7713 test: 0.7510	Class 0: 0.83; 0.82 Class 1: 0.84; 0.89 Class 2: 0.64; 0.62 Class 3: 0.72; 0.72 Class 4: 0.70; 0.76 Class 5: 0.67; 0.19 Class 6: 0.88; 1.00
Resnet 152	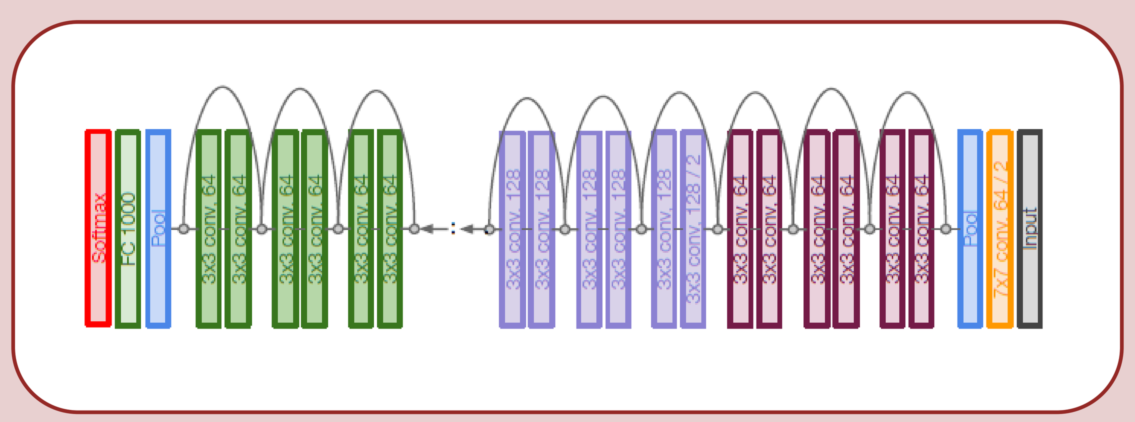			train: 0.9990 dev: 0.7429 test: 0.7200	Class 0: 0.77; 0.93 Class 1: 0.89; 0.88 Class 2: 0.63; 0.47 Class 3: 0.59; 0.67 Class 4: 0.81; 0.72 Class 5: 0.10; 0.11 Class 6: 0.85; 0.49

Discussion / Future Work

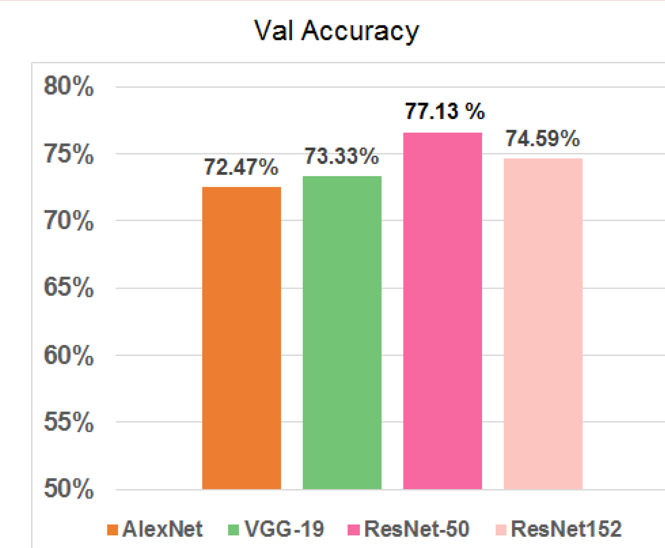

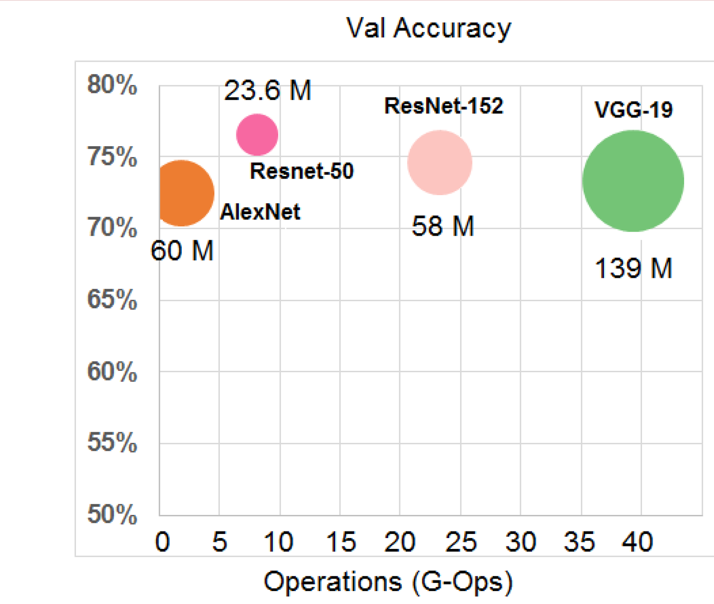

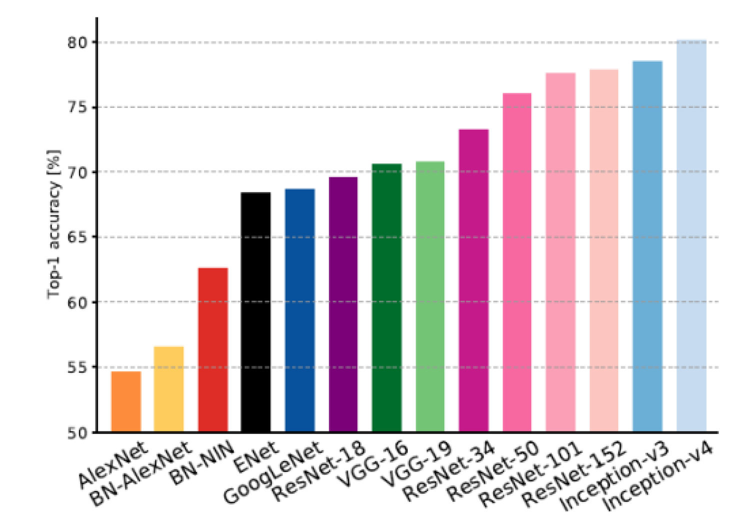

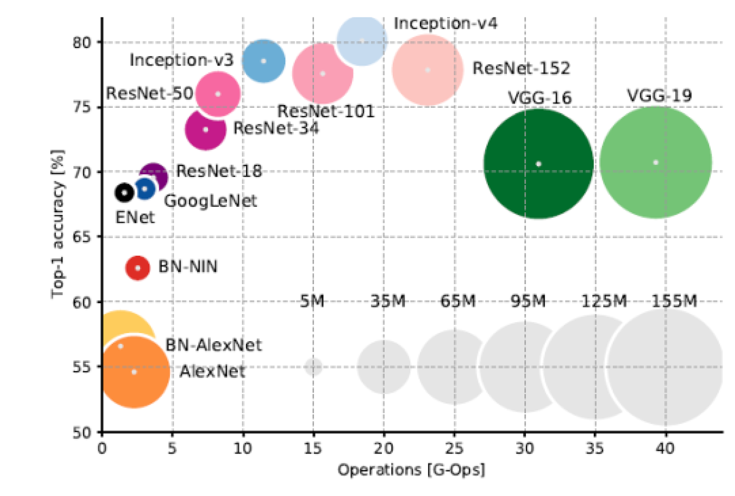

In Val accuracy (Top), the accuracies we obtained from AlexNet, VGG-19 and Resnet-50/152. For comparison, the second picture (bottom) shows accuracies, computation cost and # params of the most relevant entries submitted to ImageNet, which are similar. Whether or not this is coincidental remains to be discerned. In conclusion these models proved to improved our baseline in almost 25 percentage points. It motivates use to try out semi-supervised methods, such as meta pseudo labels, to improve even more (> 90%). Code can be accessed at: https://github.com/adamsdsit/cs231N-SP2022

References / Acknowledgments

The authors would like to thank Larissa Macedo dos Santos-Tonial (PhD) and her research team for providing us the data/soil images.