Machine Learning for Autonomous Jellyfish Tracking

Adam Wiktor
Aerospace Robotics Lab, Stanford University

Motivation

Scientists at the Monterey Bay Aquarium Research Institute need to collect observations of jellyfish over a long period of time using a remotely operated underwater vehicle. Autonomy can increase the time limit.

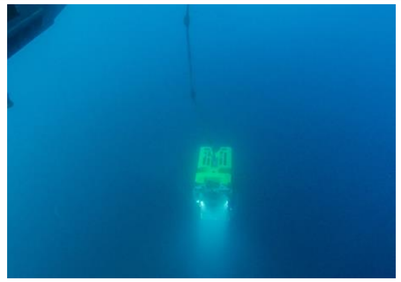

Image Processing

To isolate the jellyfish in an image, one standard technique converts the image to grayscale and applies a threshold, removing all dark pixels. A blob tracker can then be used to find the largest contiguous region. Dark regions of the jellyfish are lost in the process.

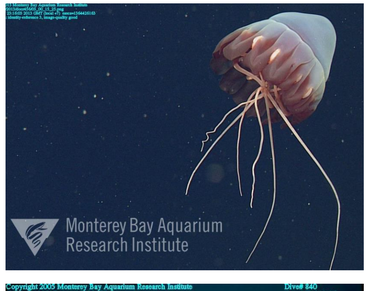

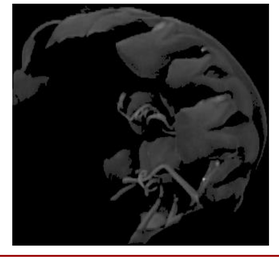

EM for Image Segmentation

The Expectation-Maximization algorithm can also be used to isolate the jelly. Assuming the image is a bimodal distribution, EM is used to estimate the mean and variance. Every pixel is a 3D data point, and by comparing pixels to both distributions the algorithm eliminates any pixels labeled as “background”. EM on the full color image preserves more features, but is slower.

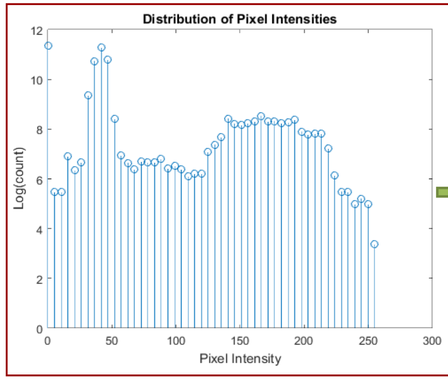

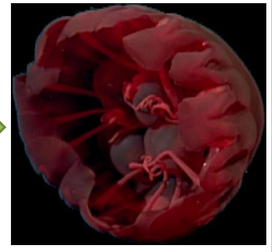

Jellyfish Orientation Determination

The orientation of each image in the training set is estimated manually, and jellies are bucketed into one of twelve orientation categories.

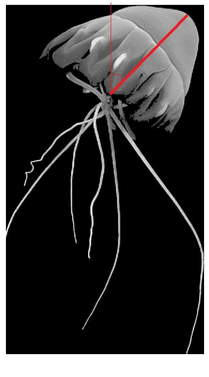

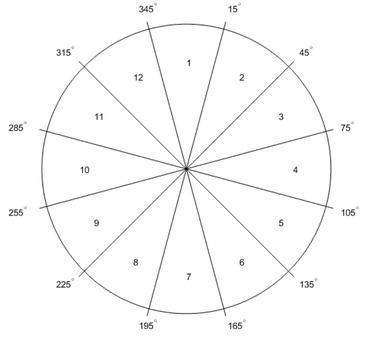

K-Means for Feature Generation

Several hundred SURF features are generated for every image, and K-Means clusters the SURF features of each category to find the most predictive features.

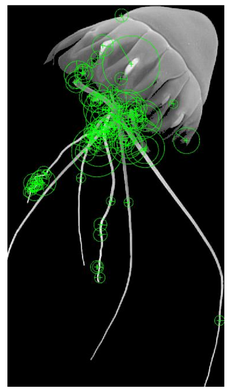

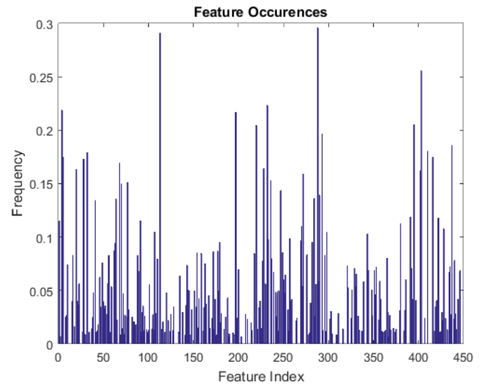

SVM Orientation Classifier

An SVM is trained on the feature vectors generated by the K-Means algorithm using a one-vs-all scheme. Cross-validation is used to ensure that the method will generalize to unseen images. The SVM classifies the orientation correctly for 96% of the training images, and is within 1 bucket for 87% of test images.

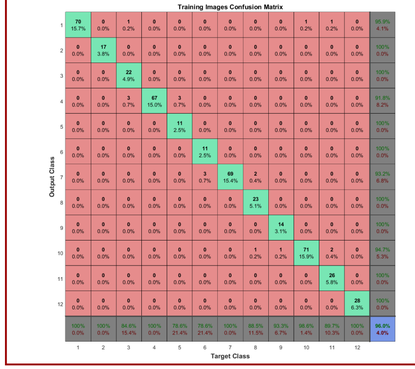

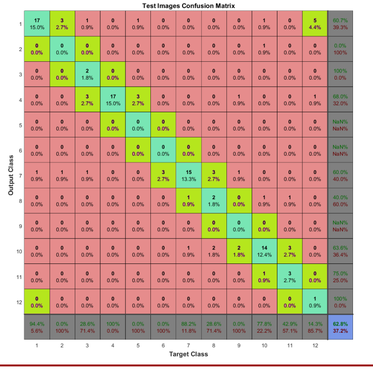

Images courtesy of Monterey Bay Aquarium Research Institute