Automated Crop Disease Classification: Examining Lightweight Transfer Learning for the Edge

Bhagirath Mehta (EE)
bmehta18@stanford.edu

Pranav Sriram (CS)
prsriram@stanford.edu

CS 231N Final Project
Mentor: Haoshen Hong

Introduction

- **Motivation & Background:**
  - Pests and diseases lead to an annual **20-40% loss** in food production around the world
    - With climate change, problem worsening
  - Automated crop disease identification is more important than ever for farmers
  - Little prior work on **approach generalized to all crops**
  - Network of cheap IoT cameras constantly gathering images of plant leafs
    - **Need for a lightweight CV approach that can run on edge**
- **Problem Statement:**
  - Evaluate the effectiveness of a lightweight CNN approach to **binary leaf disease classification**
    - **Effectiveness of transfer learning from low-quality to high-quality dataset**
  - Details:
    - Input: JPEG crop leaf images **of any species**
    - Output: Healthy/Diseased Label

Datasets

- **Low-Quality PlantVillage Dataset:**
  - 14 different crop varieties
  - Chose 9360 images with 60/20/20 split

Healthy Corn Leaf

Diseased Apple Leaf

- **High-Quality PlantLeaves Dataset:**
  - 12 different crop varieties
  - Chose 3335 images with 60/20/20 split

Healthy Chinar Leaf

Diseased Chinar Leaf

Methods

- **Architecture:**

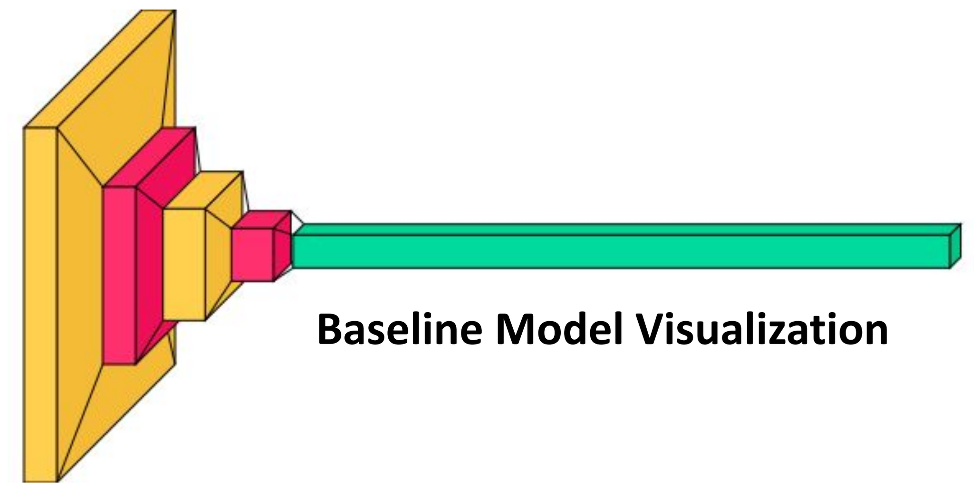

Baseline Model Visualization

  - 2 CONV-POOL
  - 2 FC Layers
  - Over **74 million** params

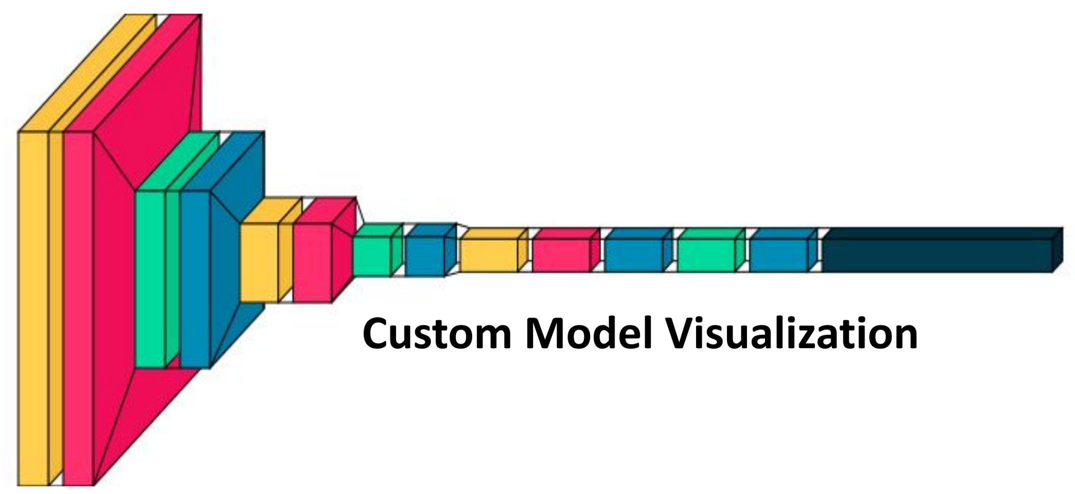

Custom Model Visualization

  - 3 CONV-POOL
  - 3 FC Layers
  - BatchNorm
  - LeakyReLU
  - Over **24 million** params

Experiments & Analysis

- **Experiments 1 and 2:** Compare performance of baseline and custom models; train and test on PlantVillage

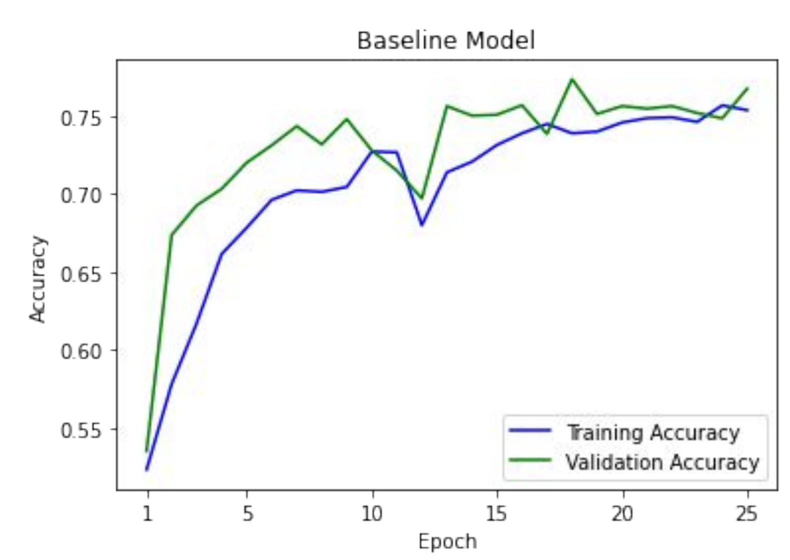

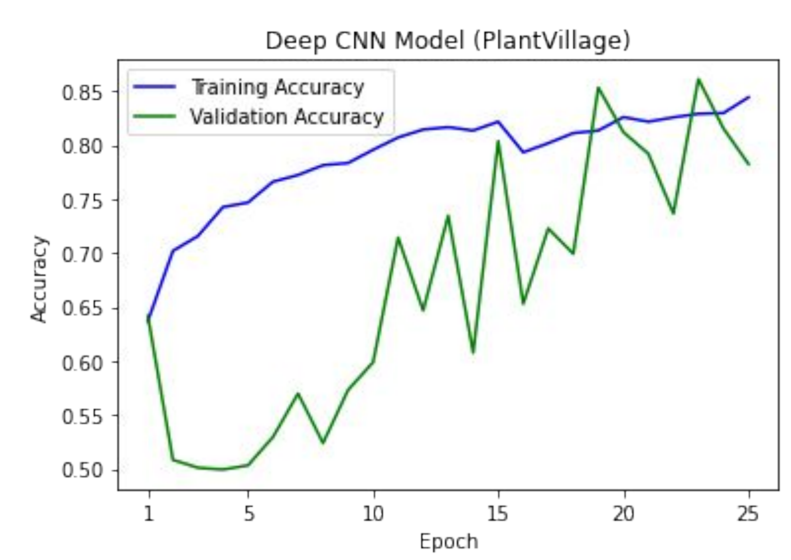

- **Experiment 3:** Using custom architecture, train and test on PlantLeaves dataset
- **Experiment 4:** Using custom architecture, train on PlantVillages and test on PlantLeaves dataset
- **Experiment 5:** Using custom architecture, pretrain on PlantVillages; finetune and test on PlantLeaves dataset

Results

Testing on PlantVillages for Experiment 1 and 2:

Model	Test Accuracy (PlantVillage)
Baseline	78.03%
Custom	**86.21%**

- We see that the custom model vastly outperforms the baseline model, by **over 8%**
  - While being far more lightweight (a third of the size)!

Testing on PlantLeaves for Experiment 3, 4 and 5:

Model	Test Accuracy (PlantLeaves)
PlantLeaves Train	**53.22%**
Only Pretrain	50.82%
Pretrain and Finetune	**52.77%**

- We see that no experiment setup performs particularly well on the binary PlantLeaves task, with each of them doing only slightly better than chance.
  - Suggests high degree of difficulty in the underlying task

Conclusion / Future Work

- Successfully constructed a **lightweight custom model for low-resolution leaf disease classification**
  - Beat heavier baseline model
- Transfer learning (from low-quality images to high-quality images) was not effective in this domain space
  - In context, even a model trained directly on the high-quality images did not perform much better
- **Future Work:** Refine approach with more data, training and complex architectures like 'squeeze-and-excitation' networks; can also reframe the problem to separate different species and diseases for ease of classification