

Convolutional Neural Network Structure for Ocean Health

Cater, Evan

Plankton are critically important to our ecosystem, accounting for more than half the primary productivity on earth and nearly half the total carbon fixed in the global carbon cycle. They form the foundation of aquatic food webs including those of large, important fisheries. Loss of plankton populations could result in ecological upheaval as well as negative societal impacts, particularly in indigenous cultures and the developing world. Plankton's global significance makes their population levels an ideal measure of the health of the world's oceans and ecosystems. Underwater research cameras take millions of photos a day, so using those images to monitor plankton populations is unrealistic for the human time-scale. An automatic image classification system would have broad applications for the assessment of ocean and ecosystem health. We trained large, deep convolutional neural networks (CNN's) to classify 30,000 varied-resolution plankton images in the Kaggle Data Science Bowl dataset into 121 different classes. We demonstrate near-human performance on the cross-validation set with our best model, which yields 92% accuracy in top-5. We explore techniques such as dropout and data-augmentation to reduce overfitting. The architectures tested are variants of a typical LeNet / Oxford Net, with promising results. Ensembling methods such as Knowledge Distillation were also explored for a balance between the large parameters of a group of models and the efficiency of a small one. Code was contributed to the Caffe GPU-Deep Learning framework for the ensemble methods. We put an emphasis on speed, because fast processing is crucial in global classification systems. We achieved 5 milliseconds per image classification.

Awards Won:

Consortium for Ocean Leadership: First Award of \$3,000