NOAA Technical Memorandum NMFS-PIFSC-83

Automated Analysis of Underwater Imagery: Accomplishments, Products, and Vision

A Report on the NOAA Fisheries Strategic Initiative on Automated Image Analysis 2014–2018

Benjamin L. Richards, Oscar Beijbom, Matthew D. Campbell, M. Elizabeth Clarke, George Cutter, Matthew Dawkins, Duane Edington, Deborah R. Hart, Marie C. Hill, Anthony Hoogs, David Kriegman, Erin E. Moreland, Thomas A. Oliver, William L. Michaels, Michael Placentino, Audrey K. Rollo, Charles H. Thompson, Farron Wallace, Ivor D. Williams, and Kresimir Williams

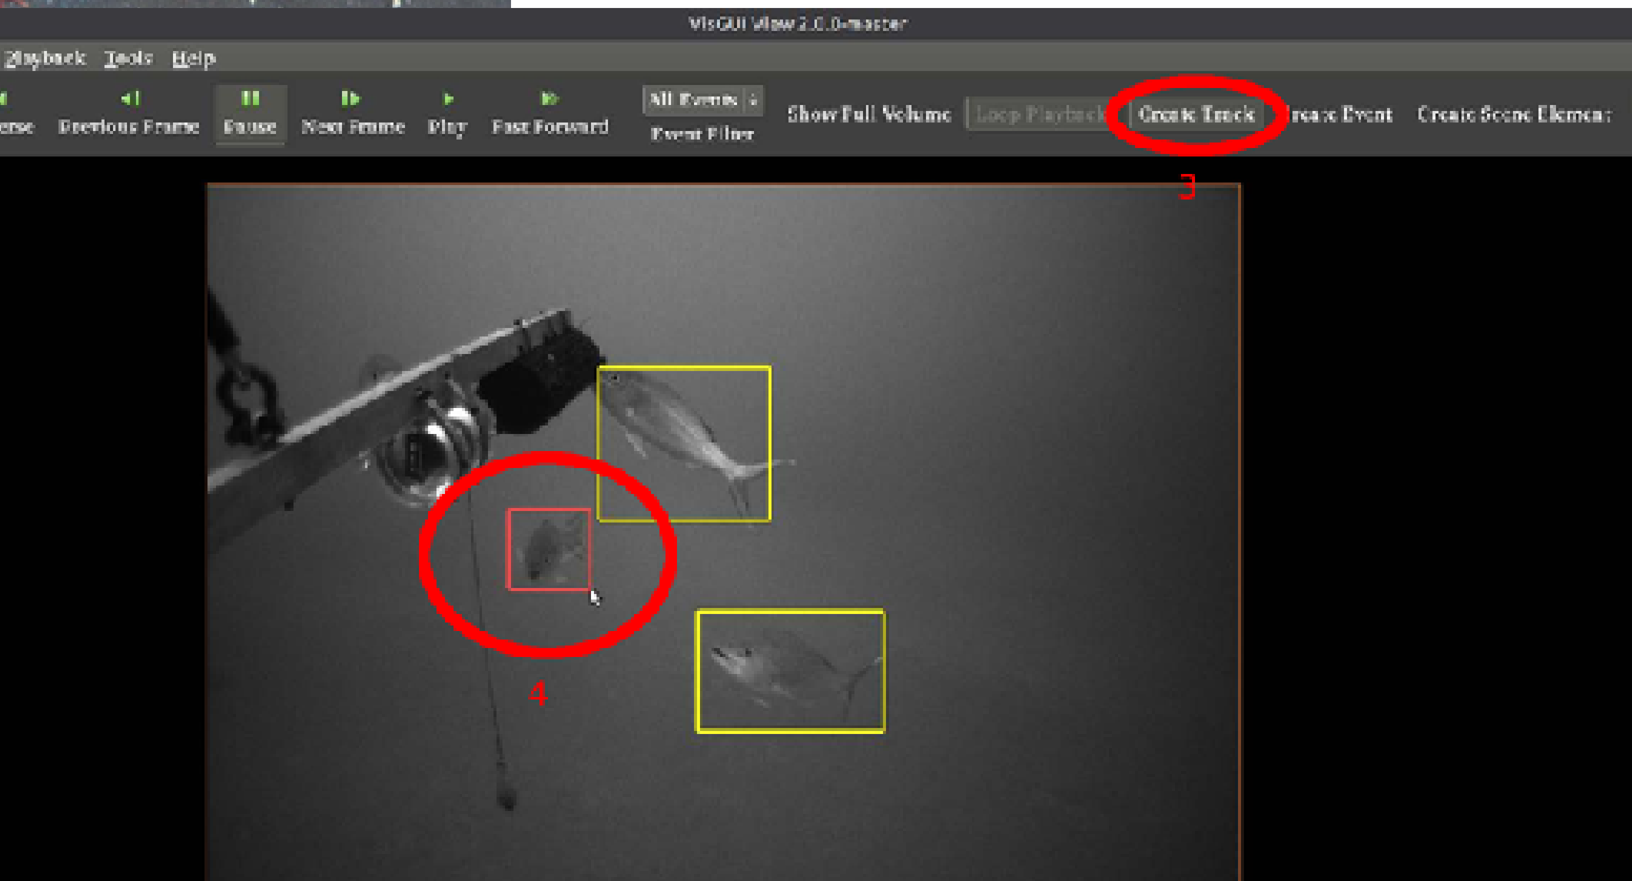

May 2019

U.S. DEPARTMENT OF COMMERCE

National Oceanic and Atmospheric Administration
National Marine Fisheries Service
Alaska Fisheries Science Center

Automated Analysis of Underwater Imagery: Accomplishments, Products, and Vision

A Report on the NOAA Fisheries Strategic Initiative on Automated Image Analysis 2014–2018

Benjamin L. Richards[1], Oscar Beijbom[2], Matthew D. Campbell[3], M. Elizabeth Clarke[4], George Cutter[5], Matthew Dawkins[6], Duane Edgington[7], Deborah R. Hart[8], Marie C. Hill[1], Anthony Hoogs[6], David Kriegman[9], Erin E. Moreland[10], Thomas A. Oliver[1], William L. Michaels[11], Michael Placentino[12], Audrey K. Rollo[1], Charles H. Thompson[3], Farron Wallace[10], Ivor D. Williams[1], and Kresimir Williams[10]

[1]Pacific Islands Fisheries Science Center
[2]nuTonomy
[3]Southeast Fisheries Science Center
[4]Northwest Fisheries Science Center
[5]Southwest Fisheries Science Center
[6]Kitware, Inc.
[7]Monterey Bay Aquarium Research Institute
[8]Northeast Fisheries Science Center
[9]University of California at San Diego
[10]Alaska Fisheries Science Center
[11]NOAA Fisheries Office of Science and Technology

NOAA Technical Memorandum NMFS-PIFSC-83
https://doi.org/10.25923/0cwf-4714

May 2019

U.S. Department of Commerce
Wilbur L. Ross, Jr., Secretary

National Oceanic and Atmospheric Administration
RDML Tim Gallaudet, Ph.D., USN Ret., Acting NOAA Administrator

National Marine Fisheries Service
Chris Oliver, Assistant Administrator for Fisheries

Recommended citation:

Richards BL, Beijbom O, Campbell MD, Clarke ME, Cutter G, Dawkins M, Edgington D, Hart DR, Hill MC, Hoogs A, Kriegman D, Moreland EE, Oliver TA, Michaels WL, Piacentino M, Rollo AK, Thompson C, Wallace F, Williams ID, Williams K. 2019. Automated Analysis of Underwater Imagery: Accomplishments, Products, and Vision. NOAA Technical Memorandum NOAA-TM-NMFS-PIFSC-83. 59 p. https://doi.org/10.25923/0cwf-4714.

Copies of this report are available from:

Science Operations Division
Pacific Islands Fisheries Science Center
National Marine Fisheries Service
National Oceanic and Atmospheric Administration
1845 Wasp Boulevard, Building #176
Honolulu, Hawaii 96818

Or online at:

https://www.pifsc.noaa.gov/library/

Cover: Photos courtesy of NOAA Fisheries.

Table of Contents

List of Tables

List of Figures

Executive Summary

Stock assessments, as required by the Magnuson-Stevens Fishery Conservation and Management Act, are the cornerstone of U.S. marine resource management. However, inadequate abundance data remains an impediment.

Increasingly, fisheries surveys are conducted using imaging systems that allow for efficient and non-lethal collection of voluminous data to fill this gap. However, the generated image data volumes exceed human analysis capacity. Automated image processing methods exist, but are nascent within the marine science community.

The NOAA Fisheries Office of Science and Technology initiated a Strategic Initiative (SI) on Automated Image Analysis with the goal of creating an open-source software toolkit allowing for automated analysis of optical data streams to provide fishery-independent abundance estimates for use in stock assessment.

The SI was directed by a research board comprising representatives from each of the NOAA Fisheries Science Centers (SC) as well as academic and private industry partners. Over the course of its five-year term, the SI developed two main products, the Video and Image Analytics for a Marine Environment (VIAME) open-source software toolkit and the CoralNet web-based solution for benthic image analysis.

CoralNet has become the operational image analysis tool for the Pacific Islands Fisheries Science Center (PIFSC) Coral Reef Ecosystem Program (CREP), accounting for more than 1 million annotations comprising more than 100,000 images.

VIAME has been released on GitHub as an open-source, publicly available software tool. Computing hardware has been procured and training sessions have been conducted at each NOAA Fisheries Science Center. VIAME is currently being used within the analysis workflow for (1) CamTrawl— AFSC Walleye Pollock assessment; (2) HabCam—NEFSC Scallop assessment; and (3) MOUSS—PIFSC Deep7 Bottomfish assessment.

Although VIAME is primarily used for underwater imagery, it is based on a generic, pipelined, deep learning-based processing system that applies to any domain. VIAME includes a graphical user interface (GUI) and modeling capabilities for users to create new automated analytics, interactively without any programming, enabling direct applicability to other NOAA imaging domains such as protected species (e.g. marine mammals, turtles), plankton, and electronic monitoring. Efforts are underway to raise awareness of VIAME and nascent collaborations exist within these domains.

VIAME and CoralNet exceeded expectations and continue to grow with increased utility spanning a broad range of programs. With major development complete, support for ongoing maintenance and customer support is needed to ensure continued utility and to support project-specific development. To maximize development imagery should be curated with a priority on access for machine learning.

Introduction

The framework for fisheries management in the United States is specified by The Magnuson-Stevens Fishery Conservation and Management Act (Magnuson-Stevens Fishery Conservation and Management Act 2007), which requires that managed fish stocks undergo periodic assessment to determine if they are overfished or are experiencing overfishing. Fishery stock assessments generated by NOAA Fisheries are the cornerstone of marine resource management in the United States. Assessments provide high-quality scientific information to marine resource managers to address (1) current stock status relative to established targets, (2) the level of sustainable catch a given stock can support and, if a stock becomes depleted, (3) what steps are required to rebuild it to health abundance levels.

A basic stock assessment requires data on fish abundance, biology (e.g. age, growth, fecundity), and catch (Quinn and Deriso 1999). While demands to continually improve stock assessments are high, a lack of adequate input data, particularly more precise, accurate, efficient and timely scientific surveys of fish abundance and their associated habitat and ecosystem, remains an impediment to accuracy and precision (Mace et al. 2001). Such input data can be derived from both fishery-dependent and fishery-independent sources. Fishery-dependent data, which generally include estimates of catch, effort, and size structure, are derived directly from the fishery through vessel and dealer reports. Fishery-independent data fall into the same general categories, but are collected independently from the fishery, often through dedicated surveys.

Recent developments in low-cost autonomous underwater vehicles (AUVs), stationary camera arrays, and towed vehicles have made it possible for fishery scientists to begin using optical data streams (e.g. still and video imagery) to generate species-specific, size-structured abundance estimates for different species of marine organisms. Increasingly, NOAA Fisheries and other agencies are employing camera-based surveys to estimate size-structured abundance for key stocks. While there are many benefits to optical surveys, including reduced inter-observer error as well as the ability to audit the observations and generate high sample sizes with reduced personnel and days at sea, the volume of optical data generated quickly exceeds the capabilities of human analysis.

Automated image processing methods have been developed and utilized in the human surveillance, biomedical, and defense domains for some time (LeCun et al. 2015; Szeliski 2010) and there are currently many open-source computer vision libraries and packages available on the internet. In the marine science environment, however, computer vision has yet to reach its full potential. Techniques for automated detection, identification, measurement, tracking, and counting fish in underwater optical data streams do exist (Chuang et al. 2014a, 2014b, 2013, 2011; Williams et al. 2016), however, few of these systems are fully automated, with all of the functions required to produce highly successful and accurate results.

Marine scientists rarely possess formal programming and development experience. Hence, existing solutions typically exist as one-off, localized applications, specific to particular analysis tasks. As such, they are generally non-transferrable as functional applications with utility across the domain. Consequently, with few exceptions (Huang et al. 2012; Williams et al. 2012; Chuang et al. 2014b; Chuang et al. 2014a; National Research Council 2014; Fisher et al. 2016; and

Williams et al. 2016) there has been little operational use of automated analysis within the marine science community.

In response to this need, in 2011, the NOAA Fisheries OST initiated a Strategic Initiative on Automated Image Analysis (SI). The mission of this SI was to develop guidelines, set priorities, and fund projects to develop broad-scale, standardized, and efficient automated tools for the analysis of optical data for use in stock assessment. The goal is to create an end-to-end open source software toolkit that allows for the automated analysis of optical data streams and in turn provide fishery-independent abundance estimates for use in stock assessment.

Methods

The NOAA Fisheries Strategic Initiative on Automated Image Analysis was envisioned by the NOAA Fisheries Science Board in 2011 as a research board consisting of representatives from each of the NOAA Fisheries Science Centers, academia, and the private sector. The rationale was to bring together a diverse group of experts from across the fisheries science, machine vision, and artificial intelligence domains to identify broad goals and solutions to span the myriad proposals received through the ST Advanced Sampling Technology Working Group (ASTWG). Ideas were solicited from across all six NOAA Fisheries Science Centers and, in 2013, a workshop was convened under the direction of the National Research Council (NRC) (National Research Council 2014).

The NRC workshop catalyzed collaboration between the marine science and computer vision communities and the solidification of the SI committee (Table 1). The SI considered both "top-down" and "bottom-up" approaches, soliciting input from each SC representative regarding key optical data streams existing at their SC that (1) could be informative to an existing stock assessment and (2) could not be fully analyzed by existing human resources. Optical data streams were categorized by physical properties of the sensor (e.g. mono vs stereo cameras, color vs grayscale, natural vs artificial light) and by the target of interest (e.g. fish underwater, fish on deck, marine mammals, corals). Finally, SC representatives were asked to identify any existing automated processing capabilities in existence at their SC.

Table 1. NOAA Fisheries Strategic Initiative on Automated Image Analysis committee members.

Name	Role	Affiliation
Benjamin L. Richards	Chair	NOAA Pacific Islands Fisheries Science Center
M. Elizabeth Clarke	Member	NOAA Northwest Fisheries Science Center
George Cutter	Member	NOAA Southwest Fisheries Science Center
Debora R. Hart	Member	NOAA Northeast Fisheries Science Center
Charles H. Thompson	Member	NOAA Southeast Fisheries Science Center
Kresimir Williams	Member	NOAA Alaska Fisheries Science Center
Clay Cuntz	Member	Google, Inc.
Alexandra Branzan-Albu	Member	University of Victoria
Duane Edgington	Member	Monterey Bay Aquarium Research Institute
Anthony Hoogs	Member	Kitware, Inc.

Name	Role	Affiliation
David Kriegman	Member	University of California, San Diego
Michael Piacentino	Member	Stanford Research Institute
Lakshman Prasad	Member	Los Alamos National Laboratories
William L. Michaels	Liaison	NOAA Fisheries Office of Science and Technology

The SI committee reviewed and ranked the information provided in terms of (1) national importance of the stock, (2) existence of pilot automated processing capabilities, (3) complexity of the optical data set. Following this ranking, three data sets were chosen as the basis for automated image analysis development: CamTrawl imagery pertaining to the Alaska walleye pollock (*Gadus chalcogrammus*) fishery, HabCam imagery pertaining to the northeast sea scallop (*Placopecten magellanicus*) fishery, and benthic imagery pertaining to coral reef ecosystems throughout the southeast and Pacific islands regions.

Data Sets

CamTrawl: Walleye Pollock (Alaska Fisheries Science Center)

Walleye pollock populations support the largest fishery by volume in the United States, and is one of the largest fisheries in the world (NMFS 2017). To date, the stock has been managed by assessing population abundance and size composition using annual acoustic surveys (Ianelli et al. 2009). Trawl samples are also obtained to identify the species and size composition of fish aggregations (Honkalehto et al. 2012). Researchers at the Alaska Fisheries Science Center (AFSC) developed a camera system (CamTrawl) that is placed within the survey trawl, allowing for increased precision in identifying fish schools, especially where different fish species or sizes occur in several distinct depth layers (Williams et al. 2010). CamTrawl collects stereo-image pairs that are analyzed to provide depth-, time-, and species-specific size-structured abundance information.

CamTrawl records stereo-images at two to four (2–4) frames per second, resulting in millions of images within a given survey. Automated image processing represents the only viable analytical approach for extracting timely abundance estimates for the stock. Fish are imaged within the constrained trawl environment and against a uniform background. Fish targets are extracted from both left and right camera images, matched, classified to species, and sized using stereo triangulation to estimate the XYZ coordinates of, and Euclidean distance between, the head and tail of each target.

Prior to the SI, AFSC scientists had collaborated with computer vision scientists at the University of Washington Electrical Engineering department to develop algorithms to process CamTrawl images (Chuang et al. 2014ab; Williams et al. 2016).

HabCam: Scallops, demersal fish, benthic invertebrates (Northeast Fisheries Science Center)

Sea scallops support one of the most valuable fisheries in the U.S. They occur mainly at depths from 30 to 120 m on the main U.S. scallop ground of Georges Bank and the Mid-Atlantic Bight. The NOAA HabCam is towed at between 5 and 7 kt, capturing approximately 6 digital still photo pairs per second. While its principal target is sea scallops, it also images demersal finfish and a wide variety of benthic invertebrates. Since 2011, about 40 million images have been captured by NOAA HabCam surveys. Only about two percent of these images have been manually annotated.

Coral Reef Benthic Imagery (Pacific Islands Fisheries Science Center)

Coral reefs support significant fisheries in both the U.S. southeast and Pacific Islands regions (Heenan et al. 2016). The scale, severity, and frequency of threats to coral reefs have increased substantially in recent years (Burke 2011; De'ath et al. 2012; Hughes et al. 2018). Given the speed of change and the increasing severity of threats to coral reefs, scientists and managers need the capability to rapidly assess coral reef status, ideally over large representative areas, and to quantify changes. While a variety of metrics are used to assess reef status, the majority of coral reef surveys and monitoring programs gather information on percent cover of benthic organisms, particularly coral cover (De'ath et al. 2012; Johansson et al. 2013). Recently, many benthic monitoring programs have transitioned from in situ measurements of benthic cover to some form of photographic survey (Heenan et al. 2016).

The imagery and benthic data used in this study come from the PIFSC, Ecosystem Science Division (ESD), Pacific Reef Assessment and Monitoring Program (Pacific RAMP), which is part of NOAA's National Coral Reef Monitoring Program. Survey sites are randomly allocated within three depth strata comprising all hard bottom habitats in less than 30 m of water and encompass substantial variability in habitat type, reef condition, and benthos, including coral assemblage and abundance. For this effort, images from 468 sites within American Samoa and 913 sites within the main Hawaiian Islands were used. At each site, 30 benthic images were captured along one or two transect lines with a total combined length of 30 m. Images were collected using digital cameras, maintained at a standard height above the substrate using a 1-m PVC monopod. No artificial lighting was used; instead, cameras were manually white-balanced immediately before each transect.

Prior to the SI, images were typically analyzed manually using point annotation software, such as Coral Point Count with Excel extensions (CPCe), photoQuad, pointCount99, PhotoGrid, or Biigle (Kohler and Gill 2006; Langenkämper et al. 2017; Porter et al. 2001; Trygonis and Sini 2012). CPCe represented a significant step forward compared to prior *ad hoc* means, in that CPCe employed an integrated interface with Microsoft Excel. Analysts could develop a unique set of target codes, and overlay points on each image in a stratified random format. CPCe would also generate a summary file for each site or set of photos. Significantly, CPCe was free, so it was widely used by researchers.

CPCe suffers two core limitations. First, it is purely a manual annotation tool, with no capacity for automation. As mentioned earlier, manual annotation of survey imagery is time consuming and expensive due to the high cost of labor, which not only limits the amount of survey data that can feasibly be analyzed, but also often leads to significant temporal lags before results become available, reducing their utility. Secondly, individual annotation data files are generated for each image. These summary data files are linked to specific image files, and the links break with any modification to the folder structure used for data management.

Software Development

Software development was undertaken by two main contractors, Kitware Inc. and the Stanford Research Institute (SRI International). Kitware was tasked with development of the overall software platform while SRI was responsible for development of discrete modules and products.

VIAME: Video and Image Analytics for a Marine Environment

Across the Automated Analysis Strategic Initiative (AIASI) and the broader marine research community, many algorithms and corresponding software modules have been developed for image and video analytics on a wide variety of data sources. Many of these algorithms address similar problems but were developed independently at different research centers, resulting in different, incompatible implementations that can be difficult to re-use and compare. In addition, many marine scientists are not programmers, and do not have access to programming resources with computer vision expertise to implement image and video analytics.

VIAME (Dawkins et al. 2017) was developed to address both of these problems through a comprehensive set of capabilities shown in Figure 1. Utilizing the open-source video analytics toolkit, Kitware Image and Video Exploitation and Retrieval (KWIVER) (Fieldhouse et al. 2014), VIAME enables the rapid integration of visual analytics modules into a pipelined architecture. Implementation challenges related to parallel processing and sequential operations are addressed and hidden from the algorithm module developer. User interfaces, databases, evaluation/scoring capabilities and other useful standalone tools are included in VIAME as well. Algorithms from multiple FSC's are integrated, such as length measurement from Alaska FSC.

To address the challenge of applying VIAME algorithms to new data sets and problems without programming, two deep-learning capabilities were developed within VIAME that enable users to create object detectors, classifiers and other analytics through user interfaces as shown in Figure [VIAME-capabilities-architecture]. Through image search and interactive query refinement, users can quickly build a complete detection and classification capability for a novel problem and then run it on any amount of imagery or video. For more challenging analytical problems, users can manually annotate images and then train a deep learning detection and classification capability specific to their problem. Both methods were successfully used by marine scientists during the VIAME training sessions to develop analytics on data sets that VIAME had not seen previously.

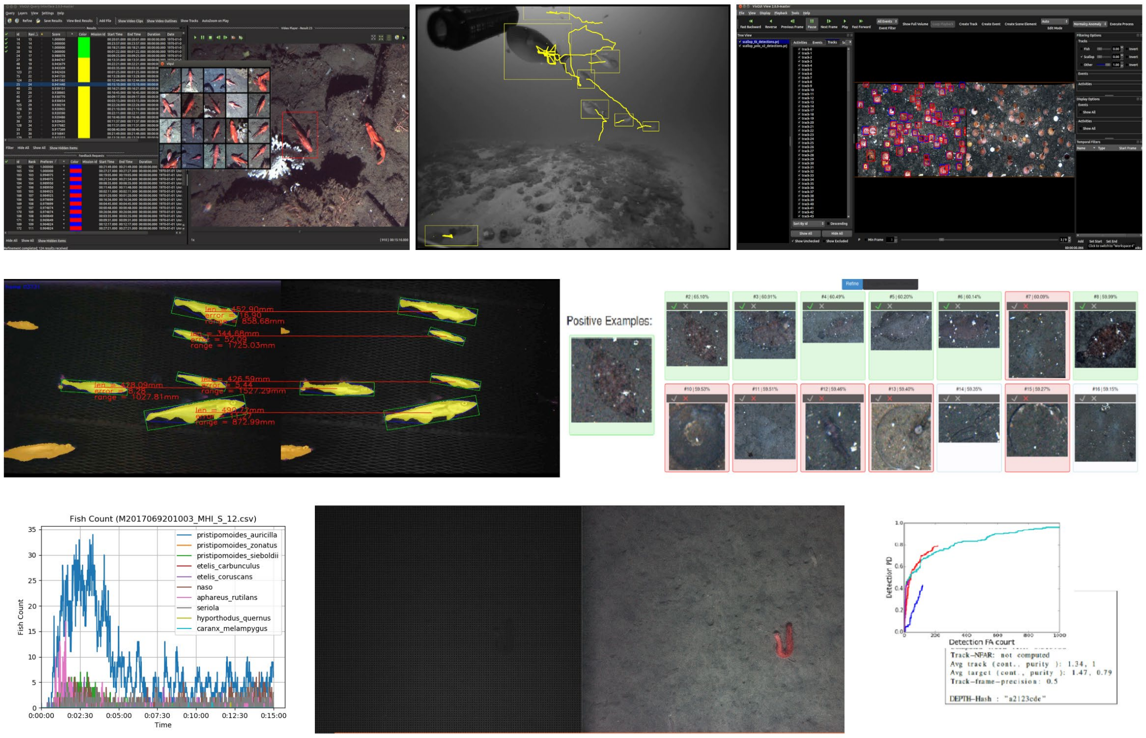

Figure 1. Visual overview of all of VIAME's functionality. From top left to bottom right: video search, object tracking, running multiple automated detectors from the annotation GUI, measurement using stereo, image search in a web browser, MaxN detection plotting, color correction, and algorithm scoring.

At the training sessions numerous requests arose for additional VIAME capabilities to address stereo video, aerial imagery mosaics, geographic information and others. While VIAME in its current form can address many NOAA problems, it seems clear that further extensions related to electronic monitoring or non-fish targets (e.g. marine mammals, seabirds, etc.) would enable a larger pool of scientists to benefit from it.

VIAME is hosted on GitHub (VIAME 2016/2018). GitHub is a public open-source software repository containing thousands of software tools. Updates and releases are made through GitHub, allowing anyone within NOAA or externally to download the software, see the source code, and modify it as desired. Binaries for common operating systems including Linux and windows are also available for users to install directly without compilation.

FLASK

FLASKS was designed as a means for scientists to rapidly count and classify fish observed in optical surveys using remote camera systems. During initial design and development, it was determined that recently available neural network capabilities for automated image processing would be beneficial. We integrated an Artificial Intelligence (AI) framework into the FLASK tools for automated fish counting and classification. During development it also became clear that many groups possessed large-image data sets without the bounding-box-level annotations required to provide usable training data for algorithm development. To meet this need, SRI developed a semi-automated rapid annotation tool that allows analysts to rapidly ingest and annotate their video as they train a novel neural network. The neural network is then used for performing rapid classification of other video sources producing a JSON or HDF5 output file containing all ROI bounding boxes with fish types. This data can then be parsed to provide fish types and counts. Recent work on FLASK has added new capabilities requested by scientists, including adding more robustness to the counts by adding temporal fish tracking and allowing for creating user definable training models based on different video sources and selectable sets of classes in each model.

CoralNet

Efforts at automated analysis of coral reef benthic imagery were advanced through the further development of the CoralNet web-based repository and a resource for benthic image analysis (Beijbom et al. 2015). CoralNet implements computer vision algorithms, which allow for fully semi-automated annotation, while also serving as a repository and collaboration platform. Unlike prior manual annotation tools, CoralNet provides a function where human analysts first identify targets manually, and these annotations are used to train a machine learning model. Once a suitable level of accuracy is achieved on a validation set, CoralNet can be used to automatically annotate new imagery, reducing time and effort required by human analysts. Through the web-based GUI, a user-defined number of points are randomly distributed across an image of the coral reef benthos. Two versions of the classification methods were developed and deployed as CoralNet Alpha and CoralNet Beta.

As benthic targets lack clear boundaries and a clear sense of shape, they are represented using texture and color descriptors. In CoralNet Alpha, images are first re-scaled to maintain a consistent pixel/mm ratio and are color-corrected using the *ColorChannelStretch* method (Beijbom et al. 2015). The Maximum Response filter bank is then used to encode rotational invariance by first filtering with bar and edge filters at different orientations and then outputting the maximum over the orientations. By cross-validating over different sizes we arrived at bar and edge filters with standard deviations of 1, 3, and 8 pixels along the short dimension, and circular filters standard deviation of 6 pixels, thus producing an 8-dimensional filter output vector. Color information is encoded by applying the filters to each color channel in the $L^*a^*b^*$ color space and then stacking the filter response vectors.

Texton maps were created using a dictionary of textons. Filter responses from each of nine classes were separately aggregated across images, and k-means clustering with 15 cluster centers was applied to each set of filter responses (Beijbom et al. 2015). Finally the cluster centers, or textons, from the different classes were merged to create a dictionary of 135 24-dimensional words. Texture descriptors are extracted by first applying the filters over a whole image, which yields a 24-dimensional feature vector for each image pixel. Filter responses are then mapped to the texton with smallest 2-norm distance, creating an integer valued texton map. The feature vector, or descriptor, is defined as the normalized histogram of textons around a patch of interest. Classification was finally performed using a support vector machine (SVM).

CoralNet Alpha was built using the algorithm described above (Beijbom et al. 2015) and ran on a deskside server. During the course of the SI, CoralNet was rewritten to support the annotation load of a government agency and to improve accuracy. AIASI support allowed further development of CoralNet and a transition from CoralNet Alpha to CoralNet Beta, which included transitioning from a desktop server to Amazon Web Services. Moving to cloud computing provides significantly greater uptime, robustness, and security, and computing resources can be elastically scaled up to a 100 machine cluster in order to process large workloads in a short time period. Accuracy was significantly improved by moving from color texton features and an SVM to a deep learning model. The particular network on CoralNet Beta is based on a convolutional neural network called VGG which was initially trained on ImageNet. We then retrained VGG, starting from the ImageNet weights, using 2.5 M annotations with 956 classes from 60,000 images that had been uploaded and manually annotated on CoralNet Alpha. For a specific source (a data set with specific set of labels), the final softmax layer of the network is trained with a modest set of manual annotations from that set. Figure 2 shows a comparison of the accuracy of CoralNet Alpha, CoralNet Beta, and human annotation. Improvements were also made to the CoralNet GUI including the addition of different annotation modes and upload methods.

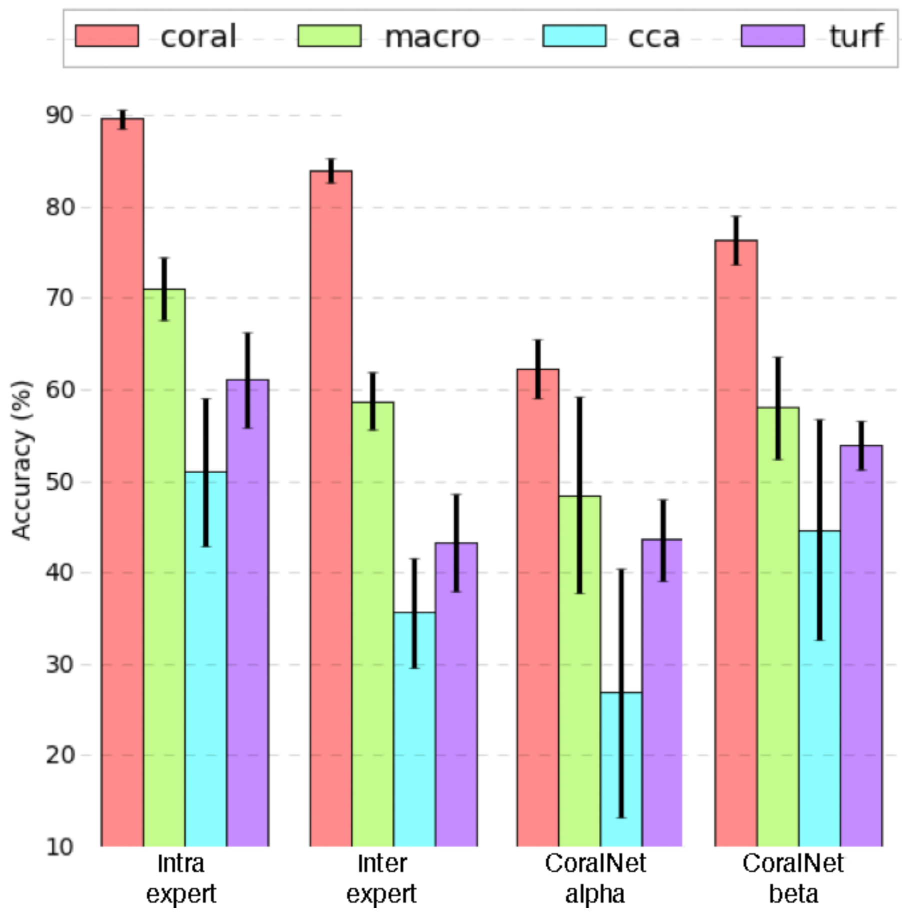

Figure 2. Accuracy of annotation using Cohen's Kappa as a metric for four classes (Coral, Macroalgae, Crustose Coralline Algae (CCA), and Turf Algae) when evaluated by (1) the Same expert re-annotating the same data (Intra Expert), a different human expert annotation (Inter Expert), Bag-of-Words classifier used in CoralNet Alpha (Texton), VGG convolutional network trained on CoralNet data (CoralNet beta).

Results

The Automated Image Analysis Strategic Initiative was successful in meeting its goal of developing an end-to-end open-source software toolkit allowing for the automated analysis of optical data streams to provide fishery-independent abundance estimates for use in stock assessment. It was also successful in its broader mission to develop guidelines, set priorities, and fund projects to develop broad-scale, standardized, and efficient automated tools for the analysis of optical data for use in stock assessment. The products created and improved by SI support offer a set of tools that are applicable not only to underwater image survey analysis, as prioritized by the SI, but also are proving useful to NOAA imagery from other domains.

VIAME: Video and Image Analytics for a Marine Environment

A key product developed under the Automated Image Analysis Strategic Initiative is the open-source computer vision software framework named VIAME: Video and Image Analytics for a Marine Environment (Dawkins et al. 2017). VIAME provides a common interface for several algorithm stages (stereo matching, object detection, object tracking, and object classification), multiple implementations of each, as well as unified methods for performance evaluation for different algorithms applied to the same task. The common open-source framework facilitates the development of additional image analysis modules and pipelines through continuing collaboration within the image analysis and fisheries science communities.

Initial specifications for VIAME included incorporation of algorithms previously developed by NOAA Fisheries personnel and implemented for CamTrawl and HabCam, as well as a fish detection module. This allows for access to existing proven methods as a core part of the toolkit and also provides examples of integrated plug-in modules, which could then be emulated for other algorithms.

VIAME is built on the concept of modular, dynamically-loadable plugins. The software can be divided into three core components: (1) the pipeline processing framework and infrastructure; (2) image processing elements that fit into the framework; and (3) auxiliary tools outside the streaming framework that provide training, graphical user interfaces (GUIs) and evaluation (Figure 3). The pipeline subsystem allows image processing elements to be implemented in the most popular languages used for computer vision (e.g. C, C++, Python, and MATLAB). VIAME provides a graphical interface (GUI) for creating new annotations (marking locations, bounding boxes, and class identity) for target objects, for visualizing individual object detections and filtering detections based on classification values, for iterative queries of image data sets that are used to create novel classifiers based on support vector machine (SVM) models, and for training of deep-learning models for detection and tracking. There are two evaluation tools included in VIAME, one for generating basic statistics for target detection performance compared with groundtruth (detection rates, specificity, false alarm rate) and a second for generating receiver operating characteristic (ROC) curves for detections which contain associated category probabilities.

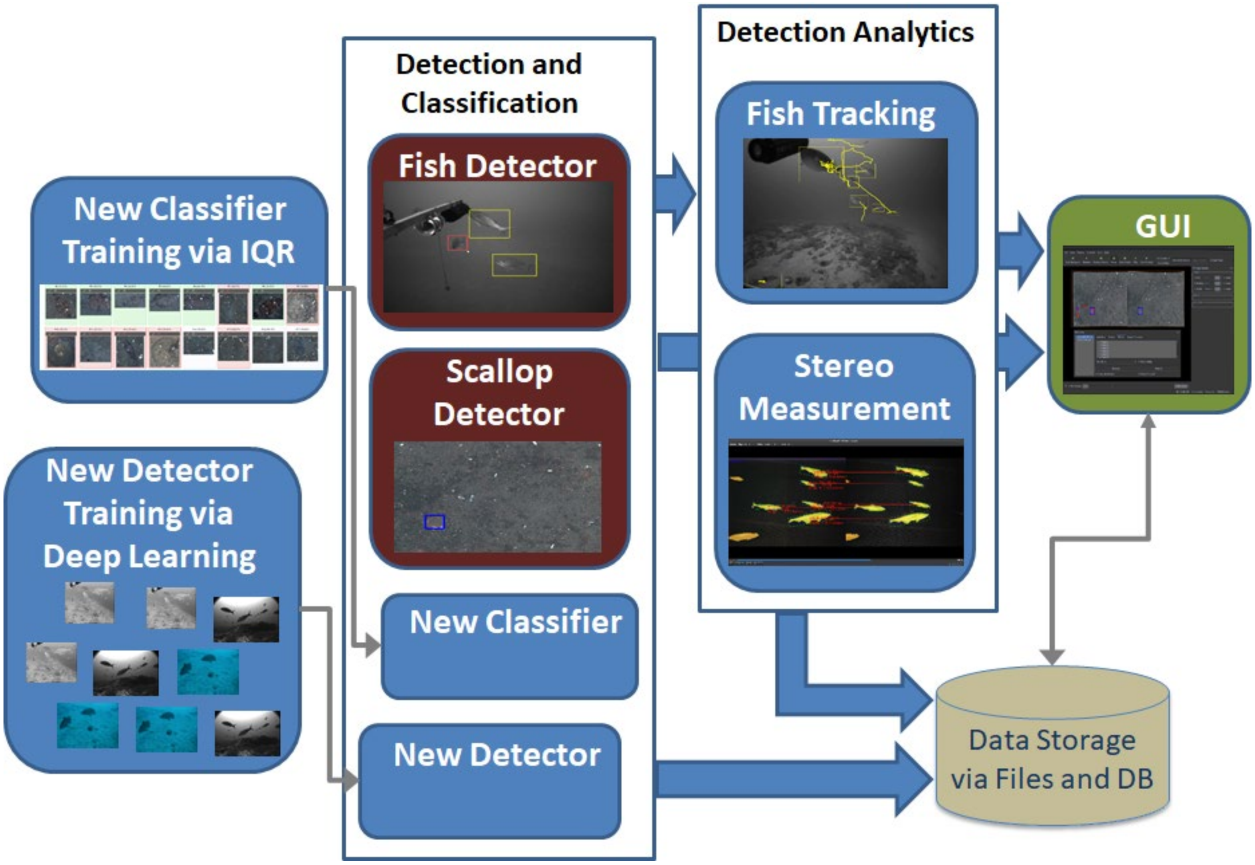

Figure 3. VIAME capabilities and data flow. Analysts can create new analytic modules for detection and classification unique to their data using GUIs and databases within VIAME.

VIAME has been released as cross-platform (Windows, Mac, Linux) open-source software on GitHub[1], with extensive documentation, tutorials, and training examples[2]. A number of initial algorithm modules have either been implemented or wrapped within the platform.

FLASK

SRI's fish detection and classification tool, called FLASK, was developed to rapidly annotate and classify fish type and fish counts from large volumes of recorded video, with minimal operator interaction. The tool has preprocessing functionality that identifies key fish features and segments fish images to allow for interactive annotation and training. The tool is able to rapidly pool similar features into clusters, which can then be viewed through a GUI that allows for rapid annotation (hundreds to thousands) of segmented objects simultaneously (Figure 4).

Following manual annotation of a typically less than 50 cluster sets, the annotation and associated metadata is input to a neural network for fine-tuning. Once the fine-tuning cycle is complete the user can (1) identify other untrained clusters in the video sources or (2) begin running video sources through the trained network, yielding automatically-annotated region of

---

[1] https://github.com/Kitware/VIAME
[2] https://viame.readthedocs.io/en/latest/

interest (ROI) bounding boxes around each fish in each video frame (Figure 4). These annotations are output in a parsable JSON file set containing a frame-by-frame record of all the fish in each video, allowing rapid fish classification and count reporting.

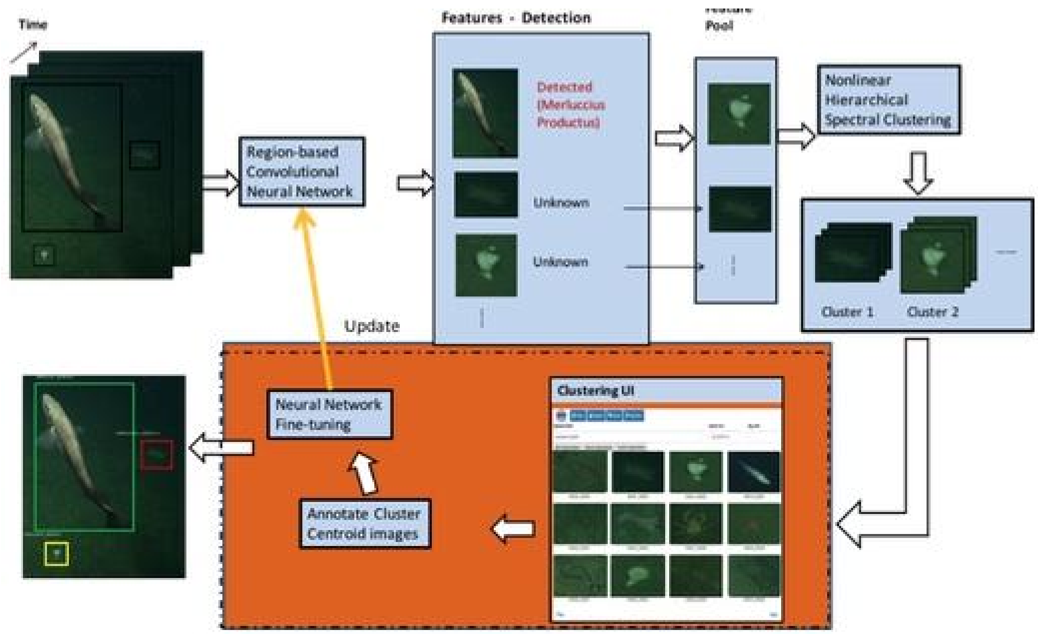

Figure 4. The orange region illustrates the GUI interface with the neural network. The users are presented with clusters of many similar objects that can be annotated by the users and then used for retraining of the neural network.

CoralNet

CoralNet is a web-based system for automatic analysis of benthic images acquired in the course of coral reef surveys. Under AIASI funding, the accuracy of CoralNet was significantly improved using deep learning, the system was transferred from a single host to a scalable distributed system on Amazon Web Services to handle large data sets, and workflows were improved to reduce the time users spend annotating images.

CoralNet preserves many desirable characteristics of CPCe, including a familiar interface, the ability for users to create a unique set of target descriptor codes, a function to overlay points randomly, and no acquisition or usage fees. The flat data structure used by CoralNet removes the inherent file structure problem in CPCe. Image metadata and annotations can be downloaded and archived and images can be randomly assigned to different analysts, a desired feature that was not possible using CPCe. The web-based deployment of CoralNet also makes it possible to easily collaborate with remote analysts.

CoralNet Alpha allowed users to upload image data sets, randomly distribute annotation points across those images, manually annotate a subset of the images using a web interface with study-specific labels (e.g., functional groups), and use those manual annotations as training data. It then automatically proposed labels for annotation points across the rest of the images and allowed users to verify and correct the proposals. In estimation of coral cover at the functional group level, CoralNet Alpha achieved a level of accuracy commensurate with human analysts (Beijbom et al. 2015), but challenges remained in identifying algal classes and many coral species. CoralNet Alpha characterized intra- and inter-expert variation, and found significant variation, particularly amongst algal classes.

The transition from CoralNet Alpha to CoralNet Beta resulted in significant improvements. Cloud-based processing significantly increased throughput, decreased latency, and reduced model training time. The transition from a bag-of-words style recognition system with texture and color features and SVM classifier to Deep Learning (deep CNN) increased classification accuracy and reduced end-to-end human analyst effort.

To date, 750,000 images have been uploaded to CoralNet Beta from 960 sources from around the globe, comprising over 27 million annotations. Currently, CoralNet supports nearly 1,000 registered users with more than 1,000 images uploaded and analyzed every day. Of the 750,000 images, more than 100,000 are from NOAA. The Pacific Islands Fisheries Science Center has been the primary NOAA Fisheries user of CoralNet. Their results are detailed in the PIFSC-specific results section below.

Region-Specific Results

Alaska Fisheries Science Center

The Alaska Fisheries Science Center (AFSC) has developed an automated image analysis protocol to process images collected by the CamTrawl system during acoustic pollock surveys. The system consists of paired still stereo images which are analyzed for fish length and species compositions. For length estimation, an automated software tool was developed in collaboration with computer vision specialists at the University of Washington (Williams et al. 2016). This code base was originally developed using the Matlab programming language and was then incorporated as a module within the VIAME framework by creating parallel routines for object detection, stereo correspondence, and triangulation using the open source Python language (Figure 5).

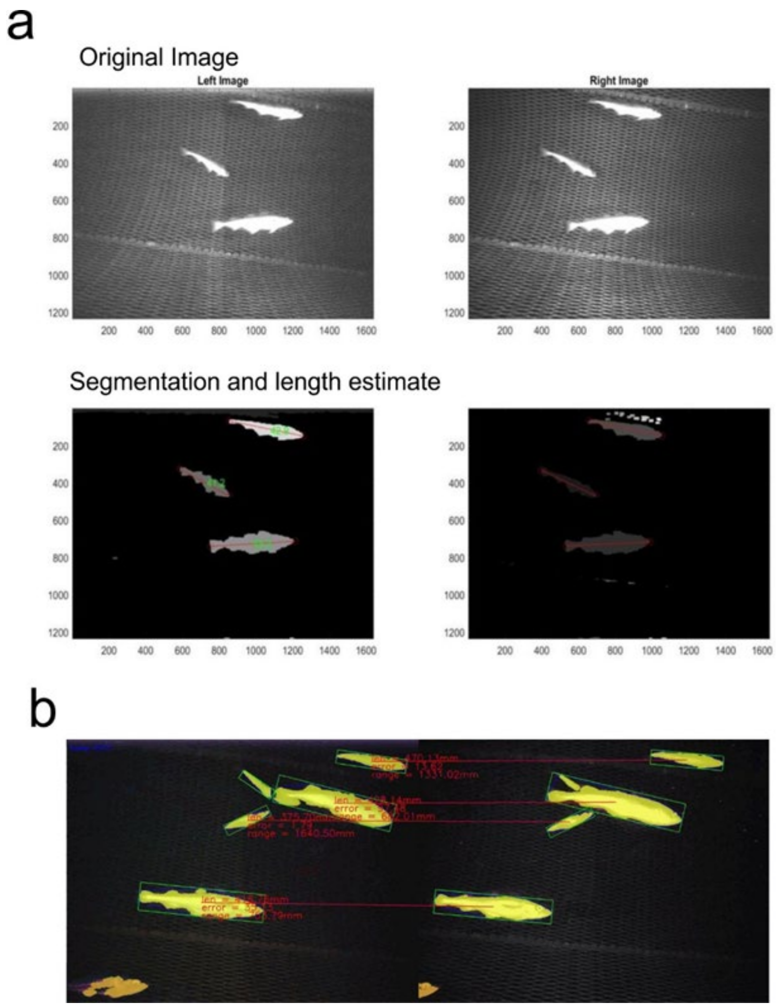

Figure 5. Automated length estimation from CamTrawl stereo-camera imagery. a) MATLAB processing routine and b) Python language routine using GMM detection (Jon Krall, Kitware). The latter process has been incorporated as a module within the VIAME automated image processing package.

The comparison of the original Matlab routine and the Python version show that the length estimations are similar (Table 2).

Table 2. A comparison of stereo image based fish length and range (distance) from camera estimates of using MATLAB- and Python-based programs. (Developed by Jon Crall, Kitware).

	MATLAB	Python (VIAME)
Fish Length (mm)	45.36 ± 4.54	44.75 ± 5.73
Fish Range (mm)	1191.41 ± 207.72	1217.51 ± 224.42
Error	4.27 ± 2.90	3.96 ± 3.51

This code base has been incorporated as one of the standard modules and analysis examples within the VIAME package. In addition, a species identification module was developed at the University of Washington and funded in part by AIASI (Figure 6; Wang et al. 2016a). This algorithm relies on Gaussian Mixture Model (GMM) to detect fish objects, a Bag-of-Features framework to extract object features and classify them using and Support Vector Machine (SVM) classifier.

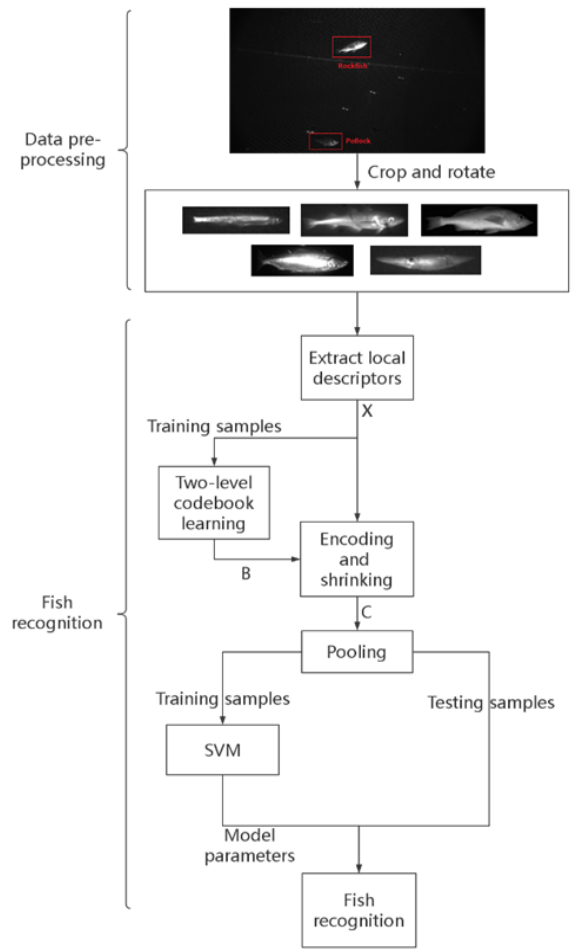

Figure 6. Diagram of fish species classification analysis for CamTrawl image data (from Wang et al. 2016a).

Results show a high degree of accuracy can be achieved with a lower number of fish categories (Table 3).

Table 3. The confusion matrix for CamTrawl data set using Bag-of-features framework and SVM classifier.

	Eulachon	Pollock	Rockfish	Salmon	Squid
Eulachon	113	3	1	2	0
Pollock	0	416	0	0	0
Rockfish	0	0	215	0	1
Salmon	1	1	1	156	0
Squid	1	5	0	0	110

An earlier version of the current identification algorithm was incorporated into VIAME in 2015. VIAME framework also includes and alternative fish detection module based on a deep "You Only Look Once" (YOLO) (Redmon et al. 2016) approach (Figure 7).

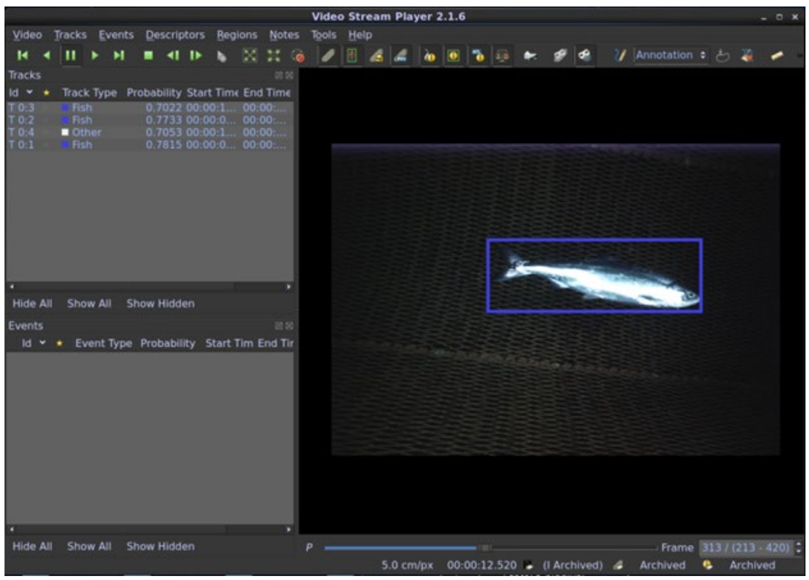

Figure 7. VIAME graphical user interface showing the results of fish detection algorithm based on deep Yolo architecture.

A collaborative project between AFSC, SWFSC, and UW with funding support from AIASI was established to develop automated methods for fish detection and tracking in ROV video. Custom routines that integrate detection and tracking components using a Deformable Part Model (DPM) for detection and multiple kernel tracking (Wang et al. 2016b). This closed-loop mechanism between detection and tracking greatly decreased the number of false detections, such as non-fish objects (Figure 8).

Figure 8. Example of a frame from automated ROV tracking using closed-loop with DPM features and motion features.

An AIASI sponsored detection/classification open challenge was carried out as part of a CVPR 2018 Workshop (Kitware Inc. 2018). AFSC provided an extensive manually annotated data set (~5,500 identified fish targets) based on still imagery.

During June 14–18, 2018, a VIAME installation and operationalization site visit was conducted in Seattle. During this meeting, the NWFSC and AFSC groups worked with VIAME representatives to develop working VIAME processing pipelines for existing image processing tasks at each center, using dedicated image processing hardware supplied through the AIASI funding. For AFSC, this included installation and running of the latest versions of the VIAME CamTrawl stereo fish length estimation, as well as investigating possible future VIAME uses such as fish detection in untrawlable habitat stereo imagery and video segments.

Additionally, representatives from the National Marine Mammal Laboratory worked with VIAME representatives to produce an initial detection model for seals from images collected during the annual ice seals aerial survey. Surveys for ice-associated seals rely on overlapping color and thermal imagery to detect and classify seals on the ice from a target altitude of 1000 ft. VIAME was not developed with this image model in mind, so training focused on preparing a training run on annotated color imagery. Augmentation modifications were implemented to address the high-resolution imagery. The model completed 4000 iterations and provided an opportunity to explore the process of reviewing and correcting results from a validation run. Approaches to improve performance were also discussed and 16-bit thermal imagery was shared with the development team.

Pacific Islands Fisheries Science Center

The Pacific Islands Fisheries Science Center (PIFSC) has implemented the CoralNet tool for operational annotation of benthic photoquadrat imagery from Reef Assessment and Monitoring Program (RAMP) surveys in the U.S. Pacific Island region. In trials using manually annotated imagery from the main Hawaiian Islands and American Samoa, a trained CoralNet Beta model was able to generate estimates of site-level coral cover that were highly comparable to those generated by human analysts (Pearson's r > 0.97, and with bias of 1% or less) (Figure 9). CoralNet Beta was also effective at estimating cover of common coral genera (Pearson's r > 0.92

and with bias of 2% or less in 6 of 7 cases), but performance was mixed for other groups including algal categories.

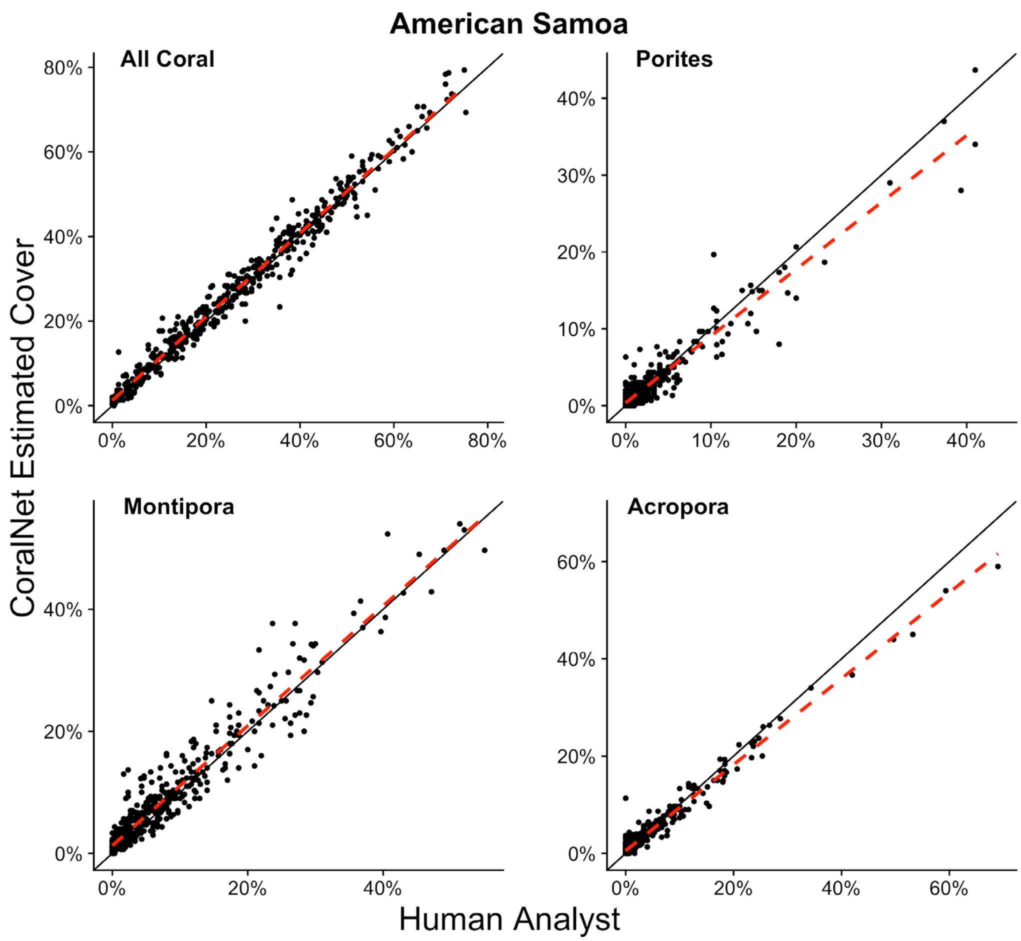

Figure 9. Site-level coral cover computed via manual (human) and automated (CoralNet) analysis for all coral and for common coral genera. Data comes from sites in American Samoa, surveyed by NOAA PIFSC in 2015. The solid black line is the 1:1 line, the dashed reline is a linear fit of the point data.

The VIAME training workshop at the PIFSC spanned 5 days and involved nearly 40 participants across 4 different divisions. A new GPS machine from image processing was set up and potential future collaboration were discussed with the cetacean and seal research groups as well as with the electronic monitoring group.

Following this workshop and training session, PIFSC has begun using VIAME to aid in annotation of modular optical underwater survey system (MOUSS) stereo-camera data from the Bottomfish Fishery-Independent Survey in Hawaii (BFISH). To assist in tuning VIAME detection and classification modules for the Hawaii Deep7 bottomfish complex (six species of deep-water snapper and one deep-water grouper), bounding boxes and track lines have been made for all species using the WAMI-Viewer semi-automated annotation module within VIAME. Annotations with track lines should assist the software to identify fish moving over complex backgrounds where they may be difficult to distinguish from the substrate in still images. These training annotations were used to tune a species-specific VIAME convolutional neural network (CNN), which is currently being tested. Work is also being initiated to develop training data to detect "heads" and "tails" of Deep7 species to aid in automated length measurement.

Northwest Fisheries Science Center

Training on VIAME was conducted at the NOAA Northwest Fisheries Science Center for analysts from the AFSC and the NWFSC on June 12–15, 2018. The workshop was attended by 20 participants representing the NWFSC, AFSC, SWFSC, University of Washington, and Harvey Mudd College. Presentations were made by Kitware and SRI on the first day of the workshop describing the software and its capabilities. On the last 3 days, hands-on training was conducted for a smaller group of about 10 analysts.

Activities during workshop included installation and running of the latest versions of the VIAME CamTrawl stereo fish length estimation as well as investigating possible future uses of VIAME, such as fish detection in untrawlable habitat stereo imagery and video segments, and examination of stereo images from the AUV to test detection of fish, corals and sponges. The focus was on ingesting still image sets, using the iterative query and refinement (IQR) tools and training of SVM models, and examination of IQR/SVM model based detections.

The newly acquired computer platforms were accessed directly by VIAME workshop attendees and the example version of the CamTrawl stereo length estimation available from the VIAME website was run with a new set of CamTrawl images to check for operability. Underwater video footage from the untrawlable habitat strategic initiative (UHSI) Channel Islands project was used as a trial for the iterative query and refinement (IQR) module in VIAME. NWFSC stereo-images collected from a bottom tracking AUV also were used as a trial for the IQR module.

While there were some initial challenges in operating the software, VIAME is now up and running for our analysts. Similar challenges were also encountered when trying to run the SRI Flask tool within CentOS (the preferred system for both the NWFSC and PIFSC). With some re-coding, this issue has also now been addressed.

As mentioned earlier, Ice seal researchers from the AFSC Marine Mammal Lab also attended the VIAME workshop at the NWFSC in order to explore the capabilities of the program and receive training in its operation.

The NWFSC goal is to use VIAME for automated detection of fish, coral and sponges in stereo still images collected by a bottom-tracking AUV. Subsequent to the workshop, more progress was made by individual analysts using the IQR tools and training SVM models. Currently, the NWFSC AUV analysts are using VIAME to develop detectors for pyrosomes in still imagery. This test case was chosen because, while pyrosomes are not a primary focus of our research, pyrosomes have unique physical characteristics well suited to automated detection. The NWFSC is continuing the development of new annotated image training data sets from still imagery and is installing a multi-user graphics processing unit (GPU) workstation, which will allow VIAME usage by a variety of research groups. The NWFSC is also exploring the use of breakaway boxes to enhance GPU capabilities of existing multiple use computers.

Southwest Fisheries Science Center

A workshop on Automated Image Analysis Workshop and VIAME training was convened at NOAA Southwest Fisheries Science Center in San Diego, CA during August 20–24, 2018. The workshop was attended by approximately 20 participants representing: NOAA SWFSC La Jolla—Antarctic Ecosystem Research Division (AERD), Marine Mammal and Turtle Division (MMTD), Fisheries Resources Division (FRD), Information Technology Services (ITS); NOAA SWFSC Santa Cruz—Fisheries Ecology Division (FED) Habitat team, Fisheries Ecology Division (FED) Biophysical Ecology group; NOAA ERD Monterey—Environmental Research Division (ERD); Monterey Bay Aquarium Research Institute (MBARI); NOAA Southwest Region Office; SRI International; University of California San Diego/DropBox; and Kitware, Inc.

The SWFSC workshop included 14 presentations on a wide variety of imaging topics. Some topics were specific to SWFSC work, but most had commonality with work being conducted throughout NOAA.

Eleven analysts attended and participated in the hands-on VIAME training for 2 to 3 days. VIAME was introduced to SWFSC image analysts during the afternoon of the first day, and within 2 days remarkable progress was made by several groups and individuals using VIAME for analysis of their own imagery and starting with no prior experience or understanding of the framework.

Participants of the VIAME training at Southwest Fisheries Science Center used VIAME on their own survey imagery or shared imagery to do the following: ingest still image sets and video imagery; use the iterative query and refinement (IQR) tools and train SVM models; examine IQR/SVM model based detections; filter classifications based on confidence statistics; create manual annotations and tracks of fish and other targets; train deep-learning based object detectors, and apply the detector to underwater and aerial image survey targets (such as fish, sea lions, and penguins) (Figure 10). Others applied default object detectors to image sequences from a lander camera system, and modified VIAME configurations to operate on computer with a sub-optimal GPU.

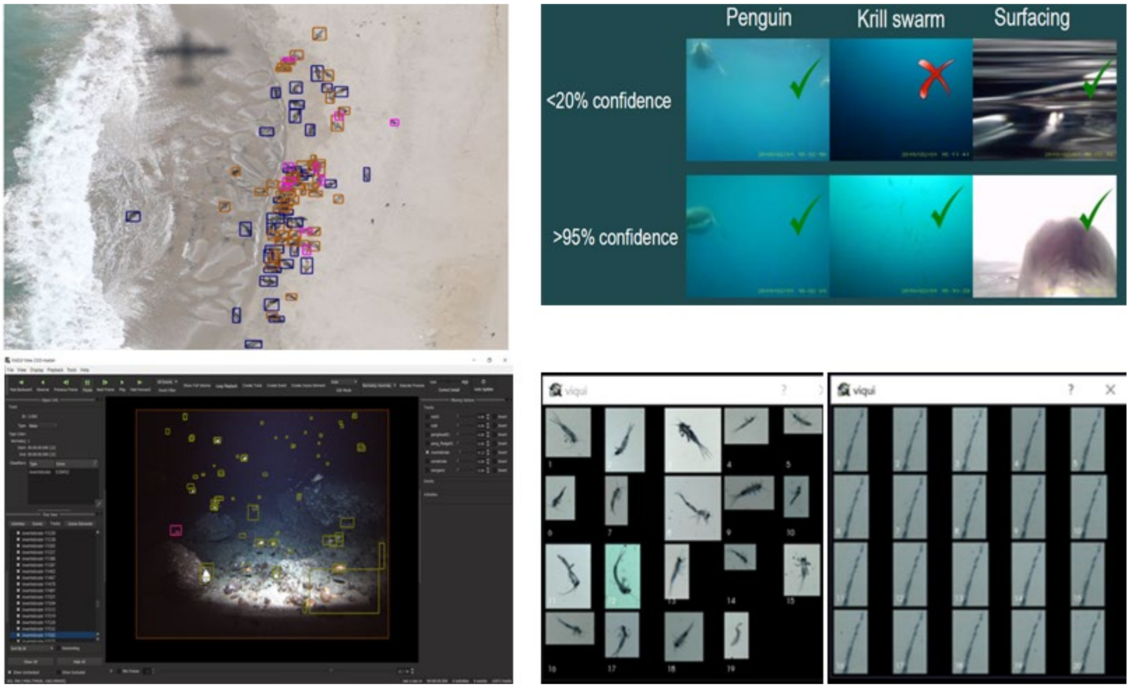

Figure 10. Examples of VIAME processes and results produced by analysts during the workshop at SWFSC.

The ability of SWFSC analysts to rapidly learn and apply the advanced techniques available in VIAME to their wide variety of imagery was possible because of the responsive development and attention by Kitware to the needs and requests from analysts who participated in previous training sessions.

VIAME has been adopted by a group at SWFSC doing censuses of pinnipeds from aerial imagery, and others are planning to use it for their analyses after realizing the potential during training. Currently there are is one desktop computer workstation at SWFSC that was procured using SI support and is suitably equipped for running VIAME's processes that require lots of computing resources and a GPU. If funds were available, SWFSC ITS would provide an enterprise computing system allowing multiple user access to VIAME and equipped with multiple GPUs to enable fast and potentially simultaneous model training.

Southeast Fisheries Science Center

The NOAA Southeast Fisheries Science Center (SEFSC) initially tested previous versions of VIAME using the incorporated default detector that had been trained using HabCam and CamTrawl images. This detector had very limited success in application to SEFSC data. Since that time, significant development, refinement, and additional features included in VIAME have produced much more promising results.

A workshop was conducted at the SEFSC Pascagoula, MS Laboratory on August 27–30, 2018. It was attended by 16 participants from SEFSC's Stennis Space Center, MS, Pascagoula, MS, Panama City, FL, and Beaufort, NC Laboratories, the NOAA Fisheries Office of Science and Technology, SRI International, and Kitware. Kitware provided an overview of VIAME (Figure

11) and SRI provided an overview of FLASK. Each of the attending laboratories presented an overview of their current methodology for acquisition and analysis of images and video that focused on reef fish surveys in the Gulf of Mexico (Pascagoula and Panama City Labs), Atlantic (Beaufort Lab), and Caribbean Sea (Pascagoula Lab). The remainder of the workshop consisted of tutorials for utilizing VIAME and hands-on training by attendees working with their own data.

Attendees acquired a good understanding of VIAME's capabilities and a comparative overview of the multiple analysis pathways that can be used. They tested VIAME's currently incorporated default detectors, facilities for manual annotation of data, Rapid Model Development using Iterative Query Refinement (IQR), and training and utilization of Deep-learning algorithms, and received training on facilities for performance evaluation and scoring of results.

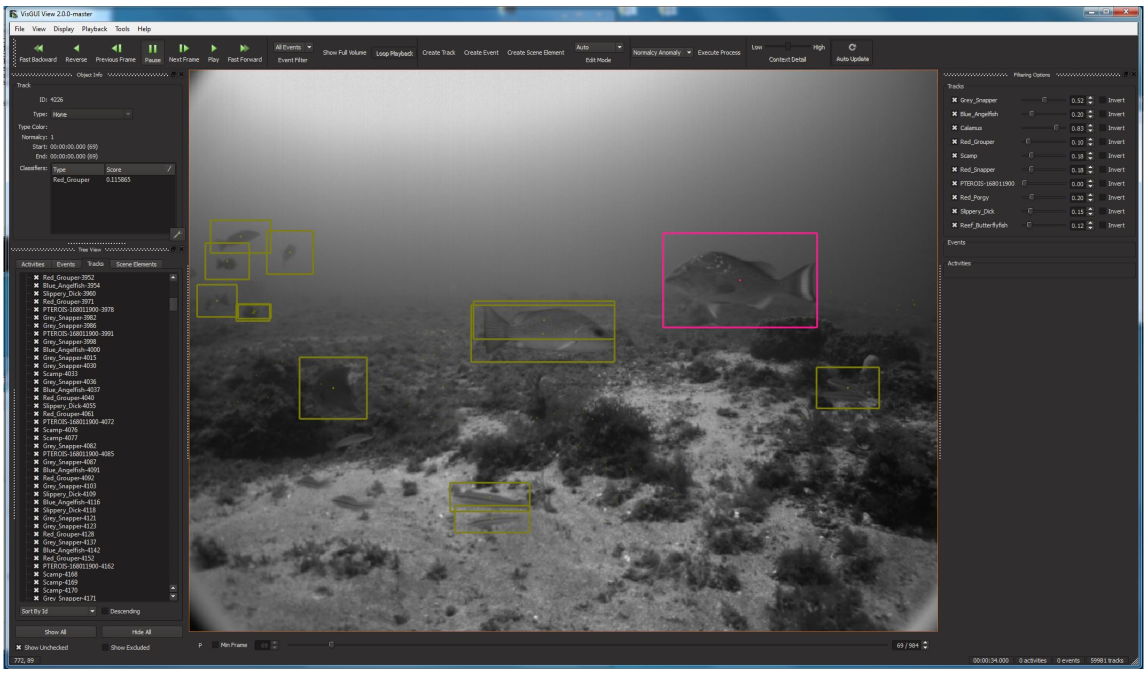

Figure 11. Example of analysis using VIAME during the SEFSC workshop. This image shows detections and identifications of multiple species in one of a sequence of images using a detector trained using VIAME’s deep learning algorithms. Not all fish in the image sequence were detected and/or correctly identified. However, results are promising, given limited training data and few observations of some species included in the model.

Although most of the workshop attendees were focused on camera surveys for reef fish, several additional applications were explored to better understand VIAME’s capability for full-frame video analysis, including automated detection and counting of turtles using trawl-mounted cameras. Potential application of VIAME for plankton image analysis is also being considered.

Recognizing the deep-learning model capabilities of VIAME as the most likely methodology to perform well for detecting and classifying the wide range of species encountered during reef fish surveys in the southeast, SEFSC is proceeding with creation of annotated image training data sets that will include a large number of ground-truth identifications of many reef fish species in diverse habitats and lighting conditions. As these image sets are produced, they will be used to train and test deep-learning models using VIAME. The creation of this training data would likely not have been possible without the IQR and other semi-automated annotation tools provided within VIAME. Where possible, SEFSC also plans to work collaboratively with other groups collecting similar data, including the Florida Fish and Wildlife Conservation Commission, to expand these training data sets and utilization of VIAME. SEFSC also plans to annotate a large number of nose and tail locations of fish in images to aid automation of length measurements from stereo cameras. Once an improved detector is compiled, we will test the video annotation against manually annotated QA/QC data sets collected in the last 2 years.

Northeast Fisheries Science Center

At the Northeast Fisheries Science Center, VIAME has been employed to analyze HabCam images for sea scallops, skates, and other fish. Currently, two percent of the images are manually annotated for scallops. Although this rate is adequate to obtain precise estimates of scallop abundance, it is insufficient for fish.

We used the northeast skate complex for a pilot study to estimate absolute abundance using HabCam images and VIAME (Figure 12). We plan to expand this study to other fish species such as red and silver hake in the future. Because catch of skates is uncertain, especially on the species level, conventional assessment methods cannot be used to estimate absolute abundance. VIAME was used to develop an automated annotator for skates using the YOLOv2 and to automatically process 4 million images from the 2016 HabCam survey. When properly calibrated, the skate annotator was used to obtain fairly precise (e.g. CV = 20–30%) spatial estimates of absolute skate abundance for 2016. This work was presented at the 2018 AFS annual meeting (Hart et al. 2018).

Automated annotators for other fish as well as scallops are also promising. Automated annotation of scallop imagery may allow for a reduction of the manual annotation rate required to obtain a desired precision (Chang et al. 2016). As manual annotation currently comprises more than 100,000 images, any reduction would result in substantial cost savings and allow for faster production of survey data for management.

Ten participants attended the NEFSC VIAME training workshop: four interested in HabCam applications, two for marine mammals (seals and whales), two for plankton, and two interested in applications related to on-deck electronic monitoring.

VIAME is currently being used to help estimate seal abundance from aerial photos. The rapid model generation module was first used to generate a potential training set, which was checked and cleaned manually, and was then used to train a CNN seal detector.

NEFSC has also been collecting plankton data using a Video Plankton Recorder (VPR). Existing software is able to segment images to extract regions of interest (ROI), but there has been no algorithm to classify the detections. Work was initiated during the workshop to develop such classification software using IQR and convolutional neural networks (CNNs).

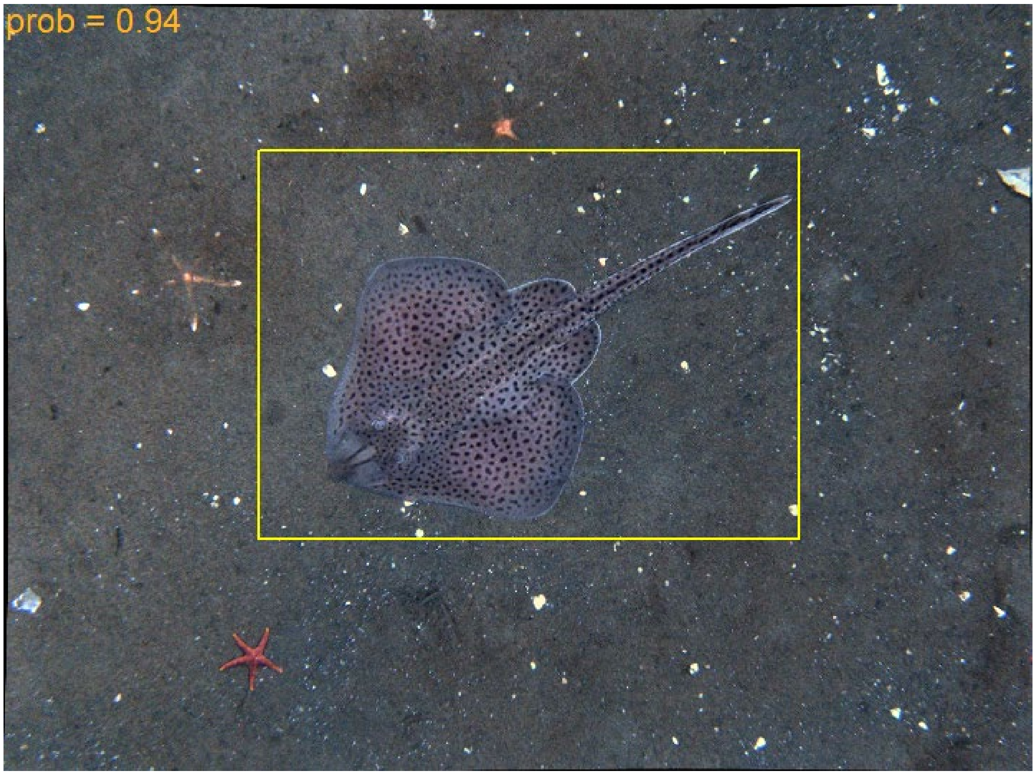

Figure 12. A HabCam image of a skate (*Leucoraja erinacea* or *L. ocellata*) with an automatically generated region of interest (ROI) outlined in yellow.

Untrawlable Habitat Strategic Initiative

While the AIASI was progressing, NOAA Fisheries was also pursuing a SI focused on developing survey methods for untrawlable habitats (Somerton et al. 2017). A central focus of the UHSI was to investigate the reaction of fish to various survey gears. Testing locations were established in the Gulf of Mexico and on the California shelf. At each location, a set of MOUSS camera systems were deployed to the seafloor. After a period of time, a variety of optical survey vehicles (e.g. SeaBed AUV, C-BASS camera sled) were deployed within the field of view of the MOUSS. The MOUSS was able to record footage of resident fish assemblages before, during, and after the passage of these vehicles, allowing researchers to investigate behavioral reactions (Somerton et al. 2017).

One of the main limitations associated with the UHSI survey has been the bottleneck associated with manual annotation of the collected video. During the 2014 UHSI survey approximately 20 terabytes (TB) of image data were acquired. Despite having metadata showing the approximate time of vehicle passage in front of the MOUSS array, significant manual effort was required to target and extract images immediately before, during, and after vehicle passage. To mitigate future manual search needs, a collaboration was established between the UHSI and AIASI to employ automated methods to help identify target image sequences (Girdhar et al. 2015).

Image annotation remains a bottleneck and, to date, only a small subset of the available UHSI imagery has been evaluated. Automated analysis tools, such as those described here, will allow for a far more intensive approach to understanding how fish acclimated to the presence of the sampling gear and to evaluate abundance and behavior trends over much longer time periods than are traditionally sampled (30 min vs 10 hr). Frame-by-frame data annotation could potentially allow for more precise understanding of exact stimuli causing fish response (e.g. vehicle noise vs visual sighting).

Workshops and Data Challenge

The Automated Image Analysis Strategic Initiative also supported four annual workshops (2015–2018) on Automated Analysis of Video Data for Wildlife Surveillance (NOAA AIASI 2017), hosted in conjunction with the Institute of Electrical and Electronics Engineers (IEEE) Winter Conference on Applications of Computer Vision (WACV) and the American Geophysical Union (AGU), Association for the Sciences of Limnology and Oceanography (ASLO), and The Oceanography Society (TOS) Ocean Sciences Meeting. Workshop attendance grew from 25 to more than 75 participants from domestic and foreign private industry, academia, as well as local, state, and national government agencies.

The UHSI and AIASI jointly hosted a session at the 2017 American Fisheries Society Annual Meeting in Tampa Bay, FL (Callouet et al. 2017). Presentations on automated image analysis were delivered by two different groups including SRI International and C-Vision Incorporated, both of which are using deep learning algorithms to identify fish in imagery. SRI the aforementioned FLASK tool that uses a clustering approach to classifying ‘like’ images, which can later be manually validated to improve classification incrementally. FLASK was used to identify habitat and fish observed using the C-BASS towed vehicle (Lembke et al. 2017), but the software was not available during the UHSI experiments for use. C-Vision demonstrated a separate deep learning tool used to automatically identify fish entering into trawls in the northeast United States. While C-Vision was not funded by NOAA Fisheries, collaborations have been initiated based on communication at this joint session.

Presentations were also given by groups integrating optics and acoustics to collect fish diversity, abundance, and biomass data. Coupled acoustic and optic approaches can provide a much broader picture of the sampled habitat but are currently limited primarily due to the manual approaches to annotating optical data. During discussion, automated methods to identify both fish and their habitat were highlighted as a strong needs, given the data volume that can be collected using integrated approaches.

The AIASI also supported a fifth workshop on Automated Analysis of Marine Video for Environmental Monitoring at the Institute of Electrical and Electronics Engineers (IEEE) Conference on Computer Vision and Pattern Recognition (CVPR) in June 2018 (“CVPR 2018 Workshop | viametoolkit.org,” n.d.). CVPR is the premier conference in computer vision, attracting more than 6000 attendees in 2018, and the workshop there reached a broader audience than at WACV.

In conjunction with the CVPR 2018 workshop, the AIASI sponsored a public data challenge on NOAA data with annotations provided by NOAA and Kitware (“CVPR 2018 Workshop Data Challenge | viametoolkit.org,” n.d.). The challenge problem is to automatically detect and classify fish and scallops in thousands of images into 10 or more classes. The challenge was announced just before the workshop, and will remain open indefinitely pending continuing support from NOAA.

Through these workshops and the data challenge, new partnerships were developed between marine and computer visions researchers; collaborations that have and will continue to bear fruit as this domain continues to grow.

Discussion

The AIASI (2013–2018) occurred during a very dynamic time in the evolution of computer vision. During this period, machine-learning tools evolved from hand-crafted, gradient-based methods into biologically-inspired, neural network-based systems. These new methods, although foreign to many analysts, outperform most previous methods for classification, localization, and detection (Krizhevsky et al. 2012).

The SI model provided a novel mechanism to integrate the needs of NOAA Fisheries with a diverse range of expertise spanning academic, non-profit, and private domains. This model, with dedicated and consistent personnel and funding, allowed a diverse group of experts to work consistently, collaboratively, and iteratively on a significant challenge for several years. The AIASI identified key optical data sets that, if they could be analyzed, could provide great benefit to national stock assessments. The AIASI also identified the breakthroughs in vision technologies, and reoriented its efforts to exploit deep learning, interactive model training, and model adaptation rather than further integration of heritage methods.

VIAME: Video and Image Analytics for a Marine Environment

Much of the existing code for automated image analysis and video analytics in the maritime domain was unique to a specific sensor, data type, or research question. The development of the VIAME deep learning architecture has led to more complete and versatile algorithmic pipelines, capable of taking the novice image analyst from imagery to data with minimal effort. The open-source, modular nature of the VIAME system facilitates the continued development of a versatile and dynamic platform capable of addressing current and future needs in automated image processing.

In addition to the ability to run and compare several state-of-the-art algorithms within operational pipelines, the VIAME platform contains multiple features, which aid in the rapid integration of new algorithms. Future work will involve the addition of new algorithm types (such as habitat classification and additional object trackers), the integration of new algorithms (e.g. detectors, trackers, classifiers), adding new GUIs to the system, and additional general system improvements. The ability to configure and change algorithm pipelines in a GUI will be a useful addition as well as a useful debugging tool. All of these capabilities would be highly valuable across a wide range of government agencies including the Defense Advanced Research Projects Agency (DARPA), Intelligence Advanced Research Projects Activity (IARPA), Air Force, Army, Navy and Department of Energy through the common KWIVER platform that underlies VIAME. Kitware develops KWIVER and builds on it for research and applications for all of these agencies.

CoralNet

The close match between manual and automated estimates of coral cover (Figure 9) pooled to the scale of island and year demonstrates the capability of CoralNet in generating data suitable for assessing spatial patterns and temporal trends. As image acquisition is relatively straightforward, the capacity for fully-automated tools to ameliorate the need for resource intensive human analysis opens possibilities for enormous increases in the quantity and consistency of coral reef benthic data available to researchers and managers.

VIAME and CoralNet exceed expectations and continue to grow with increased utility spanning a broad range of programs. Although VIAME is primarily used for underwater imagery at present, it is based on a generic pipelined vision processing system that applies to video analytics in any domain. It is hoped that the general vision community will find VIAME useful and that a vibrant open-source community will develop around the platform.

Discussions among AIASI members and representatives from other NOAA and NOAA Fisheries offices suggest that VIAME is valuable to many line offices, divisions, and programs. AIASI efforts have provided potential solutions to problems that other groups are just beginning to consider.

Electronic Monitoring

NOAA Fisheries is working with fishermen, Fishery Management Councils, and other partners to integrate technology to improve timeliness, quality, integration, cost effectiveness, and accessibility of fishery-dependent data. Electronic monitoring (EM) has clear potential to meet these challenges by incorporating cameras, gear sensors, and electronic reporting systems into fishing operations. However, as with the other domains discussed, the costs of human video review and video storage present significant barriers to moving EM programs forward. To date, NOAA Fisheries has implemented six programs, including the Atlantic Highly Migratory Species (HMS) pelagic longline fishery (2015), Bering Sea and Aleutian Island (BSAI) Non-Pollock Trawl Catcher/Processor (CP) (2007), Pollock CP (2011), Central Gulf of Alaska Rockfish Trawl CP (2012), BSAI Cod Longline CP (2013), and Small Boat Fixed-Gear (2018). On the West Coast, EM will be implemented for the whiting mid-water trawl and fixed-gear fisheries in 2018, and for the bottom trawl and non-whiting mid-water trawl fisheries in 2019. On the East Coast, the northeast groundfish fishery is targeting full implementation in the year 2020, and the mid-water trawl herring fishery in 2019.

Machine learning applications, based on image-training data sets, global positioning systems (GPS), and sensors, could substantially reduce data collection and processing costs for existing and future EM programs. As this paper describes, technological advancements in image processing and storage suggest that the automated methods described herein hold much promise for the EM community. While the AIASI focused on in-situ imagery, researchers at AFSC are currently building upon AIASI efforts, developing machine vision systems for chute and stereo camera tools that incorporate machine learning to automate image processing. Additionally, the NOAA Ship *Henry Bigelow* has recently added machine-vision camera systems in its laboratory spaces for the sole purpose of collecting images necessary to support electronic monitoring.

As in other marine science domains, few EM researchers have experience with annotation techniques or developing novel machine-learning algorithms. A great advantage of many of the algorithms described herein is that they can be re-trained using image data sets from disparate fisheries, providing a cost effective transfer of technology. Recently, several Science Centers have submitted proposals to The Fisheries Information Systems Program (FIS) to build EM image libraries. A workshop in November 2018 will discuss regional vs national needs as well as the creation of a national library, potentially using VIAME as the main interface and analysis tool. VIAME provides a framework within which to organize the many competing on-deck fish identification projects while requiring reviewers, whom are usually highly experienced at-sea observers, to learn a single annotation protocol. EM technicians have already found the VIAME

GUI easy to use and operate. VIAME could also provide wide-ranging access to a current suite of EM-specific machine learning code developed by NOAA Fisheries.

Ocean Exploration

Automated detection and classification of different kinds of marine fauna is also an important problem in the Ocean Exploration domain. Underwater robotic vehicles are routinely deployed into the ocean by various oceanographic institutions and groups around the world, and usually include video or still cameras. Today, the only reliable method to detect and identify targets of interest is to have a trained expert manually annotate the imagery. The Monterey Bay Aquarium Research Institute has spent nearly 30 years collecting and carefully curating images, cataloging more than 23,500 hours of underwater video, which has been manually annotated for more than 4,200 categories (including taxonomic species identification of observed animals) resulting in over 5.5 million annotation stored in an searchable database (Barr 2015) and Deep Sea Guide (Dalit 2016).

As an ever-increasing number of platforms are deployed for longer periods, the requirement for expert human annotation—as well as the factors such as fatigue, inconsistency between experts, and training—has become the critical bottleneck in assessing habitats and fisheries as well as in addressing scientific and societal questions of ecology, human impacts, climate change, and environmental stress. VIAME offers a framework through which recent and future advances in computer vision and pattern recognition can be applied to these challenges in ocean data analysis.

Protected Species

Cetacean (whale and dolphin) surveys often include the collection of digital photographs of dorsal fins and flukes, which are used to identify unique individuals. The photo-identification data are used to analyze movements and distribution, population and social structure, as well as to estimate abundance and other demographic parameters. A single image may contain one or more individual animals. For each group encounter, every fluke or dorsal fin is manually sorted and stored within a separate folder representing a unique individual. Each individual is assigned a distinctiveness rating, and each image of the individual is assigned a quality rating (e.g. Urian et al. 2015). Each fin or fluke is then manually matched to an existing catalog of identified individuals. This process can take several days to several months for each encounter, depending on the size of the group photographed. A typical survey can last up to 30 days, with 1–2 encounters per day.

Recently, aerial photography and photogrammetry has been used to remotely and non-invasively investigate the health of cetacean populations and individual animals (Christiansen et al. 2018; Cramer et al. 2008). Aerial photographs are collected by means of an aircraft or unmanned aerial system (UAS). Counts of individuals photographed in a group can improve group size estimates used in abundance estimation, and counts of mother-calf pairs or other age classes can serve as an index of population status. Given a known altitude, measurements of length and width at various points along the body can be made, which provides a quantitative measure of individual body condition. These processes can take several minutes per image—or longer if multiple images need to be mosaicked or considered together to achieve a full view of the group or individual.

VIAME and other tools described in this paper have the potential to improve and automate both of these efforts, drastically reducing time and cost required to process imagery. Currently, image quality hinders data collection as many images are of low contrast or contain excessive glare or reflection. Several algorithms have been developed to automate pre-processing to include color and contrast correction and glare removal (Kay et al. 2009; Hedley et al. 2005). Automated detection, segmentation, and matching of fluke and dorsal fins can be accomplished using FLASK and the IQR and/or deep learning training modules within VIAME. With modification, the length measurement software module within VIAME could automate current photogrammetry procedures and could expand them through automated calculations of additional diagnostic metrics such as body area and curvature.

Image-based surveys are also being used to assess populations of seals and turtles (Harting et al. 2004). Time-lapse cameras have been deployed on remote and often inaccessible beaches in the Northwestern Hawaiian Islands to survey the highly endangered Hawaiian monk seal (*Monachus schauinslandi*), producing more than 20,000 images per year. The goal is to use this imagery to detect and count seals as well as to distinguish individual animals. Each seal is matched to an image database through distinctive natural markings as well as tags and bleach marks. To date, these images have yet to be processed due to a lack of human resources.

UAS aerial surveys are also conducted to estimate population levels as well as to estimate body size and health of individuals through photogrammetry.

As with cetaceans, several tools developed by the AIASI—including VIAME and FLASK—could potentially automate detection and identification of individual animals and could automate the photogrammetric process.

Aerial surveys for ice associated seals in Alaska (bearded seals, *Erignathus barbatus*; ribbon seals, *Histriophoca fasciata*; ringed seals, *Phoca hispida*; and spotted seals, *Phoca largha*) shifted to an image based approach in 2012. Since then, more than 4 million images (20 TB of data) have been collected of the sea ice habitat of the Bering and Chukchi seas. The analysis of this imagery has been aided by the inclusion of thermal imagery to help detect warm bodies on the cold sea ice, but image processing is still cumbersome and time-consuming. Efforts are underway to implement machine learning as an approach to improve both detection and classification of animals on the sea ice and to reduce overall image collection during these survey efforts. VIAME modifications to accommodate and fuse thermal and color imagery will provide an avenue to develop algorithms using existing imagery. An additional effort to integrate machine learning into the image acquisition system will allow on-board testing and real-time processing with the goal of completing surveys with data in hand. This will reduce the demand on the AFSCs data storage infrastructure and support timely analysis for abundance estimation and stock assessment.

Data Accessibility

The emergence of relatively inexpensive, high-quality, optical sampling technologies have resulted in data volumes that overwhelm not only human analysts, but also traditional storage mechanisms. Until recently, the majority of optical data has been stored on individual hard drives associated with an individual researcher or project. As this became untenable, efforts were made to migrate data to local, Science Center-based clusters and—more recently—to the National

Centers for Environmental Information (NCEI), through the NOAA Fisheries Video Data Management and N-Wave Projects (NOAA NOC, n.d.). Currently, across the various Science Centers, NOAA Fisheries holds approximately 800 TB of optical data, with an anticipated annual growth of 250–900 TB. Yet, an enterprise-level operational solution to large-scale optical data storage has yet to be identified.

In this sense, storage and archiving are defined differently. Archiving of imagery is a form a deep storage that is not easily accessible, but which is maintained for an extended period of time, in accordance with the Federal Records Act. However, we should de-emphasize the archiving of data when discussing data accessibility because archives are not typically accessible quickly enough for timely processing or the development of novel analytical tools.

Storage of imagery is temporary and is less restrictive than archiving, but allows for faster and more general access to imagery for research and development of machine learning and other analytical tools. In recent years, the information technology industry has created cost-effective, scalable, cloud-based solutions allowing for storage and retrieval of large volumes of image data with minimal management effort. The NOAA Data Management Integration Team is exploring the use of cloud-based storage through cooperative research and development agreements with Amazon Web Services, Google, IBM, Microsoft and the Open Commons Consortium.

To encourage vision researchers to work in the marine domain, the AIASI sponsored the development of an image recognition challenge ("CVPR 2018 Workshop Data Challenge | viametoolkit.org," n.d.). The challenge includes a sampling of image data from multiple Fisheries Science Centers, with manual annotations of species of interest on all images. Correct and complete annotations are critical for enabling machine learning algorithm training and development, which should be carefully considered when storing and archiving data for research. Any available annotations, image metadata and other collection information should be preserved and completely cross-referenced with the original imagery.

The marine imagery challenge data is approximately 300 gigabytes (GB), which is too large for many researchers to easily and reliably download resulting in less participation and interest. Cloud storage and computation for the challenge would encourage researchers to examine the data and provide solutions to its difficult fine-grained recognition challenges.

As with any nascent research field, continued energetic development efforts are predicated on availability of data. Within the computer vision community, new algorithms are continually developed and tested using commonly accessible image libraries for training, testing, and evaluation (Krizhevsky et al. 2012). The AIASI brought a new domain of imagery to the computer vision world and catalyzed new algorithm development. However, as outlined by Margolis et al. (in prep), continued development is largely predicated on availability and accessibility of marine image data. To this end, imagery should be archived for long-term storage to meet requirements, while also being stored to allow for quick accessibility. A user-friendly, query and map-based interface should be created to promote discovery and access to imagery. To maximize development imagery should be curated in a way so that access for machine learning is prioritized.

Conclusions

The NOAA Fisheries Strategic Initiative on Automated Image Analysis was successful in its mission to develop guidelines, set priorities, and fund projects to develop broad-scale, standardized, and efficient automated tools for the analysis of optical data for use in stock assessment and in its goal to develop an open source software toolkit allowing for the automated analysis of optical data streams to provide fishery-independent abundance estimates for use in stock assessment. VIAME and CoralNet exceed expectations and continue to grow with increased utility spanning a broad range of programs. VIAME became an end-to-end system for analyzing NOAA imagery with state of the art techniques. Continued support will ensure that these tools remain state-of-the-art and applicable to NOAA's needs. The Strategic Initiative model—with specific objectives and leadership support—proved an effective model for conducting large-scale, multi-year, research and development projects to address national priorities.

Recommendations

VIAME, which has been installed at all the Science Centers, has transitioned from research and development into the initial deployment phase. Users are actively engaged with VIAME, and many requests for new features, improvements and bug fixes are being submitted. Computer vision is a rapidly developing field and new algorithms are continually being developed. For example, a new version of the YOLO CNN algorithm (v3) was just released (Redmon and Farhadi 2018). Incorporating effective algorithms, as well as support for rapidly-evolving deep learning capabilities, operating systems, processors, GPUs, and programming languages will be important to ensure that VIAME does not become obsolete. A modest level of funding (on the order of $100,000 per year) should be continued while VIAME is in active use. Such funding would be used to (1) support routine version updates to maintain compatibility with base operating systems, (2) correct software errors identified during the initial use period, and (3) provide overall customer and technical support to users. While NOAA is under no obligation to maintain support and while VIAME may continue to expand under disparate funding sources, such support would ensure that NOAA Fisheries needs continue to be a priority during future software development. While it is expected that NOAA scientists will adapt and expand VIAME to their purposes by implementing their own algorithms, by writing plug-ins, or by modifying the open source codebase, efficiencies are gained through continued support for VIAME's core developers.

Likewise, the CoralNet tool would benefit from a modest continued support package. CoralNet beta runs on Amazon Web Services cloud-based computing structure, with an annual hosting cost of $6k. Basic maintenance and technical support are estimated at $30k per year.

Hence, the total projected cost for ongoing support of AIASI-developed tools is $136k per year.

More significantly, the training sessions, workshops, and NOAA outreach exposed new user groups to the capabilities of VIAME and identified several new and important capability gaps. Addressing these new constituents will greatly enhance the range of marine science problems that VIAME can address, and consequently the impact it could have across NOAA and the broader marine science community. The most important features recommended for future development include:

- Improved animal tracking to enable species ID on tracks, robust MaxN counts, behavior identification on tracks;
- Stereo image processing, including length measurement as well as dense 3D-image depth estimation to improve animal detection, species ID, habitat classification, and habitat segmentation;
- Stereo video processing including depth-informed tracking, 3D scene reconstruction and stereo video display;
- Individual animal recognition and re-identification to track individuals within single collections (to prevent double counting) and across collections for migratory analysis;
- Behavior and event detection and classification to determine platform avoidance or attraction, predator-prey interactions, feeding and other actions of interest;

- Integration and adaptation of new and emerging deep learning capabilities such as new versions of YOLO (detection and classification), RCNN (detection and classification) (Ren et al. 2017), Mask-RCNN (semantic segmentation) (He et al. 2017), RC3D (event/behavior detection) (Ji et al. 2013);
- Improvement of detection and classification through the fusion of different camera and sensor modalities such as red, green, and blue color channels (RGB)infrared (IR), Light Detection and Ranging (LIDAR), Sound Detection and Ranging (SONAR), and others;
- Anomaly detection to identify unusual objects, animals, flora, habitats and behaviors;
- Improved user interfaces to facilitate annotation and interactive construction of new detectors; new user interfaces for added capabilities such as event detection;
- Incorporation of geospatial data associated with imagery and detections into GUIs and associated functionality;
- Cloud-enabled processing to allow VIAME to ingest and analyze data sets hosted in cloud services; and
- Integration with existing NOAA Fisheries Science Center systems, workflows and databases.

The VIAME user base continues to grow and, with it, we have seen a large increase in requests for new features, bug fixes, and general communication. We suggest a system through which users are able to enter requests for new features and bug fixes, such that they can be prioritized and tracked. Furthermore, we suggest the creation of a NOAA-hosted VIAME users group and mailing list, to enable users to communicate freely and easily. This users group would meet periodically, either physically or virtually, to facilitate and coordinate future development in automated image processing and to develop partnerships with future groups interested in automated image processing. We envision this type of future development funded via direct contract or grant from the requesting entity to developers.

Future Strategic Initiatives

As the NOAA Fisheries Office of Science and Technology considers new strategic initiatives, several considerations are worth noting. Membership from each of the Fisheries Science Centers ensured that results were widely applicable. The inclusion of representatives from academia and private industry ensured that our efforts remained on the cutting edge of technological development. Working with contractors who had already developed software code under other government research projects has allowed us to leverage additional funding and resources that would not otherwise have been available.

The initiative also benefited from clear and specific goals with respect to mission, goal, and deliverables. This clear direction, established early in the strategic initiative processes, allowed tasks to be defined and clear statements of work to be written. Each Science Center representative was able to bring to the table specific data sets and workflows in need of automation. This helped to define scope, priorities, and work plans.

Software development was completed through various funding mechanisms including contracts and grants, which supported staff at NOAA's cooperative institutes (CI) as well as developers at universities and private-sector companies. Providing funding through the Cis was straightforward, allowed support of academic affiliates, and improved university collaborations. Providing funding opportunities to private sector entities was more challenging. Restrictions also limited the initiative's ability to benefit from significant developments within the international community.

Despite these administrative challenges, the strategic initiative proved to be an effective and efficient model, allowing NOAA Fisheries to focus on large-scale goals with multi-year continuity across multiple Science Centers.

Acknowledgements

This project was funded by the NOAA Fisheries Office of Science and Technology. The authors greatly benefited from the efforts of the NOAA Fisheries Science Board, the National Research Council, and many known and unknown administrative staff.

The findings and conclusions in the paper are those of the authors and do not necessarily represent the views of NOAA Fisheries. The use of trade, firm, or corporation names in this publication is for the convenience of the reader and does not constitute an official endorsement or approval of any product or service to the exclusion of others that may be suitable.

Literature Cited

Barr N. 2015. Video Annotation and Reference System [WWW Document]. MBARI. URL https://www.mbari.org/products/research-software/video-annotation-and-reference-system-vars/ (accessed 9.28.18).

Dalit M. 2016. Deep-Sea Guide [WWW Document]. MBARI. URL https://www.mbari.org/products/data-repository/deep-sea-guide/ (accessed 9.28.18).

Beijbom O, Edmunds PJ, Roelfsema C, Smith J, Kline DI, Neal BP, … Kriegman D. 2015. Towards Automated Annotation of Benthic Survey Images: Variability of Human Experts and Operational Modes of Automation. PLOS ONE, 10(7), e0130312. https://doi.org/10.1371/journal.pone.0130312

Burke L. 2011. Reefs at risk revisited. Washington, DC: World Resources Institute.

Callouet R, Campbell MD, Michaels W, Murawski SA, Switzer TS. 2017. Integration of Technologies for Next Generation Marine Observation Systems and Fisheries Independent Surveys I - Symposium. Retrieved September 28, 2018, from https://afs.confex.com/afs/2017/meetingapp.cgi

Chang JH, Hart DR, Shank BV, Gallager SM, Honig P, York AD. 2016. Combining imperfect automated annotations of underwater images with human annotations to obtain precise and unbiased population estimates. Meth Oceanog 17: 169-186.

Christiansen F, Vivier F, Charlton C, Ward R, Amerson A, Burnell S, Bejder L. 2018. Maternal body size and condition determine calf growth rates in southern right whales. Marine Ecology Progress Series 592, 267–281. https://doi.org/10.3354/meps12522

Chuang M-C, Hwang JN, Williams K. 2014a. Supervised and Unsupervised Feature Extraction Methods for Underwater Fish Species Recognition. Computer Vision for Analysis of Underwater Imagery (CVAUI), 2014 ICPR Workshop on, pp. 33-40.

Chuang M-C, Hwang J-N, Kua F-F, Shan M-K, Williams K. 2014b. Recognizing Live Fish Species by Hierarchical Partial Classification Based on the Exponential Benefit. Presented at the IEEE International Conference on Image Processing.

Chuang M-C, Hwang J-N, Williams K, Towler R. 2011. Automatic Fish Segmentation Via Double Local Thresholding for Trawl-Based Underwater Camera Systems. Presented at the International Conference on Image Processing.

Chuang M-C, Hwang J-N, Williams K, Towler R. 2013. Multiple Fish Tracking via Viterbi Data Association for Low-Frame-Rate Underwater Camera Systems. Presented at the IEEE International Symposium on Circuits and Systems.

Cramer KL, Perryman WL, Gerrodette T. 2008. Declines in reproductive output in two dolphin populations depleted by the yellowfin tuna purse-seine fishery. Marine Ecology Progress Series, 369, 273–285. https://doi.org/10.3354/meps07606

CVPR 2018 Workshop | viametoolkit.org [WWW Document], n.d. URL http://www.viametoolkit.org/cvpr-2018-workshop/ (accessed 9.28.18).

CVPR 2018 Workshop Data Challenge | viametoolkit.org [WWW Document], n.d. URL http://www.viametoolkit.org/cvpr-2018-workshop-data-challenge/ (accessed 9.28.18).

Dawkins M, Sherrill L, Fieldhouse K, Hoogs A, Richards BL, Zhang D, Prasad L, Williams K, Lauffenburger N, Wang G. 2017. An Open-Source Platform for Underwater Image and Video Analytics. In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision. Lake Placid, New York, USA.

De'ath G, Fabricius KE, Sweatman H, and Puotinen M. 2012. The 27-year decline of coral cover on the Great Barrier Reef and its causes. Proceedings of the National Academy of Sciences, 109(44), 17995–17999. https://doi.org/10.1073/pnas.1208909109

Fieldhouse K, Leotta MJ, Basharat A, Blue R, Stoup D, Atkins C, Sherrill L, Boeckel B, Tunison P, Becker J, et al. 2014. KWIVER: An open source cross-platform video exploitation framework. In: 2014 IEEE Applied Imagery Pattern Recognition Workshop (AIPR). Washington, DC, USA: IEEE. p. 1–4. [accessed 2018 Sep 27]. http://ieeexplore.ieee.org/document/7041910/.

Fisher RB, Chen-Burger Y-H, Giordano D, Hardman L, Lin F-P. 2016. Fish4Knowledge: collecting and analyzing massive coral reef fish video data. Springer.

Girdhar Y, Cho W, Campbell MD, Pineda J, Clark E, Singh H. 2015. Anomaly detection in unstructured environments using Bayesian nonparametric scene modeling. arXiv preprint arXiv:1509.07979

Hart D, Chang J-H, Dawkins M, Hoogs A, Bergman M. 2018. Estimation of skate abundance using a towed camera survey and automated detectors. Presented at the 2018 Annual American Fisheries Society Meeting, Atlantic City, NJ.

Harting A, Baker J, Becker B. 2004. Non-metrical digital photo identification system for the Hawaiian monk seal. Marine Mammal Science, 20(4), 886–895.

He K, Gkioxari G, Dollár P, Girshick R. 2017. Mask R-CNN. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 2980–2988). https://doi.org/10.1109/ICCV.2017.322

Hedley JD, Harborne AR, Mumby PJ. 2005. Technical note: Simple and robust removal of sun glint for mapping shallow-water benthos. International Journal of Remote Sensing, 26(10), 2107–2112. https://doi.org/10.1080/01431160500034086

Heenan A, Gorospe K, Williams I, Levine A, Maurin P, Nadon M, Oliver T, Rooney J, Timmers M, Wongbusarakum S, et al. 2016. Ecosystem monitoring for ecosystem-based management: using a polycentric approach to balance information trade-offs. Journal of Applied Ecology. 53(3):699–704.

Honkalehto T, McCarthy AL, Ressler PH, Williams K, and Jones D. 2012. Results of the acoustic-trawl survey of walleye pollock (Theragra chalcogramma) on the US and Russian Bering sea shelf in June-August 2010 (DY1006).

Huang PX, Boom BJ, Fisher RB. 2012. Underwater live fish recognition using a balance-guaranteed optimized tree. In: Asian Conference on Computer Vision. Springer. p. 422–433.

Hughes TP, Kerry JT, Baird AH, Connolly SR, Dietzel A, Eakin CM, Heron SF, Hoey AS, Hoogenboom MO, Liu G. 2018. Global warming transforms coral reef assemblages. Nature. 556(7702):492.

Ianelli JN, Barbeaux S, Honkalehto T, Kotwicki S, Aydin K, Williamson N. 2009. Assessment of the walleye pollock stock in the Eastern Bering Sea. Stock assessment and fishery evaluation report for the groundfish resources of the Bering Sea/Aleutian Islands regions. North Pac. Fish. Mgmt. Council, Anchorage, AK, section, 1, pp.49-148.

Ji S, Xu W, Yang M, Yu K. 2013. 3D Convolutional Neural Networks for Human Action Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35(1), 221–231. https://doi.org/10.1109/TPAMI.2012.59

Johansson C, van de Leemput I, Depczynski M, Hoey A, Bellwood D. 2013. Key herbivores reveal limited functional redundancy on inshore coral reefs. Coral Reefs, 32(4), 963–972.

Kay S, Hedley J, Lavender S, Kay S, Hedley JD, Lavender S. 2009. Sun Glint Correction of High and Low Spatial Resolution Images of Aquatic Scenes: a Review of Methods for Visible and Near-Infrared Wavelengths. Remote Sensing, 1(4), 697–730. https://doi.org/10.3390/rs1040697

Kitware Inc. 2018. CVPR 2018 Workshop and Challenge: Automated Analysis of Marine Video for Environmental Monitoring. [accessed 2018 Sep 27]. http://www.viametoolkit.org/cvpr-2018-workshop-data-challenge/.

Kohler KE, Gill SM. 2006. Coral Point Count with Excel extensions (CPCe): A Visual Basic program for the determination of coral and substrate coverage using random point count methodology. Computers & Geosciences, 32, 1259–1269. https://doi.org/10.1016/j.cageo.2005.11.009

Krizhevsky A, Sutskever I, Hinton GE. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K.Q. Weinberger (Eds.), Advances in Neural Information Processing Systems 25 (pp. 1097–1105). Curran Associates, Inc. Retrieved from http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf

Langenkämper D, Zurowietz M, Schoening T, Nattkemper TW. 2017. BIIGLE 2.0 - Browsing and Annotating Large Marine Image Collections. Frontiers in Marine Science, 4. https://doi.org/10.3389/fmars.2017.00083

LeCun Y, Bengio Y, Hinton G. 2015. Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539

Lembke C, Grasty S, Silverman A, Broadbent H, Butcher S, Murawski S. 2017. The Camera-Based Assessment Survey System (C-BASS): A towed camera platform for reef fish abundance surveys and benthic habitat characterization in the Gulf of Mexico. Continental Shelf Research, 151, 62–71. https://doi.org/10.1016/j.csr.2017.10.010

Mace PM, Bartoo NW, Hollowed AB, Kleiber P, Methot RD, Murawski SA, Powers, JE, Scott GP. 2001. Marine Fisheries Stock Assessment Improvement Plan. Report of the National Marine Fisheries Service National Task Force for Improving Fish Stock Assessments. NOAA Technical Memorandum, NMFS-F/SPO-56 l. 69p.

Magnuson-Stevens Fishery Conservation and Management Act. 2007. 16 U.S.C. 1801.

Margolis S, Michaels WL, Alger BD, Beaverson C, Campbell MD, Kearns EJ, Malik M, Thompson CH, Richards BL, Wall CC, et al. 2019. Accessibility of Big Data Imagery for Next Generation Machine Learning Applications. Silver Spring, MD Report No.: NMFS-F/SPO-194.

National Marine Fisheries Service. 2017. Fisheries of the United States, 2016. U.S. Department of Commerce, NOAA Current Fishery Statistics No. 2016. Available at: https://www.st.nmfs.noaa.gov/commercial-fisheries/fus/fus16/index

National Research Council. 2014. Robust Methods for the Analysis of Images and Videos for Fisheries Stock Assessment: Summary of a Workshop. National Academies Press, Washington, D.C. https://doi.org/10.17226/18986

NOAA AIASI. 2017. 4th Workshop on Automated Analysis of Video Data for Wildlife Surveillance. Retrieved September 28, 2018, from http://marineresearchpartners.com/avdws2018/Home.html

NOAA NOC. n.d. N-Wave NOC NOAA N-Wave Project. Retrieved September 28, 2018, from https://noc.nwave.noaa.gov/nwave/public.html

Porter J, Kosmynin V, Patterson K, Porter K, Jaap W, Wheaton J, Hackett K, Lybolt M, Tsokos C, Yanev G, et al. 2001. Detection of Coral Reef Change by the Florida Keys Coral Reef Monitoring Project. In: The Everglades, Florida Bay, and Coral Reefs of the Florida Keys. CRC Press. [accessed 2018 Sep 28]. http://www.crcnetbase.com/doi/10.1201/9781420039412-32.

Quinn TJ, Deriso RB. 1999. Quantitative fish dynamics. Oxford University Press, New York.

Redmon J, Farhadi A. 2018. YOLOv3: An Incremental Improvement. ArXiv Preprint, (arXiv:1804.02767), 1–6. URL https://pjreddie.com/media/files/papers/YOLOv3.pdf

Redmon J, Divvala S, Girshick R, Farhadi A. 2016. You Only Look Once: Unified, Real-Time Object Detection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition

(CVPR) (pp. 779–788). Las Vegas, NV, USA: IEEE. https://doi.org/10.1109/CVPR.2016.91

Ren S, He K, Girshick R, Sun J. 2017. Faster R-CNN: towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, (6), 1137–1149.

Szeliski R. 2010. Computer vision: algorithms and applications. Springer Science and Business Media.

Somerton D, Williams K, Campbell MD. 2017. Quantifying the behavior of fish to a towed camera system using stereo optics and target tracking. Fishery Bulletin. 115:343-354

Trygonis V, Sini M. 2012. photoQuad: A dedicated seabed image processing software, and a comparative error analysis of four photoquadrat methods. Journal of Experimental Marine Biology and Ecology, 424–425, 99–108. https://doi.org/10.1016/j.jembe.2012.04.018

Urian K, Read A, Balmer B, Wells RS, Berggren P, Durban J, Eguchi T, Rayment W, Hammond PS. 2015. Recommendations for photo-identification methods used in capture-recapture models with cetaceans. Marine Mammal Science 31, 298–321. https://doi.org/10.1111/mms.12141

VIAME: Video and Image Analytics for Marine Environments - Kitware/VIAME. 2018. C++, Kitware, Inc. Retrieved from https://github.com/Kitware/VIAME (Original work published 2016)

Wang G, Hwang JN, Williams K, Cutter G. 2016a. December. Closed-Loop Tracking-by-Detection for ROV-Based Multiple Fish Tracking. In Computer Vision for Analysis of Underwater Imagery (CVAUI), 2016 ICPR 2nd Workshop on(pp. 7-12). IEEE.

Wang G, Hwang JN, Williams K, Wallace F, Rose CS. 2016b. December. Shrinking encoding with two-level codebook learning for fine-grained fish recognition. In 2016 ICPR 2nd Workshop on Computer Vision for Analysis of Underwater Imagery (CVAUI) (pp. 31-36). IEEE.

Williams K, Lauffenburger N, Chuang M-C, Hwang J-N, Towler R. 2016. Automated measurements of fish within a trawl using stereo images from a Camera-Trawl device (CamTrawl). Methods in Oceanography 17, 138–152. https://doi.org/10.1016/j.mio.2016.09.008

Williams K, Rooper C, Harms J. 2012. Report of the National Marine Fisheries Service Automated Image Processing Workshop. NOAA Tech. Memo. NMFS-F/SPO-121.

Williams K, Towler R, Wilson C. 2010. Cam-trawl: a combination trawl and stereo-camera system. Sea Technology, 51(12), pp.45-50.

Appendix 1—Terms of Reference

Terms of Reference for NOAA Fisheries Strategic Initiative on Automated Image Analysis

The demand to improve stock assessments drives a need for improved data, particularly more precise, accurate, efficient and timely scientific surveys of fish abundance and their associated habitat and ecosystem. Increasingly, NOAA Fisheries and other agencies are employing camera-based surveys to estimate size-structured abundance for key stocks. However, the volume of data produced by camera-based survey platforms quickly exceeds the capabilities of human analysis. Automated video analysis solutions are needed to extract species-specific, size-structured abundance measures from optical data streams. To affect this development, the NOAA Fisheries Office of Science and Technology (OST) has created the Strategic Initiative on Automated Image Analysis (AIASI). These Terms of References (ToRs) have been developed by the AIASI chair, with input from the AIASI members, OST, and the Advanced Sampling Technologies Working Group (ASTWG).

Objective: The mission of the NOAA Fisheries Strategic Initiative on Automated image analysis is to develop guidelines, set priorities, and fund projects to develop broad-scale, standardized, and efficient automated tools for the analysis of optical data for use in stock assessment. The goal of the AIASI is to create an end-to-end open source software toolkit allowing for the automated analysis of optical data streams to provide fishery-independent abundance estimates for use in stock assessment.

Approach: The suggested approach is to convene an international working group composed of agency, academic, and private sector representatives with the following set of tasks:

1) Identify existing technology and software to meet the stated objectives;
2) Identify research projects or beta technologies that can be easily developed, modified, or transitioned to meet the stated objectives;
3) Identify data gaps that impede development of software for automated image analysis;
4) Identify and rank the principal limitations and deficiencies in the area of automated image analysis as it related to NOAA Fisheries stocks;
5) Identify and rank research tracks for the development of automated image analysis solution to meet the stated objectives. Describe promising new technologies to improve awareness in the assessment and survey programs.
6) Fund high-ranking research projects to develop technology to meet the stated objectives;
7) Fund and organize workshops to bring together members of the computer vision, marine science, and stock assessment communities to develop technologies and research tracks to meet the stated objectives.
8) Consolidate research and development products and develop or catalyze development of an end-to-end open source software toolkit (application) allowing for the automated analysis of optical data streams to provide fishery-independent species-specific, size-structured abundance estimates for use in stock assessment

Timing: The AIASI panel shall meet at least twice a year for a three to five year term. At least one of these meeting should be face-to-face. Panel reports should be sent to OST in late February of each year, and these results can be distributed among the Science Board, stock assessment senior advisor (SASA), and ASTWG in March.

Participation: Each regional Science Director shall ensure the participation of an expert in stock assessment, survey, or sampling technologies. Additional panel members will come from academia and private industry. When feasible, the SASA, OST Director and /or national program managers will attend panel meetings to help provide national context.

Product: An end-to-end open source software toolkit allowing for the automated analysis of optical data streams to provide fishery-independent species-specific, size-structured abundance estimates for use in stock assessment.

Usage: The developed software will be used by:

1) NOAA Fisheries Regional Science Centers for the routine analysis of optical data streams to produce species-specific, size-structured abundance estimates for key assessment targets.
2) Regional, State, and Academic Partners the routine analysis of optical data streams to meet regional and local management objectives.
3) Academic and private industry partners as they continue to develop and refine automated solutions for analysis of optical data streams.

Appendix 2—Scope and Objectives

Automated Image Analysis Strategic Initiative
Scope and Objectives

Mission Statement: The NOAA Fisheries Automated Image Analysis Strategic Initiative team will develop guidelines, set priorities, and fund research to develop broad-scale, standardized automated analysis of still and video imagery for use in stock assessment.

Process

- Adopt a bottom-up and top-down approach, considering existing projects that can be scaled to the NOAA Fisheries level as well as new projects to consolidate existing work or create umbrella initiatives.
    - Conduct fact-finding both within and outside NOAA Fisheries to discover what automated image analysis projects currently exist or are in progress.
    - Determine if any of the above are helpful for SI goals
- Some portions of the image analysis workflow lend themselves well to automation. Others do not. The group should identify those portions of the workflow where NOAA Fisheries may make the most progress within the 3-5 year timeframe and budget provided for the Initiative.
- Set process, priorities and goals resulting in recommendations for rapid, automated analysis of NOAA Fisheries' still and video imagery to improve stock assessment, ecosystem based management, and scientific advice. Meetings will likely feature invited presentations outlining the current state of the art in automated image processing (still and video imagery) as it relates to NOAA Fisheries objectives, including:
        - Enumeration (within defined sampling frame)
        - Size determination
        - Species identification
- NOAA Fisheries members from each Science Center should continue the fact-finding exercise to determine what image-based data sets currently exist and what progress has been made in automating the analysis of existing data sets.
- The chair and NOAA Fisheries members should consolidate input from each Science Center on current and future image-based data streams into general data types or classes
    - Examples of image data classes
        - Still imagery of targets on static background (e.g. seals on ice, scallops on seabed)
        - Video imagery of targets on stationary background (e.g. CamTrawl, Mamigo: halibut on conveyor)
        - Video imagery of targets on moving background (e.g. stereo-video of fish from ROV/AUV)
    - each class should be evaluated based on
        - NOAA Fisheries priority

        - probability of success in automation
        - time and cost to completion
- Identify existing internal and external automated image analysis products
    - determine if existing products can be expanded to meet broad NOAA Fisheries needs.
    - develop user interfaces and data ingestion and export portals to transition algorithms and R&D products to full-scale technician-usable software products
- Make recommendations to NOAA Fisheries Office of Science and Technology and Senior Advisor for Stock Assessment regarding the most efficient and practicable pathways for automated analysis of NOAA Fisheries image-based data sets.

Appendix 3—Roadmap

A Roadmap for Development of Software Systems for Automated Analysis of Marine Optical Data

The greatest impediment to producing accurate, precise, and credible stock assessments is the lack of adequate input data (Mace et al. 2001). Increasingly, NOAA Fisheries and other agencies are employing camera-based methods for more precise, accurate, efficient and timely scientific surveys of key stocks and their associated habitat. However, the volume of optical data produced quickly exceeds the capabilities of human analysis. Automated analysis solutions are needed to extract species-specific, size-structured abundance measures from optical data streams in an accurate and timely fashion (Williams et al. 2012).

To affect such development, the NOAA Fisheries Office of Science and Technology (OST) has created a Strategic Initiative on Automated Image Analysis (AIASI) in 2013. The AIASI mission is to develop guidelines, set priorities, and fund projects to develop broad-scale, standardized, and efficient automated tools for the analysis of optical data for use in stock assessment. In addition to supporting the continued development and improvement of existing automated processing algorithms across NOAA Fisheries Science Centers, a primary goal of the AIASI is to make use of existing open-source resources (processing libraries, toolsets, extensible image processing applications, etc.) developed through existing academic partnerships and by commercial interests to create an end-to-end, open source, automated image processing software system (AIPS) allowing the marine researcher to access, evaluate, and employ a variety of automated processing tools for the analysis of optical data streams for use in stock assessment (Figure 3-1).

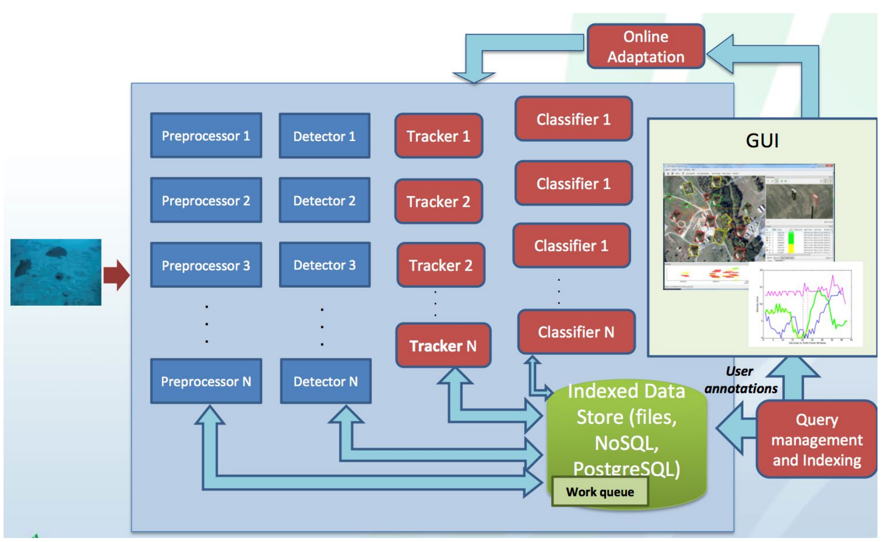

Figure 3-1. Theoretical diagram of a NOAA Fisheries Automated Image Processing Software System (AIPS) for automated analysis of underwater optical data. Optical data is ingested by AIPS (at left) and flows through several possible preprocessing modules before being fed to several possible detection, tracking, and classification modules. A graphical user interface (GUI) supports interaction with the various software modules and data extraction. (Diagram provided courtesy of Kitware, Inc.)

The initial software system will incorporate a variety of tools or software modules (both existing and to be developed) providing operational capabilities to ingest a variety of raw optical data streams, pre-process the optical data, and extract meaningful quantitative and categorical data on targets of interest. Categorical data will include discrete clusters or taxonomic information while quantitative data will include abundance and size. Incorporating disparate processing modules within an overall software framework or system will allow modules to be used synergistically, allowing for more robust processing and additional functionality and will allow existing modules to serve as building blocks for future development.

Open-source software will allow future development and extensibility and incorporation of additional software modules, providing AIPS with additional capabilities. AIPS and all related software modules, user manuals and training materials will be able available for public download from NOAA servers. NOAA Fisheries will also create a web portal for submission of new software modules as well as training/testing data, evaluation metrics, and processing results.

In addition to developing AIPS, the AIASI will continue to collaborate with the research community to support important stand-alone automated processing research and development for key image classes that may not fit within the general AIPS system. The AIASI will also support workshops at key fisheries and computer vision conferences complete with image analysis challenges and training on automated image processing solutions. The latter will serve to catalyze excitement and further development of processing techniques within the fisheries and computer vision communities external to and following the end of the AIASI.

Approach: The approach taken by NOAA Fisheries was to convene an international working group composed of agency, academic, and private sector representatives with the following set of tasks:

1) Identify existing technology and software necessary to meet the stated goal;
2) Identify research projects or beta technologies that can be easily developed, modified, or transitioned to meet the stated goal;
3) Identify data gaps that impede development of software for automated image analysis;
4) Identify and rank the principal limitations and deficiencies in the area of automated image analysis as it related to NOAA Fisheries stocks;
5) Identify and rank research tracks for the development of automated image analysis solution to meet the stated goal;
6) Fund and organize workshops to bring together members of the computer vision, marine science, and stock assessment communities to develop technologies and research tracks to meet the stated goal;
7) Fund projects to develop technology to meet the stated goal. Development will be geared to development of:
    a. An Automated Image Processing Software system (AIPS) comprising:
    b. Sub-modules and stand-alone software products responsible for:
        i. Image Preprocessing
        ii. Target Detection
            1. Unsupervised Clustering
            2. Supervised Clustering
        iii. Visualization

1. Video summarization
2. Unsupervised Clustering
3. Search by example

iv. Target Measurement
v. Target Tracking
vi. Target Enumeration
vii. Target Classification

1. Unsupervised Clustering
2. Taxonomic Identification

Roadmap:

- *Year 1 (2013)*
  - Convene AIASI working group
  - Sponsor 2014 National Academy of Sciences Workshop on Robust Methods for the Analysis of Images and Videos for Fisheries Stock Assessment http://sites.nationalacademies.org/DEPS/BMSA/DEPS_087303

- *Year 2 (2014): $725,000*
  - Fund development of software modules:
    - fish/scallop segmentation
    - fish/scallop detection
    - fish tracking
    - fish classification
    - benthic habitat classification
  - Sponsor Workshop on Computer Vision for Analysis of Underwater Imagery at the International Conference of Pattern Recognition http://cvaui.oceannetworks.ca

- *Year 3 (2015): $650,000*
  - Fund initial development of NOAA Fisheries Automated Image Processing Software System (VIAME)
    - incorporate existing modules and those in development,
    - user interface
  - Fund development of software algorithms/modules:
    - image preprocessing
    - fish/scallop segmentation
    - fish/scallop detection
    - fish tracking
    - fish classification

        - benthic habitat classification
    - Setup web server for download of beta software modules and AIPS beta
        - marineresearchpartners.com
    - Fund database translation/infrastructure to operationalize CoralNet-based automated processing of benthic habitat data for coral reef surveys at the Pacific Islands Fisheries Science Center
    - Leverage existing tools for image annotation
    - Release for evaluation, testing, and operationalization:
        - Automated processing of CamTrawl imagery for Pollock surveys at the Alaska Fisheries Science Center
    - Sponsor Workshop on Automated Analysis of Video Data for Wildlife Surveillance at the IEEE Winter Conference on Applications of Computer Vision 2015 http://marineresearchpartners.com/avdws2015/Home.html

- *Year 4 (2016): ~~$760,000~~ $680,000*

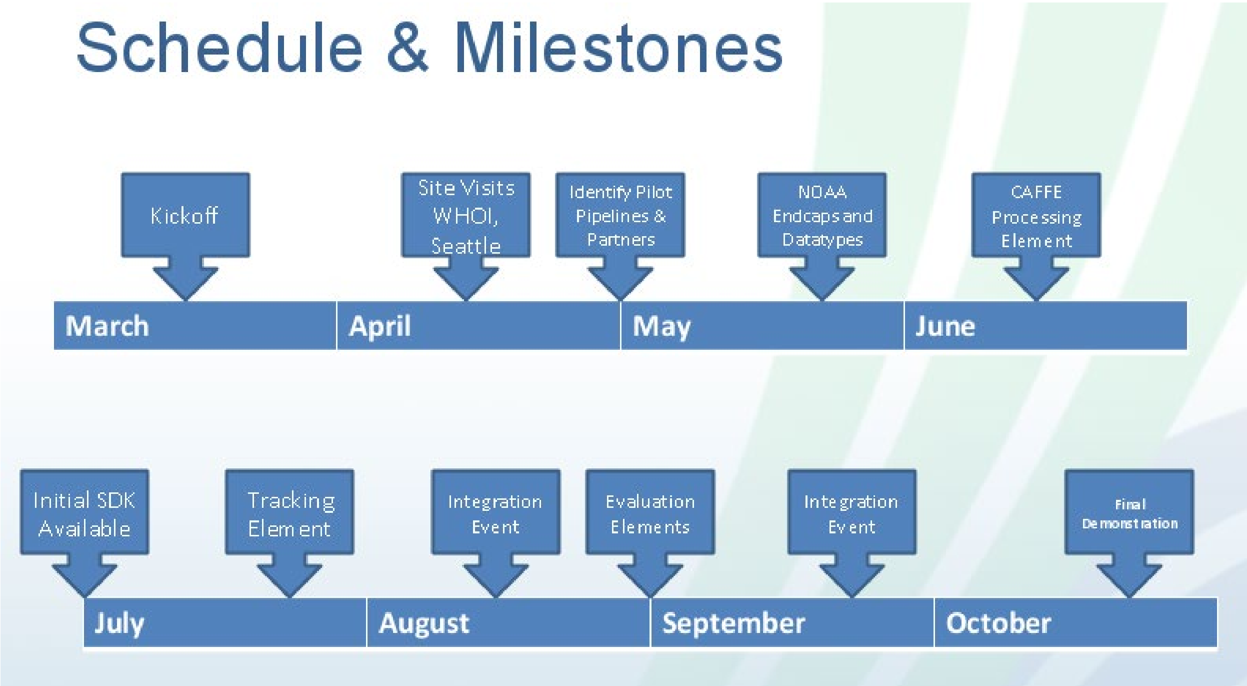

    - Continue development of NOAA Fisheries VIAME
        - incorporate existing alpha-level modules and those in development,
        - Alpha user interface
        - search by example
        - data storage
        - Alpha level stereo image processing

            - Input of paired image streams
            - synching
            - calibration/distortion-correction/rectification
        - Performance metrics and scoring
        - Release for evaluation, testing, and operationalization
    - Support development of software modules or pipelines for:
        - image preprocessing
        - unsupervised clustering / anomaly detection / video summarization
        - fish/scallop segmentation
        - fish/scallop detection
        - fish/scallop enumeration
        - fish/scallop measurement
            - length, area, 3D volume
        - fish tracking
        - fish/scallop classification
        - benthic habitat classification
    - Collaborative development hackathon (August or Sept)
    - Finalize CoralNet-based automated processing of benthic habitat data for coral reef surveys at the Pacific Islands Fisheries Science Center
    - Assemble annotated image datasets for open challenge
        - Marine environment (scallops, fish, seals, whales, dolphins, otters, others?)
        - Several thousand individuals annotated per species for training and testing
    - Sponsor Workshop on Automated Analysis of Video Data for Wildlife Surveillance at the IEEE Winter Conference on Applications of Computer Vision 2016 http://marineresearchpartners.com/avdws2016/Home.html
    - Apply for tutorial and hackathon on VIAME. Working tutorial to train developers.
        - At WACV 2017
        - At ICCV 2017, for broader community.

- *Year 5 (2017): $680,000*
    - Release for evaluation, testing, and operationalization::
        - NOAA Fisheries VIAME
            - User interface
            - image preprocessing
            - unsupervised clustering

            - search by example
            - fish/scallop segmentation
            - fish/scallop detection
            - fish tracking
            - fish classification
            - data extraction
        - Integrate MBARI VARS (Video Annotation and Reference System) http://www.mbari.org/products/research-software/video-annotation-and-reference-system-vars/ and AVED (Automated Visual Event Detection) into the VIAME Open Source Framework for Underwater Image Processing.
        - Automated processing of stereo-video data for reef fish at the Southeast Fisheries Science Center
        - Automated processing of stereo-video data for Deep7 bottomfish at the Pacific Islands Fisheries Science Center
    - Develop VIAME user manual and documentation
    - Conduct training of NOAA Fisheries personnel on use of VIAME and avenues for continued development.
    - Workshops
        - Sponsor Workshop on Automated Analysis of Video Data for Wildlife Surveillance at the IEEE Winter Conference on Applications of Computer Vision 2016[3] Sponsor CVPR VIAME Tutorial and Workshop on Automated Analysis of Video Data for Wildlife Surveillance at the IEEE Winter Conference on Applications of Computer Vision (June 2017)

[3] http://marineresearchpartners.com/avdws2016/Home.html

Appendix 4—List of Acronyms

AERD	Antarctic Ecosystem Research Division
AI	Artificial Intelligence
AIASI	Automated Image Analysis Strategic Initiative
AKFSC	Alaska Fisheries Science Center
AGU	American Geophysical Union
ASLO	American Society of Limnology and Oceanography
ASTWG	Advanced Sampling Technology Working Group
AUV	Autonomous Underwater Vehicle
BFISH	Bottomfish Fishery-Independent Survey in Hawaii
CI	Cooperative Institute
CNN	Convolutional Neural Network
CPCe	Coral Point Count with Excel Extensions
CRADA	Cooperative Research and Development Agreement
CREP	Coral Reef Ecosystem Program
DARPA	Defense Advanced Research Projects Agency
DOE	Department of Energy
ERD	Environmental Research Division
ESD	Ecosystem Science Division
EM	Electronic Monitoring
FED	Fisheries Ecology Division
FIS	Fisheries Information Systems
FLASK	Fish Labeling and Segmentation Toolkit
FRD	Fisheries Resources Division
GB	Gigabyte
GUI	Graphical User Interface
HDF5	Hierarchical Data Format version 5
IARPA	Intelligence Advanced Research Projects Activity
IEEE	Institute of Electrical and Electronics Engineers

AERD	Antarctic Ecosystem Research Division
ITS	Information Technology Services
IQR	Iterative Query and Refinement
JSON	JavaScript Object Notation
MBARI	Monterey Bay Aquarium Research Institute
MMTD	Marine Mammal and Turtle Division
MOUSS	Modular Optical Underwater Survey System
NEFSC	Northeast Fisheries Science Center
NMFS	National Marine Fisheries Service
NOAA	National Oceanic and Atmospheric Administration
NRC	National Research Council
NWFSC	Northwest Fisheries Science Center
OST	Office of Science and Technology
PIFSC	Pacific Islands Fisheries Science Center
RAMP	Reef Assessment and Monitoring Program
ROC	Receiver Operating Characteristic
ROI	Region of Interest
SEFSC	Southeast Fisheries Science Center
SI	Strategic Initiative
SRI	Stanford Research Institute
SVM	Support Vector Machine
SWFSC	Southwest Fisheries Science Center
TB	Terabyte
TOS	The Oceanography Society
UHSI	Untrawlable Habitat Strategic Initiative
VIAME	Video and Image Analytics for a Marine Environment
WACV	Winter Conference on Applications of Computer Vision
YOLO	You Only Look Once