Identifying Fungal Diseases in Growing Wheat using Convolutional Neural Networks

Okonda Joseph L.

OVERVIEW

Motivation: Wheat is the second most consumed crop in the world. The crop is susceptible to a variety of diseases which can lead to significant yield losses and thus loss of income. Currently, farmers deal with diseases by either planting cultivars resistant to specific diseases or by using fungicides. The former, while effective, may be prohibitively expensive and thus not readily available to all farmers. The latter, on the other hand, requires farmers to know when to apply the fungicide. This usually requires farmers to periodically inspect their fields to ensure that they catch any outbreak early and contain it – a tedious process that does not scale. A system that inspects fields (using UAVs, for instance) and automatically identifies any diseases, can be used to streamline large scale wheat production. Such a system would need a way to identify the various diseases.

Summary: We train 2 classifiers. First a Binary classifier to differentiate between healthy and diseased plants then a multiclass classifier for the 4* different fungal diseases.

DATA & PRE-PROCESSING

- Used Google Image Search to retrieve images for each of our classes.
- Train-Dev-Test split of 80-10-10
- Resize Image to 256x256. Augment with random horizontal flips, random brightness during training.

Disease	# of Samples
Healthy	838
Head Blight (FHB)	526
Rust (Leaf and Streak)	515
Leaf Blotch (STB)	557
Powdery Mildew (BG)*	296

References:

- Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: (2014), pp. 1–15. ISSN : 09252312. DOI : http://doi.acm.org.ezproxy.lib.ucf.edu/10.1145/1830483.1830503 . arXiv: 1412.6980 . URL : http://arxiv.org/abs/1412.6980 .
- [G. M. Murray and J. P. Brennan. "Estimating disease losses to the Australian barley industry". In: Australasian Plant Pathology 39.1 (2010), pp. 85–96. ISSN: 08153191. DOI: 10.1071/AP09064.

MODEL: BINARY.

- Experimented with 4, 5 & 6 Layer Models. Conv-1 Filters and Biases initialized from VGG-16 and frozen (Freezing other layers doesn't work as well).
- Optimizes Binary CE using Adam with learning rate=1e-3 for 20 epochs. Uses Dropout Regularization with keep-prob=0.7.
- Best model achieves 86% Accuracy on unseen test set.

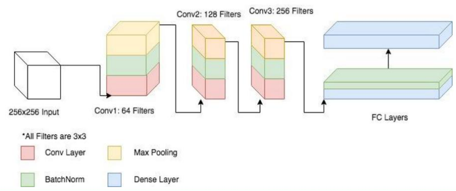

MODEL: MULTICLASS.

- Same Architecture, But:
  - Fewer/constant # of Kernels per layer (64 all through)
  - Aggressive Regularization using Dropout (0.4) and L2 Reg. ($\lambda = 0.05$)
  - Best model also uses exponential LR Decay (rate=0.9) achieves 56% accuracy.

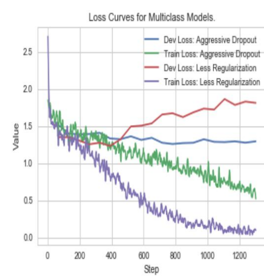

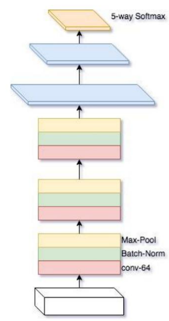

RESULTS

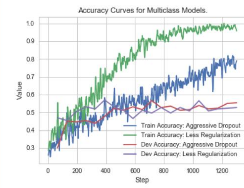

- Even with Frozen VGG-Weights, without aggressive regularization, the model still overfits.
- LR Decay, Weight Decay, and Dropout reduces the high variance alleviate this but do not improve the accuracy by much (54% - 56%)

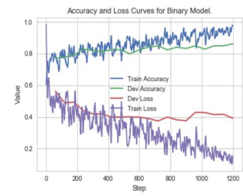

- Transfer learning helps us train on a limited dataset, resulting in fairly high accuracy and very little variance. However, we believe that this success is mainly because we have more samples for every class.

DISCUSSION AND FUTURE

Discussion

- Although transfer Learning helps our model generalize with limited data (especially in the binary case), it can only go so far. In order to do fine grained classification, more data is needed.
- When training with limited data, proper regularization and a carefully chosen learning rate annealing schedule helps in reducing the model's high variance.

Future

- The main obstacle in this project was the absence of a high quality dataset. A natural extension of this project is thus to tackle the same problem with an improved dataset.
- Moreover, the diseases explored herein and the the crop we focus could easily be replaced by any other crop, as long as data for that crop exists.