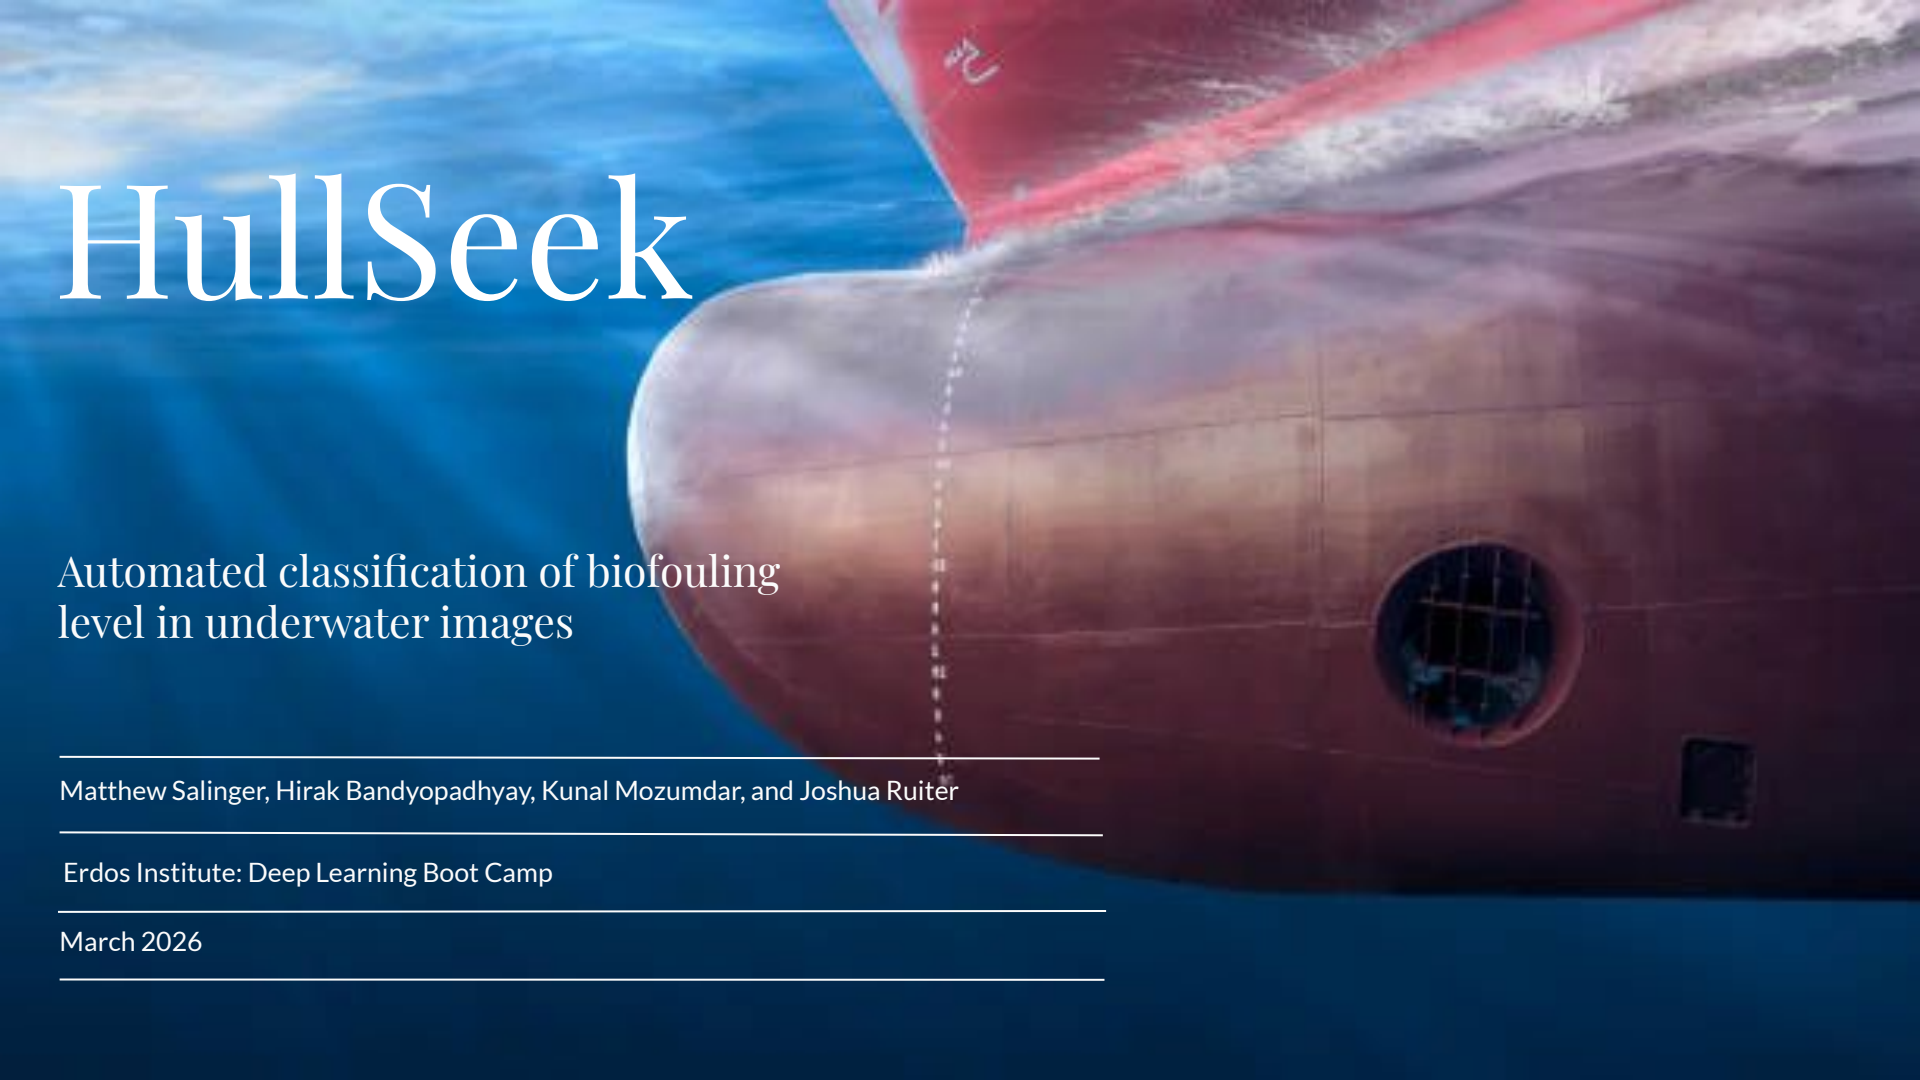

HullSeek

Automated classification of biofouling level in underwater images

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

Erdos Institute: Deep Learning Boot Camp

March 2026

Outline

01
Biofouling: causes and impacts

02
Data sets

03
EDA and baseline model (random forest)

04
ResNet, EfficientNet, CAM

05
Data augmentation with SAM3

06
YOLO

Biofouling: Causes and Impact

- Abnormal accumulation of marine life forms on submerged architectures like ship hulls, underwater pipelines, offshore platforms etc.
- Leads to transfer of invasive species, large-scale maritime infrastructure loss and fuel consumptions in vessels.

Objective :

Develop a deep-learning based model for fast automated detection of biofouling.

- Collect underwater image data from various sources.
- Annotate datasets with segmentation of biofouled regions.
- Train Object detection model like YOLOv8 to identify and classify biofouling regions.

Data sets

Biofouling (data set 1)

~10,000 underwater images

Labelled with SLoF (Simplified Level of Fouling) classification 0, 1, or 2

Heavy class imbalance (mostly SLoF 0)

Using SAM3, we annotated ~60% of the SLoF 1/2 images with YOLO-compatible bounding boxes, resulting in around 1400 usable images

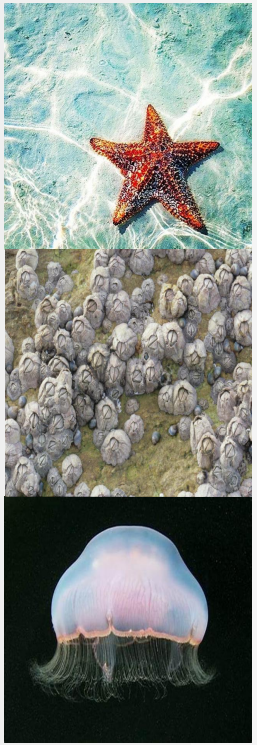

Fouling detection (data set 2)

~800 underwater images

Complete with bounding boxes data and class labellings, tailored for YOLO

Exploratory data analysis and baseline models

Feature engineering: computed some complexity-based features such as *entropy* and *gradient energy*, but none highly correlate with SLoF

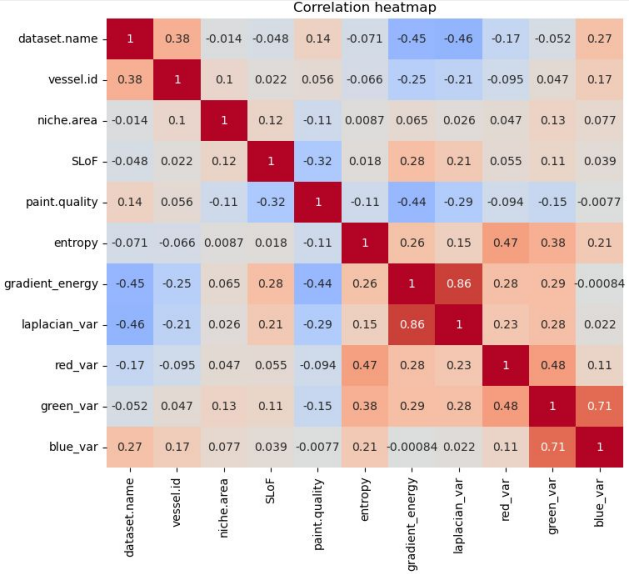

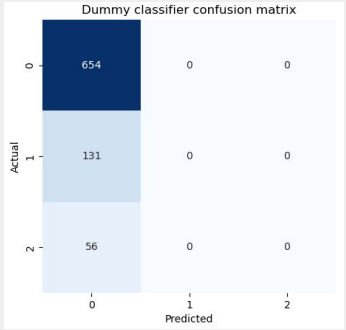

Due to class imbalance, **dummy classifier** gets high accuracy.

0.78 accuracy
0.26 precision
0.33 recall

Random forest is a marginal improvement over dummy.

0.61 accuracy
0.53 precision
0.52 recall

The highest ranking feature by importance was *paint.quality*, above all engineered features.

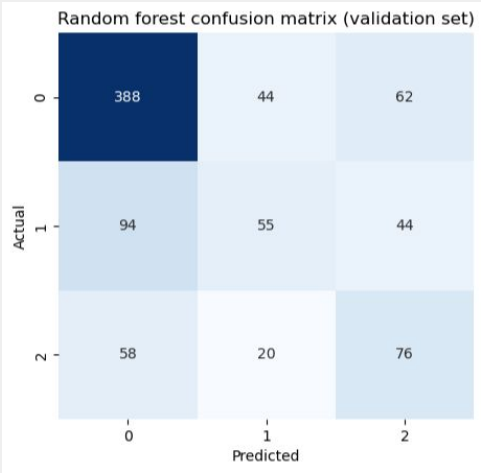

CAM-Based Biofouling Coverage Estimation

- **Model Architecture:** EfficientNet-B3 with Class Activation Mapping (CAM) for weakly supervised spatial localization
- Three-class classification (SLoF 0: Clean, SLoF 1: Patchy, SLoF 2: Heavy) + pixel-level coverage estimation
- **Key Innovation:** GAP (Global Average Pooling) enables extraction of spatial information from image-level labels, generating heat maps showing fouled regions
- Caveats : High inaccuracy in classifying SLoF 1 (medium biofouling) Category

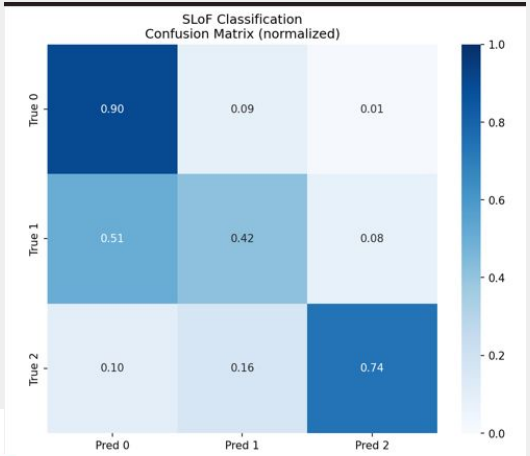

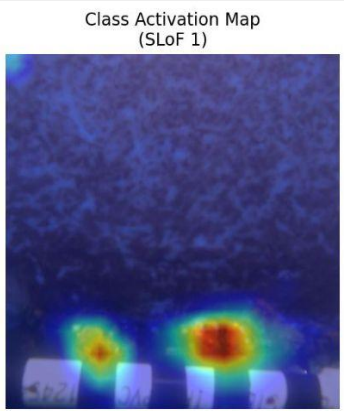

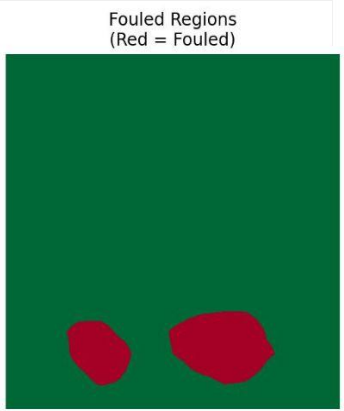

COVERAGE ESTIMATION

Predicted SLoF: 1

Coverage: 7.2%

Interpretation:

Light patchy fouling

Method: Class Activation Mapping (CAM) + threshold

CNN based classifiers: ResNet

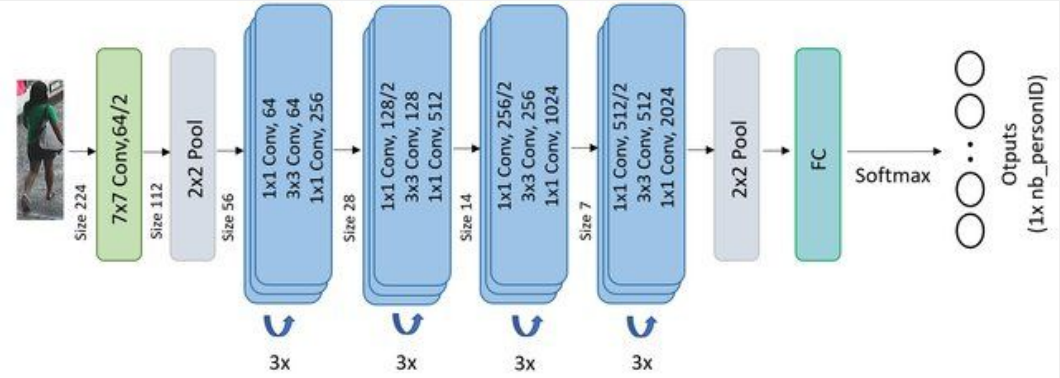

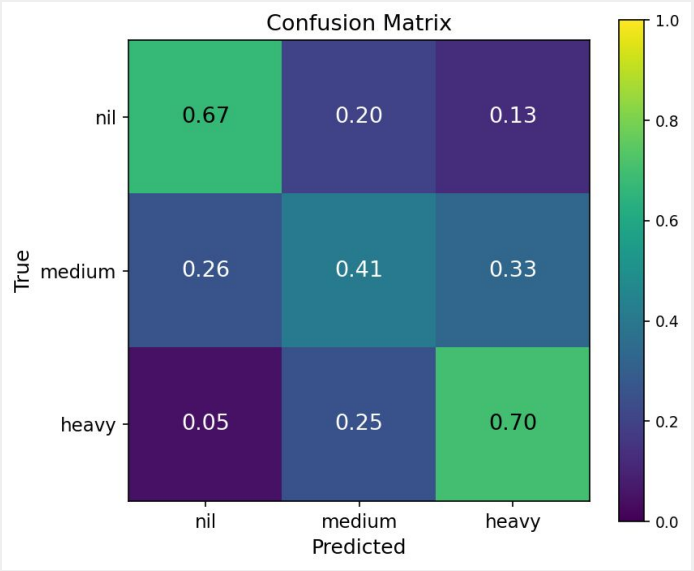

- **Model Architecture :** ResNet Classifier
- Three-class classification (SLoF 0: Nil, SLoF 1: Medium, SLoF 2: Heavy)
- Low accuracy on the medium biofouling class.
- Generally fast.

Fouling Level	SLoF 0	SLoF 1	SLoF 2
Precision	0.851	0.366	0.456
Recall	0.668	0.409	0.701

YOLO v8m Classifier :

- Decision to Merge Medium and Heavy SLoF category to a single “Fouled” category.
- Robust separation between Nil and Fouled, with optimized IOU set to 0.45.
- Better performance than ResNet and Efficient Net.

Fouling Level	Nil	Fouled
Precision	0.89	0.74
Recall	0.93	0.65
F1 score	0.91	0.69

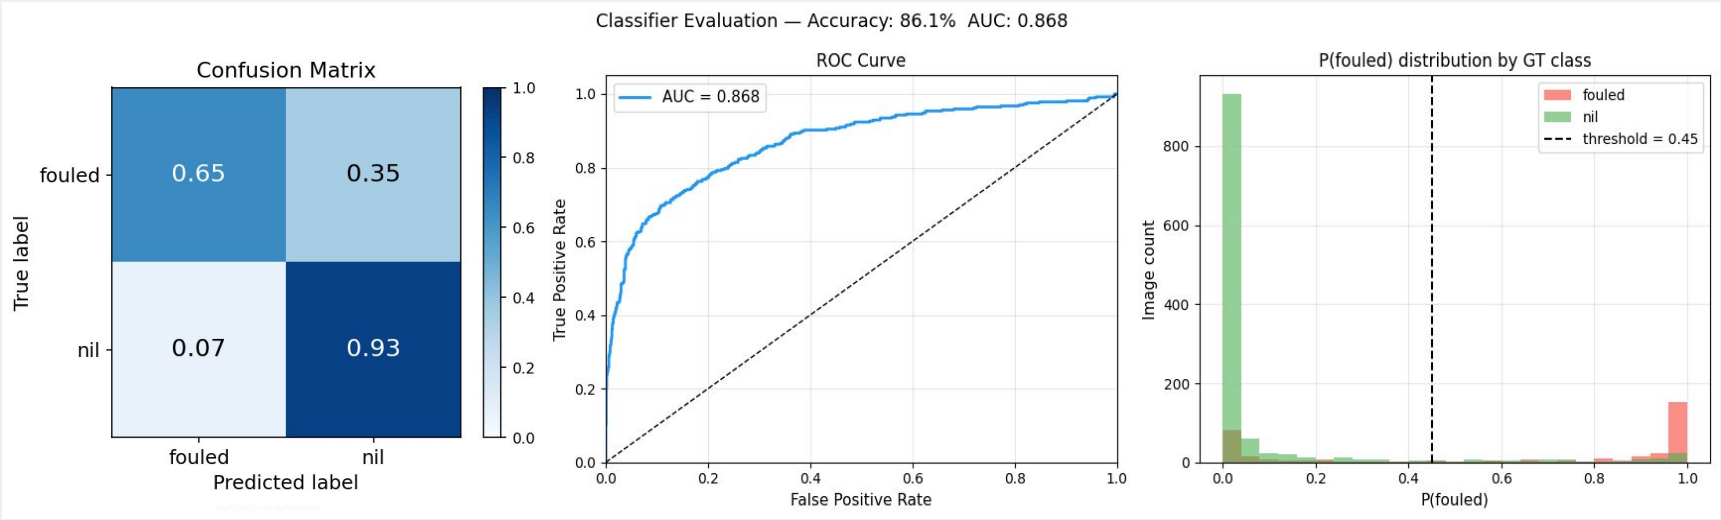

Automatic Segmentation using SAM3

- **Challenges:** To annotate Dataset-1 (10k images) with bounding boxes around biofouling regions to train a YOLO model.
- Selected only the SLoF 1 and 2 category (medium and heavy cases)
- Deployed a combination of Text based prompts and some manual prompts to create the “concept” based prompts for segmentation.

Step 1 : Text based prompts on 100 test cases

Step 2 : Failure cases were manually annotated with hand drawn boxes.

Step 3 : Merge and create concept prompts and rerun SAM3 for all data

Results : 60% of all sLoF category 1 and 2 was annotated and ready for YOLO training.

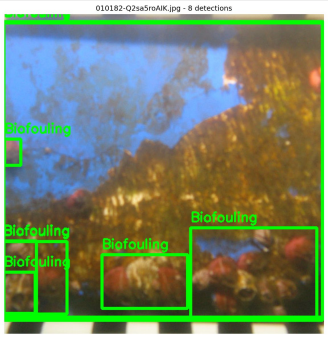

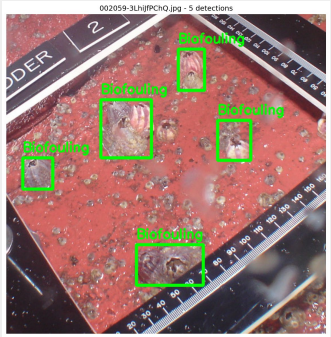

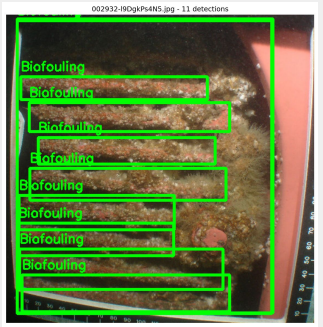

Caveat : All sLoF 1 or 2 category now just labeled as biofouling.

YOLO v8s: Object detection using bounding boxes

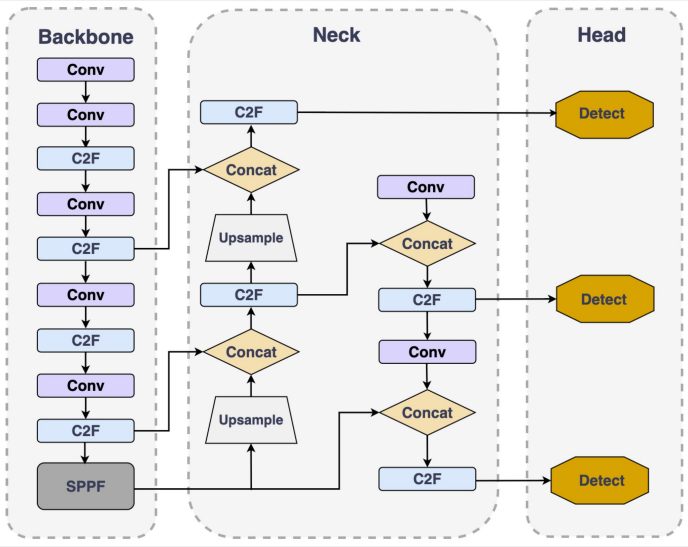

- YOLO is pretrained on general object categories from COCO dataset.
- Geometric and Photometric image augmentation implemented to improve robustness, reduce overfitting and mitigate class imbalance.
- Fine-tuned on Dataset 2: **Clean Surface** vs **Common biofouling agents** -> Barnacles, Jellyfish Polyps, Mussels, Rust, Starfish and Worms.

Precision	Recall	mAP50	mAP50-95
0.728	0.558	0.622	0.393

- Option A: Validation on SAM3 annotated Biofouling Dataset, poor result due to incompatible image features, spatial resolution.
- Option B: Retrain on annotated Biofouling Dataset, better performance on clean surface, poor recall for Biofouling.

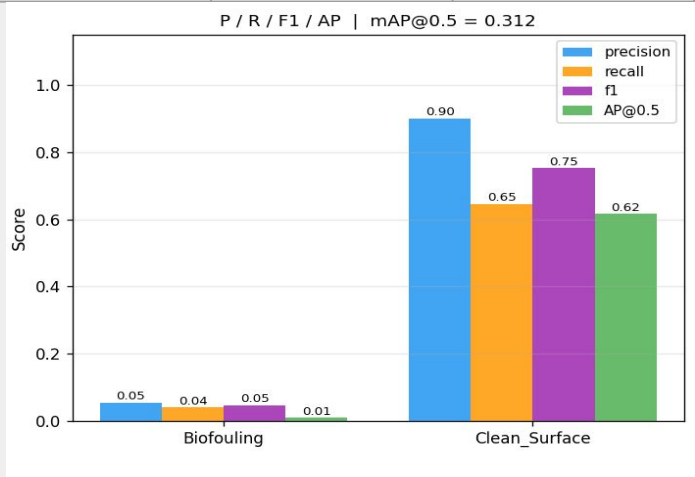

YOLO Two Task Pipeline: Classification and Detection

- **Task 1**: YOLO v8m cls model implemented for “nil” vs “biofouling” classification, and only “biofouling” classified images processed for step 2.
- **Task 2**: Training YOLO v8s model on SAM3 Annotated “Biofouling” only data (~1400 images). Then use it on “biofouling” classified images from Task1 to create bounding boxes and calculate biofouling surface area.
- High recall in “nil” means very few clean surface selected for further processing in step 2.

- Model learns to implicitly differentiate between biofouling surface vs background, i.e. clean hull, ocean water, equipments.

```
FULL PIPELINE — VALIDATION SET RESULTS
Total images          : 1540
Classification acc.   : 86.10%
fouled correct        : 237 / 366  (64.8% recall)
nil correct           : 1089 / 1174  (92.8% recall)
Task 2 ran on         : 322 images
Fouled w/ boxes found : 298 / 322
Mean fouling area     : 65.4% (fouled images only)
```

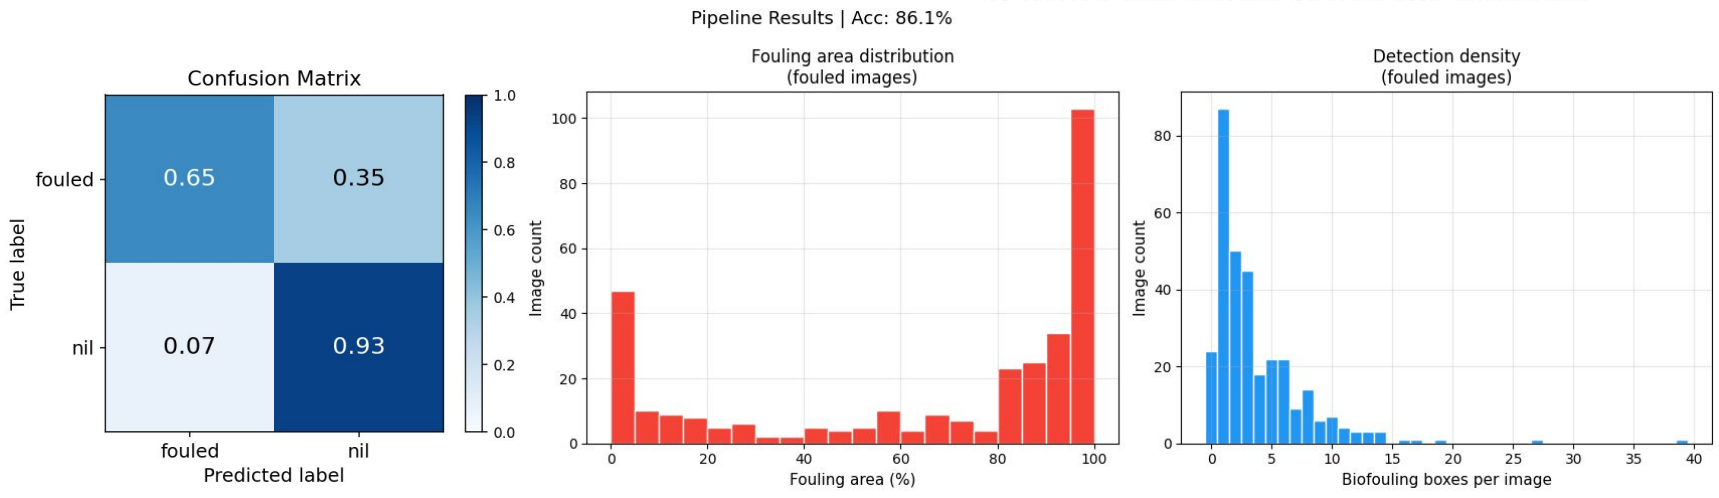

YOLO Two Task Pipeline: Classification and Detection

Clean Surface Test Set

Biofouling Test Set

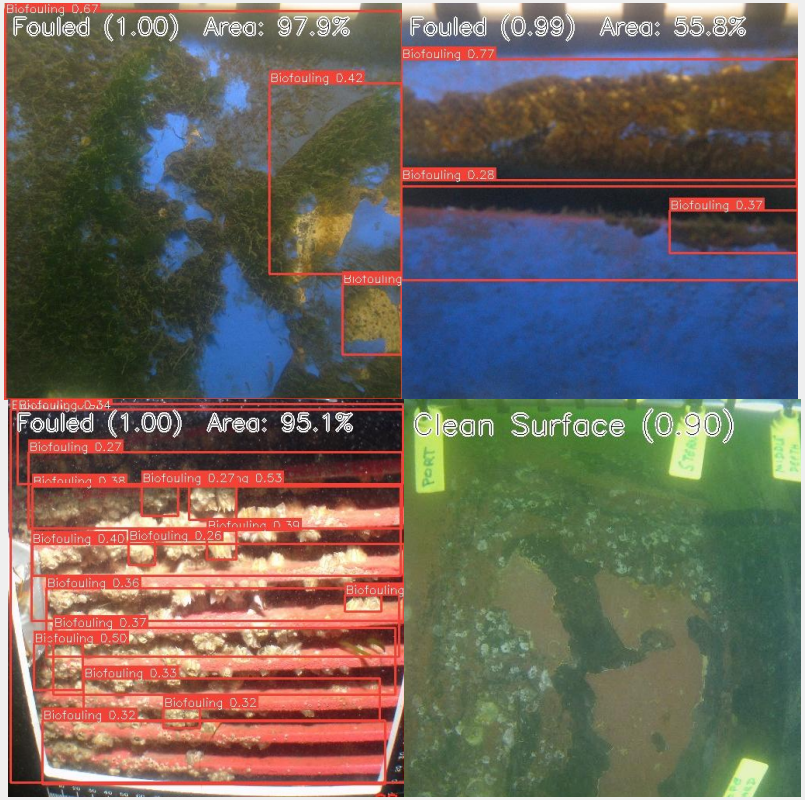

Takeaways

- Biofouling is an important issue for the shipping industry, and environmental concerns.
- Underwater image detection comes with many challenges.
- With SAM3 to augment with bounding boxes, we trained a YOLO model to successfully classify none vs some fouling.