Coastal Marsh Classification

Freya Chay, Nikhil Bhattasali, and Nolan Walsh
Department of Computer Science, Stanford University

Data Preprocessing

- Split imagery into 256X256 pixel tiles
- Subsample tiles for a dataset with 20% marsh pixels
- 80:20 train:validation split

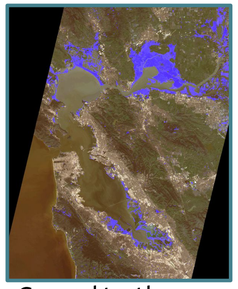

Ground truth (Bay Area)

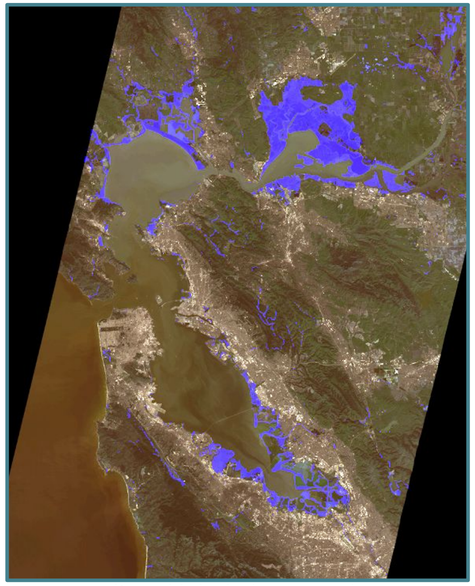

Predicted labels (Bay Area)

Model

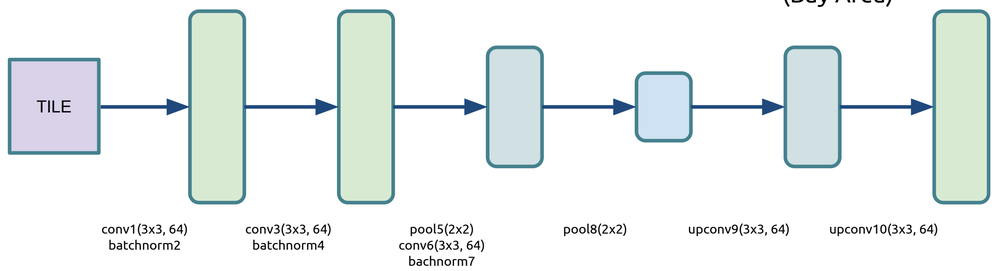

Intro

- Tidal salt marshes are one of the most effective natural habitats for sequestering carbon, but the best estimate for total land area covered by salt marsh is 2.2 - 40 million hectares.
- A more accurate estimate could facilitate political advocacy to preserve/restore salt marshes, incorporate them into carbon offset markets, and enable some tracking of area change.
- Traditional approaches for land cover classification use decision trees, random forests, and SVM; recently, neural networks have been used, but not applied to this important problem.

Problem Statement

- Estimate the total area of tidal salt marshes in the Bay Area, and develop a model that can scale to future global estimates.
- Use convolutional neural networks to perform pixelwise classification on satellite imagery with multiple spectral bands.
- Evaluate model performance using accuracy, precision, and recall for quantitative metrics, as well as visual comparison for qualitative metrics.

Data

Raw satellite imagery comes from the Planet Labs Open California dataset. The images have a resolution of 5x5 meters per pixel and have 5 data bands: Blue, Green, Red, Red Edge, and Near Infrared.

Labels were generated using the California Department of Fish and Wildlife Marine Resources GIS map of California Coastal Wetlands. In the region of consideration, only 3.5 percent of pixels are labeled as coastal marsh.

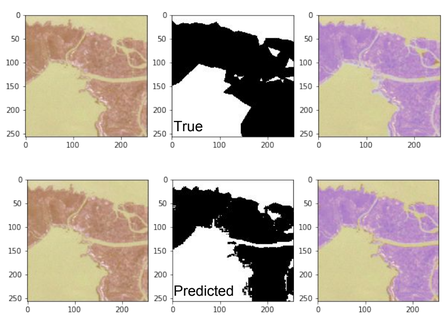

Label Noise

This validation set tile demonstrates the inherent noise of the ground truth labels available.

Experiments

Training configuration was chosen after experimentation along the following dimensions:

- Tile size
- Tile subsampling
- Model depth and filter size

Results

On the validation data set, our model achieves:

precision: 0.64757239773142627
recall: 0.57054110988974571

Using imagery from a new region (Morro Bay) as a test set, we found that the model does not perform well, indicating that our model is sensitive to perturbations in satellite imagery. The next step is to retrain the model on a dataset that is more diverse, i.e. imagery captured in different regions and by different satellites.

Ground truth (Morro Bay)

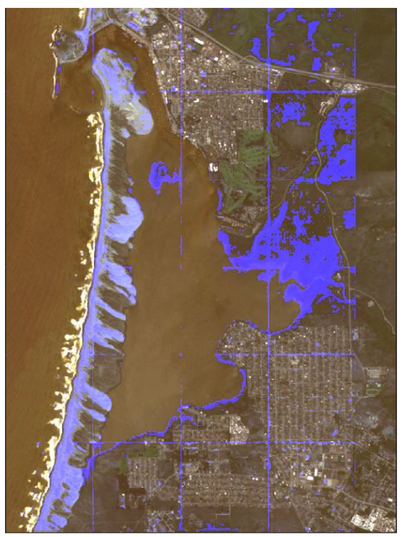

Predicted labels (Morro Bay)