


Development of the Habitat Aware Reconnaissance and Imaging Module (HARIM) with AI for Classification, Sensor Fusion and Habitat Modeling

NOAA SBIR Phase II awarded February, 2021

Scott Gallager , Steve Lerner and Cameron Fairclough
WHOI and CoastalOceanVision, Inc.



Habitat Aware Reconnaissance and Imaging Module (HARIM)

Specifications



- Vehicle: REMUS 600

Sensors

- Stereo PT Grey 10 Mpixel cameras, 12mm lenses
- TX2 6 core processor
- CTD, Chlorophyll, turbidity
- Sidescan
- Plankton imaging and classification (CPICS)


Capabilities

- On-board stereo imaging, light-field and color correction, rectification, point cloud production, and target segmentation
- Benthic target acquisition using sidescan followed by classification using stereo imaging
- On-board plankton classification
- Habitat characterization, spatial analysis, dynamic spatial sampling based on habitat type
- 10 hour deployments up to 600m depth in 3 kt current



REMUS-600 with HARIM: Phase I Sea Trials






Stereo Cameras



CPICS
plankton
camera

6 strobes





Autonomous Habitat Classification and Smart Navigation Strategy

4 x GPU Deep Learning Workstation

- Build training sets
- Build models
- Validate models



*Deploy models on
embedded hardware*




NVIDIA TX2 ARM
processor + GPUs




Run in fully autonomous
dynamic sampling mode





Two approaches

- *Holistic- complete image (substrate)
- *Targets- segmentation (scallop, seastar, fish, sand dollar, etc.)



- *8 images /sec on ship
- *4 images /sec on R600
- *15 classes (so far)
- *88-97% accuracy
- *runs on NVIDIA Jetson TX2 in HARIM


habcamrolo Validation - True vs False Positive

Legend: TP (purple line), FP (teal line), Optimum @ tp=82.0 fp=2.0 (orange vertical line).

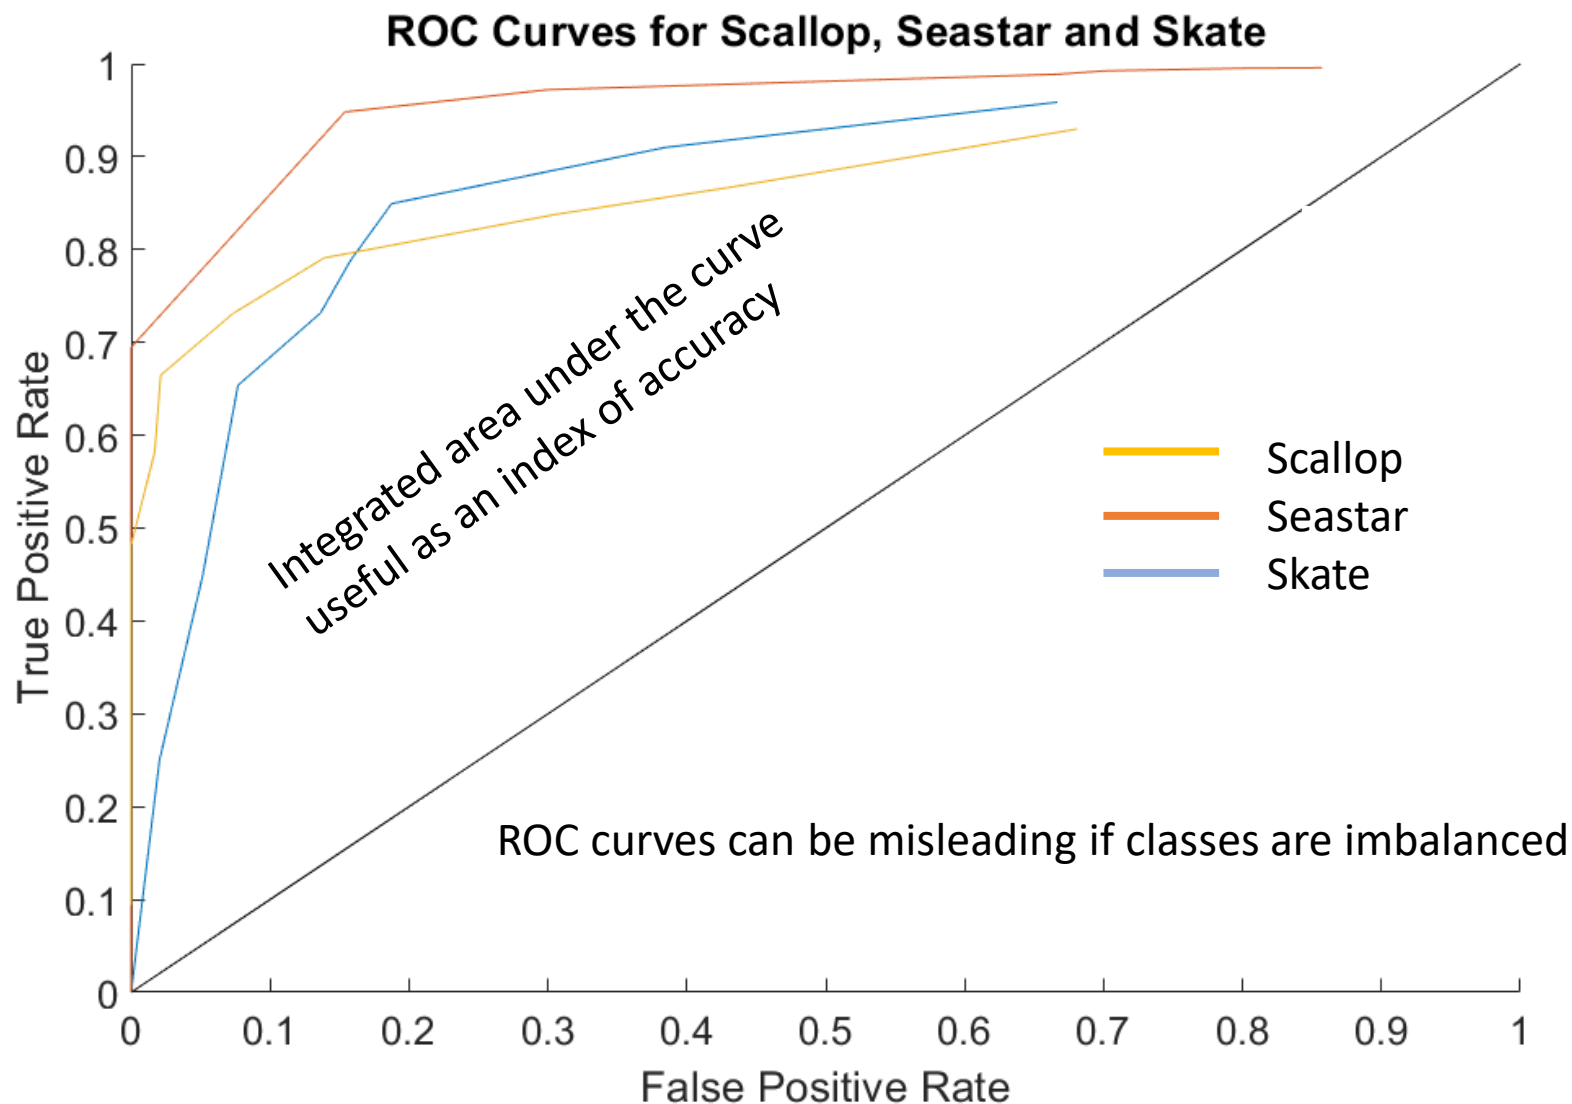
Threshold	TP (%)	FP (%)
0.1	100	100
0.2	98	45
0.3	95	25
0.4	90	18
0.5	88	15
0.6	86	12
0.7	85	8
0.8	82	2
0.9	75	0
1.0	50	0

Command Log

Automated counting and measuring scallops on the continental shelf



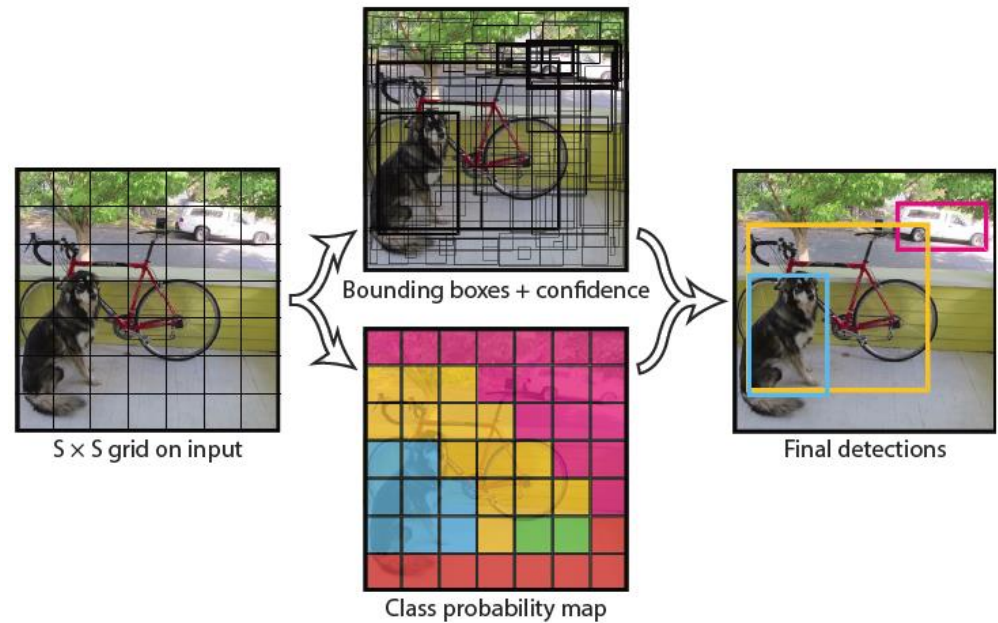
True Positive Rate = True Positives / (True Positives + False Negatives)
False Positive Rate = False Positives / (False Positives + True Negatives)



Can be 100x faster with good accuracy (mAP: 71, typical)


YOLO works in a single pass:

- * Image divided into 7x7 grid
- * Each grid predicts bounding boxes and a *confidence* score
- * x, y, w, h , IOU (Intersection over Union) of predicted box and ground truth box
- * Conditional class *probability*




The model consists of 24 convolutional layers followed by 2 fully connected layers for predicting Bbox coordinates and IoU. Each layer reduces complexity of feature space of previous layer.

Strongly based on GoogleNet replacing inception modules with 1x1 convolutions.



Codfish
Gadus morhua





Haddock
Melanogrammus aeglefinus



Cusk
Brosme brosme

Grey sole
Glyptocephalus cynoglossus



Deeplet sea anemone

Bolocera taediae



Unidentified anemone



Longfin squid
Loligo pealeii

Northern cerianthids

Cerianthus borealis


Redfish

Sebastes spp.



Sponge
Halichondria sp.







Please put on your red/cyan 3D glasses

