

Planet: Understanding the Amazon from Space

Use satellite data to track the human footprint in the Amazon rainforest

Project 26

Xuewen Yao



Motivation and Goal

Every minute, the world loses an area of forest the size of 48 football fields. And deforestation in the Amazon Basin accounts for the largest share, contributing to reduced biodiversity, habitat loss, climate change, and other devastating effects. But better data about the location of deforestation and human encroachment on forests can help governments and local stakeholders respond more quickly and effectively. The goal of this project is to label satellite image chips with atmospheric conditions and various classes of land cover/land use so that global community can better understand where, how, and why deforestation happens all over the world.

Approach


Resnet, densenet and vgg are used to train and validate. In order to save time and computation, pretrained networks from pytorch are used. However, all three nets results in 1000 classes, so a ReLU activation and fully connected layer are added at the end of the pretrained network and initialized randomly.

For resnet, the network is 18 layers deep. Batch size of 128 is used in training and 64 in validation, which results in a smoother training and validation graph. Densenet has 121 layers. The batch size is 32 for training and 16 for validation due to memory issues. Vgg shares the same batch size with densenet and it's 11 layers deep.

Dataset and Setup

A total of 40478 training images and labels are obtained from kaggle. And 10% of them were used as validation set after shuffling. An example of images with labels is shown above.

There are 17 possible labels and 4 of them are weather label and the others describe land cover/land use. The frequencies of each label are plotted below.



All data was processed using Amazon AWS EC2 instances. The details are listed below:

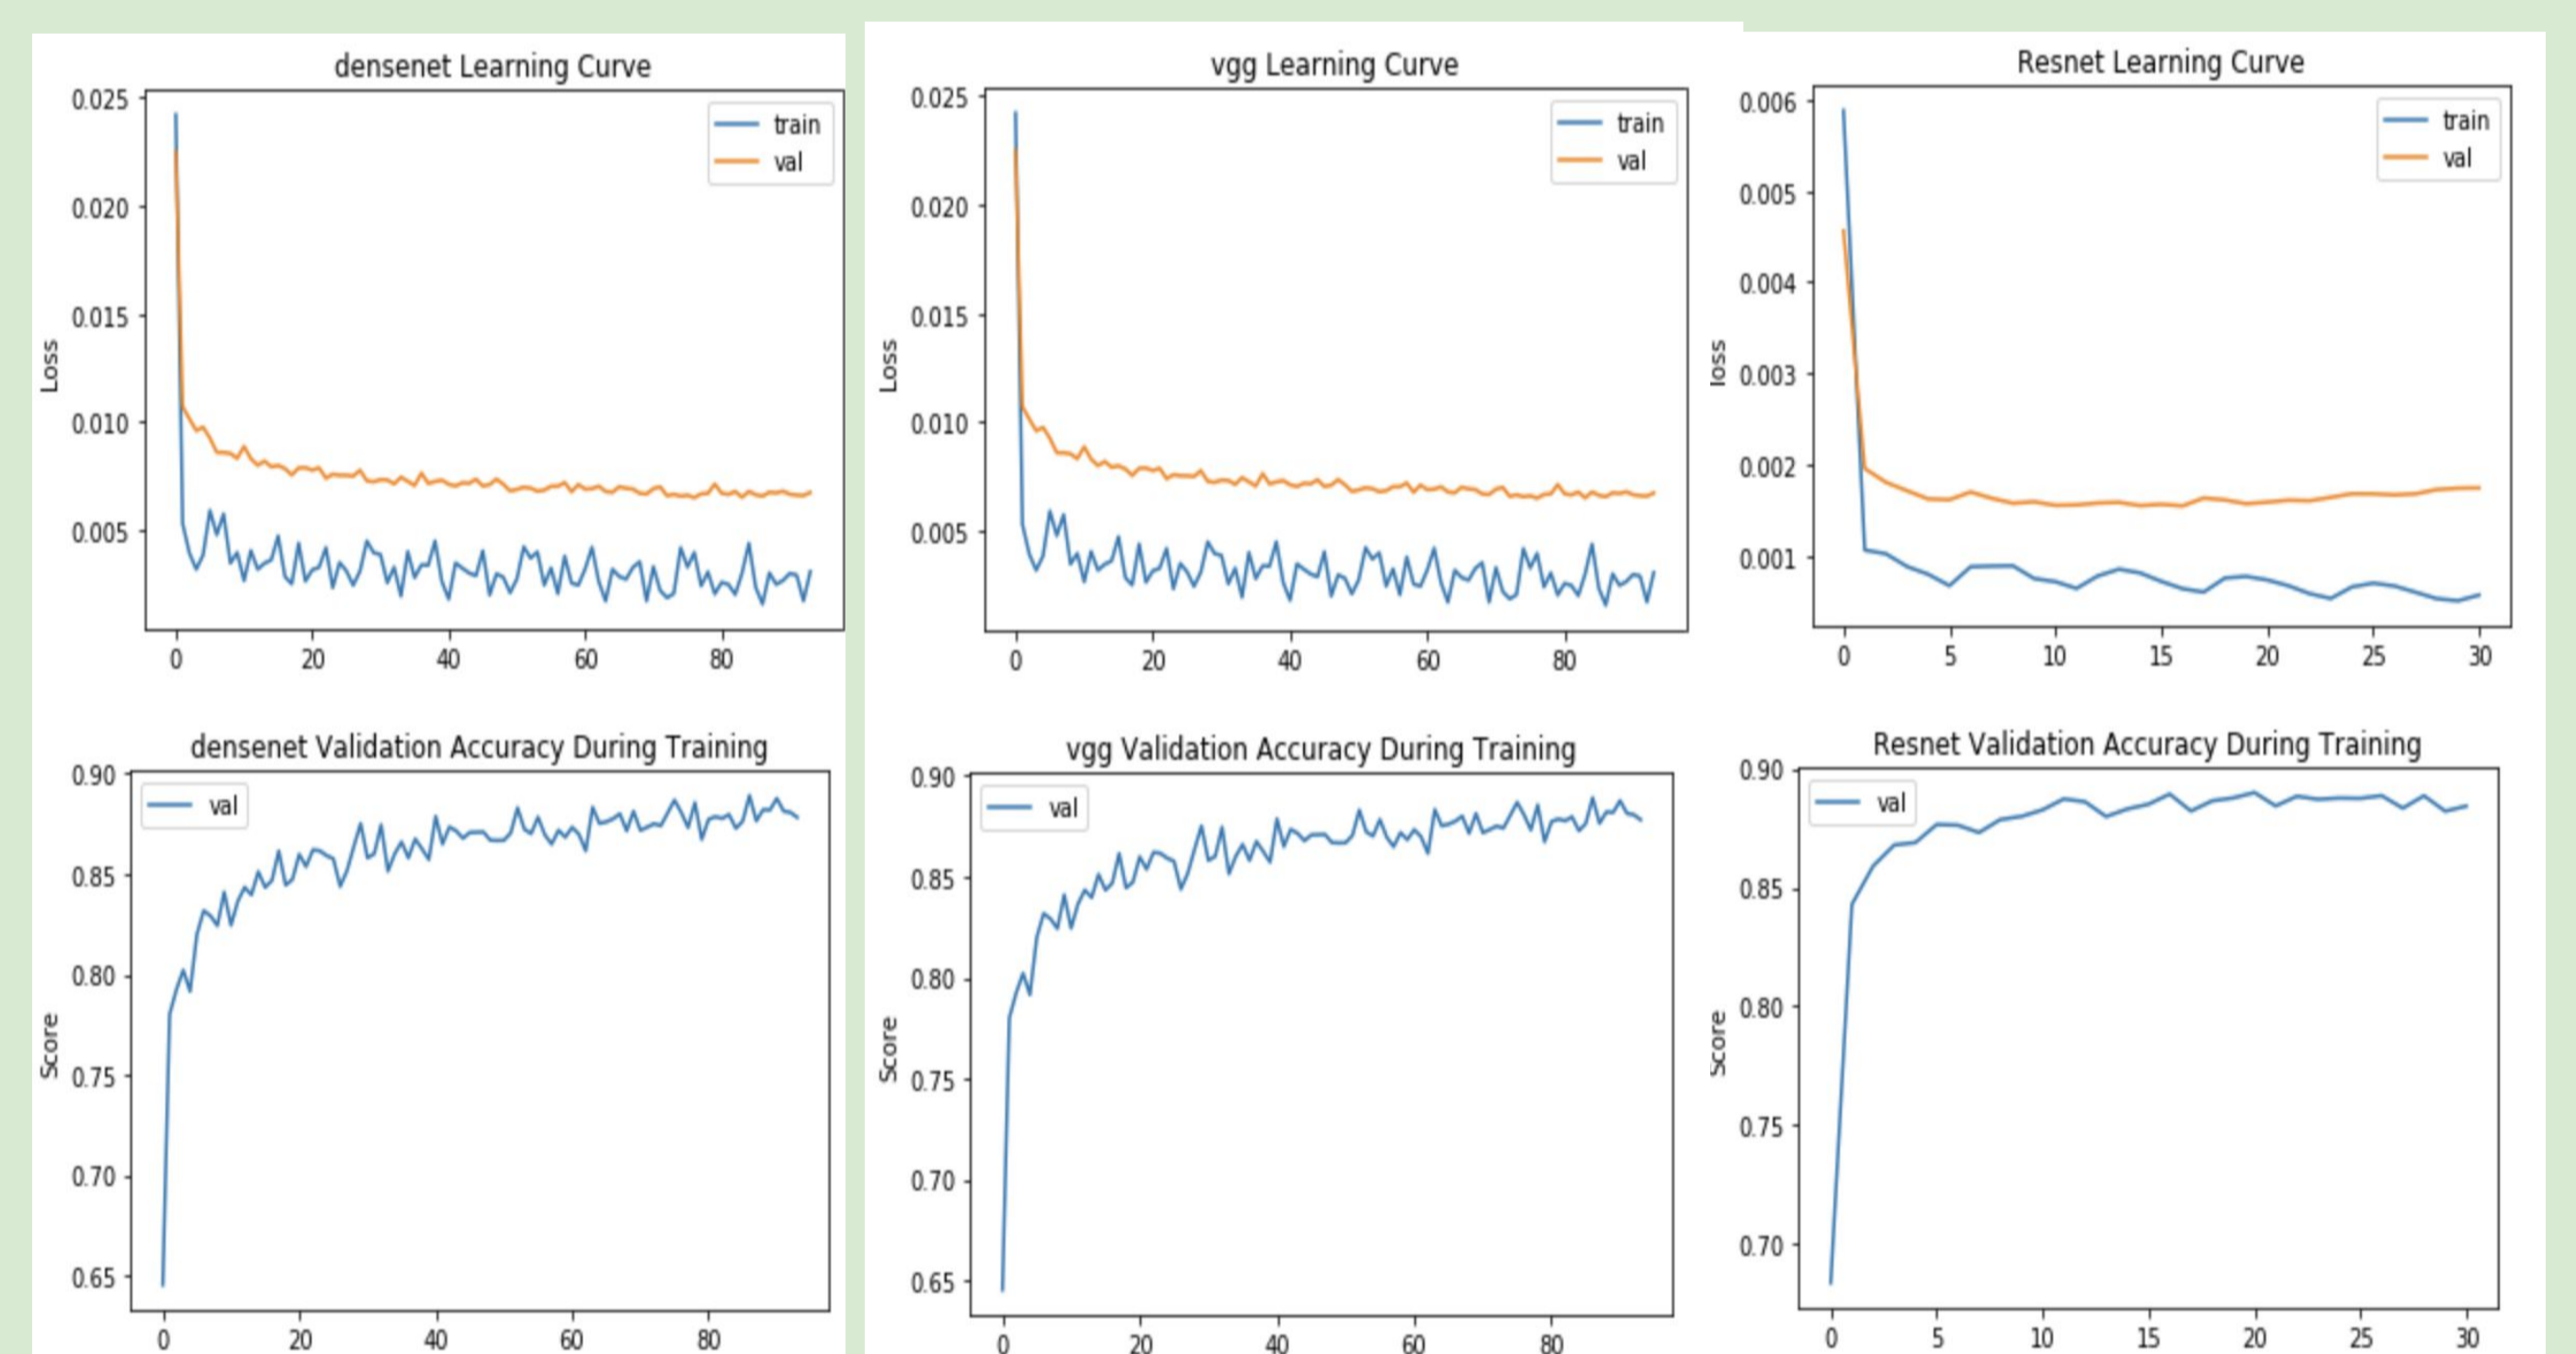
Model	GPUs	GPU Memory	vCPUs	Main Memory	EBS Bandwidth
g3.8xlarge	2	16 GiB	32	244 GiB	7 Gbps

Evaluation and Result

The mean F2 score is used to evaluate the result. It measures accuracy using weighted precision p and recall r.


$$(1 + \beta^2) \frac{pr}{\beta^2 p + r} \text{ where } p = \frac{tp}{tp + fp}, r = \frac{tp}{tp + fn}, \beta = 2.$$

All three models give same score after a number of epoches. Resnet reaches ~0.88 quickly and its loss is smaller than the other models. As mentioned before, resnet uses larger batch size so its curve appears to be smoother. Densenet uses most layers, but its performance isn't better than resnet. Vgg uses most memory among all and because of this, it runs out of memory after one epoch. So if given more memory, vgg might perform better than the other two methods.



Reference: Kaggle Competition
<https://www.kaggle.com/c/planet-understanding-the-amazon-from-space>

Images the model predicts correctly



Images the model has confusions

