Understanding the Amazon Rainforest from Space using CNNs

Shubham Chandak, Vamsi Chitters, Sagar Honnungar

CS231N
Spring 2017

INTRODUCTION

- Rapid deforestation in the Amazon basin has had a devastating impact on biodiversity and climate.
- Automated classification of satellite imagery based on land use patterns and atmospheric conditions can aid in understanding the causes of deforestation.
- With this goal, *Planet Labs*, builder of Earth-imaging satellites, has hosted a *Kaggle* contest by providing high resolution images of the Amazon basin.

PROBLEM STATEMENT

- Multi-label classification of satellite images into 17 classes:
  - Weather classes
  - Major land cover/use classes
  - Rare land cover/use classes
- Evaluation is based on *F2* score:

$$F_\beta = (1+\beta^2)\frac{pr}{\beta^2 p + r}$$

where β=2, *p* and *r* are precision and recall.

- Higher penalty for false negatives.

DATASET

- 256x256 JPEG
- ~40K training images
- ~60K test images
- Highly uneven class distribution
- Mutually exclusive weather classes

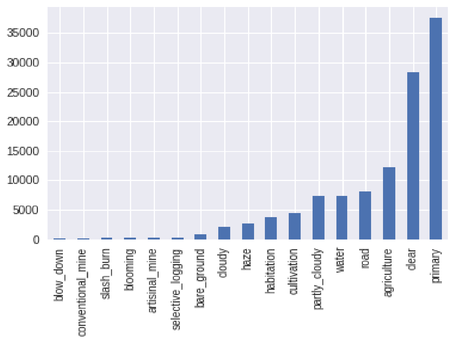

Distribution of classes in training data

METHODS

Transfer Learning

- ImageNet models:
  - ResNet18
  - InceptionV3
  - Xception
- Binary cross-entropy loss for each class.
- FC layers trained with Conv layers frozen.
- Conv layers fine tuned with low learning rate.

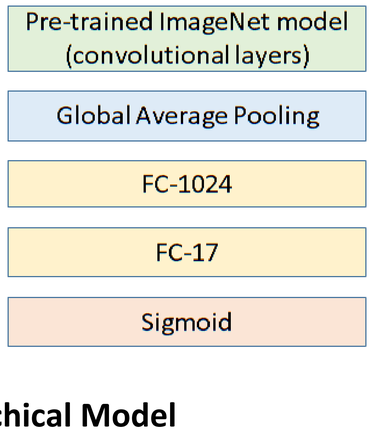

Hierarchical Model

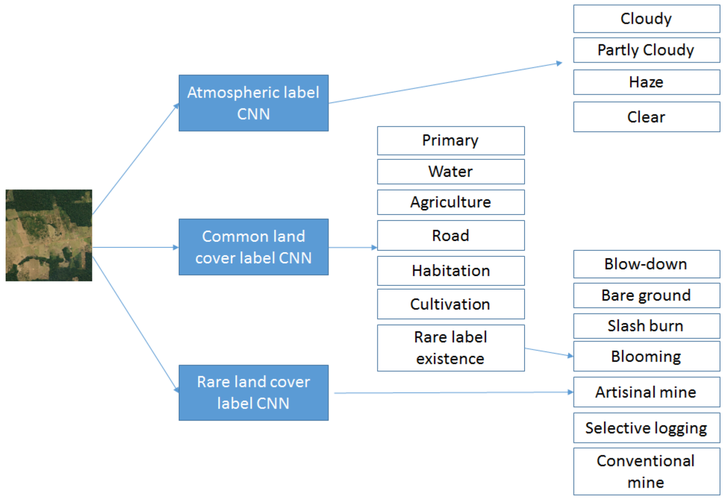

Improving Performance

- Data augmentation: random rotation, horizontal and vertical flips, zoom etc.
- Class-specific thresholds
- Ensembling different models
- Higher resolution images (32x32 → 64x64 → 96x96)
- Learning rate schedule (for Hierarchical Model)

RESULTS

Ranked **11th** among 372 contestants (as of 06/05/17)

Model	Input Size	Test F2 score
Fine-tuned ResNet18	256 X 256 X 3	0.92276
Fine-tuned InceptionV3	200 X 200 X 3	0.92398
Fine-tuned Xception	200 X 200 X 3	0.92471
Hierarchical model (1)	64 X 64 X 3	0.92323
Hierarchical model (2)	96 X 96 X 3	0.92457
Ensemble	N/A	**0.93019**

Image	True labels	Predicted labels
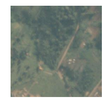	agriculture clear habitation primary road	agriculture clear habitation primary road
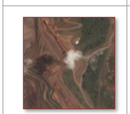	conventional_mine habitation partly_cloudy primary road	conventional_mine habitation partly_cloudy primary road
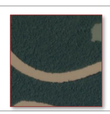	clear primary water	clear primary water
	cultivation habitation partly_cloudy primary road	cultivation habitation road *agriculture water* primary partly_cloudy
	clear primary water	water primary *partly_cloudy* clear

Saliency Maps:

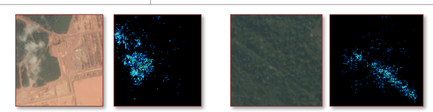

Partly cloudy Road

CONCLUSIONS AND FUTURE WORK

- CNNs proved effective for this Kaggle challenge.
- Ensembling significantly improved the F2 score.
- Hier. model performed well even for low res. images.
- Future work: utilize higher resolution images, train hierarchical model end-to-end.