Understanding the Amazon from Space

Yiqi Chen, Fanming Dong, Chuanwei Ruan

kaggle

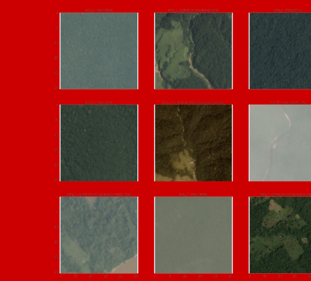

1. Introduction

The Amazon rainforest is the largest rainforest in the world. However, the deforestation of Amazon rainforest has accelerated since 1991.

In order to further understand the status quo of Amazon surface, a company called Planet released a data challenge over Kaggle, aiming to classify cloud conditions and land use phenomena from satellite images.

Previous models include gradient boost trees and other statistical methods on hand-crafted features. However, these models suffer from not allowing computers to pick features.

2. Problem Statement

We intend to classify various phenomena of interest (atmospheric conditions, land cover) in the Amazon basin, from its satellite images.

We will build various CNN architectures to predict the labels of images in the image dataset and see which architecture works the best.

We evaluate the accuracy of our approach by using average F2-score of each image in the dataset. F2-score is defined as:

$$F_2 = 5 \cdot \frac{\text{precision} \cdot \text{recall}}{4 \cdot \text{precision} + \text{recall}}$$

3. Dataset

We will use the satellite image dataset provided by Planet through Kaggle.

- Dataset size: 40,480 training samples (known labels), 61,192 test samples (unknown labels).
- Image size: 256 x 256 x 3
- Labels: the labels consist of cloud cover labels and land use labels. Each image has one cloud cover label, and 0 or more land use labels.

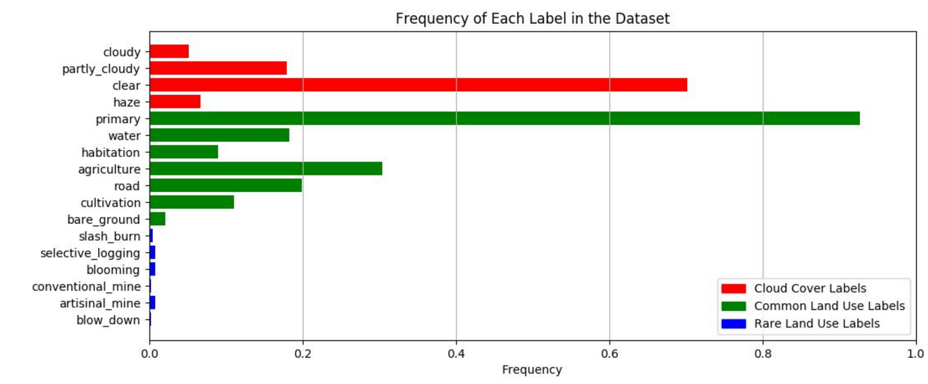

4. Methodology

1. Transfer Learning
We trained our model based on a pre-trained VGG-16 model and a ResNet-50 model. During training, we assign lower learning rates to lower layers.

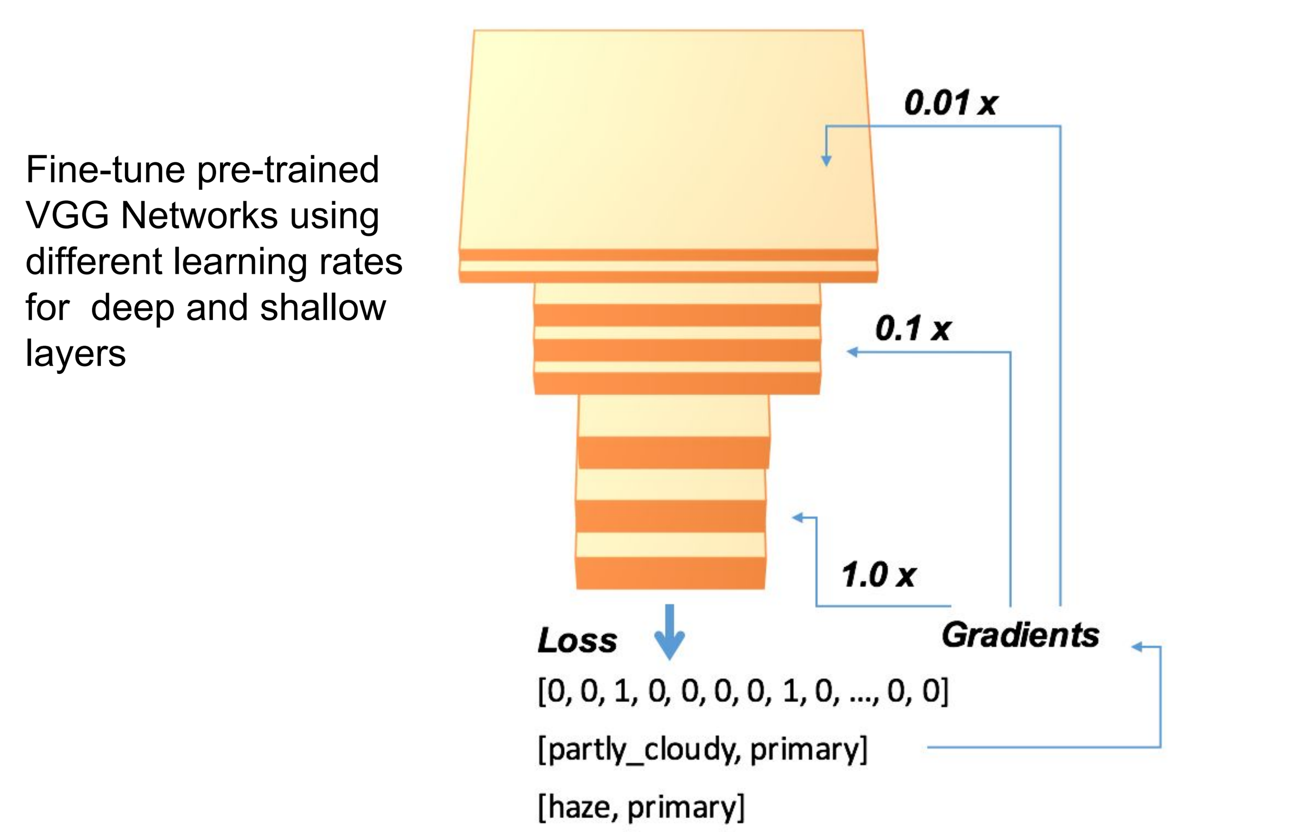

2. Data Augmentation
We transform each image into 36 images through resizing, cropping, flipping (both horizontally and vertically) and rotating.

3. Loss Functions
We designed 2 loss functions

- We compute sigmoid cross entropy loss of each label independently as a binary classification problem in each image.
- We compute one softmax loss of cloud cover labels and sigmoid cross entropy loss of each land use label.

4. Densely Connected Convolutional Neural Network

- We implemented a DenseNets model with growth rate 12 and 3 dense blocks and we trained the model from scratch on the data.
- Intuition: Labels are of different scales and semantics. DenseNets allows higher layers to access information directly from all previous layers, which makes it an ideal choice.

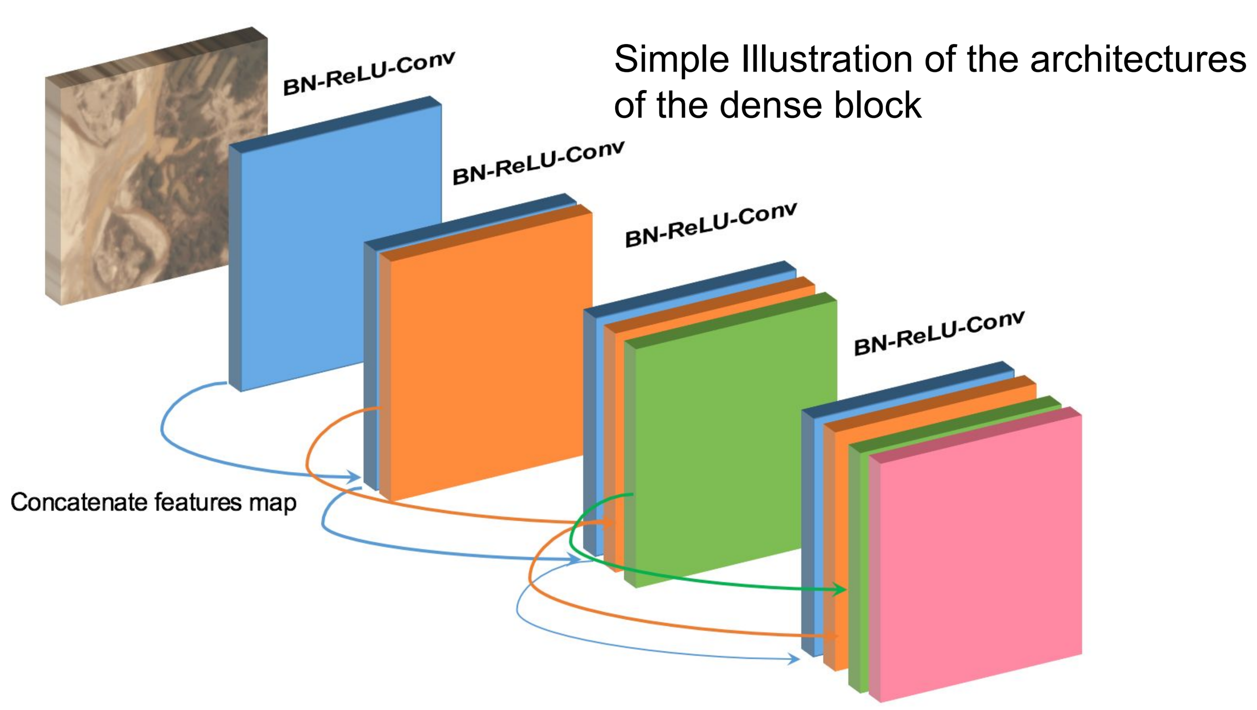

Simple Illustration of the architectures of the dense block

5. Experimental Findings

Different Models

Training v.s. Validation (DenseNets)

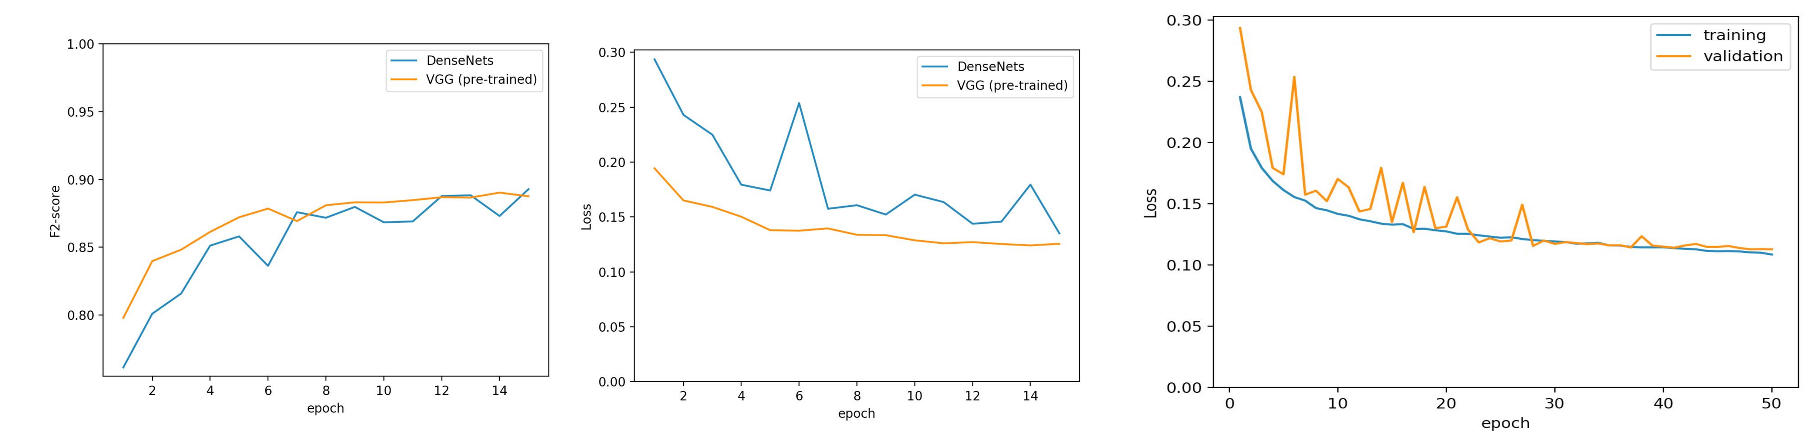

Performance Diagnostics on the Best Model

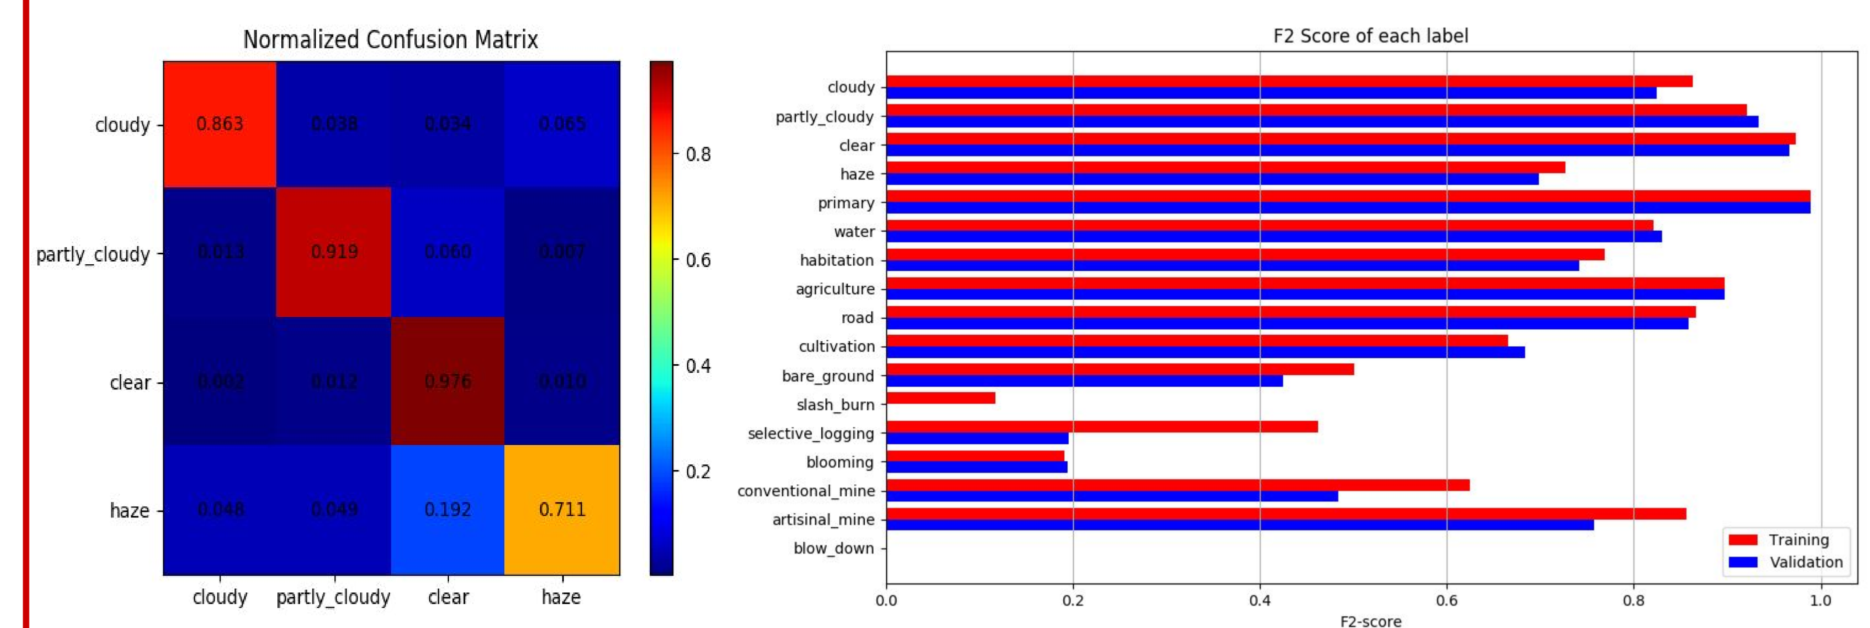

Co-occurrence Matrix

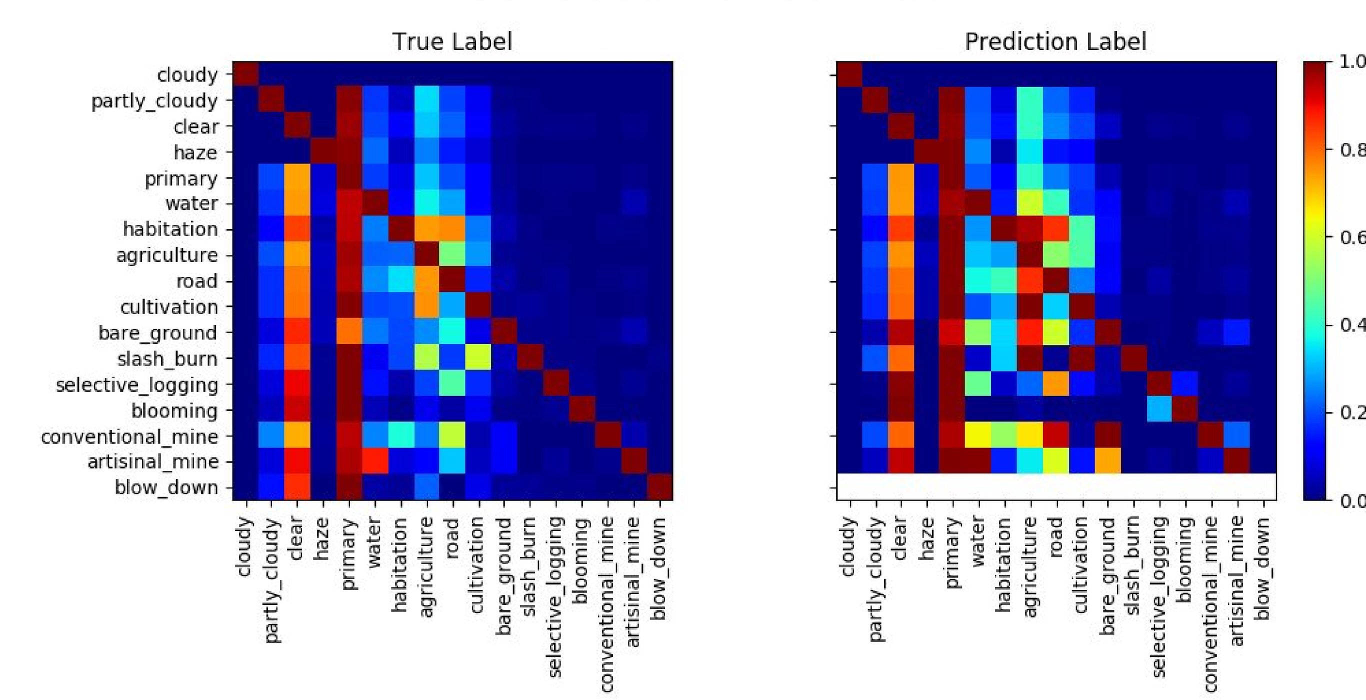

6. Discussions

- We are able to achieve > 0.9 F2-score on validation set on all of our models.
- Among all architectures we experimented with so far, fine-tuned VGG performs the best. It also achieves 0.9123 F-2 score on the test set.
- Even though the pre-trained models are trained on ImageNet, they still work well with the new satellite image dataset.
- Carefully designed architecture, such as DenseNets, can also perform well even when it's trained from scratch with limited training data.
- As a further step, we can utilize RNN to capture the co-occurrence information of the labels to get better performance.