https://doi.org/10.25923/337h-k518

Uncrewed Aircraft Systems, Machine Learning, and Polarimetric Imaging for Enhanced Marine Debris Shoreline Surveys

Final Project Report

Prepared for
NOAA Uncrewed Systems Research Transition Office



Authors

Christopher E. Parrish, W. Ross Winans, Tim Battista, Amy V. Uhrin,
Kyle Herrera, Peter Murphy, Chase Simpson, and Richard Slocum

May 2023

NOAA TECHNICAL MEMORANDUM NOS NCCOS 312

NOAA NOS National Centers for Coastal Ocean Science

Uncrewed Aircraft Systems, Machine Learning, and Polarimetric Imaging for Enhanced Marine Debris Shoreline Surveys

Final Project Report

Prepared for
NOAA Uncrewed Systems Research Transition Office

Authors
Christopher E. Parrish[1], W. Ross Winans[2], Tim Battista[3], Amy V. Uhrin[4], Kyle Herrera[1], Peter Murphy[5], Chase Simpson[1], and Richard Slocum[6]

[1] *Oregon State University, School of Civil and Construction Engineering*
[2] *ORBTL AI, LLC*
[3] *NOAA National Ocean Service, National Centers for Coastal Ocean Science, Marine Spatial Ecology Division*
[4] *NOAA National Ocean Service, Office of Response and Restoration, Marine Debris Program*
[5] *Genwest, under contract to NOAA National Ocean Service, Office of Response and Restoration, Marine Debris Program*
[6] *ARGO AI, LLC*

Prepared by
National Oceanic and Atmospheric Administration
National Ocean Service
National Centers for Coastal Ocean Science
Marine Spatial Ecology Division
1305 East-West Hwy, SSMC-4
Silver Spring, MD 20910

May 2023

NOAA Technical Memorandum NOS NCCOS 312

Suggested Citation

Parrish, C.E., Winans, W.R., Battista, T., Uhrin, A.V., Herrera, K., Murphy, P., Simpson, C., and Slocum, R. 2023. Uncrewed Aircraft Systems, Machine Learning, and Polarimetric Imaging for Enhanced Marine Debris Shoreline Surveys. Final Report. NOAA Technical Memorandum NOS NCCOS 312. Silver Spring, MD. 31 pp. doi: 10.25923/337h-k518

Acknowledgments

The project team gratefully acknowledges contributions to this project from Professor Mike Starek, Texas A&M University–Corpus Christi; Jake Berryhill, Texas A&M University–Corpus Christi; Tracy Weatherall, University of Texas Marine Science Institute; LT Lisa Campbell, U.S. Coast Guard; and Dr. Andrew Lassiter, Oregon State University.

Images throughout the report are used by courtesy of Oregon State University and NOAA NCCOS.

For more information on NOAA's National Centers for Coastal Ocean Science, please visit:
https://coastalscience.noaa.gov/

For more information on this project, please visit:
https://coastalscience.noaa.gov/project/using-unmanned-aircraft-systems-machine-learning-and-polarimetric-imaging-to-develop-a-system-for-enhanced-marine-debris-detection-and-removal/

Or direct questions and comments to:

Tim Battista
tim.battista@noaa.gov
NOAA National Centers for Coastal Ocean Science
1305 East West Highway, SSMC 4
Silver Spring, MD 20910

Disclaimer

Compiled marine debris on Padre Island, Texas. Credit: NOAA NCCOS.

Table of Contents

List of Figures and Tables

Glossary

AGL – above ground level

AI – artificial intelligence

AOLP – angle of linear polarization

CMOS – complementary metal-oxide-semiconductor

CSV – comma-separated values file

DOLP – degree of linear polarization

DSLR – digital single-lens reflex camera

endlap – the extent of overlap of successive photos along a flight strip

EXIF – exchangeable image file format

FAA – Federal Aviation Administration

focal length – the distance from a lens's rear nodal point to the point at which objects at an infinite distance come into focus

GIS – geographic information system

GNSS – global navigation satellite system

GSD – ground sample distance. The size of an image pixel projected onto the ground. It is a frequently used measure of resolution of an imaging system

image blur – blurring of imagery due to, e.g., motion of camera during time of exposure

K-NN – k-nearest neighbors supervised classification algorithm

MDMAP – Marine Debris Monitoring and Assessment Project

MDP – Marine Debris Program

ML – machine learning

MP – megapixel (one million pixels)

OPRD – Oregon Parks and Recreation Department

orthomosaic – an assemblage of multiple orthoimages (see orthoimage definition) to create a composite or "mosaic"

orthoimage– a vertical aerial photograph from which the distortions due to varying elevation, tilt, and surface topography have been removed, such that it can be used as a planimetric map (distances and areas can be measured directly)

PI – polarimetric imaging

PPK – post-processed kinematic GNSS

RGB – red, green, blue (e.g., a camera that images in RGB bands to create natural color imagery)

RTK – real-time kinematic GNSS

SDK – software development kit

sidelap – extent of lateral overlap between images acquired on adjacent flight lines

SfM – structure from motion, a relatively new type of photogrammetry with origins in the field of computer vision

TD – transformed divergence, a standard class separability index in remote sensing

UAS – uncrewed aircraft systems

UxS – uncrewed systems

VLOS – visual line of sight

Abstract

Marine debris is a pervasive and escalating environmental problem that poses health risks to both wildlife and humans, as well as degradation to our oceans and coasts. Detection and categorization of shoreline debris are currently conducted through field surveys that entail walking shoreline transects. Uncrewed aircraft systems (UAS), machine learning (ML), and polarimetric imaging (PI) are three rapidly emerging technologies with the capability to dramatically enhance marine debris shoreline surveys. This report presents findings and recommendations based on a NOAA Uncrewed Systems Research Transition Office–funded project conducted collaboratively by NOAA's National Centers for Coastal Ocean Science (NCCOS), the NOAA Marine Debris Program (MDP), Oregon State University (OSU), and ORBTL AI. Specific outcomes from the project included: 1) operationally efficient field workflows for collection of true color and polarimetric imagery from UAS; 2) deployable algorithms for automatic marine debris detection from remotely sensed true color imagery; and 3) recommended operating procedures for collecting and processing UAS imagery and communicating the locations, types, and confidence scores of potential detections of marine debris to stakeholders. This report serves two primary purposes. First, it summarizes the project methods, results, and key findings. Second, it provides guidance to stakeholders and project partners on operational use of UAS and ML for marine debris surveys.

Field testing on Texas barrier island. Credit: NOAA NCCOS.

Section 1 Introduction

The Marine Debris Research, Prevention, and Reduction Act (33 U.S.C. 33A) signed into law in 2006 and amended in 2012, 2018, and 2020, established the Marine Debris Program (MDP) with requirements to "identify, determine sources of, assess, prevent, reduce, and remove marine debris and address the adverse impacts of marine debris on the economy of the United States, marine environment, and navigation safety." Part of that role includes emergency response functions, which often necessitate rapid detection of shoreline marine debris and planning for removal by MDP partners. Thus, the ability to accurately and quickly detect and identify marine debris over large areas (tens of square kilometers), as well as areas difficult to access by foot or boat (e.g., rocky headlands, marshlands, remote coastal sites, and embayments) is crucial.

Aerial imagery collection, whether via crewed aircraft or uncrewed aircraft systems (UAS), has been the primary method for detecting large marine debris across sizable areas of shoreline and open water (Veenstra and Churnside, 2012; Brooke et al., 2015; Moy et al., 2018; Lebreton et al. 2018). Uncrewed systems (UxS) such as UAS can provide multiple advantages over crewed systems. Compared with crewed aircraft, UAS are favorable for conducting focused, small-area debris surveys with lower operational and maintenance costs, reduced technical complexity, potentially higher revisit frequencies, portability, rapid response, and easier pilot certification. Previous research has documented the benefits of UxS for marine debris mapping (Hengstmann et al., 2017; Bao et al., 2018; Martin et al., 2018; Fallati et al., 2019; Gonçalves et al., 2020); however, relatively few efforts to date have resulted in operationally viable procedures and workflows suitable for implementation by the broader marine debris community (but see Gonçalves et al., 2020).

A focus of this project was on demonstrating low-cost, easily deployed, and accurate marine debris detection workflows and tools, so as to operationalize the use of UAS by transferring the technological solutions to existing marine debris monitoring teams. A key goal was to enable more efficient periodic monitoring by providing a repeatable, systematic method of quantifying abundance and distribution of debris in survey areas to support decision-making. The scope of this work focused on routine shoreline debris surveys and did not specifically address disaster (e.g., post-hurricane) debris. However, many of the procedures developed and tested in this work may also be applicable to emergency response, and the extension of this work to such debris incidents is a recommended topic for future research.

Shoreline debris on Padre Island, Texas. Credit: NOAA NCCOS.

The project leveraged procedures, workflows, and UxS guidelines completed by the project team in previous work (Slocum et al., 2019) but also integrated new technological innovations specifically tailored to address MDP partner challenges and use cases. The novel approaches that were investigated included 1) evaluation of a polarimetric imaging (PI) camera for enhanced debris detection, as compared with conventional red-green-blue (RGB) cameras; and 2) development and testing of machine learning (ML) algorithms for automated detection and type classification of debris.

PI cameras are a relatively new type of sensor, which may greatly facilitate marine debris detection and classification. PI cameras go beyond the capabilities of typical multispectral and hyperspectral cameras, which have been widely used in a range of coastal remote sensing applications for several decades, by providing information regarding the polarization state of received light. Because different types of objects differ in how sunlight reflected from them is polarized, PI cameras can assist in obtaining information about object types, including characteristics such as surface roughness. Importantly, polarimetry has been found beneficial for detecting human-made objects and for distinguishing objects from their background (e.g., Denes et al., 1998; Islam et al., 2019). For these reasons, PI cameras are of interest for the detection and classification of marine debris on shorelines, where it is often intermixed with substrate, vegetation, and/or natural (woody) debris.

To date, identification and classification of marine debris from digital images captured by infrastructure-mounted (e.g., bridge or jetty) or UxS-mounted cameras have largely

been performed via simple spectral threshold classification algorithms that do not scale across diverse background scenes (Bao, 2018) or via direct visual inspection and assessment of collected imagery (Moy et al., 2018). Both of these methods can be labor-intensive, especially across larger survey areas, making it difficult to rapidly produce results. ML is a sub-discipline of artificial intelligence (AI) in which algorithms are trained to make classifications or predictions through an iterative process involving the use of training data to build and improve classification effectiveness and outcomes, not unlike that of the human brain. Development of these automated techniques is challenging, as the training data require representative samples of every feature that one wishes to properly classify or predict. As the size of the training dataset increases, so does ML performance. ML is currently under exploration for automated detection of marine debris from various environmental backgrounds with promising results (Martin et al., 2018; Gonçalves et al., 2020; Kylili et al., 2020; van Lieshout et al., 2020; Wolf et al., 2020; Martin et al., 2021). The ability to automatically process large collections of images spanning expansive areas of shoreline or surface water offers substantial potential for wide-area debris detection and monitoring. In this project, an ML model was developed and tested for automatically detecting debris items. When georeferenced imagery is used as an input, the output can be used in geographic information system (GIS) software to generate heatmaps representing concentrations of marine debris (number of debris items per unit area), which can guide cleanup and restoration decisions and actions.

Collectively, the workflows and software tools developed during this project serve to provide automatic and efficient detection, identification, and geolocation of debris items through an objective and repeatable process. In contrast to traditional shoreline monitoring techniques, which typically rely upon random transects and visual observations, UAS imagery can be acquired over large areas of shoreline and quickly processed using the procedures, so as to quantify specific debris item locations and distribution patterns over larger spatial areas. The debris geolocation information obtained through these workflows allows debris specialists to quickly quantify debris densities and distribution, which, in turn, can strengthen temporal change analysis from repeat surveys conducted in the same locations. The georeferenced imagery also provides a benefit over traditional monitoring methods in providing a digital record that can be revisited for additional analysis at any later time. Lastly, this approach represents a significant paradigm shift in that the allocation of personnel between field and office time in the end-to-end workflow is shifted. Specifically, the field time is reduced through highly efficient data acquisition procedures, enabling more of the analysis to be done in the office. Debris type identification is conducted efficiently through use of the imagery products collected and ML during post-processing after returning from the field.

Due to the overarching goals of this project, this report serves two purposes: 1) to document the project methods, results, and key findings and 2) to provide operational guidelines and information on how best to conduct marine debris shoreline surveys using UAS and how to use the ML model developed and tested in this work to perform auto-detection and classification of marine debris items. Readers who are more interested in the operational recommendations and workflows than in the methods and results of the study can skip directly to Section 5.

The research portion of this project was conducted through a multi-phased study designed to achieve the goals noted above, namely to investigate UAS, ML, and PI cameras for marine debris shoreline surveys. The following sections describe the methods for each phase, starting with the equipment used.

Marine debris on Padre Island, Texas. Credit: NOAA NCCOS..

Polarimetric imaging (PI) camera in helicopter over Padre Island, Texas.
Credit: NOAA NCCOS.

Section 2 Methods

2.1 Equipment

2.1.1 Uncrewed Systems

An underlying objective in this project was to keep our procedures relatively invariant to the specific make and model of remote aircraft since, due to the rapid technological maturation of UAS, any platform available today is likely to be obsolete within a few years. An additional objective was to test only platforms that are commercially available and that could be readily procured by NOAA MDP project partners who conduct shoreline debris surveys. The project team also sought to test remote aircraft of different broad types (e.g., multirotor and fixed wing), with different payload capacities and endurances (maximum flight times). With these objectives in mind, four different UAS were tested in this study: 1) a Skydio 2™ (Figure 1a), 2) a DJI Phantom™ 4 Pro (Figure 1b), 3) a Wingtra AG WingtraOne™ (Figure 1c), and 4) a DJI Matrice™ 300 (Figure 1d).

The Skydio 2 is a small (775 g), highly portable quadcopter, manufactured in the U.S., with a base price (at time of purchase) of around $1,000. It has a built-in, gimbal-mounted 12.3-MP (4056 x 3040 pixels) RGB camera, based on the Sony IMX577 complementary metal-oxide-semiconductor (CMOS) camera chip, and a 3.7-mm focal length, f/2.8 lens, as well as six navigation cameras used for obstacle avoidance. (It should be noted that, throughout this report, all focal lengths are reported as the actual focal lengths of the camera lenses, as opposed to the so-called “35-mm equivalent focal length,” which is often reported by manufacturers of small-format cameras, such as those integrated on small UxSs. If any focal length in this report appears to differ significantly from that provided on a specifications sheet from the manufacturer, it is likely because the manufacturer is reporting a 35-mm equivalent focal length. This distinction is further explained in Verhoeven, 2007.) Since the purchase of this aircraft, Skydio has released a new model, the Skydio 2+™, with extended range and flight time.

The DJI PHANTOM 4 Pro is a 1.39-kg quadcopter, equipped with a gimbal-mounted 20-MP (5472 x 3648 pixels) RGB camera with an 8.6-mm focal length lens and adjustable f/2.8–f/11 aperture. The real-time kinematic (RTK) version of the Phantom 4 Pro contains a carrier-phase recording global navigation satellite system (GNSS) receiver, capable of supporting RTK or post-processed kinematic (PPK) GNSS positioning, with accuracies on the order of a few centimeters. Although already two years old at the start of the project, the Phantom 4 Pro RTK used in this study, purchased through a third-party reseller, had an original purchase price of around $7,500.

The Wingtra One (Wingtra AG, Switzerland) is a hybrid aircraft that performs vertical takeoff and landing (VTOL) and converts to a fixed-wing configuration in flight. It stands upright on the aircraft tail for takeoff and landing, in the configuration shown in Figure 1c. It has a wingspan of 1.25 m, weighs 3.7 kg empty, and has a maximum payload of 0.8 kg. The camera is a 42-MP Sony RX1R II with a 35-mm

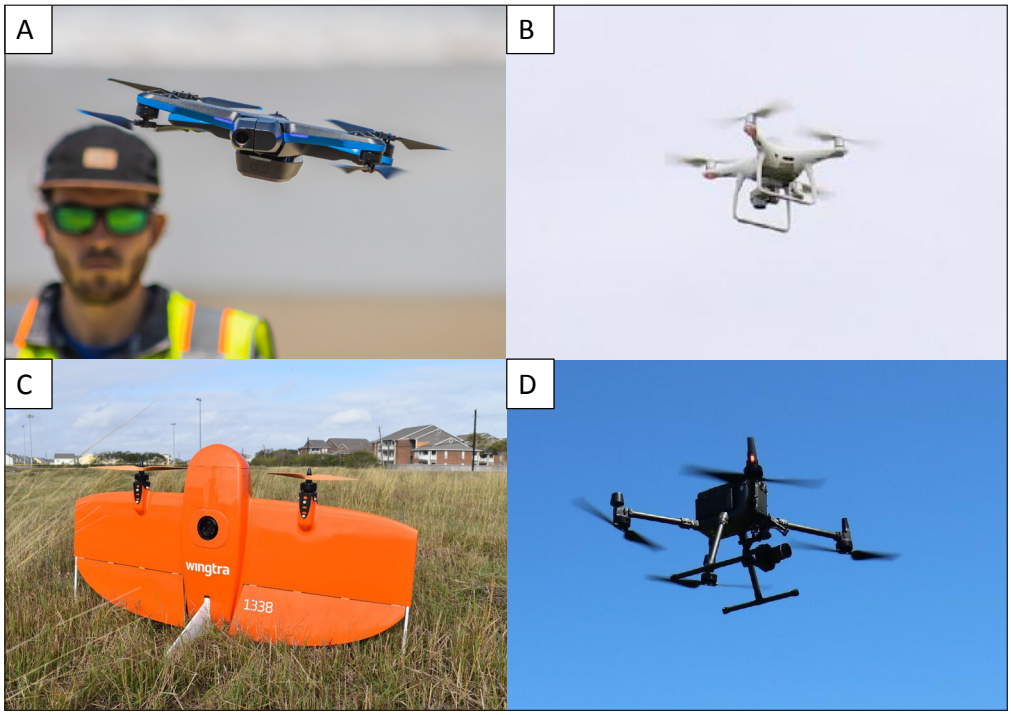

Figure 1: Unmanned aircraft systems (UAS) used in this project: (a) Skydio 2™, operated by OSU graduate student Kyle Herrera; (b) DJI Phantom™ 4 Pro; (c) Wingtra AG WingtraOne™; and (d) DJI Matrice™ 300. Credit: NOAA NCCOS (a, b, c) and Oregon State University (d).

focal length lens. Other noteworthy characteristics of this aircraft include very high endurance, with maximum flight times of nearly an hour, and a carrier-phase recording, multi-frequency, multi-constellation GNSS supporting RTK and PPK operations.

The final remote aircraft utilized in this study is the DJI Matrice 300 RTK. This quadcopter has a larger airframe (6.3 kg, including batteries) than the others tested in this project. The integrated camera is a 45-MP DJI ZENMUSE™ P1. Importantly for this project, it has a total payload capacity of 2.7 kg and supports simultaneous integration of up to three separate payloads. This capability was important for flight testing the PI camera (described below), due to the need to integrate both the camera and single-board computer on the UAS. The cost (at time of purchase) of this aircraft, including the camera and peripheral equipment, was approximately $13,000.

2.1.2 Polarimetric Imaging Camera

The PI camera procured for and investigated in this work is the FLIR Blackfly® S USB3 (RGB; Figure 2). This 5.0-MP polarimetric camera is based on the Sony IMX250MZR camera chip. Its size (<100 g, without the lens) and form factor make it well suited for integration on a UAS. Another important factor for this project is that a software development kit (SDK) was available for this PI camera. Additionally, the frame rate of 75 frames/sec was determined to be advantageous for this project. In this project, a Fujinon 12.5-mm focal length C-mount lens was used on the Blackfly-S camera. Including the lens and all components, the 2020 purchase price of this camera was around $2,700.

2.2 Field Sites

To achieve the objectives of this project, four separate field sites were selected, providing a wide range of environmental conditions and characteristics. These sites are described below.

2.2.1 Hinsdale Wave Lab

In the first phase of the study, a controlled test site on the grounds of the O.H. Hinsdale Wave Research Laboratory on the Oregon State University (OSU) campus was created for acquiring and testing marine debris imagery under varying conditions. This site contained leftover beach materials (e.g., sand and vegetation) from previous experiments, which were seeded with everyday household trash items meant to simulate marine debris. While artificial, this site enabled extensive testing of camera settings and imagery analysis in an easy-to-access, contained area. The location of this

Figure 2: FLIR Blackfly-S camera used in this project (Herrera, 2021).

site on the OSU campus was particularly beneficial, as it enabled imagery to be collected even during the winter, which is typically quite rainy in Oregon, through strategic use of brief periods of sunshine for data acquisition.

2.2.2 Neptune State Scenic Viewpoint

The next field site used for testing UAS imagery and PI camera imagery was the Neptune State Scenic Viewpoint (hereafter, simply Neptune) on the Oregon Coast (Figures 3 and 4). This site was selected both for its proximity to the OSU campus and its range of beach types, including sand, cobble, rock outcrop, steep bluff, and freshwater outflow from Cummins Creek, all within a 0.2-km^2 area. The work was conducted under Oregon Parks and Recreation Department (OPRD) permit #176. Because the Oregon coast is largely free of debris, the site was again seeded with debris items, including a mix of the household debris items used at the Hinsdale Wave Lab and some actual marine debris items obtained through working with OPRD.

2.2.3 Texas Barrier Islands

In the next phase of the study, the procedures developed at the Hinsdale Wave Lab and at Neptune were validated using data collected on the barrier islands of the Texas coast east of Corpus Christi December 8–21, 2021. Specific sites included areas on Padre Island and San José Island (Figure 5). These sites are known hotspots for marine debris accumulation due to prevailing winds, currents, and circulation patterns within the Gulf of Mexico, as well as the orientation of the coast. The barrier islands contain wide (approximately 40 m from shoreline to backing dune), sandy, dune-backed beaches. The two primary UAS utilized in the Texas data collection were the Skydio 2 and a Phantom 4 Pro (although not the RTK version). Additionally, the WingraOne hybrid UAS, owned and operated by project partners at Texas A&M University–Corpus Christi, was investigated during the Texas field data collection.

Figure 3: Neptune State Scenic Area test site on the Oregon Coast.

Figure 4: UAS imagery acquisition at the Neptune project site. Credit: Oregon State University.

Figure 5: Texas barrier island sites.

2.2.4 OSU Campus (PI UAS Tests)

The tests of the Blackfly-S camera on a UAS were conducted over a seeded debris field on the northeast end of the OSU campus (Figure 6). No special characteristics were required for these tests. Hence, the site selected was a large grass field frequently used by the OSU project team for UAS testing.

Figure 6: Seeded debris field on OSU campus. Credit: Oregon State University.

Debris in sand on a Texas barrier island. Credit: NOAA NCCOS.

Section 3 Experiments

3.1 Preliminary Testing

Preliminary testing was conducted at the Hinsdale Wave Lab site on the OSU campus in January 2021. These tests involved using the Blackfly-S PI camera on a pole mount to collect imagery over the simulated beach, which was seeded with household trash items (primarily plastics and foam), intended to simulate marine debris (Figure 7).

Figure 7: Pole-mounted PI camera set up at Hinsdale Wave Lab. Here the FLIR Blackfly S USB3 PI camera is operated from a pole in a weighted stand and connected to a field laptop (out of the photo, to the left). Credit: Oregon State University.

Simultaneously with the Hinsdale Wave Laboratory testing, project team member Ross Winans (ORBTL AI) extended a deep-learning model developed using TensorFlow during his MS thesis research at the University of Hawai'i at Mānoa (Winans, 2021) for use in this project. This work included building a deep-learning–based object-detection model delivered via an application programming interface (API) that is bundled for installation within a Docker container (Cito et al., 2017).

3.2 Field Testing

Field tests were conducted at the Neptune project site on May 4 and July 8, 2021. The primary goals of the tests at Neptune were: 1) to test imagery acquisition with UAS; 2) to test the PI camera; and 3) to test the ability to use the UAS imagery in the ML model, which was originally trained on imagery acquired in Hawai'i. While the PI camera used in this study, the FLIR Blackfly-S (Figure 2), had a sufficiently small form factor, weight, and power consumption for installation on a UAS, imagery acquisition with this camera was not sufficiently automated to enable it to be operated on a UAS in this project phase. In particular, the exposure settings for the camera needed to be set manually and adjusted frequently during acquisition. For this reason, the PI camera imagery acquired at Neptune and Texas was collected with the camera mounted on a pole.

Throughout the experiments, UAS imagery was typically acquired with a 2-cm ground sample distance (GSD) for multiple reasons. First, this GSD matches the training data used to train the ML model, which was adapted from a previous aerial debris survey conducted over the Hawaiian Island coastline on 14 separate days between August and October 2015 (Moy et al., 2018). In the Hawai'i imagery acquisition (not part of this project), a Cessna 206 aircraft equipped with two high-resolution digital single-lens reflex (DSLR) cameras (Canon EOS 5DS R) and one medium format aerial camera (Phase One P65+) was used. All camera systems were mounted on gyro-stabilized gimbals and paired with real-time differential GPS obtained through an OmniSTAR satellite-based augmentation system (SBAS). This system allowed photos to be corrected for multiple types of distortion, resulting in a within-model root-mean-square error (RMSE) of 1.5 m without the need for field-surveyed ground control points (GCPs; Moy et al., 2018). Second, this study focused on macro-debris items ≥ 2.5 cm, as this is the lower size threshold for visual shoreline monitoring methods as part of MDP's Marine Debris Monitoring and Assessment Project (MDMAP). Additionally, previous experimental work suggests that the high imagery resolutions of most UAS-deployed cameras (approximately 1.0–5.5 cm) allow for the identification of debris items in the meso- to macro-debris size range (Bao et al., 2018; Gonçalves et al., 2020). However, because a goal of the study was to test varying flight altitudes, some imagery was acquired with better (smaller) than 2-cm GSD and subsequently downsampled to 2 cm for use in the ML model. In addition to flight altitude, other acquisition variables tested during the experiments included: solar altitude, flight direction with respect to the solar azimuth, platform, and camera type.

3.3 Field Validation and Workflow Refinement

The next phase of the project entailed conducting operational tests on the Texas barrier islands sites. The fieldwork for this portion of the study was conducted December 6–17, 2021, in collaboration with project partners at Texas A&M University–Corpus Christi. Imagery was collected with the RGB cameras on the Skydio 2 and Phantom 4 Pro. Key objectives of the Texas fieldwork were: 1) to test the procedures used at the Neptune site in a location known to be a hotspot for marine debris, containing numerous actual (i.e., not seeded) debris items; and 2) to test the ability to upload the UAS imagery and run the ML model while in the field.

As in the tests at Neptune, it was not possible during the Texas barrier islands fieldwork to operate the PI camera from a UAS, due to not having developed the necessary hardware and software to operate the camera (and, in particular, to automatically adjust exposure settings) from a remote aircraft. For this reason, PI imagery was again collected primarily using a pole mount (Figure 8). To facilitate efficient acquisition of imagery with the pole-mounted PI camera, high-density debris fields were created by collecting debris items from within approximately one square kilometer and concentrating them in fields of a few tens of square meters (Figure 9).

While the tests of the PI camera over both naturally occurring debris and the high-density debris fields were successful, it was also of interest to test the PI camera from an airborne platform. The main goals were to investigate potential motion blur in the imagery, due to the moving platform, and the ability to obtain sufficient spatial resolution from a height more typical of a UAS flying height. Through a partnership with the U.S. Coast Guard (USCG), permission was obtained to operate the PI camera from a USCG MH65 Echo helicopter to simulate UAS PI imagery acquisition (Figure 10).

The experiment, led by OSU graduate student Kyle Herrera, entailed generating additional image bands from the polarimetric data. These bands included Stokes parameters (Stokes, 1852; Yan et al., 2020), as well as angle of linear polarization (AOLP) and degree of linear polarization (DOLP), which are computed from the Stokes parameters. The resulting images were visually assessed to determine whether the PI-derived information facilitated manual detection of debris items. Next, the transformed divergence (TD), a well-known separability index in remote sensing, was used to evaluate the separability of debris classes, within and without the PI-derived information included. Finally, a k-nearest neighbors (K-NN) supervised image classification was performed, with and without the PI-derived bands included, to assess the potential to increase classification accuracy through inclusion of the PI bands.

Figure 8: Pole-mounted PI camera collecting imagery of debris on the Texas barrier islands. Credit: NOAA NCCOS.

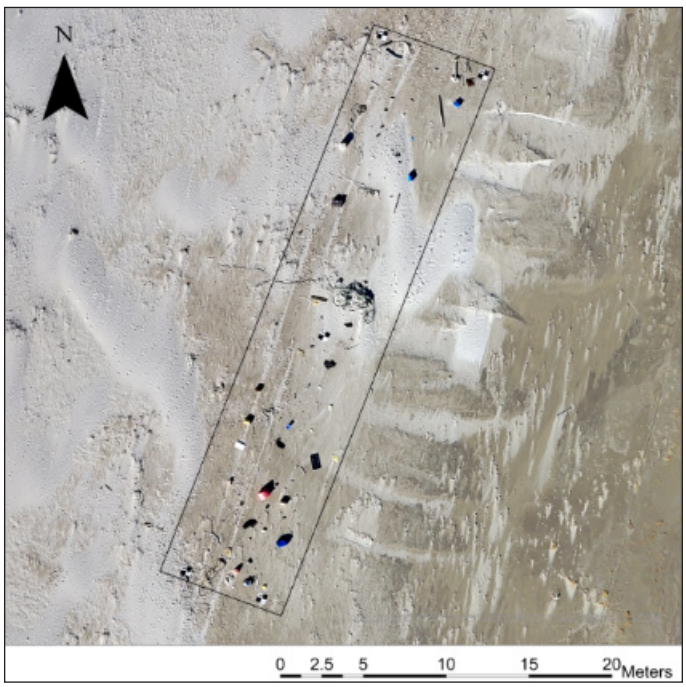

Figure 9: Seeded debris field on San José Island, Texas. Credit: Oregon State University.

Figure 10: Tests of PI camera on USCG MH65 Echo helicopter to simulate UAS polarimetric imagery acquisition. Credit: NOAA NCCOS.

3.4 UAS PI Camera Integration and Testing

Due to the fact that a computer is required to operate the Blackfly-S camera and adjust exposure settings during operations, the previous phases of the project involved manually operating the camera installed on a pole mount or USCG helicopter and connected to a laptop computer, as noted above. This phase of the project sought to address these challenges and demonstrate automated acquisition of PI imagery from a UAS. The specific goals of this project phase were to:

1. Provide hardware and software to acquire imagery using the FLIR Blackfly-S polarimetric camera and a Raspberry Pi single-board computer.
2. Demonstrate the UAS payload by acquiring data over a simulated beach scene seeded with debris items.
3. Design and test a proof-of-concept PI technique using a standard RGB camera (non-PI) with a polarizing filter, and collecting sequential, rotated images, from which PI imagery can be generated using image processing.

Software was written for a Raspberry Pi 4B running Ubuntu 20.04 to communicate with the Blackfly-S camera using the Python FLIR Spinnaker SDK. Raw imagery was logged to a USB drive and post-processed into AOLP and DOLP images using custom Python scripts, which were made publicly available on a GitHub repository.

The imaging system was integrated on the DJI Matrice 300 (Figure 11). This remote aircraft integration required testing and refinement of the vibration mounts, as the initial installation resulted in excessive vibration of the camera lens and very high levels of motion blur in the imagery. Test flights of the Blackfly-S mounted on the Matrice 300 were conducted on the OSU campus September 24–25, 2022.

The final test conducted in this project phase was to investigate whether similar results to those of the PI camera could be obtained using an off-the-shelf RGB camera and circular polarizing filter by hovering and yawing the aircraft to obtain imagery at different orientations (0-, 45-, 90-, and 135-degree azimuth) of the filter. This rotating camera concept requires only a standard, low-cost camera and could greatly accelerate the readiness and adoption of this PI.

Figure 11: Blackfly-S polarimetric camera integrated on DJI Matrice 300. Credit: Oregon Statue University.

OSU graduate student, Kyle Herrera, collecting PI imagery at Neptune State Scenic Area. Credit: Oregon State University.

Section 4 Results

4.1 Improvement in Debris Detection Using PI Imagery

Results from the PI camera was compelling. Visual analysis of the imagery indicated that the polarimetric information was useful in manual detection and recognition of debris items. Nine different examples are shown in Figure 12, where each image pair consists of an RGB image containing a particular debris item on the left, and the corresponding bivariate DOLP–AOLP on the right. As can be seen in these examples, there are a number of cases in which a debris item is difficult to distinguish visually in the RGB image, due to being partially buried and/or similar in color and texture to the background, but more readily identified in the corresponding PI image. Additional examples are provided in Herrera, 2022.

DJI Phantom Pro 4 Remote Controller in use during field test. Credit: NOAA NCCOS.

Correlation matrices revealed that PI-derived bands representing the degree and AOLP were sufficiently different from the standard RGB bands to merit inclusion in a classification algorithm. The layering of PI-derived bands with standard RGB bands improved the separability of designated debris categories (Herrera, 2022). When PI-derived bands were included in a K-NN ML algorithm, the overall accuracy of debris classification increased by 6.7–25.6 percentage points, with a mean classification accuracy increase of 15 percentage points (Herrera, 2022). Kappa coefficients (a measure of classification accuracy, as compared with chance agreement) increased by 9.1–17.6 percentage points (Herrera, 2022). There was also marked improvement in both producer and user accuracy for most categories of debris when PI-derived bands were included (Herrera, 2022). Table 1 shows the improvement in debris classification accuracy using PI-derived information for one of the field sites, while Figure 13 shows one example of the output of a supervised classification using, as input, the RGB bands only (top) and both the RGB and PI-derived bands (bottom).

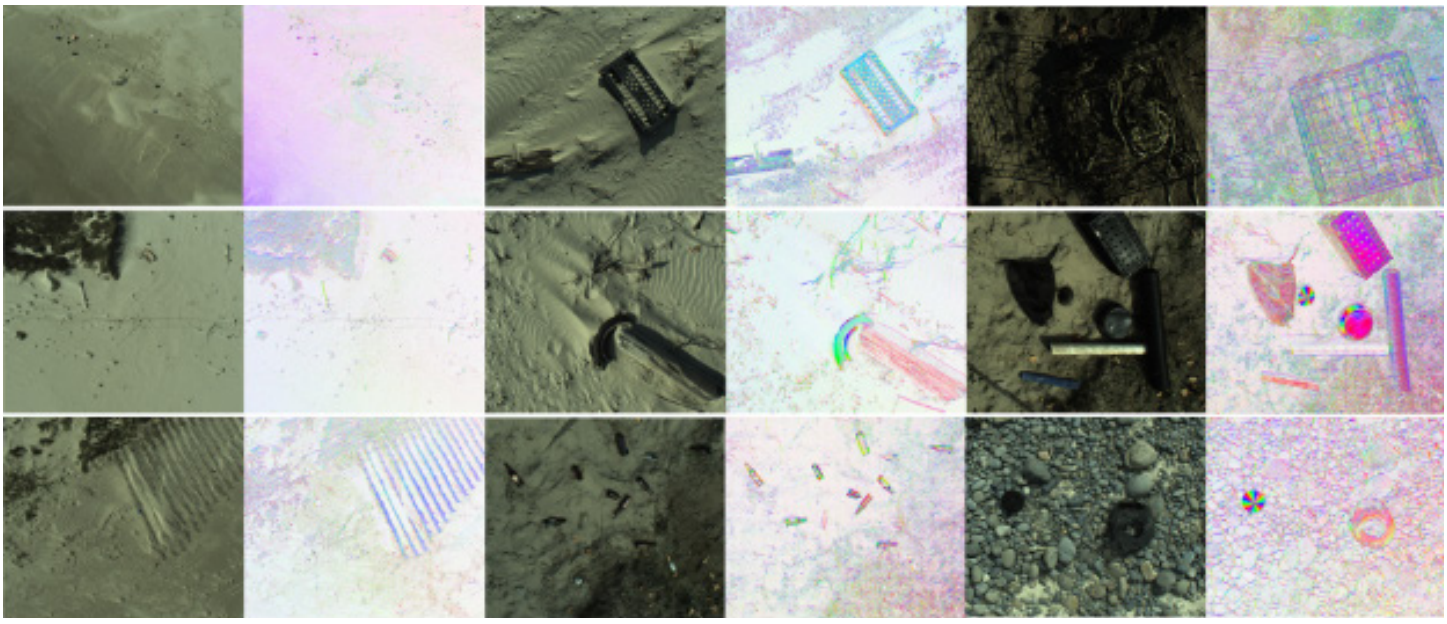

Figure 12: Visual analysis of debris items in PI imagery. In each image pair, the left image is a conventional RGB image, while the right image is a bivariate degree of linear polarization (DOLP)–angle of linear polarization (AOLP) image created from the polarimetric information.

Table 1: Improvement in debris classification enabled using PI-derived image band. It should be noted that this is one of three such tables containing results from each project site. The other two tables are available in the MS thesis of the graduate student supported through this project (Herrera, 2021). Note "TD" refers to transformed divergence, a standard class separability index. The kappa statistic is a well-known metric in remote sensing, used to evaluate classification performance, accounting for the probability of chance agreement between the classification output and reference data.

RGB (3-band)					
Material Type	*Training Accuracy (Prod, user) %*	*Test Accuracy (Prod, user) %*	*Overall/kappa train (%)*	*Overall/ kappa test (%)*	*Mean TD*
Plastic	92.2, 71.4	50.5, 79.8	64.5, 50.4	50.1, 34.2	1213
Wood	68.5, 25.7	53.2, 13.7			
Rope	72.5, 62.7	64.9, 25.7			
Substrate	55.0, 94.3	45.6, 89.8			
RGB + S0 + S1 + S2 + DOLP + AOLP (8-band)					
Material Type	*Training Accuracy (Prod, user) %*	*Test Accuracy (Prod, user) %*	*Overall/kappa train (%)*	*Overall/kappa test (%)*	*Mean TD*
Plastic	91.5, 89.7	57.8, 93.2	89.9, 83.1	75.8, 63.5	1748
Wood	71.4, 81.1	53.4, 48.7			
Rope	81.4, 83.9	77.0, 39.1			
Substrate	95.9, 93.3	90.7, 90.4			

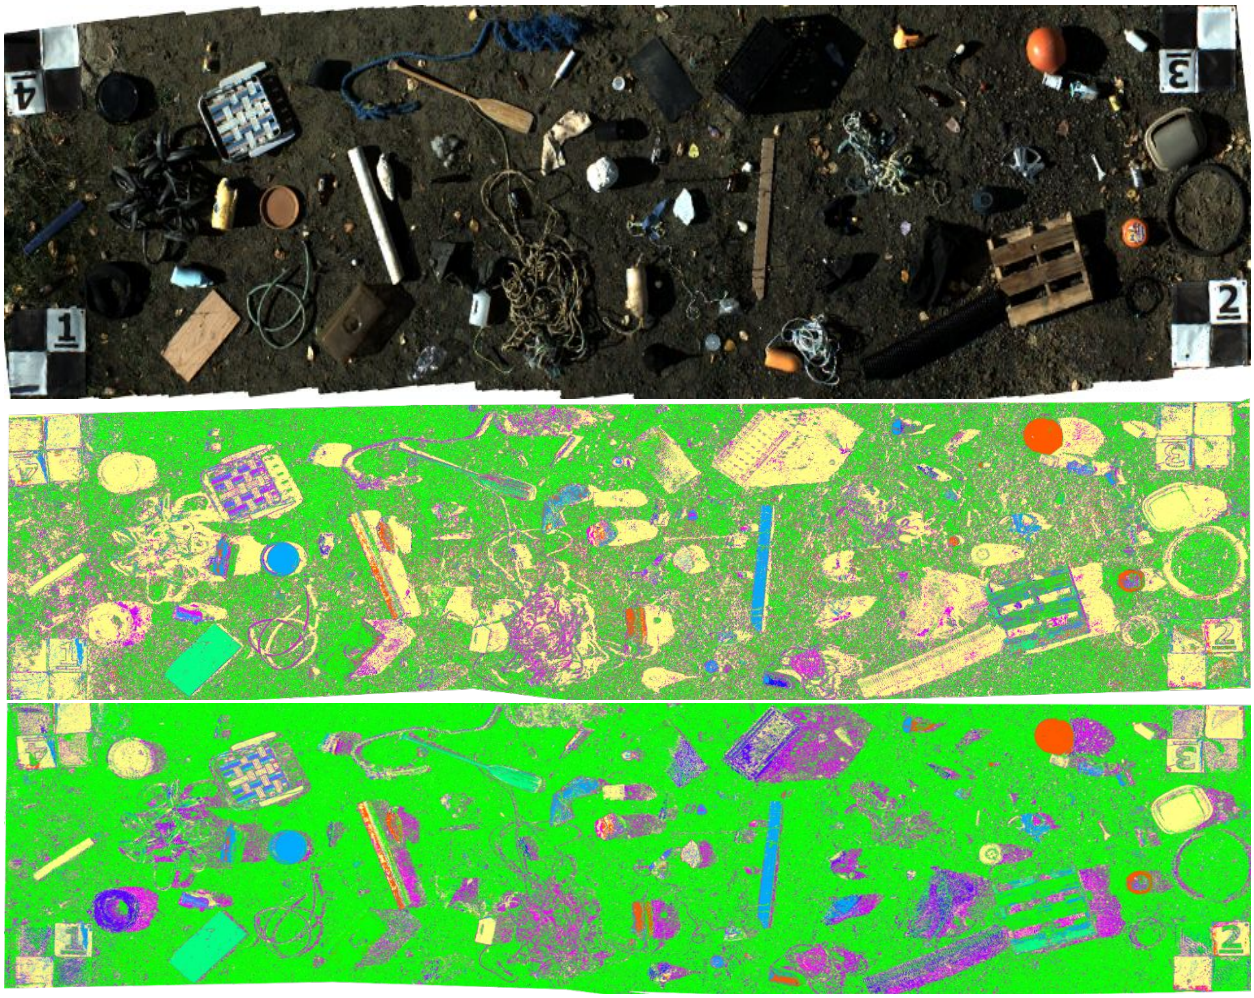

Figure 13: Comparison of classification results with and without PI-derived bands included. Top: RGB orthomosaic; middle: output of K-NN classification without including PI-derived bands; bottom: output with PI-derived bands included. The colors in the bottom two images represent the material type classification, with green representing substrate (sand).

4.2 Machine Learning Results

Results from the ML model applied to true color imagery showed that consumer-grade RGB cameras, such as those often found on consumer UAS, can serve as a potent source of information about the locations of marine debris along a shoreline (e.g., Figure 14). However, the limited spectral information limits the reliable classification of marine debris based on material type. Our first attempt at a deep-learning–based object detection system designed for the automatic discrimination of shoreline stranded macro- and mega-debris in a real-world setting returned 49% average precision (AP) when compared to a human operator on our 10-class evaluation dataset (ORBTL, 2021). The model was significantly better at detecting large and distinct debris, while class confusion was a major issue. Confusion occurred between ambiguous material classes (e.g., plastic and foam) and also naturally occurring objects (false positives such as driftwood). A well-defined classification scheme focused on item use may be more applicable to ML of mega-debris from true color images. For example, the model often achieved an AP of 70% on classes with homogeneous shape, color, and size (such as tire and fishing net; ORBTL, 2021). Small object detection is a primary challenge, and a general rule-of-thumb appears to be that reliable object detection requires a minimum of 10 pixels per object. Therefore, it is important to constrain the minimum target debris size in relation to the input ground spacing distance of the imagery. For example, our study utilized 2-cm imagery; therefore, detection of objects smaller than 20 cm was not attempted (ORBTL, 2021). While the study's bottom-line precision does not yet achieve human-level accuracy at labeling debris in true color imagery, the results do show promise based on the small size of the dataset (approximately 5,700 labels) and the limited geographic scope of that training dataset (sections of Hawaiian coastline). Additional labeled data from a variety of geographic areas, dates, sensors, lighting conditions, etc. would train more reliable deep-learning–based object detection models (ORBTL, 2021).

4.3 PI Camera UAS Integration Results

The Blackfly-S polarimetric camera system installed on the DJI Matrice 300 proved capable of acquiring well-exposed imagery over the field site at altitudes up to 100 m. The imagery was processed into polarimetric AOLP and DOLP parameters. After refining the camera mount (including adjusting the vibration isolators), the imagery had minimal image blur and clearly depicted polarization signatures of the simulated marine debris. The imagery covered the entire field site, although it was challenging to ensure the camera was acquiring imagery of the site due to the lack of a real-time display from the camera.

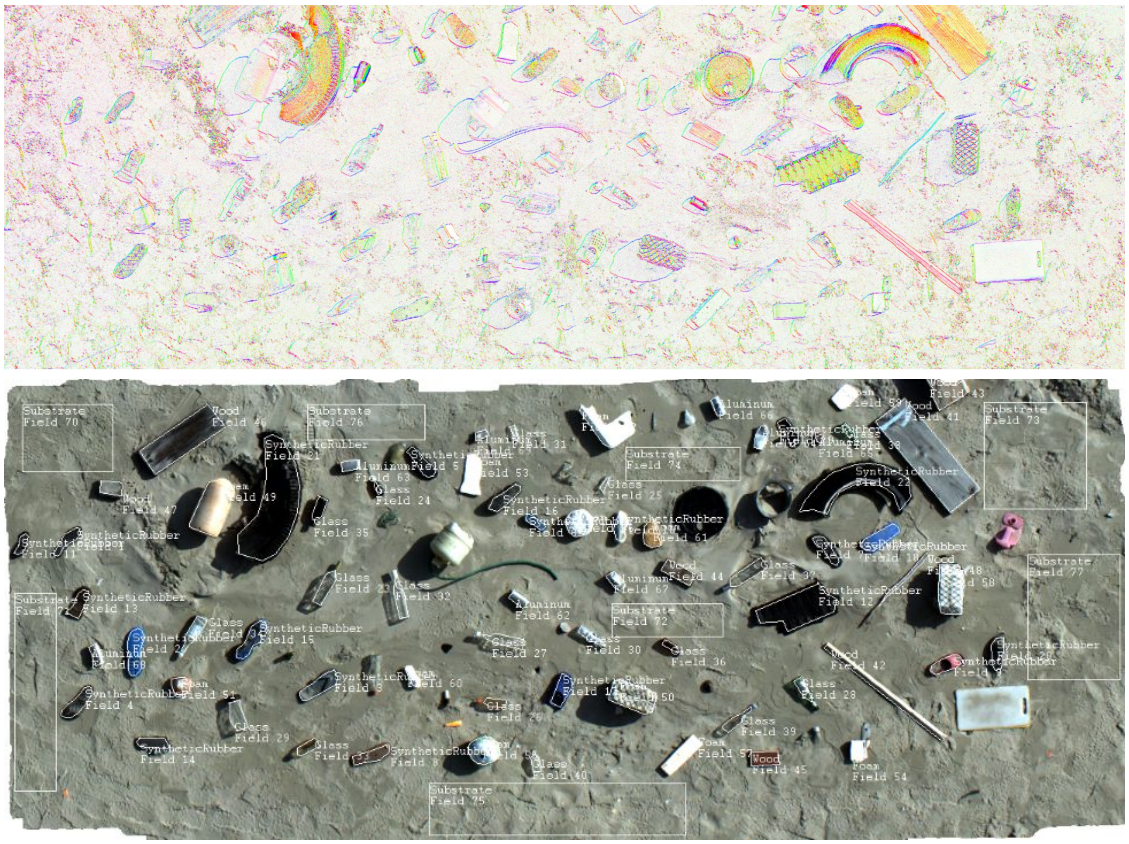

Figure 14: Auto-detection of debris items in seeded debris field. The top image shows a bivariate AOLP-DOLP image from the polarimetric camera for the high-density debris field on San José Island. The bottom portion of the figure shows the results of running DebrisScan on the RGB orthomosaic generated from the UAS imagery of the same site.

Results

The prototype auto-exposure algorithm implemented by the team was unsuccessful during a few portions of the flight, resulting in both overexposed and underexposed imagery. Future work should implement a more robust algorithm that is tuned for the specific characteristics of the camera. This is very important for this particular polarimetric camera, as it appears to qualitatively have a relatively low dynamic range. It is also important to note that commercially available PI cameras, such as that tested in this project, have poorer resolution than conventional RGB cameras. In practice, this requires the operator to fly at a lower altitude to achieve a desired GSD, which then impacts the number of flight lines and time required to complete a survey.

The imagery from the rotating RGB camera with the polarizing filter yielded very high-resolution (45-MP) images that were automatically exposed and inspected for quality using the real-time data link. The processed polarimetric imagery from these data depicted the same polarimetric signatures as the polarimetric camera but with a higher resolution (Figure 15).

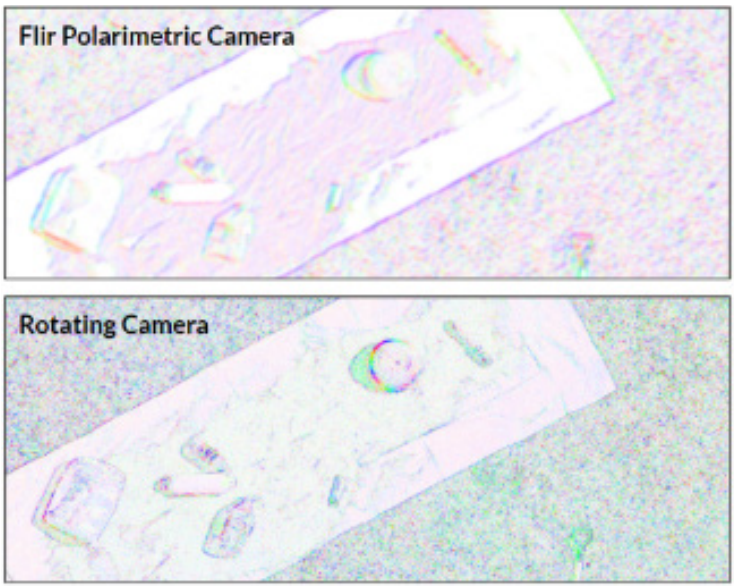

Figure 15. PI images (specifically, bivariate AOLP/DOLP images) generated from the Blackfly-S polarimetric camera (top) and rotated RGB camera with a polarizing filter (bottom).

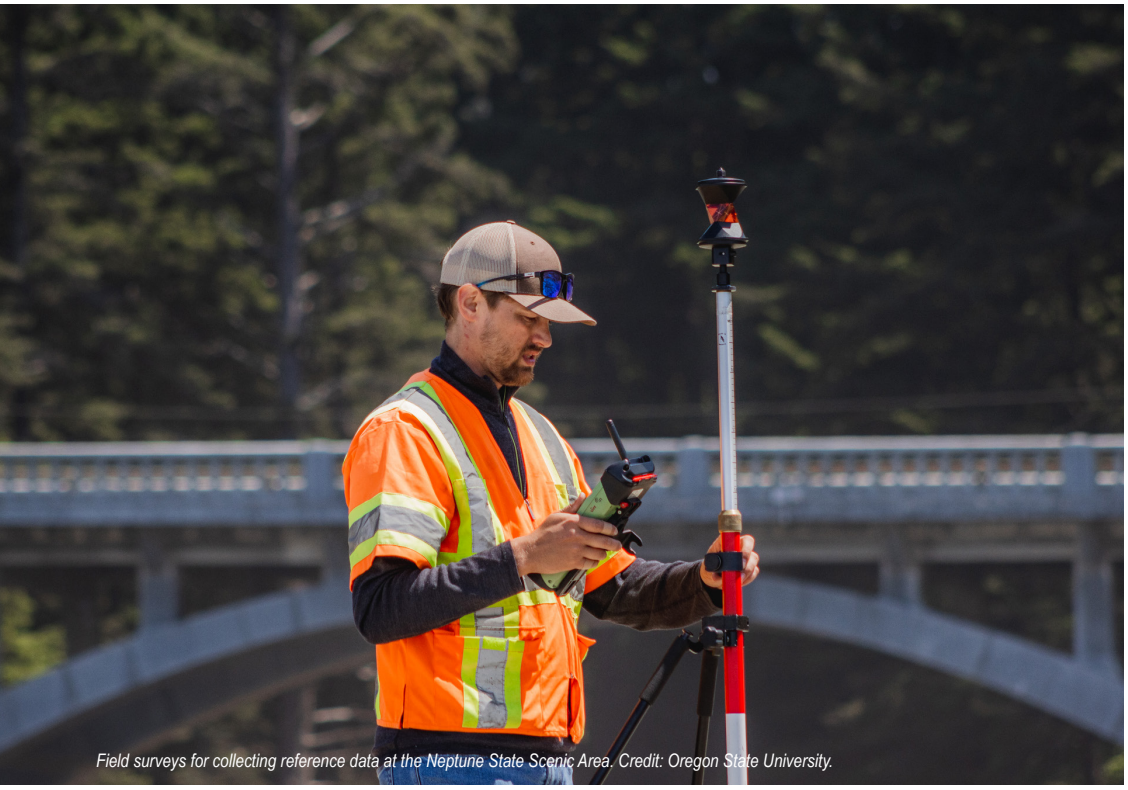

Field surveys for collecting reference data at the Neptune State Scenic Area. Credit: Oregon State University.

Tracy Weatherall of the University of Texas Marine Science Institute collecting debris on Padre Island. Credit: NOAA NCCOS.

Section 5 Recommended Operational Workflows

The following sections of this report contain operational recommendations, based on this multi-phased study, as well as experience of the project team in related studies (e.g., Slocum et al., 2019; Winans, 2021). These sections are intended to benefit researchers and programs interested in enhancing their marine debris detection abilities or programs. However, they are not intended to guide users through steps of learning to fly UAS, obtaining Federal Aviation Administration (FAA) Part 107 certification, or other regulatory aspects. There is a litany of resources available to guide interested, novice parties on initial entry into the use of UAS, from the wealth of guidance on obtaining Part 107 certification provided directly by the FAA, to the countless videos, discussion boards, blogs, and apps that cover the topics of both UAS operation and certification. (Perhaps the best vector of entry to the use of UAS is hands-on practice, beginning with recreational flights in open, safe environments.) The intended audience for this report, however, is users with moderate to extensive experience with UAS including operational basics, FAA certification, knowledge of flight permission procedures, and having sufficient experience with remote aircraft to safely collect data in a coastal environment.

5.1 Pre-Planning and Program Setup

An initial consideration is the selection of a specific remote aircraft (make and model) for conducting marine debris mapping flights. These procedures assume the use of sUAS (< 55 lb) under FAA Part 107. As noted in Section 3.2, two main categories of remote aircraft are multirotor (e.g., quadcopters, hexicopters, etc.) and fixed wing. Based on the findings of this project, it is recommended that multirotor aircraft be considered the primary choice. The advantages of fixed-wing aircraft include longer flight endurance (i.e., longer flights per battery) and the ability to efficiently cover large stretches of beach. However, it was found that the multirotor aircraft (at least those tested in this study) tended to perform better in higher winds (e.g., wind speeds on the order of 15–20 kts). In the case of the WingtraOne, the aircraft's VTOL configuration proved challenging in higher winds, as the large wing acts as a sail and can push the aircraft downwind. The other advantage of multirotor aircraft is that the small, lightweight quadcopters tested in this project are highly portable and user-friendly. This makes them ideal to carry and operate, including for personnel whose primary job function is not UAS operations but who want to use UAS periodically to enhance data collection. Additionally, multirotor aircraft are generally easier to maneuver, less expensive, and require smaller takeoff and landing areas, which is critical for operations from a small vessel.

Other factors to consider in setting up an operational UAS capability include: required IT infrastructure (e.g., disk space to process and archive large volumes of imagery), training and flight proficiency requirements and safety protocols. These aspects of establishing an operational UAS program are beyond the scope of this report. UxS standard operating procedures are available from a number of sources, including the North Carolina Department of Transportation (NCDOT, n.d.) and Connecticut Department of Transportation (Zaffetti, 2019).

5.2 Project Planning

5.2.1 Takeoff/Landing Zone

The multirotor (e.g., quadcopter, hexicopter, etc.) remote aircraft suggested for this procedure have versatile takeoff and landing options. It is recommended to have a stable, open launch zone of roughly 2 x 2 m away from impeding trees or structures. This will allow for reliable home-point recording and safe takeoffs/landings when performing automated survey flights. A flat, rigid pad (Figure 16) is recommended to prevent the UAS from tipping or being damaged from the natural ground surface beneath the pad. Having paint or markings to discern the center-of-pad helps with reliable point recording from the UAS. If a stable launch zone is unfeasible (e.g., moving/rocking vessel, rough/uneven terrain, or unsafe landing area), many UAS platforms allow for takeoff/landing by hand, where the pilot will take manual control of the device and carefully maneuver it into the positioned hand of an assisting team member (as shown in Figure 17). If conditions worsen during flight, or the UAS is unfit to land automatically, manual pilot control is recommended to safely return to the platform. Most modern UAS platforms will have an alert feature via the controller when challenges are encountered (e.g., loss of signal, poor home-point recording, or excessive wind gusts).

Figure 16: UAS takeoff and landing panel. Credit: Oregon State University.

Figure 17: Hand launch of UAS. The aircraft shown here is a Skydio 2. Credit: NOAA NCCOS.

Launch and recovery of the UAS from a vessel is also an option in some situations. However, there are a number of safety and operational considerations to be aware of with vessel launch and recovery. Because of vessel movement (roll, pitch, and heave), it is critical to have sufficient deck space for a takeoff and landing zone, larger than the amount of space that would be needed on land. On a vessel, hand launch and recovery of the UAS is often preferable to taking off and landing directly on the deck. However, the UAS manual and operating procedures should be consulted to determine whether hand launch and recovery is supported by the particular remote aircraft and to understand the manufacturer's recommended procedures. Practicing the procedures on land ahead of time is critical, as every aspect of the process is substantially more challenging on a moving vessel. A key concern is ensuring the safety of the person hand-catching a landing UAS and, in particular, minimizing the risk of being cut by the spinning propellers. Safety of the remote aircraft is also an important consideration. There are a number of situations in which the remote aircraft could inadvertently land in the water and be destroyed and/or lost. Just one example is if the remote aircraft loses communication with the controller and returns to its home point to land, but, due to movement of the vessel, the recorded home point is now over water.

5.2.2 Image Resolution

The ML model is designed to perform inference on imagery with a 2-cm GSD. Any imagery not at 2-cm resolution will be automatically up- or downsampled upon upload to the API. If auto-resampling is required, this will increase the processing time and could decrease performance. Importantly, the upsampling performed on coarser-resolution imagery to obtain a 2-cm GSD does not improve the effective resolution or quality of the imagery. For these reasons, it is recommended to avoid having the software upsample or downsample the imagery by acquiring imagery with a 2-cm GSD. The following equation can be used to precompute the GSD:

$$GSD = x_{px}\left(\frac{H'}{f}\right) \quad \text{(Eq. 1)}$$

In Equation 1, x_{px} is the physical pixel size on the camera chip, H' is flying height above ground level (AGL), and f is the camera lens's focal length. If x_{px} and f are given in units of millimeters, and H' in units of meters, the computed GSD will likewise have units of meters. Since the physical size of a pixel is fixed for a particular camera, and the focal length may be fixed for a particular lens, the AGL is most often the parameter that is varied to change the GSD.

UAS mission planning and flight management software for mapping missions (e.g., DroneDeploy, Pix4DCapture, Pixhawk Mission Planner, and senseFly eMotion) will typically automatically calculate GSD and other imagery parameters, given a project area, UAS camera model, and flying height. Furthermore, it is generally possible to reverse the process by inputting a desired GSD and allowing the mission planning software to calculate flying height and other parameters.

5.2.3 Beyond Visual Line of Sight

Operating a small uncrewed aircraft system (sUAS) under FAA Part 107 requires visual line of sight (VLOS) of the platform to ensure flight safety and situational control. This does impact coastal mapping missions where long stretches of shoreline are operationally flown. It is important to consider this range when selecting a suitable launch zone to efficiently acquire data for large areas. The VLOS requirement typically limits small quadcopters to approximately 500 m (either horizontal direction) from the ground station, as this is the maximum distance at which the remote aircraft can typically be seen, even in very clear conditions. It is recommended to select a ground station roughly in the center of the AOI to allow for maximum shoreline coverage per flight. The process of obtaining an FAA Part 107 § 107.31–Visual Line of Sight Aircraft Operation waiver involves completing and submitting a waiver application through FAADroneZone, including a description of the proposed operations and the justification for the waiver (FAA, n.d.).

5.2.4 Data Management

Appropriate data management is critical for project planning, data acquisition, and result generation. When imaging large stretches of coastline, a single survey can produce hundreds of high-resolution images, equating to gigabytes of data. One project typically includes several survey missions and gigabytes quickly turn into terabytes of imagery. As such, it is imperative to maintain consistent file organization and sufficient memory allocation. It is recommended to have several microSD cards empty and ready before beginning any project data acquisition. Most UAS will specify on the controller the remaining image capacity of the microSD card installed in the device. Subsequently, most survey planning tools will specify the number of images required to complete the mission. The project team should always remain aware of these two numbers to avoid any missing data or failed project objectives: the total images required for the mission, and image space remaining on the microSD card. It is recommended to have more microSD cards than needed in case any fail to record imagery correctly (rare but possible).

For ease of result and report generation, it is also suggested to have a reliable, consistent file structure, including an individual folder for each flight survey named according to the collection date and location. A project-dedicated portable hard drive is recommended to have project files backed up and organized in one place.

5.3 Data Collection

5.3.1 Flight Parameters

Data collection for UAS marine debris shoreline surveys may entail a combination of “mowing the lawn” (i.e., parallel flight strips with a specified amount of image overlap to provide wall-to-wall stereo coverage of the area of interest) and “hotspot” capture to acquire detailed imagery of specific debris items from varying viewing geometries and flight altitudes. Often, the mowing-the-lawn survey is conducted first, and areas of high debris concentrations are subsequently investigated using the hotspot capture mode. The following sections describe the considerations and parameters for the mowing-the-lawn portion of the mission.

Flying Height/Image Resolution

Operational flight parameters are typically dictated by the planned project scope and the target area of interest. UAS flying height AGL is determined by the desired imagery resolution and platform sensor used. For example, to achieve 2-cm GSD, the Phantom 4 Pro would collect imagery at an altitude of 68 m AGL. Alternatively, the Skydio 2 would fly at 47 m AGL to achieve 2-cm GSD. The relationship between flying height and GSD is illustrated by Equation 1. Under FAA Part 107, the maximum allowable flying height is 400 ft (120 m) AGL.

Remote Aircraft Speed

UAS flying speed is a function of the flying height and image overlap necessary for the planned mission. If the desired result is a georeferenced orthomosaic of the shoreline, it is recommended to achieve an overlap (also called endlap) and sidelap of 75%. This will ensure sufficient image overlap between image frames to accurately reconstruct the camera station and scene geometry from matched features. These settings are typically located/adjusted in survey planning settings. Once the flying height and overlap/sidelap have been specified, mission planning software will typically select the proper flying speed to obtain the best results. Like flying height/image resolution, these settings are platform specific and will vary. For instance, to acquire imagery at 2-cm resolution with 75% overlap and sidelap, the Phantom 4 Pro will fly at 7 m/s at 68 m AGL for best results. Meanwhile, the Skydio 2 would fly the mission at 15 m/s at 47 m AGL to accomplish the same objective.

Example of ground control points (GCPs) at Padre Island field site. Credit: NOAA NCCOS.

Survey Area Geometry

The area of interest for coastal mapping missions plays a large role in how these surveys are planned. Shorelines and their geomorphology vary significantly across the world and, as such, require special attention to detail. The survey area geometry strongly influences how UAS flights are designed in order to achieve the highest quality imagery. Narrow shorelines with extreme elevation change require a different survey approach than wide beaches with little to no slope. Some of the flight parameters that might vary depending on geometry are flight path orientation, gimbal angle, and flying height, while some flying conditions, such as sun angle and wind direction, will also need careful consideration. The overarching goal of imagery acquisition is to achieve the highest quality surface detail for the shoreline area of interest. Analyzing the existing geometry will aid in the determination of how to best plan a survey/mapping mission for a particular location.

Exposure–ISO, Aperture, and Shutter Speed

Exposure is of paramount importance in determining the quality of resulting imagery. Inadequate exposure levels almost always result in a loss of information at the surface level. Auto-exposure is a somewhat reliable function to gauge the relative illumination in the scene during acquisition but should not be solely relied upon. The best way to actively avoid substantial exposure problems is to avoid performing a mission during rapidly varying lighting conditions. For instance, acquiring imagery in direct sunlight and then subsequently in fully overcast conditions will severely affect image quality. Most UAS platforms allow for selection of lighting conditions prior to flying (i.e., sunny, overcast, low illumination, etc.); however, these functions serve as a primary means of balancing the incoming light levels throughout the mission. This setting is not meant to be changed frequently during a survey as the resulting imagery would vary in brightness. Therefore, it is recommended to aim for completing the planned mapping mission under consistent illumination conditions to achieve consistent image quality. This is not always possible in dynamic coastal areas, but the general concept should be adhered to as closely as possible.

Flight Conditions

UAS platforms are highly versatile in that they can successfully function under vastly different conditions. This flexibility is one aspect that makes them valuable in a shoreline scenario where conditions fluctuate regularly. There exist, however, certain conditions that allow for optimal quality of results while minimizing potential risks and hazards. The most important natural factors to be aware of when operating UAS platforms in a coastal environment are wind, rain, clouds, and sun.

Sun/Specular Reflection

Starting with the sun, there are important aspects to understand in how it impacts quality of results. Shoreline mapping missions often involve flying over or near the water surface. Water surfaces are highly reflective of natural sunlight and should be minimized as much as possible to avoid intense specular reflection in the imagery. This can be accomplished by avoiding flight over the water surface whenever possible. However, when unavoidable, specular reflection can be minimized by flying during ideal sun angles/times of day, when the sun is in the range of 20–60 degrees above the horizon. As the sun approaches its highest point in the sky, specular reflections off the water surface will be directed up toward the remote aircraft, causing hotspots and saturation in the resulting imagery. Notwithstanding this sun angle recommendation when avoiding flights over water and sun glint, it is generally recommended to collect aerial

imagery with the highest illumination possible to maximize the shoreline detail captured. Direct and consistent sunlight is desirable for result quality, but the key is consistency in illumination during each flight, as noted previously.

Weather

The most severe operational conditions are wind and rain, as these can often inhibit data collection and create an unsafe working environment. Of all the weather conditions to be aware, rain should be on the forefront of any mission planning for UAS data collection. Most UAS platforms are unfit to operate effectively during active rainfall, and this scenario should be avoided altogether. Sun or consistent overcast conditions are optimal for both image quality and equipment/personnel safety. While wind conditions can also be a limiting factor, multirotor UAS platforms are more tolerable in the amount of wind they can sustain while still accomplishing mission objectives. The maximum allowable wind gusts a UAS can withstand is generally listed in the specifications from the manufacturer. This manufacturer-provided wind tolerance is an essential parameter to keep in mind, but it should not be solely relied on for decision-making in the field. Wind begins to hamper imagery quality in shoreline mapping whenever the UAS struggles to maintain flight path consistency or camera position. In the field, this may look like the UAS swaying off center from the predetermined path or the camera gimbal being temporarily blown off the angle desired.

These effects can significantly impact the quality of results, as several images in the survey can be unusable, and the UAS may struggle to maintain sufficient overlap or sidelap for orthomosaic processing. As such, it is strongly recommended to always monitor wind conditions, both sustained speed and maximum gusts, throughout the duration of the mission. Several resources exist to provide active wind gauge readings and forecasts in coastal areas. Wind, however, can be unpredictable in shoreline environments, and it is recommended to carry an anemometer to the intended fieldwork location for careful monitoring. Of the two wind measurements, maximum gusts are the important limiting factor. If sustained winds are around 10 kts, but gusts are reaching up to 20 kts, these conditions are likely to curtail UAS operations. The two quadcopters that were utilized extensively in this study were able to effectively collect imagery in winds up to 15 kts but did encounter some instances of gimbal instability toward that upper limit. Ideal operational wind conditions were observed to be 12 kts or less to ensure high-quality imagery throughout the duration of the flight. These tolerances vary slightly between UAS platforms, so it is important to be aware of individual equipment capabilities. It is highly recommended to perform test flights in a variety of wind conditions to discern adequate confidence levels in deployment and data collection. Extreme wind gusts can be common in coastal areas and can potentially blow a UAS off course entirely, resulting in equipment damage and ceasing of mission operations. The overarching recommendation is to exercise extreme caution when planning around wind conditions and to curtail operations during instances of unpredictable gusting events.

Clear day at Neptune State Scenic Area, Oregon. Credit: Oregon State University

WiFi and Cellular Service

WiFi and cellular service availability are important utilities to consider during all phases of project planning and data acquisition. Many UAS programs for pre-planned survey missions rely on accessing the internet for both the planning and launch phases. In most instances, it is likely that the internet is available for mission planning and setting flight parameters. However, many shoreline scenarios where this procedure would be employed may have limited to no connectivity. (In this project, this was the case at the Neptune project site.) This could result in the inability to effectively launch the UAS to complete the pre-planned survey as the project files would be inaccessible. Therefore, it is recommended that the project team verify cell service availability for the area of interest. The best practice is to first open the UAS mapping software and load the intended project file prior to departing for the field site. This ensures that the pre-planned survey is effectively loaded and ready to deploy in the target shoreline location. Remote aircraft-to-controller connectivity does not typically rely on WiFi or cell service; however, the controller/phone/tablet interface requires internet data to access project files and pre-planned missions.

5.4 DebrisScan

The ML model developed and tested in this project was integrated into operationally ready software called DebrisScan. The DebrisScan application (available on GitHub: https://github.com/orbtl-ai/DebrisScan) allows end-users to utilize our pre-trained ML models for debris detection on their own imagery. DebrisScan (Figure 18) allows users to upload UAS photos and receive a zip file containing labeled images, spreadsheets, and reports about the debris in each photo. DebrisScan is a free and open-source project licensed under the permissive Apache-2.0 software license (Vendome and Poshyvanyk, 2016). This license structure allows DebrisScan to reach the widest possible audience for the largest possible impact.

This beta release contains an EfficientDet-d0 object detection model (Tan et al., 2020) from the TensorFlow Model Garden that has been pre-trained on a real-world set of 2-cm aerial imagery collected across the Hawaiian Islands and labeled by teams of experts. This model does not represent the current state of the art, but it has proven versatile over many types of terrain and imagery. It is recommended that DebrisScan users briefly review this project's EfficientDet-d0 model report to better understand DebrisScan v0.05's current capabilities (ORBTL, 2021). The EfficientDet-d0 neural architecture is also small enough to be run on a wide range of modern hardware, including personal laptop computers that lack dedicated graphics processing units (GPUs). The DebrisScan application is designed to detect only large pieces of macro- or mega-debris that are greater than 20 cm in size (this is constrained by the model's 2-cm GSD). Further, this model is designed to work on only "dry" debris that is above the waterline. While the EfficientDet-d0 model classifies debris into eight discrete categories, the performance can vary greatly across object classes. It is best to consult the linked model report above for detailed information about the data and techniques used to train and evaluate DebrisScan's primary model (ORBTL, 2021).

5.4.1 Downloading and Installing DebrisScan

Download/installation instructions can be found on the DebrisScan code repository's front page (https://github.com/orbtl-ai/DebrisScan). The entire DebrisScan application is distributed and installed using free and open-source tooling such as Git and Docker. DebrisScan supports Windows 10/11 and Linux platforms as of this writing, with plans to expand support in the future.

Running DebrisScan is relatively simple, and the entire app can be "spun up" in a single command (see instructions above). Once DebrisScan is running, it is accessed via a web browser by going to the default address of `localhost:8080` (unless changed by the user). DebrisScan v0.05's homepage is shown in Figure 19.

5.4.2 DebrisScan Preprocessing

DebrisScan allows users to bulk upload large sets of aerial imagery to be scanned for marine debris by a geospatially aware AI. In this context, "geospatially aware" means that DebrisScan can read additional georeferencing or GNSS information from imagery uploads to determine the geographic locations of debris objects, in addition to spatial characteristics such as debris size. DebrisScan's output results will be formatted with as much of this "geo aware" information as possible. If no location information is present, then DebrisScan will simply count and classify the debris without attempting to catalog debris objects in space or time.

Figure 18: DebrisScan software for running the ML model developed in this project. ORBTL AI with permission from ©Government of Japan / PICES 2015.

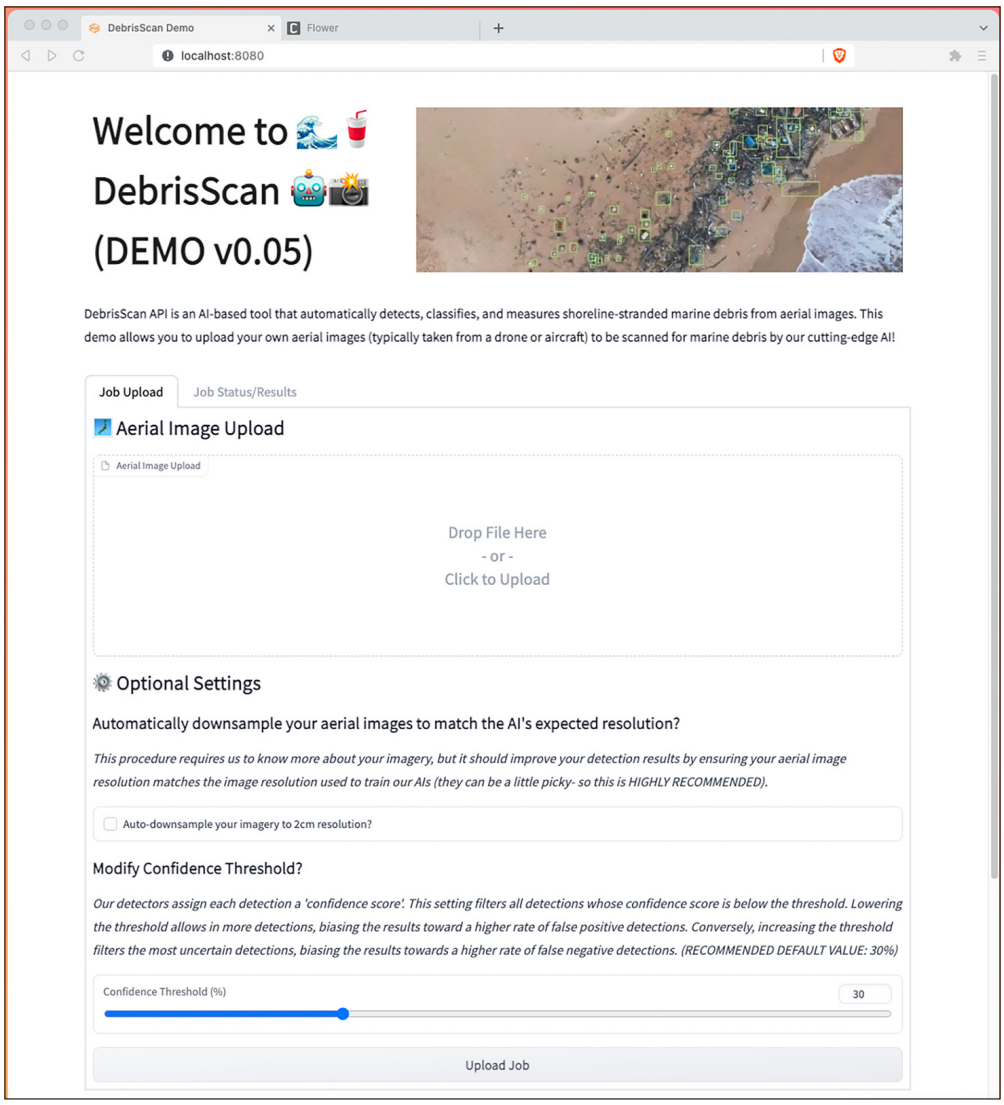

Figure 19: Image of DebrisScan homepage.

It is highly recommended that users of DebrisScan either orthorectify or georeference their aerial imagery prior to uploading. This allows DebrisScan to provide the best, most detailed results. However, not all users may need high-quality debris locations for their specific use cases. Additionally, users may not be familiar with orthorectification/georeferencing of aerial imagery. For these reasons, DebrisScan will still operate on standard single-frame images.

The following is an overview of how DebrisScan interprets, processes, and returns results for the three possible types of image uploads:

Orthorectified and Georeferenced Imagery

This type of imagery has distortions (e.g., due to topographic relief, lens distortion, etc.) removed and contains pixel-level spatial coordinates in a defined reference frame, which allows DebrisScan to relate the image pixels that comprise a debris object to a real-world location. This location information allows the results of these surveys to be communicated to assessment crews and responders to aid in the identification and removal of debris. While the steps to produce georeferenced orthoimages and orthomosaics are beyond the scope of this report, the procedures are becoming significantly more automated with new structure from motion (SfM) photogrammetry software, such as Pix4D Mapper, Agisoft Metashape, and online software, such as WebODM (Open Drone Map). (Note: these are not endorsements of any particular software but simply some of the software packages that are currently in use, including by the project team.) Despite the increasing levels of automation, generation of orthomosaics can sometimes take hours of processing time for larger projects. Additionally, the quality and positional accuracy of the outputs can vary significantly, depending on the number and

spatial distribution of GCPs, the type of aircraft GNSS (e.g., standalone pseudoranging vs. carrier-phase–based RTK or PPK), the quality of the imagery, and the specific processing procedures. The U.S. Geological Survey Open-File Report 2021-1039 (Over et al., 2021) contains useful information on recommended processing workflows and settings.

Non-georeferenced Imagery

Users of DebrisScan may be more familiar with non-georeferenced imagery. This type of imagery, which is the typical output of a UAS mapping mission, does not contain spatial coordinates at a pixel level but may contain spatial coordinates for each camera station (i.e., each image center) obtained via a GNSS receiver on the remote aircraft. DebrisScan can still count and classify debris within these standard images, but extracting accurate real-world spatial coordinates and object measurements is impossible. Fortunately, many modern digital cameras contain consumer-grade GNSS receivers and can roughly log the location an image was taken in the industry-standard exchangeable image file format (EXIF). If DebrisScan finds this image-level spatial information stored in an image's EXIF data, the coordinates will be associated with found debris objects, providing a rough location. If information about the sensor's physical characteristics and flight altitude is provided by the user (optional) or found among the EXIF data, it will be used to estimate each debris object's shape and size. The more information provided by the user (or contained in an image's EXIF data), the higher detail the results will be.

Standard Imagery

"Standard images" are any images that contain neither pixel-level nor image-level location data. Without these data, DebrisScan will operate as a standard AI model. The AI can count and classify debris objects, returning simple reports of each debris object class's rate of occurrence and image plots showing predictions. No spatial coordinates or size information will be provided.

5.4.3 Processing in DebrisScan

Starting from the DebrisScan processing interface (Figure 20), users can upload batches of aerial imagery in the primary field labeled "Aerial Imagery" and press "Submit" to run DebrisScan. At this stage, users can also provide optional information about the sensor and flight altitude (optional). This additional information allows DebrisScan to infer the GSD of each photo, resample the photos to match the expected 2-cm GSD, and potentially geolocate debris objects. These parameters are highly recommended to help DebrisScan achieve optimal results.

Finally, the user has the ability to set the "Confidence Threshold" at which DebrisScan filters the model's predictions. A confidence of 30% to 40% is ideal for most scenarios. Lowering the Confidence Threshold will bias the model toward over-counting debris. Conversely, raising the Confidence Threshold will bias the model toward under-counting debris.

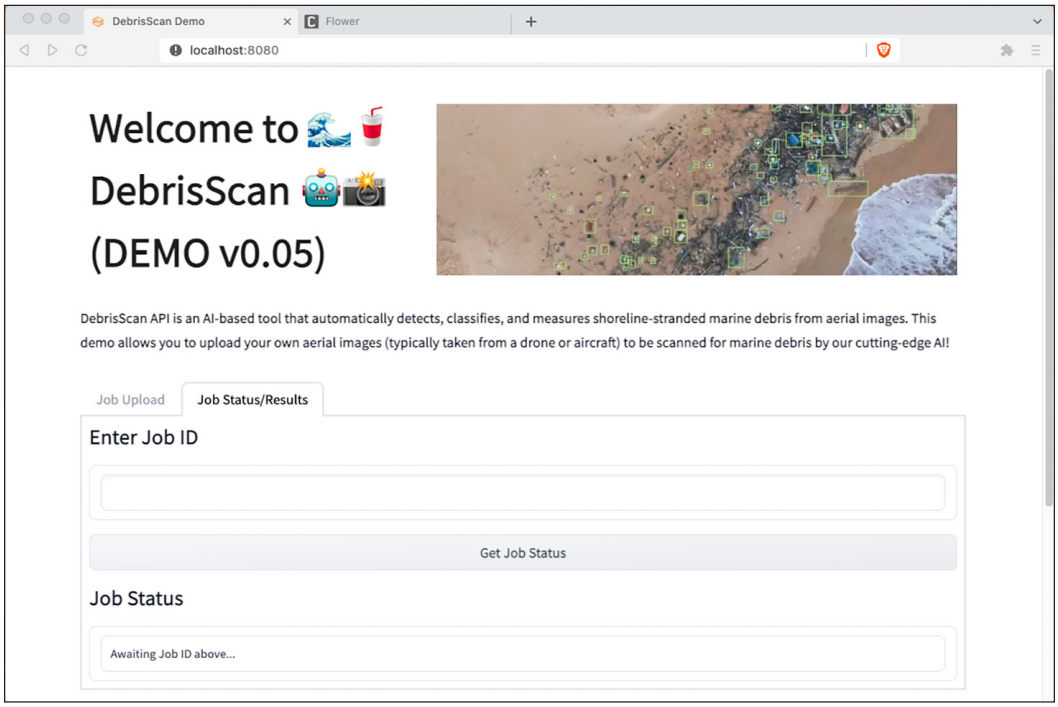

Figure 20: DebrisScan's Status and Results retrieval tab. A user-provided Job ID will return a message about the processing status.

Once the user presses "Submit," DebrisScan begins to evaluate the upload to ensure that only valid imagery is kept. If everything looks good, DebrisScan returns a unique "Job ID" to the user. The Job ID allows the user to leave the DebrisScan application, work on other tasks, and then return to retrieve results once the processing has finished. This Job ID also allows users to check on the status of their job. Retrieving the status and results from a DebrisScan processing job is done at a single "Job Status/Results" endpoint.

DebrisScan will return a zip file called "inference_results.zip," which will contain the following files:

- **Image plots** – These consist of the user's original images with debris bounding boxes drawn over all predictions that exceed the user's specified Confidence Threshold. The bounding boxes will be colored and labeled with both the object's predicted class, and the model's prediction confidence (0–100%). The word "_plot" will be appended to each of the original image names. These plots are useful for a quick qualitative review of DebrisScan's performance. If image uploads contain no spatial information, these are the primary visual outputs from DebrisScan. An example is shown in Figure 21.
- **Debris list** - A file called "all_debris_objects.csv" contains a database of all DebrisScan's predictions. This is DebrisScan's primary quantitative output, allowing users to quickly view many characteristics about their scene's debris objects. This file will contain DebrisScan's most complete report on debris, listing each prediction's location, classification, estimated size, confidence level, and associated input image. Note that these debris attributes may not all be included (specifically location and estimated size) if spatial information and/or sensor and flight information is not provided.
- **Class list** - A file called "debris_type_counts.csv" contains a list of DebrisScan's debris classes (i.e., fishing net, vessel, and buoy) in each scene, the number of debris in each category, and the total number of debris objects (across all input images). DebrisScan provides this summary file to get a high-level view of the total number and type of debris in each scene.
- **Per-image results** – DebrisScan will also create a folder labeled "per_image_results," which contains redundant copies of the above information, broken into one .csv file per image (with files named accordingly). These files will match the format of the "all_debris_objects.csv" file; however, each will contain debris

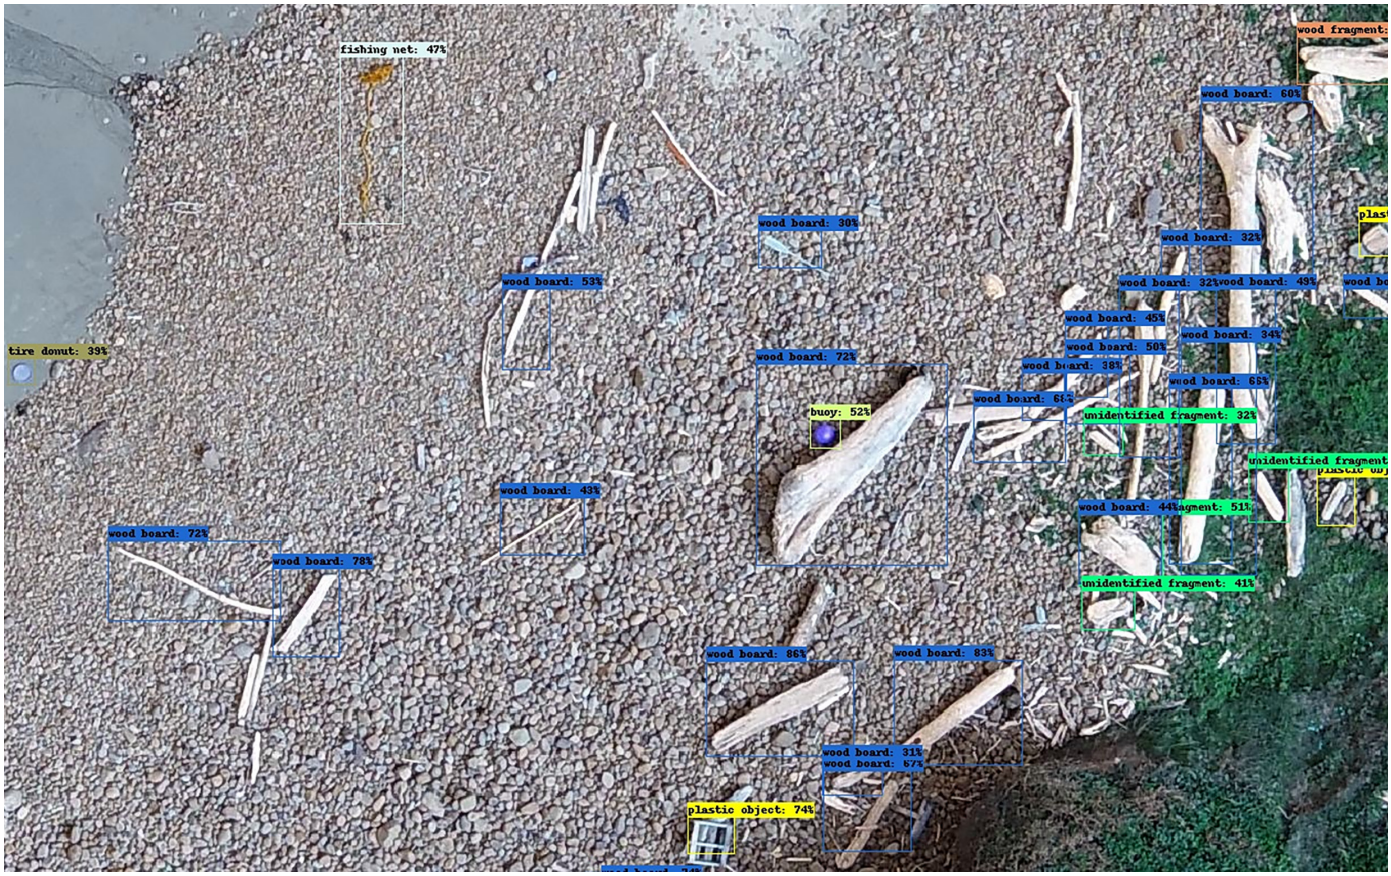

Figure 21: DebrisScan results.

objects for only a single image (as opposed to all images). Additionally, the per-image results are provided as .json files, which can be difficult for users to read but are often programmers' preferred way of working with DebrisScan's results.

5.4.4 Administrator Dashboard

DebrisScan is designed to handle multi-user production workloads. The entire DebrisScan application is built around a Celery task queue and a Flower administrative dashboard. Celery is a widely used open-source task queue, which allows jobs to "stack" and be processed in the order they were received by a pool of "workers." A Celery Worker can be either a portion of a single computer or an entire group of computers working together to handle deep job stacks or large, time-consuming user uploads. Administrators of DebrisScan will find the included Flower Dashboard helpful, which runs alongside DebrisScan by default and can be reached in the web browser by default at `localhost:5555`. See Figure 22, which shows the Flower Administrator Dashboard below with detailed information about DebrisScan and all received user uploads.

5.5 Post-Processing

To generate debris density maps, the "Point Density" tool in ArcGIS Pro can be used. It should be noted that this tool requires a Spatial Analyst toolbox license. The input to this tool consists of the locations (e.g., Universal Transverse Mercator [UTM] eastings and northings) of detected debris points from the ML model. These points can be stored as a .csv file, which is an ASCII text file containing columns of data, such as: debris_ID, utm_East, and utm_North. Please note that DebrisScan provides this information by default (provided the necessary spatial information is present in user uploads). After opening the table in ArcGIS Pro, the points can be displayed using the "Display XY data" function and then converted to a geodatabase feature class, after which they can be used as input in the Point Density tool. The tool computes the spatial density of points using a window of defined size around each cell.

To be able to consistently compare debris density across different areas, it is important to use consistent settings in running the Point Density tool. The output cell size parameter specifies the GSD of the output layer. The parameters in the "Neighborhood" portion of the window specify the shape and size of the window around each cell that is used in searching for points that are then used in the density calculation. The "Area units" parameter specifies the units of the denominator in the density calculation; for example, selecting hectares means the density will be displayed as the number of debris items per hectare. Similarly, using consistent symbology (color scheme, stretch type, and number of classes) will enable the debris density maps to be compared across different regions in a meaningful way. Once a suitable set of symbology parameters is determined, the "Apply Symbology from Layer" tool within the "Data Management" toolbox can be used to apply the same symbology to other debris density maps. Figure 23 shows an example debris density map for the San José Island, Texas, study site.

The specific settings used in ArcGIS Pro for creating the density heatmap in Figure 23 are shown in Figure 24. (Note that the same functionality exists in ArcMap 10.x.) Additional display settings used in generating the map in Figure 23 are: 1) symbology set to "equal interval" with 10 classes and a green-to-red color scheme and 2) transparency set

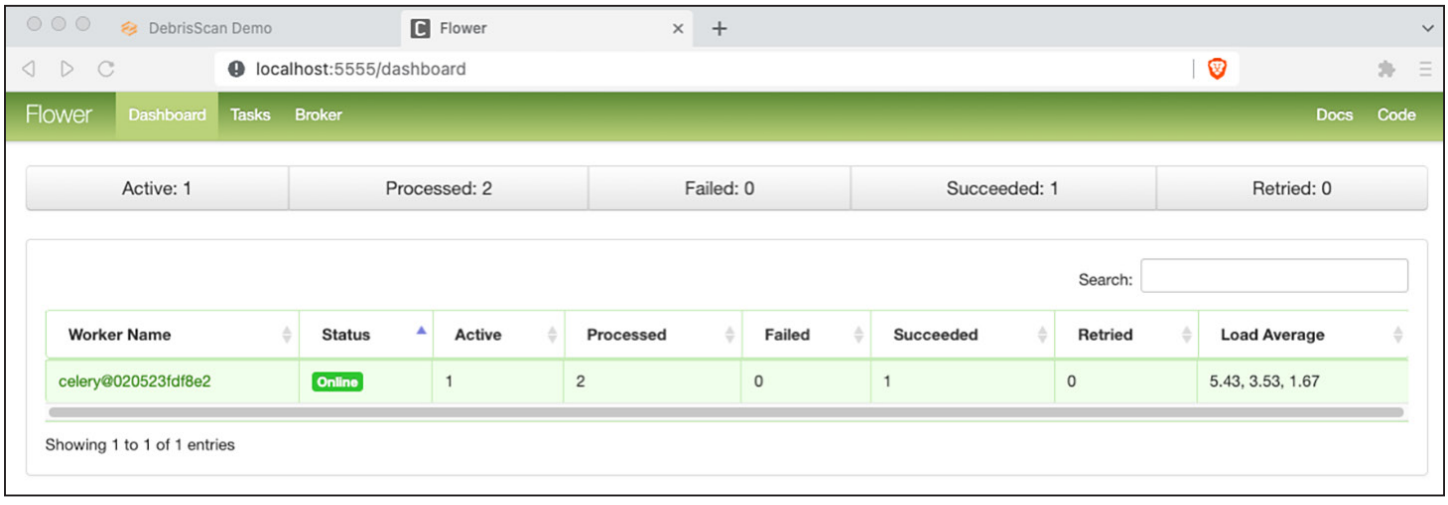

Figure 22: DebrisScan's administrator dashboard, powered by the open-source Flower package. This dashboard allows DebrisScan's administrators to see the current number of jobs, the number of celery workers, the status of DebrisScan's back-end database, and more.

to 50%. There are some edge effects with running this operation with these parameter settings on such a small area, and users may want to experiment with different settings. Again, however, it is strongly recommended to stick with consistent settings, if debris density heatmaps are going to be consistently compared across different areas and/or different time epochs.

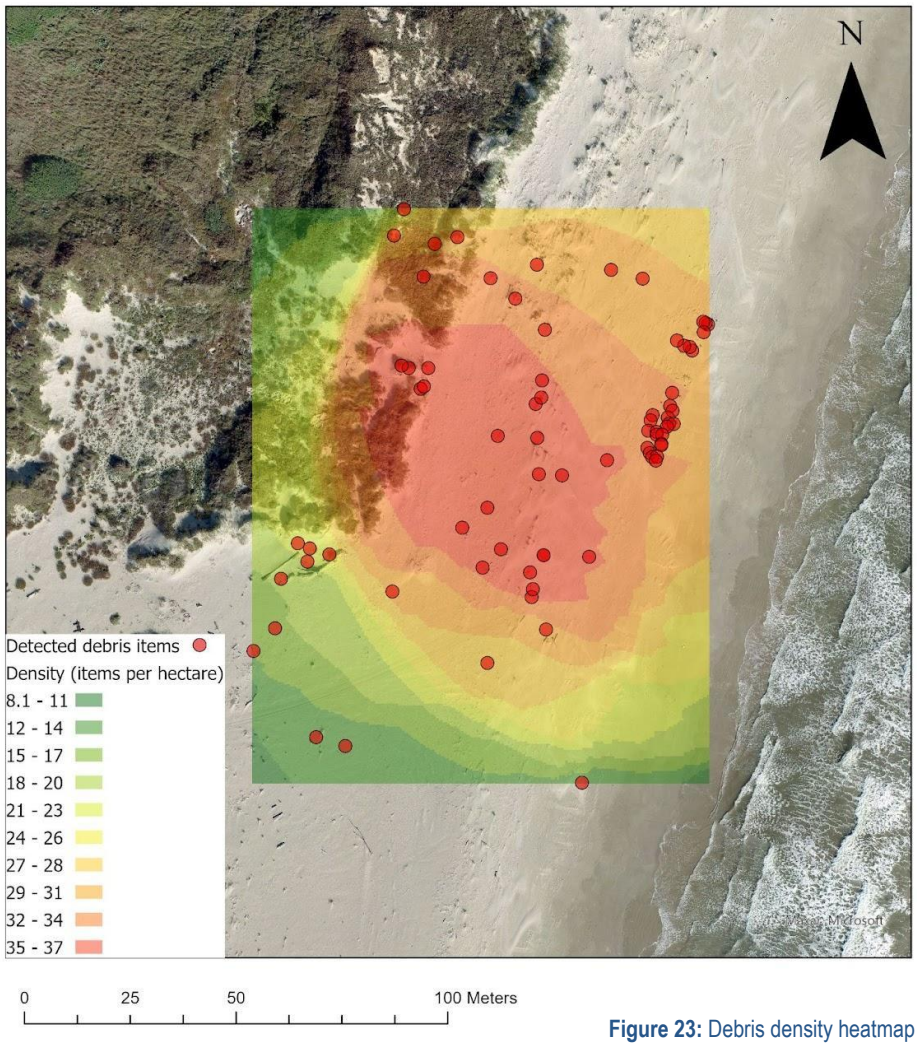

Figure 23: Debris density heatmap for San José Island, TX, site.

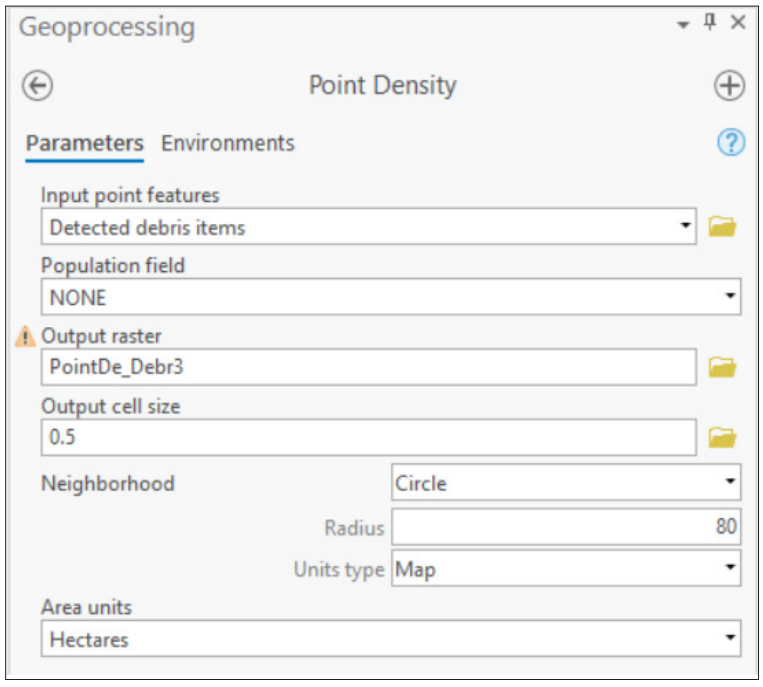

Figure 24: Settings used in ArcGIS Pro in creating the debris density heatmap.

Marine debris on San José Island, Texas. Credit: NOAA NCCOS.

Section 6 Conclusions and Recommendations for Future Work

This project demonstrated the enhancement in marine debris shoreline surveys achievable using UAS, PI, and ML. UAS were shown capable of enhancing efficiency of marine debris shoreline surveys and providing georeferenced imagery. Meanwhile, imagery from polarimetric cameras was found to substantially improve marine debris detection results over RGB spectral information only, including in both visual recognition of debris, as well as automated debris classification. The mean improvement of 15 percentage points in overall classification accuracy achieved by including PI-derived bands was noteworthy.

ML models for debris detection showed similar promise. Although not yet at the level of human detection and recognition, the ML model tested in this study performed nearly as well as human visual interpretation when detecting large and distinct debris, such as tires and fishing nets, in RGB imagery acquired from UAS. The model suffered from class confusion with small, ambiguous items, which often had two characteristics in common: 1) classification based on the object's material type (e.g., plastic, etc.) rather than use (e.g., fishing) and 2) debris items that were primarily composed of easily fragmented materials, such as plastic or foam, which resulted in a wide range of object shapes and sizes per category, confusing the ML models.

Overall, the results of this project provide strong indication that UAS with advanced imaging payloads and ML are technologies capable of complementing and enhancing existing MDP monitoring and detection methods. Notwithstanding these advances, there are still some limitations that merit discussion. Investments in equipment (including not only the remote aircraft but also data storage and other IT resources), as well as training, experience, and certifications are needed to implement a safe, successful UAS marine debris program. Beyond experience in UAS operations, personnel must also have some level of expertise in photography to acquire good-quality imagery in challenging environmental conditions, as well as experience with SfM photogrammetry software to generate orthoimagery, from which georeferenced debris items can be extracted.

Flights without an FAA waiver are currently limited to VLOS operations. Thus, each flight can typically cover a stretch of shoreline no longer than a few hundred meters in either direction from the base of operations, and frequent relocation of the base of operations along the shoreline is needed to cover sites extending a few kilometers or more.

Operation of drone. Credit: NOAA NCCOS.

Weather conditions can also limit operations, as winds can be too difficult to navigate (typically, 12–20 kt [22–37 km/h])) is a threshold, although this is highly dependent on the specific remote aircraft), or precipitation can ground flights entirely. To improve operational efficiency, ongoing improvements will be investigated for hybrid UAS, which support VTOL by switching to a fixed-wing configuration in flight. It is also of interest to investigate remote aircraft that perform well in windy coastal conditions to increase the number of potential operational days. Additionally, streamlining the process to obtain FAA waivers for beyond visual line-of-sight operations will greatly improve operational efficiency for large coastal sites.

Future improvements to the ML model and software are also recommended. The model used in this work was trained using data collected on the coastlines of Hawai‘i. While the Hawai‘i data covered a range of coastal morphologies, a larger and more geographically diverse training dataset would enable improvement in debris detection. It is highly recommended to acquire additional reference data for shorelines containing significant beach wrack and additional vegetation types not covered in this study, as well as for mixed environments that include built infrastructure, such as marinas. Reference data in saltmarshes and estuaries would also be particularly valuable for extending the model. In situ ground truthing is also recommended to validate and improve the training dataset. Ultimately, the robust training dataset developed through ongoing efforts is anticipated to

Padre Island, Texas. Credit: NOAA NCCOS.

be the most beneficial outcome of this ongoing work; new deep-learning algorithms and architectures come and go, but the underlying training data are "evergreen" and can be used to train new deep-learning frameworks as they become available. Additionally, further work is recommended to allow ML-based methods to integrate more seamlessly with the existing field of marine debris management. In particular, this will entail use of statistics, metrics, and terminology meaningful within the marine debris community, rather than those more common within the field of computer science.

Finally, it should be emphasized that this study focused on routine shoreline debris monitoring and did not address post-hurricane marine debris or floating debris. Some of the procedures developed in this study may be extensible to post-storm marine debris and potentially to floating debris as well. However, retraining of the ML model would likely be required for each of these applications. Post-hurricane imagery is often obtained at a coarser resolution, and the detected debris items may also be larger in size. Hence, training with coarser resolution imagery and larger objects (e.g., derelict vessels) would be a required step. Similarly, floating debris poses a number of additional challenges (e.g., partially submerged debris items, specular reflections from the water surface, etc.), and acquisition of suitable reference data would be a required next step.

References

Bao, Z., Sha, J., Li, X., Hanchiso, T., and Shifaw, E. (2018). Monitoring of Beach Litter by Automatic Interpretation of Unmanned Aerial Vehicle Images Using the Segmentation Threshold Method. *Marine Pollution Bulletin, 137*, 388–398. https://doi.org/10.1016/j.marpolbul.2018.08.009

Brooke, S., Graham, D., Jacobs, T., Littnan, C., Manuel, M., and O'Conner, R., (2015). Testing marine conservation applications of unmanned aerial systems (UxS) in a remote marine protected area. *Journal of Unmanned Vehicle Systems*, *3*(4), 237–251. https://doi.org/10.1139/juvs-2015-0011

Cito, J., Schermann, G., Wittern, J. E., Leitner, P., Zumberi, S., and Gall, H. C. (2017, May). An empirical analysis of the docker container ecosystem on github. In *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)* (pp. 323–333). Buenos Aires, Argentina. https://doi.org/10.1109/MSR.2017.67

Denes, L.J., Gottlieb, M.S., Kaminsky, B., and Huber, D. F. (1998, March). Spectropolarimetric imaging for object recognition. In J.M. Selander (ed.), *Proceedings SPIE Volume 3240, 26th AIPR Workshop: Exploiting New Image Sources and Sensors* (pp. 8–18). SPIE. https://doi.org/10.1117/12.300052

Fallati, L., Polidori, A., Salvatore, C., Saponari, L., Savini, A., and Galli, P. (2019). Anthropogenic Marine Debris assessment with Unmanned Aerial Vehicle imagery and deep learning: A case study along the beaches of the Republic of Maldives. *Science of The Total Environment*, *693*, 133581. https://doi.org/10.1016/j.scitotenv.2019.133581

Federal Aviation Administration (FAA), n.d. Part 107 Waiver. Online: https://www.faa.gov/uas/commercial_operators/part_107_waivers

Gonçalves, G., Andriolo, U., Pinto, L. and Bessa, F., (2020). Mapping marine litter using UxS on a beach-dune system: a multidisciplinary approach. *Science of The Total Environment, 706*, 135742.

Hengstmann, E., Grawe, D., Tamminga, M., and Fischer E. K. (2017). Marine litter abundance and distribution on beaches on the Isle of Rugen considering the influence of exposition, morphology and recreational activities. *Marine Pollution Bulletin, 115*, 297–306. https://doi.org/10.1016/j.marpolbul.2016.12.026

Herrera, F. (2022). *Enhancing Detection of Marine Debris with Polarimetric Imagery* [Master's thesis, Oregon State University]. https://ir.library.oregonstate.edu/concern/graduate_thesis_or_dissertations/1831cs647

Islam, M. N., Tahtali, M., and Pickering, M. (2019). Man-made object separation using polarimetric imagery. In *Proceedings Volume 11197, SPIE Future Sensing Technologies* (111971C). SPIE. https://doi.org/10.1117/12.2547475

Kylili, K., Hadjistassou, C., and Artusi, A. (2020). An intelligent way for discerning plastics at the shorelines and the seas. *Environmental Science and Pollution Research, 27*(34), 42631–42643. https://doi.org/10.1007/s11356-020-10105-7

Lebreton, L., Slat, B., Ferrari, F., Sainte-Rose, B., Aitken, J., Marthouse, R., Hajbane, S., Cunsolo, S., Schwarz, A., Levivier, A., Noble, K., Debeljak, P., Maral, H., Schoeneich-Argent, R., Brambini, R., and Reisser, J. (2018). Evidence that the Great Pacific Garbage Patch is rapidly accumulating plastic. *Scientific Reports, 8*, 4666. https://doi.org/10.1038/s41598-018-22939-w

Martin, C., Parkes, S., Zhang, Q., Zhang, X., McCabe, M. F., and Duarte, C. M. (2018). Use of unmanned aerial vehicles for efficient beach litter monitoring. *Marine Pollution Bulletin, 131*, 662–673. https://doi.org/10.1016/j.marpolbul.2018.04.045

Martin, C., Zhang, Q., Zhai, D., Zhang, X., and Duarte, C. M. (2021). Enabling a Large-Scale Assessment of Litter along Saudi Arabian Red Sea Shores by Combining Drones and Machine Learning. *Environmental Pollution, 277*,116730. https://doi.org/10.1016/j.envpol.2021.116730

Moy, K., Neilson, B., Chung, A., Meadows, A., Castrence, M., Ambagis, S., and Davidson, K. (2018). Mapping coastal marine debris using aerial imagery and spatial analysis. *Marine Pollution Bulletin, 132*, 52–59. https://doi.org/10.1016/j.marpolbul.2017.11.045

North Carolina Department of Transportation (NCDOT). n.d. *UxS Standard Operating Procedures: Flight Operations*. https://connect.ncdot.gov/resources/Aviation%20Resources%20Documents/NCDOT_UAS_SOP.pdf

ORBTL AI, 2021. *Final Model Report Machine Learning of Marine Debris*. ORBTL AI. https://cdn.coastalscience.noaa.gov/projects-attachments/455/FinalModelReport_MachineLearningofMarineDebris.pdf

Over, J. R., Ritchie, A. C., Kranenburg, C. J., Brown, J. A., Buscombe, D., Noble, T., Sherwood, C. R., Warrick, J. A., and Wernette, P. A. (2021). Processing coastal imagery with Agisoft Metashape Professional Edition, version 1.6—Structure from motion workflow documentation. *U.S. Geological Survey Open-File Report 2021–1039.* https://doi.org/10.3133/ofr20211039

References

Slocum, R.K., Wright, W., Parrish, C., Costa, B., Sharr, M., and Battista, T. A. (2019). *Guidelines for Bathymetric Mapping and Orthoimage Generation using sUxS and SfM, An Approach for Conducting Nearshore Coastal Mapping.* NOAA Technical Memorandum NOS NCCOS 265. Silver Spring, MD. https://doi.org/10.25923/07mx-1f93

Stokes, G. G. (1852). On the composition and resolution of streams of polarized light from different sources. *Transactions of the Cambridge Philosophical Society, 9,* 399

Tan, M., Pang, R., and Le, Q. V. (2020). Efficientdet: Scalable and efficient object detection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 10781–10790). Seattle, WA, USA. https://doi.ieeecomputersociety.org/10.1109/CVPR42600.2020.01079

van Lieshout, C., van Oeveren, K., van Emmerik, T., and Postma, E. (2020). Automated river plastic monitoring using deep learning and cameras. *Earth and Space Science, 7(*8*)*, e2019EA000960. https://doi.org/10.1029/2019EA000960

Veenstra, T. S., and Churnside, J. H. (2012). Airborne sensors for detecting large marine debris at sea. *Marine Pollution Bulletin, 65(*1–3*)*, 63–68. https://doi.org/10.1016/j.marpolbul.2010.11.018

Vendome, C., and Poshyvanyk, D. (2016). Assisting developers with license compliance. In *Proceedings of the 38th International Conference on Software Engineering Companion* (pp. 811-814). Austin, TX, USA. https://doi.org/10.1145/2889160.2889259

Verhoeven, G. (2007). Did the digital (r)evolution change the concept of focal length? *AARGnews, 34*, 30–35.

Winans, W. R. (2021). *Automatic detection of Hawai'i's shoreline stranded mega-debris using deep learning-based object detection* [Master's thesis, University of Hawai'i at Mānoa]. https://scholarspace.manoa.hawaii.edu/handle/10125/75937

Wolf, M., van den Berg, K., Garaba, S. P., Gnann, N., Sattler, K., Stahl, F., and Zielinski, O. (2020). Machine learning for aquatic plastic litter detection, classification and quantification (APLASTIC-Q). *Environmental Research Letters, 15*(11), 114042. https://doi.org/10.1088/1748-9326/abbd01

Yan, L., Chen, W., Zhang, F., Yang, B., and Xiang, Y. (2020). *Polarization Remote Sensing Physics*. Springer Singapore. https://doi.org/10.1007/978-981-15-2886-6

Zaffetti, R. P. (2019). *CTDOT Unmanned Aircraft Systems (UAS) Standard Operating Procedures*. Connecticut Department of Transportation. https://portal.ct.gov/-/media/DOT/documents/AEC/UAS/UAS_SOP_2019-04.pdf?la=en

U.S. Department of Commerce

Gina Raimondo, *Secretary*

National Oceanic and Atmospheric Administration

Richard W. Spinrad, Ph.D, *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/.