https://doi.org/10.25923/5n6d-kx34

Large-area imaging in tropical shallow water coral reef monitoring, research, and restoration:

A practical guide to survey planning, execution, and data extraction

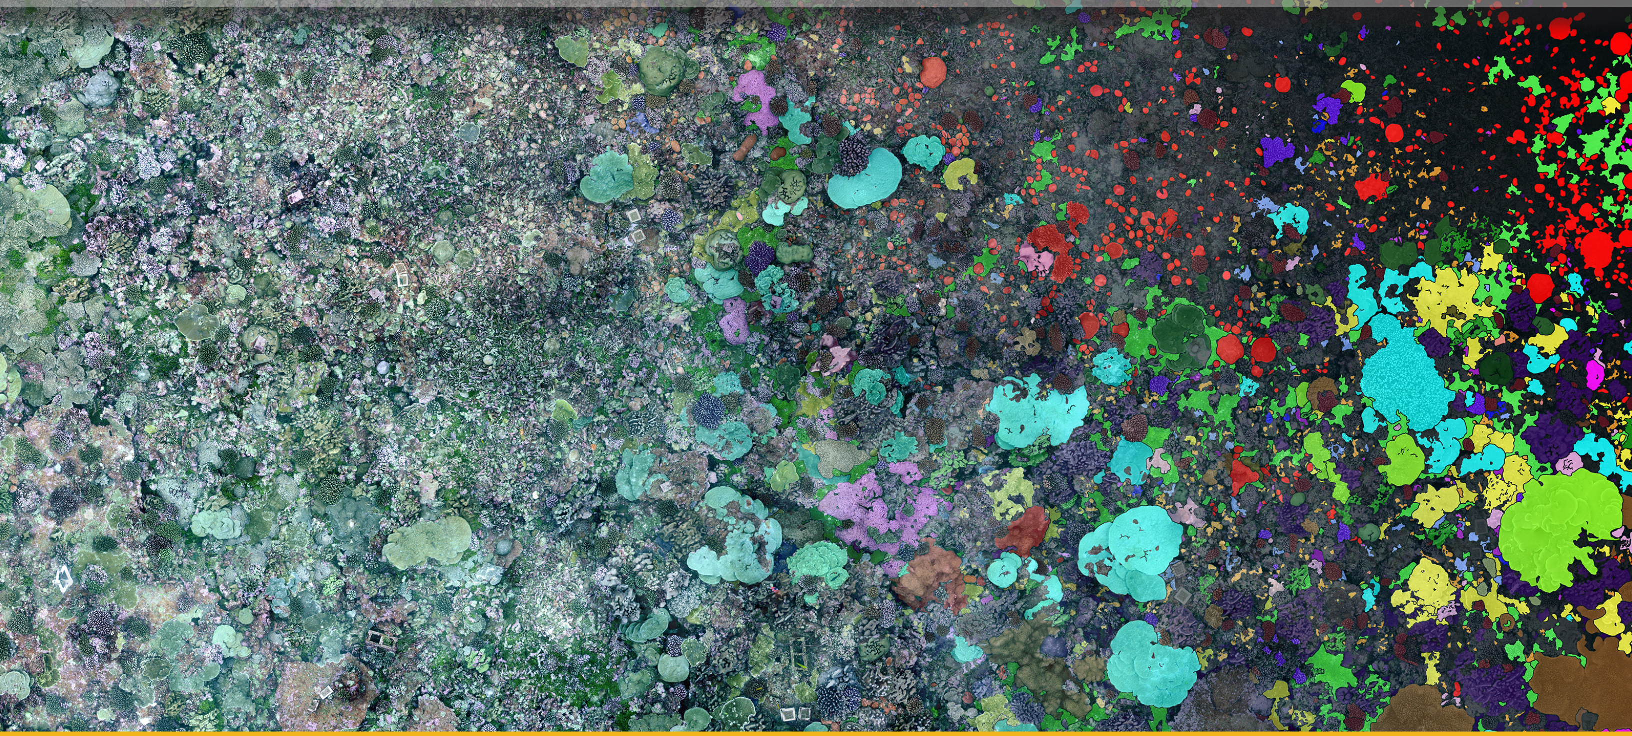

September 2023

NOAA TECHNICAL MEMORANDUM NOS NCCOS 313

NOAA NOS National Centers for Coastal Ocean Science

SUGGESTED CITATION

Edwards, C. B., Viehman, T. S., Battista, T., Bollinger, M. A., 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., Runyan, H., Sandin, S. A., and Zgliczynski, B. J. (2023). Large-area imaging in tropical shallow water coral reef monitoring, research, and restoration: A practical guide to survey planning, execution, and data extraction. NOAA National Ocean Service, National Centers for Coastal Ocean Science. NOAA Technical Memorandum NOS NCCOS 313. https://doi.org/10.25923/5n6d-kx34

ACKNOWLEDGMENTS

The authors thank all those who have contributed to the implementation and use of LAI over the past decade and the many hours of discussion and debate that have fueled the development of this approach. In particular, sincere appreciation is expressed to the countless people who have contributed to the testing of the methods introduced here, whether collecting imagery in the field, creating LAI products, spending hours extracting data in the lab, or working tirelessly to ensure the integrity of these data. This report was funded by NOAA's Coral Reef Conservation Program project 31372 to NOAA's National Centers for Coastal Ocean Science.

PHOTOGRAPHY AND FIGURES

DISCLAIMER

Large-area imaging in tropical shallow water coral reef monitoring, research and restoration:

A practical guide to survey planning, execution, and data extraction

Prepared by

NOAA National Ocean Service
National Centers for Coastal Ocean Science
Marine Spatial Ecology Division
Silver Spring, MD

Authors

Clinton B. Edwards[1,2], T. Shay Viehman[3] , Tim Battista[3], Michael A. Bollinger[1], Jonathan Charendoff[4,5], Sophie Cook[1], Ian Combs[6], Courtney Couch[4,5], Renata Ferrari[7], Will Figueira[8], Arthur C.R. Gleason[9], Sophie Gordon[7], Will Greene[10], Falko Kuester[11], Orion McCarthy[2], Thomas Oliver[5], Nicole E. Pedersen[2], Vid Petrovic[11], Sarah Rojano[1], Hugh Runyan[12], Stuart A. Sandin[2], and Brian J. Zgliczynski[2]

[1] CSS Inc., Under contract to NOAA, National Ocean Service, National Centers for Coastal Ocean Science
[2] Scripps Institution of Oceanography, University of California San Diego
[3] NOAA, National Ocean Service, National Centers for Coastal Ocean Science
[4] Cooperative Institute for Marine and Atmospheric Research, University of Hawai‘i at Mānoa
[5] NOAA, National Marine Fisheries Service, Pacific Islands Fisheries Science Center
[6] Elizabeth Moore International Center for Coral Reef Research & Restoration, Mote Marine Laboratory
[7] Australian Institute of Marine Science
[8] School of Life and Environmental Sciences, The University of Sydney
[9] Department of Physics, University of Miami
[10] Perry Institute for Marine Science
[11] Computer Science and Engineering, University of California San Diego
[12] Redwood City, CA, USA

September 2023

NOAA TECHNICAL MEMORANDUM NOS NCCOS 313

United States Department of Commerce
Gina M. Raimondo
Secretary

National Oceanic and Atmospheric Administration
Richard W. Spinrad
Under Secretary of Commerce for Oceans and Atmosphere and NOAA Administrator

National Ocean Service
Nicole LeBoeuf
Assistant Administrator

Table of Contents

List of Figures and Tables

Acronyms

AI	artificial intelligence
AIMS	Australian Institute of Marine Science
AUV	autonomous underwater vehicle
AVD	Azure Virtual Desktop
CMOS	complementary metal-oxide semiconductor
DEM	digital elevation model
DOF	depth of field
DPC	dense point cloud
DPV	diver propulsion vehicle
DSLR	digital single-lens reflex
EcoRRAP	Ecological Intelligence for Reef Restoration and Adaptation Program
GCP	ground control point
GSD	ground sampling distance
HDD	hard disk drive
HPC	high-performance computing
LAI	large-area imagery
M:IR	Mission: Iconic Reefs
ML	machine learning
MP	megapixel
NAS	network attached storage
NCCOS	National Centers for Coastal Ocean Science
NCEI	National Centers for Environmental Information
NCRMP	National Coral Reef Monitoring Program
NGO	non-governmental organization
NOAA	National Oceanographic and Atmospheric Administration
PIFSC	Pacific Islands Fisheries Science Center
PVC	polyvinyl chloride
QAQC	quality assurance and quality control
RAID	redundant array of independent disks
RAM	random access memory
ROI	region of interest
ROV	remotely operated vehicle
RTK	real-time kinematic
SfM	structure from motion
SIFT	scale invariant feature transform
SOP	standard operating procedure
SPC	sparse point cloud
SSD	solid-state drive
UC	University of California
UM	University of Miami
USBL	ultra-short baseline

Executive Summary

Clinton Edwards, CSS Inc./NOAA NCCOS and SIO, UC San Diego

This is a practical guide to the implementation of large-area imaging (LAI) for coral reef scientists. LAI refers to an approach to generate composite 3D (and derived 2D) image products from sequences of field-collected images using structure-from-motion (SfM) photogrammetry. Coral reef scientists conducting coral restoration, environmental monitoring, or research often require time series or high taxonomic resolution data to calculate metrics such as abundance or density, percent cover, species condition, and reef complexity. These data have a history of in situ collection by divers; however, the time-intensive challenge of traditional fieldwork approaches limits the volume and spatial extent of data collection and is therein a fundamental bottleneck for many restoration, monitoring, and research objectives.

Underwater photography, and more recently the proliferation of high-resolution digital cameras, enables a significant increase in the amount of data that can be collected during a single dive by allowing data collection to continue in the lab with image analysis software. Unfortunately, single photographs provide only limited snapshot views of ecosystems, making extracting information for key metrics such as colony size or population metrics difficult. LAI provides the detail available in single photographs while also enabling a spatially expansive view of underwater landscapes in a single contiguous image. LAI ultimately allows for the generation of spatially accurate, detailed maps of the benthos that can be collected easily and at a sufficient scale to capture thousands of coral colonies and other sessile benthic organisms. Multiple metrics, both ecological (e.g., details of taxonomy and the size and position of individual coral colonies) and physical (e.g., rugosity and structural complexity), can be extracted from the LAI. LAI collections from repeated sampling of permanent locations on the reef over time allow for the quantification of growth, recruitment, shrinkage, and mortality for both individual coral colonies and the population as a whole.

LAI is a flexible, computationally reliant approach that emerged from the disciplines of engineering and computer vision. For most scenarios, the tools and software needed to conduct LAI have matured to the degree that little technical expertise is needed to effectively implement the approach. However, applications of LAI for coral reef science are still relatively new, and few comprehensive instructional materials have been designed specifically for a non-technical audience. Existing documentation is often case-specific, limited in scope, or not widely available. To make the LAI process more accessible to the coral reef community, the first part of this document provides a broad overview to LAI and its implementation. The approach is divided into four steps of the “LAI pipeline”: 1- image collection, 2- model construction, 3- ecological analysis, and 4- data curation. The key technical details of LAI are summarized, focusing on those aspects of the approach where user intervention is required, and a framework is provided to help guide project planning decisions.

This guide is intended for a broad group of users, including first-time adopters as well as experienced professionals. The LAI approach described in Part I of this guide is compatible with most existing coral restoration, environmental monitoring, and research programs. While the information provided is relevant to work conducted across a spectrum of spatial scales and levels of resolution, the focus is on concepts related to the creation of LAI that allow for detailed taxonomic descriptions of benthic organisms at high levels of replication. Decisions made during both image collection and model construction can have huge implications for LAI resolution and the quality of the derived metrics. This report also provides best practices intended to ensure comparability of datasets into the future. Part II of the guide provides a series of standard operating procedures (SOPs) currently used by experienced groups for a wide range of objectives. The SOPs are briefly summarized by their respective authors with commentary designed to provide exposure to the different approaches. These SOPs are also available from the individual authors at their respective institutions, with links provided in Part II.

The computational aspects of the LAI approach represent a rapidly evolving field. While efforts have been taken to present the state of the art, many of the technical details included here, from the algorithms used to the available computer hardware, will undoubtedly evolve dramatically over the next 2–5 years. Readers are encouraged to stay up to date on the technologies presented here and routinely visit the websites where the individual SOPs are hosted, as updates to these documents will be made as they become available.

Part I:
Getting Familiar with Large-Area Imagery

Jessica Levy (Coral Restoration Foundation [CRF] and KAUST Reefscape Restoration Initiative [KRRI])

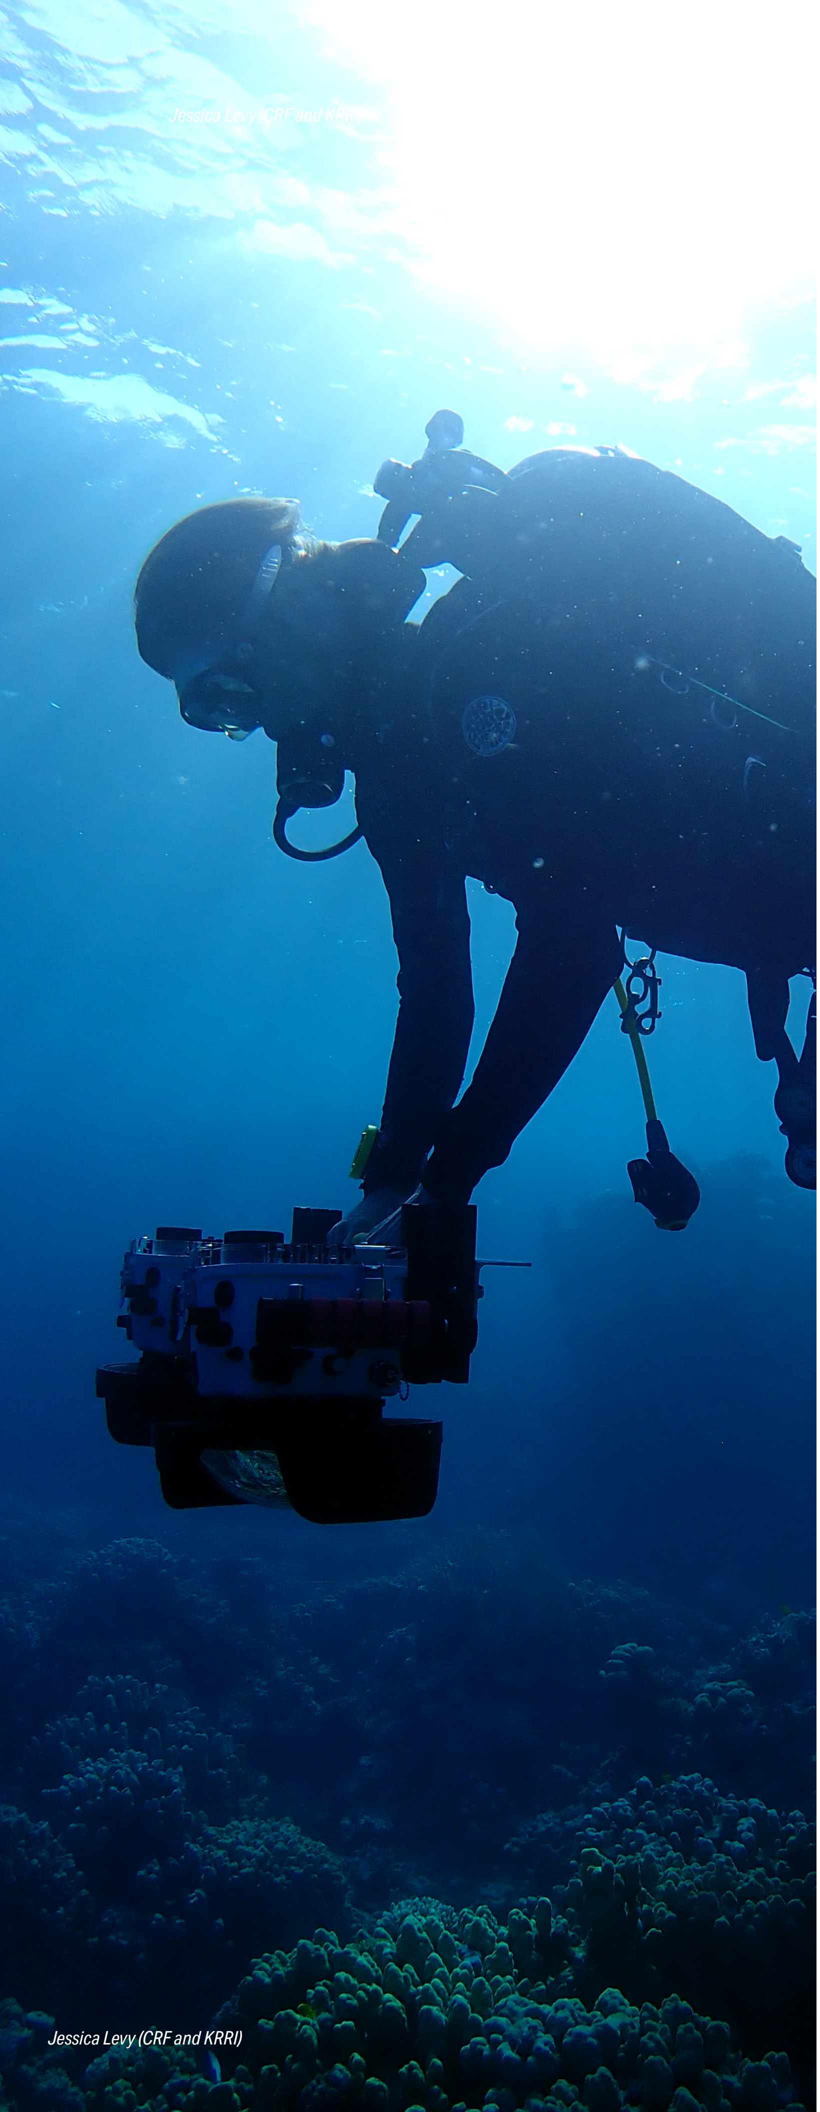

Jessica Levy (CRF and KRRI)

Chapter 1

Introduction to Large-Area Imaging

Successful coastal management efforts rely upon an empirical framework informed by well-designed research and monitoring programs that rigorously evaluate the efficacy of current interventions and adapt future plans (Ladd et al., 2018; Goergen et al., 2020). For coral reefs, monitoring data need to be collected at ecologically relevant spatial scales to develop an understanding of change in reef-building coral populations and communities. Historically, monitoring data have been collected by divers in the field, where the practical constraints of subtidal field work have limited the spatial extent of data collections. Over the last two decades, a photogrammetric approach referred to as large-area imaging (LAI; sensu Edwards et al., 2017) has been developed, which blends thousands of individual images together to create composite views of natural scenes in detail and spatially expansive views of the benthos that are larger than the individual images from which they are created. In the context of coral reef research and monitoring, LAI can be applied at any spatial scale to produce an array of visual products including spatially expansive, 2D map views of coral reef habitats or highly detailed 3D representations of complex objects.

The use of LAI is preceded by the early creation of hand-drawn maps to characterize reef structure by scientists who had only bottom trawls and depth soundings at their disposal (Darwin and Stoddart, 1962). Later researchers complemented these depictions with detailed measurements taken in the field to generate spatially explicit maps of underwater habitats (Kornicker and Boyd, 1962; Loya, 1978) to describe the physical and community structure of coral reefs (Huddell et al., 1974; Weinberg, 1981). Early uses of underwater photography for data extraction include the laborious mosaicking of composite photographs by hand to understand population trends of corals in Jamaica (Figure 1; Porter et al., 1981). These early efforts on coral reefs built upon a long history of photogrammetry, which is the process of taking measurements from any photograph that contains, or in which can be embedded, scale information. The use of imagery for this purpose is nearly as old as photography itself, beginning with early terrestrial and aerial photography for city maps (Figure 2; Jiang et al., 2008). Modern camera sensors, geopositioning technology, and computational advancements now allow for highly precise measurements to be taken from digital imagery and image-based products; however, their utility has been limited in subtidal applications. With the increasing accessibility of digital photography and high-powered computing came the development of computational approaches to automate feature detecting and matching of overlapping photographs. These computer vision innovations allow for the generation of detailed orthorectified (2D planar projection) imagery, or orthophotomosaics, (Gracias and Santos-Victor, 2001; Lirman et al., 2007) and 3D reconstructions of benthic habitats (Pizarro et al., 2009; Westoby et al., 2012) that form the basis of the LAI

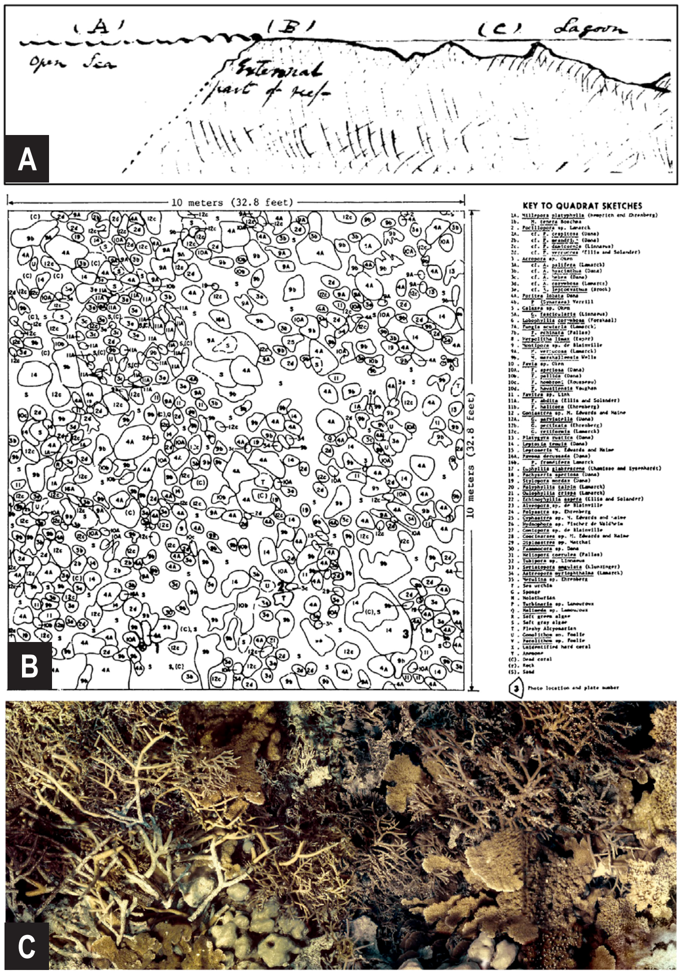

Figure 1. Hand-drawn maps and early large-area imaging (LAI). Sketch of reef cross section (A) from Darwin's 1842 work describing the formation of coral islands. With the advent of scuba, followed painstaking hand-drawn mapping (B) (Huddell, 1974). Later, underwater photographs were stitched together by hand to create photomosaics (C) (Porter, 1981).

Figure 2. Historical photogrammetry image. Sketch of early photogrammetry pioneer Aimé Lausedate conducting a photographic survey of Paris. Source: Granshaw, 2019

approach. Similar techniques have also been developed using other imaging approaches, including satellite and aerial imagery, which can provide spatially expansive views of coral reefs and many other habitats (Chirayath and Earle, 2016; Lyons et al., 2020; Li and Asner, 2023). While these remote sensing products are highly useful, they have low spatial resolution and lack the visual detail needed to describe coral community structure and change for many coral reef habitats; or they fail to deliver metrics on key processes underlying ecosystem trajectory, which are needed to inform restoration of coral reefs (Ferrari et al., 2021).

The composite products made possible through the LAI approach are typically viewed as 2D planar projection imagery (Lirman et al., 2007; Nicosevici et al., 2009) or 3D reconstructions (e.g., 3D models). The resulting digital recreations of the underwater landscape offer a substantial complement to existing in situ methodologies where search times are limited by the constraints of sampling in subtidal habitats. Researchers working on land in silico can exhaustively search large areas of digitally reconstructed benthos in a relatively rapid manner and extract nearly any form of data that do not include a census of mobile organisms or the collection of physical samples. Ultimately this approach allows researchers to generate data for everything from bulk metrics of percent cover and structural complexity to details of coral size distributions (Burns et al., 2015; Figueira et al., 2015; Murfitt et al., 2017; Sandin et al., 2020). Further, temporally replicated imagery allows for detailed tracking of change in any of these metrics (Kodera et al., 2020; Pascoe et al., 2021).

The quantitative metrics that can be extracted from LAI represent a paradigm shift in the ability to describe reef state and function. An equally important benefit of LAI is that it provides a permanent digital snapshot of the habitat as it existed on the day that the imagery was collected. The fundamental data generated by the LAI approach are inherently spatial, and users can virtually explore digitally recreated field sites as 3D models without the restrictions of static 2D views provided by video or other image products. This ability to navigate and interact with 3D models in a virtual reality environment promises the users entirely new ways of observing and studying coral reefs. Although software tools currently used to extract information from LAI remain in active development, a number of critical data streams can already be extracted from LAI and include the ability to visually overlay additional data products and field-collected information. LAI facilitates seascape-level ecological approaches, and as it continues to proliferate within applied sciences, it will undoubtedly diversify to meet the needs of an expanding community.

LAI provides a powerful communication tool well before any data are extracted. The visual nature of LAI products enables researchers to "show-not-tell," allowing the imagery to demonstrate seascape details with better clarity than traditional post hoc syntheses techniques. Visualization of LAI data products may be a necessary first step in a line of investigation but in many cases may be the only information that is needed to accurately portray ecological phenomena. Regardless of the ultimate application, it is first important to have an explicit understanding about what is generated via the LAI approach.

This guide focuses on LAI products that are generated using an approach that leverages a computational approach known as structure-from-motion (SfM) (Snavely et al., 2008; Pizarro et al., 2009; Westoby et al., 2012). For coral reef monitoring, SfM-based LAI provides a combination of visually detailed and geometrically precise reconstructions of the imaged scene, enabling the extraction of volumes of highly accurate data at nearly any desired spatial scale (Jiang et al., 2020; Ferrari et al., 2021). The LAI approach relies heavily on the fields of computer science and computer vision, but the level of expertise needed to implement it is well within the capabilities of ecological scientists.

Structure-from-motion (SfM) is a photogrammetric approach used to reconstruct 3D models of scenes from overlapping imagery. Features matched in overlapping imagery are used to estimate the relative position of images with respect to each other, through a process called camera alignment, or camera pose estimation. These camera position estimates, and subsequently the positions of matched features, are represented by what is referred to as a sparse point cloud (SPC). When raw imagery is of sufficient quality, features can be easily detected by computer software, and if there is sufficient overlap among adjacent images, these features can be reliably matched. SfM can then be conducted using overlapping imagery collected by a single camera alone, without any additional external inputs (e.g., camera positions).

There is currently a relative lack of resources designed to disambiguate the LAI approach and enable greater accessibility within the community (McCarthy, 2023). The techniques described in this report form the basis for the LAI approach but do not represent a comprehensive catalog of the research conducted in this field to date. Many aspects of the underlying technologies and tools are relatively new, remain active areas of innovation in computer science (Lindenberger et al., 2021; Bellavia et al., 2022), and will undoubtedly continue to evolve in the coming years. Despite the continued potential for innovation, many applications of the approach are mature, well vetted, and positioned for current use in ecological monitoring of benthic coral reef habitats. These products have already allowed the coral reef community an unprecedented opportunity to visualize, archive, and ultimately increase understanding of these important ecosystems. A number of groups have amassed substantial bodies of methodological and hypothesis-driven work based on this approach (Lirman et al., 2010; Burns et al., 2015; Figueira et al., 2015; Ferrari et al., 2016; Bryson et al., 2017; Edwards et al., 2017; Lechene et al., 2019; Suka et al., 2019; Sandin et al., 2020; Couch et al., 2021; Fukunaga et al., 2022a). Furthermore, the LAI approach is increasingly being incorporated into large-scale monitoring for conservation, restoration, and research (Goergen et al., 2020; Couch et al., 2021).

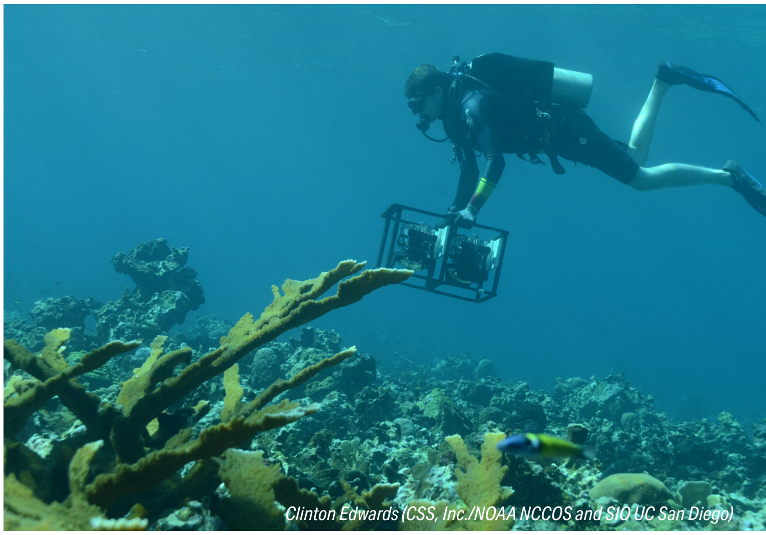

Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO UC San Diego)

Gina Madix (University of the West Indies, Mona and The Nature Conservancy)

Sarah Rojano (CSS, Inc./NOAA NCCOS)

Jessica Levy (CRF and KRRI)

Chapter 2

Overview of the LAI Pipeline

The LAI approach offers the opportunity to create accurate digital representations of natural scenes that allow for extensive ecological data collection of sessile benthic organisms and habitats. This guide focuses on LAI that is generated by the collection of highly overlapping imagery of a natural scene that is then provided to an SfM-based photogrammetric software platform to create a 3D reconstruction of the scene captured in the imagery. The computational approach taken by SfM-based software is to first identify features in raw imagery and then to match these features in overlapping photographs. The matches are then used to estimate the position of the camera and the matched features, allowing for the digital reconstruction of the original scene. Therefore, the fundamental requirement of the LAI approach is the collection of raw imagery of sufficient quality for computer software to easily detect features and with sufficient overlap among adjacent images so that these features can be reliably matched, for the entirety of the survey area.

Feature detection and matching refers to the computational process of identifying features and matching features visible in overlapping imagery. There are a variety of approaches used, with the most common being the scale invariant feature transform (SIFT), which identifies features in imagery that are (largely) independent of the scale, illumination, and orientation of the source image. When these features, often referred to as key points, are visible in overlapping images, they can be used as correspondences, also known as tie points, to match the overlapping images together.

The requirements and constraints of raw image collection are present regardless of the visual detail or spatial scale desired from LAI products. It is important to recognize that the comprehensiveness and level of detail available from the raw imagery and downstream products directly constrain the analyses that can be conducted on the resultant LAI products (see Part I, section 3.2). Similarly, the availability of analytical tools for data extraction should also be considered, as well as the substantial archival responsibilities associated with raw imagery and composite LAI and associated products (see Part I, sections 2.2.4 and 2.2.5). In many cases, requirements at later stages of the creation and use of LAI limit the available options at the initial stages. For instance, when taxonomically specific or geometrically precise information is needed, raw imagery must be highly detailed, and often more of it must be collected (see Part I, section 2.1.2). As the number of raw images used to create composite LAI products increases, the computational demands to both store and create these products also increases, and computer infrastructure must be scaled accordingly. Further, when time series data are collected, there is a clear need for not only additional storage but also a prioritization of data accessibility to facilitate temporal comparisons. Similarly, decisions about ecological analysis should guide model building parameters and data curation planning. In each case, the needs at these later data processing stages should guide decisions about earlier data collection in the field, including the selection of the specific camera used and the spatial extent over which imaging is to occur.

The interdependency of the decisions needed to complete the LAI workflow can present as logistically complicated, particularly for those groups with broad mandates and significant data demands. To facilitate the decision-making process, the LAI workflow is divided into a conceptual pipeline with four main steps: 1– image collection, 2– model construction, 3 – ecological data extraction, and 4 – data curation (Figure 3). The LAI pipeline is designed to guide decisions at each step with respect to logistical capacity at the others. While these steps are not formally sequential, decisions required at each step should ideally be made before any project begins, particularly with respect to the data products that will be needed to reach project objectives.

Part I of this guide provides a brief overview of the major concepts and processes associated with each of the steps in the LAI pipeline. Part II then provides a series of standard operating procedures (SOPs) currently used by a variety of groups pursuing diverse objectives, to allow users exposure to the decisions that have been made by others in the past. The computational aspects of the LAI approach represent a rapidly evolving field. While efforts have been taken to present the current state of the art, many of the technical details included here, from the algorithms used to the available computer hardware, will undoubtedly evolve over the next 2–5 years. Users are encouraged to stay up to date on the technologies and visit the websites where the individual SOPs are hosted to check for updated versions.

The ultimate goal of this guide is to streamline decisions on what imaging and analytical approaches can be applied in different settings and under different constraints. This will help groups accomplish their stated objectives for coral reef monitoring for conservation, management, or restoration.

2.1. Image collection

2.1.1. Camera selection

The fundamental requirement for the creation of LAI is sharp (i.e., in-focus) and highly overlapping imagery of a scene, collected at the level of detail needed to meet visualization and analysis objectives. LAI products will be no more detailed, and generally less so, than the raw images they are created from, so it is important to make decisions regarding image quality at the beginning of any LAI project. In digital cameras, images are captured by a light sensor that is divided into discrete units of area, or pixels. The total number of pixels a sensor is divided into is known as the megapixel (MP) resolution of a given camera. The overall detail of an image is a function of the quality of the camera used, the degree to which the image is in focus, and the spatial footprint represented by each pixel in the image, also known as the spatial resolution or ground sampling distance (GSD). However, the quality of a camera system is dictated not only by the sensor, its size, or the total number of MPs available but by a wide variety of factors including the optics of the lens and underwater housing, details of onboard image processing, and the manner in which the camera is operated by the user.

A high-resolution image is most useful for generating LAI if it is in focus, as the ability of SfM software to identify and match features is reduced as focus quality degrades. Cameras with larger sensors and higher quality optics generally have a higher capability to perform better over a wider range of lighting conditions than smaller form factor cameras (e.g., with cropped-sensor). The ability to collect sharp in-focus imagery is also a function of various camera settings, and depending on the specific scenario, different combinations of camera settings can produce imagery of similar quality. Additionally, the manner in which the camera is operated under different combinations of settings can affect the quality of imagery. For instance, if the camera is moved too quickly and the shutter speed is too low, motion blur can occur, and imagery will be blurry and out of focus. However, while faster shutter speeds can reduce motion blur, they also limit the amount of light that will reach the sensor and it becomes important to compensate with light sensitivity and aperture settings.

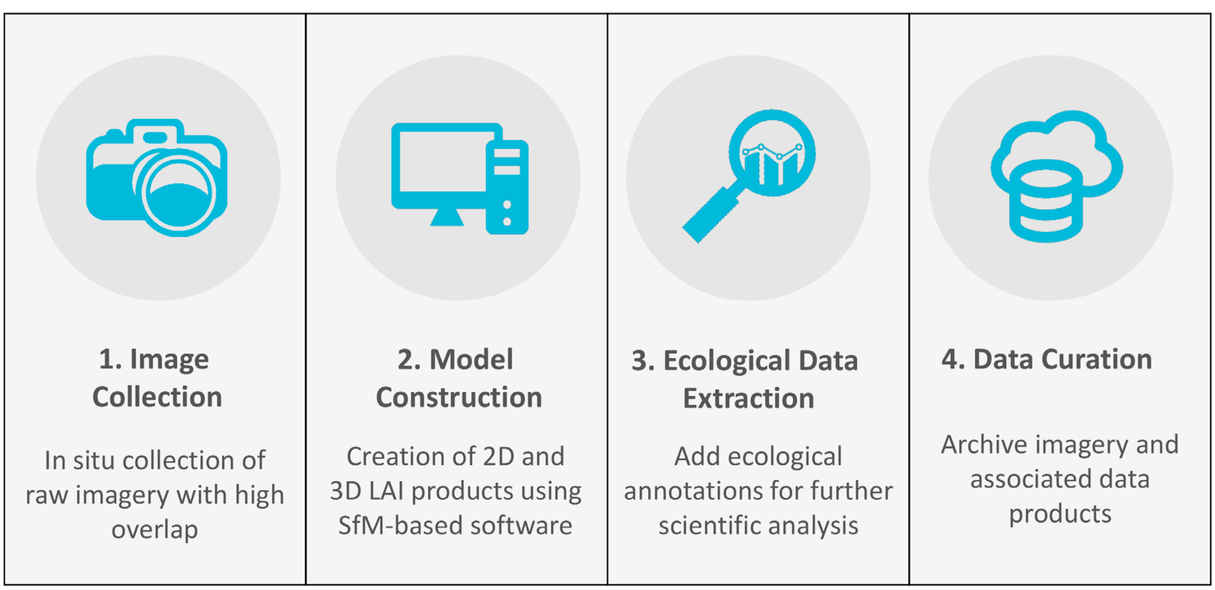

Figure 3. Large-area imaging (LAI) pipeline schematic. The four steps of the pipeline are not strictly sequential with the structure of each steps dependent on the structure of the other three. However, in practice, steps 1–3 occur in sequence, with data curation activities occurring throughout.

Given the dynamic environmental settings of most coral reef ecosystems, the challenge is to select a camera with a combination of available settings and image modes that will work well under a variety of conditions. As lighting conditions can be particularly variable, it is generally recommended to select a camera with a wide range in light sensitivity (e.g., dynamic range) available in a programmable, or similar, mode, so that settings do not need to be changed during the dive. Similarly, the camera should have a wide aperture range, also available in a programmable or other suitable mode. While the opacity of water imposes practical constraints on how far from the benthos highly detailed imagery can be collected, higher-end cameras will also tend to produce higher quality imagery at the same distance from the bottom, regardless of the number of available MPs. Guidance on different camera settings and configurations under different operating conditions can be found in Part II.

2.1.2. Image resolution and spatial extent

Image resolution is a function of the manufacturer specifications of the camera (e.g., number of MPs), the focal length of the camera lens, and the distance from the bottom at which the camera is operated. To capture a higher-resolution image, the camera must be operated closer to the bottom or be outfitted with a longer focal length lens. As more detailed imagery generally has a smaller GSD, more images are required to cover the same area at higher resolution than lower resolution (Figure 4). For example, generating 3D models with the detail needed for precise measurements of corallite volume requires collection of raw imagery with even higher levels of detail than needed in the model (GSD approximately 0.1 mm). Images at this level of resolution will cover very limited spatial extents but still require substantial overlap to successfully build the 3D model. The overlap requirements for SfM are 60%–80% in both side-to-side and front-to-back adjacent images; and for underwater applications where navigation is imprecise, it is strongly recommended to err on the side of caution. For example, imaging a single 20-cm-radius coral colony for sub-corallite–level detail will require several hundred images. The standard operational unit to capture species-level taxonomy at mid-depths (8–12 m), in a single dive, ranges from 30–400 m^2 (Edwards et al., 2017; Hernández-Landa et al., 2020; Couch et al., 2021). Collecting imagery sufficient for corallite-level analysis in plots of this size would require tens of thousands of images. The time needed to both collect such imagery and process it into LAI would be substantial and likely unfeasible if many such plots were required. In contrast, generating LAI to accurately identify taxonomy and colony boundaries within the same 100-m^2 plot would require far fewer (<3,000) and lower-resolution images (GSD approximately 1 mm) that could be collected in under an hour and processed into an adequate data product in a relatively short amount of time.

While high levels of overlap among adjacent images are necessary to generate LAI, maintaining low GSDs has been shown to greatly improve reconstruction quality (Mosbrucker et al., 2017; Marre et al., 2019). Further, cameras with larger sensors, higher quality lenses, and other optical characteristics will have a greater ability to differentiate between pixels and produce detailed and accurate imagery (Figure 5) when operated at the same distance from the bottom. As a result, these larger form factor cameras tend to produce more accurate LAI (Thoeni et al., 2014, Nocerino et al., 2019). However, details of how the camera is operated, including

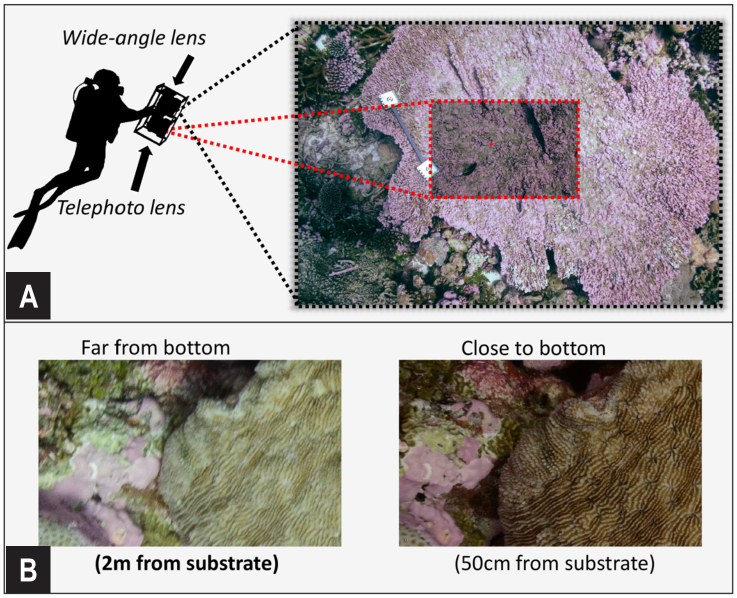

Figure 4. Spatial footprint as a result of distance from benthos and lens focal length. Differences in spatial footprints of images collected with different focal length lenses operated at the same distance from the benthos can be large. (A) shows the spatial footprint of imagery collected with a camera equipped with an 18-mm focal length lens (black dashed lines) and a 55-mm focal length lens (red dashed lines). Similarly, (B) shows portions of two images collected with the same camera operated at different distances from the bottom. For some applications, lower-resolution imagery will prevent detailed taxonomic assignments during data extraction.

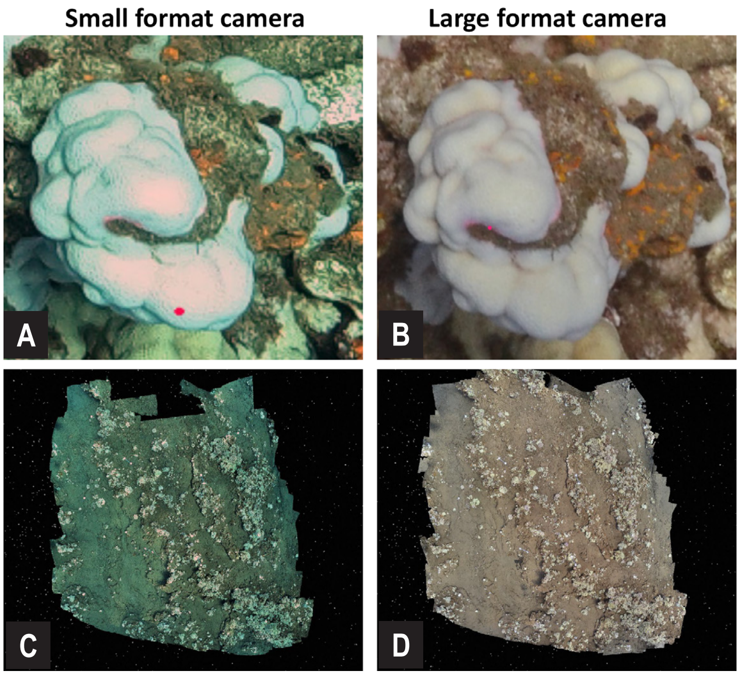

Figure 5. Large versus small format camera comparison. Raw imagery and resulting 3D models collected with a small format camera (GoPro 9, sensor size: 28.07 mm^2, [A] and [C]) and a large format camera (Nikon D780, sensor size: 864 mm^2, [B] and [D]). Cameras were operated simultaneously approximately 1.5 m above the benthos. Manual color correction was available only on the Nikon, leading the observed differences in color. The GoPro produced detailed imagery; however, it is grainy, less in focus, and overexposed relative to the imagery from the Nikon. The core areas of both 3D models are geometrically comparable; however, plot margins are better reconstructed using the higher quality imagery available from the larger format camera.

the lenses and underwater housings used, can dramatically affect how a given camera performs under different scenarios. Cameras with small sensors can produce imagery of sufficient quality when operated with care, and conversely, high-end cameras with large sensors will produce poor imagery when used inappropriately. While minimizing the GSD enables a wider array of cameras to be used, since the spatial footprint of imagery is reduced, more images must be collected to adequately cover a given plot.

2.1.3. Image collection pattern

The LAI approach is incredibly flexible, requiring only in-focus overlapping imagery collected over the entirety of the study area. However, the quality of LAI is largely a function of the quality of imagery and degree of overlap, and it is important to ensure consistency in the imaging approach across the entirety of plot being imaged. For small areas (<10 m^2), or single coral colonies, often no systematic approach is needed to comprehensively image the desired area. In smaller areas, divers can rely on memory to cover the plot with sufficient overlap, though a systematic approach is recommended to maintain consistent coverage. As the spatial extent of the plot expands, a systematic imaging pattern becomes necessary to reliably cover the entire area. The most common approach to obtain raw imagery for LAI is systematic collection of downward-facing (nadir) imagery acquired in a gridded pattern (though see Part II, section 5.6). The gridded approach is commonly utilized as it provides robust assurance that the study area has been consistently imaged and also provides views of the substrate from multiple directions. To aid navigation during image acquisition, visual markers such as transect lines, temporary markers, or floats are typically distributed throughout the plot. Additionally, an underwater compass for navigation and a depth gauge are particularly useful in cases of limited visibility or in other conditions where divers can more easily lose track of their positioning in the plot.

In many cases, imagery collected with the top-down gridded survey approach will result in accurate 3D reconstructions of the imaged scene, as the periphery of even moderately wide-angle lenses will capture side-on views that are otherwise not visible (i.e., occluded) from the top-down camera view (Figure 6). Thus, for most settings, estimates of percent cover, topographic complexity, or colony surface area can be accomplished with LAI generated from nadir imagery alone. However, nadir imagery alone is insufficient for some applications. Examples include highly complex sites with dramatic changes in depth or when complete closed 3D surfaces are required for analysis (e.g., for volumetric measurements or more detailed structural metrics) (Figure 7A). In these cases, oblique (i.e., side-on view) imagery can be collected to capture more views of the sides of object.

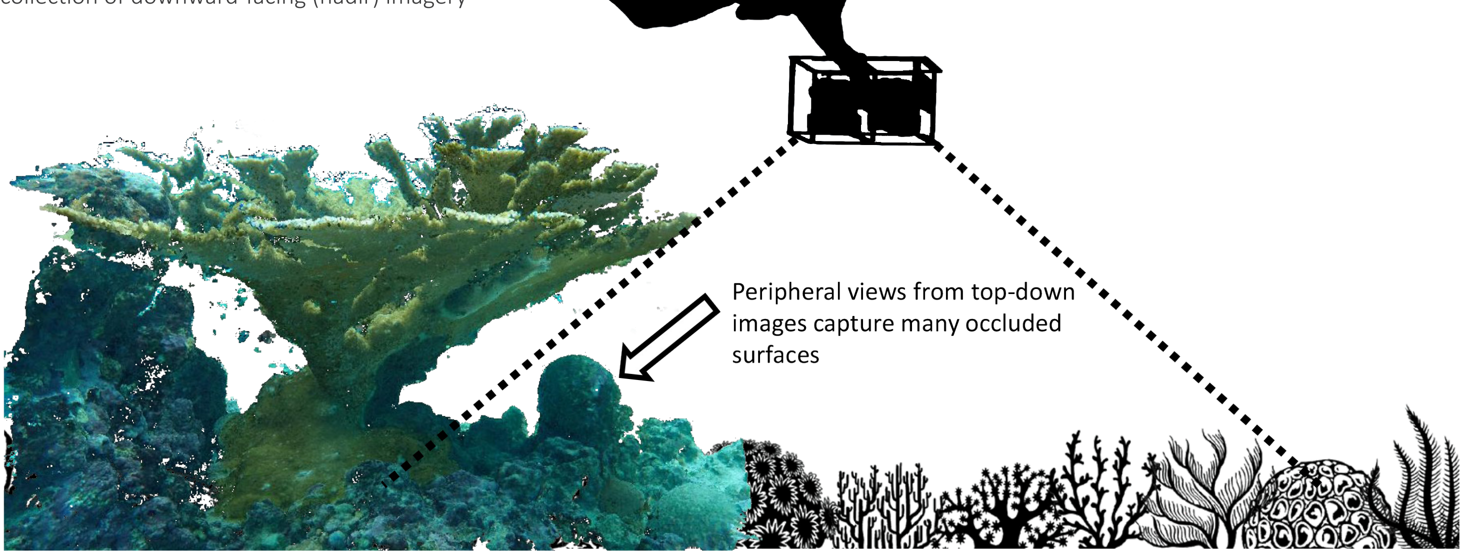

Figure 6. Example of views offered by peripheral views in nadir imagery. Image collection patterns that rely on top-down (i.e., nadir) views can create surprisingly complete 3D models. The periphery of nadir images collected adjacent to areas occluded from the top down can often "see" under these occlusions. When a sufficient number of images capture these views, the resulting 3D reconstruction will often include these features.

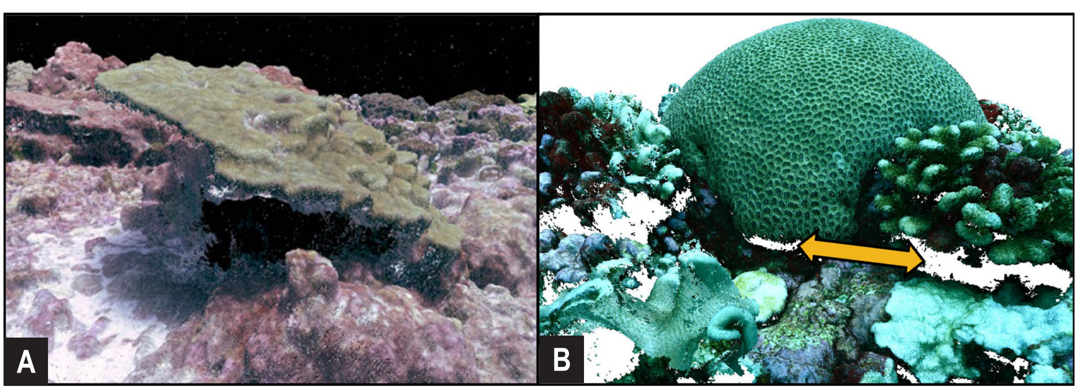

Figure 7. Example of additional area captured by oblique imagery. Nadir perspective imagery is insufficient to capture highly occluded views, such as the underside of the *Porites* colony shown in (A), preventing accurate measurements of colony volume. Collection of oblique imagery can allow for more complete reconstructions better suited to the extraction of volumetric data (B). However, habitat complexity can limit the positions at which the camera can be operated, thus limiting the comprehensiveness of 3D models, particularly for the undersides of colonies and more complex branching species.

Ultimately, however, even when care is taken to comprehensively image a scene, some views of the benthos simply cannot be captured due to practical limitations of operating cameras in complex environments (Figure 7B).

The capture of background scenery can negatively impact 3D model quality, and care must be taken when collecting oblique imagery. Due to light attenuation, items in the background will suffer from shifts in color balance relative to items in the foreground, with the effect becoming more pronounced the further from the camera these objects are. Under most camera settings, these background items will also tend to be out of focus or at lower resolution than items in the foreground, which can also lead to imprecise feature matching during SPC generation. At a minimum, color shifts or reduced sharpness can lead to a reduction in LAI visual detail of the dense point cloud (DPC) and at the worst can cause dramatic geometric distortion in LAI products (Figure 8). These errors can often present as floating blue points above an otherwise well-reconstructed surface (arrows). Some data extraction workflows can be influenced by these errors, and point confidence filtering tools available in Metashape or other software platforms can be used to remove or reduce their impact. When collecting oblique imagery, it is generally recommended to reduce the minimize the capture of background scenery by operating the camera closer to the bottom. However, imaging in this manner will require a greater number of images to achieve sufficient overlap and thus will also take significantly longer. Further, the complexity of reef habitats prevents divers from positioning the camera such that objects can be imaged from all angles, limiting the comprehensiveness with which they can be imaged, regardless of the camera being used (Figure 7B). Therefore, the collection of imagery to generate comprehensive 3D surfaces should be limited to appropriate physical settings and over limited spatial extents or with lower levels of replication. For discussion on the caveats associated with the processing of oblique imagery, see Part I, section 2.2.6.

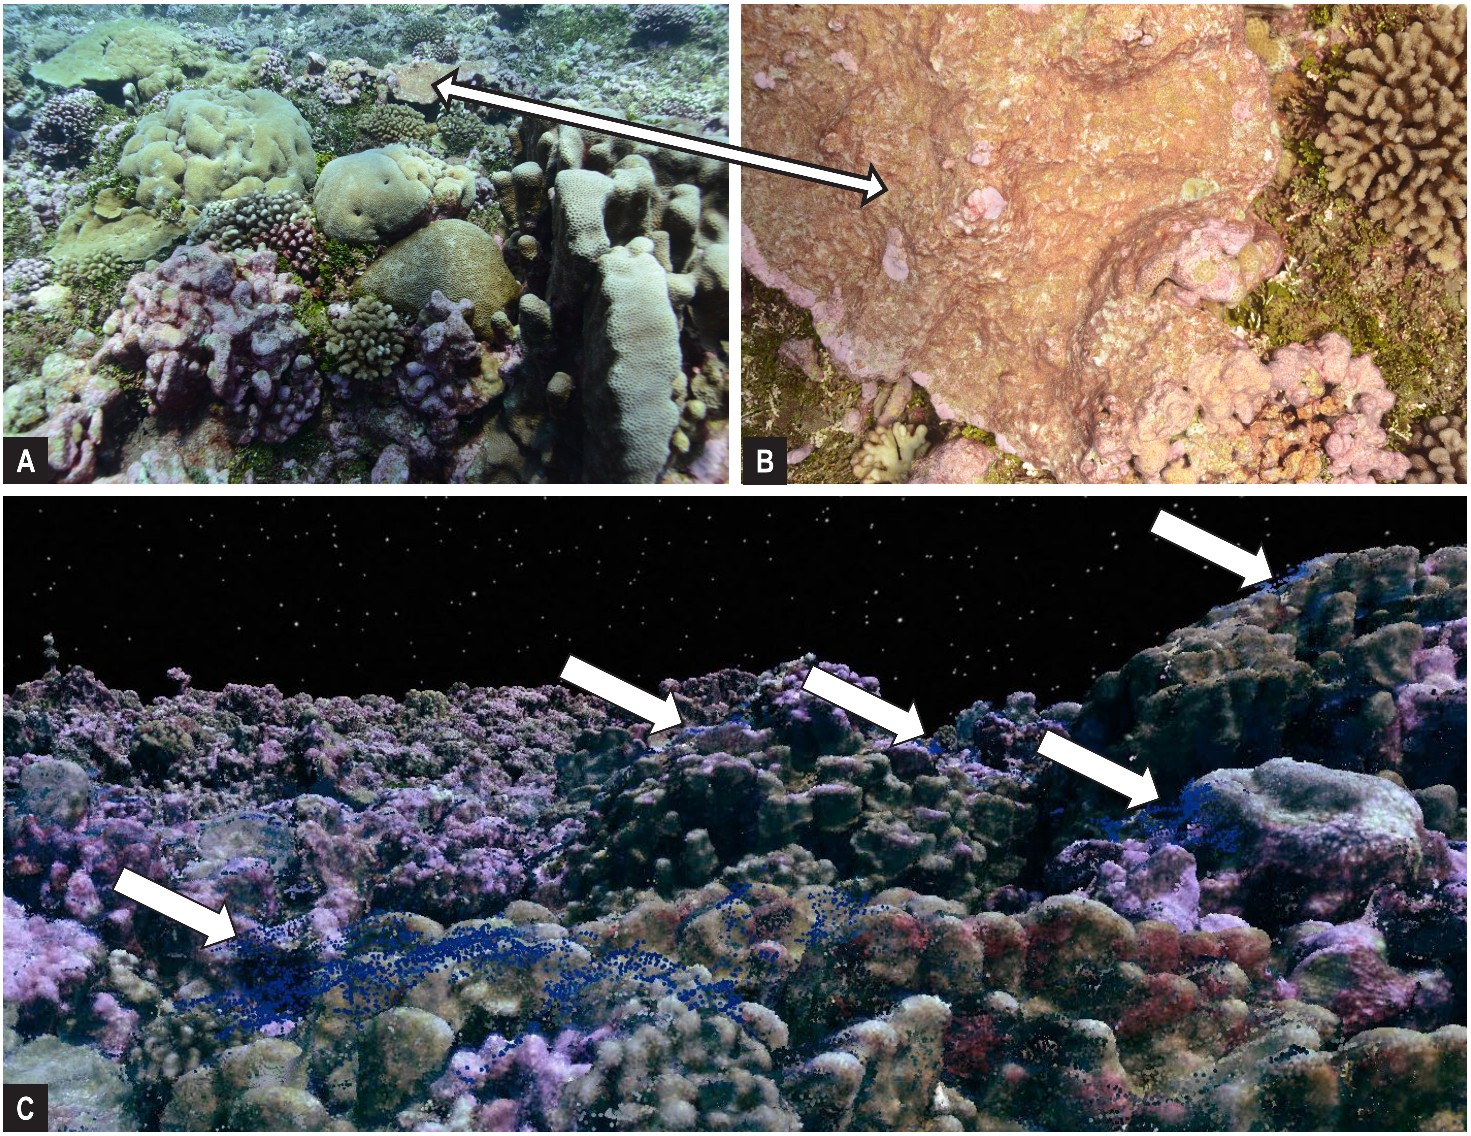

Figure 8. Example of light shift for items collected in nadir imagery, relative to the same item collected in the background of oblique imagery. In (A), a feature (arrow) is captured by an oblique image from a distance of over 5 m. (B) shows the same feature captured directly overhead from a distance of approximately 1.5 m. Large differences in resolution and color balance of this feature will lead to visual artifacts and reconstruction errors for these portions of the model as shown here in the DPC (C).

2.1.4. Metadata

Another fundamental aspect of the in-water approach of LAI workflows concerns the collection of associated metadata, most importantly, aspects of scale, orientation (i.e., the position with respect to gravity), and geographic location. Creation of accurate SfM-based LAI can be conducted without explicit assignment of scale or orientation. However, while the coordinate system of the resulting model will be internally consistent, scale and orientation must be applied at some point to derive most ecological metrics of interest. Providing scale is straightforward and is accomplished first by placing items of known size (e.g., scale bars) in the plot during image collection. Similarly, orientation with respect to gravity is accomplished by providing the depth of known objects within the scene. Geolocation is of fundamental importance to any ecological survey and is particularly useful when conducting time series analyses or any other effort that requires spatially co-located data. At a minimum, GPS coordinates should be obtained at known locations within the imaged plot or, when possible, at multiple permanent locations on the reef denoted by physical markers. These locations (often referred to as ground control points [GCPs]) can be temporary or permanent, but it is important that their location within the imaged area can be represented with a high degree of precision.

The use of GCPs can be complemented, or even replaced, when position information can be provided for each image. Currently, the most popular approach to provide georeferencing information to individual photographs is using ultra-short baseline (USBL) acoustic positioning systems paired with high-precision surface GPS information (i.e., real-time kinematic [RTK] GPS). This approach provides highly detailed information that can also be used to aid underwater navigation and speed model construction (see Part I, sections 2.2.2, 3.8, and 3.11). Some additional expertise is needed to operate RTK GPS and USBL devices and post-process the resulting data before they can be useful for model construction (Gerke and Przybilla, 2016; Benassi et al., 2017). Further, underwater acoustic communication might not be appropriate in settings where line of sight between surface and subsurface units cannot be maintained. As a result, use of RTK GPS and USBL devices should be limited to those situations where the information they provide is required to meet project objectives or where the capacity to effectively use these devices already exists or can be obtained without sacrificing capacity at other steps in the LAI pipeline.

Scale, orientation, or geolocation information obtained during image collection can also be used during the later steps of model construction. In cases where objects within the scene can be scaled and georeferenced with high levels of precision, or when georeferencing is available for individual images, this information can be used to optimize the reconstruction process, reduce processing times, and increase the geometric precision of LAI products (Crandall et al., 2011; Eltner and Schneider 2015; James et al., 2017). In some instances, it is more straightforward to apply this information to LAI products in later steps of data extraction, depending on the choice of software for model construction or analysis. The benefits of using image-level GPS information or GCPs during model construction are addressed briefly in Part I, section 2.2.2. Approaches and specific instructions for applying this information and conducting the optimization process are explained in Part II, sections 4.6 and 5.6.

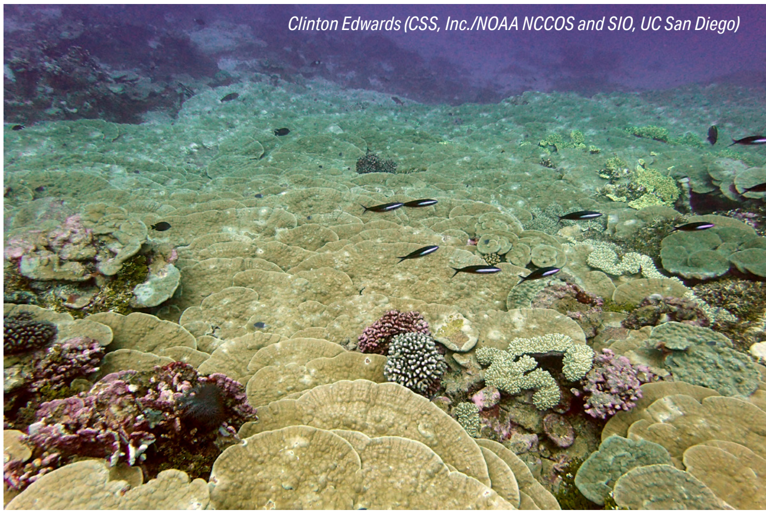
Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, UC San Diego)

2.2. Model construction

2.2.1. Overview

With in-focus overlapping imagery acquired across the entire scene, accurate LAI can be generated. The key steps of LAI model generation that are of general importance to coral reef scientists are the creation of the primary 3D products (e.g., sparse and dense cloud reconstructions and meshed models), and the derived 2D products (e.g., orthophotomosaics and digital elevation models [DEMs]). In practice, these products can be generated in a single SfM-based software platform. Commercial software packages currently available (e.g., Agisoft Metashape, Pix4d, and VisualSFM) have broad functionality and are well vetted by industry standards (Jiang et al., 2020). Combinations of different commercially available or open-source tools can be used to create customized processing pipelines (e.g., Carbonneau and Dietrich 2017; Forsmoo et al., 2019); however, this requires considerable computer science expertise. A major advantage of the approach taken by these programs is that accurate LAI can be generated from the images alone. Scale, orientation, and position information can be used in some cases to speed model creation or to improve model geometry but is not required. The generation of 3D models is a computationally intense process, and the time needed to process raw imagery into the desired LAI products is a function of the number of images that were collected (and their size), the specific settings used during model creation, and the computer configurations. While many of the details discussed here are relevant regardless of the particular software used, this document will refer almost exclusively to the program Metashape (Agisoft LLC, St. Petersburg, Russia). Metashape offers an easy-to-operate interface, consistently rates at the top of both speed and accuracy tests under a variety of conditions, and is widely used in various fields including the environmental sciences and archaeology (Sona et al., 2014; Nouwakpo et al., 2016; Kingsland 2020).

2.2.2. Sparse point cloud reconstruction

The first product that is generated as a part of the LAI approach is a 3D model known as a sparse point cloud (SPC). The generation of the SPC in Metashape begins with the detection of features (key points) that are then matched in overlapping raw images to create tie points, which are used to estimate the camera position (pose estimates) in relation to the imaged scene. Camera pose and tie point position estimates are refined through a process known as bundle adjustment. Once the locations of key points have been estimated, they can be visualized as a 3D model, known as the SPC. The accuracy of these estimates is largely a function of the precision of matched features across image pairs, as well as the degree of overlap and the number of high-quality matched features between image pairs. The more precise the identification of these features, and the more similar they are across image pairs, the better the tie point estimates will be. For example, if a given feature is precisely located in some overlapping images but less precisely in others (Figure 9), the estimated position of this feature is likely to be less precise than for a feature that is precisely identified across all images containing that feature. The quality of key point estimates along with camera and tie point positions thus set the geometry of the scene. As a general rule, image quality with respect to matched features, and the degree of overlap and quality of tie points among adjacent images, will largely drive the quality of the reconstruction.

One of the overarching benefits of SfM is that it can be used to create accurate 3D models using only overlapping imagery of a scene. However, when imagery collected in the field is georeferenced, this position information can be used as an initial guess of camera position. Continuous position information collected with a USBL device (or other means) is post-processed to provide each image with a position estimate that can be embedded directly into the image file or provided directly to Metashape as a separate file. These initial estimates of position are then used to reduce the number of possible image pairs during the tie point matching step. Particularly when working with large image datasets (>10,000 images) this approach will dramatically reduce processing times. In some cases, these reductions may be necessary to achieve project timelines and can also decrease the required computational resources. Image-level positioning information can be particularly important to improve the reconstruction of large datasets, as small reconstruction errors can accumulate over space, particularly for high-aspect-ratio plots (e.g., long rectangular transects), resulting in systematic geometric distortion over large image sequences or spatial scales (Figure 10; Lhuillier 2012).

Along with georeferenced imagery, objects of known position or scale in the scene (e.g., GCPs), and measurements between them, can also be used to guide or correct SPC construction through a process known as optimization. The spatial information can be provided to the software as scale bars, GPS locations (or positions in some other internally consistent coordinate system), or depths of GCPs or other objects that can be identified in raw

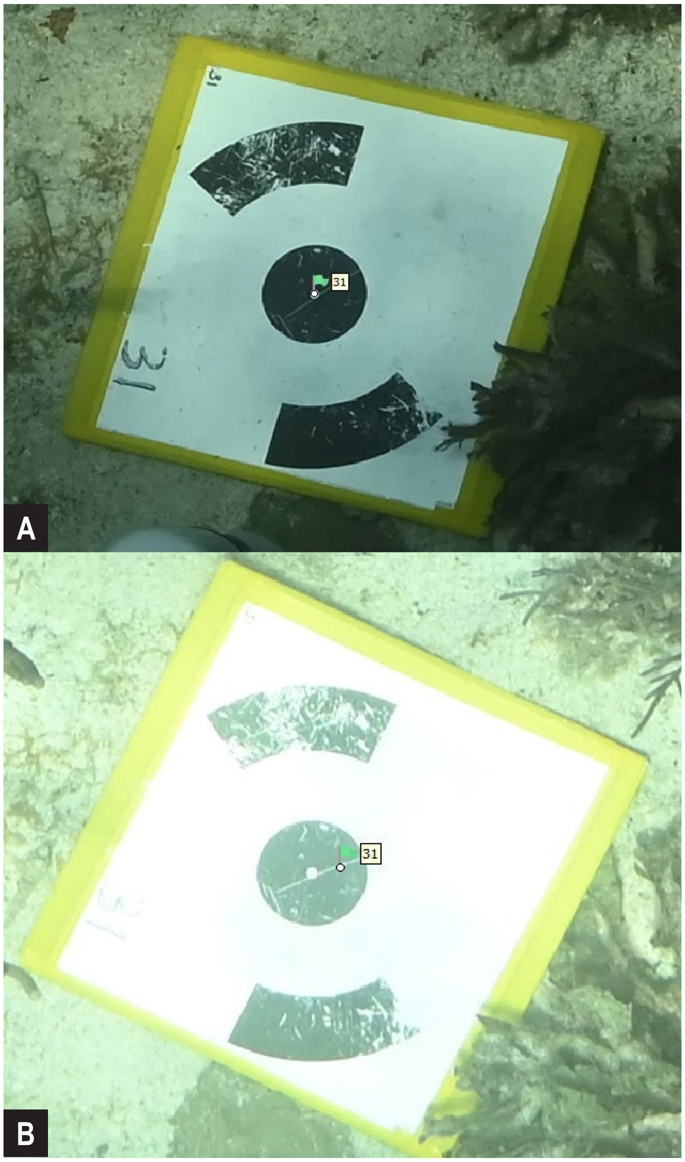

Figure 9. Error in location of matched features as a result of image quality. Image pair showing varying accuracy of a key point for a feature identified during SPC generation in Metashape, in this case, the center of a Metashape target that is already known to the system. The first image (A) is crisp and well exposed, and the center of the feature is identified with high precision (green flag). The second image (B) is overexposed and less crisp, and while the center of the feature has again been identified, the key point is displaced from the true location of the feature.

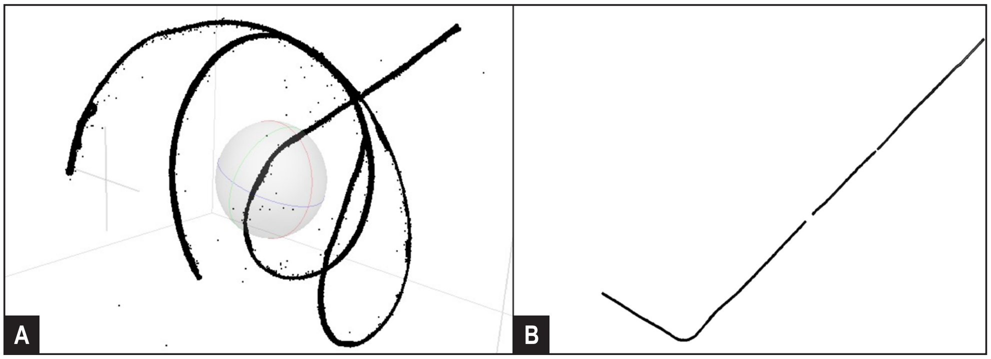

Figure 10. Reconstruction error associated with high-aspect-ratio plot size. Example of accumulated reconstruction error that can occur in high-aspect-ratio rectangular plots. The SPC shown in (A) was created from a sequence of 50,000 images of a 5-mi-long sewage outflow pipe. Images were collected along a single linear swath, and the SPC was generated without any external inputs of position. While SPC geometry is “reasonable” over short spatial scales (<5 m), accumulated errors over greater lengths lead to highly inaccurate geometry over longer scales. (B) shows the corrected SPC generated from the same set of images but with the addition of USBL positioning data for each image.

imagery. Metashape (or other software) then uses this spatial information to validate the geometry of the SPC and make any needed corrections. There are a variety of approaches to add this information in Metashape, including semi-automated protocols using automatic detection of coded targets that are specifically designed to provide highly accurate feature detection and matching for the images containing these targets. Key points can also be manually edited to improve their accuracy before optimization.

When image quality or overlap is poor, even if only for a portion of the imaged area, optimization may be required to correct SPC geometry. Currently, workflows for adding GCPs and performing optimization procedures can be time consuming, particularly for datasets containing thousands of images. While obtaining image-level position data can require some additional buildup of infrastructure and expertise, once this information is in hand, it can be used for SPC generation with little additional effort. In all cases, when imprecise GPS or GCP information is used during SPC generation or optimization, it can lead to relative reductions in the geometric accuracy (James et al., 2017). Regardless of the amount or quality of external information that is provided, when there is too little overlap among adjacent images, some portions of the imaged area will not be reconstructed into a 3D model. To avoid these situations or the need for labor-intensive post-processing steps, it is strongly encouraged to develop a robust image collection approach that maximizes the quality of the image dataset. For more information on the optimization procedure, please see Part II, sections 4.6 and 5.6.

Dense point cloud reconstruction follows image alignment, camera pose estimation, and generation of the SPC, using information from individual images to generate a more detailed 3D model, commonly referred to as the dense point cloud (DPC). DPC generation relies on information from the SPC to generate depth maps from pairs of overlapping images, which provide estimates of each source pixel's depth (distance from camera) based on stereo triangulation (similar in principle to biological binocular depth perception). This per-pixel depth information is then used to essentially project pixels from the raw imagery to 3D points and fill in the SPC with texture and color, allowing for the generation of 3D models with high levels of visual detail.

2.2.3. Dense point cloud reconstruction

While SPCs contain the bulk of the spatial information of the imaged scene, they lack the detail needed for visualization and many subsequent ecological analyses. In order to produce the visually detailed LAI products of interest, a DPC must also be generated. Once the SPC has been generated, pairs of overlapping images are processed with a stereo reconstruction algorithm to yield per-pixel depths for each image, which are then projected onto the SPC using the camera pose and lens parameter estimates. As a result, though information has been added, the geometry established by the SPC itself will remain unchanged during DPC generation. However, when key point and tie point estimates are themselves imprecise, the geometric quality of the DPC will also suffer. These imprecisions present as a "rough" or "fuzzy" surface texture on what should otherwise be smooth, high-resolution surfaces. DPC reconstructions can further suffer when adjacent overlapping images have different exposures or color balance, as these differences will not only affect the precision of feature matching but will also lead to visual artifacts in the DPC (Figure 8). SPCs and DPCs can be viewed separately or at the same time; however, as they contain redundant information, often only one of these products is needed for subsequent analysis.

2.2.4. Derived LAI products: meshes and DEMs

A number of additional LAI data products can be generated after the sparse or dense point clouds have been reconstructed. One of the commonly used LAI products is a DEM, which is typically presented as a 2D raster image where depth (or elevation) values are stored in each pixel. For visualization purposes, these depth values are frequently converted to RGB or grayscale values (Figure 11). DEMs provide useful 2.5-dimensional (2.5D) representations of scenes that allow for straightforward visualization of depth gradients and topographical relief and can also be used for calculations of structural complexity or surface area. The SPC, DPC, or mesh model can be used as the source data for the DEM, and precision and accuracy of the DEM is a function of both the specified DEM grid cell size and the resolution of the source data. The particular approach used to generate the DEM will also determine how the depth value is reported for each pixel, and in some cases can further obscure important structural features.

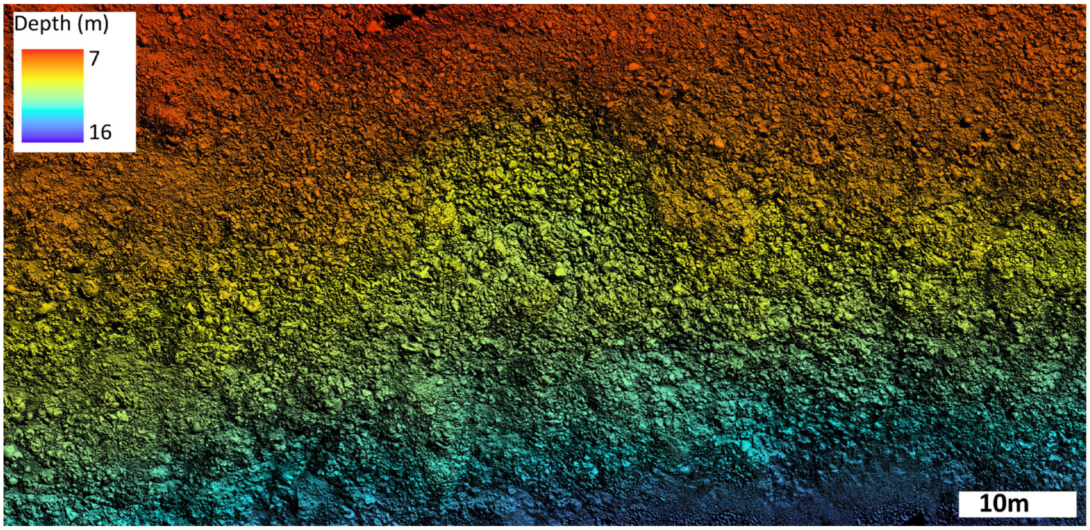

Figure 11. Example of digital elevation model. DEM of a coral reef on the island of Curacao. The color ramp goes from warmer colors (shallower depths) to cooler colors (deeper depths), with the gradient clearly showing the reef slope at this location.

Digital elevation models (DEMs, also known as digital surface models, or DSMs) represent sites' 2D raster images, where each pixel or grid cell has a depth (or elevation) value. Since only a single depth value can be specified for each horizontal 2D location, DEMs cannot represent overhanging 3D structures and are thus commonly referred to as 2.5-D products. Depth values can be used to assign color values for visualization (using an elevation color ramp or topographic relief shading) or used directly for a variety of structural analysis, such as slope and curvature, or to compute 3D surface area measurements. DEMs are generated using a number of different technologies, including acoustic and lidar bathymetric mapping, as well as satellite-based remote sensing approaches. In the context of the LAI approach, DEMs can be generated directly from SPCs and DPCs, as well as from meshed 3D models. As a result, the resolution of the DEM is dependent both on the specific settings used during the generation of the DEM, as well as on the resolution of the source data.

3D meshes are a popular 3D model product comprising a set of flat surfaces (often triangles) with image textures optionally applied to provide greater visual detail than would be afforded by mesh geometry alone. A major advantage of meshes is that they can represent continuous "water-tight" 3D surfaces, which provide bounded closed volumes and are mathematically more tractable than point clouds for extraction of volume or surface area metrics. Meshes can be constructed from dense or sparse point clouds, or directly from depth maps, with the number of constituent faces generally constraining how well the mesh geometrically represents the source data. Textured meshes with relatively low face counts are frequently constructed to be used as lightweight, more easily processed—but lower-fidelity—site representations.

For instance, for DEMs representing structurally complex habitats but exported at relatively low resolution, single pixels may cover a range of depths but will report only a single value. Thus, decisions on whether to present the maximum, mean or median elevation can result in dramatically different depth values for a given pixel. Further, depending on the number of points in the source point cloud and the desired spatial resolution of the DEM, assumptions must be made to discard noisy points and, in cases where points are spatially sparse, interpolate between them.

Another commonly used derived LAI product is a meshed 3D model, which is represented as a 3D model with a closed surface. Meshes can be derived from the sparse or dense point clouds, or from DEMs, depending on the detail needed in the mesh and that available in the source data. A variety of mostly automated approaches can be used to generate meshes, and when using Metashape, the user needs to supply only the number of surfaces to be created from the source data to generate the mesh. Following mesh creation, users can specify the texturing mode, or coloring overlay and resolution, both of which are largely derived from the raw imagery. When meshes are generated from high-density point clouds, depending on the settings used during mesh generation, the meshing process can result in large amounts of discarded data. This reduction in data is particularly large when the face count of the mesh is low, or a large DPC is used as the source data. Importantly, the meshing procedure is required to create closed surfaces and enable important measurements such as coral colony volume. Further, as meshes typically have much smaller file sizes than the point clouds from which they are generated, they tend to be much more tractable for visualization purposes on a variety of widely available software platforms. Moreover, there are numerous details to consider when generating meshes, including processing time, mesh model assumptions, and decisions regarding needed detail. Readers are encouraged to reference the SOPs provided in Part II of this document for more guidance on the mesh generation procedure.

2.2.5. Derived LAI products: Orthophotomosaics

For many applications, the most useful products generated with the LAI approach are orthophotomosaics, which provide 2D map views of the imaged scene. Depending on the approach taken when collecting imagery, orthophotomosaics can offer expansive and highly detailed map-like views of underwater habitats. Corals and other benthic organisms can be identified with high levels of taxonomic specificity, their locations mapped and their sizes accurately measured in a relatively straightforward manner with a number of commonly used software platforms. As a result, orthophotomosaics enable a wide variety of analyses including investigations of coral demography and spatial patterns and have been one of the more widely used LAI products to date for coral reef monitoring and research. As with the other derived products, the quality and detail provided by orthophotomosaics rely on the quality of the underlying raw imagery and a series of assumptions that should be explicitly considered before they are generated.

Orthophotomosaics are generated by stitching together overlapping raw images that have been corrected for scale. In standard perspective imagery, the perceived scale at any point in the image is a function of the distance of the object from the camera. As a result, two objects of the same size but at different distances from the camera lens will appear to be of different sizes. If all objects in the scene are at very large distances from the camera (e.g., satellite imagery), the effect is negligible. However, at smaller distances and in topographically complex habitats, perspective distortion can lead to dramatic differences in geometry and relative size of similar objects within a single photograph. Using information from the structure of the 3D model, this distortion can be corrected using a transformation technique called orthorectification. After images have been corrected, adjacent and overlapping images can then be blended and stitched together to create an orthophotomosaic (Habib et al., 2007; Lirman et al., 2007; Nicosevici et al., 2009). Individual images collected with a typical high-end full-frame camera measure 6,000 pixels × 4,000 pixels. However, stitching together many images of this size will create much larger orthophotomosaics. Once images become larger than 30,000 pixels × 30,000 pixels, file sizes become large (>1 GB) and their use can be unwieldy in many software applications. Therefore, when orthophotomosaics are generated from thousands of raw images, significant downsampling

Orthophotomosaics refer to 2D map views of 3D models and are among the most commonly used LAI products in ecological analysis. A key step in the generation of orthophotomosaics is the process of orthorectification (also known as orthoprojection), which uses the structure of the 3D model to correct for perspective-based distortion in the sizing of objects. Together with estimates of camera positions, individual images are orthorectified and blended together to generate an "orthophotomosaic," a 2D map view of the scene. These images can provide visually detailed top-down map views of the benthos that allow for a variety of metrics (e.g., coral size, abundance, and spatial patterns) to be extracted for analysis.

(e.g., compression) of source imagery is needed. Alternatively, orthophotomosaics can be generated with less downsampling if they are subdivided into smaller chunks and analyzed separately. However, this workaround requires subsequent reassembling of data, and given the overall complexity and number of steps already involved, caution is recommended before any chunking of data.

Orthophotomosaics present some challenges as a result of the perspective distortion of raw imagery captured in topographically complex areas (Figure 12). The orthorectification process attempts to correct for perspective-based distortion, but when a single photograph contains dramatic differences in vertical relief, there are limits to the correction process. As a result, orthophotomosaics will sometimes suffer from distortion in topographically complex regions where there must be significant amounts of transformation to raw imagery before adjacent images can be blended together, particularly when adjacent images are subject to different degrees of change. An alternative approach is to directly orthorectify the DPC itself, thus avoiding the potential errors associated with orthophotomosaic creation. Orthoprojected DPCs can be captured as image files and analyzed in the same manner as orthophotomosaics (Kodera et al., 2020; Sandin et al., 2020; Couch et al., 2021). This process removes the need to blend and distort the 2D source images by directly using the 3D point cloud, which provides the most accurate geometric reconstruction of the reefscape. However, the visual detail of point cloud orthoprojections is dependent on the quality of the DPC rather than the high-resolution source imagery, and as a result, this approach comes at the cost of reduced resolution, though the effect tends to be minimal (approximately 1-mm reduced GSD). The point cloud orthoprojection approach is an area of active development, and efforts are underway to improve the visual detail of the resulting orthoprojections. Ultimately, the difference in resolution between an orthoprojected point cloud and an orthophotomosaic generated from the same underlying 3D model will depend on the visual detail of the DPC and the level of downsampling during orthophotomosaic construction. Regardless, any lack of visual detail in point cloud orthoprojections or orthophotomosaics is offset by the ability to spatially query and inspect the raw imagery for any location in the orthophotomosaics or point cloud orthoprojection. Using the raw imagery allows for more detailed taxonomic or segmentation decisions to be made. For more information, see Fox et al. (2019), Electronic Supplementary Materials 1; and Part II, sections 4.6 and 7.6 of this report.

2.2.6. Practical limitations to model construction

The generation of sparse and dense point clouds are computationally intense tasks and in most cases must be conducted on hardware allocated solely to that purpose during the model

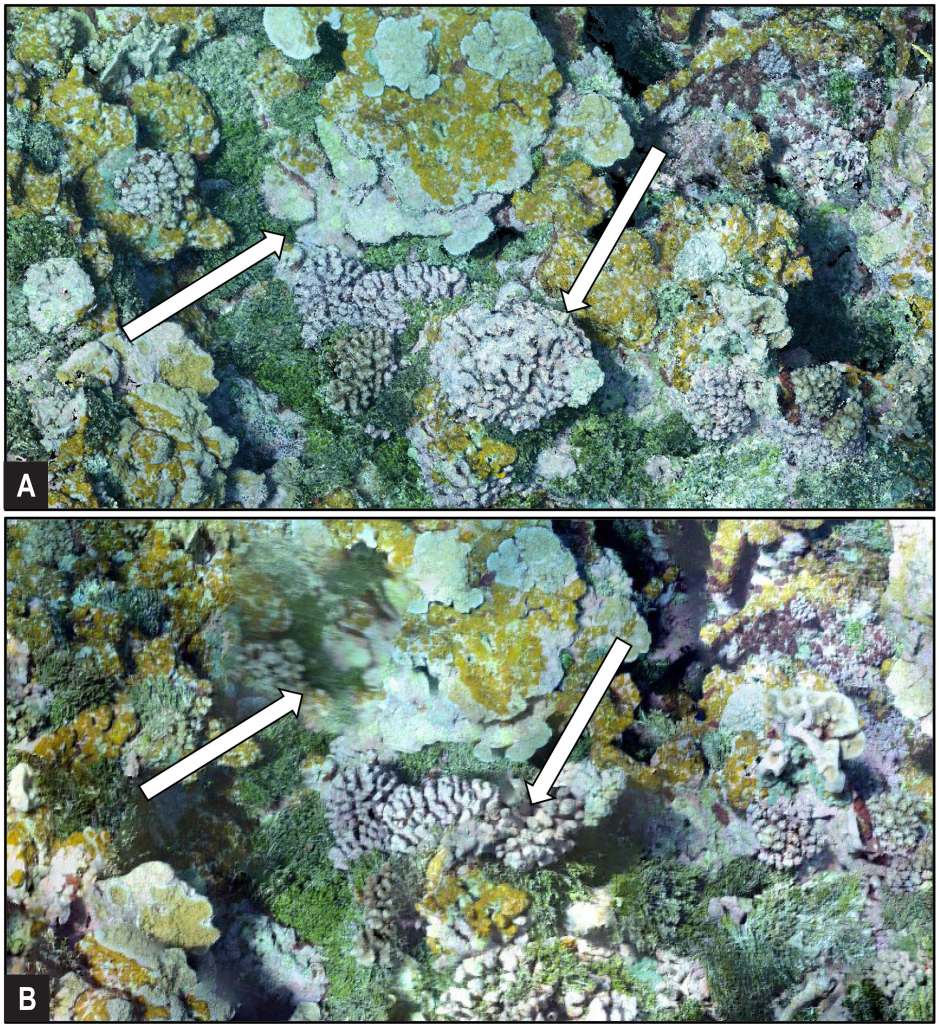

Figure 12. Example of distortion resulting from blending of raw imagery during orthophotomosaic construction. Comparison of an orthoprojection generated from a 3D point cloud (A) and an orthophotomosaic generated in Metashape using blended orthorectified imagery (B) in a topographically complex area. The point cloud orthoprojection is contiguous and without distortion, while stretching and blurring in the orthophotomosaic (B) is visible in the high-relief areas (arrows), with some features missing and others dramatically deformed or reduced in size.

construction procedure. The time needed to generate the SPC is primarily a function of the specifications of the computer hardware used (e.g., number of CPU and GPU cores), the number of images used (and their size), and the specific key point and tie point settings used. The maximum number of images that can be processed together at a single time is primarily limited by the amount of available system memory (i.e., random access memory [RAM]). Metashape is not currently optimized for super-computer architectures, and purpose-built high-performance computing (HPC) platforms tend to achieve the best performance (https://www.agisoft.com/downloads/system-requirements/). Metashape provides an easy-to-use networking interface that allows a single Metashape project to be run across multiple HPCs connected via a high-speed LAN to reduce the total required processing time (see Part II, section 7.3). There are limits to the number of images that can be handled by Metashape's graphical interface, and when collections exceed 60,000 images of 24.5 MP each, some functions, particularly those that require use of image thumbnails, will become unreliable. However, given the time needed to collect such image datasets and extract meaningful data from them, care should be taken before attempting imaging at this scale. Regardless of system configuration, or the total number of images used, the time required for feature detection and image matching can be significantly reduced when high-quality raw imagery is collected with a high degree of overlap. More generally, the computational approach underlying LAI is relatively new and will inevitably continue to evolve in the coming years, as will the computer hardware available to execute the approaches outlined here.

The SfM-based approach to generate LAI is incredibly robust, yet there are important practical constraints to the process that should be considered. While generating 3D models and other derived LAI products requires relatively little user intervention, in some cases, it might be necessary to perform filtering and color correction on imagery collected in the field before model construction can begin. As mentioned in section 2.2.3, when imagery is collected in an oblique fashion, it is important to avoid capturing features in the background that are also imaged at a closer range (Figure 8). When a given feature is imaged at a range of distances, there can be large differences among images in the relative GSD of the pixels associated with that feature, as well as a greater likelihood for differences in focus between images for the feature in question. There will also be a relative color shift between these images, which, together with the differences in resolution and focus, will reduce the precision of feature identification and matching across these image pairs. Similarly, differences in the exposure of a feature can reduce the precision of feature matching, and for this reason, it is generally recommended to use ambient lighting to avoid exposure artifacts due to shadowing associated with the use of artificial lighting. For the vast majority of tropical shallow-water applications, the light sensitivity of most cameras is more than adequate to preclude the need for artificial lighting. If lighting is used, care must be taken to minimize the presence of shadows and avoid any illumination of any large particulate matter in the water column, both of which can be challenging when collecting imagery in large areas and when time is limited. The impact of minor differences in GSD, focus, color balance, and exposure are relatively limited; however, when these differences are pronounced, they can lead to distortion in the portions of the model containing this imagery (Figure 8).

There is no solution to compensate for differences in the focus or resolution of a feature, though color balance and exposure can be adjusted before model generation using image editing software. When imagery is collected with sufficient overlap and redundancy, it may be possible to manually filter or mask those images with too much background scenery, though this can require considerable effort for collections with large volumes of images. When manual white balance is conducted with a color card underwater and imagery is collected in a raw file format (e.g., CR2, RAW, SRF, or NEF), image processing software can be used for color correction to minimize the impact of shifts in color balance. Color corrections can also be particularly important when there is significant depth variation within plots to maintain consistency in DPC color. Recent advances in the ability to apply information from image depth maps and other position information may soon expedite the ability to perform color correction and other needed post-processing to imagery (Akkaynak and Treibitz, 2019). Currently, however, image processing, particularly performing color corrections to raw image formats, is a time-consuming process, and care should be taken to design an image collection approach that minimizes the need for post-processing of images collected in the field.


Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, UC San Diego)

Mobile organisms and other features of the scene that are not stationary (e.g., sea fans, fish, and fleshy algae) cannot be represented with consistent tie point positions. This is important because the SfM-based approach for generating LAI requires matching of multiple stationary features in overlapping images (Bryson et al., 2017). When items shift subtly during imaging, the reconstruction of this portion of the plot will tend to appear "fuzzy" due to imprecise geometry in these regions. In other cases, the impact of these non-stationary items will be sparse or incomplete LAI in these areas. In the more extreme case, items which are imaged, move, and are then subsequently reimaged, can lead to severe geometric deformations in the reconstructed 3D model, such as duplications or "ghosting," appearing twice or in the wrong location altogether (Figure 13A). The presence of non-stationary items will additionally result in some images, or portions thereof, being lost from the reconstruction process, leading to reductions of overlap and further degradation of reconstruction accuracy in these portions of the model.

LAI products are generally generated from a collection of spatially contiguous, yet finite, set of overlapping images. For applications where adjacent plots are combined into a single contiguous LAI product, imagery collected at the margins of each plot must still have sufficient overlap, which can be challenging when there is an interval between surveys or no fixed markers present. Imaging an area in multiple surveys can also lead to differences in the lighting of overlapping images at these plot margins, which can lead to the previously mentioned reductions in the precision of feature identification and matching, and overall lowered model quality. If surveys must be conducted over multiple days, this interval should be minimized to reduce the possibility that items on the bottom (e.g., survey markers or loose substrate) move during the survey interval. Further, when possible, multiday surveys of a site should be conducted at the same time of day to minimize differences in lighting.

2.3. Ecological data extraction

2.3.1. Overview

The first two steps of the LAI pipeline described previously, 1- image collection and 2- model construction, prepare the imagery data for the third step of the pipeline: 3- ecological data extraction. LAI products represent digital snapshots of sites at the time of image collection, allowing nearly unlimited digital access to sessile benthic habitats and collection of virtually any information not involving the collection of physical samples. The LAI approach is most suited to produce three categories of data: 1- community composition data, the benthic cover, abundance, and size of organisms; 2- spatial data, the position of organisms across the landscape and with respect to each other; and 3- physical data, the structural attributes of the surveyed area.

The objectives of ecological analyses can vary extensively, as can the approaches used to generate the data needed for analyses. For some projects, summary-level information (e.g., whether corals are present or not) is sufficient, while other projects may require information in addition to image survey products. Workflows exist in several commercially available and custom software platforms to extract classic data metrics, such as percent cover and coral size abundance from both 3D and 2D LAI products.

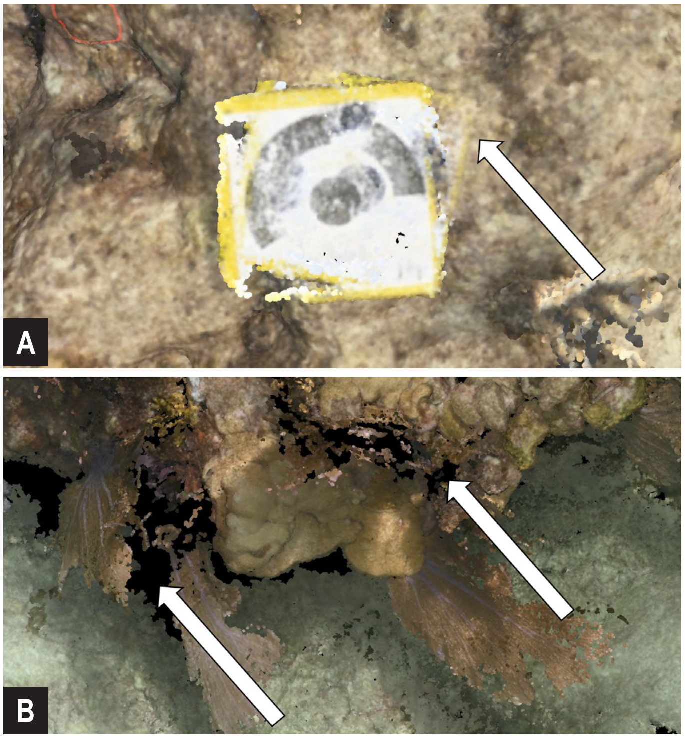

Figure 13. Reconstruction error arising from non-stationary objects during imaging. When items are imaged, shift subtly, and are then captured in additional images, they can appear duplicated or incomplete in the resulting reconstruction (A, arrow). Similarly, non-stationary items, such as the sea fans shown in (B, arrows), will often reconstruct if they are captured in enough images but will do so incompletely. Further, adjacent stationary features that are intermittently occluded by non-stationary features captured in the same images can fail to reconstruct as a result.

There are also opportunities to develop software that allow for additional or novel metrics to be extracted, such as structural complexity (Figueira et al., 2015; González-Rivero et al., 2017; Pascoe et al., 2021) or spatial analytical approaches (Burns et al., 2016b; Edwards et al., 2017; McNamara et al., 2019). The development of these tools is highly reliant on collaboration with other disciplines such as engineering and computer science. This collaboration represents an influx of talent and skills that alleviates some of the burden of technological development from coral reef scientists, allowing them to focus their energies toward research objectives. Because LAI archives both the raw data (the view of the ecosystem) and the data extracted for analyses, additional post-processing or analyses can be incorporated at any time as new lines of inquiry arise. However, LAI is not a panacea, and once data have been extracted, the classic challenges of scientific hypothesis testing, data analysis, and summarization remain. Guidance is provided on the approaches to producing detailed and well-organized ecological data.

More generally, LAI represents digital snapshots of study sites that allow users to visualize and explore the system without the logistical time constraints associated with in-water work and with the advantage of computationally enhanced tools. 3D models and other derived LAI products are often referred to as raw data, as it is only once they are in hand that ecological data extraction and subsequent analysis can be conducted. In some cases, however, these visual datasets are all that are needed to draw conclusions about reef state and change. When working with the public, policy makers, or funding agencies, visual LAI data are some of the most important pieces of information to communicate what is needed. Graphs cannot convey the beauty and complexity of coral reefs in the way that visualizations of the reefs can themselves, and LAI can be a vehicle to share this critical information. Arguably, these digitally enhanced data, with their ability to allow virtual exploration, can reveal never-before-seen observational insights by letting researchers see the world in ways that were previously not possible without ever having to enter the water.

2.3.2. Estimates of percent cover

The percentage of the benthos occupied by corals or other sessile animals is a simple and robust metric of ecosystem status and is one of the most commonly used metrics in coral reef science. Historically, benthic cover is measured in the field using the point intercept method, either along a transect line or inside a gridded quadrat (Jokiel et al., 2015; CRCP 2022a). This method produces highly accurate data but is dependent on the ability of the expert observer to identify taxonomy in situ and requires considerable field effort or limited replication when applied over large spatial extents. More recently, image analysis software has been increasingly used to randomly sample points on a planar digital image and provide a class label to these points (i.e., the photoquadrat method). With the photoquadrat method, multiple planar images can be taken of a site and analyzed later in the lab where ample reference information can be consulted when making designations (Kohler and Gill 2006; Jokiel et al., 2015).

While a robust method, photoquadrat image collections must maintain consistent top-down views of the substrate, and the spatial extent that can be analyzed is limited to individual images. As with any image-based method, the level of detail with which taxonomy can be assigned is restricted by the quality of the imagery collected in the field.

The LAI approach has similarities to both in situ and photoquadrat approaches and also offers significant improvements. Both standard photoquadrat imagery and LAI (orthoprojections or 3D models, depending on the approach used) can be randomly sampled with points that can be labeled by taxonomic class to produce percent cover estimates (Fox et al., 2019). However, with LAI, cover estimates can be determined over a much larger spatial footprint. All points in the 3D model are constructed from many overlapping raw images, each with a slightly different perspective of the benthos. These complementary views allow multiple images and perspectives to be referenced when assigning taxonomy at each sample point, providing a key advantage over the photoquadrat approach (Figure 14). The ability to access the raw imagery provides redundancy in the event that some of the images are not of sufficient quality and also provides higher-resolution views than are generally available from orthophotomosaics, meshes, or point clouds. Additionally, access to multiple high-resolution views of different regions of LAI products enables a suite of other data streams to be extracted,

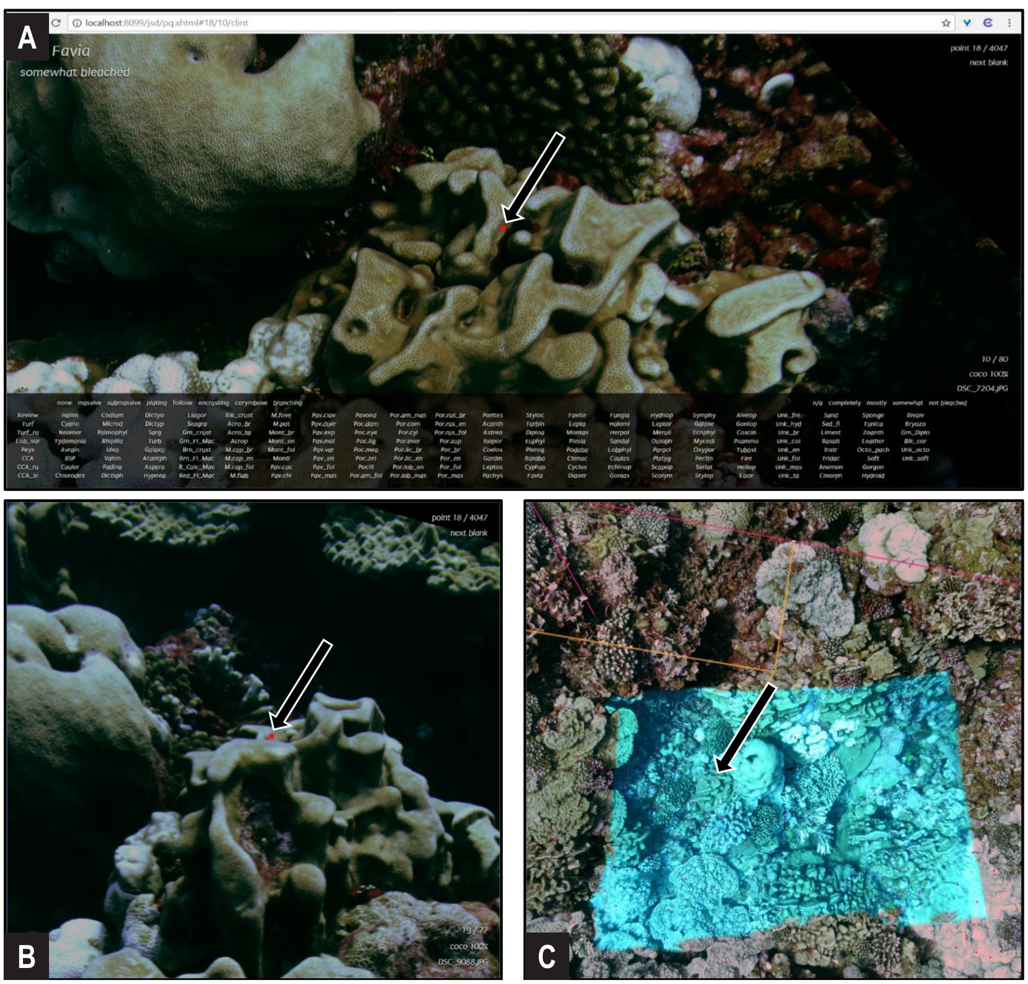

Figure 14. Example of raw imagery access and multiple views made available. The visual-analytical software platform Viscore can be used to provide ready access to high-resolution imagery in a separate window (A) and (B) or draped directly over the 3D model (C). This approach allows multiple views of the same location in the model (A)–(C), arrow/red dot), allowing the best possible information during the digitization process. Note, color corrections are applied to imagery in (A) and (B) but have been removed in (C) to highlight the embedded image. Figure adapted from Fox et al. (2019).

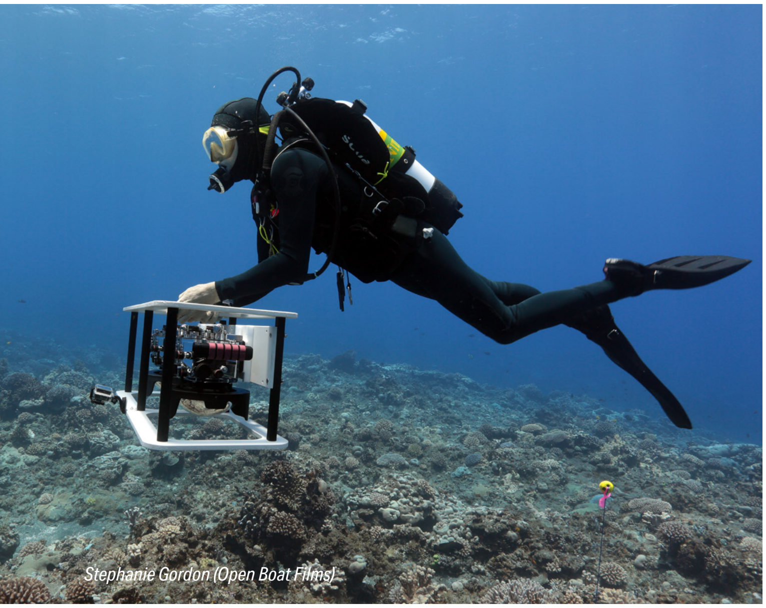

such as assessments of colony condition (e.g., bleaching status or presence of disease) or recruit surveys (see Fox et al., 2019; and Part II, section 7.6 of this report). Finally, conducting this work in a virtual environment provides the potential to accelerate workflows with artificial intelligence (AI) tools (Beijbom et al., 2015, Ditria et al., 2022, Pavoni et al., 2022, Runyan et al., 2022).

2.3.3. Demographic data

LAI can be used to generate comprehensive size structure data for corals and other benthic organisms. The demographic approach allows identification of ecological differences such as the presence of larger and faster growing, or more resilient, individual coral genotypes or temporal shifts in recruitment or mortality rates. In particular, when LAI is replicated across time at permanent locations on the reef, it can facilitate the generation of demographic data critical to evaluation of population trends in both natural and restored habitats (Hernández-Landa et al., 2020; Kodera et al., 2020; Sandin et al., 2020; Ferrari et al., 2021; Rodriguez et al., 2021). For the coral restoration community in particular, data on the growth and survivorship of individual coral colonies (and genets) is of fundamental importance to tracking and maximizing intervention success (Drury et al., 2017; Goergen et al., 2020). Evaluating context-dependent changes in colonies, populations, and communities requires an approach to track many individual coral colonies though time and extract accurate and precise measurements of coral size. Current efforts to track coral growth require in situ colony measurements such as total linear extension or colony dimensions to approximate ellipsoid area (Huntington and Miller, 2014; Lirman et al., 2014; Pratchett et al., 2015; CRCP 2022b). These data are difficult to collect once, and finding and remeasuring individual colonies over time can be particularly time intensive. Therefore, diver-based collection of individual colony measurements is not scalable to meet the objectives of restoration or monitoring efforts for large areas or with robust replication (Goergen et al., 2020; Couch et al., 2021).

To generate demographic data with LAI, corals are digitally enumerated, classified, and measured using image analysis software through a process known as digitization, also referred to as instance, or semantic segmentation in computer science fields (Garcia-Garcia et al., 2017; Pavoni et al., 2020; Pierce et al., 2021). In practice, this process is carried out by tracing colony boundaries using a mouse or digitizing tablet and then providing class labels (e.g., taxonomic designation). Additionally, existing in situ approaches to measure coral growth, including widely used metrics such as linear extension (Pratchett et al., 2015), total linear extension (Johnson et al., 2011), and maximum diameter or ellipsoid area (Kayal et al., 2015), can be replicated using LAI. A variety of commercial software packages currently available have been used to conduct this work (Edwards et al., 2017; Sandin et al., 2020; Couch et al., 2021; Million et al., 2021; Rodriguez et al., 2021; Urbina-Barreto et al., 2021), and specific workflow examples are provided in Part II of this report. An additional advantage of the LAI approach for generating such information derives from the collection of overlapping images to create LAI, allowing a given coral colony to be viewed in multiple images and angles. These additional views enable more detailed determinations of colony condition (Figure 14), such as whether live tissue connects adjacent patches of a colony or whether disease is present. In addition, the raw data used to make these decisions can be reviewed and reanalyzed as needed, rather than relying on non-repeatable determinations made in the field.

2.3.4. Coregistration

Generation of demographic and other time series data from LAI is made easier when image collections are repeated at the same location through time and then spatially aligned in a process known as coregistration. Coregistration refers to the process of placing two different datasets into the same coordinate system. LAI, whether 2D imagery or 3D models, can be coregistered by matching a minimum of 3 fixed (non-collinear) positions, often referred to as control points, that are identified in both datasets and used to match the images to each other (Figure 15). When georeferenced datasets use the same projected or geographic coordinate system, they are already effectively coregistered and can be viewed as overlapping layers depending on the choice of visual–analytical software being used. However, in some cases, errors in georeferencing might lead to small offsets in the coregistration that must be corrected. Numerous software platforms are available to coregister images and other datasets, including widely used geospatial software such as the commercially available ArcGIS (Esri, USA) and open-source QGIS (https://qgis.org/en/site/). Similarly, coregistration of temporally replicated LAI can be accomplished with or without georeferenced information using open-source programs such as TagLab (https://github.com/cnr-isti-vclab/TagLab) and Cloud Compare (http://www.cloudcompare.org), or the custom software Viscore (Petrovic et al., 2014; Sandin et al., 2020).

Coregistration refers to the process of aligning layers of an image or other spatial data that contain overlapping information. Image layers can be coregistered by visually matching features, while geospatial layers can be coregistered using geographic position information. Coregistration of overlapping layers can aid temporal data comparisons.

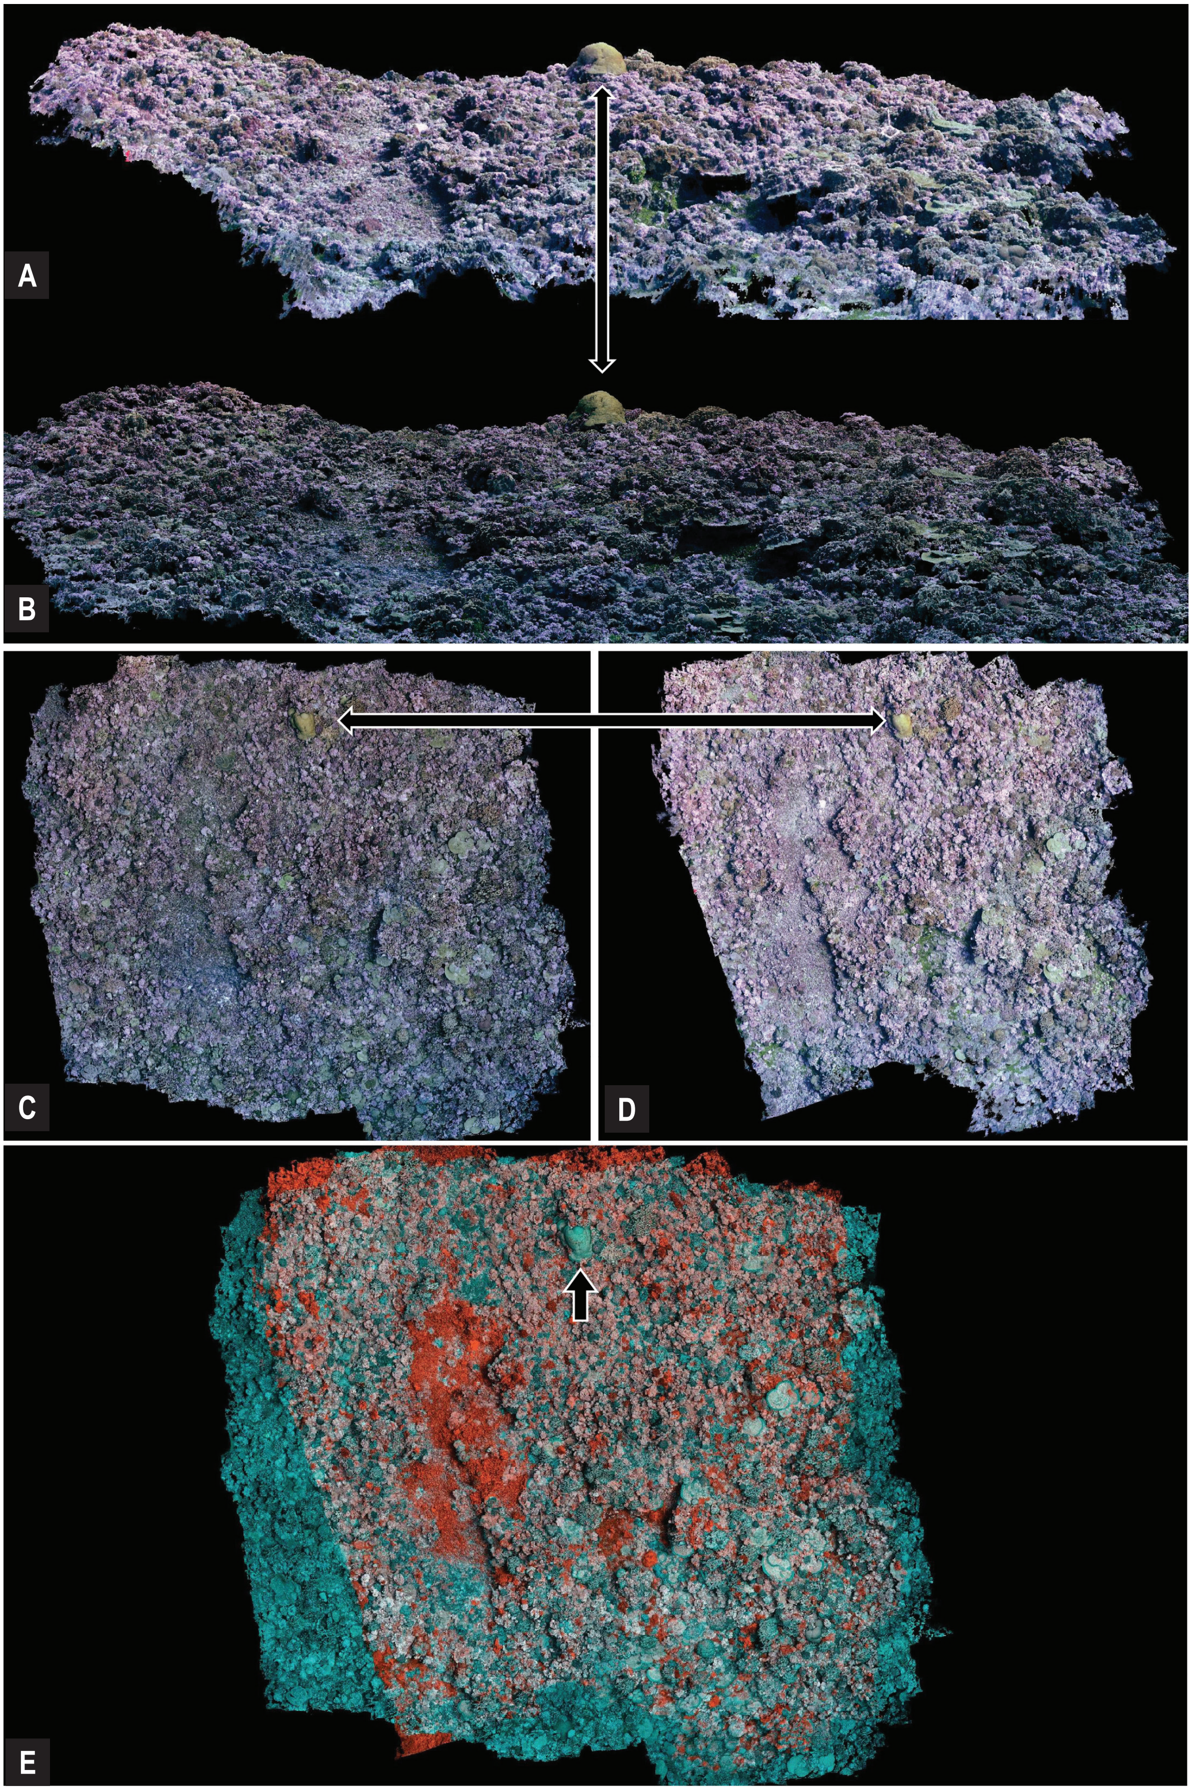

Figure 15. Example of coregistration and colony tracking. Conceptual example of coregistration. Using matched features (example shown with arrows), the models can be translated and rotated vertically, (A) and (B), or horizontally, (C) and (D), as needed for successful coregistration (E). A single feature is shown with the arrows here, but most software platforms require a minimum of three matched features to complete coregistration. In some cases, coregistration may also require rescaling in order to properly align spatially overlapping models. (E) shows the resulting coregistration of the models from (A) and (B), with red areas showing portions of the plot that were lost to erosion and other physical processes between the two surveys, or areas where the two models do not overlap.

Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, UC San Diego)

Overlapping coregistered LAI allows for straightforward digital tagging and tracking of individual coral colonies, preventing the need to relocate them underwater. The unavoidable difficulty of relocating colonies underwater limits the total number of colonies that can be tracked and therefore hinders achieving statistical robustness and ecological representativeness. However, due to the dynamic nature of underwater habitats and complexity of coral growth, along with error inherent in LAI, automated extraction of time series data is discouraged, particularly with respect to measurements of colony-level growth or partial mortality. In such cases, care must be taken to measure overall changes in colony area, as coregistration error prevents fully spatially explicit comparisons. Further, current tools for coregistration require some degree of manual intervention, limiting the scalability of time series investigation. However, as the LAI approach continues to mature, increased opportunities for automation and increases in precision of coregistration will undoubtedly become available. Regardless of the specifics of the approach, coregistration of time series image data is a powerful strength of the LAI approach. Examples of workflows to conduct coregistration or colony tracking are provided in Part II of this guide.

2.3.5. Colony measurements in 3D

To date, depictions of coral demography have largely been based on 2D estimates of surface area, whether using measurements collected in the field or via image-based methods, to measure colony size and survivorship through time (Hughes, 1984; Bak and Nieuwland, 1995; Edmunds 2015; Pratchett et al., 2015; Couch et al., 2021). Corals, however, clearly grow along multiple axes, and while numerous studies have taken this into account, they have largely been operationally limited to lower sample sizes, which preclude robust descriptions of demographic patterns (Pratchett et al., 2015). Ideally, comprehensive data should be collected to reflect the entirety and complexity of coral habitats, and there is clear motivation to move toward 3D descriptions of change on reefs, particularly with respect to coral growth. Comparisons of LAI-derived 3D estimates to existing laboratory and in situ methods have shown broad support that LAI can produce equivalent, or more robust, data relative to existing approaches (Ferrari et al., 2022; Curtis et al., 2023). Having such information in hand allows improvements to be made to estimates that rely on it, such as calcification budgets (Perry et al., 2012; Lange et al., 2022), and enables novel descriptions of biological and ecological processes (Burns et al., 2015; Zawada et al., 2019). Studies have also shown that the LAI approach will be particularly useful to replicate or replace existing techniques used to measure growth of branching corals, such as *Acropora cervicornis*, a commonly used restoration species in the Caribbean, allowing data to be collected at previously unavailable spatial scales (Million et al., 2021). Approaches to extract 3D measurements represent a rapidly developing segment of the LAI approach and include widely used programs such as Meshlab (Cignoni et al., 2008; Aston et al., 2022), along with more recently available platforms such as CloudCompare (Lange et al., 2022). For guidance on the extraction of 3D measurements, please see Part II, section 1.6.

This guide largely focuses on an approach for image collection at a spatial scale that enables robust depictions of reef state, whether through descriptions of physical structure or percent cover and coral abundance. However, this approach is currently not well suited to obtain the imagery needed to extract 3D metrics such as surface area and volumetric data at scale. First, the raw imagery needed to generate complete 3D surfaces on the undersides of colonies is time consuming and difficult to obtain (Figure 7), particularly at the spatial scale desired for most ecological investigations. Next, important practical limitations arise with respect to the use of LAI to extract 3D metrics. Even when image collection is designed to minimize occlusions, portions of the 3D model will invariably be incomplete due to the complexity of the reef surface. When 3D meshes are built, these poorly reconstructed areas will require interpolation, preventing accurate measurements of volume in these regions. Important decisions must also be made regarding what volume of the model should be measured. Unless a coral has been tracked from its initial settlement, there is no reliable method to bound a given

coral colony for measurement of volume. In other words, while it is possible to determine where a coral colony ends, it is difficult to correctly measure where it starts. While this can prevent measuring total colony volume, measurements of change can still be conducted. For more targeted objectives, it is certainly possible to image individual coral colonies comprehensively with opportunities to manage decisions related to the meshing process or how to bound volume measurements. However, collecting 3D measurements of corals from LAI remains a labor-intensive workflow, available in only a handful of software programs, and accomplishing this work at the scale of most monitoring or restoration programs is not currently tractable. In the future, as the LAI approach continues to progress, the opportunity to comprehensively describe change in coral communities in 3D will become increasingly possible.

2.3.6. Structural metrics

Despite the challenges of working directly with 3D products, there are ample opportunities to describe the 3D structure of LAI. Some of the most common applications of LAI are descriptions of the structural complexity using DEMs (Figure 11), which offer a 2.5D representation of the bathymetry of the imaged area. From DEMs, various metrics can be automatically computed, such as surface complexity, slope, or roughness, which can then be related to other biological variables, including coral percent cover or fish and invertebrate biomass from in situ surveys (González-Rivero et al., 2017; Helder et al., 2022; Swanborn et al., 2022). More recently, authors have used DEMs together with digitized orthoprojections (when both have been generated using the same projection of the 3D model) to investigate how particular coral morphologies drive reef complexity or how a loss of live coral cover results in reductions in complexity (Burns et al., 2019; Fukunaga et al., 2022b; McCarthy et al., 2022; Carlot et al., 2023). As described in Part I, section 2.2.4, there are caveats associated with DEM generation that must be carefully considered as they can obscure important details of fine-scale 3D structure. Workflows also exist for extracting structural information directly from the point cloud itself without creating a DEM. One example of this approach is to use a digital approximation of a profile gauge, through which virtual "poles" sample the depth of the point cloud along a virtual transect (McCarthy et al., 2022). The depth at which the virtual poles intersect the point cloud can be used to calculate linear rugosity at a set level of resolution, which can serve as a direct analogue to traditional in situ approaches.

Regardless of the specific approach for extracting structural information, LAI offers substantial improvements over historical approaches to describe reef structure that were largely limited to large-scale bathymetric surveys, which tend to provide highly detailed, but coarse (>1-m pixel spacing) data that are not amenable to many questions of ecological interest. In situ methods (e.g., chain and tape) to describe reef complexity, on the other hand, can provide high-resolution data but are extremely difficult to scale spatially or to higher levels of replication. For more information on approaches used to extract structural information from LAI, please see Part II, sections 1.6, 5.6 and 7.6 of this guide.

2.3.7. Optimizing data extraction

Transferring the bulk of ecological data extraction from the field to the laboratory enables an increase in the detail and volume of data that can be collected. However, this presents an important challenge and frequent misconception regarding data extraction from LAI: it is more time consuming than comparable approaches in the field. For instance, percent cover estimates can be generated nearly instantaneously using in-water line-point intercept approaches (yet field data sheets must still be transcribed and data quality checked). Similarly, while collecting demographic data in situ is limited in scale by the effort needed to measure corals underwater, the data are available for analysis as soon as data sheets are transcribed and quality checked. On the other hand, using the LAI approach requires a significant time investment before data extraction can begin (e.g., for 3D model and orthophotomosaic generation and use of computer software). However, once LAI is available, experienced users can analyze thousands of sample points in less than a day to generate robust estimates of percent cover. Further, extracting demographic data from LAI simply involves a user tracing over the image with a mouse or digitizing tablet. The specific measurements are then done by the software, allowing hundreds of corals measurements to be collected by a single observer, dramatically outpacing what can be generated by dive teams collecting measurements underwater. Other approaches to collecting data, such as that needed for spatial analyses, are sufficiently difficult to conduct underwater that they have seldom been attempted, regardless of levels of replication (though see: Lewis, 1970; Dana, 1976; Carlon and Olson, 1993; Karlson et al., 2007; Marhaver et al., 2013; Deignan and Pawlik, 2015). However, with the LAI approach, once colonies have been mapped, the measurements needed for spatial analysis can be collected almost instantaneously. Similarly, many structural complexity analyses can be conducted immediately after models have been constructed. Indeed, the process can become considerably more time consuming than the in situ approach when there are no realistic limits to the amount of data that are extracted. However, for well-designed studies, what emerges is that the LAI approach represents a fundamental step forward in the ability of expert ecologists to obtain data and ask questions about coral reef function in ways that were previously unfeasible due to logistical limitations.

Whether for purposes of evaluating coral reef ecosystems or restorations, a major goal for ecological analysis is to maximize the efficiency of information collection in the context of project objectives, meaningful summary metrics, and needed sample sizes. The time needed to create LAI products from raw imagery represents an unavoidable increase in the time needed to generate data relative to historical in-water workflows; however, these increases are exacerbated by unconstrained data extraction. Data efforts should thus be limited to the sample sizes needed to answer stated objectives. In the future, however, AI approaches hold the promise of dramatically accelerating many of the data extraction activities discussed here (Figure 16). Currently, most applications of AI first require human-generated data to “train” the system, and there similarly must be “validation data” available to check the performance of the algorithm. The principal advantage of well-trained AI algorithms is in the dramatic increases in speed over human-driven efforts, yet they lack the precision of well-trained human experts (Beijbom et al., 2015; Pavoni et al., 2022). In particular, humans are able to achieve a much higher level of detail with taxonomic classifications, particularly for rare classes, or during segmentation of complex colony boundaries (Pavoni et al., 2020). Machine-assisted tools enable the time-consuming step of segmentation to be handled by the computer, allowing the more detailed but less time-intensive steps of corrections to the computer-generated predictions to be overseen by expert human observers. Currently, and likely in the near future, the best overall performance will arise from human-supervised, machine-accelerated workflows such as that offered by the program TagLab (Figure 17; Pavoni et al., 2022). With TagLab, the expert observer remains in the loop, while AI algorithms are easily interchangeable, allowing for refinement and testing of new algorithms as they become available, such as the popular Segment Anything Model (Kirillov et al., 2023). Moreover, most applications of AI are still in development with respect to their use in the field of coral reef science, and given the current lack of readily available public training data, most groups should plan for investment in human-driven data extraction workflows.

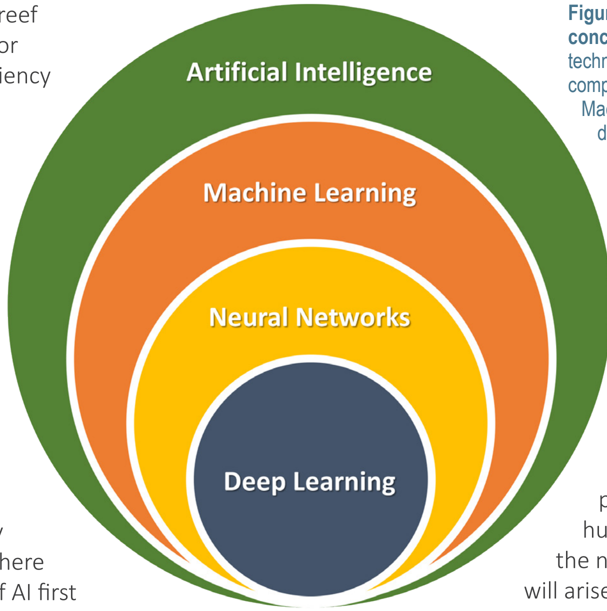

Figure 16. Artificial Intelligence (AI) terminology conceptual diagram. AI encompasses a range of techniques that utilize computer algorithms to perform complex tasks, such as identifying corals in reef imagery. Machine learning (ML) is a subfield of AI, wherein historical data (often referred to as training data) are employed to construct predictive models. Many of the more recent and successful ML techniques utilize increasingly complex neural networks such as deep learning.

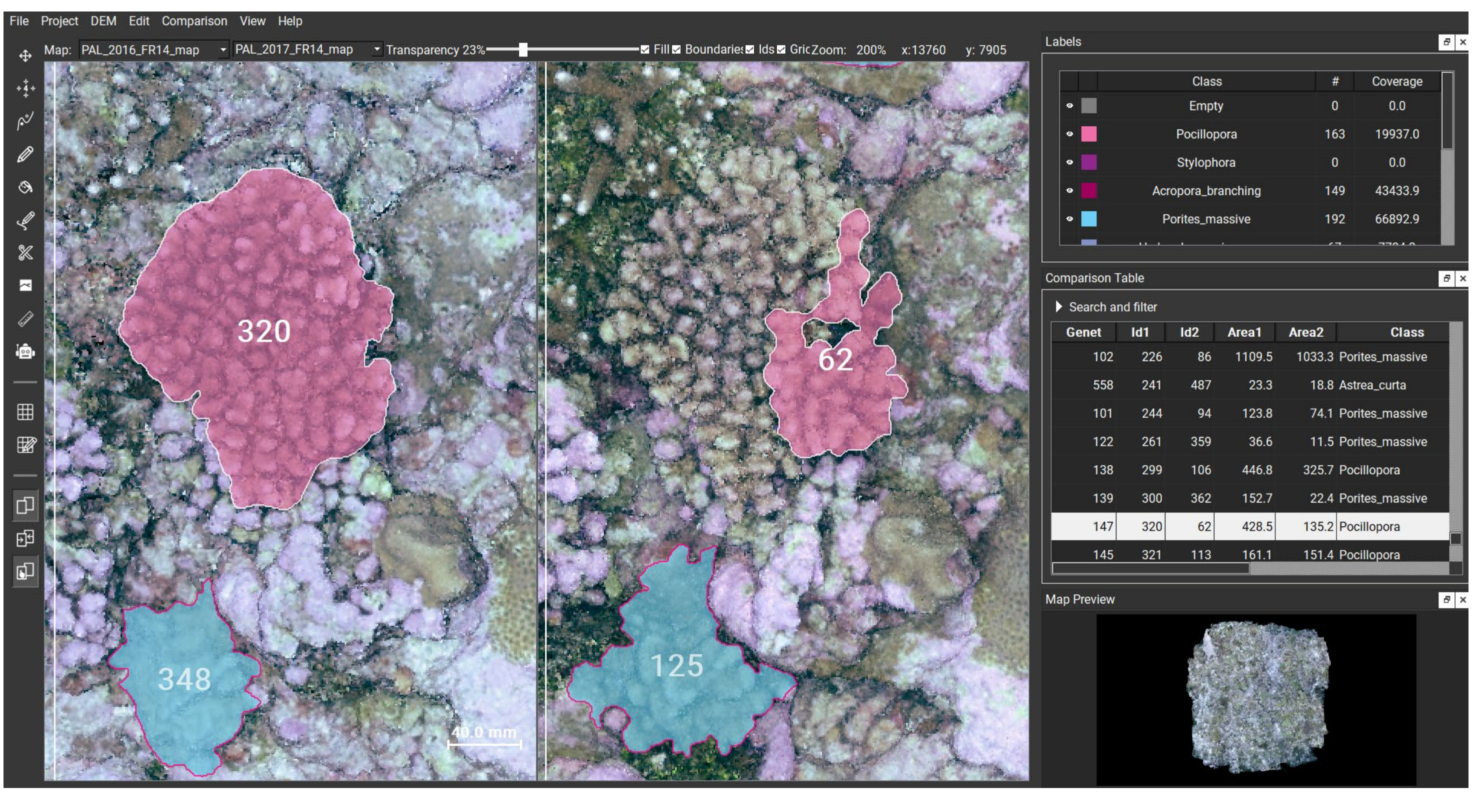

Figure 17. TagLab user interface showing a coregistered time series with segmented coral colonies. A segmented colony (ID 320) is shown in 2016 on the left and again in 2017 on the right (ID 62). Colonies featuring a minimum of 50% overlap between time points can be matched using Taglab’s automatic matching feature or manually as needed. Colony size is shown each year for the matched colonies in the bottom left, showing the highlighted colony decreased in size from 428.5 cm2 to 135.2 cm2 from 2016 to 2017.


Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, UC San Diego)

2.4. Data curation and access

2.4.1. Overview

Data curation is presented as the last step in the LAI pipeline but operationally should occur before, during, and after each of the previous steps for any successful long-term LAI based project. Data curation responsibilities fall into two major categories: 1- decisions regarding organizational structure and 2- decisions regarding hardware infrastructure and architecture. The volumes of data generated for even a single instance of LAI are substantial, both in terms of the number of files and file types generated, as well as the cumulative disk space needed to securely store these products. For instance, a standard unit of sampling for species- or genus-level analysis with LAI is a 100-m^2 plot, typically requiring the collection of anywhere from 2,000–8,000 images, depending on the particulars of the imaging approach (see Edwards et al., 2017; Couch et al., 2021; and Part II of this report). When using a high-resolution camera and collecting images as compressed JPEGs, a single survey comprises anywhere from 20–80 GB of raw imagery. If imagery is collected in a raw image format (e.g., CR2, RAW, SRF. or NEF), file storage requirements can be 2–3 times larger (see Part I, section 3.4). Depending on the settings used during model construction, particularly during any optimization steps or during DPC generation, the Metashape project can be of the same file size as the collection of image files, though derived products (e.g., orthoprojections or DEMs) will be much smaller. Metashape provides an increasing number of analytical tools, but if another program is used for data extraction, the 3D model must be exported, and the file size of the exported point clouds can be of equivalent size to the Metashape project, while meshed models tend to be much smaller. Even for models consisting of relatively few photos (approximately 300), the combined data products can require up to 10 GB of storage space. All told, a single instance of LAI will be represented by many gigabytes of data and multiple file types that need to be consistently named and cross referenced. Further, for temporal and spatial comparisons, considerable portions of these data need to be readily searchable and accessible. The process of managing and serving these data should, however, not be perceived as a barrier to adoption of the LAI approach. With strong organizational structure, and even moderately equipped infrastructure, data can be made accessible and secure well beyond immediate project needs.

2.4.2. Data management best practices

The greatest challenges to the effective use of the LAI approach are related to longstanding best practices regarding data organization. Ideally, these organizational tasks begin before leaving the laboratory, starting with decisions regarding the collection of metadata, and in particular, decisions regarding naming conventions. Well-structured metadata are essential for maintaining efficient analytical workflows and achieving project objectives and are reliant on consistent and thorough naming conventions. Whether for metadata fields or file and folder names, while brevity is encouraged, it is critical to prioritize names that are intuitive, human readable, sortable, and non-repetitive. When codes or abbreviations are used in place of full names, they should be intuitive, with readily accessible definitions. To the degree possible, metadata cataloged in the field on physical data sheets or digitally in the laboratory should follow the same naming conventions. Moreover, once the process of LAI construction has begun, not to mention later steps of ecological data extraction, renaming multitudes of files to be consistent across data collections, and updating any metadata that catalogs these names, is not only burdensome but a potential and likely point of failure.

After naming conventions have been decided, the next challenge is to develop a well-structured organizational system for file storage. Even for those institutions with enterprise-level database infrastructures, the first, and often final, destination for most data products tends to be folders stored on personal computers. To the degree possible, it is recommended to maintain a simple and clearly labeled folder structure that facilitates discoverability for later stages of organization onto other more sophisticated platforms. For instance, overreliance on the use of subfolders for organization can reduce discoverability and, in some cases, might result in file path character limits being exceeded. Most importantly, the tendency to fall into the trap of working too quickly must be avoided. The mere seconds saved by not writing out full names is trivial compared to the hours needed to disambiguate nonintuitive abbreviations. Further, best practices can be encouraged by maintaining rigorous quality assurance and quality control (QAQC) standards throughout data curation activities. Ultimately, by adhering to consistent and straightforward data structuring practices, the multitude of data streams, including raw imagery, metadata, and summary data products, can be easily and rapidly accessed, strengthening the ability to meet project requirements.

2.4.3. Infrastructure

With the basics of data management considered, the next challenge of data curation concerns obtaining the physical infrastructure for file storage and transfer. As noted earlier, the data volumes associated with LAI can be considerable, and storage infrastructure should be designed to accommodate not only the magnitude of storage that will be needed, but also the rate at which that data can be ingested, accessed, and disseminated. This process begins with the seemingly simple task of transferring imagery collected in the field from cameras to initial storage locations. If a limited number of total images are collected, data transfer and organization can be relatively painless, but when multiple cameras are used to collect many thousands of images across several sites a day, transfers can last overnight, creating backlogs that represent a fundamental challenge to good data management practices, particularly when working in field environments. In the lab, the challenge is to create a data storage infrastructure that can scale to accommodate project needs. In some cases, simple solutions such as using high-speed memory cards in cameras and high-speed, high-capacity hard drives with high-speed connectivity will suffice to meet project needs. A particular challenge of the LAI approach is the large files sizes of some of the products like orthoprojections, which can frequently exceed 1.5 GB, or other products that include large numbers of small files. In both cases, files can be burdensomely slow to open or transfer on certain database architectures, including any system that is supported by standard mechanical hard disk drives (HDDs). One increasingly available solution is to use solid-state drives (SSDs), which have dropped dramatically in price in recent years and are over an order of magnitude faster than HDDs. SSDs can offer transfer up to 3 GB/s, a substantial increase from the 0.1 GB/s available on standard HDDs. However, the realized speed can vary substantially based on drive type and connectivity, and it is important to pair SSDs with other high-end peripherals, such as fast camera memory cards (150 MB/s read/write) and USB 3.0 or higher connectivity (e.g., Thunderbolt, USB-C). While the financial investments to update to the most modern hardware is not trivial, the savings in labor cost associated with this faster hardware far exceeds any savings gained from purchasing less expensive but slower hardware.

2.4.4. Data accessibility

Those groups needing to provide multiple users access to data, or with high levels of site replication or long time series, will require storage systems designed to allow significant portions of the total data to be simultaneously available. In such cases, collections of individual disk drives may not be up to the task of securely or efficiently serving such volumes of data. Some institutions may have access to long-term archival databases with ample storage capacity, which often have the added benefit of rigorous and transferable metadata standards. Unfortunately, the architecture of these systems tends not to be amenable to the ready access needed by groups with active programs. Consumer-grade network attached storage (NAS) systems are an increasingly available and popular solution to this challenge. These systems can be designed to accommodate tens of terabytes of data, and usually feature a redundant array of independent disks (RAID) storage to ensure data integrity in the case of partial device failure. NAS systems are available with SSD storage, though these configurations can be costly. Alternatively, hybrid systems are available that combine SSD caching with HDD storage as a more economical solution. The speed of the local network that a NAS is mounted on can be a rate-limiting step; however, many facilities are increasingly equipped with 1 GB/s ethernet, which is sufficient for most activities. If high-speed internet is not available, NAS systems can be used over a local ethernet or as standalone drives connected via USB, making them essentially high-capacity disk drives. An advantage of NAS systems is that they can be configured for remote access

Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, Uc San Diego)

to make data more broadly available, though considerations of local internet bandwidth will again be a rate-limiting step to data accessibility. Regardless of the architecture that is ultimately chosen, it is critical to ensure the safety of data products, whether via RAID or the maintenance of fully redundant and geographically distributed copies.

Considerations of data transfer continue throughout the LAI pipeline and can pose a particular challenge when transferring raw imagery and LAI products between colleagues. Data repositories such as DataOne, NOAA's National Centers for Environmental Information (NCEI), and NSF's Biological and Chemical Oceanography Data Management Office (BCO-DMO) are well-known solutions to making data widely available as they use best practices for metadata structure and usually have embedded file transfer protocols designed for efficient data transfer. However, access to data in these repositories is still limited as architecture tends to be optimized for archival activities rather than daily access, and data transfer will still be limited by local internet bandwidth. Further, the requirements for data ingestion into these repositories are strict and tend to be optimized for completed rather than ongoing or early-stage efforts. Cloud computing and storage platforms such as Google Drive or Amazon Web Services have become increasingly available and offer tiered data-storage plans with nearly unlimited storage. While these platforms offer more flexibility in data structure, data transfer remains limited by local internet bandwidth, and the most economical archival storage tiers tend to not be optimized for regular access.

An increasingly popular option for data access and sharing is the use of remote workspaces, such as Microsoft Azure Virtual Desktop (AVD) or Amazon WorkSpaces. With these platforms, users can remotely access high-speed computing infrastructure with access to customizable processing and storage architecture supported by considerable resources. Geographically distributed users can thus work from the same file system, precluding the need for large file transfers. Internet speed tends to be less limiting for remote computing platforms, as they tend to provide access to the remote computer via video streaming (10 Mbps for HD; 25 Mbps for 4K). Data upload and download to local hardware are still limited by local internet speeds, though some platforms provide mail-in services that allow data to be sent directly to and from the storage system for rapid data transfer and are limited only by transportation times. Regardless of the storage or transfer system that is ultimately used, any successful LAI project will have dedicated and well-planned investments for infrastructure design and maintenance. As with most aspects of the LAI approach, available technologies are rapidly changing, and users are encouraged to stay up to date on the latest developments.

2.4.5. Data standards and comparability

Ultimately, there is importance in disseminating raw LAI data products and analytical tools to the broader marine science community (McCarthy, 2023). To the degree possible, emphasis should be on ensuring that data are publicly accessible and well structured so that they are both discoverable and sharable with anyone interested in these important shared global resources. As mentioned previously, LAI can be created with very few constraints. However, to conduct ecological inquiry and to facilitate cross-study comparisons, a minimum set of data collection and reporting standards must be met (Ferrari, 2022). The first standard to be met applies to any set of ecological data and requires that imagery is collected from a known geographic location. The technical specifications of any hardware or software used to gather additional geographic positioning information should be reported. For instance, when using a USBL to georeference individual images or a high-precision surface GPS to mark the location of GCPs, the specifications of the device should be reported, as well as any post hoc corrections to the data that were necessary. Next, details of the image collection approach should be reported including the specification of the camera and settings that were used, as well as the lens type and the distance from the benthos at which the camera was operated. The general survey pattern should also be reported, including the spatial extent of the area that was imaged, the elapsed survey interval, and the number of images collected.

There must also be clear reporting standards associated with image collection and the settings used during the generation of the 3D model and other associated LAI products. The selection of settings during sparse and dense point cloud reconstruction should be reported, as well as any optimization procedures that were conducted. Information must also be provided regarding the settings used during the generation of meshed 3D models, DEMs, orthomosaics or any other derived products, as well as the export resolution of these products. For DEMs and orthomosaics in particular, it is important to report how the model was oriented, as this can dramatically affect the reported depth values as well as the projected area of items in later steps of colony segmentation. Reporting of SfM-based reconstruction accuracy and precision is difficult without ground-truth comparisons. However, information such as the proportion of photos that were successfully aligned and the error between scale bars can provide at least some understanding of the reconstruction quality. Once LAI products have been prepared for ecological analysis, standard scientific guidelines should inform reporting standards. Importantly, broad adoption of data-reporting standards facilitates the aggregation of data for large-scale data syntheses or optimization of AI-driven data extraction. As with many other aspects of the LAI workflow, reporting standards are rapidly evolving, and any group implementing the approach is advised to stay up to date with recent developments.

Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, Uc San Diego)

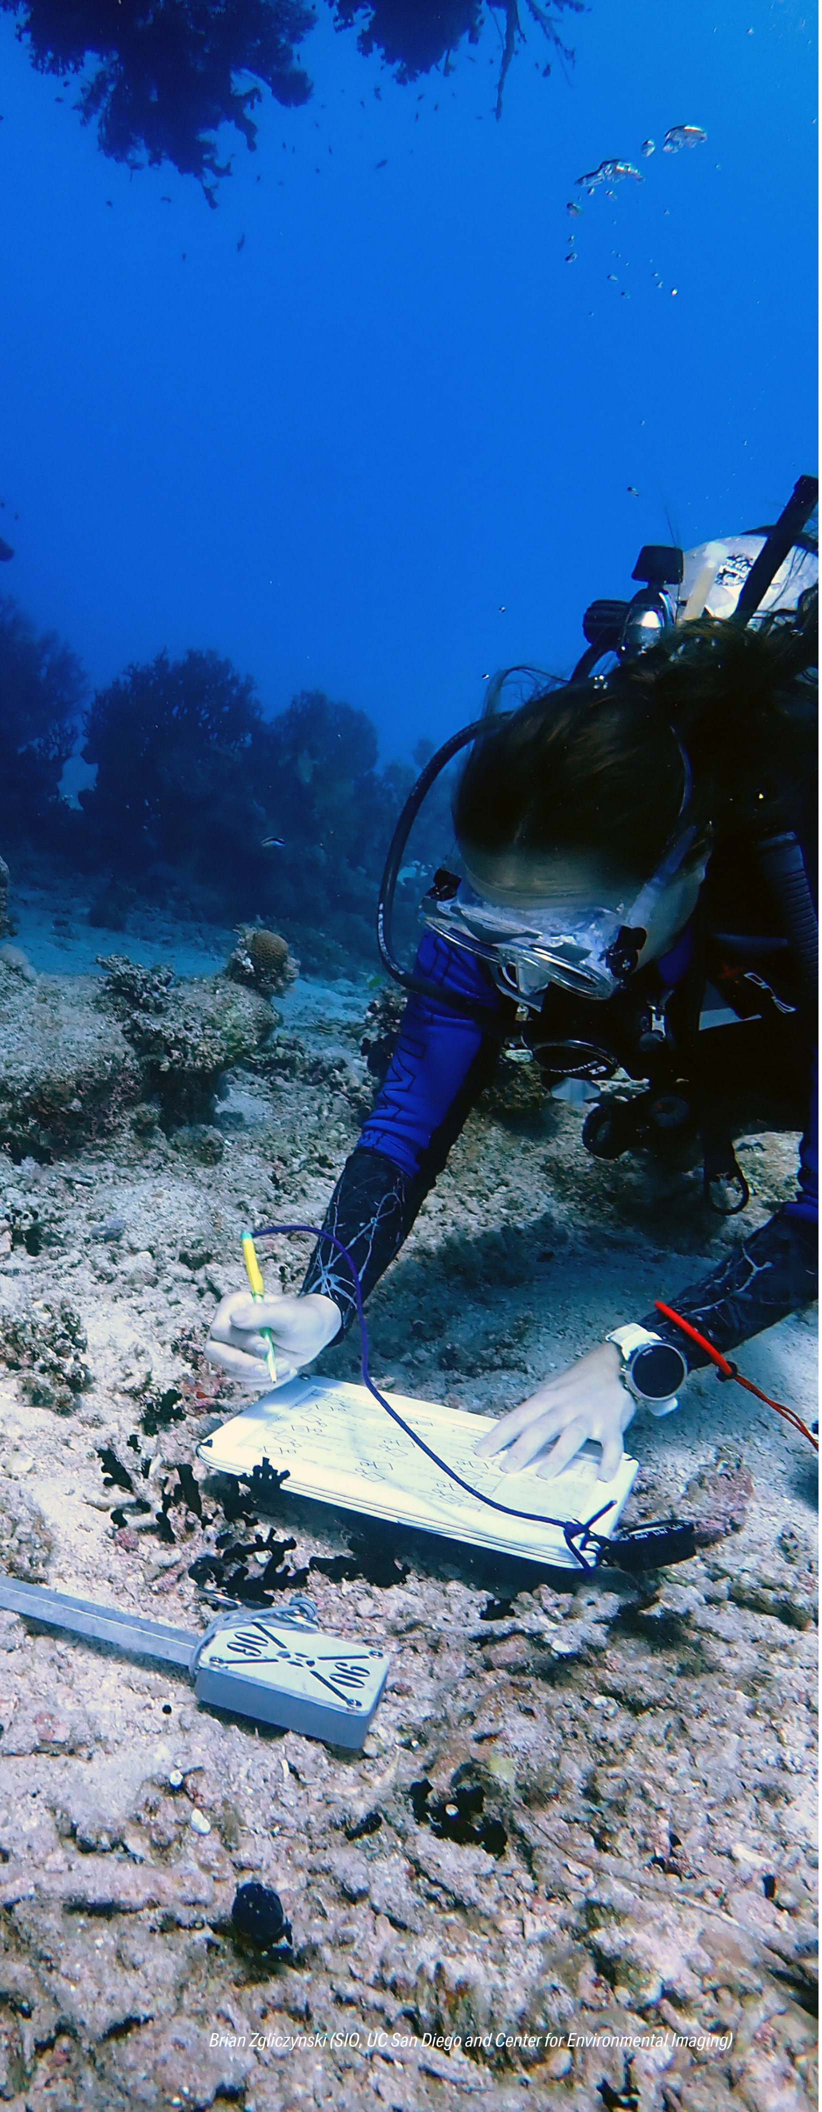

Brian Zgliczynski (SIO, UC San Diego and Center for Environmental Imaging)

Chapter 3
Planning and Implementing LAI

3.1. Overview: putting together the pieces of the pipeline

When designing an LAI workflow, it is important to consider both the practical constraints of the approach along each step of the pipeline and the investments that will be needed at each of these steps (Figure 18). Before any decisions regarding implementation can be made, the project goals and the data that will be needed to achieve these objectives should be clearly articulated. Next, the required spatial and temporal extent and resolution of data should be determined, including levels of replication, in order to answer specific ecological questions. These decisions are critical to defining the capacity that will be needed for both image collection in the field and model processing, as more detailed imagery and greater spatial extents generally require more extensive survey planning, sophisticated equipment, and time. Further, collecting greater volumes of raw imagery takes longer not only to collect in the field but also to process into LAI and will require more computational resources to do so in a timely manner. Similarly, while some ecological metrics can be generated almost immediately after LAI products are available, others will require significant investments of time and human effort to compile. Ultimately, decisions will have to be made along each step of the pipeline, and the challenge is to balance available resources with project goals and timelines to optimize distribution of resources along the LAI pipeline and minimize any rate-limiting bottlenecks.

3.2. Planning for image detail

The most important questions to address when designing a specific LAI workflow are: 1 – What level of detail is needed in imagery? 2 – What spatial extent will imagery be collected over? and 3 – What levels of replication and frequency of image collection are needed? Decisions regarding needed image detail should be made with respect to the ecological objectives. The resolutions of LAI products, and thus the data that can be extracted from them, are limited by the detail of the raw imagery. Similarly, the level of detail that is needed in the raw imagery should be based on the granularity of the data needed for analysis. For instance, obtaining species-level taxonomy for corals will require much more detailed imagery than that needed to obtain genus-level information. For cryptic organisms, such as algae or sponges, even more detailed imagery will be required. Analyses of colony-level spatial patterns often require only the position of coral colony centroids, which can be accurately derived from relatively low-detail LAI alone (Zvuloni et al., 2009; Edmunds et al., 2018; Price et al., 2021). Alternatively, others such as species-level interactions might require explicit descriptions of colony shape, which will in turn require much more detailed LAI products and raw imagery (Burns et al., 2016a; Edwards et al., 2017; George et al., 2021; Corso et al., 2022). Investigations of structure can also vary widely in the level of needed detail. At one end of the

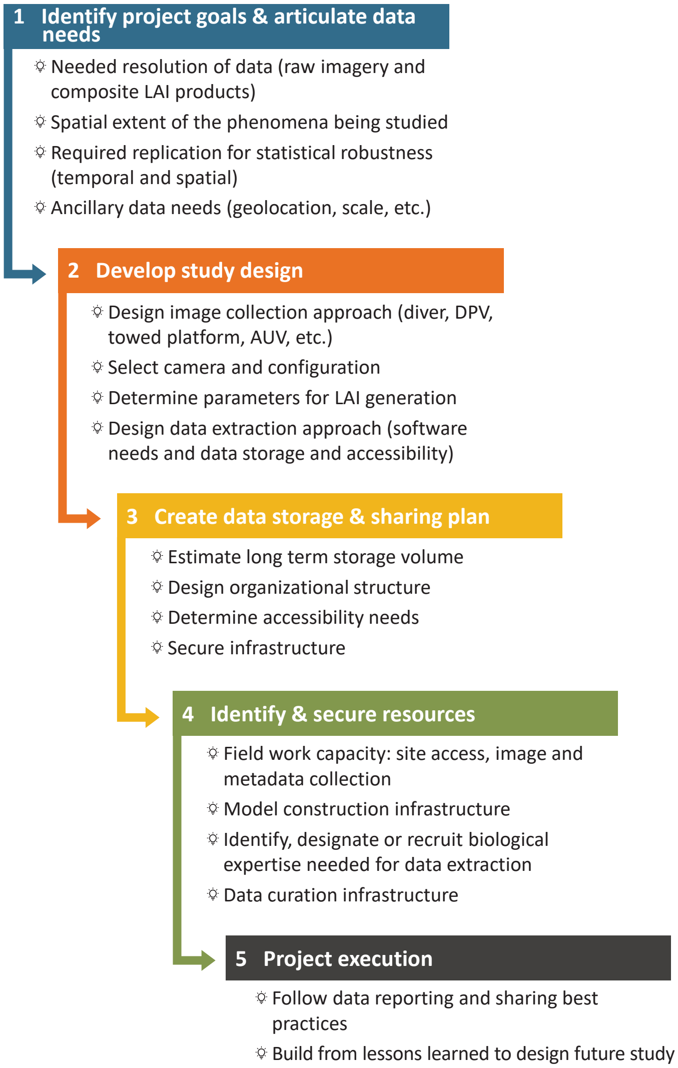

Figure 18. Conceptual diagram for planning and implementing LAI. While the LAI pipeline itself is not sequential, a sequence of interdependent decisions must be made in order to construct a functioning pipeline. LAI = large-area imaging; DPV = diver propulsion vehicle; AUV = autonomous underwater vehicle.

spectrum, the mean depth or slope of a plot might be all that is needed, which can be derived from relatively low-detail LAI. On the other hand, those interested in describing differences in surface complexity as a function of different coral morphologies will require much more detailed LAI.

Importantly, more detailed imagery, such as that needed to examine the surface complexity of individual coral heads, must be collected closer to the bottom and will cover a limited footprint of the bottom (Table 1). Different combinations of cameras and lenses can be used to image equivalent areas of the benthos. For example, when operated at the same location the same scene can be photographed with either a camera with an APS-C sensors and an 18-mm focal length lens, one with a 1/2.3 sensor (6.17 mm × 4.55 mm) and a 4.5-mm focal length lens, or a full frame camera with a 24-mm lens will produce images covering roughly the same spatial extent when operated at the same distance from the bottom. Generally, however, higher-end cameras will collect more detailed imagery at the same distance from the benthos. Further, regardless of the camera type used, collecting more detailed imagery within increasingly large plots will increase the complexity of the LAI approach at all stages in the pipeline.

3.3. Planning for spatial extent

Once the desired resolution has been determined, the next key decision regards the necessary spatial extent (e.g., the plot size) over which images are to be collected, as well as the number of replicate plots that will be needed. These choices should be based on the ecological context of the study system and the amount of data that are needed to capture the desired ecological signal with statistical robustness. For instance, to obtain estimates of coral population demographics, surveys must be designed to capture a representative sample of corals. In some settings, a single 10 m × 10 m plot might capture hundreds of corals; however, in locations where corals are less abundant, several such plots would be required to obtain an equivalent sample size (Figure 19). Determining plot size is particularly important for restoration activities, as there is an explicit expectation that the spatial footprint will change over time. The initial location of coral restoration (e.g., outplanting corals) should be included, with the plots ideally designed to accommodate the growth of corals as well as any expected movement (e.g., due to dislodgement or fragmentation) over the desired course of study (Goergen and Gilliam, 2018; Goergen et al., 2020). As long as the core area is retained, plot size can easily be increased over time. However, implementing a larger area at the outset will allow insights into how restoration activities alter the surrounding area, which is critical to evaluating restoration success at the ecosystem level (Ladd et al., 2018; Goergen et al., 2020). Finally, because the accuracy of 3D reconstructions relies on overlap in the raw imagery, it is important to plan for a buffer around the margins of the plot in order to ensure that the core plot area is well reconstructed.

Table 1. Ground sampling distances at different camera elevations. Spatial resolution of images collected at different distances above a flat surface with a 24-megapixel full-frame camera (sensor dimensions: 36 mm × 24 mm) outfitted with a 24-mm lens.The horizontal field of view (HFOV), vertical field of view (VFOV), total imaged area, and GSD are provided for a range of distances above the substrate. The greatest proportional increase in the area imaged occurs between distances of 0.5 to 1.5 m above the substrate. Note, however, the GSD is also a function of the number of available megapixels.

Vertical distance from bottom (m)	HFOV (m)	VFOV (m)	Imaged area (m²)	GSD - 24MP (mm/px)
0.5	0.75	0.5	0.375	0.124
1.0	1.5	1	1.5	0.248
1.5	2.25	1.5	3.375	0.372
2.0	3	2	6	0.496
2.5	3.75	2.5	9.375	0.620
3.0	4.5	3	13.5	0.744
3.5	5.25	3.5	18.375	0.868
4.0	6	4	24	0.992
4.5	6.75	4.5	30.375	1.116
5.0	7.5	5	37.5	1.240
5.5	8.25	5.5	45.375	1.364
6.0	9	6	54	1.488
6.5	9.75	6.5	63.375	1.612
7.0	10.5	7	73.5	1.736

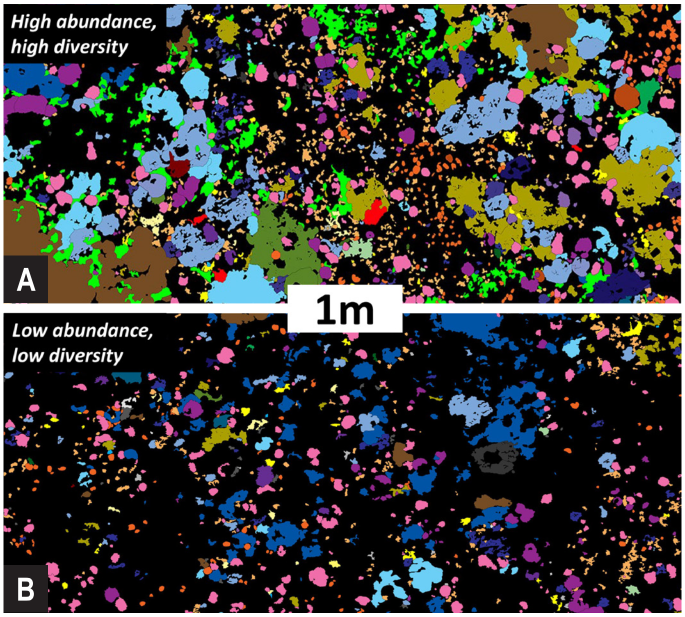

Figure 19. Example of difference between high and low abundance / diversity plots. The level of sampling intensity required to obtain representative demographic samples varies between locations. The orthoprojection shown in (A) contains over 1,500 coral colonies representing 37 species, and each species is represented by over 30 colonies, while the second example (B) contains fewer than 300 colonies of 9 species, and some taxa are present only a handful of times. The sample sizes required for the target analysis should be based on a priori assumptions of variance and effect sizes, and verified with power analyses.

When LAI is collected for purposes of spatial analyses or to extract structural metrics such as rugosity, imagery must be collected at a scale commensurate with the ecological signal of interest. Some phenomena, such as changes in complexity as a result of coral growth, can occur at relatively small spatial scales (<10-m linear extent). Others, such as differences in complexity between spur-and-grove vs. patch reefs, occur over much larger spatial scales (>100-m linear extent) and will thus require much larger plot sizes to be adequately characterized (McCarthy et al., 2022). Similarly, the level of replication needed is also a function of the experimental or monitoring design of a particular project. In some situations, it might be appropriate to achieve needed replication through subsampling of a single larger plot. However, for some ecological processes, it will be necessary to establish spatially separate plots so as to maintain sample independence and avoid pseudo-replication. Decisions regarding the spatial extent and needed replication of plots should be based on the ecological questions of interest; however, it is important to consider the impact of these decisions across the LAI pipeline.

3.4. Camera selection and operation

Once the necessary detail of imagery and spatial extent over which images will be collected has been determined, the next decision involves selecting a camera type and how it will be used. When deciding which camera is most appropriate, it is critical to distinguish between the concepts of image resolution and image sharpness, both of which ultimately contribute to the level of detail in a given image. As discussed briefly in section 2.1.1, image resolution is a function of the number of pixels available on the camera sensor and the physical extent of a pixel in an image, which is determined by the focal length of the lens and the distance of the camera from the bottom when the image was taken. Sharpness, on the other hand, refers to the degree to which an image is in focus or the portion of that image that is in focus. LAI can be collected over a wide range of resolutions, but images must be sharp to be useful for feature detection during model construction, as well as for human interpretation during stages of data extraction. Full-frame and other high-end cameras have large light sensors and a wider range of available settings, giving them the capacity to produce in-focus images under a variety of lighting conditions. Point-and-shoot and action-style cameras with smaller cropped sensors can perform well in some settings, but generally produce less-sharp images than higher-end cameras under the same conditions. Many cameras are equipped with sensors featuring increasingly high pixel counts; however, the ability to produce sharp imagery largely remains a function of the physical size of the sensor, the settings used, and other optical characteristics of the camera, lens, and underwater housing. Metashape can provide estimates of image sharpness (https://www.agisoft.com/downloads/user-manuals/) as a function of pixel contrast. However, as this measure is sensitive to compression and other camera settings, it should not be used as an approach to evaluating image focus.

When any camera is operated closer to the bottom, or when a higher focal length lens is used, images will cover a smaller spatial footprint and thus have smaller GSD regardless of the size of the camera sensor. However, the amount of light that reaches the camera sensor poses a practical limitation to the distance from the benthos at which imagery can be collected. Generally speaking, larger form factor cameras will be able to produce detailed images over a greater range of distances from the bottom, as they are able to capture more light and have a greater ability to resolve lighting differences. Advantages of small form factor cameras are generally based on cost or related to their ease of use and lower profile in the water. For these reasons, they can be particularly useful to obtain imagery over large areas for purposes of describing reef structure or functional-level community gradients. However, when highly detailed imagery is required, smaller form factor cameras must be operated closer to the benthos to collect imagery at the same level of detail as larger format cameras. Such considerations are critical, as the spatial footprint of each individual image together with the desired plot size will dictate the total number of images that will be needed and will also inform how best to collect those images.

Regardless of camera type, a major cause of lack of focus results from motion blur. The primary solution to motion blur is to move the camera more slowly over the benthos; however, camera settings can also be optimized to reduce motion blur. In particular, higher shutter speeds will generally reduce motion blur but also reduce the amount of light that reaches the camera sensor resulting in underexposed images. To compensate, higher light sensitivity and/or a wider aperture (i.e., low f-stop values) can be used. Care should be taken, however, as greater light sensitivity (i.e., higher ISO-values) can help to reduce motion blur but can result in graininess, which reduces detail in imagery. Similarly, wide apertures can result in a reduced depth of field (DOF), or the distance range within which an image is relatively sharp and in-focus. It is typically important to maximize the available DOF, particularly in complex habitats where some parts of an image are closer to the camera than others. Higher-end cameras with larger sensors typically feature greater ranges in available aperture and light sensitivity but can also suffer from reduced available field of view. However, the wider variety of available shooting modes can optimize these trade-offs in settings under different scenarios. The issues highlighted here become increasingly problematic in more complex habitats or the closer to the benthos the camera is operated, and as noted, the most reliable approach to collecting in-focus imagery is to reduce the speed at which the camera is moved over the area to be imaged area.

Care should also be taken when selecting a lens configuration, as well as camera settings, as these can impact the level of image detail that is achieved. Extra-wide-angle “fish-eye” lenses, as are common on action and smaller form factor cameras, suffer from significant geometric distortion and a lack of focus at image margins. High focal length zoom lenses also suffer similar distortion, and importantly, interactions between the lens and the camera housing can cause further distortion, regardless of the lens being used (Menna et al., 2017). While Metashape and other SfM software used to generate LAI can correct for this distortion using models of camera and lens parameters, the LAI generated from distorted imagery will tend to be of lower overall quality (Carbonneau and Dietrich, 2017; O’Connor et al., 2017; Eltner and Sofia, 2020; Lochhead and Hedley, 2021). While computational approaches can compensate for geometric distortion to a degree, they cannot correct areas of images that are out of focus. Further, as noted above, it is important to consider how lens optics and camera settings interact with each other. Increasing the focal length of the lens will reduce the amount of light that reaches the camera sensor, and settings such as aperture and light sensitivity must be allowed to change in order to compensate. Further, different combinations of lens and camera housings can affect the available DOFs, regardless of the settings used. Similarly, when a wider aperture is used (i.e., lower f-stop values), for instance to help reduce motion blur as mentioned above, the DOF can be reduced as a result. Moreover, higher-end cameras will have greater flexibility in settings, and available lenses, maximizing the ability to obtain in-focus imagery with a wide range of settings. An illustration of the trade-offs between aperture, shutter speed, and light sensitivity on image quality is shown in Figure 20. Regardless, it is recommended to thoroughly test available options before full-scale deployment of any camera system. For more details regarding camera specifications, see Part II.

A seemingly simple but critically important aspect of camera specification regards the manner in which the camera shutter is actuated. Depressing the shutter button manually requires the diver to operate a trigger on the housing, which can be fatiguing over time and difficult in some environmental settings. In cases where a limited number of photos need to be taken, such as for single-colony models or small plots, manual shutter operation can suffice. However, to standardize the rate at which images are collected and allow the diver to focus on buoyancy and navigation, it is generally recommended to select a camera model with a built-in interval timer (i.e., intervalometer) that allows images to be taken at a specified interval. It is typically recommended to use a frequency of one image per second, as this allows the camera’s autofocus time to adjust to the changing scene as the diver moves along the bottom while allowing imagery to be conducted with sufficient overlap and at easy-to-manage speeds. Some cameras can be outfitted with a remote shutter release that can be programmed to trigger the shutter at a specified interval. In most cases, using a remote shutter release underwater requires custom underwater housings or aftermarket modifications to consumer housings. Operating electronics underwater is inherently risky, particularly when working in dynamic settings such as from small boats, as many scuba diving activities frequently do. Camera housings should thus be selected or designed to be as robust as possible, with a minimum number of potential failure points, and it is strongly recommended to select a camera with the necessary features built in and to avoid custom solutions with

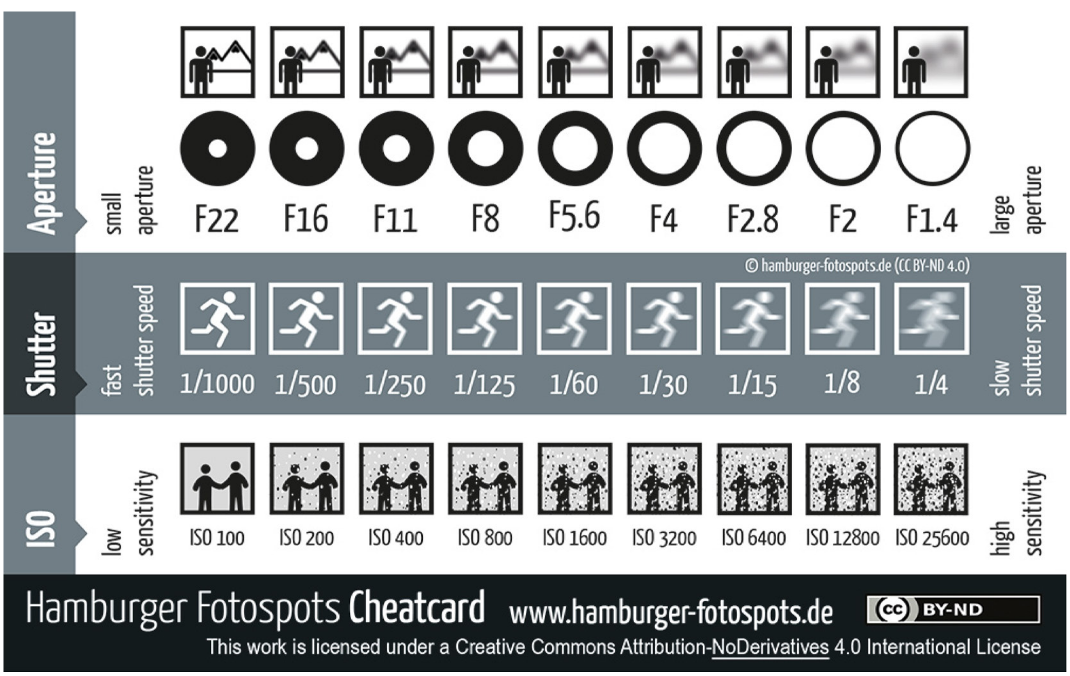

Figure 20. The aperture, shutter speed and ISO triangle. Various camera settings affect the degree to which an image is in focus. Among the most important settings are aperture, shutter speed, and light sensitivity (e.g., ISO). Larger apertures (i.e., lower f-stop values) allow more light to reach the sensor but can decrease the DOF. Slower shutter speeds allow more light to reach the sensor but can result in motion blur for moving objects or when the camera itself is moving. Finally, higher light sensitivity can compensate for low f-stop values or shutter speeds but can result in grainy images that lack detail. Image credit: hamburger-fotospots.de

additional failure points. Unfortunately, intervalometers are not available on most camera models, though most feature burst or continuous shooting modes that will allow the camera to rapidly collect images. These modes will collect approximately 5 images/s with the diver continuously depressing the shutter, which is less fatiguing than pressing the shutter for each image. However, less time for autofocus between images can result in many blurry or out-of-focus images that need to be removed either manually or via tools in Metashape during model processing (Part II, section 5.6). Moreover, to minimize operational complexity in the field, it is recommended to select cameras not only for their optical qualities but for their practical capabilities.

It is also critical that cameras be equipped with adjustable white balance settings as well as the ability to collect imagery in a raw file format (e.g., CR2, RAW, SRF, or NEF) in addition to the standard JPEG file format. While many cameras are equipped with automatic white balance modes and presets for underwater white balance, it is generally recommended to manually white balance underwater using a color card. In some locations, there can be considerable depth variation within a single plot; however, it is generally not feasible to change white balance during the course of a single dive. To avoid significant variation in color correction, it is thus recommended to conduct the white balance at the median depth of the plot. When possible, it is also recommended to collect imagery in a raw image file format in addition to standard JPEG. The JPEG file format is useful for most viewing purposes but represents a file compression that includes the white balance, therefore limiting flexibility in color correction using image editing software. Raw image file formats provide direct access to the image as it was collected by the sensor and before white balance is applied, thus allowing full control of color correction. RAW file formats also allow considerable flexibility in adjustments to exposure, particularly for underexposed imagery and low-light settings in general. However, in most cases, there is no correction for overexposed imagery, and care must be taken to use appropriate settings, particularly when operating in shallow, well-lit environments. When selecting a camera, it is important to consider its functionality and how it can be operated with respect to stated project objectives. Ultimately, the goal is to collect imagery at the desired level of detail, and any camera type or combinations of settings can be used as long as this requirement is met.

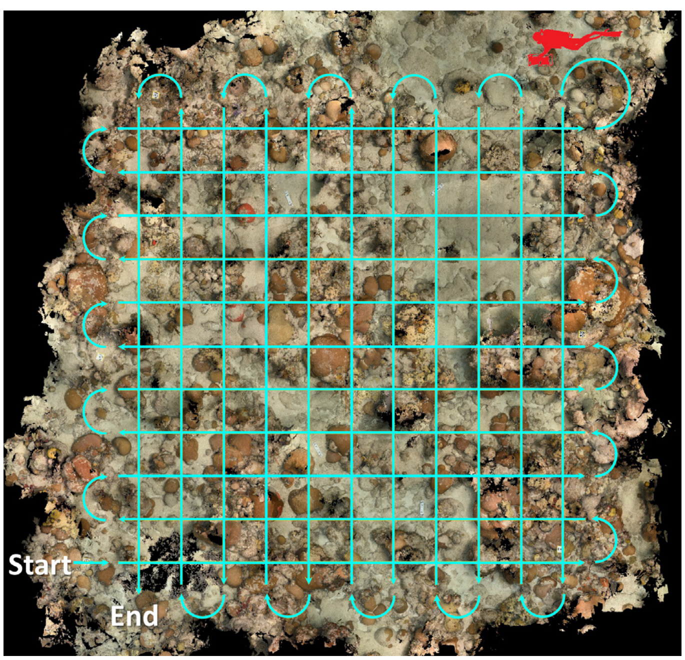

Figure 21. Example of gridded swim pattern. The double-lawnmower image collection pattern (two sets of passes perpendicular to each other) is an easy-to-implement approach that provides robust assurance that the area has been sufficiently imaged. The distance between adjacent passes and height the camera is operated from the bottom are designed to achieve 60%–80% side-to-side overlap. Similarly, divers should swim along each pass at a speed that allows for 60%–80% front-to-back overlap (see Figure 22).

3.5. Selecting an image collection approach: 100-m^2 plot example

Once decisions have been made regarding plot size and the camera to be used, the next task is to design a field approach and image collection pattern. To do so first requires an estimate of how long it will take to image the desired area. It is important to remember that when the spatial footprint of each image is smaller, more images will be needed per unit area to obtain sufficient overlap and the longer it will take to image the area. Moving the camera faster is not a viable solution, as this can lead to motion blur. To develop an understanding of how to develop an image collection plan, consider the example of a 100-m^2 plot (10 m × 10 m core area with a 1-m buffer on each side), imaged at a level of detail that allows corals to be identified to the species level. This example uses a Nikon digital single-lens reflex (DSLR) camera with a complementary metal-oxide semiconductor (CMOS) sensor and an 18-mm lens and is operated by a single diver in a "double-lawnmower" pattern (Figure 21). Images are collected once per second, moving at a speed that produces 60%–80% forward-backward overlap on adjacent images (Figure 22). The camera is operated 1.5–2 m above the bottom, and parallel passes are separated by 0.75 m, to again produce 60%–80% side-to-side overlap. Navigation is accomplished by diver memory aided by the use of temporary floats at each of the four corners of the plot. Using this approach, approximately 2,500 images can be collected

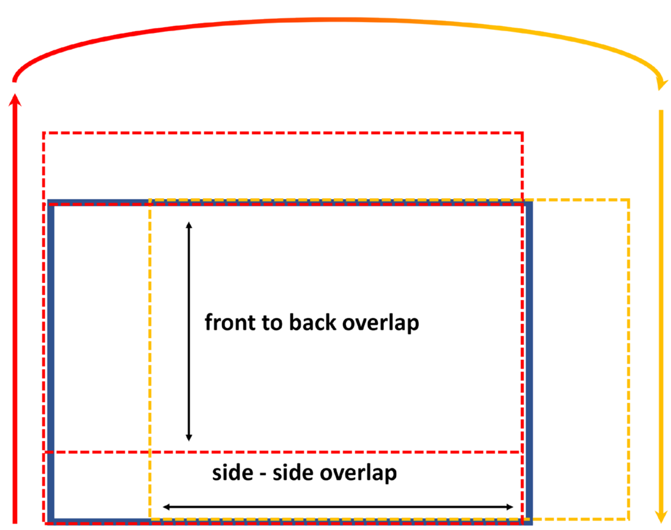

Figure 22. Conceptual example of image overlap. The first image (blue) has 80% front-to-back overlap with the second image (red). The third image (yellow) is collected on the return pass and has 80% side-to-side overlap with the first and second images.

in 45 min with a single camera, providing ample overlap between images and an additional 15 min for other dive activities including plot setup/cleanup and metadata collection. When working at mid-depths (8–12 m), this approach will allow sufficient time to operate safely and comprehensively image the plot; however, adverse environmental conditions such as currents or high wave activity can reduce available dive times (see Part II, section 7.6). At shallower depths, this approach will easily allow for a greater area to be imaged with similar image density and detail, and conversely, the areal coverage will be limited at depths greater than 12 m. In the latter case, and more generally, such limitations can be overcome in a variety of ways, including imaging the plot in multiple dives, having additional divers operating cameras, or through the use of a multi-camera array. However, once the area becomes sufficiently large, new approaches must be designed to image the area in a timely manner.

3.6. Selecting an image collection platform

The image acquisition approach outlined above is ideal for areas that can be imaged relatively quickly (e.g., during a single dive) by small teams of divers operating a single camera by hand. As mentioned in section 2.1.4, it is preferable to minimize the need for multiple surveys of contiguous areas and for the interval between subsequent surveys of adjacent areas. Once areas become larger than what can be collected by divers in a single day, alternative approaches are strongly recommended.

When imaging must be conducted at depths beyond 12 m, available bottom time will be limited for diver-based imagery acquisition. To reduce total dive time and limit the need to conduct multiple dives, a single diver can operate multiple cameras in a linear array (Figure 23). Cameras can be positioned along the array to achieve the required overlap at the expected minimum distance from the bottom and allow for a single pass by a diver to cover a wider swath of the benthos. Diver-operated camera arrays are ideal for situations where there is limited time available for imaging or to increase the spatial extent that can be imaged. However, such arrays can be cumbersome for divers to operate, particularly in current or surge. In such cases, the use of large camera arrays can instead reduce the total dive time available, negating the benefit of the wider swath. To compensate, additional equipment, such as diver propulsion vehicles (DPVs) can be used to extend dive time for skilled divers using large camera arrays. DPVs can allow for larger reef areas to be imaged (or to overcome effects of current) across depth strata. Further, DPVs can be equipped with camera arrays with essentially no penalty on dive time or ease of use. DPVs can be costly, and many dive programs require special training before they can be used. Further, few commercially available DPVs are pre-equipped with camera mounts, so most require custom modifications. However, DPVs provide skilled divers greater operational flexibility over a wide range of environmental conditions (Figure 23).

Figure 23. Linear camera arrays and DPV images. Linear camera arrays (A)–(C) outfitted with multiple cameras can allow divers to image a greater spatial extent in a single dive without sacrificing image resolution. DPVs (C) allow even greater flexibility, by allowing divers to operate a larger array and add additional instruments such as USBL receivers or tablet displays for navigation. Credit: (A and B) Jessica Levy (Coral Restoration Foundation and KAUST Reefscape Restoration Initiative); (C) Brian Zgliczynski (SIO, UC San Diego and Center for Environmental Imaging)

For areas that are large and relatively shallow, image acquisition can be accomplished via surface-vessel-towed or surface-vehicle-based camera arrays that enable dramatic increases to the available operating time (Figure 24; Raber and Schill, 2019; Hatcher et al., 2020). Surface arrays can be configured with multiple cameras in a linear array with fixed overlap between cameras. The distance between cameras should be set so that sufficient overlap can be achieved between cameras when operated at the expected minimum distance from the bottom (Figure 25). Towed surface arrays are typically operated behind surface vessels, allowing navigation via GPS. Care must be taken to operate the vessel and position the array in a manner that eliminates propeller wash underneath the camera. Towed surface arrays are typically operated from the surface and are thus ideally utilized in shallow habitats with low turbidity. When operated in areas where the depth exceeds 3 m, image resolution will be reduced to the degree that detailed, species-level taxonomy will not be possible. For depths of 3–10 m, functional-level and, in some situations, genus-level information can be extracted from imagery if the water is clear. There is more flexibility when imagery is collected for the purposes of obtaining structural metrics or habitat classification at the functional level (e.g., hardbottom vs. sand). Operating times of surface arrays are still limited to available camera battery; however, there are more opportunities to develop robust custom systems that provide constant power supply to cameras. Surface-towed arrays also benefit from more straightforward GPS collection that can be used to create detailed navigation and image collection plans. This information can also be used to georeference imagery; however, it is critical that any offset between the GPS device and the camera array be minimized or explicitly accounted for to facilitate post-processing of georeferencing information.

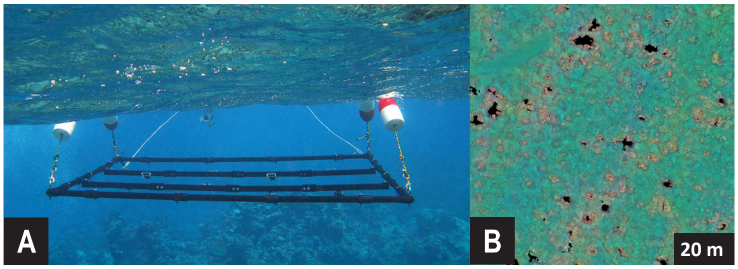

Figure 24. Towed camera array and orthoprojection example. Towed surface array (A). The surface vessel is operated at a minimum speed in order to ensure adequate front-to-back overlap and is towed at a distance to further avoid propeller wash. Towed arrays can be operated slightly below the sea surface in order to minimize camera movement due to surface conditions. The array shown here was used to image a 65,000-m^2 plot at Palmyra Atoll in 2014 (22,500 m^2 orthoprojection shown in [B]). Note, black spots are holes in the model resulting from extremely shallow spots on the reef (<1-m depth) where the camera array could not be operated safely. Credit: (A) Gareth Williams (SIO, UC San Diego and Bangor University)

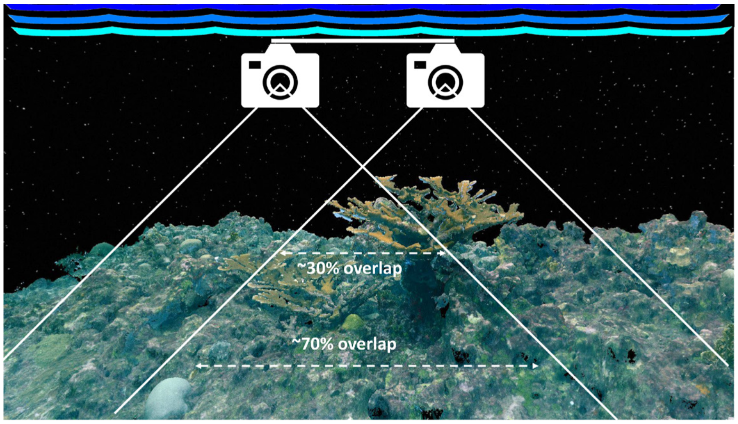

Figure 25. Designing overlap with camera arrays. When a towed surface array is being operated, cameras should be positioned to ensure adequate side-to-side overlap based on the depth of the imaged area. In the conceptual example shown here, camera spacing would be adequate for most features but insufficient to capture complex and important features whose shallower depth prevents sufficient overlap in images taken from the camera mounted to the array. If such cases are common, cameras should be spaced more closely to obtain the needed overlap. Alternatively, if such features are relatively rare, conducting more closely spaced passes will ensure sufficient overlap in these regions of the plot.

Finally, autonomous underwater vehicle (AUV) and remotely operated vehicle (ROV) platforms have essentially no depth limitation, and some devices have long endurance (e.g., 12 hr) or the ability to hot-swap batteries, offering the potential to dramatically increase the area that can be imaged daily. Currently, there are only a handful of readily available AUV or ROV platforms that can be easily deployed in shallow-water coral reef environments and feature the dynamic positioning capabilities needed to operate in high-energy environments or sufficiently close to the benthos on highly complex reefs to produce highly detailed LAI (Pieterkosky et al., 2017; Raber and Schill, 2019; Lesser and Slattery, 2021; Price et al., 2021). However, several models featuring such capabilities are currently in field trials (Figure 26). Regardless of these innovations, current platforms are well suited to capture relatively lower-resolution data over large spatial extents. Moreover, while AUVs and ROVs hold the promise to dramatically expand the area that can be imaged daily, they are costly and require significant amounts of planning and support infrastructure, making them accessible only to technologically advanced or well-resourced organizations.

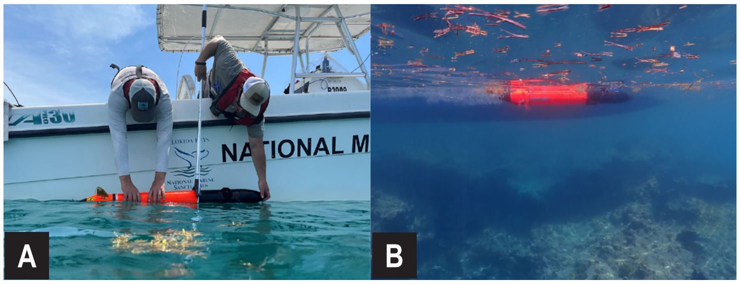

Figure 26. Example of micro AUV deployment and operation. Deployment of a BAE Riptide MK II micro-AUV with an Arctic Rays Swordfish camera payload at Eastern Dry Rocks, Key West, using standard small boat operations in May 2023 (A). Imagery was collected inside four separate 4,400-m^2 areas over the course of 12 days of operation. Existing bathymetric maps were used for mission planning, and the AUV navigated successfully in water depths from 4–9 m at heights 2–3 m above the benthos (B). Credit: Sarah Rojano (NOAA/CSS)

3.7. Limits to spatial extent

Imaging increasingly large areas has practical limitations, both with respect to the burden of collecting the imagery in the field and during the model generation process. A single instance of LAI is created from a group of contiguous raw images that overlap.

A viable approach for plots that cannot be imaged in a single bout due to logistical limitations is to subdivide and image them separately in the field and later process them as a single collection of imagery. However, it is critical to ensure that there is sufficient overlap among subplots to generate a continuous LAI product. More importantly, the images of subdivided plots should be taken as closely together in time as possible. To capture a snapshot of a given reef at a given moment in time, it is important to complete imaging before any ecological change can occur. Thus, imaging should occur in a timely manner such that the landscape does not change in any substantive way over the course of image collection. As described in Part I, section 2.1.4, LAI reconstruction requires hard matching of features, and if objects move around within the plot during imaging, it can lead to inconsistencies in the reconstruction that at the minimum will lower accuracy and in the worst case prevent reconstruction completely.

Regardless of whether or how the plot is subdivided, there are also computational limits to plot size that occur as a function of the total number of images that can be processed at once by Metashape (see Part I, section 2.2.5; and Part II, section 3.6). When image collections exceed 60,000 images of 24.5 MP each, some Metashape functions, particularly those that require use of image thumbnails, will become unreliable. Alternatively, image collections can be subdivided into separate "chunks" that can be processed separately and then combined into a single continuous model. Chunk alignment greatly benefits from increased overlap among chunks, and it is recommended to devise an image collection pattern in the field to ensure needed overlap.

3.8. Developing a navigation plan

Prior to image acquisition in the field and regardless of the level of detail, spatial scale, or imaging platform being used, it is important to develop a systematic navigation plan to ensure that the entire area is imaged with consistent overlap and detail. The first element of navigation regards maintaining a fixed distance from the bottom in order to achieve a constant level of detail and spatial footprint of imagery across the plot, the latter of which is critical to ensuring consistent overlap (Figure 27). In complex landscapes, the diver will need to adjust their depth to maintain distance from the substrate, which requires buoyancy adjustments and can be challenging on reef slopes or other areas with high relief. A greater challenge is to reliably navigate in a manner that ensures the entire plot is imaged with adequate overlap and similar levels of image density.

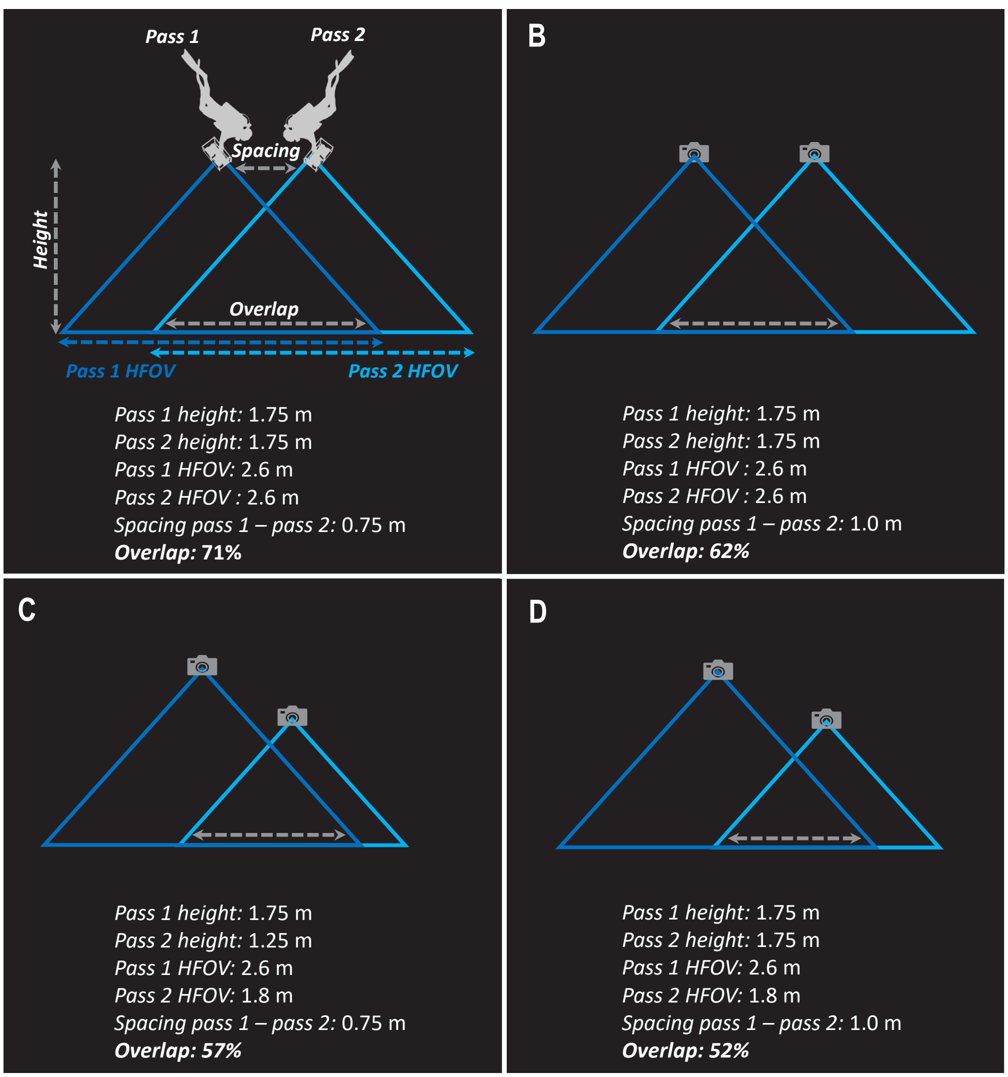

Figure 27. Effect of swim height on overlap and model quality. The degree of overlap between images from adjacent passes is a function of spacing between passes and the horizontal field of view (HFOV), or the horizontal swath of the benthos covered by an individual image. In (A), passes are conducted 1.5 m above the bottom, and pass spacing is 0.75 m. The HFOV of each camera is 2.6 m (see Table 1) and the overlap between adjacent passes is 71%, which is adequate for LAI generation. In (B), passes are conducted at the same height, but spacing increases to 1 m. Consequently, overlap drops to 62%, near the minimum needed for accurate LAI. In (C), the first pass is again conducted at a height of 1.5 m above the bottom and covers a 2.6-m HFOV; however, the second pass is now conducted at a height of 1.25 m, and HFOV drops to 1.8 m. Overlap between these passes falls to 57%. In the final scenario (D), the second pass is again conducted 1.25 m from the bottom; but pass spacing increases to 1 m, and overlap between passes drops to 52%. Given the inherent unpredictability of navigating underwater, these final scenarios would most likely lead to inaccurate and incomplete LAI.

The most straightforward collection pattern is a systematic grid (e.g., double-lawnmower) (Figure 21). Consistent overlap in sequential images is straightforward to achieve by maintaining a constant swim speed that ensures 60%–80% forward-backward overlap. Ensuring adequate side-to-side overlap between passes is much more difficult and first requires moving a fixed distance between adjacent passes and then maintaining that spacing along the length of the pass. Moving a fixed distance between passes can be aided if plot edges are delineated with transect tapes, though this approach can be unreliable in practice as the tape is rarely easy to read while navigating and operating the camera. Most commonly, spacing between passes is estimated by the diver using their body, objects on the ground, or the camera itself. It is generally recommended to be conservative with pass spacing, as overestimating spacing can lead to insufficient overlap, reduced model accuracy, and in the worst case, non-continuous LAI (Figures 27 and 28). Overly conservative spacing should also be avoided, as it will result in a greater number of passes, potentially exceeding available bottom time and leading to incomplete imaging of a portion of the plot. For an alternative example of image collection patterns, please see Part II, section 5.6.

With consistent spacing, adequate overlap between passes can be achieved. The next challenge is to maintain this spacing and navigate straight lines underwater, which is notoriously difficult and becomes more so the longer the distance being navigated. When the plot is small enough or visibility is sufficient to see across the plot, it is possible to navigate by line of sight. In such situations, it is also helpful to use objects on the bottom to maintain direction, though care must be taken as ambiguous features in homogenous landscapes can be misleading. Establishing landmarks along a pass can also help maintain spacing between the initial and subsequent passes. In general, maintaining a mental map of the benthos during imaging is extremely helpful to ensure that the entire area has been imaged. However, in some cases, particularly when diving conditions are rough, maintaining buoyancy and situational awareness, operating the camera, swimming straight, and remembering and finding specific corals or features on the bottom can be overwhelming, even for highly experienced divers. Regardless of the setting, it is recommended to mount a high-quality compass in a visible location on the camera to aid navigation. A compass provides critical backup and reduces the overall burden on the diver but is not a replacement for situational awareness. Compass operation should be complemented with the use of landmarks, as small errors in navigation can accumulate over longer distances. Further, when navigating across distances of 10 m or further, it is strongly recommended to use additional visual navigation aids, such as temporary floats. Floats placed at regular intervals can be used not only to complement compass navigation but also to estimate spacing between passes. Care should be taken when incorporating markers such as weighted floats, as deployment and retrieval can be time consuming and add operational complexity.

While the low-tech solutions described here are viable in a variety of settings, underwater positioning technology has become increasingly accessible in recent years and, in some situations, may be requisite. The most straightforward and popular approach employs USBL acoustic communication, which uses a surface base station transceiver and one or more transponders mounted to the camera(s) in order to track position. This information can be fed to divers via a tablet display, providing the ability to correct any navigation errors or redo a poorly navigated pass (Figures 23 and 28). When properly calibrated and used in appropriate environmental settings, these systems can provide precise real-time location and tracking information that can be used for navigation. As mentioned in section 2.1.4, image-level GPS positions can also be used to reduce reconstruction times and improve model accuracy (Figure 10; also see Part I, section 3.11). However, the cost of USBL and similar devices can range from several thousand to well over ten thousand dollars and require some degree of expertise to operate. Further, the rate that information can be fed to the subsurface unit is a function of the distance between the transponder and receiver; however, the required direct line of sight between them is not always possible, particularly in shallow or physically complex habitats. Intermittent signal lags can also lead to erroneous information being sent to the subsurface unit. However, when used properly, these devices can provide critical navigation information which would otherwise be unavailable.

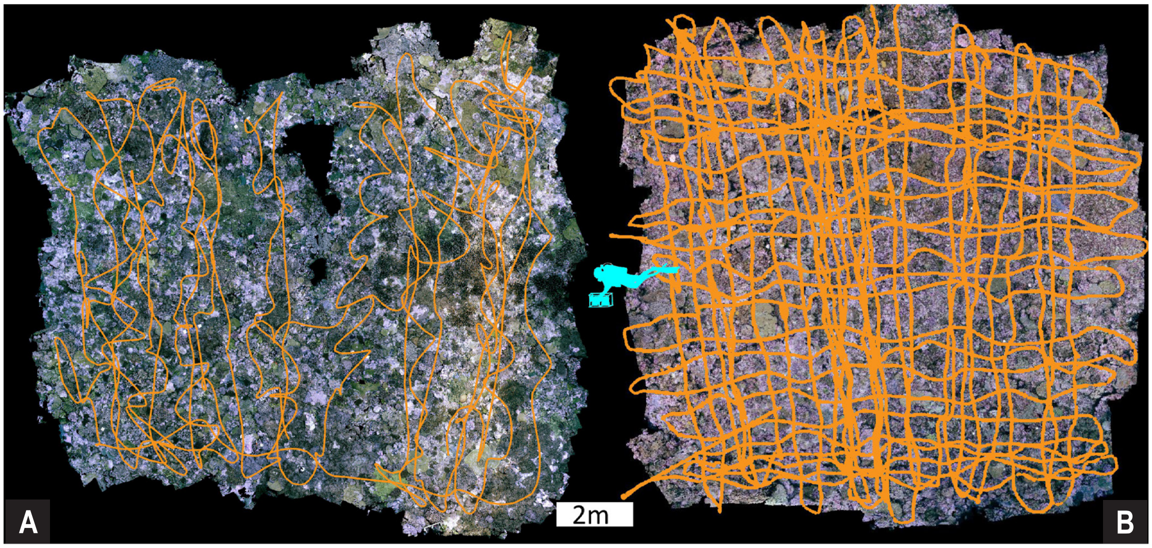

Figure 28. Comparison of image collection. Comparison of a plot conducted with a single set of widely spaced passes (A) relative to a more thoroughly imaged plot using a gridded pattern and close pass spacing (B). With fewer navigation aids available, passes were not straight, and together with the wide spacing, the resulting reconstruction was incomplete.

3.9. Navigation example

To illustrate the challenge of navigation and the potential impact on model construction, examples are provided in Figure 28. In the image shown in (A), the diver used a single set of long parallel passes to systematically image the plot, while in the image shown in (B), the diver used the double-lawnmower pattern to collect images. In the first example, the plot was delineated only with transect tapes around the perimeter, while in the second case, a series of floats and markers were placed around the perimeter of the plot. Transect tapes along the perimeter are useful to guide spacing between passes but can be difficult to see in complex landscapes or when the camera is operated close to the bottom. Floats can be seen from across the plot and can thus be used by divers to maintain heading. Note that, in the second example, passes are more closely spaced, more numerous, and from different directions, thus providing ample overlap and several different views of each location of the benthos. As a result, it is not surprising that the model from the first example is incomplete, and, while a model was generated, without consistent overlap, it suffered from several large gaps in the reconstruction. In both examples, the boundaries of each plot are ragged, which is an inevitability of LAI, as pictures on plot margins will overlap with fewer other images. As a result of these lower levels of overlap, reconstruction quality is also lower at plot margins, and to ensure the target area is well reconstructed, a buffer of 2 or 3 camera passes (approximately 1–2 m outside the core plot area) should always be included in imaging plans.

3.10. Metadata: scale and depth

A model generated without any additional information will have an internally consistent yet arbitrary coordinate system without scale or a fixed orientation; however, both are required in order to extract most quantitative metrics from LAI products (Figure 29). To convert these arbitrary units into an interpretable distance, some scale information must be present in the imaged scene or embedded in the image metadata. Depth information collected in the field is used to orient the model with respect to the sea surface and thus defines the plane of projection for creating orthophotomosaics (see Part I, section 2.25). To obtain and apply this information, scale bars and depth markers must be deployed in the field and present in the raw imagery and fully reconstructed model. Depth information must be collected in at least three locations to orient the model, as long as those locations do not fall on a single line (in which case at least one additional measurement must be collected). Though at least four points are required to estimate error in plane fitting, six or more such locations are recommended in practice. Virtually any item of known length can be used for scale, and depth markers can similarly be of any design. In both cases, the only requirement is that they are clearly visible in raw imagery and are accurately reconstructed in the final model. This necessitates that any markings are clearly legible in raw imagery so that they can be easily identified and that the item remains stationary during imagery so that they reconstruct accurately. Keeping markers stationary can be problematic in high-energy environments, and to prevent the markers from moving during imaging, it is recommended that markers and scale bars have a low profile and are sufficiently weighted (1–2 lb). Further, scale bars and depth markers should be distributed throughout the plot, both to provide redundancy in the event one or more markers does not reconstruct well (e.g., Figure 13) and to improve the precision of orientation and scaling. Markers have an added advantage of being useful to aid navigation, but care should be taken to avoid overuse of weighted items, as the added weight can pose considerable operational complexity and dive safety considerations in some settings.

Scale and depth information can be provided to Metashape during the model reconstruction process through manual and semi-automated processes (see Part II, sections 4.6 and 6.6) or subsequently in other visualization and analysis software (see Part II, section 7.6). When supplied directly to Metashape during the reconstruction process, this information can be used for model optimization procedures. Such workflows can be valuable to correct model geometry in cases of poor image quality or low overlap, though they can add moderate hands-on time during model processing (see Part II, section 4.6).

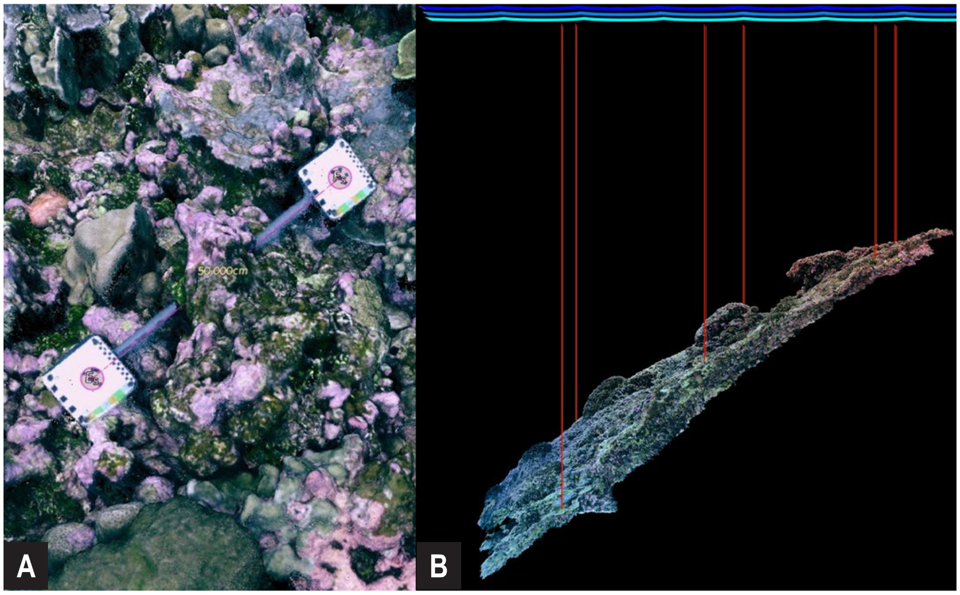

Figure 29. Scale and orientation image. In order to provide meaningful scale and orientation for LAI, measurements must be collected in the field, and scale bars (A) must be deployed during imaging. Multiple scale bars should be placed throughout the plot, in secure locations to ensure they are well reconstructed. Depth information should be collected at a minimum of three locations around the plot. In (B), depth measurements were collected at six locations around the plot (red lines) to allow orientation with respect to the plane of the sea surface.

3.11. Metadata: geolocation

Geolocation information is a critical piece of metadata that must be collected in the field. There is a wide degree of flexibility in how this information can be collected, depending on the specific analysis being conducted. Regardless of the objective, GPS location information should always be collected for each study site, as this is not only standard reporting information but necessary to relocate plots and repeat surveys. In most shallow-water settings, location information can be obtained from a single point using a consumer-grade handheld GPS from the surface. When models are properly scaled and oriented, precise measurements can be collected within a given model without the need for more thorough georeferencing. However, questions such as how communities respond to directional forcings (e.g., wave exposure and current) require map orientation of the plot (Corso et al., 2022) in addition to scale, orientation, and geolocation. This information can be easily obtained with a compass heading between known locations in the model, such as depth markers. If more precise location information is required, for instance to precisely measure distances between locations in different LAI plots or between LAI plots and other mapped environmental features, advanced positioning tools such as RTK GPS devices and USBLs are necessary. Such systems provide not only positional information to each image but also up to centimeter-level georeferencing for LAI.

When precise image-level position information is available, it can be used by Metashape as an initial "guess"' of camera position. These preliminary pose estimates reduce the number of image pairs that must be searched by the software during feature matching and can thus dramatically reduce processing times. Further, in cases where an image collection involves tens of thousands of images collected over large areas, such comprehensive information can be critical for optimization procedures. Care must be taken with the use of USBL and RTK GPS devices, as systematic offsets can lead to increases in model reconstruction time and decreases in model accuracy and therefore must be corrected before model construction. As mentioned, these tools require substantial financial, logistical, and intellectual investment, and thus, their ultimate utility to research questions should be explicitly considered before investment.

3.12. Collecting time series imagery

Most evaluations of coral reef ecosystems or restoration success require time series information to track change in coral communities. The LAI approach is well suited to temporal studies, as it provides views of sites that enable simultaneous tracking of multiple locations and individuals without the need for painstaking relocation underwater (Figure 30). In practice, there are no substantive differences between surveys of sites that are intended to be imaged only once, or the initial collection of imagery at a site that will be studied through time. Location, scale, and depth information should always be collected, and the imaging approach is the same, whether during the initial survey or subsequent ones. The precision of many consumer grade handheld GPS units is often better than ± 3 m (Shamshiri and Ismail, 2013), which is sufficient to get "close enough" for divers to enter the water and relocate a site. Next, when permanent markers are installed during the initial survey, they also can be used to aid relocation of the exact plot location. Markers should be installed in well-consolidated locations where they will not be expected to be dislodged, using high-grade steel (316) that will not deteriorate between sampling events. LAI from the initial survey can be used to relocate the plot and any permanent markers by printing the 2D map view orthoprojection on underwater paper. These maps can then be used by swimmers on the surface or taken underwater by divers to relocate the plot. Such maps can be extremely useful as they can also be annotated with depths and the locations of other key features from the site that are particularly useful for relocation, such as a sand patch or large coral head. A caveat to the use of maps for relocation occurs when significant change occurs in the plot area, such as that due to prolific growth or storm damage, which can reduce the utility of the map. In such cases, both precise GPS information and previously installed permanent markers can be critical to relocation. It is important that key metadata such as scale and depth are captured during every survey of a time series plot, and not only in the first survey. However, if unexpected events prevent this information from being collected, software tools exist that allow this information to be derived from the initial survey.

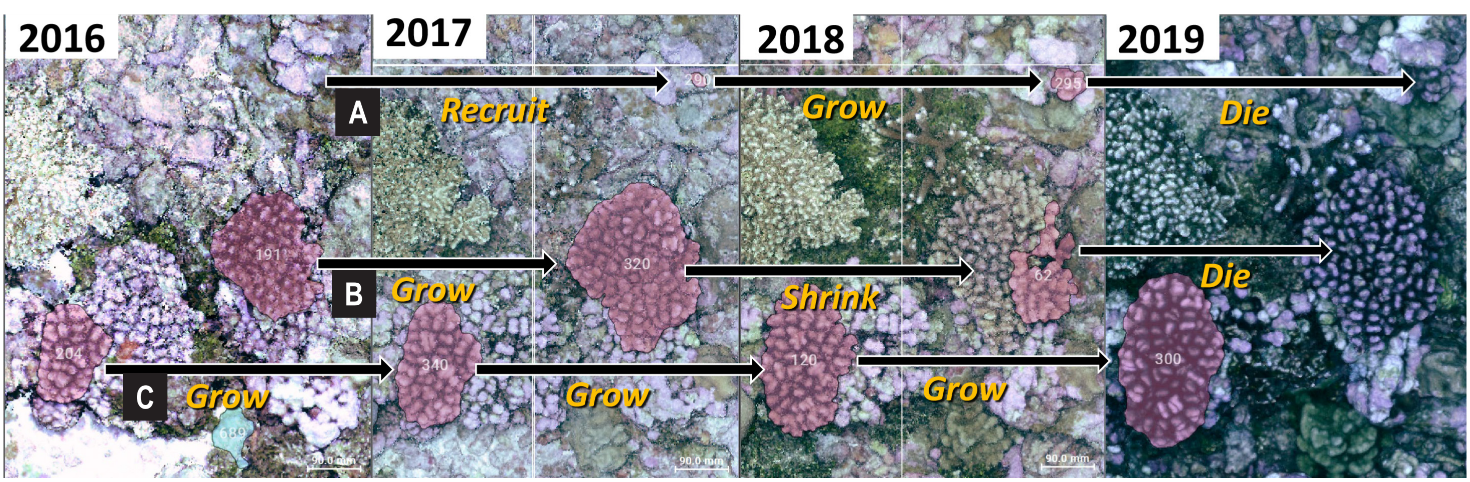

Figure 30. Example of diversity of demographic fates. A four-year time series from 2016 to 2019 illustrating the various demographic outcomes that can be observed with coregistered time series. In the first sequence (A), a colony recruits and grows for two years before suddenly dying. The colony in the second sequence (B) grows for a year before suffering partial mortality and eventually dying. The final sequence (C) shows a colony that grows consistently through the time series. Over 400 colonies were tracked from 2016–2019 in this time series collected at Palmyra Atoll, USA.

3.13. Model construction infrastructure

After estimation of the volume of raw imagery to be collected and any associated metadata, a plan can be developed for the computational infrastructure that will be needed to meet project objectives. The primary trade-off to be considered is between the timeline on which LAI products must be generated and the resources that can be devoted to building the infrastructure needed to meet these timelines. Depending on the number of images collected in each plot and the number of total plots collected, the model construction process can be a major rate limiting step for LAI workflows. In general, the speed at which a given model can be generated is a function of the number (and speed) of available CPU and GPU cores, and the more processing power that is available, the faster a model can be generated from a given collection of images. The maximum number of images that can be processed in a single project is largely a function of the available RAM. When using the same computer configuration, model generation times will increase non-linearly as the number of images increases. For example, building a model from a collection of 200 24-MP images on a quad-core laptop equipped with a high-end GPU and 16 GB of RAM will require approximately 60 min of processing time. Processing a collection of 2,000 images on this same system would require over a week of runtime to complete model reconstruction. On the other end of the spectrum, a custom HPC platform equipped with a 24-core CPU, a high-end gaming GPU, and 256 GB of RAM can reconstruct a model from the same 2,000 images in less than a day. Importantly, the needed resources are not only a function of the number of images but also the size of those images, such that 24-MP images will require more resources than 12-MP images.

Cloud computing services offer an alternative to building hard infrastructure and can be an ideal solution for some groups. These services are offered via subscription, typically charging users per hour of usage, and a variety of hardware configurations are available. Cloud computing is ideal for those groups who might have only intermittent needs for HPC infrastructure, as the more costly configurations only need to be used for complex processing jobs, while more economical and lower-power configurations can be used for the vast majority of daily computational tasks. Regardless of the configuration being used, local internet bandwidth limits upload and download speeds, which may cause substantial delays in some cases. For programs collecting large volumes of imagery at regular intervals, it ultimately may be more cost effective to secure physical resources. Some academic or governmental groups might have access to supercomputing platforms (which are essentially local cloud computing services); however, these architectures are generally not optimized for LAI model generation. IT staff and Metashape documentation should be consulted and available configurations thoroughly tested before finalizing plans for use of cloud or supercomputing infrastructure. Again, all choices should be made with respect to project goals, delivery timelines, and available resources.

For groups regularly conducting LAI work at the scale of tens of thousands of images a day, obtaining dedicated computer hardware may be preferred, particularly as local processing does not rely on local internet bandwidth. HPCs are commercially available from a variety of vendors, including the major computer manufacturers. A commercially available HPC capable of processing large datasets (<10,000 images) relatively quickly can be purchased for $4,000–$8,000. Where expertise is available, an equivalent system can be custom built at a savings of 25%–30% relative to commercial products, and similarly, some consumer-grade systems can be modified aftermarket to add memory and storage at significant savings. For groups collecting large numbers of models and working under short delivery timelines, several HPCs can be networked together to create an HPC cluster. The time saved by networking scales with the amount of processing power available; however, so too does the maintenance needed to make sure that processing is not interrupted by configuration errors. Part II of this guide presents several examples of infrastructure used by groups with extensive experience in model construction and provides context of what can be done with various levels of investment. Ultimately, decisions about needed computational infrastructure should be balanced with available resources and the rate at which products need to be generated. While startup costs of building computer infrastructure are significant, they are fixed, and a well-maintained computer will last for at least several years and be capable of processing millions of images and thousands of models over its lifetime. When onboarding any new technology, it is recommended to pilot test configurations if possible or work with colleagues with previous experience to design the optimum system for specific applications.


Yoan Eynaud (SIO, UC San Diego)

3.14. Planning for ecological data extraction

Perhaps most important when designing and implementing the LAI approach are considerations of the effort that will be needed at the step of ecological analysis. As mentioned in Part I, section 2.3.4, the expectation is that AI tools will eventually accelerate data extraction steps; however, such workflows are still in active development. Plans should therefore be designed with the expectation that data extraction will be a human-driven step and will require dedicated investments of financial and institutional resources in the near term. Extracting ecological data from LAI can be conducted using a variety of free, commercially available, and custom software platforms. Nearly all of these software platforms can be operated at full capacity on modern consumer-grade PC laptops, though several do not currently run on Apple operating systems. Before data extraction can begin, time must be devoted to building the expertise needed to properly use the various analytical tools and workflows needed in the LAI approach, many of which are new and not commonly used. Thankfully, after this expertise has been built, some data streams, such as structural complexity, require minimal human intervention and can be conducted relatively quickly. Other workflows, such as those needed to derive estimates of benthic cover or to extract demographic data, can require substantially greater investments of person-time, depending on the level of effort needed to capture the ecological signal of interest. For instance, robust estimates of percent coral cover require a sampling density of approximately 25 points/m^2, but capturing rare or cryptic benthic species might require double that sampling density (Brown et al., 2004; Dumas et al., 2009). Further, the time needed to designate each point is a function of the underlying diversity of the plot, image quality, and expertise of the observer. Generally, conducting this work within a 100-m^2 plot at a density of 25 points/m^2 will require anywhere from 5–15 hr of effort for a single well-trained observer to produce. However, when balanced against other steps of the pipeline, the rate at which structural complexity or even more time-intensive data streams such as percent cover is generated, greatly exceeds the pace at which models can be generated even by a well-equipped HPC.

Extraction of segmented coral boundaries for spatial or demographic analyses is one of the most exciting and promising applications of the LAI approach (Figures 17, 19, and 30). However, care should be taken when designing workflows to produce these data, as they require substantial levels of effort that depend on the quality of the raw imagery and the comprehensiveness of the data being extracted. For instance, tracking yearly growth in slow-growing species with high levels of partial mortality and complex colony boundaries will require sub-centimeter-level precision, far more than is needed for spatial analyses, which rely simply on mapping colony centroid locations. Further, not all demographic analyses require complete colony segmentation. For instance, when a census of survivorship is the primary data required, such information can be extracted relatively quickly from hundreds to thousands of corals without the need for full colony segmentation. Additionally, the ability and time required to produce highly precise segmentation data are functions of both the quality of composite LAI products and of the underlying raw imagery. The lower the detail of LAI, the more difficult it is to segment LAI in the first place, and the more extensively raw imagery must be inspected to produce precise segmentation data. Moreover, time can be saved during data extraction by collecting more detailed raw imagery in the field, thus producing more detailed composite LAI products and reducing the reliance on the underlying raw imagery.

The time needed for data extraction is a function of not only plot size but also underlying ecological setting, including the abundance of corals in the plot and the complexity of their growth forms. For instance, segmenting 200 healthy hemispherical Colpophyllia colonies can be achieved in less than a day by a single well-trained observer. However, 200 colonies of Orbicella with complex colony perimeters resulting from partial mortality and fragmentation will require far more care to segment and more thorough inspection of raw imagery and, as a result, could take up to a week to segment with the same accuracy (Figure 31).

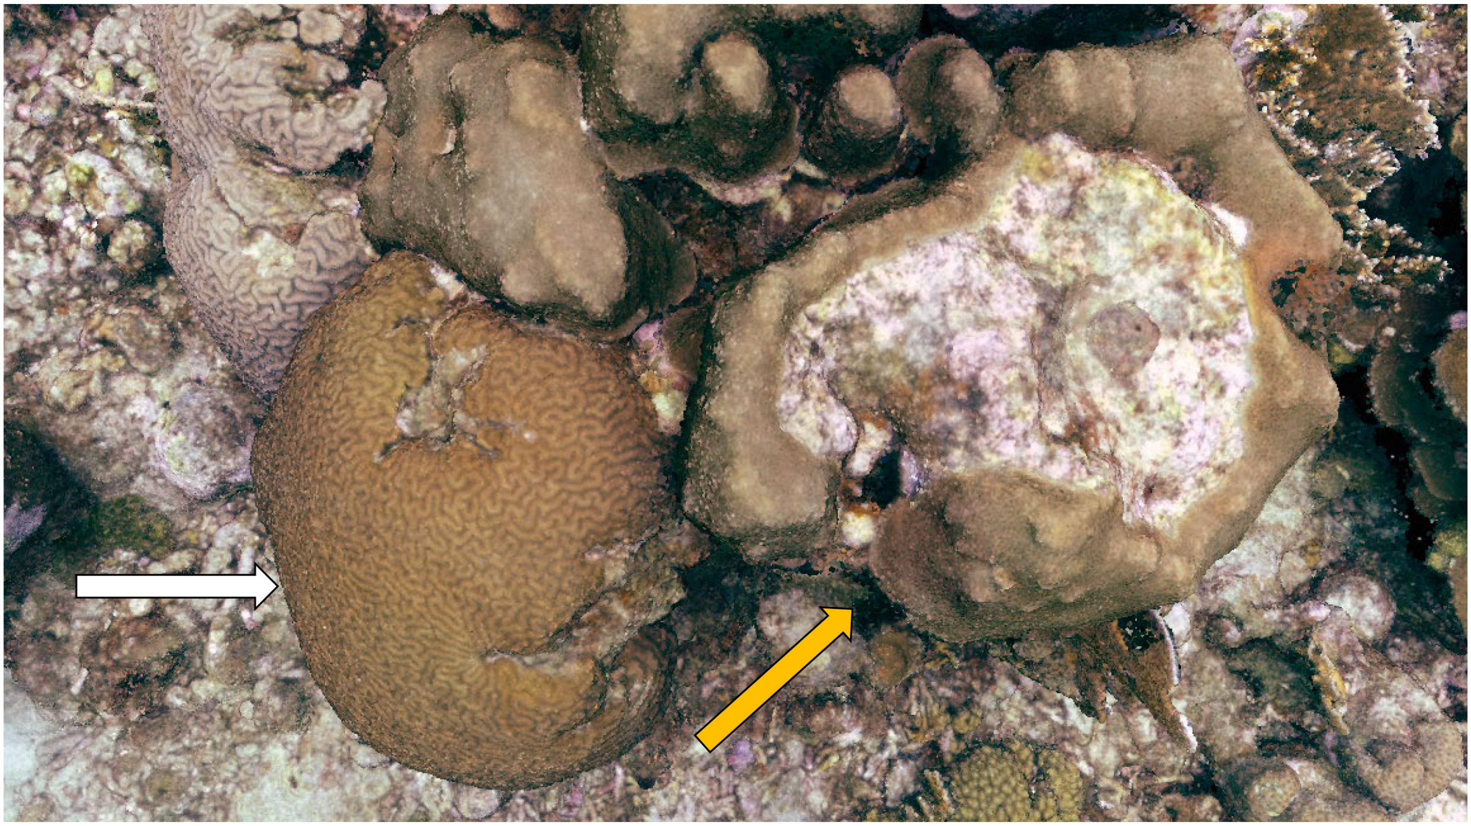

Figure 31. Simple vs. complex colony shapes. Corals with large corallites and simple growth forms, such as the *Colpophyllia* colony on the left (white arrow) are relatively straightforward to segment. Other groups, with smaller corallites and more complex fission and fusion patterns, such as the *Orbicella* colony on the right (yellow arrow), can require significantly more effort to accurately segment, including the need for more detailed raw imagery to confirm whether to combine or separate adjacent patches of live tissue.

Further, selection of minimum size thresholds and the abundance of corals in these smaller size classes can dramatically impact the time needed to extract demographic data. Similar to colonies with complex growth patterns, corals in smaller size classes can be difficult to accurately segment and require additional inspection of raw imagery. Other workflows that rely extensively on the use of raw imagery, such as abundance surveys for cryptic organisms, can require anywhere from 15 min/m^2 to 2 hr/m^2. When conducting any analysis, whether using LAI or any other method, it is important to remember that it is not necessary to sample comprehensively to robustly describe the ecological phenomena of interest; to this end, statistical power analyses are suggested. It is also recommended to conduct pilot data extraction to determine the sample size needed to capture the desired ecological signal with the necessary statistical robustness.

3.15. Data curation infrastructure

The last step in the pipeline is at once the most straightforward to conceptualize and the most difficult to execute. In the most basic sense, planning for data curation simply involves estimating the total amount and type of storage space that will be needed and securing the necessary infrastructure to accommodate those needs. However, data must also be well organized and accessible, both of which are requirements that become increasingly difficult when either the size of data collection or the number of people requiring access to them increases. As a result, simply buying additional external hard drives as storage demands increase is not a viable solution for high-volume projects. However, in the context of project planning, it is important to consider how activities at one step of the pipeline affect each of the other steps to balance resource investment across the pipeline (Figure 18). Image collection is the easiest step to increase in scale but has the greatest direct impact on data curation and storage needs as even if it is never processed or used for data extraction, raw imagery still needs to be organized and stored for future use. This can be a distinct challenge for groups new to the LAI approach who are unfamiliar with the substantial data volumes that will be generated. As described in Part I, section 2.4, there are a variety of approaches to maintaining well-structured and accessible data, and the challenge is to identify which approach best suits the scale of data that will be collected with the resources available. Data storage can be costly, and it is strongly recommended to estimate data needs and associated costs early in any project, as these demands often exceed the costs of other needed physical infrastructure.

3.16. Putting it all together

While plans and mandates should always be expected to change, the most successful implementations of LAI, whether for purposes of hypothesis testing, environmental monitoring, or coral restoration efforts, will balance clearly articulated project objectives with available logistical capacity (Figure 18). The most effective LAI project planning will begin by first identifying the derived data products and delivery timelines that will be needed. With this information in hand, decisions can then be made regarding the resolution needed in both the raw imagery and derived LAI products used to generate the desired data. The sampling design must also be described, including the temporal frequency of sampling and the number of replicate plots that will be needed, the size and shape of those plots, and their distribution within or across locations. Together, these decisions will dictate the volume of raw imagery that must be collected, which can then be used to guide decisions regarding the computational infrastructure needed to generate models at a rate that accommodates project timelines and, ultimately, the storage system that will be needed to store and share everything.

Depending on the question of interest, the level of investment along with the various steps of the LAI pipeline can vary considerably even for the same location. For example, a targeted restoration research project involving yearly sampling of a handful of long-term monitoring plots will require different levels of resources along the LAI pipeline as contrasted with a project comparing seasonal changes in structural complexity of those same plots. When considering investments in effort along the LAI pipeline, the first case will require more, and higher-resolution, imagery that will take longer to acquire within each plot, but the second will require higher frequency sampling and more overall time in the field. In the first case, more raw imagery means LAI will take longer to generate but only need to occur once per year. Data demands are, however, substantial. Human-driven image analysis efforts might continue throughout the year, in turn reducing the rate at which models need to be generated in the first place. In the second case, output metrics can be derived from LAI as soon as it is available, but many more models are collected per year, with model construction being continuous. Far less human effort will be required for data extraction in the second case, but more computational infrastructure will be needed to generate models in a timely fashion. Further, data curation activities will differ substantially between these two products with the first case

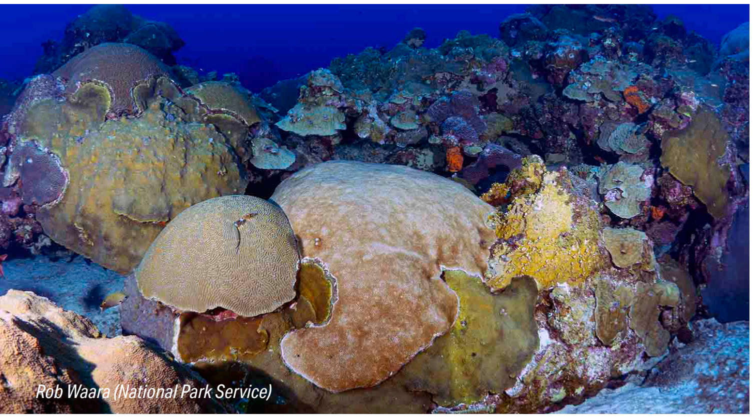
Rob Waara (National Park Service)


Samantha Clements (SIO, UC San Diego)

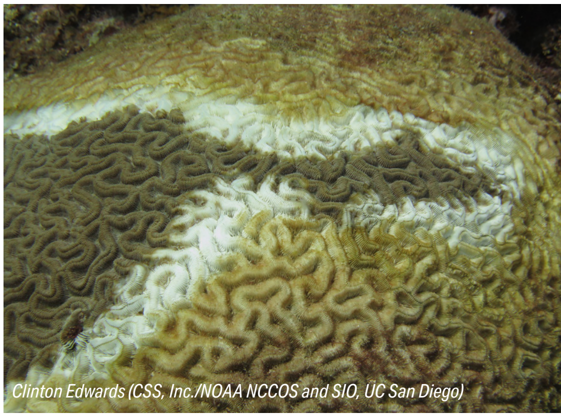
Clinton Edwards (CSS, Inc./NOAA NCCOS and SIO, UC San Diego)

requiring far less volume of total storage but more sophisticated database management due to the greater volumes of ecological data that will be extracted. Clearly articulating the needs for each of these cases will help optimize investment of resources along the steps of the LAI pipeline and thus maximize the capacity to meet the objectives of each project. While guidance is provided on the time needed to produce various data metrics, it is strongly recommended to conduct pilot testing of workflows to make sure that needs can be met before significant investments are made. Importantly, infrastructure should always be designed such that additional capacity can be added into existing infrastructure without the need to redesign from the ground up.

When designing each step of the pipeline, it is important to identify where capacity can be added to achieve the next set of project goals. If additional capacity can be added without jeopardizing success in the near term, then it is worth considering where it will have a maximum future impact. For instance, the price difference between a point-and-shoot camera and a high-end DSLR is on the order of $2,000–$3,000. However, the smaller camera will produce lower-quality LAI, and for some workflows, this might result in more time interrogating raw imagery, while LAI produced with the higher-quality camera might be of sufficient detail that raw imagery does not need to be inspected for many data streams. In this case, the initial savings on the cost of the camera is eventually lost to increased payroll costs associated with longer data extraction time. Further, there will be some operating conditions that the smaller form factor camera simply is not suited to, as well as analyses that cannot be conducted on the lower-quality imagery.

Thankfully, many aspects of the infrastructure needed at each step of the pipeline can be designed in such a manner that it is straightforward to add additional modules and increase capacity. However, awareness of the demands at each of these steps is recommended. Simply because it is possible to collect imagery over large spatial extents does not mean that it is the most appropriate approach to capturing the ecological signal of interest. Likewise, every pixel of an orthoprojection holds ecological data, but data need only be extracted from a fraction of those pixels to achieve project goals. It is recommended to consider how activities at one step of the pipeline affect each of the other steps to balance resource investment across the pipeline. Investments in both hard infrastructure and human effort are required at each step and must be made with respect to each other to build a functioning workflow.

The goal of this document is not to simply provide a series of step-by-step instructions but to promote a general understanding that will allow innovation and adaptation of the LAI approach as new settings and challenges arise. Part II of this document contains a series of example SOPs demonstrating how these concepts have been integrated into successful workflows in the past. Even for those groups who have contributed to these SOPs, there is no single set of instructions, nor is it possible to adequately capture the myriad variations that have been implemented to adapt to changing situations.


Sarah Rojano (CSS, Inc./NOAA NCCOS)

Part II:
Case Studies: LAI Standard Operating Procedures

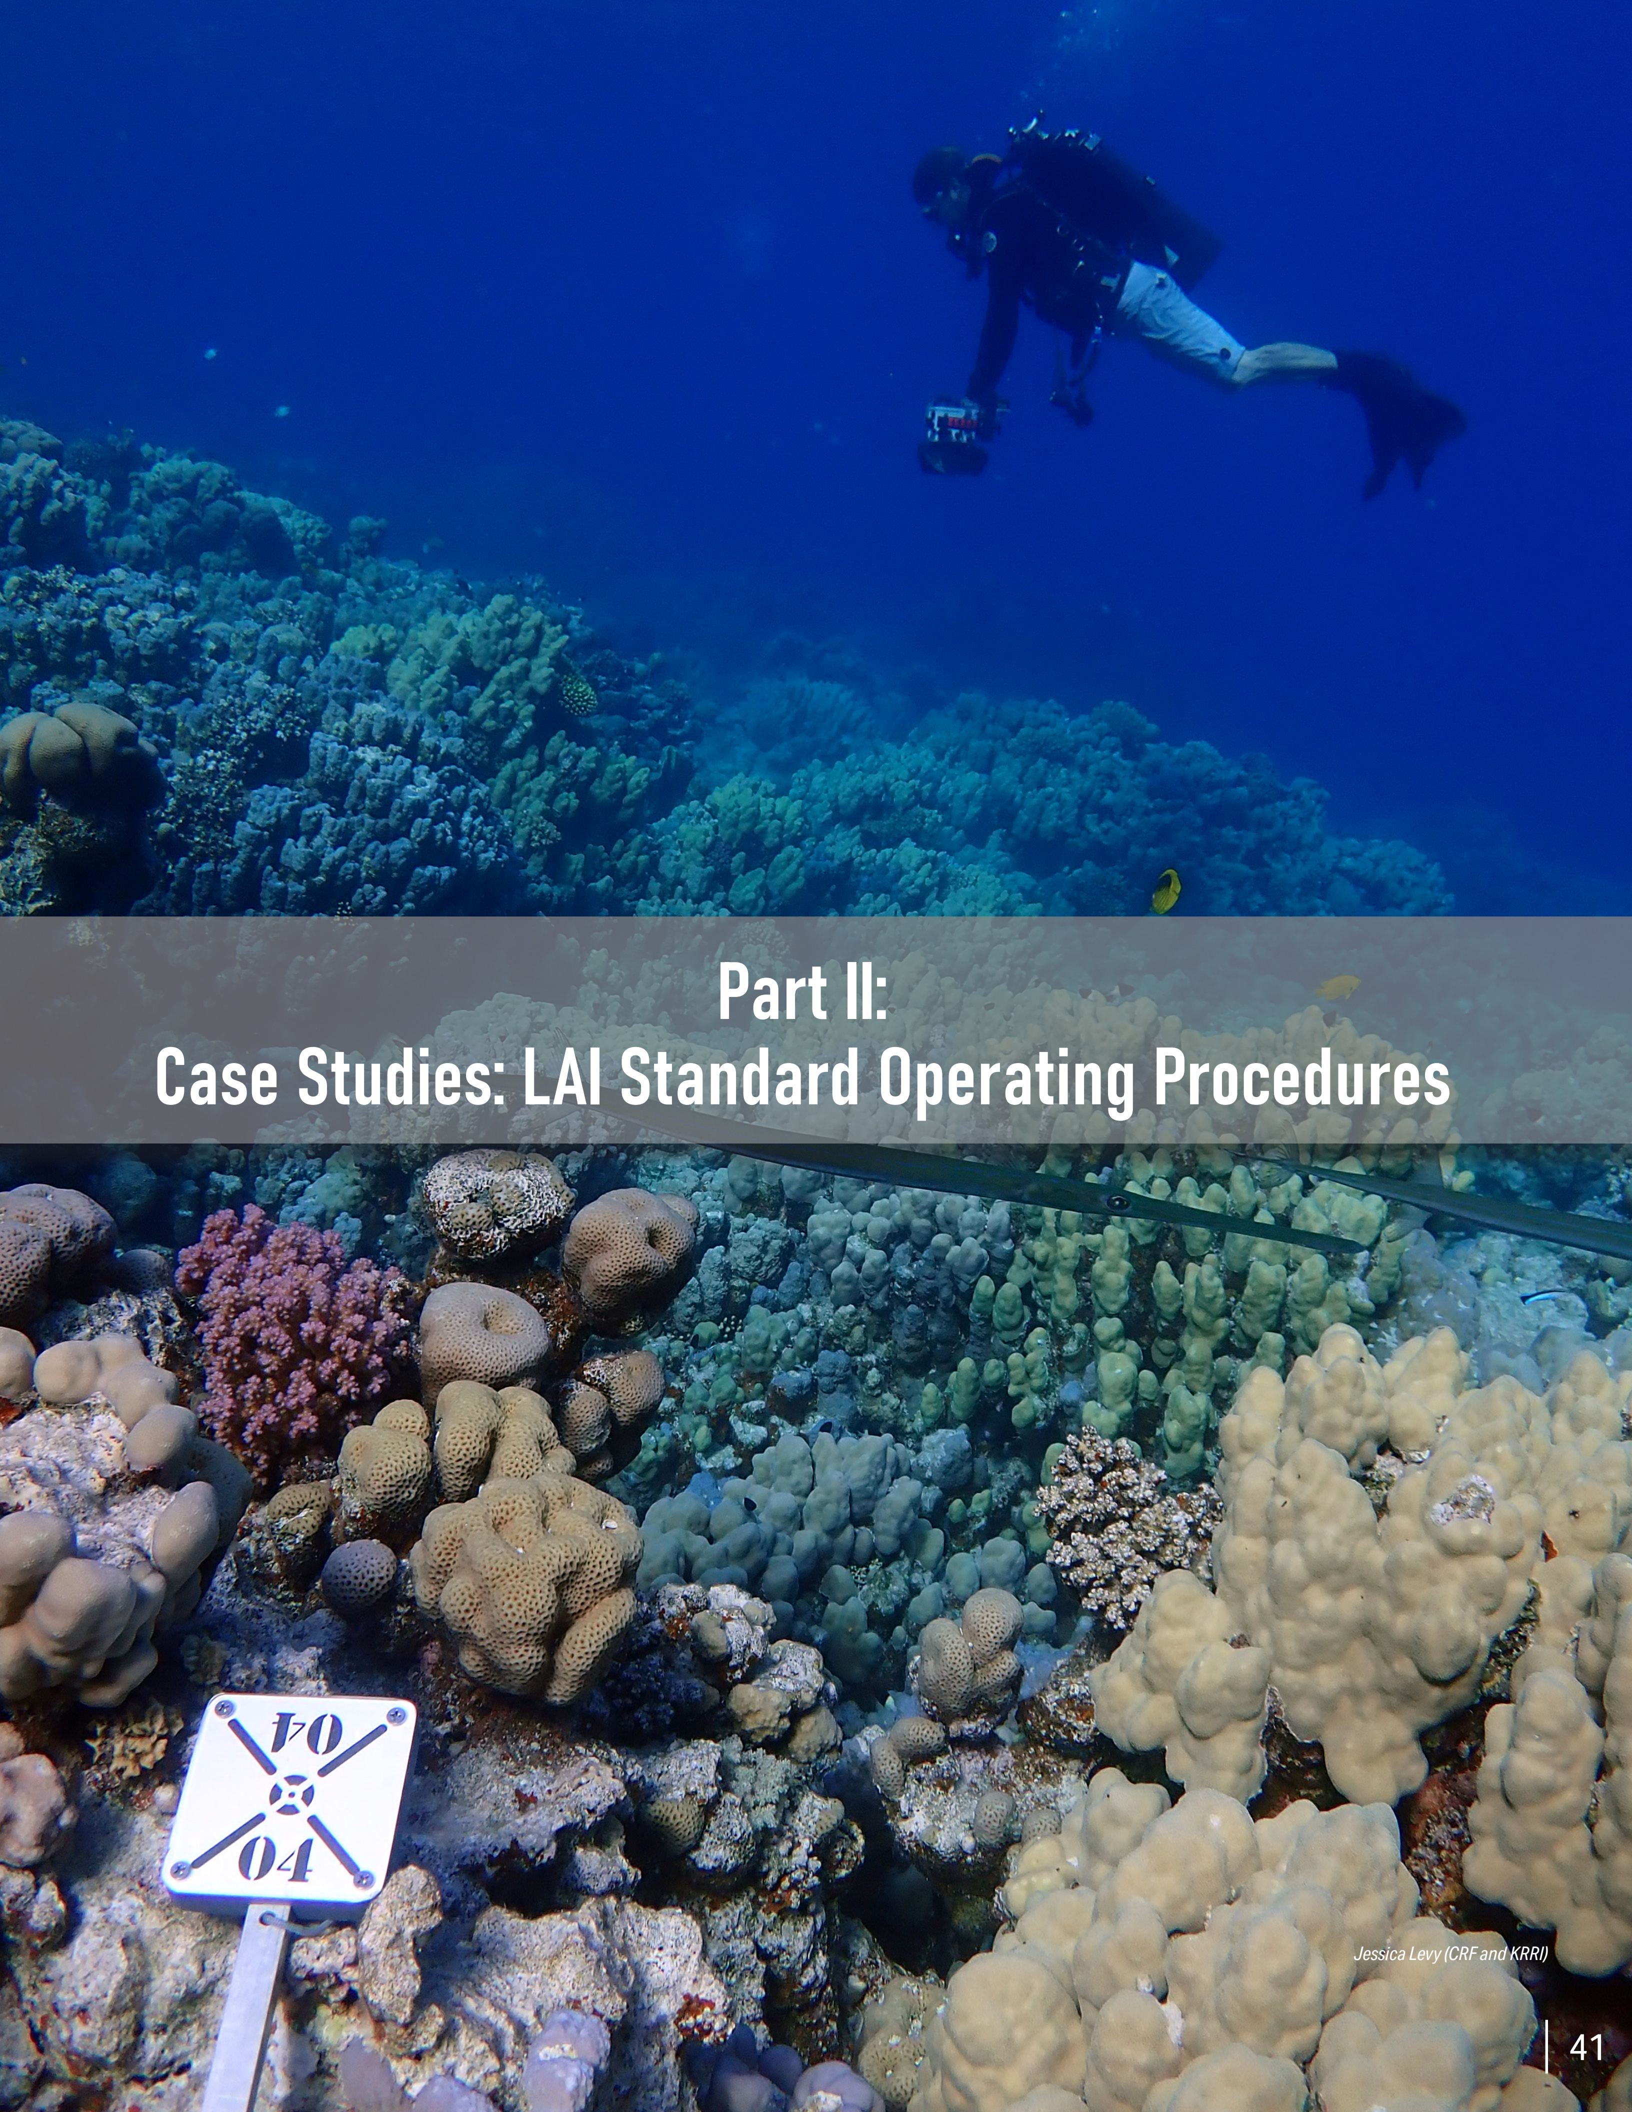

Jessica Levy (CRF and KRRI)

Case Studies: Overview

The goal of Part I of this guide has been to foster the understanding needed to conceptualize, design, and implement the LAI approach. To provide additional details needed to execute the LAI approach, Part II presents specific examples of how these concepts have been successfully implemented in real-world settings. These SOPs are used by a variety of independently operating groups, each with a different collection of expertise and project objectives. Each example includes a brief summary of the approach taken at each step of the pipeline and the rationale for the decisions that were made by each group. The full text of each SOP is provided as supplementary hyperlinks, and each is publicly available. Readers are encouraged to use information in these SOPs as needed to select the combination of tools and approaches that best fit their specific project objectives.

The groups included in this section have varying levels of experience and expertise with the LAI approach but share a track record of independently building and implementing an LAI workflow. However, while each has built their LAI pipeline independently, none of the projects have been done in isolation. Each of these groups represents individual nodes of a broader informal LAI network and has collaboratively worked in research settings and on a number of formal and informal training seminars. These collaborative efforts have made possible the articulation of the LAI pipeline, with the ultimate goal of encouraging and facilitating the expansion of this network and application of LAI in coral reef science. A brief summary of each program is provided below, with subsequent sections dedicated to each group.

Contributing organizations

Australian Institute of Marine Science (AIMS). AIMS has led coral reef monitoring efforts on the Great Barrier Reef for over 35 years. AIMS implements LAI as part of several programs, including the Long-Term Monitoring Program, the Ecological Intelligence for Reef Restoration and Adaptation Program, and ReefScan, across marine ecosystems in Australia and the Pacific. AIMS also leads or collaborates on developing technologies, such as ReefCloud and TagLab, to accelerate the uptake, analyses, and utility of LAI products for coral reef management and conservation.

Mote Marine Laboratory's Elizabeth Moore International Center for Coral Reef Research & Restoration. Operating North America's largest land-based coral nursery, Mote Marine Laboratory's International Center for Coral Reef Research & Restoration is a long-time leader in coral restoration throughout Florida's coral reefs. Mote utilizes LAI as a coral restoration and ecosystem monitoring tool to quantify the effect of restoration and potential ecosystem recovery.

Center for Environmental Imaging (CEI), LLC. CEI provides cutting-edge imaging and analytical tools to academic, governmental, and private clienteles. CEI's goal is to remove the financial and logistical barriers limiting the application of environmental imaging technologies and to help clients leverage image-based data to maximize the impact of their work.

National Center for Coastal and Ocean Science (NCCOS), National Oceanographic and Atmospheric Administration (NOAA). A long-time lead in benthic monitoring and remote sensing in Florida and U.S. Caribbean, NCCOS implements LAI to evaluate the ecological progress of coral restorations, such as for the Mission: Iconic Reefs restoration in the Florida Keys National Marine Sanctuary (FKNMS). NCCOS works closely with the Sandin and Kuester Labs at University of California (UC) San Diego on the implementation and development of LAI tools and workflows.

Pacific Islands Fisheries Science Center (PIFSC), NOAA. PIFSC is responsible for coral reef ecosystem monitoring across the U.S. Pacific Islands. PIFSC uses LAI to monitor coral community demographics as a part of the Pacific National Coral Reef Monitoring Program. PIFSC collaborates with the Sandin Lab at UC San Diego, Hawaii Institute of Marine Biology, and the University of Hawaii at Hilo to implement LAI-based monitoring.

Perry Institute of Marine Science (PIMS). As a small non-governmental organization (NGO), PIMS is focused on ecosystem monitoring, conservation, restoration, fisheries research, and education surrounding marine environments in The Bahamas and the Caribbean. PIMS has used LAI since 2013 for the monitoring of natural and restored reef systems and maintains a network of more than 150 long-term monitoring sites throughout The Bahamas. A principal aim is to develop methods to enhance the scalability, ease of implementation, and utility of LAI for organizations throughout the region. PIMS collaborates with numerous NGOs, research labs, and government agencies to assist in reef monitoring efforts.

Sandin Laboratory, Scripps Institution of Oceanography (SIO), UC San Diego. The Sandin Lab is a community ecology lab focusing on coral reef ecosystems, considering the status, trends, and interactions of fish and benthic taxa. The Sandin Lab has used LAI to study benthic coral reef ecosystems since 2012, using both geographic and time series sampling at a variety of locations throughout the nearshore tropics. Working closely with the Gleason Lab, University of Miami; the Kuester Lab, UC San Diego; and the Visual Computing Lab, Consiglio Nazionale delle Ricerche (Pisa, Italy), the Sandin Lab has contributed to the design and development of the software platforms Viscore and TagLab in applications of coral reef science.

Gleason Laboratory, University of Miami. The Gleason laboratory uses remotely sensed data to study coral reefs and related tropical ecosystems. The lab is a leader in the development of novel instrumentation and techniques with a focus on the use of underwater imagery for high-spatial resolution mapping and satellite vicarious calibration. LAI has been at the center of this work since 2003, beginning with the development and deployment of 2D image mosaicking techniques in collaboration with the Underwater Vision and Imaging Lab at University of Miami and the Underwater Vision Lab at the University of Girona.

Case Study 1: Australian Institute of Marine Science

1.1. Overview

AIMS uses LAI across multiple programs since 2016. The first program to implement it was the Long-Term Monitoring Program. This program uses LAI across plots of 20 m^2 to quantify small-scale structural complexity and relate it to different metrics of reef state and trajectory. AIMS also uses LAI through the Ecological Intelligence for Reef Restoration and Adaptation Program (EcoRRAP; https://gbrrestoration.org/program/ecorrap/), the ReefSong project (https://www.aims.gov.au/information-centre/news-and-stories/scientists-broadcast-reef-songs-underwater-replenish-reefs), and the ReefScan project (https://www.aims.gov.au/research/technology/reefscan), as well as multiple higher-degree student projects. The approach used by EcoRRAP is summarized here, and links to all SOPs are included at the end of this section.

EcoRRAP uses LAI to quantify structural complexity, benthic communities, and demographic rates of coral reefs across spatial and temporal scales. This program quantifies natural rates of ecological and genetic reef recovery and adaptation in response to acute and chronic disturbances, as well as key environmental variables related to different coral reef communities. EcoRRAP study sites encompass a range of environmental conditions within the Great Barrier Reef and Torres Strait, spanning latitudinal temperature gradients and cross-shelf gradients in water quality and wave exposure across >350 permanent plots (72-m^2 plots), nested within 88 (2,000-m^2) zones. Two key outputs are created from the images collected by EcoRRAP: 3D DEMs and 2D orthomosaics. The 3D DEMs are used to quantify landscape metrics, such as structural complexity and the demographic rates of complex coral morphologies (i.e., staghorn coral). The 2D orthomosaics are used to quantify benthic community composition and demographic rates of simpler coral morphologies (e.g., tabulate corals). This information is used to inform the Reef Restoration and Adaptation Program restoration interventions, the largest reef restoration program in the world as of 2020.

1.2. Image collection

Two photogrammetry techniques were used to describe differences in ecological information and spatial scales. Four EcoRRAP plots (72-m^2) are embedded within zone-scale plots (approximately 2,000 m^2 per zone). At the scale of EcoRRAP plots, DSLR cameras and associated in-water imaging techniques were used to construct high-resolution outputs (0.3 mm per pixel in orthomosaics) of reef areas. In contrast, medium-resolution action cameras and less complex in-water imaging techniques were employed to image reefs at the zone scale. Details of imaging techniques are provided in Table 2 and Figure 32. A variety of GCPs (Table 3; Figure 32) were deployed to scale and orient models and allow temporal coregistration of models between annual sampling events.

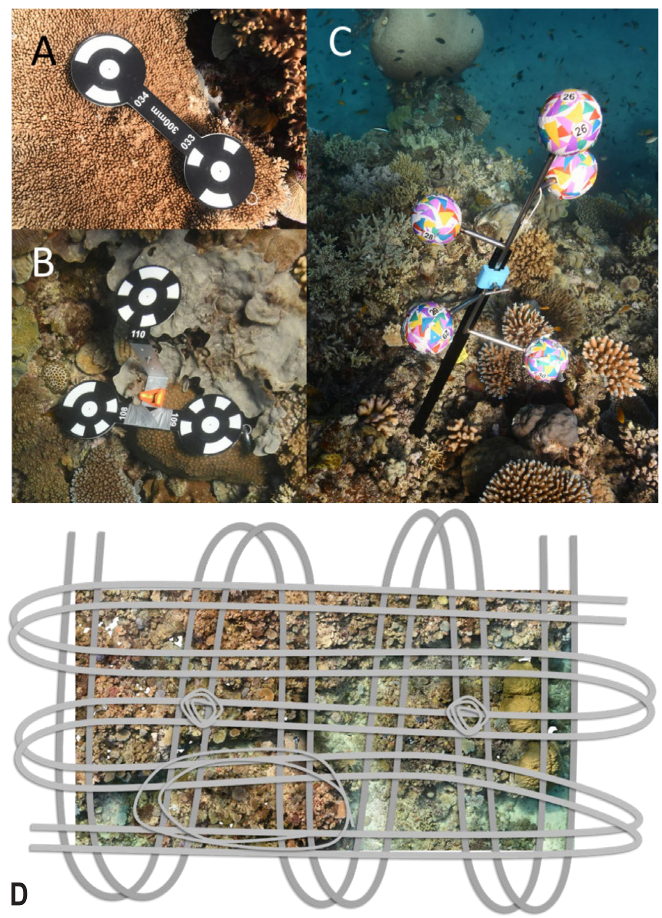

Figure 32. AIMS scale bars and swim pattern used for image collection. Photogrammetry ground control points used in the current workflow. (A) dumbbells, (B) triads, (C) sphere tree attached to permanent star picket on reef, and (D) depiction of flight path over high-resolution plot area. Image: Stephanie Gordon (AIMS)

Table 2. EcoRRAP photogrammetry imagery techniques.

Scale	Cameras	Camera spacing	Flying altitude	No. photos collected	Swim pattern
12 m × 6 m "Plot"	2 × Nikon D850 DSLR in Nauticam housing, 20-mm prime lens, 8-in dome port	57 cm	1.0–1.5 m	1,700–2,700	• Flight path: ▪ 5 longitudinal passes in-line with plot ▪ 6 passes perpendicular to the plot ▪ "Spirals" or "slices" around sphere trees mounted on permanent markers ▪ Additional passes made around complex structures • Horizontal passes are used to capture nadiral images (relative to reef), while perpendicular passes aim to capture 3D oblique imagery • Passes should extend approximately 2 m past the plot extent of interest (e.g., 14 m × 8 m imaged)
100 m × 15 m "Zone"	3 × GoPro Hero in GoPro housing	90 cm	1.5–3.0 m	5,000-7,000	• Divers swim two horizontal passes in-line with transect spaced side-by-side to ensure sufficient overlap between cameras • Imagery is nadiral relative to the reef and extends approximately 2 m past the zone extent

Table 3. Ground control point (GCP) types used in the current workflow and their applications.

GPC name	Construction	Application
Dumbbell	Flat aluminum shape consisting of two 12-bit markers spaced at a known distance	• Enables 2D and 3D model scaling
Triad	Standing aluminum shape consisting of three 12-bit markers spaced at known distances, with a bi-directional bubble level mounted on the base	• Enables 2D and 3D model scaling • Provides orientation relative to gravity
Sphere tree	Five spherical, stainless-steel shapes with patterned decals attached to a branched central pole. Sphere trees are temporarily attached to permanently deployed star pickets	• Spatial reference for coregistration of 3D models between years • GPS coordinates of every picket • 2 pickets per plot

1.3. Model construction

Models are constructed using Agisoft Metashape Professional (v. 1.8.), following the steps described below in Table 4. Additional coregistration techniques using CloudCompare v2.13 (GLP software, 2022) are used prior to final orthomosaic and textured DEM generation. Temporal coregistration is achieved with centimeter-level precision by using a combination of permanent markers and temporary spherical targets used during imaging surveys. These targets can be used to coregister 2 or more models collected at different time points by using the point/sphere-picking tool in CloudCompare. Specifically, at each permanent plot, two sphere trees were attached to permanent pins hammered into the reef substrate along the central line of each plot and roughly 6 m apart. Each sphere on the sphere trees is used as a coregistration target. The trees and clamps are designed to ensure only one position was possible during tree setup between surveys. Final reference DEMs are oriented using a cross check level in the field (Figure 32B), and subsequent DEMs are coregistered to reference DEMs by using the sphere center pair picking function in CloudCompare. Once DEMs are coregistered, the transformation matrix of the registered DEM is exported and applied to the original model in Agisoft Metashape Pro. Finally, the orthomosaic is produced for the registered model (Lechene et al., in prep). Chain 1 is run in the field, mostly for data quality and assurance purposes; chains 2 and 3 are run back on land using AIMS HPC.

Table 4. EcoRRAP model processing steps.

Chain no.	Step no.	Job name
1	1	Quality check
	2	Lowest quality alignment
	3	Detect markers
	4	Add and check scale bars
	5	Import depths
	6	Generate log
2	1	High-quality alignment
	2	Calculate scale bar error
	3	Resize bounding box
	4	Duplicate chunk
3	1	Crop point cloud
	2	Initial camera optimization
	3	Filter and re-optimize cameras
	4	Build depth maps
	5	Build mesh

1.4. Ecological data extraction

Once final DEM and Orthomosaics are produced, the EcoRRAP workflow splits them for data extraction. Three main datasets are routinely extracted from EcoRRAP's LAI as well as other LAI at AIMS: 1- habitat structural complexity metrics; 2- coral demographics, size frequency distributions, and seascape ecology metrics (i.e., cluster analyses); and 3- benthic cover, community composition, and structure data. Below, each is briefly commented on.

1.4.1. Habitat structural complexity

An in-house set of scripts is applied to all DEMs to automatically extract up to 50 different metrics of structural complexity, such as surface rugosity and slope (Friedman et al., 2012). These metrics are systematically saved and backed up as CSV files.

1.4.2. Hard coral colony vital rates and seascape ecology

The EcoRRAP orthomosaics are loaded into TagLab, where the predominant species of hard coral taxa are digitized and quantified to morphological-taxon level (e.g., tabulate *Acropora* spp.), depending on the difficulty to identify species reliably from the images. The raw images are also used for classification, and a trained taxonomist from the North Queensland Museum visited all plots during 2022 (second year of survey) for ID validation. The ID validation assessment was used to fine-tune which taxa could be reliably identified from orthomosaics and images and to which taxonomic resolution. Once enough colonies were manually digitized in TagLab, automated classifiers were trained and applied to the remaining orthomosaics. These colonies are then segmented and identified on orthomosaics every year.

1.4.3. Benthic cover, community composition, and structure

The EcoRRAP images are classified using ReefCloud (Gonzalez-Rivero et al., 2020) by a team of trained benthic ecologists using the Collaborative and Annotation Tools for Analysis of Marine Imagery and video (CATAMI) (Althaus et al., 2015) classification scheme to the highest possible taxonomic resolution, which varies widely across benthic groups (e.g., hard corals vs. algae). ReefCloud is an online platform to upload images into the cloud, which allows the user to manually classify points on images, and it then uses these points to automatically classify other pixels in an image. This software platform can replicate expert observations with 80%–90% confidence to produce accurate estimates of coral reef benthic composition 700 times faster than manual assessment. The labels on the images are then projected onto the orthomosaics and DEMs to calculate benthic community composition and structure, as well as benthic cover estimates.


Erich Bartels (Mote Marine Lab)

1.5. Data curation

A Microsoft Access database is used to record all metadata in the field. The database records depth (measured in situ) for every GCP, GPS coordinates for every picket (on initial deployment), time of day, environmental conditions (e.g., tide), and other important metadata (see detailed SOP for details). The same database is used to record initial processing settings (i.e., percentage of images successfully aligned after chain 1 is processed). Data are backed up on a daily basis while in the field.

All data generated through the project are available through AIMS metadata records, and final 3D outputs are available through SketchFab. All associated scripts are available through GitHub.

1.6. Links to detailed SOPs

Field Photogrammetry in 4D SOP (Gordon, 2023):
https://doi.org/10.25845/SE7T-PS86

Reef monitoring sampling methods | AIMS | https://www.aims.gov.au/research-topics/monitoring-and-discovery/monitoring-great-barrier-reef/reef-monitoring-sampling-methods#SOPreefmonitoring

Long-term monitoring habitat reconstruction SOP (González-Rivero, 2020):
https://www.aims.gov.au/sites/default/files/2021-07/20210323_LTMP%20SOP12%20-%20Habitat%203D-Final_AIMS.pdf

Examples of 3D models:
https://sketchfab.com/AIMSEcoRRAP

AIMS metadata records:
Ecological Intelligence for Reef | Restoration and Adaptation Project (EcoRRAP) | AIMS metadata | https://apps.aims.gov.au/metadata/view/fe9659f1-12d6-4acf-ab89-67acdd37efe5

Case Study 2:
Mote Marine Laboratory

2.1. Overview

Mote Marine Laboratory leverages SfM-based LAI to answer a variety of scientific questions regarding Mote's coral restoration efforts throughout the Florida Keys region of Florida's Coral Reef. Mote employs two general strategies for generating LAI. The first of these strategies is a more focused, higher-resolution image capture and data extraction workflow aimed at answering specific research questions. This workflow is adapted directly from the methods outlined by the Sandin Laboratory in this document (Sandin Lab, 2023). The second strategy complements in situ coral restoration monitoring efforts and is aimed at generating ecologically relevant data such as percent cover, outplant fusion/fragmentation, and growth rate data that are unattainable with traditional in situ monitoring. The data generated from this strategy directly inform Mote's outplanting and restoration efforts and are detailed within this document.

2.2. Image collection

Instructions are provided for the configuration of the GoPro imaging system, in situ site configuration, and image capture. Image capture for outplant monitoring uses two GoPro Hero 10 cameras on a polyvinyl chloride (PVC) frame, although any newer GoPro model will produce good reconstructions (Hero 8 Black or newer). Cameras are spaced roughly 0.6 m apart and secured to the PVC frame using GoPro's handlebar mounts. This maximizes the capture area, thus minimizing the amount of swimming needed on these large outplant events where thousands of corals are being outplanted in areas often exceeding 200 m^2. These cameras were chosen for this specific strategy because they are lightweight, durable, and easy to use. This is a considerable advantage over other camera systems because image capture occurs concomitantly with other restoration activities, such as outplanting, site prep, and in situ monitoring, where team members are conducting multiple duties and objectives on a single dive.

Site selection and preparation are species dependent but generally fall into two larger categories: *Acropora cervicornis* and *Acropora palmata* together, and boulder species. Differences in site setup between the two outplant strategies are explained in greater detail within the SOP; however, equipment needs, metadata collection, and navigation all follow similar specifications. Temporary floats are placed throughout the scene where no discernible landmarks exist (e.g., ledges or ends of spurs), and scale bars with Metashape markers are placed regularly throughout the plot. The cameras are set to the "Time Lapse" setting and capture 1 image/s. The diver swims roughly 1.5 m from the benthos in a double-lawnmower pattern. Outplant sites are typically rectangular, and the diver swims many passes with a high degree of overlap in the "short" direction and fewer passes in the "long" direction. Image capture usually takes 30–40 min, and images range from 500–2,000 per camera depending on the size and complexity of the outplant location.

2.3. Model construction

LAI products are constructed using Agisoft Metashape Pro. Many of the steps can be performed using the Standard Edition, and it is noted within the SOP where the Professional Edition is explicitly required. Products are constructed on a custom-built PC with 32-core CPU, 128 GB of RAM, and an NVIDIA GeForce RTX 2080Ti GPU. Briefly, the processing steps used to generate 3D DPCs are marker detection, image alignment, and DPC generation, with the additional steps of building a triangulated mesh and orthomosaic generation for 2D orthomosaics. Minimal QAQC occurs between steps; the most notable being initial image QAQC and the filtering of low-confidence points after the DPC is rendered. The DPCs are exported as a polygon file format (PLY) for downstream ecological analysis.

2.4. Ecological data extraction

This SfM pipeline was developed specifically to bolster outplant monitoring capabilities to support Mote's coral restoration efforts throughout Florida's Coral Reef. As such, the data of particular interest include more spatial metrics like percent cover, change in area (growth), and fragmentation/fusion. The most efficient, accurate, and reliable software to use for these goals is TagLab. TagLab is AI-powered segmentation software designed to support the analysis of large orthographic images generated through the photogrammetric pipeline. Orthoprojections are rendered using the custom software package Viscore, and those workflows can be found in Part II, section 7.6 of this document. Once imported, outplanted colonies are segmented where the aforementioned spatial metrics are extracted. Additionally, with the customizable classification dictionary within TagLab, outplants can be segmented and classified as specifically as genotype or as broadly as outplant species, depending on the monitoring needs.

2.5. Data curation

SfM photogrammetry and LAI is still in its nascent stage at Mote Marine Laboratory. As such, many workflows and SOPs are still under development and are regularly updated. Currently, the thousands of images, Metashape projects, and downstream products are stored on NAS devices that can be accessed via the local network or through attaching external drives. SSDs are used and recommended for all data transferring as they are much faster, allowing work to be conducted directly from the drive without needing to download the products locally. Data management is supported by using thoughtfully curated, consistent naming schemes and file hierarchy. Most data generated from the data extraction protocols are also stored via a GitHub repository, including the most up-to-date SOPs.

2.6. Links to detailed SOPs

The SOPs for Image collection, model construction, and data extraction can be found on **GitHub:**
https://github.com/Mote-Coral-Reef-Restoration/MoteSOPs

Case Study 3:

Center for Environmental Imaging

3.1. Overview

The SOP used by CEI builds upon the methods used by the Sandin Laboratory at Scripps Institution of Oceanography, UC San Diego, with modifications made to create digital twins of increasingly large areal extents. As the field of LAI progresses, it has become evident that there is a need to expand the capabilities of large-area imagery to scales greater than a traditional survey area of 100 m^2. There are a number of challenges associated with scaling up areas of operation that span all steps of the large-area imagery pipeline. The goal of this SOP is to address some of these challenges, building from insights gained through current efforts in the field.

3.2. Image collection

The field approach was designed to support the collection of high-resolution, large-area imagery within 2,500 m^2 from shallow-water (<30-m depth) coral reef habitats. The protocol is designed to capture a relatively large area while meeting two design constraints: 1 – to complete surveys in approximately half of one diving day (2 or 3 dives) and 2 – to use commercially available cameras, cyberinfrastructure, and software approaches to create LAI products without the need for subdivision of the imagery collection (<80,000 images). The approach has been designed to minimize the time needed to collect imagery over large areas, while providing additional metadata needed to optimize model construction. The approach uses a Nikon D780 camera, selected for quality of imagery, extended battery life, and a built-in intervalometer. In order to minimize dive time, the approach employs multiple divers simultaneously operating a three-camera array and mounted to a DPV.


Brian Zgliczynski (SIO, UC San Diego and CEI)

For all large-area imagery projects, regardless of scale, it is imperative that images are collected to ensure 60%–80% overlap between images. Cameras are mounted to a sturdy aluminum bar at a spacing (1 m) designed to provide a minimum of 60% overlap when operated 2 m from the bottom. In order to minimize diver fatigue and expand operational flexibility (e.g., in cases of current or increased depth), the array is mounted to a DPV with a custom camera bracket. The use of DPVs or towed systems is not meant to speed up the collection of imagery along a given pass, but rather to facilitate the use of larger arrays with multiple cameras to increase the area imaged per pass. On average, it takes a team of three divers operating a three-camera array approximately 80 min to complete the imaging, with about 4,000–5,000 images collected per camera during the survey period.

Navigation, and in particular maintaining consistent spacing along extended parallel passes, is challenging across longer distances. To aid navigation, the 2,500-m^2 plot is divided into two subsections, each measuring 25 m × 50 m and delineated with transect tapes. Weighted floats are placed at regular intervals to provide visual reference, and an underwater compass is mounted to each camera array. Divers operate in a staggered manner, maintaining a minimum separation to maintain necessary side-to-side overlap between adjacent arrays, and staggered along each pass to avoid the diver being captured in imagery (Figure 33). Additionally, approaches to providing real-time position information to divers via USBL and a heads-up display are discussed.

3.3. Model construction

Similar to smaller plot areas (e.g., 100 m^2), models are constructed using Agisoft Metashape Professional Edition using a networked cluster of custom-built HPCs or other supercomputing resources. The configuration of each build varies slightly, but each has an 8- to 16-core CPU, one to three high-end gaming GPUs and 128–256 GB of RAM. Processing a single operational unit (2,500 m^2) of 35,000–45,000 images requires about eight days using a single HPC cluster node. This computer configuration can process a maximum of about 80,000 images per operational unit. This number can increase with the addition of camera position information linked to each image (e.g., GPS or USBL) or an increase in the memory allocation on each HPC. Processing time will also decrease with advances in computing resources and improvements in Metashape. A robust image collection procedure with high image overlap has led to minimalistic settings to be used during processing, which balances high-quality and efficient processing. Scale bars and depth information for markers placed throughout the plot are referenced in Metashape to add an additional alignment optimization step to reduce the impact of errors that accumulate over larger areas. Camera pose estimates are exported as an XML file along with a JSON file, which references image file paths. The DPC is exported as a PLY for subsequent visualization and analysis. Additionally, a georeferenced orthomosaic is generated for visualization purposes using an intermediate step of a depth map–derived mesh.

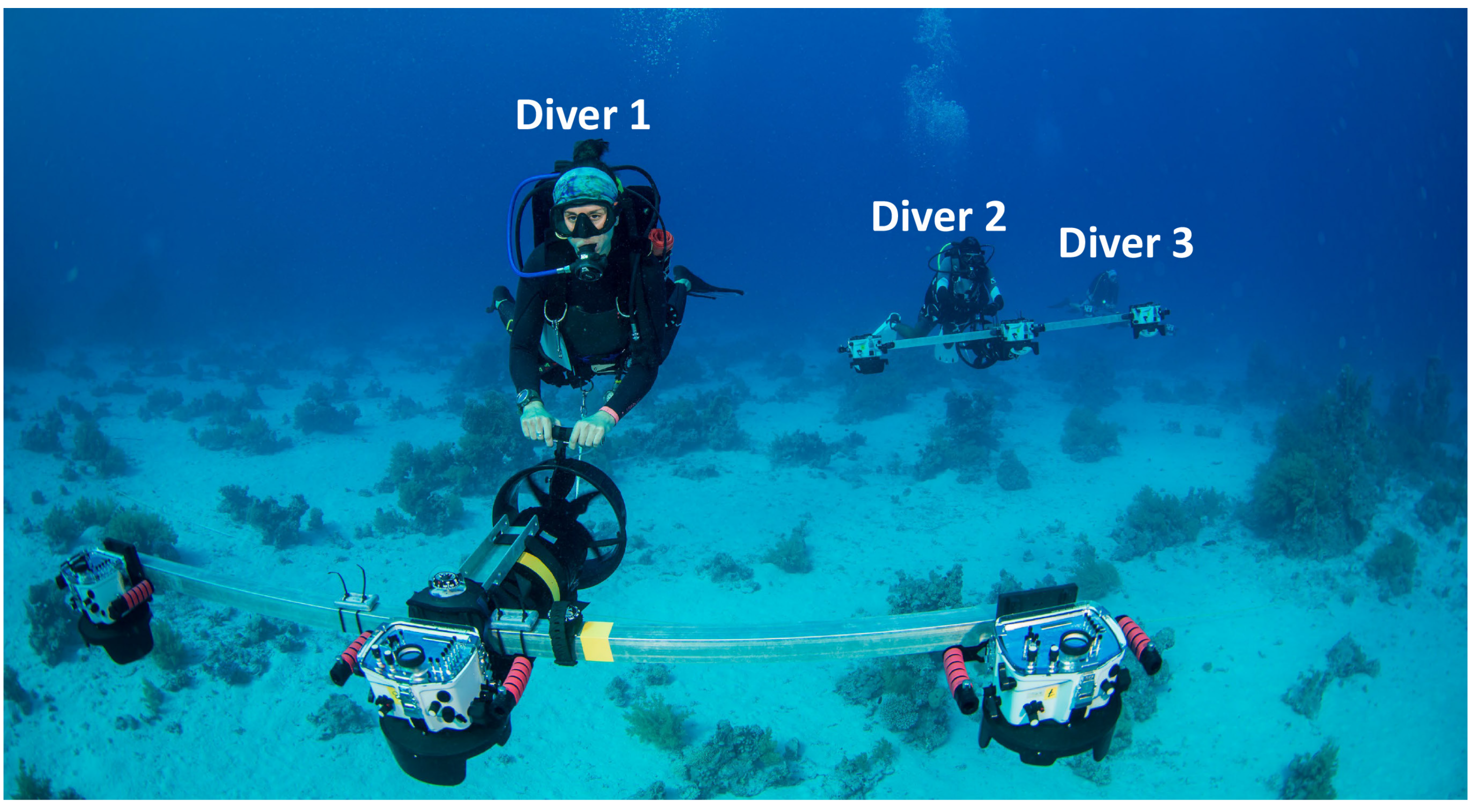

Figure 33. Example of staggered diver formation. Image of a team of divers collecting imagery using DPVs. Three divers operate in parallel, each using a three-camera array. Cameras are separated by 1 m, and divers are similarly offset from each other by 1 m in order to maintain consistent 1-m spacing between each of the nine cameras. Divers conduct passes in a staggered manner to avoid capturing the adjacent diver in imagery, particularly when adjusting depth over reef slopes or topographically complex features. Navigation is accomplished using an array of visual markers and a high-quality underwater compass mounted to the center camera. Credit: Brian Zgliczynski (SIO, UC San Diego and Center for Environmental Imaging)

3.4. Ecological data extraction

The processing of increased spatial scale plots (>2,500 m^2) follows the same workflow as more traditionally sized plots (approximately 100 m^2). Image collection is optimized for the collection of high-resolution LAI necessary to extract precise colony size and shape information, with taxonomic identification to the species or genus level. Multiple data extraction workflows can be used for considering ecological patterns, with few functional differences in post-processing methodologies.

However, given the increasingly large volume of data available from surveys of expansive areas, two important considerations are noted for ecological data extraction. First, for many ecological questions, such imagery datasets may be best explored through subsampling. To answer particular questions, it may be more time- and cost-efficient to create a sampling protocol for analysis that is comparable to field-based efforts to gather sufficient data to reach levels of statistical power commensurate with the goals of the study. Instead, if there is interest in comprehensive quantification of all features within the surveyed area, the potential of computer-assisted efforts of ecological data extraction are highlighted here. By focusing investments on the development of machine learning and other computer-assisted protocols, there are means to accelerate data extraction procedures. Note that it is recommended to maintain some human oversight within any such applications of computer-assisted analyses, with so-called "human-in-the-loop" protocols offering the means to error check post-processing protocols to maintain data accuracy and precision.

3.5. Data curation

The data volumes associated with scaling up LAI projects can easily move into the realm of petabytes (PB, 1,000 terabytes). LAI efforts will continue to grow through time as methods and technology improve. Traditional approaches of storing data using external hard drives will not support the data volumes of these larger-scale efforts. It is now common for large-scale efforts to rely on enterprise storage systems managed by university facilities or supercomputing centers as storage needs can exceed 1–3 PB/year. Strict file naming protocols as well as extensive quality control procedures must be followed to manage these extensive data volumes and associated metadata. Raw data layers, including the collected images and associated metadata are stored in duplicate on separate storage systems to protect against data loss of irreplaceable files. Efforts are being made to upload these raw data layers to stable and open digital collections for permanent archival and accessibility to the general public.

3.6. Link to detailed SOPs

Detailed SOPs from the Center for Environmental Imaging are available online at the **UC San Diego Library Digital Collections (Sandin Lab, 2023):** https://library.ucsd.edu/dc/collection/bb43111847

Case Study 4:

National Centers for Coastal Ocean Science, NOAA

4.1. Overview

NCCOS applies LAI in a variety of settings, including diver-based approaches, shallow-water AUV surveys, and mid-depth surface-operated drop-camera approaches. The specific SOPs provided here have been developed in support of evaluating coral restorations, such as Mission: Iconic Reefs (M:IR; https://www.fisheries.noaa.gov/southeast/habitat-conservation/restoring-seven-iconic-reefs-mission-recover-coral-reefs-florida-keys), a large-scale restoration project currently underway at seven reefs in the Florida Key National Marine Sanctuary (FKNMS). NCCOS has developed a series of SOPs for the creation and use of LAI for the benthic monitoring and restoration evaluation and has applied them extensively to support M:IR in collaboration with FKNMS. M:IR is a long-term restoration project, and workflows have been optimized to support multi-temporal comparisons as well as the multi-geography nature of the project.

Laughlin Siceloff (CSS, Inc./NOAA NCCOS

4.2. Image collection

The SOP provided here has been used since December 2021 by NCCOS and FKNMS staff to collect imagery in the field for the generation of LAI within 100-m^2 plots. The seven reefs included in M:IR are distributed widely along the Florida Keys reef tract, covering a range of depths (2–15 m) and environmental exposures, including both relatively protected back reefs and exposed oceanic fore-reef locations. As a result, the image collection approach is designed to be robust and rapidly implemented. The selection of camera type was based on the desire for an easily replaceable and commercially available camera system, requiring no aftermarket or otherwise custom modifications. Further, it was necessary that the camera be available with a high-capacity battery capable of several hours of operation in order to prevent the need to swap batteries while working from surface vessels. A built-in intervalometer was also required. The Nikon D780 was selected as it features the desired combination of settings, has economical housings available, and has been thoroughly tested by other groups implementing LAI at scale (Part II, section 7.6).

Ben Edmunds (NOAA ONMS/Florida Keys National Marine Sanctuary)

During field operations, the camera is mounted in a sturdy high-density plastic frame, designed to protect the camera and also to provide a stable easy-to-operate platform during imaging. The frame is designed to be slightly negatively buoyant (approximately 1 lb) with cameras mounted and comfortable to operate. Several attachment points are provided for mounting navigation instruments, such as a compass or dive computer. The frame is also useful to protect the camera during transport on surface vessels, provides protection underwater during plot setup, and also allows divers to complete other tasks. Of note, the camera and housing setup used here are themselves buoyant, and care is taken when removing the camera from the frame underwater.

A plot size of 100 m^2 is used to accommodate both baseline and continued ecosystem monitoring for the M:IR project, as well as to enable tracking of coral restoration activities. Conducting imaging within 100-m^2 plots allows rapid collection of baseline data, including robust descriptions of percent cover and extraction of percent cover data. Further, plots of this size allow assessments of short-term survivorship of outplanted corals, as well as initial growth of restored corals. However, to address restoration questions at the reefscape scale, larger plot sizes will be used to track the impact of restoration activities. NCCOS is currently field-testing image acquisition using AUVs to support these expanded spatial efforts into the future. As these approaches mature, additional SOPs and method documents will become available, and readers are encouraged to visit the NCCOS website or contact the authors for the most up-to-date information.

Plots are established with two permanent steel pins, and GPS positions are taken from the surface at each pin. Four weighted scale bars (45 cm) are placed throughout the plot, along with six weighted tiles at which depth measurements are collected with a handheld depth gauge. Divers navigate the plot in a gridded pattern with the assistance of a compass and visual markers, including weighted floats placed at each corner of the plot along with scale bars and depth tiles that are deployed around the plot. A compass and depth gauge are also used for navigation, the former which is particularly valuable during low-visibility conditions. A buffer of 2 or 3 passes (approximately 1–2 m) around the plot is imaged to ensure that the core study area is well represented in the final LAI products.

The imaging approach used here is adapted from the SOP used by the Sandin Lab, UC San Diego (see Part II, section 7). However, as only a single camera equipped with a relatively wide-angle lens is used, imagery is collected closer to the benthos in order to maintain the image detail needed for high levels of taxonomic specificity during ecological analysis. The camera is typically operated 1.0–1.5 m from the bottom, and 2,500–3,500 images are collected during a single dive, allowing sufficient time for plot setup and breakdown to occur. Prior to imaging, cameras are white balanced at the median depth of the plot using a gray card, and images are collected as JPEGs.

4.3. Model construction

3D models and other LAI products are created with Agisoft Metashape Professional Edition using Microsoft AVD. With AVD, users have access to a wide variety of available virtual computer configurations that can be tailored to meet specific project needs. Using AVD prevents the need to purchase and maintain hardware, and users pay only for active use time of the virtual machine. The total number of 100-m^2 plots associated with the M:IR project is currently 132, with anywhere from 0–20 plots imaged per month. It is critical to provide rapid feedback to the field image collection team regarding whether or not the area had been sufficiently imaged, so a relatively high-powered virtual machine was configured. Processing is conducted on a single virtual machine consisting of a 48-core CPU, 448 GB RAM, and 4 NVIDIA Tesla M60 GPU's with 32 GB of video RAM. The general processing workflow includes construction of an SPC, optimization and error-reduction steps, automatic marker detection, manual entry of scale bars and depth measurements, along with GPS information, and finally, construction of a DPC, orthophotomosaic, and DEM. Model construction generally takes 3–5 days, and most steps of the workflow are accomplished using Python scripting; however, as mentioned, some steps require limited manual intervention. Camera pose estimates along with the DPC are exported for later post-processing steps using the custom visual-analytical software Viscore.

4.4. Ecological data extraction

Data extraction activities are designed to support the monitoring and restoration objectives of the M:IR project. Image collection and approaches to 3D model processing were thus designed to provide highly detailed raw imagery and LAI, allowing for extraction of highly specific taxonomic information and precise estimates of coral size and growth. The primary data streams extracted in support of these activities are estimates of percent benthic cover, structural complexity, and coral colony size data, including tracking of coral colonies for demographic data. The custom software platform Viscore is used to generate percent cover and structural complexity data, coregister time series LAI and export orthomosaics. Coral size estimation, as well as tracking of colonies and generation of demographic data, is conducted in the program TagLab using (coregistered) orthoprojections exported from Viscore. Of note, rather than using orthophotomosaics generated in Metashape to create 2D map views needed for generating coral size data, the DPC itself is directly orthoprojected. As described in Part I, section 2.2.5, the advantage of orthoprojecting directly from the DPC is greater geometric accuracy relative to orthophotomosaics, which can be subject to distortion in some cases. Additionally, Viscore orthoprojections can be easily generated in the same region as other data metrics, facilitating comparison of data.

4.5. Data curation

NCCOS uses the extensive file storage options available via the AVD ecosystem. Raw imagery and associated metadata collected by the field team is remotely uploaded to the AVD file system directly on site and immediately following image collection. Once file upload is complete, NCCOS staff can directly access imagery and begin the LAI construction process. Similarly, model construction is performed on the AVD file system, and output products are directly available to ecological data analysts who also conduct ecological data extraction using AVD resources. AVD allows geographically distributed users access to the same file system, and while upload of imagery is still limited by local internet bandwidth, all other steps of the LAI pipeline can be conducted directly on AVD. This approach facilitates rapid data access and sharing, particularly important given that over 10 TB of data were generated in the first year of the project and with data volumes expected grow exponentially over time. Strict metadata standards and file naming protocols are maintained throughout and are currently managed by a simple non-relational database.

As a federal public institution, data curation and archival activities are critical to the NCCOS and broader NOAA mission. Archiving and providing broad access to LAI products, however, presents a challenge due to the volumes of data, file types, and the relative novelty of these data. Staff at NCCOS work closely with NOAA's National Centers for Environmental Information and NOAA's Coral Reef Conservation Program to develop a data archival structure that will not only facilitate the long-term storage of LAI and associated data products but also encourage discoverability and access.

4.6. Links to detailed SOPs

LAI SOP (Cook et al., 2023): https://doi.org/10.25923/w8h9-4z75

Karen Bohnsack, NOAA ONMS/FKNMS

Case Study 5:

Pacific Islands Fisheries Science Center, NOAA

5.1. Overview

PIFSC is using the LAI approach to assess the status and trends of benthic communities, quantify coral vital rates, and measure coral reef structural complexity across the U.S. Pacific Islands and Territories. NOAA's Ecosystem Sciences Division within PIFSC, formerly the Coral Reef Ecosystem Division, has been monitoring Pacific coral populations and benthic communities since the early 2000s as part of the Pacific National Coral Reef Monitoring Program (NCRMP). PIFSC is charged with monitoring the status and trends of coral reefs across 40 primary islands, atolls, and shallow banks in the Hawaiian Archipelago (including Papahānaumokuākea Marine National Monument), the Mariana Archipelago (Guam and the Commonwealth of the Northern Mariana Islands, including the Marianas Trench Marine National Monument), American Samoa (including Rose Atoll Marine National Monument), and the Pacific Remote Island Areas Marine National Monument (Wake, Johnston, Palmyra, and Kingman Atolls and Howland, Baker, and Jarvis Islands). Historically, PIFSC has used in situ visual assessments of coral communities and benthic photoquadrat imagery to generate coral demographic metrics (colony density, size structure, partial mortality, and prevalence of disease and compromised health states), and percent cover of benthic taxa, which is conducted across hundreds of sites each year. The program has partnered with UC San Digeo, University of Hawai'i at Hilo and Hawai'i Institute of Marine Biology to develop methodological approaches for extracting benthic metrics from large-area image mosaics (60–130 m^2) derived using SfM techniques.

5.2. Image collection

Although LAI can be generated using almost any camera, collecting imagery underwater requires special field logistics and equipment considerations. The cost, compactness, portability, and available technical support for a camera system are important considerations. Given the logistical constraints of NCRMP surveys, the cameras need to be able to collect the desired imagery in 10–25 min depending on survey type. The Ecosystem Sciences Division tested a number of underwater camera systems (Canon G9x, GoPro5, Nikon D90, Sony A6300, and the Canon SL2). The combination of higher quality lenses, image resolution, superior white balancing, user control of camera settings, portability, continuous shooting, and affordability makes the Canon SL2 the camera of choice for PIFSC NCRMP surveys. In 2022, PIFSC switched to using the upgraded successor Canon SL3 because the SL2 model is no longer available for sale. SL2 and SL3 cameras are essentially the same camera with small differences in software and camera body. Through a series of underwater tests, determination was made of the camera settings that would produce the highest quality images for SfM processing over a broad range of operating conditions (e.g., overcast versus bright sun, clear shallow water versus deeper murky water, etc.), allowing the photographer to focus on swimming the plot with control instead of adjusting camera settings underwater. Cameras are equipped with an 18-mm lens inside an underwater housing with a 6-in dome port, and divers swim 1 m above the substrate to capture 5 images/s with a substrate footprint of 1.03 m × 0.69 m per image. Cameras are white balanced at the beginning of each dive using an 18% gray card.

PIFSC uses two image collection methods. The first method (belt survey) is used at each PIFSC benthic survey site (between 0–30 m depth) where an SfM survey is conducted along one 20-m transect containing four 2.5 m × 1 m segments beginning at each of the 0-, 5-, 10-, and 15-m marks. This survey approach also allows replication of the same area that has been visually assessed historically by the Ecosystem Sciences Division for coral demography metrics. Images are taken by swimming three paths along each side (six total passes) of the transect line. Given that historical in situ benthic surveys were conducted along belt transects, a number of square- and rectangular-shaped survey plots of different sizes were tested. Through this exercise, it was determined that the optimal survey plot shape is a rectangle of 3 m × 20 m. This sampling method allows work to remain within the operational time constraint (approximately 10 min) and capture of most large colonies, with the exception of very large coral thickets. 1,300–2,000 images are collected depending on underwater conditions.

The second method is used at mid-depth (15-m) fixed sites where an SfM survey is conducted across a 12-m diameter plot. The plot is marked by reference stakes across the diameter. An 8-in diameter reel is placed in the center of each plot, and 6 m of line is unspooled from the reel and attached to the camera at the edge of the plot. The plot is imaged in a spiral pattern by one diver swimming in toward the center of the plot followed by the second diver swimming outward. Image collection lasts 15–25 min, and 2,500–4,000 images are collected.

5.3. Model construction

Models are constructed using Agisoft Metashape Professional on several GPU-accelerated servers/high-powered workstations housed at PIFSC. The configuration of the servers/workstations varies, but each has an 8-core CPU, one or two high-end gaming GPUs, and 64–160 GB of RAM. A pilot project is currently being conducted to process models within the cloud for more dynamic processing capacity, but this has not been operationalized as of May 2023. A 2,000-image plot takes approximately 6 hr to process. Four to eight plots are typically processed simultaneously on each computer. PIFSC has developed a custom Python script that batch processes model building steps, requiring only prefiltered images and depth and scale metadata Briefly, processing steps include removing low-quality images based on Metashape's image-quality algorithm and then aligning photos to create an SPC. The poorly reconstructed points in the SPC are then filtered and removed through three steps using the Gradual Selection and Optimization tools to improve model geometry. Models are also scaled within Metashape at this point. A medium-quality DPC is built, and poorly reconstructed points are again removed using a point confidence threshold. A human checkpoint is built into the processing script here to add depth orientation and identify any outstanding issues. A 2.5D DEM is created from the scaled and oriented DPC. A high-quality (0.5 mm/pixel) orthomosaic is constructed from the DEM surface. For fixed sites, DPCs from multiple time points are coregistered within Agisoft Metashape before DEM and orthomosaic products are created and exported for analysis of time series data. The DPC is exported as a PLY file with camera positions exported as XML and JSON files for referencing raw imagery during data analysis steps. DEMs and orthomosaics are exported as TIFF files for subsequent analysis.

5.4. Ecological data extraction

Presently, PIFSC follows several workflows to extract data from image collection efforts. All models are imported into the custom software Viscore (Petrovic et al., 2014) primarily for referencing raw imagery to ensure the highest-quality data are recorded across methods. Orthomosaics created from Agisoft Metashape are imported into ArcGIS Pro to record features within each model and add data that are associated with each feature, such as species ID, condition, morphology, etc. Feature data are stored within a previously created ArcGIS geodatabase for each survey effort (e.g., cruise and project). For belt surveys, detailed measurements of coral colonies on the orthomosaics follow traditional benthic survey protocols to properly identify colony characteristics that replicate in-water surveys used by PIFSC. Refer to Winston et al. (2020) for the latest SOP for these in-water methods. Data extraction from fixed sites includes subsampling the plots with sampling quadrats and preparing a geodatabase in ArcGIS Pro. TagLab is used to semi-automatically annotate coral colonies within the specified sampling region of each orthomosaic and to semi-automatically link live coral colonies across time points. TagLab coral colony species IDs are merged and linked to an ArcGIS Pro geodatabase to estimate coral vital rates. Other data products, such as structural complexity and colony surface area, are directly analyzed in R using the DEM (Torres-Pulliza et al., 2020; Asbury et al., 2023) or within ArcGIS Pro.


Ari Halpern (NOAA NMFS/Pacific Islands Fisheries Science Center)

Mia Lamirand (NOAA NMFS/PIFSC)

5.5. Data curation

Good data management and data management plans are a crucial component of this workflow and were implemented during early stages of methods development. With imagery and data moved multiple times between servers and software programs, there are many opportunities for errors. The guidelines listed below were developed specifically to meet PIFSC needs, but the framework could be easily adapted by other organizations. Images collected are downloaded directly from the SD cards to a "Camera Downloads" back-up drive and are organized by project and date but are not sorted by site. The photos in the Camera Downloads server are the only assurance that a backup exists in case of a problem with the imagery in the months and years after image collection. Images are sorted from Camera Downloads into site folders within an organization scheme on a PIFSC internal server for model generation and data extraction. Data layers such as Metashape and ArcGIS Pro project files are stored in a subfolder within each site folder. Metadata information, such as marker and scale bar numbers and depths, and processing information are currently stored within a tracking spreadsheet on Google Drive; however, PIFSC is exploring options for more stable databases to store these vital data. Data products extracted from these models are uploaded to repositories, and current efforts are underway to upload imagery to NCEI for long-term archival and for use by the general public.

5.6. Links to detailed SOPs

Old Belts SOP (Suka, 2019): https://repository.library.noaa.gov/view/noaa/22753/; updated version currently in review

Benthic REA Survey SOP (Winston, 2020): https://repository.library.noaa.gov/view/noaa/23951

Vital Rates (Fixed Site) SOP (Rodriguez, 2021): https://repository.library.noaa.gov/view/noaa/32739

TagLab Technical Memo (Amir, 2023): https://repository.library.noaa.gov/view/noaa/50840

Case Study 6:

Perry Institute of Marine Science

6.1. Overview

PIMS is a small NGO doing coral reef work primarily in the Commonwealth of The Bahamas. They have developed their own SOP that is aimed at smaller organizations looking to efficiently utilize time series LAI for monitoring of both natural and restored reefs at any depth in a way that is relatively easy to implement in the water, makes use of relatively inexpensive photographic equipment, and can be processed by relatively inexpensive computers that are more often available to small organizations. The total startup cost of this system can be less than $5,000, including camera/lens/housing, three sets of plot markers, homemade scale bars, GPS device, software at educational pricing, and a computer and portable hard drive, though more can be invested to increase the efficiency and scale of monitoring.

6.2. Image collection

Before image collection, permanent markers are installed at the four plot corners. These markers are 90 mm × 95 mm flexible tiles with an automatically detectable 12-bit circular-coded Metashape target printed on a special non-toxic anti-fouling material patch. They are fixed to the substrate using four 2-in masonry nails or bolts per marker. Installing four of these markers usually requires 10–15 min, and they are designed to stay in place and resist biotic growth for decades, providing consistent GCPs with which to automatically line up time points for successive monitoring. The standard plot is 100 m^2, but the process can be used for any reasonable plot size. After installation, the depth of each marker is recorded with a dive watch to the nearest 10 cm. GPS points are collected at the surface above each marker using a Bluetooth GPS device (e.g., Bad Elf GNSS Surveyor or Garmin GLO2) and a smartphone, which inputs the locations and depths directly into a custom survey in the ESRI Survey123 app. In addition to the markers, scale bars consisting of an acrylic bar and printed Metashape targets a known distance apart are placed within the plot and weighed down with dive weights to prevent them from moving. Markers are installed, and reference data are collected only at the initial setup for a given plot; subsequent monitoring only requires photos to be collected, with no other reference information necessary.

Photos are collected in JPEG+RAW format using a single Sony a6400 or a6600 camera (APS-C sensor size) and a Samyang/Rokinon 12-mm f/2 lens (18-mm full-frame equivalent) inside a SeaFrogs polycarbonate housing with a 6-in dome port (available in either glass or acrylic). By using a wider-angle lens with approximately 90º field of view, the likelihood of data gaps is reduced, and the 3D structure of the reef is more accurately represented due to the sides of the photos providing better coverage of adjacent vertical surfaces. A single camera allows for a cheaper and less complex kit that is easier to set up and manage underwater. Prior to image collection, the camera is manually white balanced on a gray card at the median plot depth. Photos are taken 1.5–2.0 m from the substrate at 1-s intervals (using the camera's built-in intervalometer function) and are collected in a double-lawnmower pattern, with navigation performed visually by the operator, who is either on scuba or snorkel depending on reef depth. Rather than always holding the camera straight down (nadir), the operator varies the camera angle based on orientation of the substrate below. In the first-pass direction, the operator aims for tightly overlapping passes about 1 m apart, leaving no data gaps, and in the second-pass direction (at a 90º angle to the first), emphasis is less on side overlap and more on "stitching" back over the plot in a z-pattern, thus eliminating some redundancy and reducing the number of overall photos, which enables quicker processing (Figure 34). For a typical 10 m × 10 m plot, image acquisition takes 20–40 min and totals around 1,200–2,400 images. With the camera set up as described above, the final mosaic resolution is approximately 0.5 mm/pixel, allowing for expert identification of corals to the species level and most other sessile benthic organisms to the genus and sometimes species level.

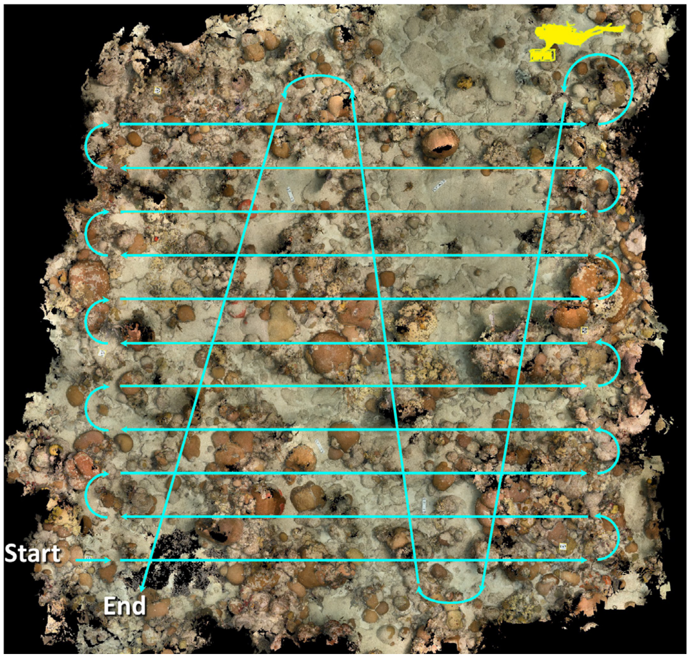

Figure 34. Example of half-lawnmower with tie lines image collection. Rather than use the full double-lawnmower pattern, PIMS image collection begins with a set of parallel passes with standard spacing. After these passes are complete, a z-pattern consisting of at least three angled passes is conducted over the first set of passes.

6.3. Model construction

Once collected, the photos are imported onto local hard drives for storage and processing. RAW images are archived for potential future use, and JPEG images are immediately used for LAI processing. PIMS has developed a set of custom Python scripts that interact with Agisoft Metashape Professional to automate the entire LAI processing pipeline. Three pieces of data are input into the script's graphical user interface window: a folder containing the photos, an exported CSV of the plot's corner marker GPS locations and depths, and a text file containing the measurements of the scale bars used. Processing naming conventions, filenames, and storage locations are set up through the window, shielding the user from needing to navigate Metashape's own interface until the end of the process for QAQC. At that point, errors can be corrected and the script run again to pick up from any point in the process where the error may have occurred, saving time and effort by reprocessing only what is necessary.

The majority of the processing decisions have been automated. Data for each plot processed using this script are thus standardized, promoting consistency of datasets. The script generates the SPC, then detects markers and inputs the reference and scaling information before optimizing the SPC. Next, depth maps are calculated and used to generated a 3D mesh, which is then used to build the orthophotomosaic and DEM. To reduce computational demands and processing time, the DPC is not generated. Finally, the Python script generates a processing report and exports the orthomosaic and DEM as georeferenced TIFF files and the region of interest (ROI) as a shapefile.

Will Greene (Perry Institute of Marine Science [PIMS])

PIMS created a secondary Python script that automates the alignment of subsequent time-point data (organized in the same Metashape file as a new "chunk") to the first, using the locations of the four corner markers. Next, the original Python script is run again, creating pixel-aligned outputs across time points with almost no manual effort. PIMS currently processes data using Agisoft Metashape on a large PC with a 14-core Intel i9 processor and three 11-GB NVIDIA GTX 1080Ti graphics cards, and a secondary Mac Mini with an 8-GB AMD graphics card housed inside an eGPU. A typical 1,500-image plot takes approximately 4 hr on the first, more powerful system, while the less powerful Mac system takes around 8 hr.

6.4. Ecological data extraction

PIMS output orthomosaics are fully georeferenced (including depth) and standardized at 0.5-mm resolution, which allows for the extraction of ecological, physical, and geographical information. The specific data extracted depends on the project. For general statistics about the reef's benthic cover, PIMS uses the automatically created ROI to generate random points in QGIS, and then manually identifies the benthic cover at each point. A QGIS environment is set up specifically for this task, enabling the rapid assessment of points in a similar fashion to software similar to Coral Point Count (Kohler and Gill, 2006) but with the added


Valeria Pizarro (PIMS)

benefit of unlimited file sizes and extremely fast loading. To gain more in-depth information about coral species/size distribution and changes over time, the orthomosaic and DEM are directly exported from Metashape, cropped to the ROI, in a format that is ready for loading and analysis in TagLab. Corals are outlined semi-automatically using the available tools in TagLab, identified to species, and tracked through time to analyze changes. Geographic and physical (e.g., rugosity) data are extracted either directly in Metashape, where a full 2D to 3D surface area ratio can be easily calculated, or in other GIS software like QGIS. With the exception of Metashape Pro, the PIMS process uses only free, open-source, and publicly available software for analysis.

6.5. Data curation

PIMS's LAI program currently totals around 30 TB of photos, processing files, and data products. As a small non-profit, PIMS lacks access to large, university-scale storage solutions, and thus a combination of local hard drive storage and cloud storage has been the most cost-effective method for curating data up to this point. Original photos (RAW and JPEG) are organized onto 14-TB hard drives with a backup copy created at regular intervals on a separate drive. Processing files (orthomosaics, point clouds, meshes, etc.) are stored on separate drives and are similarly backed up. Products ready for analysis (finished orthomosaics, DEMs, and processing reports) are uploaded to Dropbox cloud storage for sharing within the organization. Analysis files (TagLab JSON files, shapefiles, spreadsheets, etc.) are similarly stored in Dropbox alongside the mosaics, DEMs, and reports. As PIMS grows, this system is expected to change, perhaps into a more fully cloud-based curation workflow as is needed.

6.6. Link to detailed SOPs

The complete SOP and Metashape Python script (Github):
https://github.com/Perry-Institute

Case Study 7:

Sandin Lab, UC San Diego

7.1. Overview

The SOP used by the Sandin Lab was originally developed in 2012 in consultation with Dr. Arthur Gleason at the University of Miami. Over the following 10 years, the protocol has been updated to reflect lessons learned and innovation to various aspects of the workflow. This approach has most frequently been implemented in the context of the 100 Island Challenge as well as associated research efforts conducted by the Sandin lab and collaborating partners. As of the end of 2022, this SOP has guided the collection of over 4,000 LAI data products and informed the data and analyses of over 10 peer-reviewed publications.

7.2. Image collection

This SOP was designed to support the collection of LAI from shallow-water (<20-m depth) coral reef habitats, principally from exposed, oceanic fore reef locations. As such, the approach has design characteristics to be time efficient (45–50 min) with equipment and procedures that are robust to the variable environmental conditions that can be expected across reef environments (e.g., swell and current).

The design criteria for the camera system included elements of availability, robustness, and functionality. The aim was to identify an off-the-shelf camera system that would support the collection of LAI and require little to no customization of equipment. The main features guiding the selection of a camera system included a built-in intervalometer and extended battery capacity. The team has conducted a number of camera tests over the years to identify camera systems most suitable for supporting LAI efforts. Initially, the team identified a cropped-sensor camera (i.e., Nikon D7000 with a CMOS sensor) as a robust solution to support the collection of imagery during multiple dives throughout a day. Over time, advancements in camera technology and competitive pricing led to the use of a full-frame DSLR camera (i.e., Nikon D780). Both cameras are robust and easy to use and feature built-in intervalometers.

The SOP identifies a two-camera setup, an approach that allows for the simultaneous collection of complementary imagery by a single diver. Of the two cameras, the first is outfitted with a wide-angle lens to ensure a collection of imagery with high inter-image overlap needed for model generation. The second camera is outfitted with a higher focal length lens to collect imagery in greater detail for improved identification of objects within the imaged area (e.g., species-level identifications). The 18-mm lenses originally used with the Nikon D7000 cameras were replaced with a 24-mm fixed focal length after switching to the Nikon D780 in 2021. Due to the differences in sensor size, the lenses cover the same spatial footprint when operated at the same distance from the bottom; however, the D780 is a more advanced camera with a full-frame sensor and produces higher-quality imagery at this distance. Cameras are mounted in a sturdy, high-density plastic frame, and the entire camera setup is trimmed to be slightly negatively buoyant.


Brian Zgliczynski (SIO, UC San Diego)


Wayne Andrew (One Reef)

In an effort to capture the entire borders of hundreds of coral colonies per LAI dataset, the target plot size was defined as 100 m^2, with an approximately 1-m buffer margin also imaged around the plot. The plot is delineated with a series of temporary floats at the four corners and in the center of two of the sides. Navigation is accomplished visually by the diver, with the aid of a compass and depth sensor, as needed. Imagery is collected from a distance of 1.5–2.0 m from the bottom, and the plot is imaged in a gridded pattern, with approximately 12–15 passes conducted per direction. Depending on the complexity of the site and environmental conditions, image collection lasts for 30–50 min, and 2,000–3,000 images are collected (per camera). Cameras are white balanced at the beginning of each dive using a gray card (18% gray) at the median plot depth.

7.3. Model construction

Models are constructed with Agisoft Metashape Professional Edition using a networked cluster of custom-built HPCs. The configuration of each build varies slightly, but each has an 8- to 16-core CPU, one to three high-end gaming GPUs and 64–256 GB of RAM. Processing a single set of 2,000–3,000 images (per camera) has taken up to seven days on a single machine, with processing time decreasing through the years (now averaging 1 or 2 days) due to improvements in Metashape. A maximum of seven HPCs has been used as a cluster to accommodate processing needs when image collection has reached upward of approximately 550 models in a single year or when rapid processing of a collection of images is required. Processing steps include image alignment, which uses both sets of images, and the building of a high-quality DPC, which uses images from only the wide-angle lens camera (as there are differences in lighting between the two cameras that can affect the visual quality of the DPC). A robust image collection procedure with high image overlap has led to minimalistic settings to be used in Metashape during processing, which balances high-quality and efficient processing. No scaling or optimization is done during the model construction process. Camera pose estimates are exported as an XML file along with a JSON file, which references image file paths. The DPC is exported as a PLY file for subsequent analysis. Additionally, an orthomosaic is generated for visualization purposes using an intermediate step of meshing the SPC.

7.4. Ecological data extraction

Image collection is optimized for high-resolution LAI necessary to extract precise colony size and shape information, enabling the most resolved taxonomic identification possible (typically species or genus level) for the given image resolution. A variety of workflows are used in the custom software platform Viscore to facilitate data extraction. This includes point-based annotations to estimate percent cover and geometric measurements of scaled point clouds for structural complexity analyses. Detailed measurements of the size and shape of features (e.g., coral colonies) are created by generating orthoprojections directly from the DPC, which are exported directly to the program TagLab to facilitate the segmentation and annotation process. The segmentation process is largely manual but is aided by the use of semi-automated segmentation tools in TagLab. Given that orthoprojections can often offer a view that is visually ambiguous, especially in defining sub-centimeter border features, this SOP includes a workflow employing TagLab in tandem with Viscore. A major benefit of Viscore is the ability to fetch raw imagery, easily and interactively, from any location on the model, which can allow a user to access the highest-resolution data (i.e., the original imagery) to improve the precision of segmentation data products recorded using TagLab. Additionally, Viscore is used to coregister point clouds collected through time, which enables additional analyses exploring change over time for any of the data extraction workflows.

7.5. Data curation

The data volumes collected by the Sandin Lab have grown exponentially through time as methods and technology have improved. Initial efforts to store data began with external hard drives and then expanded to NAS devices. Now, the lab uses enterprise storage systems, managed by the UC San Diego Research Data Library and the facilities of the San Diego Supercomputer Center; by the end of 2022, storage needs exceeded 500 TB. Strict file naming protocols as well as extensive quality control procedures are followed to manage these extensive data volumes and associated metadata. Raw data layers, including the collected images and associated metadata, are stored in duplicate on separate storage systems to protect against data loss of irreplaceable files. Current efforts are being made to upload these raw data layers to the UC San Diego Library Digital Collections to permanently archive these data and make them available to the general public. Generated data layers that are not heavily used, such as the Metashape project files, are stored on a limited access storage share. Heavy data layers that are more readily used during data extraction, such as the Viscore and TagLab files, are typically copied from their centralized network storage location to a local drive for data extraction. Processed files are then re-uploaded to the centralized storage location with updated naming to reflect any work history in order to prevent data loss through overwriting newer files when modified by multiple users. Data extraction activities are also extensive, with size information extracted for over 100,000 coral colonies to date, in addition to nearly 500,000 point-level annotations.

7.6. Link to detailed SOPs

Detailed SOPs from the Sandin Lab are available online at the **UC San Diego Library Digital Collections (Sandin Lab, 2023)**: https://library.ucsd.edu/dc/collection/bb43111847

Lindsay Bonito (SIO, UC San Diego)

Brian Zgliczynski (SIO, UC San Diego)

Clinton Edwards (SIO, UC San Diego)

Case Study 8:

University of Miami

8.1. Overview

The University of Miami (UM) group differs from most or all of the other teams included in this summary in the sense that rather than maintain a standard set of operating procedures, they work closely with other groups to develop a tailored approach specific the desired application. The team at UM focuses more broadly on techniques, hardware design, and software development rather than any specific type of ecological analysis. UM partners with many other colleagues to use LAI in applications as diverse as monitoring coral restoration success, coral demographic/bleaching/disease monitoring, damage assessment from hurricanes or vessel groundings, documentation of submerged cultural resources, and assessment of legacy underwater military munitions. Since both the needs of these diverse applications and the resources of these various partners vary widely, there is not a single SOP that is appropriate for all of them. Nevertheless, there are general principles that form the starting point in every case.

8.2. Image collection

As discussed in Part I, section 3, planning for image collection starts with the questions to be asked of the data. This process determines the spatial resolution, areal extents, and temporal revisit frequency needed, which in turn guide camera and site setup choices. Budget and shipping or travel logistics also factor into equipment decisions, of course. Generally speaking, it is recommended to use DSLR or interchangeable-lens mirrorless cameras due to the optical quality of the lenses and housing ports available, the high bit-depth and low noise of the sensors, the large batteries, and the ability to precisely control shutter speed. Also, generally speaking, multiple cameras are recommended for redundancy and to simultaneously collect data at multiple spatial resolutions, thereby allowing collection of imagery with both high overlap and fine detail of the seabed (Gintert et al., 2008).

For "standard" plot sizes (up to 500 m^2), it is suggested to use two DSLR cameras with different focal length lenses if allowed by budget and logistical constraints. Otherwise, a single DSLR camera paired with a small format action camera (e.g., GoPro) can be used to achieve the same end. For surveys covering larger areas (>500 m^2), it is recommended to use multiple cameras spaced about 1 m apart on a pole. If budgets are limited, or a diver will be swimming the array, GoPro cameras are used. If funds allow and the pole will be towed by a surface vessel or operated by DPV, multiple DSLR cameras should be used. Arrays with as few as two and as many as eight cameras at one time have been successfully tested and implemented in various applications.

Regardless of the camera chosen, data collection must always prioritize a) high overlap, b) sharp in-focus images, c) at least some "tie lines" if not full double-lawnmower coverage, and d) scale bars of some kind. Also, measuring of the depth of known points and marking the site with suitable stakes are required to facilitate repeat surveys.

8.3. Model construction

Between 2003 and 2015, custom in-house software developed by Nuno Gracias was used to generate large-area imagery in the form of orthophotomosaics (Gracias and Santos-Victor, 2001; Gracias et al., 2003; Lirman et al., 2007; Edwards et al., 2017). Since 2016, Agisoft Metashape has been used for model construction. When needed, older image sets were reprocessed using Metashape, which has worked well for all except the earliest datasets that were collected with still images extracted from interlaced video. Models are processed using an HPC cluster consisting of seven workstations running Linux (CentOS). The CPUs on these machines are old, but the GPUs have been upgraded within the past few years.

Data are prefiltered to remove blurry or extremely oblique images, and a contrast stretch and color balance are applied to multiple cameras when using a large array. These steps are all done using image processing functions built into MATLAB. After making the sparse cloud, scale bars are located, and scaling transform is applied directly within Metashape. Camera parameter optimization is conducted before construction of the DPC, DEM, and orthophotomosaic. At a minimum, the DPC and orthophotomosaic are exported for visualization (see Part II, section 8.5).

8.4. Ecological data extraction

The recommended approach to data extraction has also varied over time as available technologies have increased and in order to best suit various project needs. For estimations of percent cover, workflows have been developed using Coral Point Count (Kohler and Gill, 2006), custom MATLAB scripts, and Viscore. For fate tracking (colony demographics, bleaching, disease assessment, etc.), a combination of software platforms has been used including GIS packages such as ArcMap, QGIS, Global Mapper, or similar. Methods to relate the individual component images of the LAI to the raw images collected in the field have been used extensively (usually Zoomify or custom MATLAB scripts). More recently, Viscore has been used extensively for these functions. Simple segmentations of large polygons (e.g., tracing the extent of a damaged area in vessel grounding surveys) are straightforward to accomplish via GIS packages, and as these data are often required to be delivered in GIS format, they are still commonly conducted in this manner. See examples of data extraction in (Lirman et al., 2010; Griffin et al., 2015).

8.5. Data curation

Data curation activities are conceptualized in two distinct phases. The first phase consists of management of data that are actively used in the model construction or data extraction steps of the LAI pipeline. The second phase consists of data that are, for the time being at least, no longer active and are therefore considered archived. Data that are active are stored on shared NAS devices. Linux, Windows, and Mac clients can connect to these NAS devices as needed for processing and data extraction. Each NAS device has a RAID-6 array comprising spinning magnetic platter hard drives. Although these are slower to read/write than modern SSDs, the trade-off in speed has been acceptable given the lower cost of these configurations. The drives are connected to hard-wired gigabit ethernet, however, since the speed boost relative to Wi-Fi connections is noticeable. Every night, each NAS device is mirrored to an identical unit located in a different building for redundancy.

Once the data are not being actively used to create or analyze models, they are migrated to "archived status." Web-viewable versions of orthomosaics and DPCs are created and posted to a public website. All of the raw and processed data are then copied onto individually formatted (i.e., not RAID) hard drives, of which duplicated copies are made. As archive hard drives become full, they are taken offline and stored in a secure location. If a given dataset needs to be reprocessed in the future, these archived disks can be mounted and the data copied back to the active arrays.

8.6. Link to detailed SOPs

The link provided here is not for a specific detailed SOP for a given survey objective but rather a collection of guidelines applicable to many different survey goals. Links are also provided to a variety of data products, and readers are encouraged to contact Art Gleason for any inquiries.

Overview of Gleason Laboratory LAI approach: https://web3.physics.miami.edu/~agleason/mosaics/

Example datasets: https://web3.physics.miami.edu/~agleason/mosaic_results/

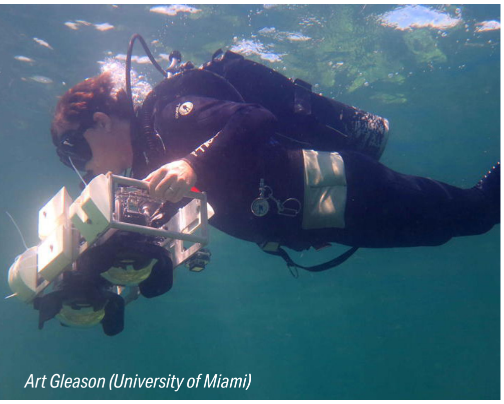

Art Gleason (University of Miami)

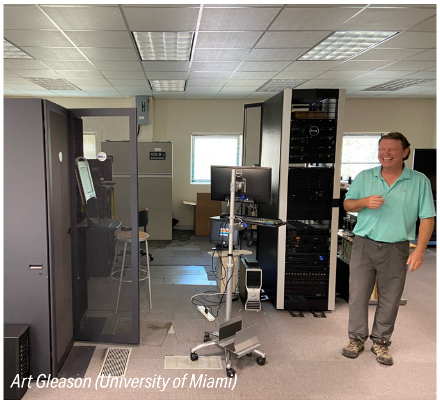

Art Gleason (University of Miami)

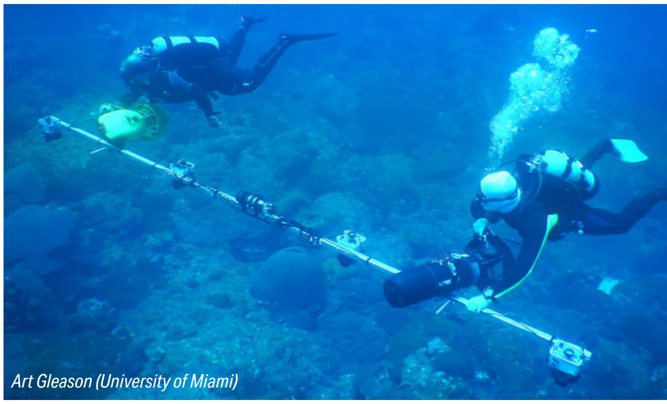

Art Gleason (University of Miami)

References

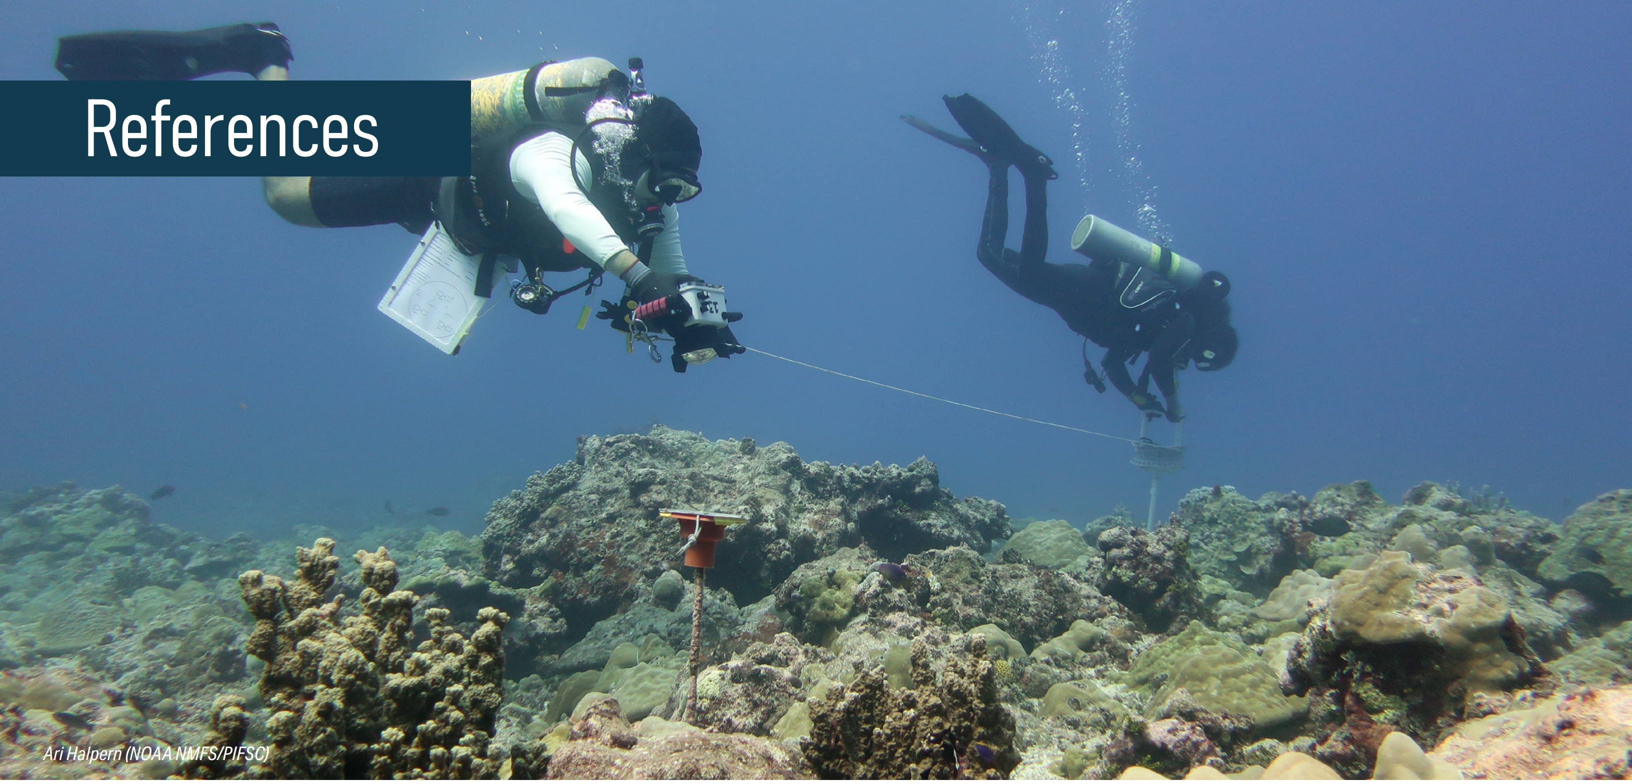

Ari Halpern (NOAA NMFS/PIFSC)

Akkaynak, D., and Treibitz, T. (2019). *Sea-Thru: A method for removing water from underwater images*. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 1682–1691). https://doi.org/10.1109/CVPR.2019.00178

Althaus, F., Hill, N., Ferrari, R., Edwards, L., Przeslawski, R., Schönberg, C. H. L., Stuart-Smith, R., Barrett, N., Edgar, G., Colquhoun, J., Tran, M., Jordan, A., Rees, T., and Gowlett-Holmes, K. (2015). A standardised vocabulary for identifying benthic biota and substrata from underwater imagery: The CATAMI classification scheme. *PLOS ONE, 10*(10), e0141039. https://doi.org/10.1371/journal.pone.0141039

Amir, C., Lamirand, M., Couch, C. S., and Oliver, T. (2022). *Measuring coral vital rates using TagLab semi-automatic coral segmentation and temporal linking across fixed sites: standard operating procedures and time savings estimates*. NOAA National Marine Fisheries Service, Pacific Island Fisheries Science Center. NOAA Technical Memorandum NMFS-PIFSC-139. https://doi.org/10.25923/s2ym-tn10

Aston, E. A., Duce, S., Hoey, A. S., and Ferrari, R. (2022). A protocol for extracting structural metrics from 3D reconstructions of corals. *Frontiers in Marine Science, 9*, 854395. https://doi.org/10.3389/fmars.2022.854395

Bak, R. P. M., and Nieuwland, G. (1995). Long-term change in coral communities along depth gradients over leeward reefs in the Netherlands Antilles. *Bulletin of Marine Science, 56*(2), 609–619.

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., and Kriegman, D. (2015). Towards automated annotation of benthic survey images: Variability of human experts and operational modes of automation. *PLOS ONE, 10*(7), e0130312. https://doi.org/10.1371/journal.pone.0130312

Bellavia, F., Colombo, C., Morelli, L., and Remondino, F. (2022). Challenges in image matching for cultural heritage: An overview and perspective. In Mazzeo, P.L., Frontoni, E., Sclaroff, S., Distante, C. (Eds.), *Image analysis and processing, ICIAP 2022*. *Lecture notes in computer science*, vol 13373 (pp. 210–222). Springer, Cham. https://doi.org/10.1007/978-3-031-13321-3_19

Benassi, F., Dall'Asta, E., Diotri, F., Forlani, G., Morra di Cella, U., Roncella, R., and Santise, M. (2017). Testing accuracy and repeatability of UAV blocks oriented with GNSS-supported aerial triangulation. *Remote Sensing, 9*(2), 172. https://doi.org/10.3390/rs9020172

Brown, E. K., Cox, E., Jokiel, P. L., Rodgers, S. K. u., Smith, W. R., Tissot, B. N., Coles, S. L., and Hultquist, J. (2004). Development of benthic sampling methods for the Coral Reef Assessment and Monitoring Program (CRAMP) in Hawai'i. *Pacific Science, 58*(2), 145–158. https://doi.org/10.1353/psc.2004.0013

Bryson, M., Ferrari, R., Figueira, W., Pizarro, O., Madin, J., Williams, S., and Byrne, M. (2017). Characterization of measurement errors using structure-from-motion and photogrammetry to measure marine habitat structural complexity. *Ecology and Evolution, 7*(15), 5669–5681. https://doi.org/10.1002/ece3.3127

Burns, J. H. R., Alexandrov, T., Ovchinnikova, E., Gates, R. D., and Takabayashi, M. (2016a). Investigating the spatial distribution of growth anomalies affecting *Montipora capitata* corals in a 3-dimensional framework. *Journal of Invertebrate Pathology,* 140, 51–57. https://doi.org/10.1016/j.jip.2016.08.007

Burns, J. H. R., Alexandrov, T., Ovchinnikova, K., Gates, R. D., and Takabayashi, M. (2016b). Data for spatial analysis of growth anomaly lesions on *Montipora capitata* coral colonies using 3D reconstruction techniques. *Data in Brief, 9*, 460–462. https://doi.org/10.1016/j.dib.2016.09.009

Burns, J. H. R., Delparte, D., Gates, R. D., and Takabayashi, M. (2015). Integrating structure-from-motion photogrammetry with geospatial software as a novel technique for quantifying 3D ecological characteristics of coral reefs. *PeerJ, 3*, e1077. https://doi.org/10.7717/peerj.1077

Burns, J. H. R., Fukunaga, A., Pascoe, K. H., Runyan, A., Craig, B. K., Talbot, J., Pugh, A., and Kosaki, R. K. (2019). 3D habitat complexity of coral reefs in the northwestern Hawaiian Islands is driven by coral assemblage structure. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-2/W10*, 61–67. https://doi.org/10.5194/isprs-archives-XLII-2-W10-61-2019

Carbonneau, P. E., and Dietrich, J. T. (2017). Cost-effective non-metric photogrammetry from consumer-grade sUAS: Implications for direct georeferencing of structure from motion photogrammetry. *Earth Surface Processes and Landforms, 42*(3), 473–486. https://doi.org/10.1002/esp.4012

Carlon, D. B., and Olson, R. R. (1993). Larval dispersal distance as an explanation for adult spatial pattern in two Caribbean reef corals. *Journal of Experimental Marine Biology and Ecology, 173*(2), 247–263. https://doi.org/10.1016/0022-0981(93)90056-T

Carlot, J., Vousdoukas, M., Rovere, A., Karambas, T., Lenihan, H. S., Kayal, M., Adjeroud, M., Pérez-Rosales, G., Hedouin, L., and Parravicini, V. (2023). Coral reef structural complexity loss exposes coastlines to waves. *Scientific Reports, 13*(1), 1683. https://doi.org/10.1038/s41598-023-28945-x

Chirayath, V., and Earle, S. A. (2016). Drones that see through waves- Preliminary results from airborne fluid lensing for centimetre-scale aquatic conservation. *Aquatic Conservation: Marine and Freshwater Ecosystems, 26*, 237–250. https://doi.org/10.1002/aqc.2654

Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F., and Ranzuglia, G. (2008). Meshlab: An open-source mesh processing tool. In Scarano, V., De Chiara, R., and Erra, U. (Eds.), *Eurographics Italian Chapter Conference* (pp. 129–136). The Eurographics Association. https://doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136

Cook, S., Rojano, S. G., Edwards, C. B., Bollinger, M. A., Pierce, J., and Viehman, T. S. (2023). *Standard operating procedures for the use of Large-area imaging in coral reef science: Applications for the Florida Keys Mission Iconic Reefs Restoration Project*. NOAA National Ocean Service, National Centers for Coastal Ocean Science. NOAA NCCOS Technical Memorandum 320. https://doi.org/10.25923/w8h9-4z75

Coral Reef Conservation Program (CRCP) (2022a). *National Coral Reef Monitoring Program (NCRMP) Benthic Community Assessment Survey Field Protocols for U.S. Atlantic: Florida, Flower Garden Banks, Puerto Rico, U.S. Virgin Islands 2022.* NOAA Office of Coast Management, Coral Reef Conservation Program. https://doi.org/10.25923/0708-8333

Coral Reef Conservation Program (CRCP) (2022b). *National Coral Reef Monitoring Program (NCRMP) Coral Demographics Survey Field Protocols for U.S. Atlantic: Florida, Flower Garden Banks, Puerto Rico, U.S. Virgin Islands 2022*. NOAA Office of Coast Management, Coral Reef Conservation Program. https://doi.org/10.25923/9a1r-m911

Corso, J., French, B. J., Edwards, C. B., Pedersen, N. E., Zgliczynski, B. J., Planes, S., Pacala, S., and Sandin, S. A. (2022). Non-random orientation of *Pocillopora* colonies on forereefs of Moorea, French Polynesia. *Marine Ecology Progress Series, 693*, 177–182. https://doi.org/10.3354/meps14078

Couch, C. S., Oliver, T. A., Suka, R., Lamirand, M., Asbury, M., Amir, C., Vargas-Ángel, B., Winston, M., Huntington, B., Lichowski, F., Halperin, A., Gray, A., Garriques, J., and Samson, J. (2021). Comparing coral colony surveys from in-water observations and structure-from-motion imagery shows low methodological bias. *Frontiers in Marine Science, 8*, 647943. https://doi.org/10.3389/fmars.2021.647943

Crandall, D., Owens, A., Snavely, N., and Huttenlocher, D. (2011). Discrete-continuous optimization for large-scale structure from motion. In *CVPR 2011* (pp. 3001–3008). https://doi.org/10.1109/CVPR.2011.5995626

Curtis, J. S., Galvan, J. W., Primo, A., Osenberg, C. W., and Stier, A. C. (2023). 3D photogrammetry improves measurement of growth and biodiversity patterns in branching corals. *Coral Reefs, 42*(3), 623–627. https://doi.org/10.1007/s00338-023-02367-7

Dana, T. F. (1976). Reef-coral dispersion patterns and environmental variables on a Caribbean coral reef. *Bulletin of Marine Science, 26*, 1–13.

Darwin, C., and Stoddart, D. R. (1962). Coral Islands with introduction, map and remarks by D. R. Stoddart. *Atoll Research Bulletin, 88*, 1–20. https://doi.org/10.5479/si.00775630.88.1

Deignan, L. K., and Pawlik, J. R. (2015). Perilous proximity: Does the Janzen–Connell hypothesis explain the distribution of giant barrel sponges on a Florida coral reef? *Coral Reefs, 34*(2), 561–567. https://doi.org/10.1007/s00338-014-1255-x

Ditria, E. M., Buelow, C. A., Gonzalez-Rivero, M., and Connolly, R. M. (2022). Artificial intelligence and automated monitoring for assisting conservation of marine ecosystems: A perspective. *Frontiers in Marine Science, 9*, 918104. https://doi.org/10.3389/fmars.2022.918104

Drury, C., Manzello, D., and Lirman, D. (2017). Genotype and local environment dynamically influence growth, disturbance response and survivorship in the threatened coral, *Acropora cervicornis*. *PLOS ONE, 12*(3), e0174000. https://doi.org/10.1371/journal.pone.0174000

Dumas, P., Bertaud, A., Peignon, C., Léopold, M., and Pelletier, D. (2009). A “quick and clean” photographic method for the description of coral reef habitats. *Journal of Experimental Marine Biology and Ecology, 368*(2), 161–168. https://doi.org/10.1016/j.jembe.2008.10.002

Edmunds, P. J. (2015). A quarter-century demographic analysis of the Caribbean coral, *Orbicella annularis*, and projections of population size over the next century. *Limnology and Oceanography, 60*(3), 840–855. https://doi.org/10.1002/lno.10075

Edmunds, P. J., Nelson, H. R., and Bramanti, L. (2018). Density-dependence mediates coral assemblage structure. *Ecology, 99*(11), 2605–2613. https://doi.org/10.1002/ecy.2511

Edwards, C. B., Eynaud, Y., Williams, G. J., Pedersen, N. E., Zgliczynski, B. J., Gleason, A. C. R., Smith, J. E., and Sandin, S. A. (2017). Large-area imaging reveals biologically driven non-random spatial patterns of corals at a remote reef. *Coral Reefs, 36*(4), 1291–1305. https://doi.org/10.1007/s00338-017-1624-3

Eltner, A., and Schneider, D. (2015). Analysis of different methods for 3D reconstruction of natural surfaces from parallel-axes UAV images. *The Photogrammetric Record, 30*(151), 279–299. https://doi.org/10.1111/phor.12115

Eltner, A., and Sofia, G. (2020). Structure from motion photogrammetric technique. In Tarolli, P., and Mudd, S. M. (Eds), *Developments in Earth Surface Processes, 23(2020)*, pp. 1–24. https://doi.org/10.1016/B978-0-444-64177-9.00001-1

Ferrari, R., Bryson, M., Bridge, T., Hustache, J., Williams, S. B., Byrne, M., and Figueira, W. (2016). Quantifying the response of structural complexity and community composition to environmental change in marine communities. *Global Change Biology, 22*(5), 1965–1975. https://doi.org/10.1111/gcb.13197

Ferrari, R., Lachs, L., Pygas, D. R., Humanes, A., Sommer, B., Figueira, W. F., Edwards, A. J., Bythell, J. C., and Guest, J. R. (2021). Photogrammetry as a tool to improve ecosystem restoration. *Trends in Ecology and Evolution, 36*(12), 1093–1101. https://doi.org/10.1016/j.tree.2021.07.004

Ferrari, R., Leon, J. X., Davies, A. J., Burns, J. H., Sandin, S. A., Figueira, W. F., and Gonzalez-Rivero, M. (2022). Advances in 3D habitat mapping of marine ecosystem ecology and conservation. *Frontiers in Marine Science, 8*, 827430. https://doi.org/10.3389/fmars.2021.827430

Figueira, W., Ferrari, R., Weatherby, E., Porter, A., Hawes, S., and Byrne, M. (2015). Accuracy and precision of habitat structural complexity metrics derived from underwater photogrammetry. *Remote Sensing, 7*(12), 16883–16900. https://doi.org/10.3390/rs71215859

Forsmoo, J., Anderson, K., Macleod, C. J. A., Wilkinson, M. E., DeBell, L., and Brazier, R. E. (2019). Structure from motion photogrammetry in ecology: Does the choice of software matter? *Ecology and Evolution, 9*(23), 12964–12979. https://doi.org/10.1002/ece3.5443

Fox, M. D., Carter, A. L., Edwards, C. B., Takeshita, Y., Johnson, M. D., Petrovic, V., Amir, C. G., Sala, E., Sandin, S. A., and 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

Friedman, A., Pizarro, O., Williams, S. B., and Johnson-Roberson, M. (2012). Multi-scale measures of rugosity, slope and aspect from benthic stereo image reconstructions. *PLOS ONE, 7*(12), e50440. https://doi.org/10.1371/journal.pone.0050440

Fukunaga, A., J. H. Burns, K. H. Pascoe, and R. K. Kosaki. (2022a). A remote coral reef shows macroalgal succession following a mass bleaching event. *Ecological Indicators, 142*, 109175. https://doi.org/10.1016/j.ecolind.2022.109175

Fukunaga, A., Pascoe, K. H., Pugh, A. R., Kosaki, R. K., and Burns, J. H. R. (2022b). Underwater photogrammetry captures the initial recovery of a coral reef at Lalo Atoll. *Diversity, 14*(1), 39. https://doi.org/10.3390/d14010039

Garcia-Garcia, A., Orts-Escolano, S., Oprea, S., Villena-Martinez, V., Martinez-Gonzalez, P., and Garcia-Rodriguez, J. (2018). A survey on deep learning techniques for image and video semantic segmentation. *Applied Soft Computing, 70*, 41–65. https://doi.org/10.1016/j.asoc.2018.05.018

George, E. E., Mullinix, J. A., Meng, F., Bailey, B. A., Edwards, C., Felts, B., Haas, A. F., Hartmann, A. C., Mueller, B., Roach, T. N. F., Salamon, P., Silveira, C., Vermeij, M. J. A., Rohwer, F., and Luque, A. (2021). Space-filling and benthic competition on coral reefs. *PeerJ, 9*, e11213. https://doi.org/10.7717/peerj.11213

Gerke, M., and Przybilla, H.-J. (2016). Accuracy analysis of photogrammetric UAV image blocks: Influence of onboard RTK-GNSS and cross flight patterns. *Photogrammetrie - Fernerkundung - Geoinformation, 2016*(1), 17–30. https://doi.org/10.1127/pfg/2016/0284

Gintert, B., Gracias, N. R., Gleason, A. C., Lirman, D., Dick, M., Kramer, P. A., and Reid, R. P. (2009). Second-generation landscape mosaics of coral reefs. In *Proceedings of the 11th International Coral Reef Symposium* , pp. 577-581. Ft. Lauderdale, Florida.

Goergen, E. A., and Gilliam, D. S. (2018). Outplanting technique, host genotype, and site affect the initial success of outplanted *Acropora cervicornis*. *PeerJ, 6*, e4433. https://doi.org/10.7717/peerj.4433

Goergen, E. A., Schopmeyer, S., Moulding, A. L., Moura, A., Kramer, P., and Viehman, T. S. (2020). *Coral reef restoration monitoring guide: Methods to evaluate restoration success from local to ecosystem scales*. NOAA National Ocean Service, National Centers for Coastal Ocean Science. NOAA Technical Memorandum NOS NCCOS 279. https://doi.org/10.25923/xndz-h538

González-Rivero, M., Harborne, A. R., Herrera-Reveles, A., Bozec, Y. M., Rogers, A., Friedman, A., Ganase, A., and Hoegh-Guldberg, O. (2017). Linking fishes to multiple metrics of coral reef structural complexity using three-dimensional technology. *Scientific Reports, 7*(1), 13965. https://doi.org/10.1038/s41598-017-14272-5

González-Rivero, M., Beijbom, O., Rodriguez-Ramirez, A., Bryant, D. E. P., Ganase, A., Gonzalez-Marrero, Y., Herrera-Reveles, A., Kennedy, E. V., Kim, C. J. S., Lopez-Marcano, S., Markey, K., Neal, B. P., Osborne, K., Reyes-Nivia, C., Sampayo, E. M., Stolberg, K., Taylor, A., Vercelloni, J., Wyatt, M., and Hoegh-Guldberg, O. (2020). Monitoring of coral reefs using artificial intelligence: A feasible and cost-effective approach. *Remote Sensing, 12*(3), 489. https://doi.org/10.3390/rs12030489

González-Rivero, M., Bray, P., Jonker, M., and Ferrari, R. (2020) 3D habitat reconstructions of benthic communities. Long-Term Monitoring of the Great Barrier Reef. *Standard Operational Procedure 12*. Australian Institute of Marine Science, Townsville. https://doi.org/10.25845/m1f0-p935

Gordon, S., Aston, E., Lechene, M., Harianto, J., Bray, P., Figueira, W., Gonzalez Rivero, M., Ferrari Legorreta, R. (2023) Field photogrammetry in 4D. Reef Restoration and Adaption Program (EcoRRAP). Standard Operational Procedure Number 1: Overview and in-field workflow. Australian Institute of Marine Science, Townsville. https://doi.org/10.25845/SE7T-PS86

Gracias, N., and Santos-Victor, J. (2001). Underwater mosaicing and trajectory reconstruction using global alignment. *MTS/IEEE Oceans 2001. An Ocean Odyssey. Conference Proceedings* (IEEE Cat. No.01CH37295)*, 4*, 2557–2563. https://doi.org/10.1109/OCEANS.2001.968403

Gracias, N. R., van der Zwaan, S., Bernardino, A., and Santos-Victor, J. (2003). Mosaic-based navigation for autonomous underwater vehicles. *IEEE Journal of Oceanic Engineering, 28*(4), 609–624. https://doi.org/10.1109/JOE.2003.819156

Granshaw, S. I. (2019). Laussedat bicentenary: origins of photogrammetry. *The Photogrammietric Record, 34*(166), 128-147. https://doi.org/10.1111/phor.12277

Griffin, S. P., Nemeth, M. I., Moore, T. D., and Gintert, B. (2015). Restoration using *Acropora cervicornis* at the T/V MARGARA grounding site. *Coral Reefs, 34*(3), 885. https://doi.org/10.1007/s00338-015-1310-2

Habib, A. F., Kim, E.-M., and Kim, C.-J. (2007). New methodologies for true orthophoto generation. *Photogrammetric Engineering and Remote Sensing, 73*(1), 25–36. https://doi.org/10.14358/PERS.73.1.25

Hatcher, G. A., Warrick, J. A., Ritchie, A. C., Dailey, E. T., Zawada, D. G., Kranenburg, C., and Yates, K. K. (2020). Accurate bathymetric maps from underwater digital imagery without ground control. *Frontiers in Marine Science, 7*, 525. https://doi.org/10.3389/fmars.2020.00525

Helder, N. K., Burns, J. H. R., and Green, S. J. (2022). Intra-habitat structural complexity drives the distribution of fish trait groups on coral reefs. *Ecological Indicators, 142*, 109266. https://doi.org/10.1016/j.ecolind.2022.109266

Hernández-Landa, R. C., Barrera-Falcon, E., and Rioja-Nieto, R. (2020). Size-frequency distribution of coral assemblages in insular shallow reefs of the Mexican Caribbean using underwater photogrammetry. *PeerJ, 8*, e8957. https://doi.org/10.7717/peerj.8957

Huddell, H. D., Willett, J. C., and Marchand, G. (1974). Nearshore currents and coral reef ecology of the west coast of Guam, Mariana Islands. Naval Oceanographic Office, Technical Report AD-780 435.

Hughes, T. P. (1984). Population dynamics based on individual size rather than age: A general model with a reef coral example. *The American Naturalist, 123*(6), 778–795. https://doi.org/10.1086/284239

Huntington, B. E., and Miller, M. W. (2014). Location-specific metrics for rapidly estimating the abundance and condition of the threatened coral *Acropora cervicornis*. *Restoration Ecology, 22*(3), 299–303. https://doi.org/10.1111/rec.12057

James, M. R., Robson, S., d'Oleire-Oltmanns, S., and Niethammer, U. (2017). Optimising UAV topographic surveys processed with structure-from-motion: Ground control quality, quantity and bundle adjustment. *Geomorphology, 280*, 51–66. https://doi.org/10.1016/j.geomorph.2016.11.021

Jiang, R., Jáuregui, D. V., and White, K. R. (2008). Close-range photogrammetry applications in bridge measurement: Literature review. *Measurement, 41*(8), 823–834. https://doi.org/10.1016/j.measurement.2007.12.005

Jiang, S., Jiang, C., and Jiang, W. (2020). Efficient structure from motion for large-scale UAV images: A review and a comparison of SfM tools. *ISPRS Journal of Photogrammetry and Remote Sensing, 167*, 230–251. https://doi.org/10.1016/j.isprsjprs.2020.04.016

Johnson, M. E., Lustic, C., Bartels, E., Baums, I. B., Giliam, D. S., Larson, E. A., Lirman, D., Miller, M. W., Nedimyer, and K., Schopmeyer, S. (2011). Caribbean *Acropora* restoration guide: Best practices for propagation and population enhancement. The Nature Conservancy. Retrieved from https://nsuworks.nova.edu/occ_facreports/71

Jokiel, P. L., Rodgers, K. S., Brown, E. K., Kenyon, J. C., Aeby, G., Smith, W. R., and Farrell, F. (2015). Comparison of methods used to estimate coral cover in the Hawaiian Islands. *PeerJ, 3*, e954. https://doi.org/10.7717/peerj.954

Karlson, R. H., Cornell, H. V., and Hughes, T. P. (2007). Aggregation influences coral species richness at multiple spatial scales. *Ecology, 88*(1), 170–177. https://doi.org/10.1890/0012-9658(2007)88[170:AICSRA]2.0.CO;2

Kayal, M., Vercelloni, J., Wand, M. P., and Adjeroud, M. (2015). Searching for the best bet in life-strategy: A quantitative approach to individual performance and population dynamics in reef-building corals. *Ecological Complexity, 23*, 73–84. https://doi.org/10.1016/j.ecocom.2015.07.003

Kingsland, K. (2020). Comparative analysis of digital photogrammetry software for cultural heritage. Digital *Applications in Archaeology and Cultural Heritage, 18*, e00157. https://doi.org/10.1016/j.daach.2020.e00157

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., and Girshick, R. (2023). Segment anything. *arXiv:*2304.02643. https://doi.org/10.48550/arXiv.2304.02643

Kodera, S. M., Edwards, C. B., Petrovic, V., Pedersen, N. E., Eynaud, Y., and Sandin, S. A. (2020). Quantifying life history demographics of the scleractinian coral genus *Pocillopora* at Palmyra Atoll. *Coral Reefs, 39*(4), 1091–1105. https://doi.org/10.1007/s00338-020-01940-8

Kohler, K. E., and 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 and Geosciences, 32*(9), 1259–1269. https://doi.org/10.1016/j.cageo.2005.11.009

Kornicker, L. S., and Boyd, D. W. (1962). Shallow-water geology and environments of Alacran Reef Complex, Campeche Bank, Mexico. *Bulletin of the American Association of Petroleum Geologists, 46*(5). 640–673. https://doi.org/10.1306/BC743851-16BE-11D7-8645000102C1865D

Ladd, M. C., Miller, M. W., Hunt, J. H., Sharp, W. C., and Burkepile, D. E. (2018). Harnessing ecological processes to facilitate coral restoration. *Frontiers in Ecology and the Environment, 16*(4), 239–247. https://doi.org/10.1002/fee.1792

Lange, I. D., Molina-Hernández, A., Medellín-Maldonado, F., Perry, C. T., and Álvarez-Filip, L. (2022). Structure-from-motion photogrammetry demonstrates variability in coral growth within colonies and across habitats. *PLOS ONE, 17*(11), e0277546. https://doi.org/10.1371/journal.pone.0277546

Lechene, M. A. A., Haberstroh, A. J., Byrne, M., Figueira, W., and Ferrari, R. (2019). Optimising sampling strategies in coral reefs using large-area mosaics. *Remote Sensing, 11*(24), 2907. https://doi.org/10.3390/rs11242907

Lesser, M. P., and Slattery, M. (2021). Mesophotic coral reef community structure: The constraints of imagery collected by unmanned vehicles. *Marine Ecology Progress Series, 663*, 229–236. https://doi.org/10.3354/meps13650

Lewis, J. B. (1970). Spatial distribution and pattern of some Atlantic reef corals. *Nature, 227*(5263), 1158–1159. https://doi.org/10.1038/2271158a0

Lhuillier, M. (2012). Incremental fusion of structure-from-motion and GPS using constrained bundle adjustments. *IEEE Transactions on pattern analysis and machine intelligence, 34*(12), 2489–2495. https://doi.org/10.1109/TPAMI.2012.157

Li, J., and Asner, G. P. (2023). Global analysis of benthic complexity in shallow coral reefs. *Environmental Research Letters, 18*(2), 024038. https://doi.org/10.1088/1748-9326/acb3e6

Lindenberger, P., Sarlin, P.-E., Larsson, V., and Pollefeys, M. (2021). Pixel-perfect structure-from-motion with featuremetric refinement. *2021 IEEE/CVF International Conference on Computer Vision (ICCV),* pp. 5967–5977. https://doi.org/10.1109/ICCV48922.2021.00593

Lirman, D., Gracias, N., Gintert, B., Gleason, A. C. R., Deangelo, G., Dick, M., Martinez, E., and Reid, R. P. (2010). Damage and recovery assessment of vessel grounding injuries on coral reef habitats by use of georeferenced landscape video mosaics. *Limnology and Oceanography: Methods, 8*(3), 88–97. https://doi.org/10.4319/lom.2010.8.0088

Lirman, D., Gracias, N. R., Gintert, B. E., Gleason, A. C. R., Reid, R. P., Negahdaripour, S., and Kramer, P. (2007). Development and application of a video-mosaic survey technology to document the status of coral reef communities. *Environmental Monitoring and Assessment, 125*(1-3), 59–73. https://doi.org/10.1007/s10661-006-9239-0

Lirman, D., Schopmeyer, S., Galvan, V., Drury, C., Baker, A. C., and Baums, I. B. (2014). Growth dynamics of the threatened Caribbean staghorn coral *Acropora cervicornis*: Influence of host genotype, symbiont identity, colony size, and environmental setting. *PLOS ONE, 9*(9), e107253. https://doi.org/10.1371/journal.pone.0107253

Lochhead, I., and Hedley, N. (2021). Dry-lab benchmarking of a structure from motion workflow designed to monitor marine benthos in three dimensions. *The Photogrammetric Record, 36*(175), 224–251. https://doi.org/10.1111/phor.12370

Loya, Y. (1978). Plotless and transect methods. In Stoddart, D. R., and Johannes, R. E. (Eds), *Coral reefs: Research methods* (pp. 197–218). UNESCO Press.

Lyons, M. B., Roelfsema, C. M., Kennedy, E. V., Kovacs, E. M., Borrego-Acevedo, R., Markey, K., Roe, M., Yuwono, D. M., Harris, D. L., Phinn, S. R., Asner, G. P., Li, J., Knapp, D.E., Fabina, N. S., Larsen, K., Traganos, D., and Murray, N. J. (2020). Mapping the world's coral reefs using a global multiscale earth observation framework. *Remote Sensing in Ecology and Conservation, 6*, 557–568. https://doi.org/10.1002/rse2.157

Marhaver, K. L., Vermeij, M. J. A., Rohwer, F., and Sandin, S. A. (2013). Janzen-Connell effects in a broadcast-spawning Caribbean coral: distance-dependent survival of larvae and settlers. *Ecology, 94*(1), 146–160. https://doi.org/10.1890/12-0985.1

Marre, G., Holon, F., Luque, S., Boissery, P., and Deter, J. (2019). Monitoring marine habitats with photogrammetry: A cost-effective, accurate, precise and high-resolution reconstruction method. *Frontiers in Marine Science, 6,* 276. https://doi.org/10.3389/fmars.2019.00276

McCarthy, O. S., Smith, J. E., Petrovic, V., and Sandin, S. A. (2022). Identifying the drivers of structural complexity on Hawaiian coral reefs. *Marine Ecology Progress Series, 702*, 71–86. https://doi.org/10.3354/meps14205

McCarthy, O. S., Contractor, K., Figueira, W. F., Gleason, A. C., Viehman, T. S., Edwards, C. B., and Sandin, S. A. (2023). Closing the gap between existing large-area imaging research and marine conservation needs. *Conservation Biology*. https://doi.org/10.1111/cobi.14145. Epub ahead of print.

McNamara, D. E., Cortale, N., Edwards, C., Eynaud, Y., and Sandin, S. A. (2019). Insights into coral reef benthic dynamics from nonlinear spatial forecasting. *Journal of The Royal Society Interface, 16*(153), 20190047. https://doi.org/10.1098/rsif.2019.0047

Menna, F., Nocerino, E., and Remondino, F. (2017). Flat versus hemispherical dome ports in underwater photogrammetry. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-2/W3*, 481–487. https://doi.org/10.5194/isprs-archives-XLII-2-W3-481-2017

Million, W. C., O'Donnell, S., Bartels, E., and Kenkel, C. D. (2021). Colony-level 3D photogrammetry reveals that total linear extension and initial growth do not scale with complex morphological growth in the branching coral, *Acropora cervicornis*. *Frontiers in Marine Science, 8*, 646475. https://doi.org/10.3389/fmars.2021.646475

Mosbrucker, A. R., Major, J. J., Spicer, K. R., and Pitlick, J. (2017). Camera system considerations for geomorphic applications of SfM photogrammetry. *Earth Surface Processes and Landforms, 42*(6), 969–986. https://doi.org/10.1002/esp.4066

Murfitt, S. L., Allan, B. M., Bellgrove, A., Rattray, A., Young, M. A., and Ierodiaconou, D. (2017). Applications of unmanned aerial vehicles in intertidal reef monitoring. *Scientific Reports, 7*(1), 10259. https://doi.org/10.1038/s41598-017-10818-9

Nicosevici, T., Gracias, N., Negahdaripour, S., and Garcia, R. (2009). Efficient three-dimensional scene modeling and mosaicing. *Journal of Field Robotics, 26*(10), 759–788. https://doi.org/10.1002/rob.20305

Nocerino, E., Neyer, F., Gruen, A., Troyer, M., Menna, F., Brooks, A., Capra, A., Castagnetti, C., and Rossi, P. (2019). Comparison of diver-operated underwater photogrammetric systems for coral reef monitoring. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-2/W10*, 143–150. https://doi.org/10.5194/isprs-archives-XLII-2-W10-143-2019

Nouwakpo, S. K., Weltz, M. A., and McGwire, K. (2016). Assessing the performance of structure-from-motion photogrammetry and terrestrial LiDAR for reconstructing soil surface microtopography of naturally vegetated plots. *Earth Surface Processes and Landforms, 41*(3), 308–322. https://doi.org/10.1002/esp.3787

O'Connor, J., Smith, M. J., and James, M. R. (2017). Cameras and settings for aerial surveys in the geosciences. *Progress in Physical Geography: Earth and Environment, 41*(3), 325–344. https://doi.org/10.1177/0309133317703092

Pascoe, K. H., Fukunaga, A., Kosaki, R. K., and Burns, J. H. R. (2021). 3D assessment of a coral reef at Lalo Atoll reveals varying responses of habitat metrics following a catastrophic hurricane. *Scientific Reports, 11*(1), 12050. https://doi.org/10.1038/s41598-021-91509-4

Pavoni, G., Corsini, M., Callieri, M., Fiameni, G., Edwards, C., and Cignoni, P. (2020). On improving the training of models for the semantic segmentation of benthic communities from orthographic imagery. *Remote Sensing, 12*(18), 3106. https://doi.org/10.3390/rs12183106

Pavoni, G., Corsini, M., Ponchio, F., Muntoni, A., Edwards, C., Pedersen, N., Sandin, S., and 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://doi.org/10.1002/rob.22049

Perry, C. T., Edinger, E. N., Kench, P. S., Murphy, G. N., Smithers, S. G., Steneck, R. S., and Mumby, P. J. (2012). Estimating rates of biologically driven coral reef framework production and erosion: a new census-based carbonate budget methodology and applications to the reefs of Bonaire. *Coral Reefs, 31*(3), 853–868. https://doi.org/10.1007/s00338-012-0901-4

Petrovic, V., Vanoni, D., Richter, A. M., Levy, T. E., and Kuester, F. (2014). Visualizing high resolution three-dimensional and two-dimensional data of cultural heritage sites. *Mediterranean Archaeology and Archaeometry, 14*, 93–100. Retrieved from https://escholarship.org/uc/item/58z4j8k2

Pierce, J., Butler, M. J., Rzhanov, Y., Lowell, K., and 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*, 706674. https://doi.org/10.3389/fmars.2021.706674

Pieterkosky, S., Ziegwied, A., Cavanagh, C., and Thompson, L. (2017). BIV meets ASV: Bio-inspired fish drones and autonomous surface vehicles for coral reef monitoring. In *OCEANS 2017 - Anchorage,* pp. 1-5. IEEE. Anchorage, Alaska.

Pizarro, O., Eustice, R. M., and Singh, H. (2009). Large area 3-D reconstructions from underwater optical surveys. *IEEE Journal of Oceanic Engineering, 34*(2), 150–169. https://doi.org/10.1109/JOE.2009.2016071

Porter, J. W., Woodley, J. D., Jason Smith, G., Neigel, J. E., Battey, J. F., and Dallmeyer, D. G. (1981). Population trends among Jamaican reef corals. *Nature, 294*(5838), 249–250. https://doi.org/10.1038/294249a0

Pratchett, M. S., Anderson, K. D., Hoogenboom, M. O., Widman, E., Baird, A. H., Pandolfi, J. M., Edmunds, P. J., and Lough, J. M. (2015). Spatial, temporal and taxonomic variation in coral growth- implications for the structure and function of coral reef ecosystems. *Oceanography and Marine Biology: An Annual Review, 53*, 215–295.

Price, D. M., Lim, A., Callaway, A., Eichhorn, M. P., Wheeler, A. J., Lo Iacono, C., and Huvenne, V. A. I. (2021). Fine-scale heterogeneity of a cold-water coral reef and its influence on the distribution of associated taxa. *Frontiers in Marine Science, 8*, 556313. https://doi.org/10.3389/fmars.2021.556313

Raber, G. T., and Schill, S. R. (2019). Reef rover: A low-cost small autonomous unmanned surface vehicle (USV) for mapping and monitoring coral reefs. *Drones, 3*(2), 38. https://doi.org/10.3390/drones3020038

Rodriguez, C., Amir, C., Gray, A., Asbury, M., Suka, R., Lamirand, M., Couch, C. S., Oliver, T. (2021). *Extracting coral vital rate estimates at fixed sites using structure-from-motion standard operating procedures*. NOAA National Marine Fisheries Service, Pacific Islands Fisheries Science Center. NOAA Technical Memorandum NMFS-PIFSC-120. https://doi.org/10.25923/a9se-k649.

Runyan, H., Petrovic, V., Edwards, C. B., Pedersen, N., Alcantar, E., Kuester, F., and Sandin, S. A. (2022). Automated 2D, 2.5D, and 3D segmentation of coral reef pointclouds and orthoprojections. *Frontiers in Robotics and AI, 9*, 884317. https://doi.org/10.3389/frobt.2022.884317

Sandin Lab, Scripps Institution of Oceanography, UC San Diego (2023). *Standard Operating Procedure Documents for Coral Reef Ecological Monitoring. In Center for Environmental Imaging Collection*. UC San Diego Library Digital Collections. https://doi.org/10.6075/J0Z31ZTW

Sandin, S. A., Edwards, C. B., Pedersen, N. E., Petrovic, V., Pavoni, G., Alcantar, E., Chancellor, K. S., Fox, M. D., Stallings, B., Sullivan, C. J., Rotjan, R. D., Ponchio, F., and Zgliczynski, B. J. (2020). Considering the rates of growth in two taxa of coral across Pacific islands. *Advances in Marine Biology, 87*(1), 167–191. https://doi.org/10.1016/bs.amb.2020.08.006

Shamshiri, R., and Ismail, W. I. W. (2013). Exploring GPS data for operational analysis of farm machinery. *Research Journal of Applied Sciences, Engineering and Technology, 5*(12), 3281–3286. https://doi.org/10.19026/rjaset.5.4568

Snavely, N., Seitz, S. M., and Szeliski, R. (2008). Modeling the world from internet photo collections. *International Journal of Computer Vision, 80*(2), 189–210. https://doi.org/10.1007/s11263-007-0107-3

Sona, G., Pinto, L., Pagliari, D., Passoni, D., and Gini, R. (2014). Experimental analysis of different software packages for orientation and digital surface modelling from UAV images. *Earth Science Informatics, 7*(2), 97–107. https://doi.org/10.1007/s12145-013-0142-2

Suka, R., Asbury, M., Couch, C., Gray, A., Winston, M., and Oliver, T. (2019). *Processing photomosaic imagery of coral reefs using structure-from-motion standard operating procedures*. NOAA National Marine Fisheries Service, Pacific Islands Fisheries Science Center. NOAA Technical Memorandum NMFS-PIFSC-93. https://doi.org/10.25923/h2q8-jv47

Swanborn, D. J. B., Stefanoudis, P. V., Huvenne, V. A. I., Pittman, S. J., and Woodall, L. C. (2022). Structure-from-motion photogrammetry demonstrates that fine-scale seascape heterogeneity is essential in shaping mesophotic fish assemblages. *Remote Sensing in Ecology and Conservation, 8*(6), 904–920. https://doi.org/10.1002/rse2.290

Thoeni, K., Giacomini, A., Murtagh, R., and Kniest, E. (2014). A comparison of multi-view 3D reconstruction of a rock wall using several cameras and a laser scanner. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XL-5*, 573–580. https://doi.org/10.5194/isprsarchives-XL-5-573-2014

Urbina-Barreto, I., Garnier, R., Elise, S., Pinel, R., Dumas, P., Mahamadaly, V., Facon, M., Bureau, S., Peignon, C., Quod, J.-P., Dutrieux, E., Penin, L., and Adjeroud, M. (2021). Which method for which purpose? A comparison of line intercept transect and underwater photogrammetry methods for coral reef surveys. *Frontiers in Marine Science, 8*, 636902. https://doi.org/10.3389/fmars.2021.636902

Weinberg, S. (1981). A comparison of coral reef survey methods. *Bijdragen tot de Dierkunde, 51*, 199–218.

Westoby, M. J., Brasington, J., Glasser, N. F., Hambrey, M. J., and Reynolds, J. M. (2012). 'Structure-from-Motion' photogrammetry: A low-cost, effective tool for geoscience applications. *Geomorphology, 179*, 300–314. https://doi.org/10.1016/j.geomorph.2012.08.021

Winston, W., Couch, C., Huntington, B., and Vargas-Ángel, B. (2020). *Ecosystem Sciences Division standard operating procedures: data collection for rapid ecological assessment benthic surveys, 2019 update*. NOAA National Marine Fisheries Service, Pacific Islands Fisheries Science Center. NOAA Technical Memorandum NMFS-PIFSC-96. https://doi.org/10.25923/ws5s-km69

Zawada, K. J. A., Madin, J. S., Baird, A. H., Bridge, T. C. L., and Dornelas, M. (2019). Morphological traits can track coral reef responses to the Anthropocene. *Functional Ecology, 33*(6), 962–975. https://doi.org/10.1111/1365-2435.13358

Zvuloni, A., Artzy-Randrup, Y., Stone, L., Kramarsky-Winter, E., Barkan, R., and Loya, Y. (2009). Spatio-temporal transmission patterns of black-band disease in a coral community. *PLOS ONE, 4*(4), e4993. https://doi.org/10.1371/journal.pone.0004993

U.S. Department of Commerce

Gina M. Raimondo, *Secretary*

National Oceanic and Atmospheric Administration

Richard Spinrad, *Under Secretary for Oceans and Atmosphere*

National Ocean Service

Nicole LeBoeuf, *Assistant Administrator for National Ocean Service*

The mission of the National Centers for Coastal Ocean Science is to provide managers with scientific information and tools needed to balance society's environmental, social and economic goals. For more information, visit http://www.coastalscience.noaa.gov/.