Road Detection Using Satellite Imagery

David Kwok, Soumya Patro, and Kate Wharton
{kwokd, sopatro, kateow} @stanford.edu

CS230
Deep Learning
Spring 2018

Introduction

- **Big picture goal:** estimate population of refugee camps
- This project focuses on **road detection**, a first step in estimating population
- **Satellite imagery** is key input because it is low-cost, open source, and offers near real-time visibility

Data & Augmentation

- **Satellite Images:** Publicly available satellite images from Rotterdam
  - 10,000 RGB tiles of 256x256x3
  - Resized images to 64x64x3
- **Roads Ground Truth Label:** Shapefiles indicating coordinates of roads for the same geography -- pixel-level boolean mask

Satellite images with road labels

Methodology & Architecture

1. **Presence or Absence of Road in an image tile using 14-layer ResNet and binary cross-entropy loss function**
2. **Pixel-Level Road Detection using 14-layer ResNet with 4,096 output units (64x64)**

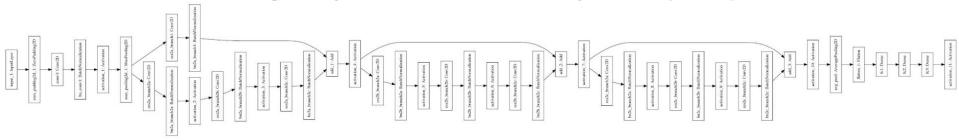

Weighted cost function to address class label imbalance (17% road pixels, 83% other):

$$-\frac{1}{p*m}\sum_{i=1}^{m}\sum_{j=1}^{p}\left(\lambda y_j^{(i)}\log\left(a_j^{[L](i)}\right)+(1-\lambda)\left(1-y_j^{(i)}\right)\log\left(1-a_j^{[L](i)}\right)\right)$$

Results

1. Binary Classification

Accuracy: 80% on dev set

ROC Curve

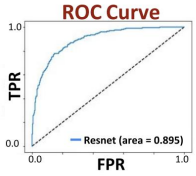

Misclassified Examples

Actual: No Road
Predicted: Road

Actual: Road
Predicted: No Road

2. Pixel-Level Road Detection

Precision: 0.15; Recall: 0.34

Input Image

Ground Truth

Predicted Labels

Conclusions

Discussion

- Model recognizes **narrow, straight edges** as roads; weak at recognizing **wide, bent roads**
- Automatic creation of labels yielded **large dataset**, but at the cost of **precise labels**
- ResNet14 performed much better than deeper networks

Future Work

- Hand label roads; data augmentation
- Investigate differences in performance between shallow and deep networks

References:

He, K., Zhang, X., Ren, S., & Sun, J. (2016). *Deep Residual Learning for Image Recognition*. Proceedings of CVPR, pp. 770-778.

Taspinar, A. (2017). *Using Convolutional Neural Networks to detect features in satellite images.* Blog post and Github repository. <https://github.com/taspinar/sidl>.