

Lewis Fishgold
Azavea

The code for this project is open source and is located at
<https://github.com/azavea/keras-semantic-segmentation>.

Entry 1

This entry uses a simplified version of U-Net which is an encoder-decoder architecture for semantic segmentation. The model was trained from scratch without the use of pretraining, and was not combined with any other models as part of an ensemble. We trained the model for 100 epochs with 4096 samples per epoch with a batch size of 8 using the Adam optimizer. The learning rate started at 0.001 and was divided by 10 after the validation loss did not improve for ten epochs. The input to the network included red, green, blue, infrared, NDVI, and normalized depth channels which were concatenated together. The training data consisted of 80% of the labeled data, and was randomly augmented using 90 degree rotations and horizontal and vertical flips. The training process took ~12 hrs on an NVIDIA Tesla K80 GPU.

The experiment configuration file for this entry is located at
https://github.com/azavea/keras-semantic-segmentation/blob/develop/src/experiments/3_21_17/unet_all.json

Entry 2

This entry uses a ResNet50 pretrained on ImageNet which was converted into a Fully Convolutional Network (FCN) architecture. It contains three skip connections, from the last 32x32, 16x16, and 8x8 layers, which are upsampled using bilinear interpolation. We trained the model for 100 epochs with 4096 samples per epoch with a batch size of 8 using the Adam optimizer with a learning rate of 1e-5. The input to the network only included infrared, red and green channels. The training data consisted of 80% of the labeled data, and was randomly augmented using 90 degree rotations and horizontal and vertical flips. The training process took ~8 hrs on an NVIDIA Tesla K80 GPU.

The experiment configuration file for this entry is located at
https://github.com/azavea/keras-semantic-segmentation/blob/develop/src/experiments/4_7_17/fcn/fcn_resnet_iirg_cross_4.json

Entry 3

This entry uses cross-validation over 5 models (each trained with a different 20% validation set), and combines them together as an ensemble. Each model is the same type as described in Entry 2, and the predictions of the five models are combined by averaging the individual output probabilities.

The experiment configuration files for this entry are located at
https://github.com/azavea/keras-semantic-segmentation/blob/develop/src/experiments/4_7_17/fcn/