Labeling Satellite Imagery with Atmospheric Conditions and Land Cover (Kaggle)

Shakti Sinha (shakti@cs.stanford.edu) CS231n 2017

Background

Deforestation in the Amazon Basin contributes to reduced biodiversity, habitat loss, climate change, and other devastating effects.

Detection and understanding of markers of human activity over large areas will enable faster and more effective response to activity that indicates deforestation

Problem Statement

Input: 256x256 satellite image in 3-channel JPG or 4-channel TIF

Output: One or more labels from a set of 17 possible labels that denote atmospheric conditions, land cover or land use

Objective: Correctly label satellite images, as measured by the F2 score.

Architecture 1- Effect of Input Size

[Conv – max_pool – dropout]*n – [affine – dropout] – affine

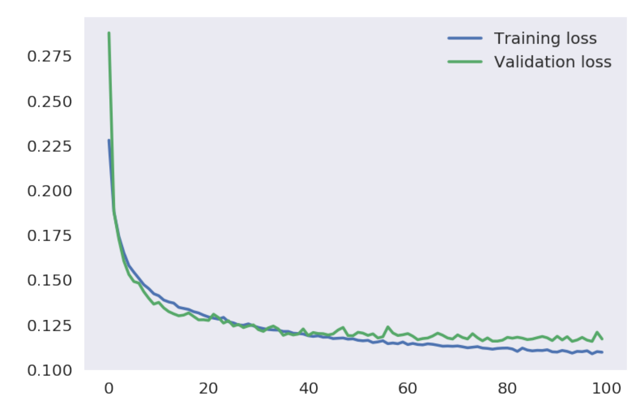

64x64

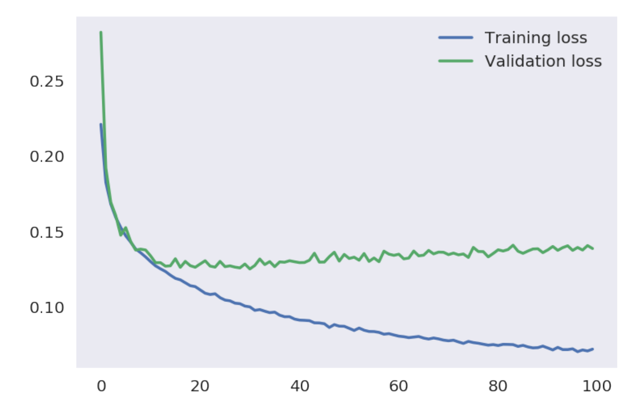

128x128

Larger input leads to overfitting because of larger dense layer

Architecture 2 – Refine Pre-trained Models

Inception V3([0:172]frozen – [172:]trained – 512 Dense)

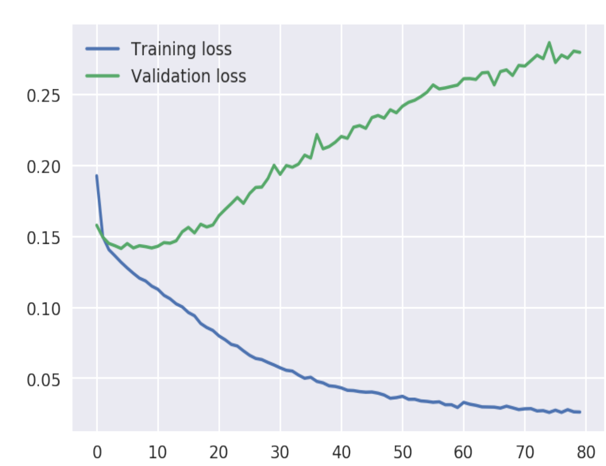

1024 Dense layer overfits

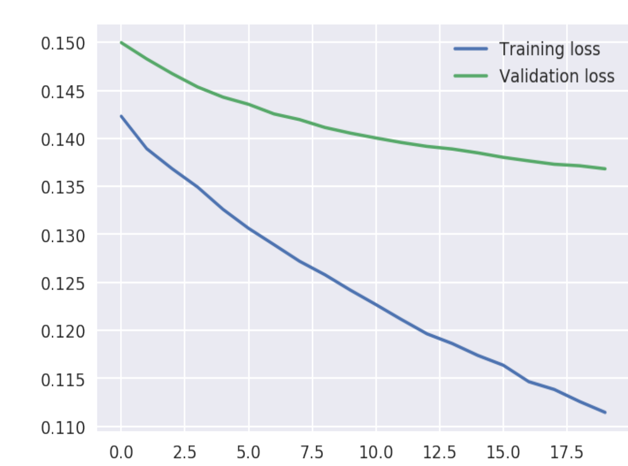

Training [172:] layers can improve further

Dataset

- 40479 labeled satellite images
- Labeling has been performed using Cloudflower
- JPG images have 3 channels – R, G, B
- TIF images have 4 channels – R, G, B, IR
- Test data with labels withheld for public leaderboard evaluation
- Unreleased data for private leaderboard evaluation

Label Co-occurrence

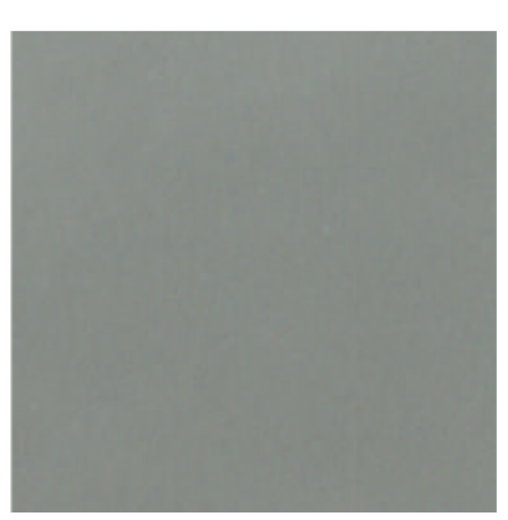

Noisy labels: Model detects unlabeled rainforest under clouds

Model Errors

Using Label Co-occurrence

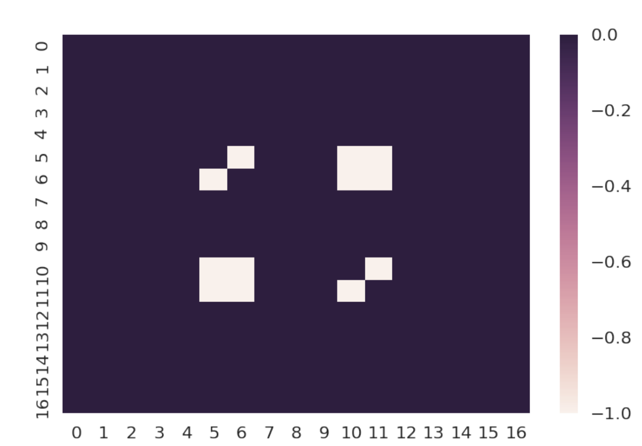

Parameters discourage multiple atmospheric labels

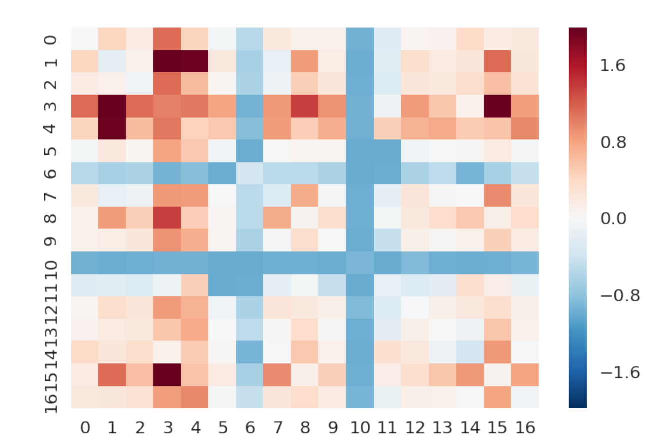

Effect of co-occurrence loss: Blue regions co-occur less often in predictions

Co-occurence of (cloudy, haze, partly cloudy) with clear:

No co-occurence loss: [6.81e-3, 3.51e-2, 2.60e-2]

With co-occurence loss: [8.27e-05, 2.88e-04, 1.99e-04]

Conclusion and Future Directions

- Best F2: 0.8925, with lot of room to improve
- Label correlations are important in multi-label classification
  - CNN-RNN approach from CVPR 2016
- Models trained on Imagenet can be useful on very different tasks
  - First few layers are likely broadly applicable
- Robustness to labeling noise is complicated
  - Model might be more accurate than labelers, but be penalized