Shallow Cloud Classification in Satellite Images

Timur Reziapov (timrez@stanford.edu)

Motivation

Shallow clouds make up an important part of Earth's climate but are poorly understood and represented in current models. Clouds play a critical role in climate because they reflect a significant part of sunlight back into space which helps cool Earth down. Understanding shallow cloud organizations can help improve our long-term climate projections and even short-term dangers such as hurricanes and typhoons.

Problem definition

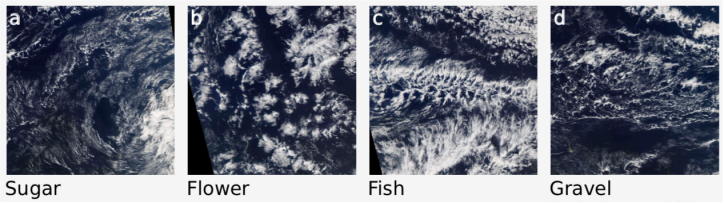

- Detect, classify and segment cloud formations from 4 classes
- Evaluate on Dice coefficient (DSC)
- Training loss:

$$DSC = \frac{2|X \cap Y|}{|X| + |Y|}$$

$$\sum_{c \in \{Sugar, Flower, Fish, Gravel\}} BCE_c(y_t, y_p) + (1 - DSC(y_t, y_p))$$

Dataset

- 10,000 (RGB) images downloaded from NASA Worldview
- Original resolution is 1400 by 2100 pixels
- Run length encoding used for mask segments
- Subjective crowd-sourced label masks
  - Each is a union of boxes from 3 different scientists
- A lot of blank spaces: water or black bans from orbit image stitches
- Class distribution is fairly well-balanced
  - Sugar (31.7%) and Gravel (24.8%) most frequently observed
- Most of the images have 2 labels/class masks

Methods and models

- Architectures: UNet and FPN
- Alternative encoding backbones: ResNet34, EfficientNetB3

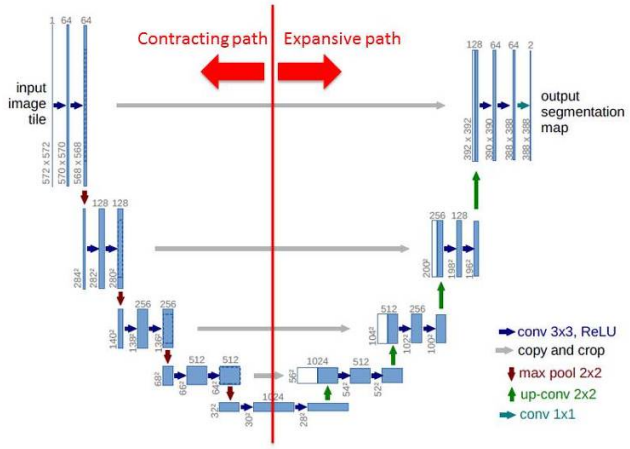

- Inputs are scaled down to 320x480 pixels
  - Motivation: Fit GPU and speed up training
  - Batches of 4,8 or 16
- Image augmentations used for reducing overfitting:
  - Rotations, flipping, shifting
- Backbones for better feature extraction
- Pre-Trained backbone models for faster training convergence

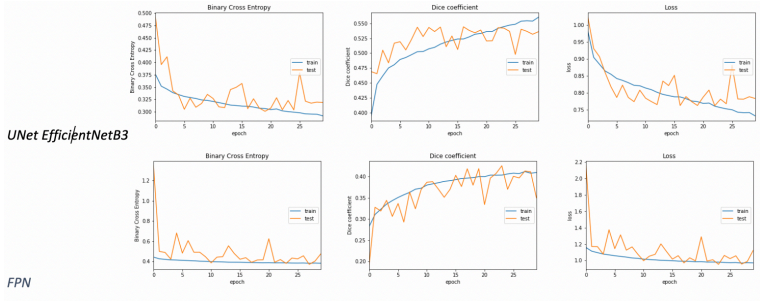

Training (blue) and validation (orange) metrics: BCE, Dice, Loss

Model	Train BCE	Validation BCE	Train DICE	Validation DICE	Test DICE
UNet	0.3184	0.3273	0.5052	0.5192	0.514
UNet ResNet34	0.306	0.316	0.533	0.543	0.539
UNet EfficientNetB3	0.311	0.306	0.554	0.544	0.602
FPN	0.3840	0.3790	0.4039	0.4259	0.404
FPN ResNet34	0.3735	0.3954	0.4282	0.4472	0.495
FPN EfficientNetB3	0.3641	0.3695	0.4419	0.4504	0.516

Results and discussion

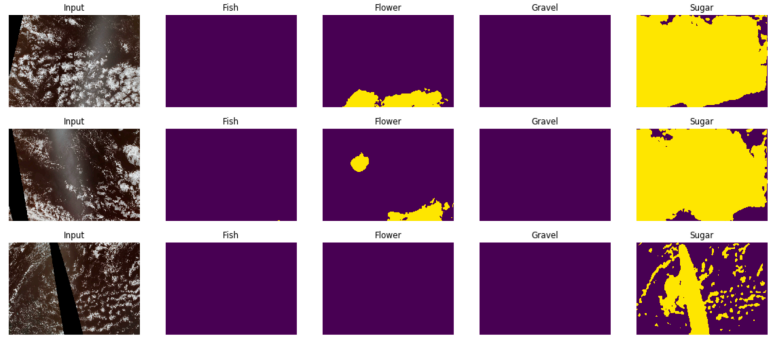

- Models learned to created segmented masks even when training on boxes
- The best model, UNet with EfficientNet B3 backbone, achieved a Dice score of 0.602 and is able to show clear group segments without any thresholding

Future work and improvements

- Separate classification and segmentation performance by using and training separate models
- Need to do a sweep for optimal thresholds on the predicted class labels and segmentation masks