Projects in Computer Vision for Coral Reef Monitoring

At ECEO, we are developing a new method for monitoring coral reefs from underwater videos. Videos from coral reef sites collected as part of the Transnational Red Sea Center's expeditions in Israel, Jordan, and Djibouti. The videos are analyzed using frame-wise semantic segmentation, and simultaneous localization and mapping (SLAM) is used to create 3D point clouds from reef sites visited, where each point has both its RGB color and its semantic class (e.g. Rock, Sand, Live Coral, Dead Coral, etc.) attached.

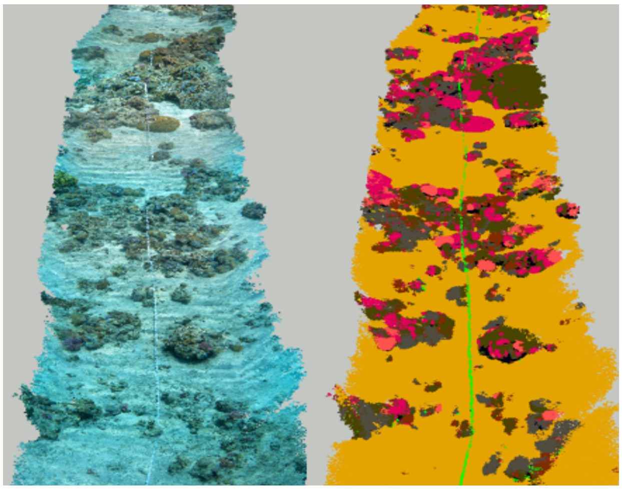

The ease with which such videos can be collected promises to increase the scalability of coral reef monitoring methods by an order of magnitude - provided the analysis of the point clouds can be reliably automated. There are a number of projects for students available in the machine learning / computer vision domain in the line of this project:

Project Idea 1: Correspondence Search in 3D Point Clouds

To make the 3D point clouds efficient for coral monitoring purposes, multiple 3D point clouds of the same area should be comparable. As the method for generating point clouds from the videos is noisy, finding correspondences within the point clouds is challenging, as there is no single rigid transformation that overlays two point clouds (see example picture below).

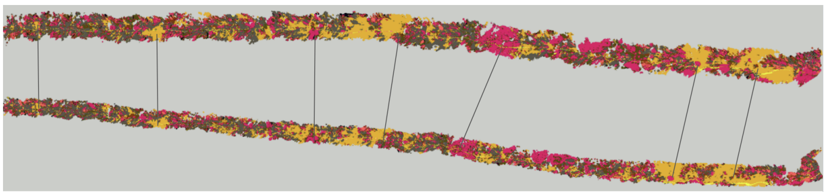

The project aims at establishing a baseline for ortho-projected 2D image matching using conventional methods like SIFT [6], and comparing this baseline to more involved methods.

More involved methods could make use of the 3D geometry at hand (e.g. using established point cloud matching methods [4]), as well as the semantic class information. One

EPFL

Jonathan Sauder
EPFL ENAC IIE
Rue de l'Industrie 17
CH–1951 Sion
Switzerland

+41 77 472 7507
jonathan.sauder@epfl.ch
https://people.epfl.ch/jonathan.sauder

idea could be to use neural networks to learn 3D correspondences in a self-supervised manner [3, 8]. This project is aimed at students that are proficient in Python, experienced with deep learning and PyTorch. Background knowledge in 3D geometry processing or 2D image matching is certainly beneficial.

Project Idea 2: Semantic Segmentation

In this project, the goal is to develop the best possible setup for determining the benthic classes of interest from the video frames. For this, a dataset of annotated image frames has been created, with a held-out test-set that allows benchmarking varying setups for training neural networks.

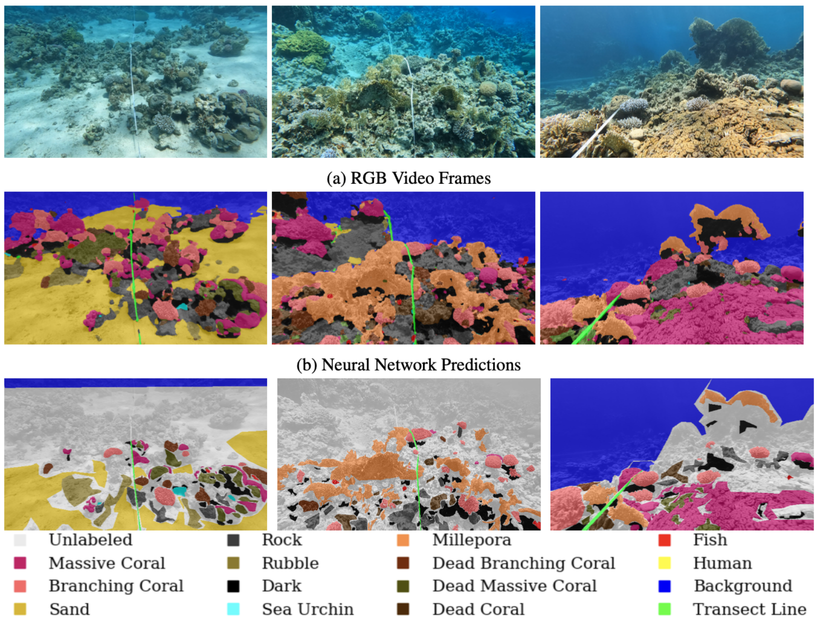

One possible direction would be to use the Segment Anything Model [5] to transform publicly available datasets of sparse annotations [2] into dense semantic segmentation labels, aiming to improve upon existing approaches based on super-pixels [1, 7].

This project is aimed at students that are proficient in Python, and have some background in machine learning. Experience with deep learning and PyTorch is beneficial.

Contact

Jonathan Sauder (jonathan.sauder@epfl.ch)

References

[1] Inigo Alonso, Matan Yuval, Gal Eyal, Tali Treibitz, and Ana C Murillo. Coralseg: Learning coral segmentation from sparse annotations. *Journal of Field Robotics*, 36 (8):1456–1477, 2019.

[2] Oscar Beijbom, Peter J Edmunds, David I Kline, B Greg Mitchell, and David Kriegman. Automated annotation of coral reef survey images. In *2012 IEEE conference on computer vision and pattern recognition*, pages 1170–1177. IEEE, 2012.

[3] Zan Gojcic, Caifa Zhou, Jan D Wegner, and Andreas Wieser. The perfect match:

EPFL

Jonathan Sauder
EPFL ENAC IIE
Rue de l'Industrie 17
CH–1951 Sion
Switzerland

+41 77 472 7507
jonathan.sauder@epfl.ch
https://people.epfl.ch/jonathan.sauder

3d point cloud matching with smoothed densities. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5545–5554, 2019.

[4] Dirk Holz, Alexandru E Ichim, Federico Tombari, Radu B Rusu, and Sven Behnke. Registration with the point cloud library: A modular framework for aligning in 3-d. *IEEE Robotics & Automation Magazine*, 22(4):110–124, 2015.

[5] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.

[6] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1150–1157. Ieee, 1999.

[7] Matan Yuval, Iñigo Alonso, Gal Eyal, Dan Tchernov, Yossi Loya, Ana C Murillo, and Tali Treibitz. Repeatable semantic reef-mapping through photogrammetry and label-augmentation. *Remote Sensing*, 13(4):659, 2021.

[8] Yiming Zeng, Yue Qian, Zhiyu Zhu, Junhui Hou, Hui Yuan, and Ying He. Corrnet3d: Unsupervised end-to-end learning of dense correspondence for 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6052–6061, 2021.

EPFL

Jonathan Sauder
EPFL ENAC IIE
Rue de l'Industrie 17
CH–1951 Sion
Switzerland

+41 77 472 7507
jonathan.sauder@epfl.ch
https://people.epfl.ch/jonathan.sauder