Looking at the Amazon Through Computer Vision

Sneha Kudli, Benjamin Pastel, Steven Qian

Motivation

The rainforests are a precious ecosystem, home to an abundance of wildlife and vegetation. Unfortunately, it is a treasure that is quickly disappearing, as every minute we are seeing the world lose an area of forest the size of 48 football fields. This deforestation accounts for many damaging effects including habitat loss, reduced biodiversity, and even climate change. We hope that the use of collected satellite imagery reinforced with computer vision may help organizations saving the rainforest to respond more quickly and efficiently.

Figure 1. Training Data [1]

The Challenge

Kaggle has presented the challenge of labeling satellite imagery based on the atmospheric conditions, common land cover/land use, and rare land cover/land use depicted in the chips.

We will be testing multiple different Neural Net models to tackle this problem and will be evaluating our progress through the test set and correct labels provided by Kaggle.

Dataset

The provided dataset is composed of 4-band tif satellite images. This effectively means we are given image data across 3 channels, the traditional RGB channels and a near infrared channel. The most common labels in the dataset are rainforest, agriculture, rivers, and towns/cities.

Figure 1 shows an example of a few given data-point and their labels.

Methods

Different Models: See **Figure 4** for the layout of the current best model that is composed of custom layers on top of Squeezenet. We have also tried building original models and augmenting Inception and Resnet.

Data Bootstrapping: Augmented images with infrequent labels by flipping/rotating and add the new images to the data set with the original label.

Threshold Optimization: Used an iterative algorithm to find the threshold probabilities for each class to maximize the F2 score.

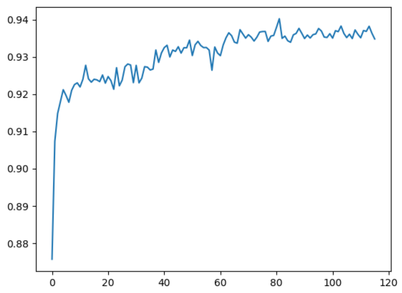

Figure 2. Training F2 Scores

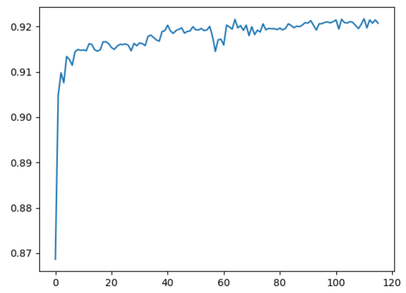

Figure 3. Validation F2 Scores

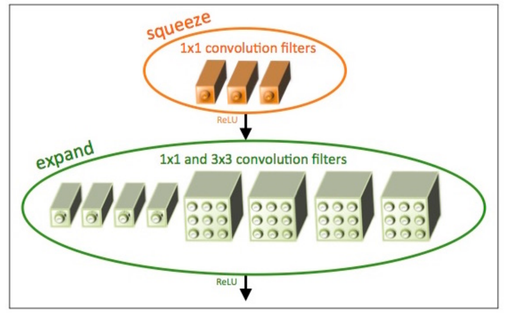

Figure 5 Squeezenet Fire Module [2]

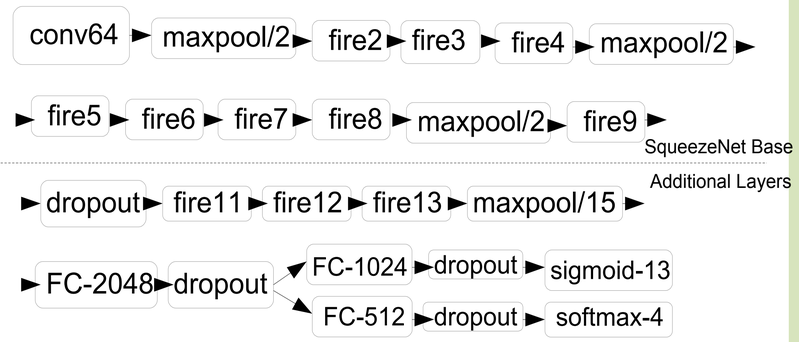

Figure 4 Current Model

Results

Currently, the highest test accuracy we have achieved is: **0.922**. This has put us up to the **33rd** position on the Kaggle leaderboard for this challenge.

Work Cited

[1] K Inc. Planet: Understanding the amazon from space, 2017.
[2] SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <1MB model size. Iandola et. al.arXiv:1602.07360, 2016