HullSeek

Kunal Mozumdar, Matthew Salinger, Joshua Ruiter, and Hirak Bandyopadhyay

We developed a deep learning based toolkit to identify biofouling regions from submerged architectures. Biofouling, the accumulation of marine lifeforms on submerged surfaces, results in degradation of maritime infrastructure and contributes to the spread of invasive species. It presents major challenges to both efficiency and environmental impact. Our model aims at early detection of biofouling regions from underwater images aiding maritime engineering and shipping industry.

Our computer vision model pipeline can automatically detect biofouling regions and estimate the extent of the biofouling by analyzing underwater images. The model uses a fine-tuned YOLO (You Only Look Once) model that is retrained on not only different biofouling agents but also real-world expert categorized samples of biofouling. We also developed efficient data pre-processing pipelines to segment fouling regions in murky underwater images and class imbalance due to the large category of biofouling agents. Our model shows improved performance compared to traditional deep learning tools which rely on large training datasets requiring expert labeling of fouling regions.

Objectives

- **Data Preprocessing:** Collection of underwater biofouling images, and annotation with bounding boxes using SAM3, to enable YOLO usage.
- **Baseline Models:** Initial approach is developing simple CNN based classifiers to predict extent of biofouling as 3 classes - Nil, Medium, and Heavy and perform comparison to a baseline Random Forest model.
- **Spatial Detection:** Develop a modern deep learning based approach that detects the spatial location and extent of the fouling on a given image using a pre-trained YOLO model.

Data specifications and preprocessing

We trained our model using an expert labelled real-world dataset (Dataset-1/Nature Scientific Report). This dataset consists of underwater images classified into nil, medium, and heavy fouling. Utilizing ResNet and EfficientNet, we develope initial baseline classifiers and used a Class Activation Map (CAM) to identify elements of each image that contribute to classification. Most models misclassified level 1 as level 0 due to the data and class imbalance. In addition, we used a smaller dataset (Dataset-2) with bounding box annotations for 7 different classes including bounding boxes usable by YOLO.

Automated data annotation pipeline

Dataset-1 lacked bounding boxes around fouled regions that are essential for training object detection models like YOLO. For that purpose we developed a fully automated pipeline creating bounding boxes using a segmentation model, SAM3 using a combination of text based prompts and manual annotation. This annotated data was used to fine-tune the YOLO model. Model is fully deployable at any dataset (providing a fouling level increases the efficiency of the code).

YOLO

We constructed a YOLO model that can detect and spatially locate biofouling on a specific image. We fine-tuned the YOLO model on both the datasets using different architectures and obtained a model accuracy of 60% on noisy real-world data, significantly outperforming baseline models. Our model also removes the issue of class imbalance and accurately is able to classify the no biofouling regions.

Future directions

- Automation of data annotation pipeline by removing the need for manual prompting
- Extension of scope to classify biofouling type, to aid mitigation strategy
- Real-time detections from video data

GitHub