

## Integrating Widespread Coral Reef Monitoring Tools for Managing both Area and Point Annotations

Gaia Pavoni<sup>1</sup>, Jordan Pierce<sup>2</sup>, Clinton B. Edwards<sup>2</sup>, Massimiliano Corsini<sup>1</sup>, Vid Petrovic<sup>3</sup>, Paolo Cignoni<sup>1</sup>

<sup>1</sup> Istituto di Scienza e Tecnologia dell'Informazione A. Faedo, Pisa, Italy  
(gaia.pavoni, massimiliano.corsini, paolo.cignoni)@isti.cnr.it

<sup>2</sup> CSS Inc., Under contract to NOAA, National Ocean Service, National Centers for Coastal Ocean Science,  
(jordan.pierce, clinton.edwards)@noaa.gov

<sup>3</sup> Computer Science and Engineering, University of California, San Diego  
vid@ucsd.edu

**Keywords:** automatic recognition, classification, segmentation, coral reefs, structure-from-motion, large-area imaging

### Abstract

Large-area image acquisition techniques are essential in underwater investigations: high-resolution 3D image-based reconstructions have improved coral reef monitoring by enabling novel seascape ecological analysis. Artificial intelligence (AI) offers methods for significantly accelerating image data interpretation, such as automatically recognizing, enumerating, and measuring organisms. However, the rapid proliferation of these technological achievements has led to a relative lack of standardization of methods. Remarkably, there are notable differences in procedures for generating human and AI annotations, and there is also a scarcity of publicly available datasets and shared machine-learning models. The lack of standard procedures makes it challenging to compare and reproduce scientific findings. One way to overcome this problem is to make the most used platforms by coral reef scientists interoperable so that the analyses can all be exported into a common format. This paper introduces functionality to promote interoperability between three popular open-source software tools dedicated to the digital study of coral reefs: TagLab, CoralNet, and Viscore. As users of each platform may have different analysis pipelines, we discuss several workflows for managing and processing point and area annotations, improving collaboration among these tools. Our work sets the foundation for a more seamless ecosystem that maintains the established investigation procedures of various laboratories but allows for easier result sharing.

### 1. Introduction

Automatic recognition tasks for marine species can be quite diverse, ranging from patch-based point classification to semantic segmentation of entire individuals. In some cases, these studies are conducted on single images; in others, involving information on habitats, on orthoimages coming from photogrammetric acquisitions. Often, the organisms or substrates of interest do not have clearly defined boundaries that are suitable for segmentation (such as encrusting algae that cover large regions), and it is more meaningful to quantify their coverage via random point sampling. On the other hand, pixel-wise semantic segmentation is necessary for other taxa, such as corals, to quantify area coverage changes and infer the associated growth rate. Popular AI-assisted software solutions, like CoralNet (Beijbom et al., 2015, Chen et al., 2021), ReefCloud, TagLab (Pavoni et al., 2022), VIAME (Kitware, 2017), BIIGLE (Langenkämper et al., 2017) can handle point annotations, or area annotations, or both, and have web-based or local workflows. However, common workflows for sharing data and results among these platforms are still missing. In this work, we propose to create an ecosystem of tools for managing both area and point annotations, making three popular platforms in this application fields, Coralnet, TagLab, and Viscore (Petrovic et al., 2014), interoperable.

Viscore is a custom visualization and analytical software platform specifically designed for handling large 3D point-based datasets (e.g., dense point clouds) (Fox et al., 2019, Edwards et al., 2023). Viscore features a suite of analytical tools, including iView, which is used to generate orthoimages as orthoprojections from reconstructed point clouds to provide access to the

component images from which point clouds are built (Fox et al., 2019). This exploitation of raw images from multiple views in Viscore guarantees to disambiguate the most complicated classification cases (Figure 2). However, most laboratories do not use orthoprojections of point clouds but orthoimages generated by photogrammetric reconstruction software.

CoralNet is one of the first and most widely used web-based platforms for automatic point annotations of coral reef imagery. CoralNet users have contributed over 3500 source repositories, comprising 3.32 million images and 123 million individual point annotations. CoralNet utilizes an EfficientNet-B0 backbone, pre-trained using patches extracted from approximately 16 million points labeled on over 1200 classes of marine organisms (Chen et al., 2021). After users upload a new source of annotated images, the EfficientNet-B0 is used as a feature extractor, and CoralNet trains a multi-layer perceptron (MLP) classifier to recognize classes from the user's source. CoralNet efficiently performs the patch-based classification task but does not support semantic segmentation models. Further, it does not handle large images, and the sampling strategy is limited to a single area per image.

TagLab is an AI-assisted labeling tool for fast segmenting and analyzing geo-referenced data derived from photogrammetry, such as orthoimages and digital elevation models. This advanced interactive system was designed to support large-scale, repetitive underwater ecological investigations on coral reefs, allowing scientists to process larger amounts of imagery and ultimately facilitate a better understanding of habitat changes on coral reefs. TagLab follows a human-centric AI approach integrating different automation criteria: it fastens the interactive




Figure 1. The TagLab interface has been equipped with several features for managing point annotations: the label panel, layers, and data tables keep a record of both points and regions annotation data on the right-hand side. The *Sampling Settings* panel, which is open above the main view, has two sections. The left section of this panel contains the *Sampling Area* positioning strategy options - manual, single, or random inside a Working Area, or multiple along a transect - along with their corresponding parameters. The right section provides all the sampling information for each Sampling Area, such as its dimensions and the number of sampled points.

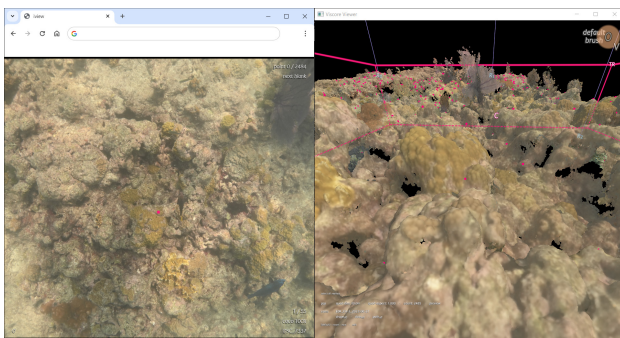


Figure 2. Viscore's iView feature. Component images used in the SfM reconstruction process are shown to users (Left) with points randomly sampled on the dense point cloud (Right) superimposed. Users are able to cycle through all images containing specific points, providing additional vantages of the region of interest and aiding in taxonomic classification.

annotations by proposing pixel-wise AI-assisted region labeling tools; it has a built-in learning pipeline that allows the user to train their own custom semantic segmentation models, and it allows the quick navigation and editing of the geo-referenced predictions by way of semi-automatic tools. This human-in-the-loop strategy enables users to attain an accuracy not achievable through standard machine-learning methods alone, considerably speeding up scientists' annotation work. Moreover, TagLab automates the extraction of demographic information and supports the comparison of multi-temporal surveys, offering tools for automatic/supervised tracking of colonies and calculation of growth and shrinkage, as well as mortality and recruitment rates. With regards to interoperability, TagLab is already interoperable with common GIS platforms; it supports importing and exporting orthoimages, DEMs, and shapefiles

while preserving geo-referencing.

In an effort to make Viscore, CoralNet, and TagLab interoperable, this study worked to add features to the latter, solely due to its open-source nature. TagLab primarily focuses on area-based annotations. Therefore, the first objective was to add support for point annotations and the ability to exchange point annotations between Viscore and CoralNet. Furthermore, TagLab, unlike CoralNet, works on orthoimages rather than on sets of images; one of the most interesting aspects of this work is the integration of different sampling strategies at the habitat scale, which simulates the work that is usually carried out *in situ* in ecological investigations. This allows users an improved flexibility in their workflow and enables integrated analysis of regions and points. Finally, automatic point annotation is achieved with the CoralNet *Deploy API* as a way of the CoralNet-Toolbox, which allows users to invoke CoralNet classifiers. In summary, the main contributions of this work are:

1. We provide an ecosystem of tools for managing area and point annotations.
2. We describe the benefits of using both areas and point annotation in the coral reef restoration activity.

Section 2 details the modifications made to TagLab for this study. In particular, 3 describes the various workflows that can be adopted with the new software ecosystem. Section 4 reports the advantages introduced by this ecosystem solution from a case study using data collected via the Mission Iconic Reef project, a large-scale coral restoration effort based in the Florida Keys National Marine Sanctuary (FKNMS) and led by the National Oceanographic and Atmospheric Administration (NOAA).




Figure 3. A *Sampling Area* of  $5 \times 5$  m sampled following different point sampling criteria: uniform grid, random, stratified.

## 2. Methodologies

This section introduces the functionalities of TagLab developed for the purposes of this study. Generally, annotation points can be manually placed or sampled following the same labeling, visualization, and navigation paradigm of annotation regions. The *Layers* panel displays all the plots within the project; each Plot has its own *Annotations* layer, which contains annotations for Regions and Points in two separate sublayers. Each point annotation has a unique *id*, and its properties are visible from the *Info and Attributes* panel. The *Data Table* panel, synchronized (and navigable) with the *Main View*, shows an overview of the annotations. The class can be assigned from the *Labels* panel, which interactively counts the points labeled per class. The annotation process in TagLab is made easier through the synchronized navigation of both the main view and data tables, equipped with a sort-and-search mechanism. When users click on a row, they can zoom in for a closer look at a specific point, and vice versa; clicking on any point will highlight its associated information in the table, allowing for easy access to all relevant data. Finally, points automatically classified in CoralNet and imported into TagLab preserve the confidence information for the most probable classes, which are displayed in the *AI-suggestions* panel (see Figure 1).

### 2.1 Sampling Strategies

Taglab's sampling strategies are designed to improve the flexibility of analysis on large-area imaging. Points are sampled in *Sampling Areas*; similarly to Viscore and CoralNet. The sampling might follow different methods: uniformly, randomly, or stratified. We refer to Random Sampling, when the points are sampled randomly in a Sampling Area, to Uniform Grid Sampling, when they are assigned to the centers of a regular grid with a constant step defined according to the number of points required in the given area, and to Stratified Random Sampling when a random point is sampled for each cell of the same grid (see Figure 3). Sampling Areas can be manually placed or sampled inside a TagLab's Working Area. The Working area is primarily used when only a part of the orthoimage is being annotated, as the statistics and the import/export functionalities are restricted to it. Additionally, multiple Sampling Areas can be defined along a transect: these can be equally spaced along it or positioned randomly without overlap. The transect is assumed to follow a straight line on the orthoimage; the user needs to specify its start and endpoints (see Figure 1 and Figure 4).

### 2.2 CoralNet-Toolbox

The CoralNet-Toolbox is an open-source coding library specifically tailored to support scientists and restoration practitioners working on coral reefs, particularly in the context of image analysis for studying changes in community composition over time (Pierce et al., 2023). Initially developed for NOAA scientists, the toolbox has been made publicly available to assist the broader community of researchers and practitioners.

The toolbox contains a collection of Python scripts that extend the capabilities of image analysis for coral reefs, and can be accessed through a user-friendly interface or command-line. These scripts primarily focus on point-based annotation techniques that are commonly used in methodologies such as Coral Point Count (CPCe) (Kohler and Gill, 2006) and CoralNet (Beijbom et al., 2015, Chen et al., 2021), but also extends to semantic segmentation for images and 3D models. A key feature of the CoralNet-Toolbox is its ability to programmatically interact with the CoralNet platform using the browser automation library Selenium. This integration enables users to 1) upload and download data to and from CoralNet and 2), provides the ability to interact with the CoralNet *Deploy API* allowing users to obtain predictions for selected points using their own source's model, or any other source's model that is publicly available. The toolbox also contains scripts for pre-training and fine-tuning models locally for tasks including patch-based image classification and semantic segmentation in Pytorch (Paszke et al., 2019). By having the ability to train models locally in addition to those trained on the CoralNet platform, researchers are empowered to develop custom models tailored to their specific needs and requirements, while retaining ownership of the model. Furthermore, the toolbox uses the Metashape (Agisoft LLC, St. Petersburg, Russia) Python API to automate the typical SfM workflow enabling researchers to create high-resolution 3D models and orthomosaics of coral reefs. Combined with segmentation masks, researchers can create classified 3D models using the multi-view semantic segmentation technique outlined in (Pierce et al., 2021).

In this work, we add the tools related to interacting with the CoralNet platform to TagLab as functionalities to enable data exchange, and to facilitate automatic point-based classification in future studies. Further details are given in the next sections (3.2 and 3.3).




Figure 4. *Sampling Area* placing examples. Four Sampling Areas are randomly sampled inside a predefined Working Area (Left) or placed equally spaced along a transect (Right).

### 3. Cross-platform Workflows

This section describes the different working pipelines that users can adopt by combining the three software: sampling points in Viscore and importing derived data into TagLab, sampling points in TagLab and creating a custom classifier in CoralNet, sampling points in TagLab and utilizing an existing classification model from a CoralNet source (or directly importing classified points from platforms such as CoralNet and CPCe into TagLab).

#### 3.1 Exploiting high-res raw images in Viscore to disambiguate the most complicated classification cases

Viscore leverages the output of the typical SfM workflow, which includes the component images, estimated camera positions and their transformation matrices, and the resulting dense point cloud. In Viscore, users are able to generate randomized point samples from the dense point cloud within specified regions. These sampled points are then manually annotated by users to estimate the relative abundance of the benthic organisms of interest using the iView feature. After initially sampling points from the dense point cloud, iView uses the camera transformation matrices to present users a view displaying each point re-projected onto every image that it appears in. Users can then cycle through each image the unoccluded point appears in, obtaining views from different vantages, assisting in making the correct taxonomic designation that would otherwise be too difficult when using only the dense point cloud as reference (see Figure 2).

Viscore supports the ability to create and export rasters, including orthoimages and digital elevation models (DEM) of the dense point cloud, as well as the annotated points relative to the rasters or to each of the component images (i.e., CoralNet format). After exporting, the point annotations and rasters can be imported into TagLab to be visualized and analyzed in synergy with the region-based annotations.

#### 3.2 Creating CoralNet point-based classifiers from TagLab

Each CoralNet source (i.e., project) contains images, annotations, a label-set, and, with enough labeled data, a trained classifier. The label-set is a user-made ontology specifying the class categories that can be found within the project's images; this label-set is also used to define the class categories the source's model is trained on. In CoralNet, labels are initially defined globally and made public, allowing users to share label definitions within each others source's label-sets.

The CoralNet Deep Learning Engine, exploitable by the users, provides an initial backbone trained on 16 million labeled patches of benthic-related categories. Users can leverage this backbone as a feature extractor, training a MLP classifier to recognize classes from their source using image patches ( $224 \times 224$  pixels) cropped on the center of each point annotation. If users choose to make their source public, their patch-based image classifiers, which are often trained on an extensive number of coral taxonomic groups and other common benthic class categories, may be accessed by others via the CoralNet *Deploy API*. These two features make CoralNet an invaluable resource regarding point-based image classification in coral reef monitoring.

Users who work with large ortho-mosaics but need the ability to sample or flexibly annotate their dataset with point annotations can now do so entirely in TagLab. Point annotations can be imported, made, sampled, labeled, and also exported as a file in CSV format compatible with CoralNet. Imagery data easily can be transferred from TagLab to CoralNet by exporting the defined areas on the orthoimage as cropped tiles in JPG format with dimensions up to  $8000 \times 8000$  pixels, compatible with the maximum dimensions supported by CoralNet. The metadata necessary to reassemble these tiled images is stored as the filename following a standard naming protocol.

Additionally, to train a custom CoralNet model with data annotated in TagLab, users must first create a source on CoralNet

and define a label-set for the project. Afterward, within TagLab, users can use the CoralNet-Toolbox feature to automatically upload labeled point annotations and tiled orthoimages to the CoralNet source. The toolbox uses Selenium, an open-source library for browser automation, to automate the process of logging the user in, driving to the applicable URLs, and uploading data to the specified source through the CoralNet platform. Once a source contains at least 20 labeled images, within a 24-hour period CoralNet will initialize a model training session, providing an ID for the trained model that can be accessed through the CoralNet *Deploy API*.

### 3.3 Classifying ortho-images in TagLab exploiting existing CoralNet sources




Figure 5. CoralNet Toolbox GUI. Available classifiers, searchable by sub-string (Left), login and source set up (Center) label-set mapping with import and export options (Right).

Currently, the only way to take advantage of the public models within the CoralNet model catalog is to access them through the CoralNet *Deploy API*. In TagLab, the ability to access these models to perform automated point classification has been made available through the CoralNet-Toolbox.

Typically, to use the CoralNet *Deploy API*, users must provide their CoralNet API token, URLs for images stored online, point coordinates represented in JSON format, and the ID of the source containing the desired model to perform inferencing. However, two major pain points of the API are 1), it can only be accessed through REST commands and 2), the images must be self-hosted within a cloud-based storage account (e.g., Drop-Box, Google Drive, Amazon AWS) and be made publicly accessible via URL; these URLs are passed to the CoralNet *Deploy API*, which are then downloaded and used in the inferencing pipeline. Unfortunately, these requirements are difficult to satisfy by many of those who wish to access the CoralNet's extensive model catalog due to either lack of experience with REST syntax, or the inability to host their own data through a cloud-based storage provider.

In an effort to make accessing different models more accessible, the CoralNet-Toolbox abstracts the process substantially. Within TagLab, a graphical user interface allows users to input the necessary parameters, which are then utilized to construct and submit REST commands on their behalf. For image storage, the CoralNet-Toolbox removes the need to self-host by leveraging CoralNet's existing cloud-based storage system. After uploading images to a source for training or inference, the toolbox automatically identifies and records the images' publicly accessible URLs by crawling the source's pages. The toolbox then constructs and submits REST commands and data payloads containing image URLs and point coordinate through the

CoralNet *Deploy API*, periodically querying CoralNet on progress competition. CoralNet currently allows for a user to have up to five active jobs at any given time, where one job represents up to 100 images, each containing up to 200 point annotations; the toolbox manages job submissions through a queue system to maximize throughput at any given time. Figure 6 shows the inference in action.

Once CoralNet has completed its inference pipeline on all queued jobs, the predictions are downloaded as a file in CSV format and imported into TagLab. To help facilitate the use of all sources, when initially selecting a source's model, TagLab incorporates a label-set mapping tool that allows the user to choose a mapping between the model's output class categories to those that most closely align with the defined categories currently present in their TagLab project (see Figure 5). In summary, predicted labels are made in CoralNet, imported, then mapped and displayed to the user in TagLab, enabling seamless integration of the external model's predictions into the user's existing project workflow.

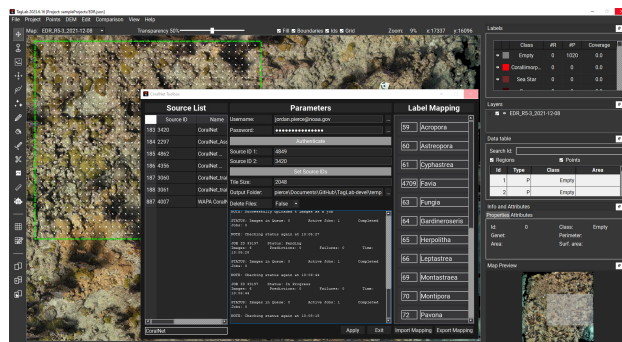


Figure 6. CoralNet Toolbox GUI. Inference in action after the upload operation in TagLab (see the center part of the GUI).

## 4. Mission: Iconic Reefs Coral Restoration Use Case

As climate change impacts are predicted to become more frequent and severe, coral restoration interventions are becoming increasingly necessary to maintain and restore ecosystem functions in many locations. When conducting coral restoration projects, coral stocks are often maintained inland, and in field-based coral nurseries and then outplanted onto the reef; these are activities which require substantial investments of human-driven effort. After outplanting, restoration plots must also then be monitored to assess their effectiveness and track progress towards project goals.

Given that effective restoration includes anywhere from hundreds to tens-of-thousands of corals outplanted into plots spanning hundreds of square meters to several acres, monitoring requirements can themselves require significant investments of effort. Monitoring has typically been conducted by divers *in-situ*, limiting the spatial scale over which information can be collected. More recently large-area imaging is increasingly being used to quickly collect imagery data in the field, which are then used to generate accurate 3D models and orthomosaics within the laboratory, essentially transferring the burden of monitoring and data extraction *ex-situ*.

Although this transition from diver-based surveys in the field to annotating orthomosaics in the lab by way of software helps to tackle the problem of scale, the types of annotations that are

made for the underlying imagery plays a key role in determining what information can be obtained for analysis. Best practice monitoring efforts primarily focus on continuous tracking of the growth of outplanted (and wild) corals, as well as tracking changes in the overall relative abundance of corals and other benthic organisms (e.g., turf and macroalgae) within restoration and control plots (Goergen et al., 2020). As a result, the two most common types of annotations used to fulfill these analytical needs are areas (i.e., polygons) and points.

Area annotations are useful for tracking “things”, such as the growth and survivorship of individual corals with well defined boundaries; however, many other classes consisting of “stuff” that do not have discrete boundaries (including turf and macroalgae) cannot be easily tracked with area annotations. In these cases, the conventional technique has been through randomized point sampling (typically stratified random), which is the most straightforward way to track changes in relative abundance (i.e., percent coverage) through time. With regards to annotation times, point annotations typically take approximately 10s per point for experts, with sampling densities of 25 points per square meter or higher needed to generate accurate data. Moreover, the time required to create area annotations varies significantly depending on the complexity and density of the coral visible in the imagery, but typically ranges from one to ten minutes per square meter. Thus, as the spatial extent of restoration increases, manual efforts to annotate data represents a fundamental bottleneck that can only be overcome with the use of some automated or AI-assisted data labeling tools. In this regard, fortunately in the context of coral restoration, the target corals of interest often belong to a limited taxonomic group; this results in highly replicated training data being available. In contrast, other benthic classes, such as algae and substrate, are typically lumped together. The benefit being that obtaining a machine learning model that meets accuracy requirements for coral identification is sometimes easier to attain than in some other applications.

In this study, we present a use case from the Mission: Iconic Reefs coral restoration project in the Florida Keys (Edwards et al., 2023). To date, baseline (pre-outplanting) imagery has been collected within 105 monitoring plots. Plots measure 10×10m, and within each plot approximately 2000 images are collected 1.5-2m above the benthos using a full frame camera equipped with a 24mm lens. Images are then processed into 3D models (dense point clouds) using Metashape Pro (Agisoft LLC, St. Petersburg, Russia). Each 3D model is then imported into Viscore to generate orthoimages, which are then imported into TagLab, where all coral colonies greater than 3cm of diameter are segmented using TagLab’s semi-automated tools (Pavoni et al., 2022, Cook et al., 2023). Random point samples are also generated in Viscore and labeled to the finest taxonomic resolution using Viscore’s iView (Fox et al., 2019). Point samples are then exported and converted into CoralNet format within TagLab, where they can be overlaid and used as visual guides while segmenting orthoimages with area annotations.

All sites have been analyzed for percent cover using random point sampling, and 75 of 105 sites have been segmented to generate demographic information for wild corals, which will be used as a reference to compare to restored coral demographic information once that becomes available. Annotations are currently being used as training data to learn multiple machine-learning models, with the goal of expediting the annotation of future imagery. Point annotations are being used to train a patch-based image classification model by utilizing the

CoralNet-Toolbox to export and upload data to CoralNet. Post training, the learned model will be accessed through the CoralNet *Deploy API* to make suggestions for each point sampled on the orthoimage in TagLab. Similarly, area annotations are being used to train a custom semantic segmentation model within TagLab, which will be used to assist in segmenting future orthoimages. Both models will help to relieve human annotators of identifying and segmenting more common classes.




Figure 7. Mission: Iconic Reefs orthoimage. Point annotations in pink (Left) and area annotations in multiple colors (Right) overlaid on an orthoimage.

## 5. Conclusions and Future Works

This paper proposes an ecosystem to annotate points and regions on photogrammetry-derived data, supporting interoperability between Viscore, CoralNet, and TagLab.

Future works will directly incorporate additional machine learning and computer vision pipelines into TagLab. Due to its niche in coral reef ecology, except for CoralNet, very few annotation platforms offer patch-based image classifiers for images, let alone large area imagery-derived data products such as orthoimages. By providing the functionality to do so, users could train their own classifiers to use them in or outside of TagLab as an AI-assisted tool to enhance the accuracy and efficiency of data labeling. This enhancement could empower users to customize machine learning models tailored to their specific research needs, helping to facilitate more precise and comprehensive analysis.

Another future direction of work involves integrating key Viscore features into TagLab by way of the Viscore API; key features including iView allows for more detailed user validation of benthic classes when creating area annotations. While orthoimages are beneficial for ecologically relevant sampling, their classification capability is limited due to lower ground sampling resolution. To overcome this limitation, camera pose information and Viscore’s iView feature could be utilized to access raw imagery for any spatial position in the dense point cloud or orthoimage. The iView functionality also offers the ability to perform multi-view classification, which, in preliminary tests, has shown promising results in significantly improving patch-based image classification accuracy for many cases.

In the long term, this study lays the foundation for designing an infrastructure dedicated to the digital research of coral reefs, including a repository for sharing datasets, resources, and machine learning pipelines. Introducing such infrastructure would greatly promote adopting common standards and analysis tools, ensuring image-based data interoperability, scalability, and quality assessment (all of which are core pillars of the

FAIR principles of open science), fostering greater collaboration and advancement within the scientific community. However, achieving this principle necessitates overcoming obstacles related to data homogeneity and navigating data governance issues and access rights across diverse cultural and national contexts.

**Acknowledgments** This work was supported in part by the Italian Ministry of Foreign Affairs and International Cooperation, grant number US23GR14, "ReefSurvAI: Verso un'infrastruttura web che supporti l'uso dell'intelligenza artificiale per il monitoraggio delle barriere coralline".

## References

- Beijbom, O., Edmunds, P. J., Roelfsema, C., Smith, J., Kline, D. I., Neal, B. P., Dunlap, M. J., Moriarty, V., Fan, T.-Y., Tan, C.-J., Chan, S., Treibitz, T., Gamst, A., Mitchell, B. G., Kriegman, D., 2015. Towards Automated Annotation of Benthic Survey Images: Variability of Human Experts and Operational Modes of Automation. *PLOS ONE*, 10(7), 1-22.
- Chen, Q., Beijbom, O., Chan, S., Bouwmeester, J., Kriegman, D., 2021. A new deep learning engine for coralnet. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 3693–3702.
- Cook, S., Rojano, S., Edwards, C. B., Bollinger, M. A., Pierce, J., Viehman, T. S., 2023. Standard operating procedures for the use of large-area imaging in tropical shallow water coral reef monitoring, research and restoration: Applications for mission: Iconic reefs restoration in the florida keys national marine sanctuary. Technical report, NOAA Technical Memorandum NOS NCCOS 320.
- Edwards, C. B., Viehman, T. S., Battista, T., Bollinger, M., Charendoff, J., Cook, S., Combs, I., Couch, C., Ferrari, R., Figueira, W., Gleason, A. C. R., Gordon, S., Greene, W., Kuester, F., McCarthy, O., Oliver, T., Pedersen, N. E., Petrovic, V., Rojano, S., H., R., Sandin, S. A., Zgliczynski, B., 2023. Large-area imaging in tropical shallow water coral reef monitoring, research and restoration: A practical guide to survey planning, execution, and data extraction. Technical report, NOAA Technical Memorandum NOS NCCOS 313.
- Fox, M. D., Carter, A. L., Edwards, C. B., Takeshita, Y., Johnson, M. D., Petrovic, V., Amir, C. G., Sala, E., Sandin, S. A., Smith, J. E., 2019. Limited coral mortality following acute thermal stress and widespread bleaching on Palmyra Atoll, central Pacific. *Coral Reefs*, 38(4), 701–712. <https://doi.org/10.1007/s00338-019-01796-7>.
- Goergen, E., Schopmeyer, S., Moulding, A., Moura, A., Kramer, P., Viehman, S., 2020. Coral Reef Restoration Monitoring Guide: Methods to evaluate restoration success from local to ecosystem scales.
- Kitware, 2017. VIAME toolkit. <https://www.viametoolkit.org>. [Online; accessed 8-February-2023].
- Kohler, K. E., Gill, S. M., 2006. Coral Point Count with excel extensions (CPCE): A visual basic program for the determination of coral and substrate coverage using random point count methodology. *Computers amp; Geosciences*, 32(9), 1259–1269.
- Langenkämper, D., Zurowietz, M., Schoening, T., Nattkemper, T. W., 2017. Biigle 2.0-browsing and annotating large marine image collections. *Frontiers in Marine Science*, 4, 83.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E. Z., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *CoRR*, abs/1912.01703. <http://arxiv.org/abs/1912.01703>.
- Pavoni, G., Corsini, M., Ponchio, F., Muntoni, A., Edwards, C. B., Pedersen, N., Sandin, S., Cignoni, P., 2022. TagLab: AI-assisted annotation for the fast and accurate semantic segmentation of coral reef orthoimages. *Journal of Field Robotics*, 39(3), 246-262. <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.22049>.
- Petrovic, V., Vanoni, D., Richter, A., Levy, T., Kuester, F., 2014. Visualizing high resolution three-dimensional and two-dimensional data of cultural heritage sites. *Mediterranean Archaeology and Archaeometry*, 14, 93-100.
- Pierce, J., Butler, M. J., Rzhanov, Y., Lowell, K., Dijkstra, J. A., 2021. Classifying 3-D Models of Coral Reefs Using Structure-From-Motion and Multi-View Semantic Segmentation. *Frontiers in Marine Science*, 8. <https://www.frontiersin.org/articles/10.3389/fmars.2021.706674>.
- Pierce, J., Edwards, C., Vieham, S., Rojano, S., Cook, S., Costa, B., Sweeney, E., Battista, T., 2023. Coralnet-toolbox. <https://github.com/Jordan-Pierce/CoralNet-Toolbox>. GitHub repository.