Automatic Image-Based Coral Polyp Analysis through Multi-View Instance Segmentation

S. Dutta[1], G. Pavoni[1], M. Corsini[1], F. Ganovelli[1], P. Cignoni[1] P. Rossi[2], E. Cenni[2], R.Simonini[2], F. Grassi[2] D. Cassanelli[2], S. Cattini[2], L. Rovati[2], A. Capra[2], and C. Castagnetti[2]

[1]CNR-ISTI, Pisa, Italy, [2] University of Modena and Reggio Emilia, Modena, Italy

Abstract

We present an automated framework for counting and measuring the polyps of Cladocora caespitosa, a Mediterranean reef-building coral. To our knowledge, the most practical method for counting polyps currently involves ecologists' visual inspection of a 3D model. However, measuring polyps from the model can lead to inaccuracies due to distortions in the reconstruction. Our method integrates deep learning-based instance segmentation on 2D images with 3D models for unique polyp identification, ensuring precise biometric extraction. The proposed pipeline automates polyp detection, counting, and measurement while overcoming the limitations of manual in situ methods. Laboratory validation demonstrates its accuracy and efficiency, paving the way for scalable, high-resolution phenotyping, and field monitoring of Mediterranean coral populations.

CCS Concepts

• ***Computing methodologies*** → ***Object recognition; Image segmentation; Shape analysis;***

1. Introduction and Motivations

Ecological monitoring is essential for assessing underwater ecosystem health and the impact of global changes on marine species. Detecting temporal variations in biogenic three-dimensional structures, such as coral colonies, requires high accuracy and fine-scale resolution, yet traditional non-destructive methods are time-consuming and imprecise. Underwater close-range photogrammetry has been widely applied in coral studies, mainly focusing on tropical species, while Mediterranean reef-building corals remain less examined.

This study tackles the challenge of automating polyp counting and measurement in Cladocora caespitosa through multi-view small object instance segmentation. Due to their dense and repetitive structure (Figure 1), ecologists typically rely on 3D models for counting, as direct image-based methods are impractical. However, the transparency and motion of polyps limit reconstruction accuracy, making biometric estimates from geometry and texture unreliable. Manual counting on 3D models is also time-consuming. To overcome this, we introduce a hybrid method that uses 2D image-based segmentation for measurement, leveraging 3D models only for identifying and disambiguating repeated instances. Laboratory tests in controlled tank environments demonstrate that our method effectively automates fine-scale biometric extraction, paving the way for improved coral monitoring. Future work will focus on adapting this methodology for in-field monitoring of Cladocora populations.

2. Methodology

In Cladocora caespitosa, key health indicators include the number of living polyps, oral disc length (LP), width (WP), and surface area of each corallite. Typically, this requires segmenting each polyp on a 3D model. Our pipeline automates this by performing instance segmentation from raw images, reconstructing a 3D model via Structure-from-Motion (SfM), transferring segmentations to the model for unique counting, and estimating biometric metrics from a selected view of each polyp. To our knowledge, no deep learning-based automatic recognition networks exist in the literature for the segmentation of coral polyp instances. In [LBK*21], the authors designed an SVM classifier for the semantic segmentation of living corals in multispectral images taken in the laboratory. Due to the lack of annotated datasets, we created a dataset for training an instance segmentation network to segment corallite oral discs using the open-source interactive annotation software TagLab [PCP*22]. TagLab has several tools that facilitate semi-automatic disc labeling, including the foundational agnostic segmentation model SAM [KMR*23]. TagLab efficiently manages training datasets. To use it, we only needed to add an exporter for the input format of the chosen automatic segmentation network, the Yolo11 [JQC23]. Yolo11 (the nano version) has a lightweight architecture (originally designed to work in real-time), allowing us to fully segment all the polyp disks from the 3D reconstruction image set in a short time. The network has been trained on about 8,000 examples of oral disk instances. On the 2,834 oral disk instances of the test dataset, the fine-tuned YOLO11 scores an average precision of 0.902 (mAP50) and 0.593 (mAP50-95) for the object localization.

DOI: 10.2312/egp.20251022

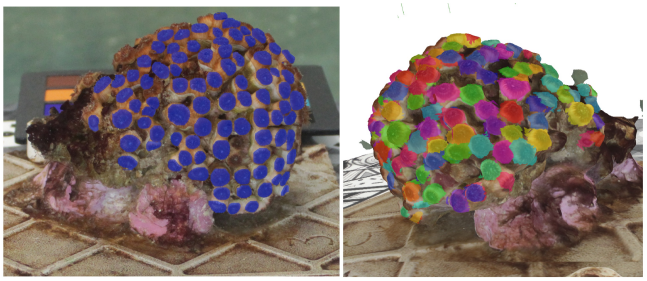

Figure 1: *The polyp disks as they appear segmented on the original image and the 3D model. Instances on the 3D models have been segmented from multiple views.*

2.1. Multi-view segmentation transfer

As discussed earlier, using the acquired images is the most effective approach for segmenting polyps. However, to effectively segment all polyps, multiple images are required, meaning that each polyp will typically appear in several images. To ensure the unique identification of each polyp, we employ image-based 3D reconstruction using Structure-from-Motion (SfM) as a common reference domain for all images. Since SfM estimates the camera parameters, we can reproject each segmented pixel region onto its corresponding part of the 3D model. We chose to associate each segmented polyps the covered mesh triangles with the number of pixels per triangle, which makes it easy to estimate the overlapping. When two or more regions overlap in the 3D domain, we can confidently determine that they belong to the same polyp, ensuring that polyps are not counted multiple times (see Figure 1). A second advantage of 3D reconstruction is that it allows us to identify the image capturing each segmented polyp from the most frontal perspective, ensuring the most accurate segmentation. This allows us to use that image and the corresponding label to determine the polyp's biometric measurements reliably. Figure 2, displays the polyps, the corresponding labels, and the extracted biometric measurements in TagLab.

3. Result

For our preliminary studies, we tested our measuring methodology on colonies of Cladocora caespitosa transplanted in water tanks. The colonies were photographed using a standard photogrammetric acquisition procedure, and the models were reconstructed and scaled with Metashape software. Marine ecologists manually counted polyps directly into Metashape to provide us with ground truth. In the current version of the automatic counting tool, the number of polyps detected automatically aligns with the manually counted ones with an average accuracy of 96.4%. For instance, in the example shown in Figure 1, our system counts 196 polyps, compared to the ground truth of 191. This discrepancy is primarily due to the approximations in the 3D model, which lead to improper surface coverage, particularly in areas near the silhouette of the projection. Considering that a small margin of error is acceptable, as some dead polyps are challenging to distinguish even for experts. Given that this significantly reduces working time, as the current procedure takes some hours per model, while our automatic counting tool processes it in just a few minutes. Additionally,

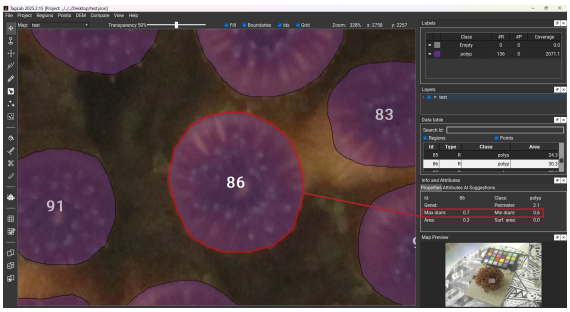

Figure 2: *Metrics are extracted from the best view of a disc and displayed within TagLab. The maximum diameter measurement represents the oral disc length (LP), while the oral disc width (WP) corresponds to the minimum diameter measurement. The disk area is indicated on the line below.*

we automatically provide the width and height of the oral section of each polyp, saving additional time of work. Moreover, these measurements are taken on the most frontal view, avoiding the need for the experts to decide where to get some measures (on the images), or to inspect the 3D model for a long time the 3D model to get them directly in 3D space.

4. Conclusion

Our framework supports coral monitoring on the field, as the instance segmentation network is robust enough to recognize Cladocora caespitosa polyps in their natural environment. At the end of the project, we will release the pre-trained network, the polyp disk dataset, and a Python-based tool for the automatic counting and measuring of polyps from multi-view stereo reconstruction for the benefit of the community.

5. Acknowledgements

This work was developed within the research project "EUPHROSYNE - Enhancing Underwater PHotogrammetRy, fluOreScence imageryY and deep learning solutions for monitoring the health status of mediterraNEan corals" (funded by Next Generation EU - PNRR - Mission 4 - Component 2 - CUP E53D23003980006).

References

[JQC23] JOCHER G., QIU J., CHAURASIA A.: Ultralytics YOLO, Jan. 2023. URL: https://github.com/ultralytics/ultralytics. 1

[KMR*23] KIRILLOV A., MINTUN E., RAVI N., MAO H., ROLLAND C., GUSTAFSON L., XIAO T., WHITEHEAD S., BERG A. C., LO W.-Y., DOLLÁR P., GIRSHICK R.: Segment anything. *arXiv:2304.02643* (2023). 1

[LBK*21] LIU H., BÜSCHER J. V., KÖSER K., GREINERT J., SONG H., CHEN Y., SCHOENING T.: Automated activity estimation of the cold-water coral lophelia pertusa by multispectral imaging and computational pixel classification. *Journal of Atmospheric and Oceanic Technology 38*, 2 (2021), 141–154. 1

[PCP*22] PAVONI G., CORSINI M., PONCHIO F., MUNTONI A., EDWARDS C., PEDERSEN N., SANDIN S., CIGNONI P.: Taglab: Ai-assisted annotation for the fast and accurate semantic segmentation of coral reef orthoimages. *Journal of Field Robotics 39*, 3 (2022), 246–262. doi:https://doi.org/10.1002/rob.22049. 1