Amazing Amazon

Aaron Loh (MSCS ’17), Kenneth Soo (Stats ’17)

Background and Problem Statement

Background: Deforestation is an urgent problem in our world today, as it contributes to reduced biodiversity, habitat loss, climate change, and other devastating effects. Using satellite data, we can track deforestation and better understand where it happens.

Current tracking efforts within rainforests largely depends on coarse-resolution imagery from Landsat (30 meter pixels) or MODIS (250 meter pixels) which have limited effectiveness in detecting small-scale deforestation.

Problem Statement: This is a multi-label classification problem, aiming to label satellite chips corresponding to land cover, land use, as well as atmospheric conditions, using CNNs to generate predictions, and the F-2 Score for evaluation.

Dataset

The data contains 40k images in the Amazon, and each image contains 1 or more of the following labels. To visualize the spread of the labels, a similarity matrix is plotted:

Sample images with labels

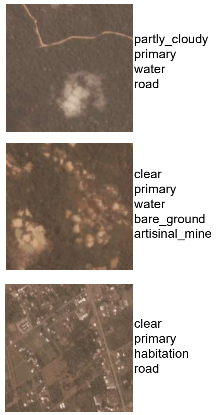

Label count	Label title
28431	clear
7261	partly_cloudy
2089	cloudy
2697	haze
7411	water
100	conventional_mine
862	bare_ground
339	artisinal_mine
37513	primary
98	blow_down
340	selective_logging
332	blooming
4477	cultivation
209	slash_burn
12315	agriculture
3660	habitation
8071	road

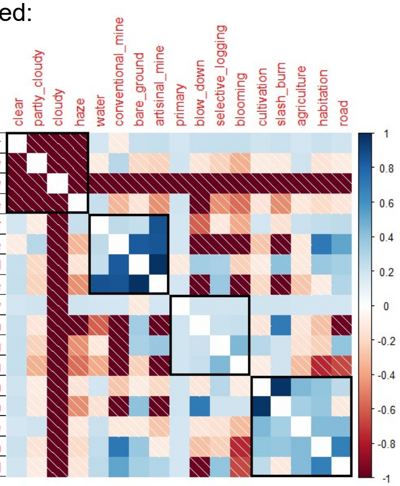

Models

Simple CNN model

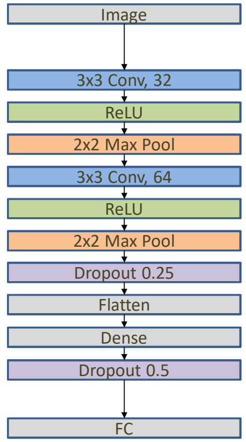

Resnet model

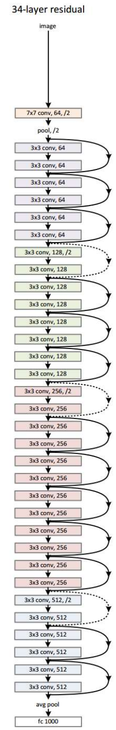

Hyperparameter	Value
Optimizer	Adam
Loss function	Binary Cross Entropy
Learning Rate	1e-3
Batch Size	64
Epochs	5

We experimented with two different network architecture, as shown above. The first used regular convolutional, max-pooling, and dropout layers, while the second used transfer learning with the ResNet model. The models end with a fully connected layer and outputs that tell us the probability of each label occurring.

Results and Analysis

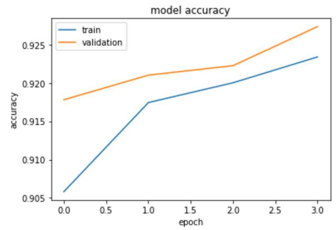

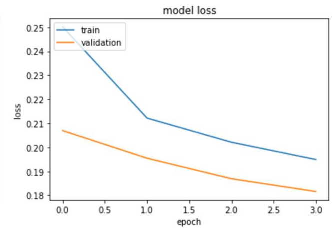

Evaluation

We use the F2 score to evaluate our model. This is based on the F score and it weighs recall higher than precision.

$$F2 = \frac{5pr}{4p+r}, \text{ where } p = \frac{tp}{tp+fp}, \text{ and } r = \frac{tp}{tp+fn}$$

Results

Model	F2
ResNet with Transfer Learning	0.769
Simple CNN Model	0.823

Analysis

We are able to outperform the ResNet model with transfer learning with a simple CNN model, but that may suggest that we need to run the ResNet for more epochs, as it is deeper and harder to train. Our accuracy has also not plateaued, so running for more epochs may help.

Conclusion/Future Directions

We found that CNNs are adept at predicting atmospheric conditions from images, but that further hyperparameter tuning and experimentation with architectures may be necessary to improve F2 Score. We aim to also include data augmentation and visualizations of the layer activations in the future, so as to better understand what the model is learning.