NOAA FISHERIES

Machine Learning to Improve Marine Science for the Sustainability of Living Ocean Resources

Report from the 2019 Norway - U.S. Workshop

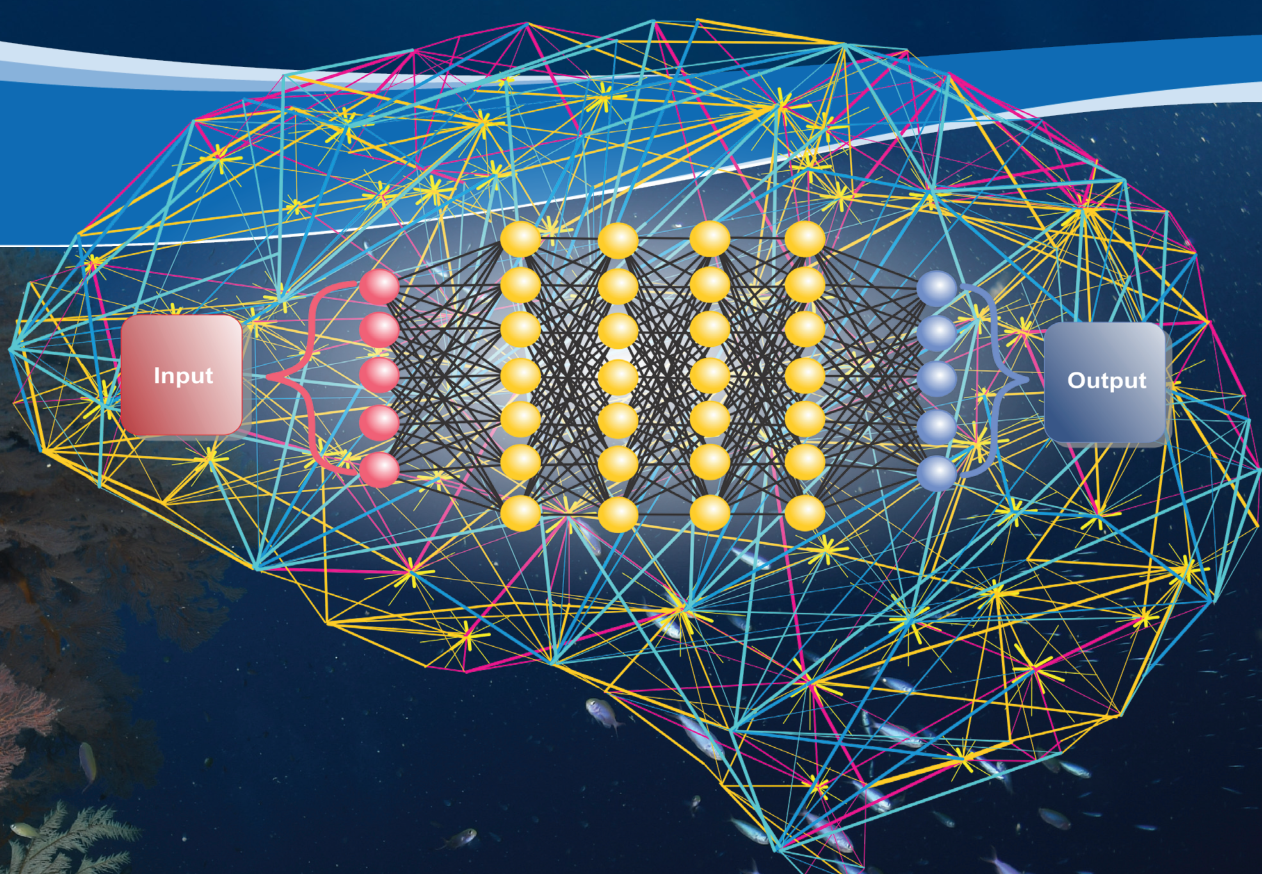

NOAA Technical Memorandum NMFS-F/SPO-199
November 2019

Machine Learning to Improve Marine Science for the Sustainability of Living Ocean Resources

Report from the 2019 Norway - U.S. Workshop

Edited by William L. Michaels, Nils Olav Handegard, Ketil Malde, Hege Hammersland-White

Contributors (*alphabetical order*): Vaneeda Allken, Ann Allen, Olav Brautaset, Jeremy Cook, Courtney S. Couch, Matt Dawkins, Sebastien de Halleux, Roger Fosse, Rafael Garcia, Ricard Prados Gutiérrez, Helge Hammersland, Hege Hammersland-White, William L. Michaels, Kim Halvorsen, Nils Olav Handegard, Deborah R. Hart, Kine Iversen, Andrew Jones, Kristoffer Løvall, Ketil Malde, Endre Moen, Benjamin Richards, Nichole Rossi, Arnt-Børre Salberg, , Annette Fagerhaug Stephansen, Ibrahim Umar, Håvard Vågstøl, Farron Wallace, Benjamin Woodward

November 2019

NOAA Technical Memorandum NMFS-F/SPO-199

U.S. Department Of Commerce

Wilbur L. Ross, Jr.
Secretary of Commerce

National Oceanic and Atmospheric Administration

Neil A. Jacobs, PhD
Acting Under Secretary of Commerce for Oceans and Atmosphere and NOAA Administrator

National Marine Fisheries Service

Christopher W. Oliver
Assistant Administrator for Fisheries

Recommended citation:

Michaels, W. L., N. O. Handegard, K. Malde, and H. Hammersland-White (eds.). 2019. Machine learning to improve marine science for the sustainability of living ocean resources: Report from the 2019 Norway - U.S. Workshop. NOAA Tech. Memo. NMFS-F/SPO-199, 99 p.

Available online at https://spo.nmfs.noaa.gov/tech-memos/

Cover photographs were taken by William L. Michaels.

Contents

ACKNOWLEDGEMENTS

The Norwegian Institute of Marine Research (IMR) and U.S. Department of Commerce's National Ocean and Atmospheric Administration (NOAA) have worked collaboratively for decades in fisheries and marine science, including in the development of ocean technologies. Dr. Geir Huse (IMR Research Director) and Dr. Cisco Werner (NOAA Fisheries' Chief Scientist) are acknowledged for their continued support of the IMR-NOAA partnership to advance science for the sustainability of living ocean resources. We are grateful for the sponsorship and technical contributions to the workshop by Helge Hammersland and Hege Hammersland-White of Scantrol Deep Vision AS. The steering committee of internationally recognized experts (Hege Hammersland-White, Nils Olav Handegard, Ketil Malde, and William Michaels) were instrumental in developing the framework of this workshop, assuring the invited participants provided an inclusive diversity of perspectives on the current state and application of machine learning for marine and fisheries science. Most importantly, the expertise and recommendations of the participants contributed to the overall success of the workshop that we hope will result in long-lasting contributions for the timely delivery of high-quality scientific products for the sustainability and health of our ocean resources.

List of Acronyms

AI: Artificial intelligence

API: Application programming interfaces

COGMAR: Norwegian Computing Center's Cognitive computer vision for marine services

CNNs: Convolutional neural networks

CPU: Central processing units

eDNA: Environmental DNA

EM: Electronic monitoring

GPU: Graphics processing unit

GUI: Graphical user interface

IMR: Norwegian Institute of Marine Research

LSSS: Large Scale Survey System acoustic software

ML: Machine learning

NOAA: National Oceanic and Atmospheric Administration

NMFS: NOAA's National Marine Fisheries Service, also referred to as NOAA Fisheries

VIAME: Video and Image Analytics for Marine Environments software

USV: Unmanned surface vehicle

UxS: Unmanned systems

VCS: Version control system

EXECUTIVE SUMMARY

The Norway - U.S. Workshop entitled “Machine Learning to Improve Marine Science for the Sustainability of Living Ocean Resources” was held in Bergen, Norway on 23-25 April 2019. The goal of the workshop was to exchange information on the current state of development, progress, and applications of computer vision and machine learning (ML) analytics. The three-day workshop was held at the Institute of Marine Research (IMR), with sponsorship by IMR, National Ocean and Atmospheric Administration (NOAA), and Scantrol Deep Vision AS.

The first day provided an overview of ongoing research efforts and progress pertinent to the various applications of machine learning for fisheries and marine science. Day 2 of the workshop provided a more technical focus on the research for applying analytical ML methods and to define priorities for a collaborative roadmap to effectively advance the use of ML. Invited contributions provided diverse overviews and case studies on the application of computer vison and ML classifiers for imagery and acoustic data collected from underwater and aerial surveys, including detection-classification of plankton, fish, and marine mammals. ML applications involved data collections from traditional survey platforms, autonomous platforms, monitoring-classification systems of fisher trawl operations to reduce bycatch, and electronic monitoring of fishing vessel operations and catch. The final phase of the workshop provided hands-on training on GitHub and the newly released open source Video and Image Analytics for Marine Environments (VIAME) toolbox that utilizes computer vision and machine learning analytics.

The goal of the workshop was to bring perspectives together to understand the complexity of implementing ML applications, end-users’ statistical requirements, ethics and confidentiality, and considerations in building partnerships. Collective expertise and inclusive perspectives of workshop participants from governmental, academic, and private sectors presented an opportunity to understand how to strengthen partnerships. There was consensus that ML analytics are readily available, and deep learning has already revolutionized how we will process and analyze scientific data. This report highlights some of the key requirements for transitioning ML into effective human-ML collaborations. Recommendations include improving data accessibility and training datasets to ensure acceptable model performance for the end user, the need to improve the workflow to effectively use open source tools and cloud computing, and the importance of partnerships to build ML capacity. There was also considerable discussion on the importance of building proficiency in ML through recruitment, training, scientific exchange, and collaboration across various sectors. Scientific exchange and collaborations are critical for remaining knowledgeable and ensuring the integrity of scientific products as the tools and resources for ML applications rapidly evolve. These recommendations will help to improve

organizational efficiencies in the application of ML to cost-effectively deliver high quality and timely scientific products for policy decisions on the sustainability of ocean resources.

1 INTRODUCTION

1.1 PROGRESS IN MACHINE LEARNING

For the purpose of this report, ML is defined as the algorithms and computational methods to learn information from data without predetermined equations and explicit instructions, thereby improving model performance for automated detection, classification, and predictions. AI has a broader meaning that includes other capabilities, such as the interaction of ML with sensors, autonomous vehicles, and computer-based reasoning. The concept of ML has been around for decades, and was introduced as a discipline of computer science that enables computers to learn with explicit programming (Samuel 1959).

Although ML does not require an explicit description for solving a problem, it is instead heavily dependent on input data. As ML evolves to address more complex problems, larger datasets can be analyzed, in turn requiring more computing power. The recent advances in computing power and data accessibility are important reasons why the interest and application of ML has dramatically increased in recent years. Additionally, scientists strive to apply ML algorithms to complex problems such as automated detection and classification under a range of environmental conditions, and predictive forecasting through environmental data fusion. This transformation is also taking place in marine science (Malde et al. 2019), placing a priority on improving data accessibility to enable the application of ML for environmental science (Margolis et al. 2019).

While the increase in data accessibility and computing power has empowered the use of automated detection, classifiers, and other cognitive capabilities of ML resulting in significant cost savings in the processing and analyses of data, more data does not necessarily result in better performance of the ML algorithms (Schneier 2016). Throughout this workshop, there was an underlying theme in the ML case study presentations and discussions that the input data should be relevant and compatible to the scope of the model, and of the importance of independent testing of the model performance to assure the integrity, reliability, and credibility of the science.

There are two learning approaches for ML computations discussed with this report. Unsupervised ML uses the input training dataset to train the ML algorithm, but the training data does not contain information about the required output. Unsupervised learning is used for analysis between input variables (e.g., cluster analysis, anomaly detection, and multivariate analysis) to discover hidden patterns or model the distribution of data. Supervised ML uses input training data that are labeled with desired outputs, and this allows more rigorous statistical analysis (e.g., regression analysis, linear discriminant analysis, decision trees, vector machines, k-nearest neighbor, neural networks) to evaluate model performance and bias (Russel and Norvig 2010; Bengio et al. 2013).

In this report, many of the case studies utilized deep learning methods such as artificial neural networks. Deep learning is a subset of ML algorithms with hierarchical learning that has received considerable attention in recent years (Stone 2019). The term “deep” is in reference to multi-layered neural network architecture, typically trained with labeled training datasets (Fig. 1-1).

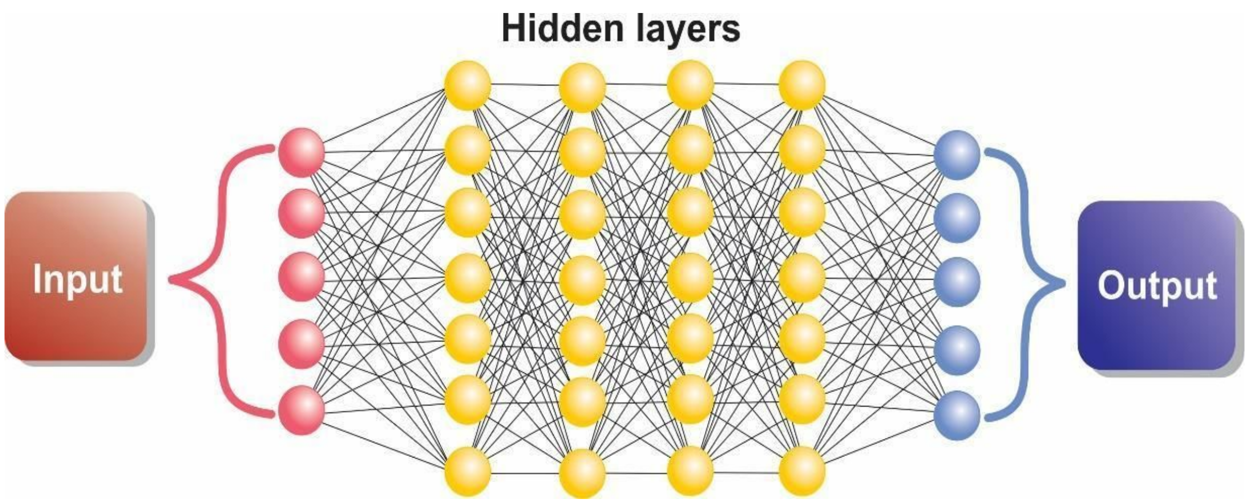

Figure 1-1. Deep learning is in reference to the hidden layers of the neural network’s mathematical computations.

Inspired by the convolutional architecture of the human brain’s visual system, a class of neural networks referred to as convolutional neural networks (CNNs) has been developed. This is where several recent advances in deep learning have occurred, beginning with the five-layer LeNet1 to classify hand-written zip codes (LeCun et al. 1989). However, after the deep convolutional network by Krizhevsky et al. (2012) which made a leap forward in image classification accuracy on the ImageNet Large Recognition Challenge (ILSVRC; Russakovsky et al. 2015), CNNs started to become the standard in image analysis. Today, most image-based deep learning applications use CNNs (Aggarwal 2018; Goodfellow et al. 2016).

There are currently no exact rules on how to set up a deep neural network. If the data is linearly separated, then hidden layers are not required. The hidden computational layers are only required if data must be separated non-linearly (Breimen 2001). One of the challenges in creating neural networks is deciding the number of hidden layers and associated neurons for each layer. The ML deep learning method tends to perform better when well-structured large data sets are available for iterative training. The weakness in deep neural networks is the requirement for the large amount of input data and computational resources for training and the black-box dilemma that obscures the decision process. While neural networks are “black box” systems, where the exact path to the output is difficult to analyze, more traditional ML methods are more transparent.

In addition to neural networks, there are other categories of ML algorithms (e.g., nearest neighbor, naive Bayes, decision trees, linear regression, support vector machines). Some of the case studies in this workshop discussed that decision trees are easier to interpret in regard to feature selection. While decisions trees tend to overfit data (exaggerate minor fluctuations in data), this method can be regularized, for instance by removing unnecessary structure (pruning) after construction. Random forests (Breiman 2001, Kong and Yu 2018) were also discussed during the workshop, which is a decision tree approach to minimize random error or noise of overfitting by randomizing the binary question selection (rather than a greedy approach) by constructing multiple trees (Fig. 1-2) to be used as an ensemble. Further information on ensemble-based modeling is provided by Aggarwal (2018) and Cai et al. (2015).

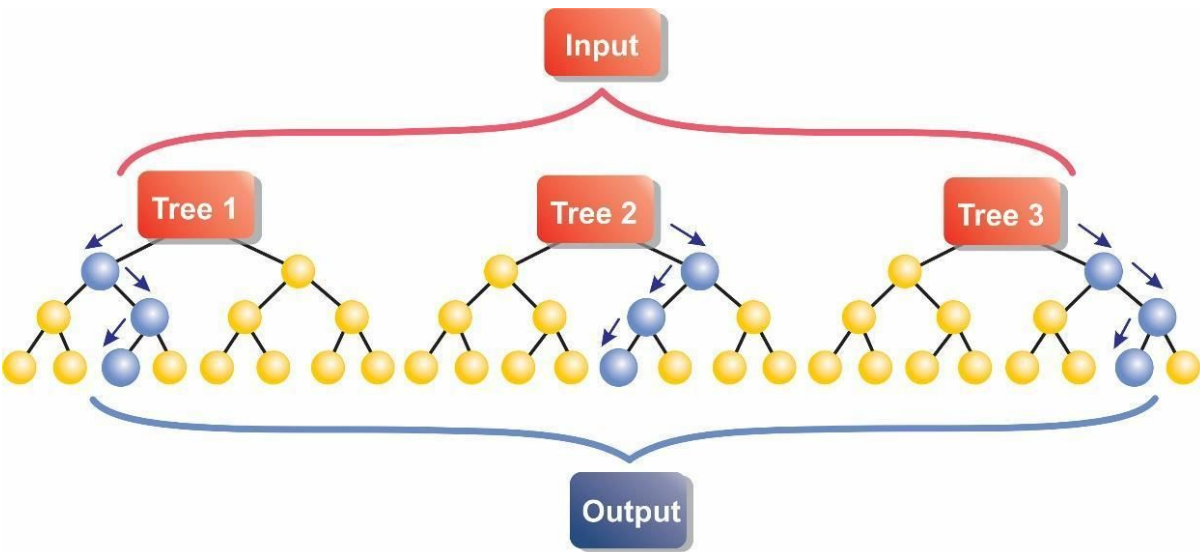

Figure 1-2. Random Forests deep learning method randomizes the binary question at each node resulting in the construction of multiple trees that split at each node until the node is terminated.

While the introduction provides background on the ML analytics discussed in the report, there were also discussions on the importance of open source tools advancing ML research and applications. User-friendly programming languages such as Python and R that are available from free and open source distributors like Anaconda enabled many to build deep learning algorithms. A number of participants also access TensorFlow and Keras open source libraries for ML algorithms that run efficiently on computers using multiple CPUs and GPUs (refer to Section 8.4).

Training was provided for GitHub (see Section 7), including continuous integration, as a code hosting and testing platform that empowers scientists to work collaboratively and on ML specifically. Training was also provided for the Video Image Analytics for Marine Environment (VIAME) software, openly available through GitHub, for a user-friendly end-to-end pipeline for ML which includes access control, feature requests, and task management. Clearly, these open source tools have also significantly contributed to the progress in ML applications.

1.2 Recent Interest in Machine Learning

The recent surge in interest to utilize machine learning can in part be attributed to the increase in availability of data and computing power. In comparison to the general-purpose central processing units (CPUs), graphics processing units (GPUs) provide a highly parallel, high performance computing architecture well suited for deep learning technologies such as CNNs. The recent collaborative efforts of researchers have increased the availability of open source libraries and user-friendly coding which advanced the application of ML by the wider community. An example is the open source VIAME toolkit designed specifically for automated object detection and classification for applications in fisheries and marine science (refer to Sections 2, 3, and 7).

The recent interest in ML became most prevalent when Internet services began utilizing machine learning analytics around 2014. Today, a wide range of Cloud AI platforms have become available with deep learning analytical tools, such as Amazon AWS SageMaker, Google Cloud Machine Learning Engine, IBM Watson Machine Learning, and Microsoft Azure Machine Learning (refer to Section 8.4). The Cloud AI platforms provide the opportunity to build partnerships to transition ML research into operations even as this technology continues to evolve. The scope of this workshop is focused on how to build upon the recent interest, progress, and partnership opportunities for the application of ML in fisheries and marine ecosystem science.

With the recognition that the dramatic increase in data collections can be attributed to the recent emergence of sampling technologies, the case studies addressed the need to more effectively process and analyze data from imaging, acoustic, and environmental sensors. This analysis bottleneck (Malde et al. 2019) is often referred to in conjunction with big data, and results in increasingly massive volume of data that cannot be processed manually or with traditional software. Discussions throughout the workshop included the importance of improving data accessibility and workflows for utilizing ML tools to resolve the dramatic increase in data collections. Many agencies and institutions have recently begun strategic initiatives to transition ML into organizational efficiency by significantly reducing the processing time and costs, as well as improving the predictive capabilities of models. For fisheries and marine science, it is timely that experts were invited to a workshop from governmental, academic, and private sectors to provide perspectives on building partnerships to transition ML into operations.

1.3 Key Challenges in Machine Learning

As data and analytical tools become more commonplace to research and scientific operations in the marine environment, there is an increasing need for fisheries and marine scientists to build collaborative partnerships with experts in machine learning, computer vision and artificial intelligence community. The emergence and routine deployment of ocean technologies, such as remote sensing and unmanned platforms, have resulted in a dramatic increase in data collections that have exceeded our ability to process in a timely manner using conventional manual processing methods. Therefore, the urgency to utilize AI and ML analytics has become a priority to provide more accurate and timely scientific products for ocean policy decisions.

Research and development of AI and ML is evolving rapidly not only due to data accessibility and computing power, but also because of the recent advances in open source tools and cloud platforms. In recent years, consensus has been developing on the use and benefits of ML methods, yet collaborative work is still required to provide guidance on enhancing the data accessibility, workflow, and metrics to fully utilize the analytical tools that are becoming more readily available. The overarching challenge is how to transition these AI and ML advances into enhanced data processing and workflow pipelines, and how to best build collaborative partnerships to optimize AI/ML capacity. The case studies and discussions of the workshop provide insight on the need for collaborative research on ML algorithms, data accessibility and training datasets, validation and model performance, training, and partnerships. The challenges will be discussed further in the Conclusions (Section 8).

1.4 Workshop Objectives

The goal of the workshop was to exchange information on the current state of development, progress, and applications of computer vision and ML analytics. The collective expertise and inclusive perspectives of 33 invited workshop participants (Fig. 1-3; Appendix B) from governmental, academic, and private sectors presented an opportunity to solicit insight on the requirements for ML applications, data accessibility, workflow, validation and statistical accuracy, ethics and confidentiality, training, and partnership considerations. The first day provided an overview of ongoing research efforts and progress pertinent to the various applications of machine learning for fisheries and marine science.

The second day had a more technical focus on the research for applying the analytical methods of ML and to define the priorities for the collaborative roadmap to effectively advance the use of ML. Invited contributions provided diverse overviews and case studies on the application of computer vision and ML classifiers for imagery and acoustic data collected from underwater and aerial surveys, including detection-classification of plankton, fish, and marine mammals. ML applications involved data collections from traditional survey platforms, autonomous platforms, monitoring-classification systems of fisher trawl operations to reduce bycatch, and electronic monitoring of fishing vessel operations and catch.

Figure 1-3. Group photograph of participants at the Norway – U.S. Workshop on Machine Learning to Improve Marine Science for the Sustainability of Living Ocean Resources held at the Institute of Marine Research in Bergen, Norway on 23-25 April 2019.

The workshop also included hands-on training sessions on standard tools widely used in ML and the newly released open source Video and Image Analytics for Marine Environments (VIAME) toolbox that utilizes computer vision and ML algorithms. Refer to Appendix A for the terms of reference and agenda of the workshop.

The workshop concluded with panel discussions to address the challenges associated with the dramatic increase in data from ocean technologies and the need to develop a collective roadmap for the implementation of ML for timely scientific products (Fig. 1-4). The following trigger questions were introduced to begin the discussions:

- What should the data management enterprise consider for increasing the accessibility of enriched annotated metadata for ML analytics?
- How to develop a roadmap for implementing the rapidly developing ML analytics?
- Is it better to create purpose-built AI solutions or modular tailored capabilities that can be plugged into an AI multi-processing pipeline framework?
- What resources are needed to enable the application/deployment of ML analytics?
- How do we build partnerships and workforce competence to implement ML?

- How to accelerate the contributions of the private sector to advance AI capabilities?

Figure 1-4. During the Norway – U.S. Workshop, diverse expertise and perspectives were provided from representatives of governmental, academic, and private sectors working on the application machine learning analytics to improve fisheries and marine science for sustainable living ocean resources.

2 OVERVIEW OF MACHINE LEARNING PROJECTS

Machine learning analytics to enhance the value and timeliness of scientific products from big data collected in the marine environment.

William L. Michaels[1] and Benjamin Richards[2]

[1] NOAA Fisheries, Office of Science and Technology, Silver Spring, MD 02910, USA

[2] NOAA Fisheries, Pacific Islands Fisheries Science Center, Honolulu, HI, USA

Extended Abstract

NOAA Fisheries initiated the Automated Image Analysis Strategic Initiative (Richards et al. 2019) to resolve the increasing cost and backlog of big data imagery collected from NOAA surveys. In partnership with Kitware Computer Vision Inc., the Video and Image Analytics for Marine Environments (VIAME; Dawkins et al. 2017) software was delivered in 2018. VIAME is an end-to-end open-source software package for automated image analysis of marine and fisheries science data that utilizes advanced computer vision and machine learning (ML) analytics for automated object detection, tracking, and classification (Fig. 2-1; see Sections 6 and 7).

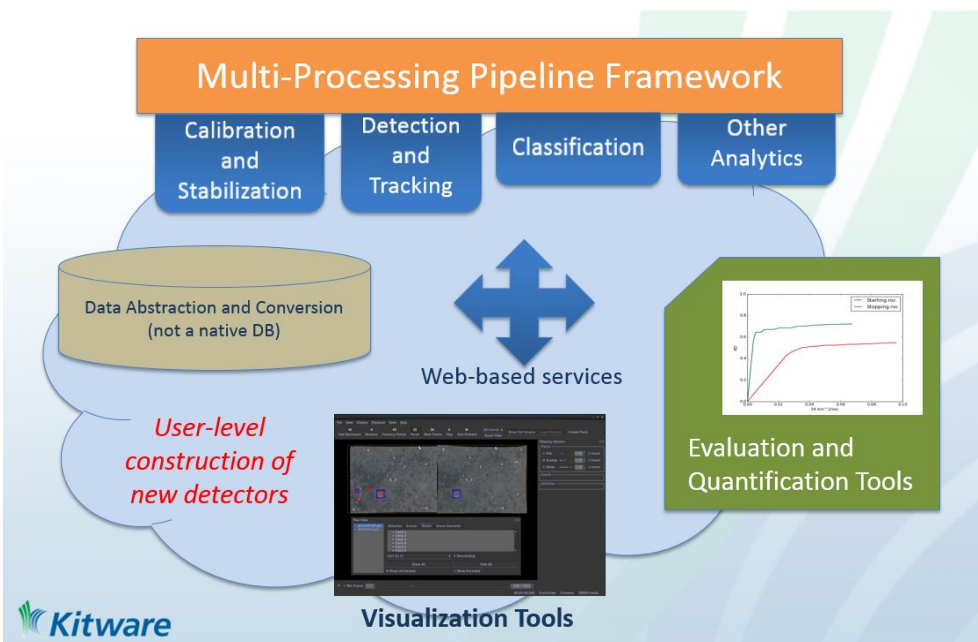

Figure 2-1. The Video and Image Analytics for Marine Environments (VIAME) open source software provides an effective multi-processing workflow for automated object detection and classification for a variety of applications for fisheries and marine science.

The use of VIAME was initially applied to underwater fisheries surveys to improve the quality and timeliness of abundance estimates for stock assessments. Its computer vision and ML capabilities streamlined the processing of still and video imagery data (Fig. 2.2), resulting in 25-75% cost-savings for various programs.

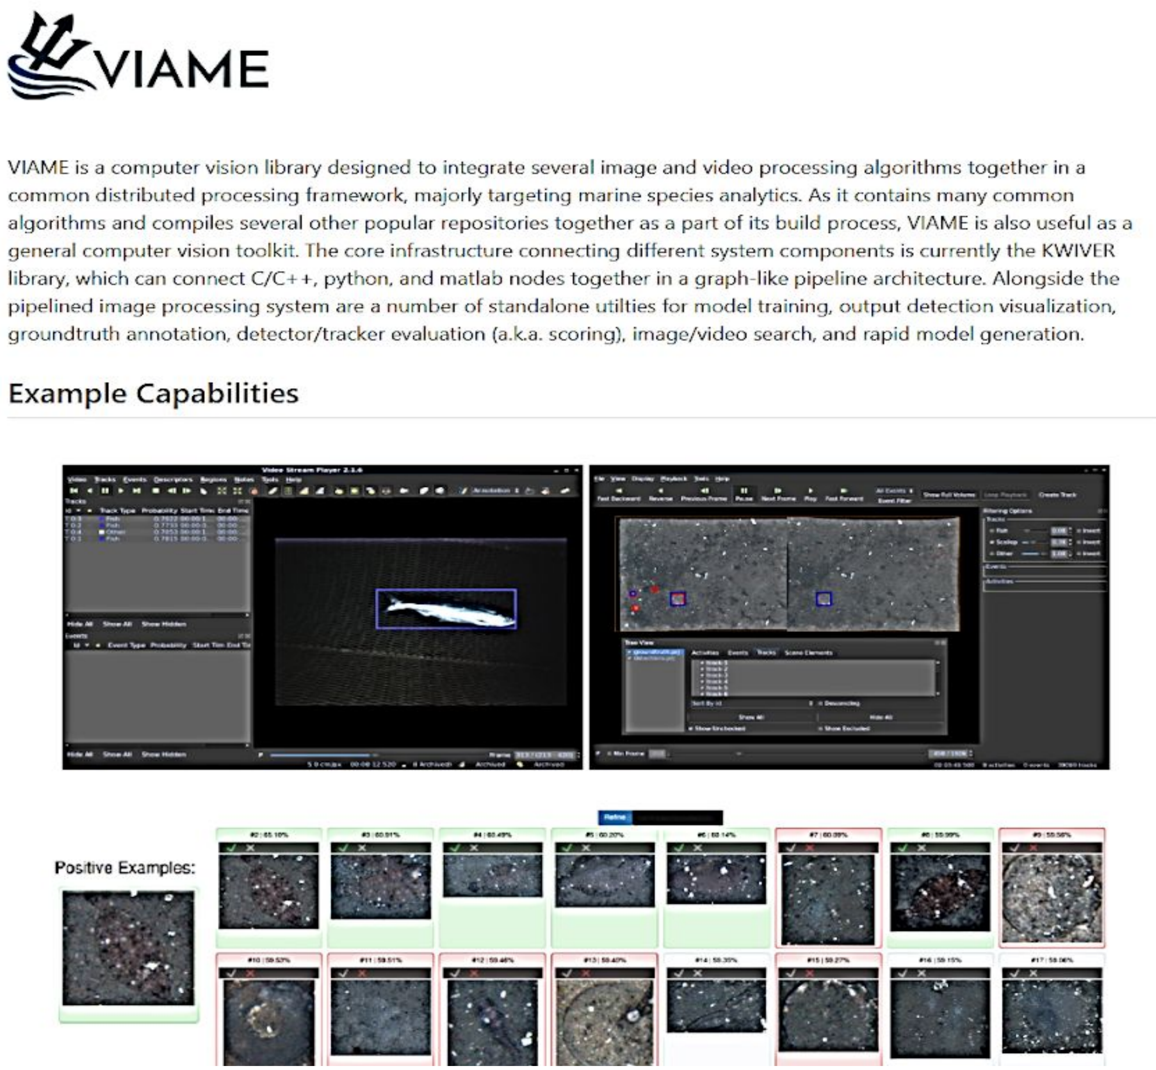

Figure 2-2. VIAME computer vision library and machine learning algorithms are utilized in a graphical pipeline architecture for developing automated object detection and classification models in the marine environment.

During the past year, the applications of VIAME have expanded for rapidly processing data imagery from optical scallop surveys along the U.S. northeast continental shelf (Fig. 2-3; refer to Section 4), Bering Sea pollock survey (Fig. 2-4), to the Gulf of Mexico reef fish survey (Fig. 2-5) and U.S. west coast aerial seal surveys (Fig. 2-6).

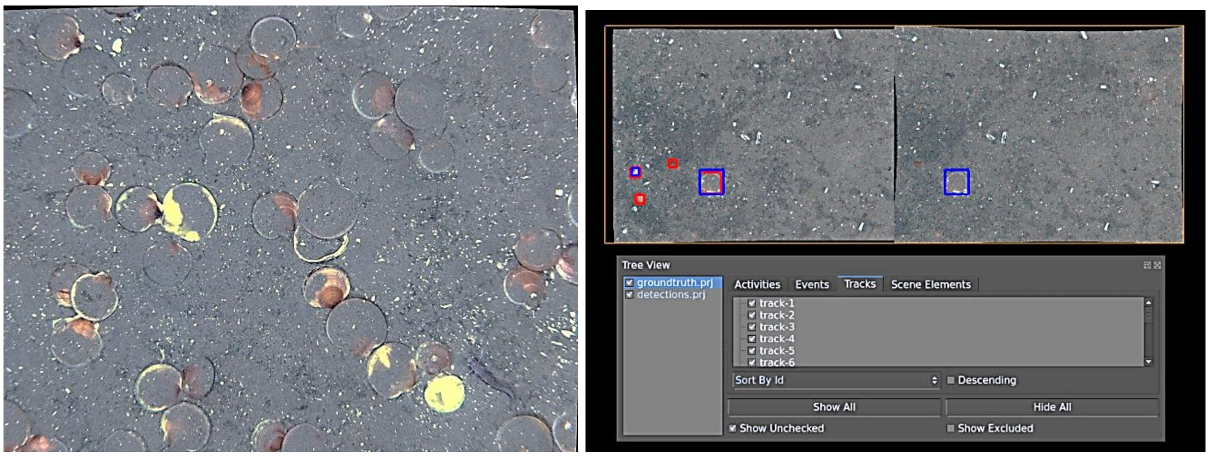

Figure 2-3. VIAME is used for automated scallop detection and measures to provide accurate and timely assessment for fishery management. This has resulted in significant cost savings with a 50% reduction in the manual processing of images. Contact: deborah.hart@noaa.gov

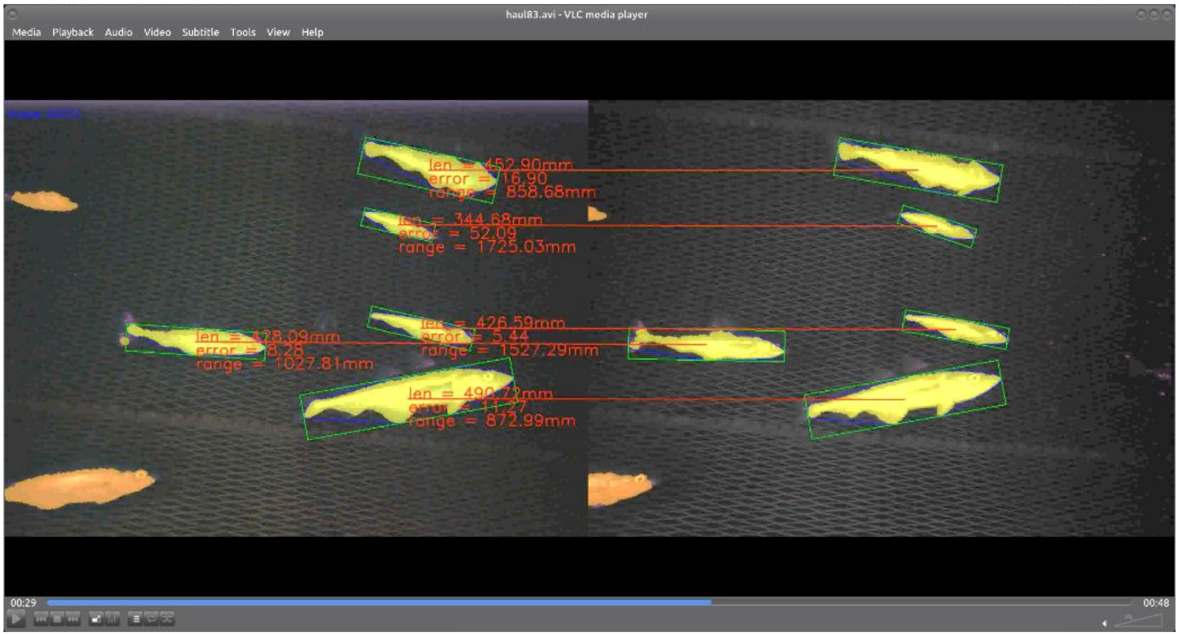

Figure 2-4. VIAME is used to automate the processing of stereo imagery collected during the Bering Sea pollock survey, resulting in improved automation of species identification with accurate length measurements for stock assessments. Contact: kresimir.williams@noaa.gov

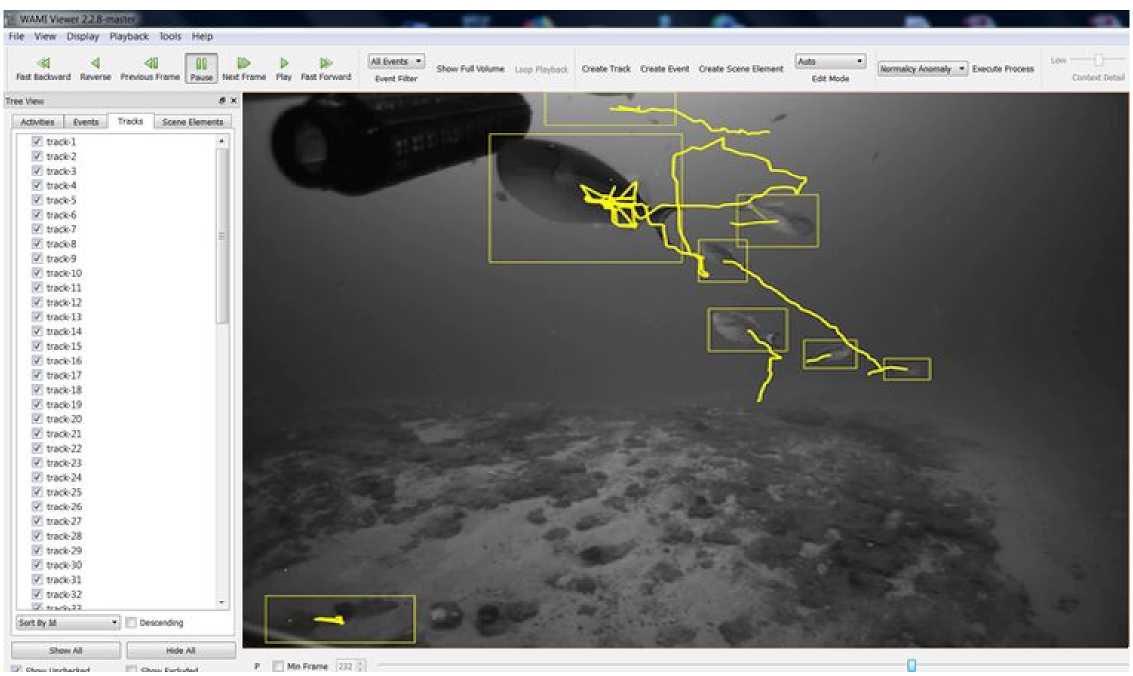

Figure 2-5. VIAME provides automated detection, tracking and classification of fish, and is used to streamline the processing of video collected during the NOAA Southeast Fisheries Science

> Centers reef fish surveys. Progress has been made with developing image training datasets for red snapper and gray triggerfish, and the goal is to build annotated libraries for training classifiers for other visual fish surveys in other regions. Contacts: matthew.d.campbell@noaa.gov and zeb.schobernd@noaa.gov

VIAME uses GitHub to create a user-friendly end-to-end pipeline for ML that enables computer vision and graphical representation for deep learning, feature requests, and task management. Clearly, the availability of user-friendly open source tools has significantly contributed to the progress in ML applications. The Kitware Image and Video Exploitation and Retrieval (KWIVER)[1] includes a repository of computer vision tools and deep learning algorithms that are continuously updated. Kitware Inc.[2] continues to build custom solutions to ML applications and build partnerships to transition research into implementation.

In addition to VIAME, NOAA scientists are evaluating the use of other sources of ML tools. NOAA Fisheries has a new initiative to deploy electronic monitoring systems aboard fishing vessels, and apply ML tools (refer to Section 5). The following example is the application of using Google AI technology to automate the detection, recognition, and annotation of whale songs from passive acoustic data (Fig. 2-7).

[1] KWIVER, https://www.kitware.com/platforms/#kwiver

[2] Kitware Inc., https://www.kitware.com/

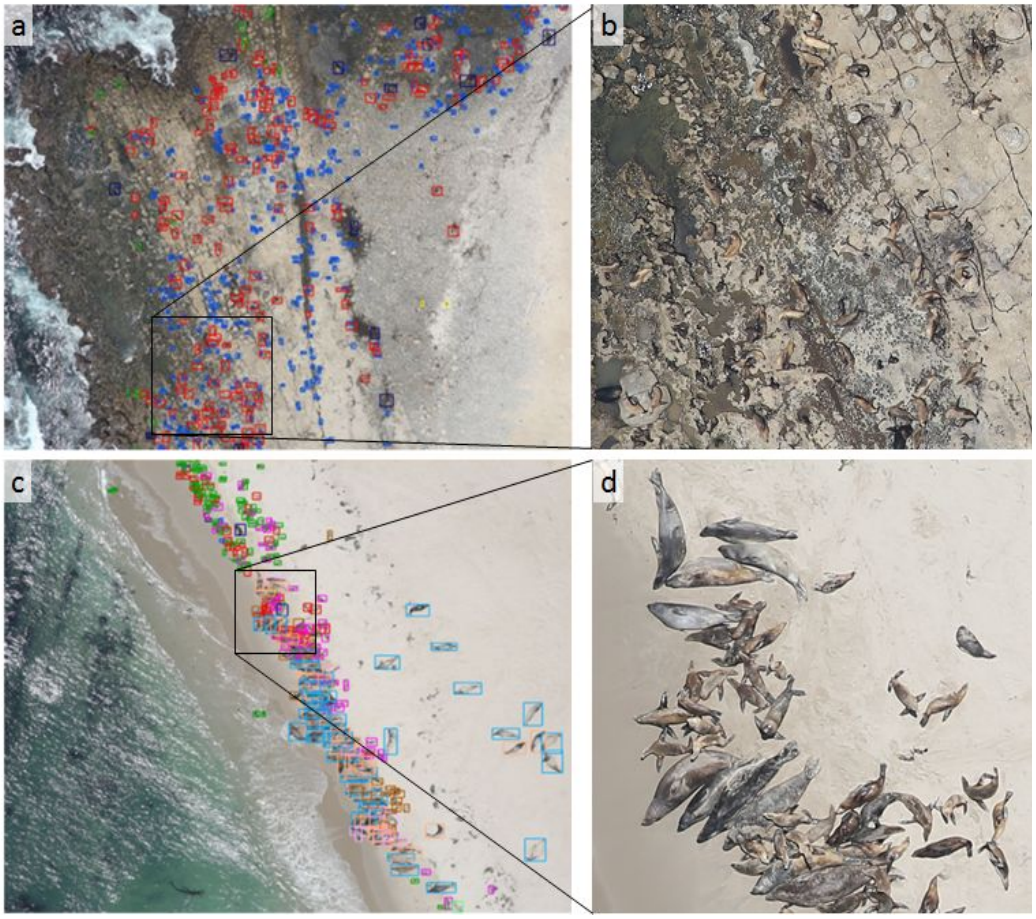

Figure 2-6. The applications of VIAME have recently expanded to aerial surveys. In the example, the NOAA Southwest Fisheries Science Center is using VIAME to detect and classify California sea lions and other seals by species, age, and gender. There is recent success in single detector pinniped models (training with 27 annotated images of 5,428 pinnipeds). Work continues to improve the GUI annotation module of VIAME, and to improve the performance from a single detector model to a multi-class detector model. Contacts: beth.jaime@noaa.gov and george.cutter@noaa.gov

The ML is also being applied to quantify percent cover of coral reef benthic communities using point classification on benthic images to develop the CoralNet[3] software product, which provides an end-to-end web-based tool for the automated processing of benthic photo-quadrats (Fig. 2-8; refer to Section 4).

[3] CoralNet, https://coralnet.ucsd.edu/

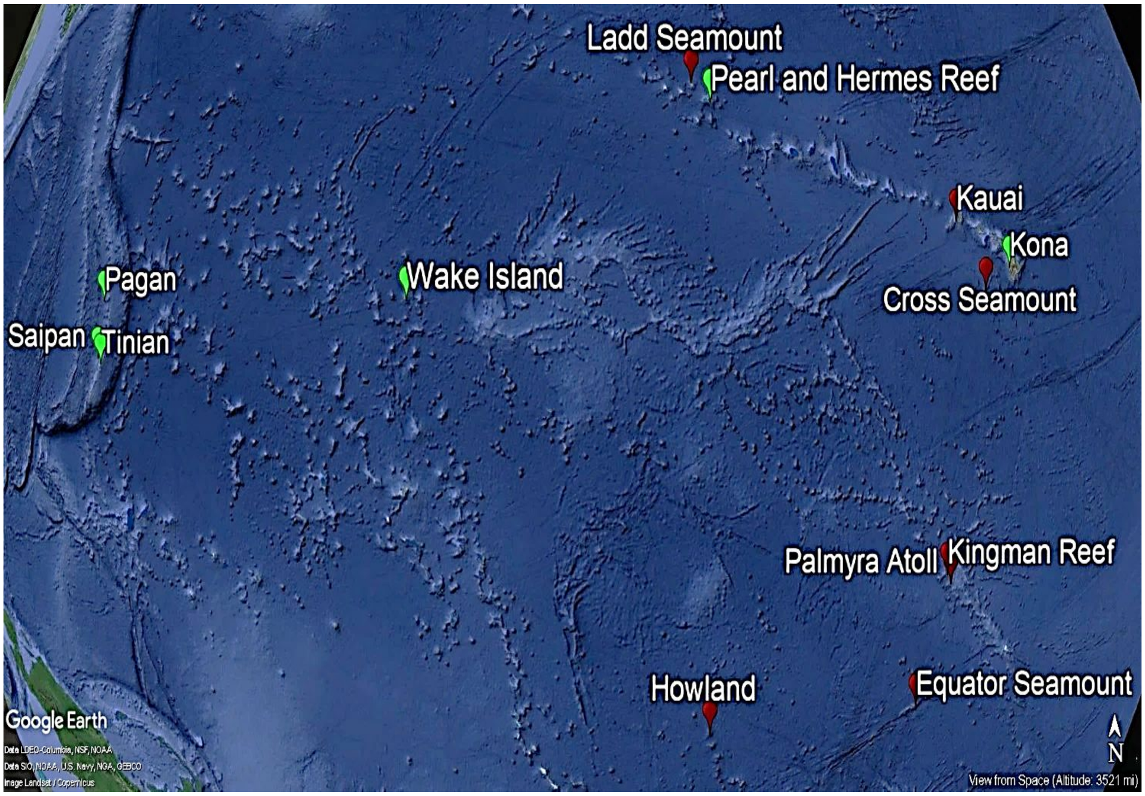

Figure 2-7. The NOAA Pacific Islands Fisheries Science Center's Cetacean Research Program surveys marine mammals, and collected >170,000 passive acoustic recordings from monitoring instruments throughout the Pacific Islands. Google developed a machine learning model to recognize and annotate humpback whale songs, resulting in significant savings in processing with 90% in precision and recall. The effort is presently expanding to automated detection and recognition for other whale species. Contacts: ann.allen@noaa.gov and erin.oleson@noaa.gov

NOAA has recently identified artificial intelligence (AI) as a cross-functional mission priority for the agency, and to work with NOAA line offices to build partnerships to advance the applications of ML analytics. The initiative is closely linked with NOAA's strategic initiatives on unmanned systems (UxS), NOAA Big Data Program, NOAA Enterprise Data Modernization, and NOAA Fisheries' Survey Optimization Strategic Initiatives. The rapid technological advancements in data collection platforms, big data accessibility, and processing efficiencies require an integrated approach with collective partnerships for our next generation of integrated surveys and ocean observations systems. As the accessibility of big data with annotated training data sets becomes more available to ML analytics, the improved quality, timeliness, and discovery capacity will provide added socioeconomic value to science-based products.

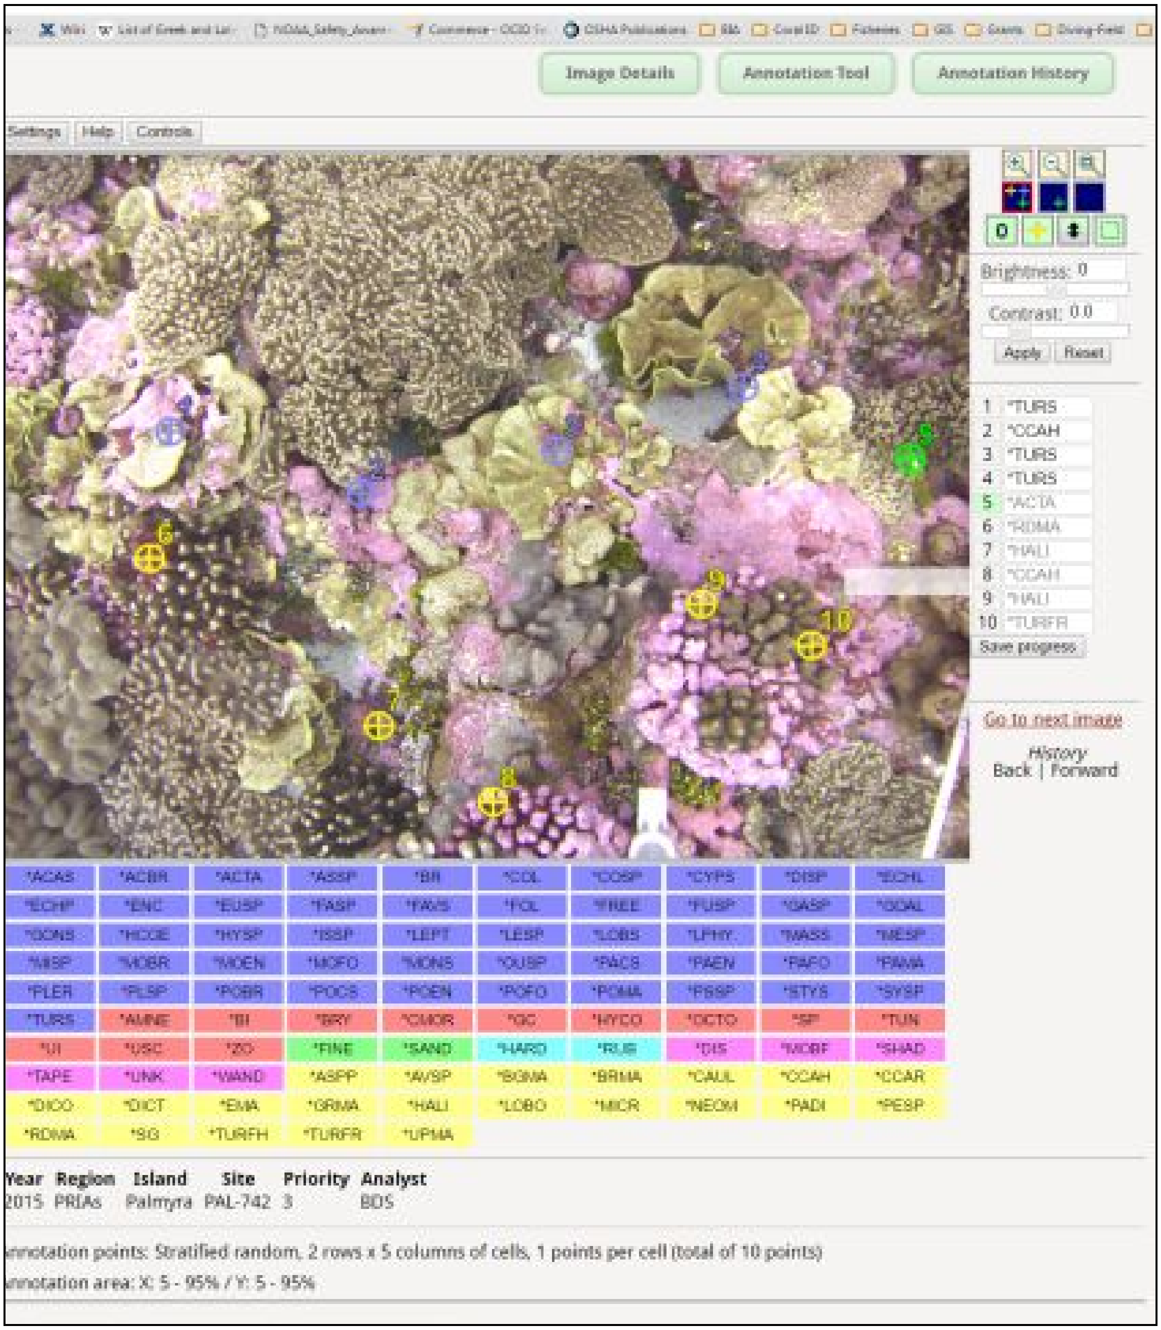

Figure 2-8. NOAA Pacific Islands Fisheries Science Center utilizes the CoralNet software, developed by experts at the University of California San Diego, as a web-based image classification system for coral reef habitats. CoralNet has significantly reduced PIFSC manual processing of coral benthic images.

Discussions: The open source VIAME toolkit that utilizes machine learning for automated detection and classification was recently released, and the number of users within NOAA has dramatically increased during the past year. Some of the examples discussed have demonstrated the processing of imagery data from fisheries surveys has been reduced from months to days, with 93-98% cost savings in cost and time. Integrity of the science used for policy decisions is critical, and in some cases, efforts to refine the labelled training data sets have resulted in 93-98% accuracy in the model detection and classification. NOAA during the upcoming year will be developing an implementation to improve operational efficiencies through its agency using artificial intelligence and machine learning.

Machine intelligence and the data driven future of marine science.

Ketil Malde[1], Nils Olav Handegard[1], Line Eikvil[2], and Arnt Børre Salberg[2]

[1] Institute of Marine Research, Bergen, Norway

[2] Norwegian Computing Centre, Oslo, Norway

Extended Abstract

Technological advances in sensor technology, autonomous platforms, and information and communications technology now allow marine scientists to collect data in larger volumes than ever before. But our capacity for data analysis has not progressed comparably, and the growing discrepancy is becoming a major bottleneck for effective use of the available data, as well as an obstacle to scaling up data collection further. Recent years have seen rapid advances in the fields of artificial intelligence and machine learning, and in particular, so-called deep learning systems are now able to solve complex tasks that previously required human expertise. This technology is directly applicable to many important data analysis problems and it will provide tools that are needed to solve many complex challenges in marine science and resource management (Malde et al. 2019). Here, I highlight some of the recent achievements and why this prepares the ground for further advances into fields beyond data analysis (Fig. 2-9).

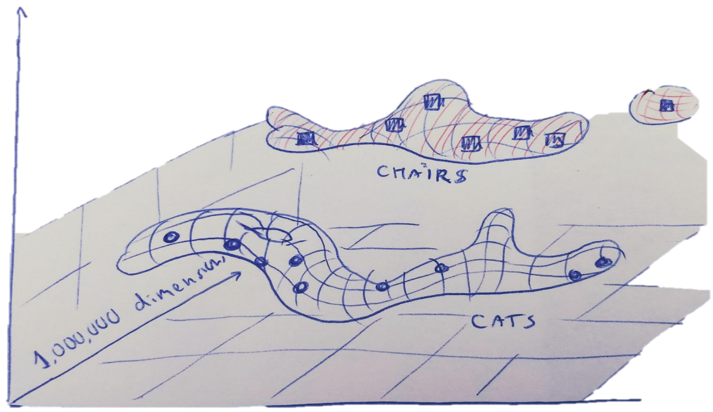

Figure 2-9. Machine learning definition: Identifying complex structures in a vast and mostly empty space, based only on a few examples. Here: the high dimensional space of images, and the regions representing images of chairs and cats (artist's rendition).

Discussions: Scientists at IMR are working along several lines to better utilize ML. The data organization is a key bottleneck and several data types are currently being exposed through application programming interfaces (API) that allow efficient data access. This process is ongoing, and will continue making more data available. IMR are also building competence in openly available machine learning frameworks such as Keras, Pytorch, and TensorFlow. This will allow IMR to better interact with industry partners that can offer competence in machine learning.

The COGMAR project.

Arnt Børre Salberg[1], Line Eikvil[1], Ketil Malde[2], Nils Olav Handegard[2]

[1] Norwegian Computing Center, Bergen, Norway

[2] Institute of Marine Research, Bergen, Norway

Extended Abstract

Deep learning has been called the revolutionary technique that quietly changed machine vision forever, but is at present mainly applicable to standard RGB images of natural scenes or objects, or otherwise only for other types of imagery when a substantial amount of annotated data is available, which is seldom the case.

The COGMAR[4] project aims at enabling this technology for computer vision problems anywhere, by developing easy-to-use cognitive solutions for complex marine data images and thereby extending the use of autonomous cognitive computer vision systems to solve key big data computer vision challenges in the marine sector. The overall concept of the project is to exploit the power of deep learning that tends to obtain improved performance when the amount of training data increases. However, the major challenge is that annotated data is hard to get. The COGMAR project aims to develop new deep learning solutions for learning necessary to classify, localize, and segment objects in non-standard, sparsely labeled, image data. Motivated by the method's ability to generalize and the fact that unlabeled data is often inexpensive to acquire, our approach for solving this will be based on three main concepts: cross-domain transfer learning, semi-supervised learning, and data augmentation and simulation.

Currently, the project focuses on developing automated solutions for analyzing acoustical echosounder data, fish species identification from in-trawl camera images, and estimating fish age from images of otoliths (refer to abstracts in Section 4 by Brautaset, Handegard, and Moen, respectively).

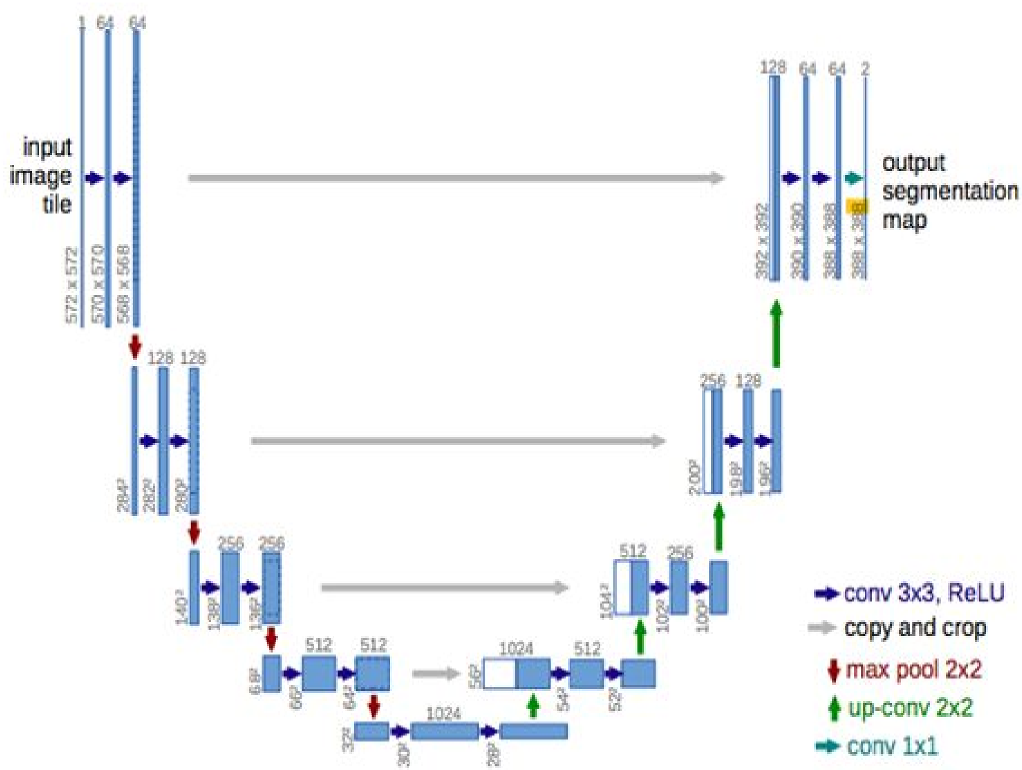

[4] COGMAR, https://www.nr.no/en/projects/cogmar-ubiquitous-cognitive-computer-vision-marine-services

Figure 2-10. For analyzing acoustical echosounder data, the popular UNet architecture for deep learning has proved successful. This semantic segmentation shows each blue box corresponds to a multi-channel feature, while each white box represents copied feature maps.

Major challenges the COGMAR project focuses on include

- Developing methods that generalize to new survey data. This is particularly challenging if the new surveys are from new areas, different depths, and different times of year as the training data.
- Handling class imbalance. We are often looking for small objects, or objects that are sparsely represented in the training data.
- Understanding the predictions. Why is the model predicting that outcome?

Solutions from the project will contribute innovations for industries manufacturing solutions for automated monitoring of fish and marine environments. This effort will continue during 2017-2022, as the COGMAR project expanded its partnership with the Norwegian Computing Center, Institute of Marine Research, University of Tromsø, Scantrol Deep Vision AS, Royal Institute of Technology (Sweden), and University of Maryland.

Discussion: In less than a decade, deep learning analytics has revolutionized AI technology. The COGMAR project will continue to develop easy-to-use cognitive solutions using deep learning for complex marine problems, such as classifiers for acoustic backscatter, fish images, ageing otoliths, identifying plankton, and benthic habitat characterization. When using the transitional classification algorithms, the classification models become more saturated as more data becomes available. On the other hand, deep learning continues to improve classifier models as more data becomes available. Transferring deep learning from the semi-supervised phase of using unlabeled data to the augmentation-simulation phase is powerful. Participants discussed the unique challenges of applying deep learning to acoustic data collection. While image data has challenges with the attenuation of light through the water, acoustic has the challenge of validation of acoustic backscatter to species. There is agreement that deep learning is critical to address the processing bottleneck of the increasing volume of acoustic data, and will likely provide immediate benefits of reducing the human bias associated with the manual post-processing of acoustic backscatter.

NORCE and smart oceans.

Annette Fagerhaug Stephansen

NORCE Norwegian Research Center, Bergen, Norway

Extended Abstract

NORCE[5] is an interdisciplinary research institute with a focus on applied research and innovation. The key areas of research are energy, health, climate, environment, society, and technology, and the institute includes around 1000 employees spread over 10 Norwegian cities.

NORCE has an in-house big data cluster, where the current system counts 400 CPU-cores with 2000 GB of RAM and 500 TB of internal storage. An in-house cluster gives independence and control of the data as well as the competence to deal with it. It is important to provide high functionality to the researchers to avoid the vulnerability which comes with each researcher adopting their own system and using local hard drives.

The 5 Vs of Big Data are volume, variety, velocity, veracity, and value. Smart sensors can increase the velocity part, as in giving the ability to work in real-time with the data being captured. There is also the possibility of using machine learning locally on the sensor, impacting the data gathering in a positive way. NORCE has been working with smart sensors within various disciplines, like O&G and transport, and has a solid competence in working with drones and sea-going autonomous vessels.

NORCE has developed their own decision support, Nlive, for their drones, and we are investigating using machine learning to make the operation of the drones more autonomous. NORCE is the supplier for the European Maritime Safety Agency for using drones to detect oil spills and improve operations.

The sail buoy (Fig. 2-11) was the first unmanned surface vehicle that completed an Atlantic crossing. It can be equipped with a large number of sensors, can be used as a communication relay station, and was recently used for doing a krill survey in the Antarctic.

[5] NORCE, https://www.norceresearch.no/en/

Figure 2-11. The Sail Buoy is the first unmanned surface vehicle that completed an Atlantic crossing.

To investigate the veracity and increase the value of big data, both visualization and decision support are important. The visualization tool Enlighten has been developed to permit the analysis of large data sets and in particular to examine different types of data variety at the same time. The tool has been used in various projects, including the European Plate Observatory System and projects involving the Sail Buoy.

As an example of our decision support development skills, we highlight NORCE's decision support SARA[6] used by the Norwegian Joint Rescue Coordination Centres since 2000 and which is still in operation. For marine stock assessment and research, the Large Scale Survey System (LSSS)[7] has been developed by NORCE in collaboration with IMR and has been used since 2007. It consists of two parts: a pre-processor (e.g., noise reduction, bottom detection, school detection, species identification, plankton analysis) and an interactive interpretation module (Fig. 2-12).

[6] SARA, https://www.cmr.no/projects/10341/sara/

[7]SIMRAD Large Scale Survey System (LSSS) post-processing and analysis acoustic software https://www.simrad.com/www/01/NOKBG0240.nsf/AllWeb/F90228B1C9B6F0EBC12574AA00490DD5?OpenDocument

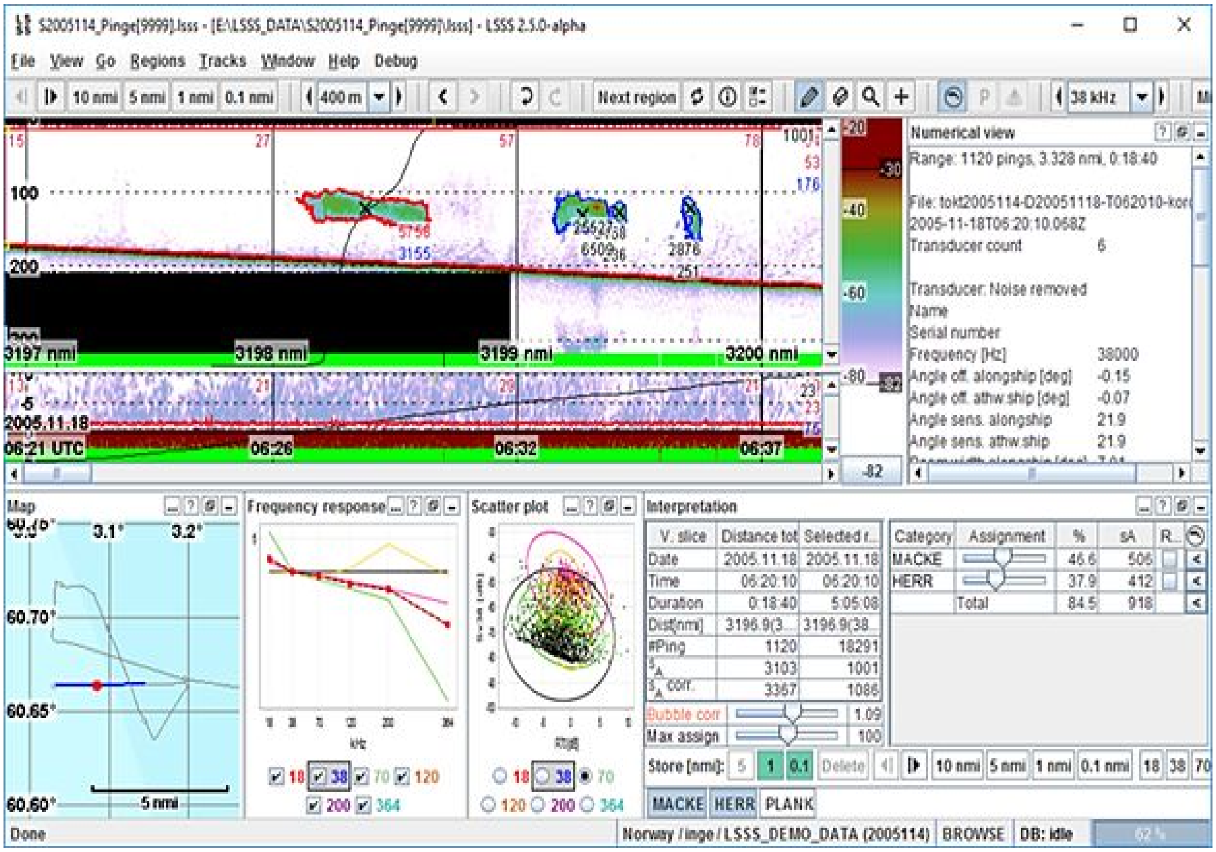

Figure 2-12. The Large Scale Survey System (LSSS) acoustic post-processing and analysis software was developed by NORCE in collaboration with the Institute of Marine Research (IMR).

NORCE has a long experience with the use of machine learning applied to various disciplines and markets. Some concrete examples as applied to the marine environment are estimates of biomass relevant for fisheries management, automatic mapping of plastic, automatic fish detection, lice counting, and automatic prediction of fish growth and health of farmed salmon (Fig. 2-13).

Figure 2-13. NORCE has used machine learning for automatic fish detection, lice counting, and prediction of fish growth and health of farmed salmon.

Contacts: Annette Fagerhaug Stephansen (anst@norceresearch.no), Jeremy Cook (jeco@norceresearch.no), Junyong You (juyo@norceresearch.no), Klaus Johannsen (kljo@norceresearch.no), Alla Sapronova (alsa@norceresearch.no) and Inge Eliassen (inel@norceresearch.no).

Discussion: NORCE uses their own architectural infrastructure for their innovative research, believing independence is important. Discussions focused on the challenges of the 5 Vs of big data: volume (must optimize storage), variety (structured vs. unstructured data), velocity (speed of creation and transfer), veracity (trustworthiness of data), and value (value of cross-functional and discovery products). There were also discussions on the trade-offs of cloud services and hybrid approaches, and concerns on the bandwidth requirements for big data sets. One option is to use in-house big data cluster architecture which provides more flexibility. However, while researchers need some degree of independence (e.g., regional servers) when analyzing data, they should never be allowed to store data on their own hard-drive storage devices. There was recognition that sampling technologies, including unmanned platforms, have dramatically increased the volume of data collection. It's easier to add computing power to the data, than to process and analyze the data. Processing efficiencies using ML are highly dependent on the metadata quality and accessibility. Smart sensors aboard unmanned platform operations can check data collections in real-time. Efforts are underway to incorporate ML into drone operations for automated detection and real-time visualization of large data collection by extracting clusters of data sets with Enlighten software.

3 Case Studies: Deep Vision and Image Analysis

State of Deep Vision image analysis - Deep Vision Project.

Helge Hammersland and Hege Hammersland-White

Scantrol Deep Vision AS, Bergen, Norway

Extended Abstract

The Deep Vision[8] system has been developed over a period that spans more than a decade. Deep Vison is a subsea vision system for identifying and measuring fish underwater. Today, Deep Vision is sold for marine research purposes as a system that can identify the species, length, and location of fish from images taken in the trawl (Fig. 3-1). Under development is a version that will be launched as a sorting device for commercial trawlers where analysis of size and species (Fig. 3-2) will trigger a sorting mechanism in the trawl so that only desired fish are caught in the trawl whilst the rest are released back into the sea.

Figure. 3-1. The participants of the Norway-US workshop on machine learning were provided a demonstration of Scantrol's Deep Vision stereo camera system.

[8] Deep Vision, https://deepvision.no/deep-vision/deep-vision

Figure. 3-2. Image of fish and fish length measures from the Deep Vision system.

Key to the functioning of the system is the Deep Vision analysis software. The software has been developed by Scantrol Deep Vision AS[9], together with key resources in their field from the research community. In 2017, Girona Vision Research[10] was co-founded by the company in order to access highly specialized resources for underwater robotics and vision development.

Primary challenges in the years to come, and to facilitate the employment of the system for commercial trawl fisheries, involves further developing the software for automatic length measurement and species recognition. Online communication via acoustic link from the Deep Vision trawl system to the vessel is currently being tested and will be further developed in the years to come. Finally, all the electronics added to the system will have to be developed and tested for use in rough environments, both in the sea and on deck.

Discussions: Challenges include species classification, need for more compact and low-cost electronics, and improved underwater communications. Presently the company is working with

[9] Scantrol Deep Vision AS, https://deepvision.no/
[10] Girona Underwater Vision and Robotics, https://cirs.udg.edu/

Kongsberg to improve the acoustic communication link, and with Girona Vision Research to build ML expertise and capability. Deep Vision works closely with fishers to help reduce bycatch during commercial pelagic and bottom trawl operations, and this can provide economic benefits to their industry.

State of deep vision image analysis - Deep Vision project.

Håvard Vågstøl

Scantrol Deep Vision AS, Bergen, Norway

Extended Abstract

When talking about classifier systems, abundant high-quality data is a good predictor of classification accuracy. With an increasing number of Deep Vision systems being deployed around the world, there is great potential for mutual gain to be had from the aggregation of data from different sources. Expert users ensure the quality control of data, and a centralized model generation allows for continuous improvement of classification models. This is the philosophy behind the Deep Vision Bio Base, the opt-in third platform of the Deep Vision ecosystem - where expert knowledge is combined into quality classifier models for optimal results for the end user (Fig. 3-3).

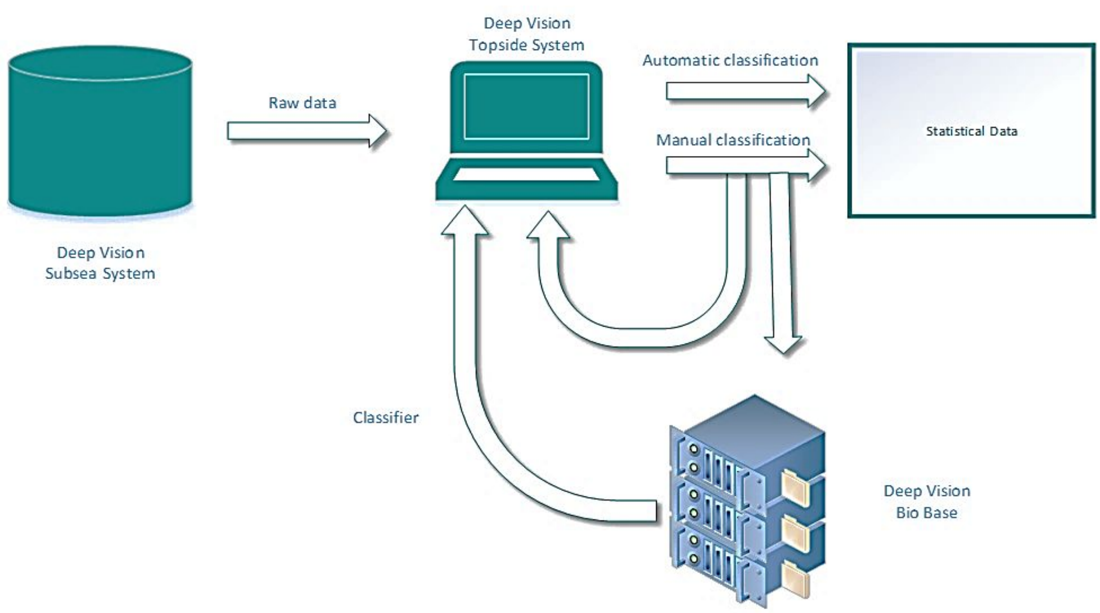

Figure 3-3. Workflow of the Deep Vision Bio Base, where the opt-in third platform of the Deep Vision ecosystem and expert knowledge is combined into quality classifier models for optimal results for the end-user.

Summary of the user perspective of the DVBB: In the field, users work offline with a preloaded classifier parameter set. When working with analysis, the classifier suggests species for the approval of the user. Approved data are considered expertly labeled, and stored with object cutout/bounding box and metadata. Unknown species may be entered manually, thus providing a training example. The local system may include features for single-example or few-example training and classification of new classes. Ideally the quality of such should increase with more examples. The labeled data set may be uploaded to the training server, and updated classifier data may be downloaded from the training server.

Deep Vision – From concept to production model.

Kristoffer Løvall

Scantrol Deep Vision AS, Bergen, Norway

Extended Abstract

The Deep Vision stereo camera system is a result of collaboration between numerous companies, institutions and groups through its journey from concept to finished research version (Fig. 3-4). A large group of engineers, biologists, computer scientists, and others have worked together for more than a decade to form a system that is well suited for consistently taking high quality images in rough and varying conditions. Along the way more than a few prototypes have been built, a large number of images have been taken and, most importantly, a vast amount of experience has been accumulated. The presentation aims to give the audience a brief review of the development timeline, as well as a technical explanation of the resulting system and its use.

Figure 3-4. Deployment of the Scantrol Deep Vision stereo camera system.

The Deep Vision Research Version has gone through extensive prototyping and testing. The information and experience gained from this have been consolidated into a production model that has proven to be a great resource for IMR as well as being planned for use on numerous upcoming research vessels. The system in its standard version is designed for 8 hours of operating time with replaceable batteries for continuous use and takes 5 stereo images per second at a 1.4 MP resolution in a controlled and illuminated environment (Fig. 3-5). The frame is robust and modular - the studio can be removed from the frame and replaced for any reason (i.e. for custom duct shapes, different/added illumination, different background). Current development work is on improving the analysis software as well as making prototypes of the selection device that is to be fitted between the Deep Vision Studio Frame and the cod-end. This is to be a motorized device, able to sort out the catch based on catch composition, time and depth.

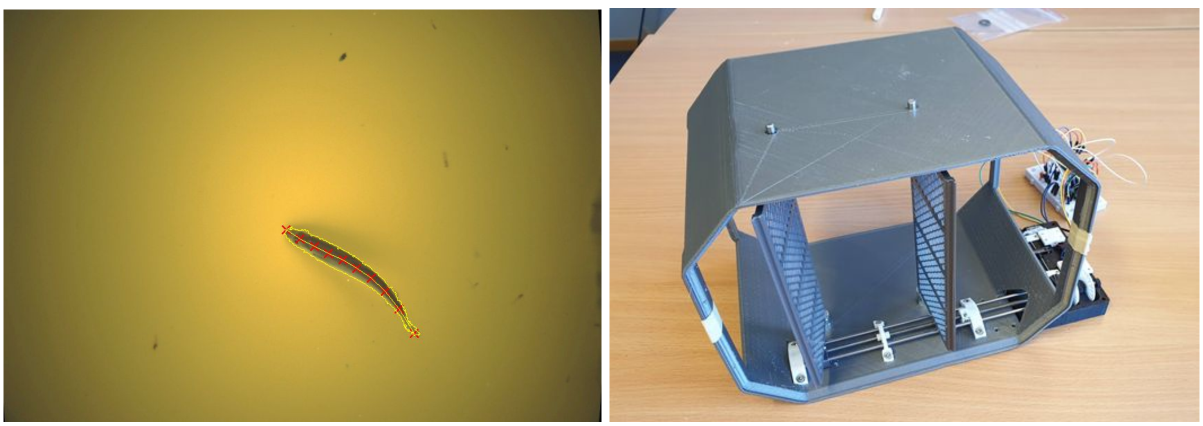

Figure 3-5. The Scantrol Deep Vision Research Version provides high resolution (1.4 MP) images and image analysis.

Discussion: Scantrol collaborated with IMR with the early system that imaged fish on sorting belts aboard the research vessel G.O. Sars. The system was expanded to the CRISP project to refine the fish imaging system in the trawl cod-end. The IMR-Scantrol partnership contributed to the success of the Scantrol Deep Vision Research version which provides automated high-resolution images of fish passing through the trawl. Stereo imaging is presently being refined to improve fish length measurement, which is critical for age-based stock assessments. One of the main challenges today is the need for better annotation subscription software to enable rapid image classification with ML. There were also discussions about data ownership trade-offs. Is it better for the private sector to take ownership of the data management and development of annotated libraries for their systems, and then simply provide the classifiers?

Scientists typically do not want to conduct research and surveys using a ‘black box’ situation, so the open source solutions are often preferred.

Segmentations and measurements.

Ricard Prados Gutiérrez

Girona Vision Research, Spain

Extended Abstract

The Scantrol Deep Vision Subsea Unit developed by Scantrol Deep Vision AS features a high-resolution stereo camera system able to acquire color pictures with a given frequency and under controlled conditions, when attached to a fishing trawl.

The goals of the Deep Vision system include fish species identification, tracking for counting purposes, and measurement. The previously used manual measurement of fish procedure, requiring the user to click at least two points (front and back) of a given specimen in both stereo pair images, suffers from several problems such as subjectiveness and repetitiveness. The specifications of the system (featuring a specifically selected background color and an appropriate lighting setup) have allowed us to develop an optimized segmentation algorithm that exploits these constrained acquisition conditions, and is able to deal with challenging artifacts such as non-uniform illumination, shadows, and reflections. This accurate fish segmentation enables precise measurement using the information provided by the stereo calibration of the system. Then, an automated fish measurement pipeline, involving the approximation of the fish pose using a RANSAC-based curve estimation method, has been developed, aiming to reduce human intervention while improving accuracy.

The full measurement pipeline starts with a precise segmentation in both images of the stereo pair of a manually-clicked fish on the right one (Fig. 3-6). Then, the morphological skeleton of the right fish segmentation is computed and refined, and the extracted points are used as coordinates to estimate a curve. A set of control points, including front and back, are distributed along the estimated curve on the right image along the fish segmentation, and mapped later on into the left one. These control points define segments whose lengths are computed by means of epipolar geometry, and finally added to provide the fish measurement.

The results obtained by the segmentation pipeline, as well as those obtained by the automatic fish measurement method, can be used as labelled data to train deep-learning based classification algorithms.

Figure 3-6. Result of the automatic fish segmentation and measurement pipeline. The specimen is segmented analyzing the image to background saturation ratio. The measurement is later on performed by computing a third order polynomial using RANSAC approximating the shape of the fish-blob skeleton morphological operation. The length of the fish has been estimated as 267 mm.

Discussions: There were discussions on improving the segmentation pipeline, and the need to conduct segmentation of fish and object detection on complex backgrounds.

VIAME toolkit and potential deep learning applications.

Matt Dawkins and Anthony Hoogs

Kitware Computer Vision, Inc., Clifton Park, NY, USA

Extended Abstract

Seafood sustainability is predicated on healthy fish and shellfish populations. Recent developments in the collection of large-volume optical survey data by autonomous underwater vehicles, stationary camera arrays, and towed vehicles has made it possible for fishery scientists to generate species-specific, size-structured abundance estimates for different species of marine organisms via imagery (Fig. 3-7). The immense volume of data collected by such survey methods quickly exceeds manual processing capacity and creates a strong need for automatic image analysis. To address these challenges, we have created the Video and Image Analytics for Marine Environments (VIAME) toolkit[11]. VIAME is an open-source computer vision software platform designed to integrate common image and video analytics, such as stereo calibration, object detection, and object classification, into a sequential data processing pipeline that is easy to program, multi-threaded, and generic.

[11] Video and Image Analytics for Marine Environments (VIAME), http://www.viametoolkit.org/

The system provides a cross-language common interface for each of these components, multiple implementations of each, as well as unified methods for evaluating and visualizing the results of different methods for accomplishing the same task. Most recently, the ability to measure fish and to rapidly train models for novel detection tasks has been integrated into the platform. Sponsored by the Automated Imagery Analysis Strategic Initiative of the United States National Oceanic and Atmospheric Administration's (NOAA)[12] National Marine Fisheries Service (NMFS)[13], VIAME will be deployed at multiple NOAA Fisheries Science Centers to continue improving scientific data that support stock assessments. VIAME is freely available to the global community of marine researchers (Richards et al. 2019).

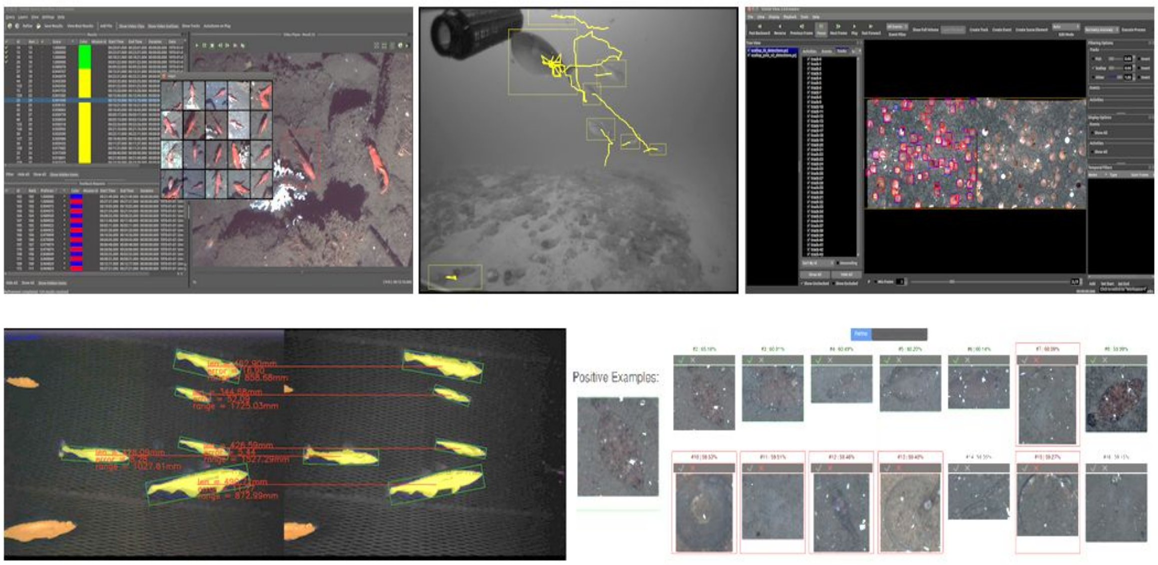

Figure 3-7. Example of VIAME capabilities and desktop version graphical user interfaces.

Discussions: There were discussions on the general categories of ML applications; standard object detection classifiers, box or polygon detection classifiers, and full frame pixel-based classification. The challenges often include the lack of annotated data, need to use established format standards for image data collection, and more user-friendly installation and GUI improvements.

[12] National Oceanic and Atmospheric Administration (NOAA), https://www.noaa.gov/

[13] NOAA Fisheries, National Marine Fisheries Service (NMFS), https://www.fisheries.noaa.gov/

4 Case Studies: Machine Learning Applied to Surveys

Automated species recognition using CNNs.

Vaneeda Allken, Nils Olav Handegard, Shale Rosen, Tiffanie Schreyeck, Thomas Mahiout, Ketil Malde

Institute of Marine Research, Bergen, Norway

Extended Abstract

Marine stock management and the environmental monitoring of marine life can be done through acoustic-trawl surveys. Commonly encountered pelagic species in the area we target are herring, blue whiting, and mackerel, and it can be challenging to separate these species acoustically. In a recent acoustic trawl survey, underwater images were taken with the Deep Vision (Scantrol Deep Vision AS, Bergen, Norway) system at a frequency of 10 frames per second, resulting in millions of images. Our primary aim is to develop a system for automatic fish species identification to support acoustic trawl surveys, using deep learning methods such as convolutional neural networks (CNNs) for image classification. In this study, we explore the potential of the use of synthetic data as a way to generate the large data sets necessary for training deep learning classifiers. One recurring problem in deep learning is the lack of sufficient amounts of labelled training data as the process of labelling by a human expert is often labor intensive. Effective simulation methods to generate synthetic data thus opens up new fields to analysis by deep learning methods. Generating synthetic data is also ideal for object location, in our case identifying the position and size of each fish, in addition to determining which species it belongs to (Fig. 4-1). Annotating images for this type of task is especially labor intensive as it involves drawing a box around every individual object, but synthetic data provides those annotations "for free". We find that training deep neural networks on synthetic data works reasonably well for fish detection tasks and can give near-human level performance on fish classification tasks (Allken et al. 2019).

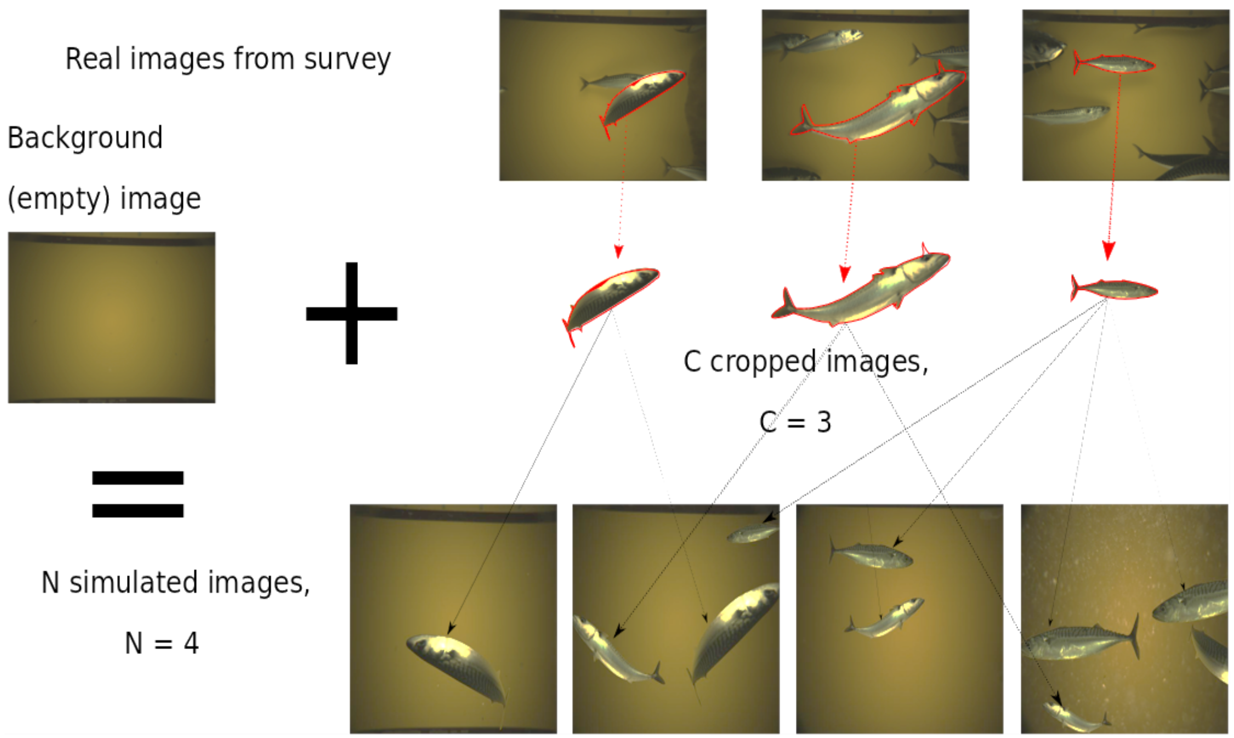

Figure 4-1. Generating synthetic images from fish crops for training.

Discussions: There were discussions on the performance of the classifiers required by the end users. For example, stock assessment scientists require a high degree of accuracy and precision in detection, classification, and measures from ML models when addressing harvest control rules and other fishery management policy decisions.

The IMR deep vision pipeline for the herring survey, DV + CNN + LSSS.

Nils Olav Handegard[1], Petter Rossebø[1], Shale Rosen[1], Vaneeda Allken[1], Håvard Vågstøl[2], Rolf Korneliussen[1] and Ketil Malde[1]

[1] Institute of Marine Research, Bergen, Norway

[2] Scantrol Deep Vision AS

Extended Abstract

We are working on implementing the Deep Vision system on the Norwegian Sea Herring surveys in May. The objective is to have automated classification of herring, mackerel, and blue whiting along the trawl track visible in the Large Scale Survey System (LSSS) software that is used to annotate acoustic data. There are several steps and pieces of software that need to work together before a successful system deployment. These steps include practical handling of the system by the crew, data logistics, running the Deep Vision software, development and deployment of a computer vision system, integrating the data flow into our LSSS acoustic interpretation tool, ensuring that all the components are connected, and that the work flow is sufficiently efficient. Figure 4-2 presents the current implementation, and good strategies for stitching the software components together were discussed. There are pros and cons between a single monolith software and several smaller pieces, but it is hard to strike the correct balance.

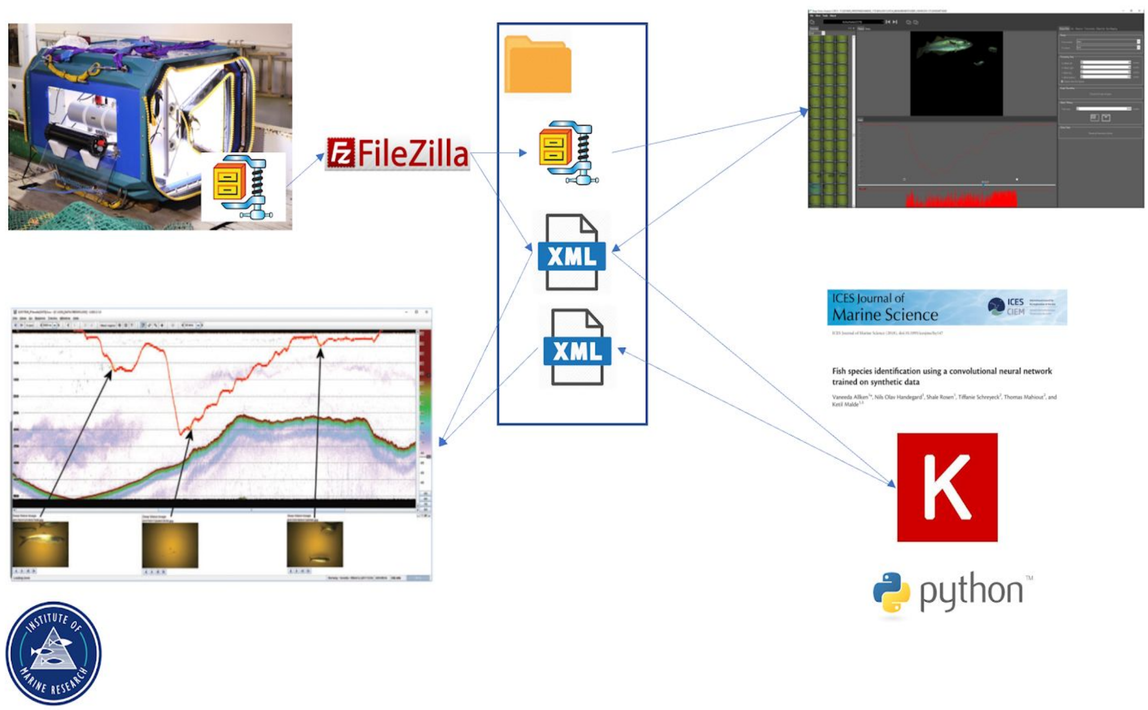

Figure 4-2. The different components in our implementation of Deep Vision in acoustic trawl surveys.

Discussions: There were discussions on dynamic post-processing and the need to improve segmentation using multi-frequency response from acoustic backscatter. The importance of accuracy was discussed when deriving abundance estimates for stock assessments. The primary challenge is species classification from acoustic data; however, there is agreement that ML can automate and reduce the human bias associated with post-processing acoustic data.

Automated acoustic data processing.

Olav Brautaset[1], Anders Ueland[1], Arnt Børre Salberg[1] and Nils Olav Handegard[2]

[1] Norwegian Computing Centre

[2] Institute of Marine Research, Bergen, Norway

Extended Abstract

Assessing the amount of fish from acoustic images currently involves several manual steps, in particular the manual allocation of acoustic energy to species or species classes. We propose a method for automating this by means of image segmentation using deep learning.

Our dataset consists of annotated echograms, which are multi-frequency echosounder images (Fig. 4-3). Fish appear in schools, and a single pixel in an echogram can contain a large number of individuals. The schools typically comprise a small portion of the pixels in an echogram. To classify fish schools in echograms, we develop a segmentation model for separating pixels of fish from background pixels. We obtain this by training a convolutional neural network (UNet architecture) on annotated echograms.

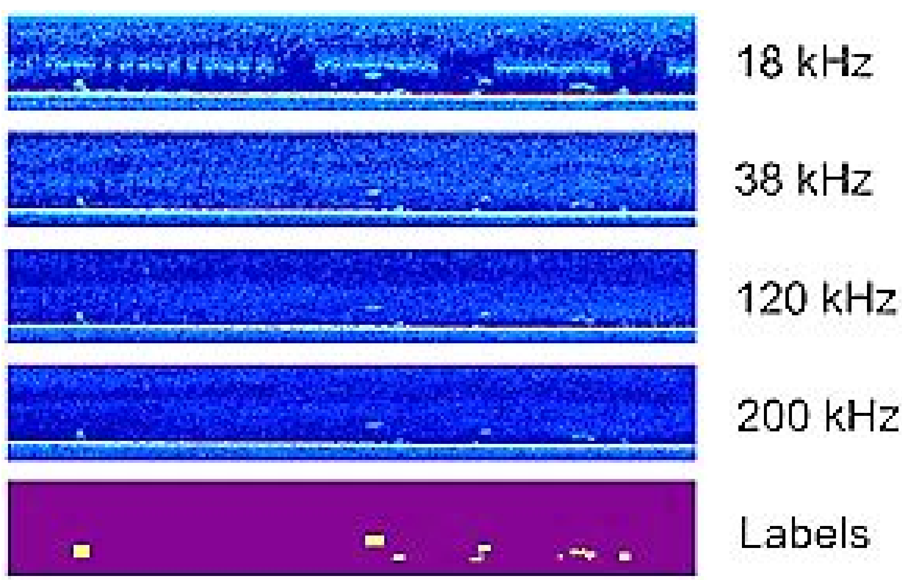

Figure 4-3. Dataset consists of annotated echograms, which include multi-frequency echosounder images. Fish appear in schools, and a single pixel in an echogram can contain a large number of individuals.

When training our segmentation model, we incorporate an image preprocessing step as part of the neural network itself (Fig. 4-4). We replace the commonly used decibel transform by a parametrized log-sigmoid function, applied to the individual pixels of the input data. These parameters are channel-specific and are optimized together with the parameters of the UNet network (Ronneberger et al. 2015). We find this speeds up training and increases the performance of the final segmentation model.

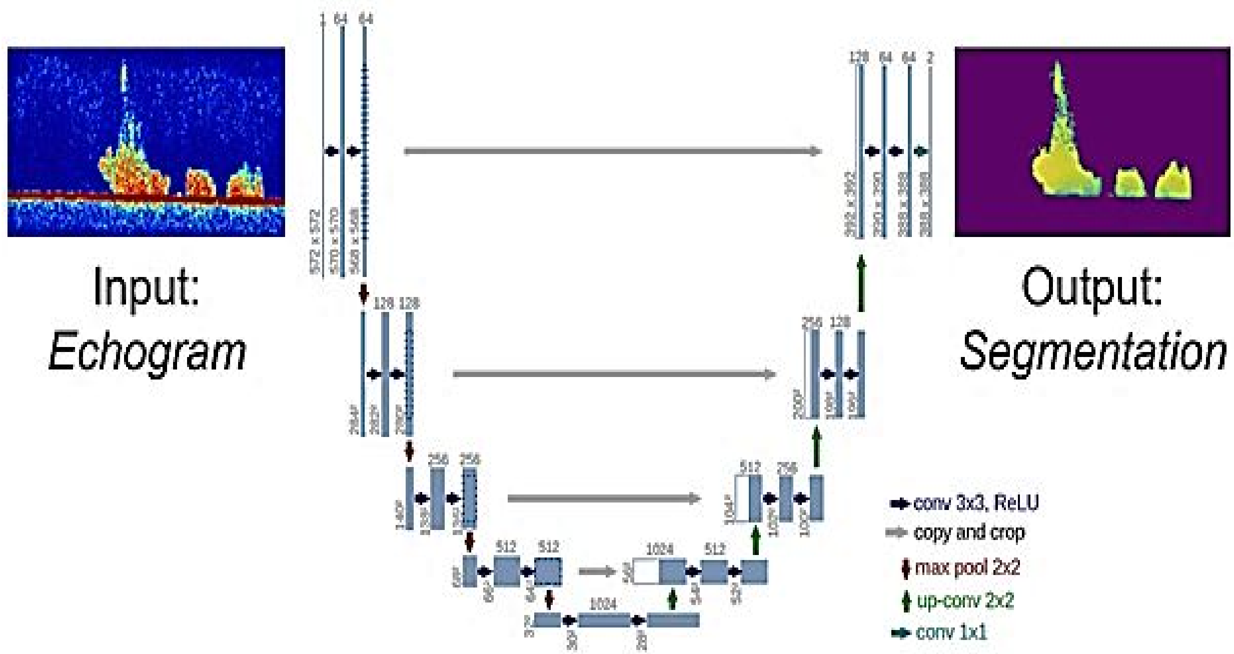

Figure 4-4. Segmentation models are used to separate pixels of fish from background pixels, and this is used to train a convolutional neural network (UNet architecture) on annotated echograms.

We also aim at discriminating between fish species, which in principle can be obtained directly by applying the UNet architecture as a multi-class pixel classifier (background, fish type A, fish type B, etc.). The relative frequency response in the acoustic data should contain sufficient information to discriminate between fish species. As of today, we are not able to obtain this with the UNet architecture. Currently we are investigating whether our dataset may contain inconsistent annotations or different instrumentation settings leading to non-homogeneity between subsets of our data. We will explore various normalization techniques to overcome these issues.

Discussion: The UNet architecture uses computer vision with semantic segmentation for pixel classification, and then applies the fully convolutional networks (FCNs) model. The CNN cannot process the whole image at once, so FCNs process small segments of pixels each time for image classification. Further information on CNN for semantic segmentation is provided by Shelhamer et al. (2017). This deep learning method is commonly used for various applications, such as biomedical imaging, autonomous vehicles, and geoscience.

Population estimation by combining automated and manual image annotations.

Deborah R. Hart, Jui-Han Chang, Burton V. Shank

NOAA Fisheries, Northeast Fisheries Science Center, Woods Hole, MA, USA

Extended Abstract

Surveys have been conducted in the Georges Bank region using the Habitat Mapping Camera System (HabCam) to provide abundance estimates for stock assessments[14]. The Video and Image Analytics for Marine Environments (VIAME) software[11] has been used to streamline image data processing with deep learning detection and classification of scallops, and efforts are underway to apply this to fish (Fig. 4-5).

[14] Habitat Mapping Camera System (HabCam), https://habcam.whoi.edu/

Figure 4-5. Imagery from the HabCam survey operations with an example of skate detections using VIAME/YOLO v2.

1

Comparisons between the manual and automated detection of skates during survey operations on Georges Bank provides an example of ML application to the skate complex in the northeast U.S. using data from NOAA HabCam data (Fig. 4-6).

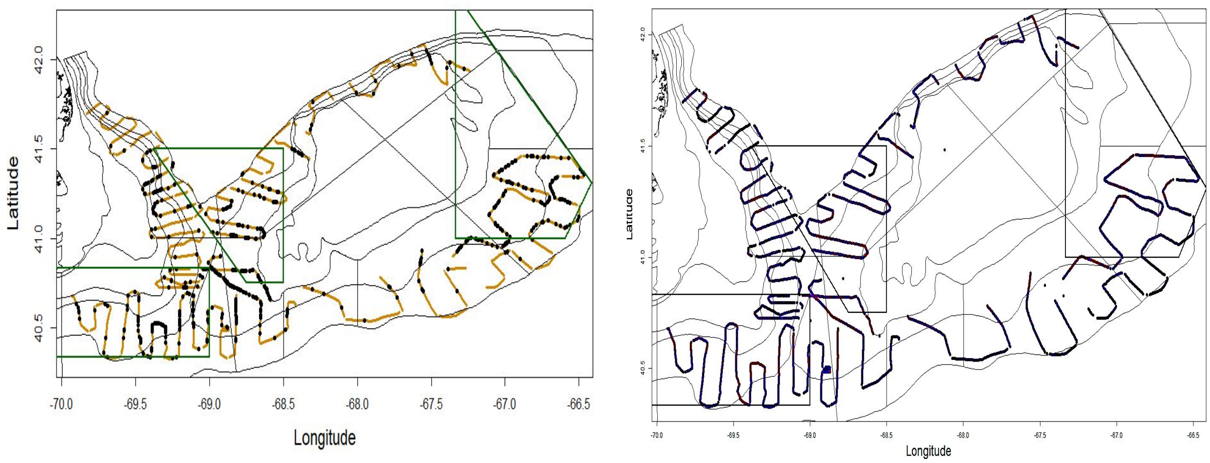

Figure 4-6. Comparison of manual (left) and automated (right) detections of skates on Georges Bank. The black dots indicate manual detections, where the colors represent densities of skates according to the automated annotations (blue: low density, red: high density).

Although automated image analysis is rapidly improving and becoming more popular, the final step of estimating population abundance from the automated counts has received less attention. Automated annotators can make errors, and the error rates are often autocorrelated and vary spatially. Therefore, the population estimates based on the raw computer automated counts can be biased. We evaluate and discuss methods to combine automated annotations with manual annotations of a sample of the images to obtain unbiased population estimates that can be more precise than the manual annotations alone. These methods include stratifying the images based on the automated results, and various forms of generalized regressions. In Figure 4-7, our results indicate that the probability of a false negative of automated object detector is a key factor in determining the best method (Chang et al. 2016).

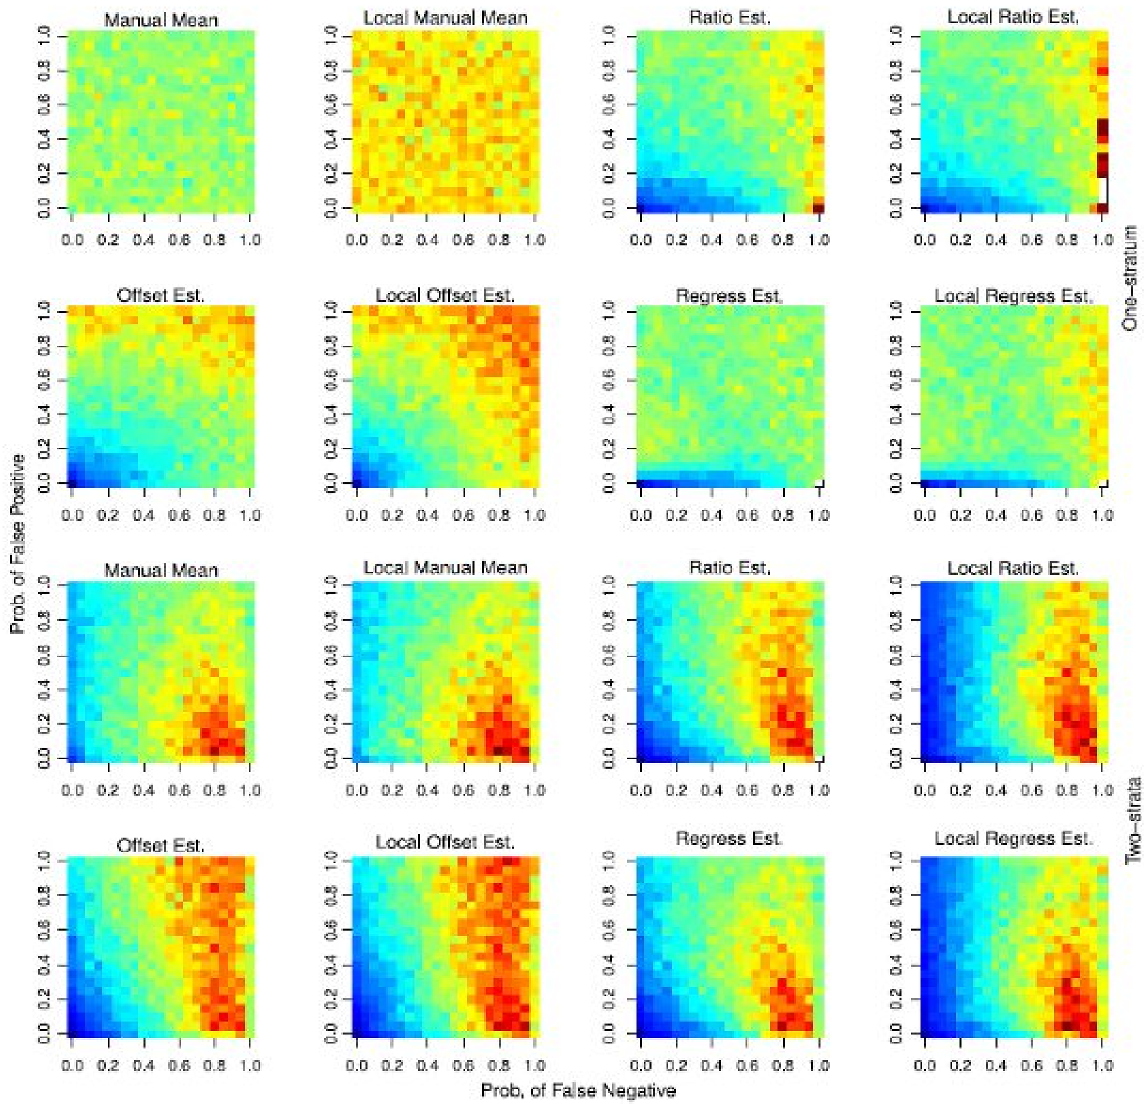

Figure 4-7. Precision of abundance estimates using various candidate estimators for combining automated and manual annotations. Blue denotes high precision; red low precision. Local ratio or regression estimators with stratification appear to generally give the best results (Chang et al. 2016).

Discussion: Independent validation of the ML model performance and statistical accuracy using labeled test data is critical. The degree of precision and accuracy required by the end-user can vary depending on the application. For example, the detection and classier accuracy of 90% may not be sufficient for stock assessment purposes because a high degree of trustworthiness of the scientific information is expected for policy decision results.

Scaling up coral reef monitoring through imagery and machine learning: Advances and next steps.

Courtney S. Couch[1,2], Thomas A. Oliver[1], Ivor D. Williams[1], Rhonda Suka[1,2], Vid Petrovic[3], Oscar Beijbom[4,5], Bernardo Vargas-Angel[1,2], Stuart Sandin[6], John Burns[7], Falko Kuester[3], Russell E. Brainard[1]

[1] Ecosystem Sciences Division, Pacific Islands Fisheries Science Center, National Oceanic and Atmospheric Administration, Honolulu, USA

[2] Joint Institute for Marine and Atmospheric Research, University of Hawaii at Manoa, USA

[3] Cultural Heritage Engineering Initiative, Qualcomm Institute, University of California San Diego, USA

[4] Global Change Institute, University of Queensland, Australia

[5] Berkeley Artificial Intelligence Research, University of California Berkeley, USA

[6] Scripps Institution of Oceanography, UC San Diego, USA

[7] Department of Marine Science, University of Hawaii at Hilo, USA

Extended Abstract

As the threats to coral reefs mount, reef managers are leaning more heavily on digital imagery to increase the scale and scope of coral reef monitoring. With imagery becoming easier to collect, our growing challenge is rapidly converting imagery into timely, management-relevant metrics of reef condition (Fig. 4-8). NOAA's Ecosystem Sciences Division (ESD) has conducted monitoring at 40 islands across the Pacific since the early 2000s and recently assessed the capacity of the trained and widely-used machine-learning image analysis tool CoralNet to generate fully-automated benthic cover estimates from small-scale photo-quadrat imagery.[3]CoralNet 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). CoralNet was generally effective at estimating cover of common coral genera, but performance was mixed with less common coral taxa and benthic algae (Lozada-Misa et al. 2017).

To improve efficiency and spatial coverage of benthic field assessments beyond benthic cover estimates, ESD is collaborating with partners to develop a processing pipeline for extracting benthic metrics (e.g. coral density, surface area, partial mortality, and health status) from large-area imagery using structure-from-motion. While significant progress has been made in

developing a software pipeline for generating 3D dense point clouds and extracting reef metrics, timely data generation is impeded by the lack of a cloud processing framework and inadequate semi-automated segmentation of coral colonies (Figs. 4-9 and 4-10).

The development of a modular and scalable processing pipeline that leverages machine learning and cloud processing has the potential to transform the way we assess and ultimately manage coral reefs worldwide. These approaches are also likely transferrable to other benthic habitat and marine ecosystem types. Further information on leveraging machine learning to improve science on coral reef status can be found in Williams et al. (2019).

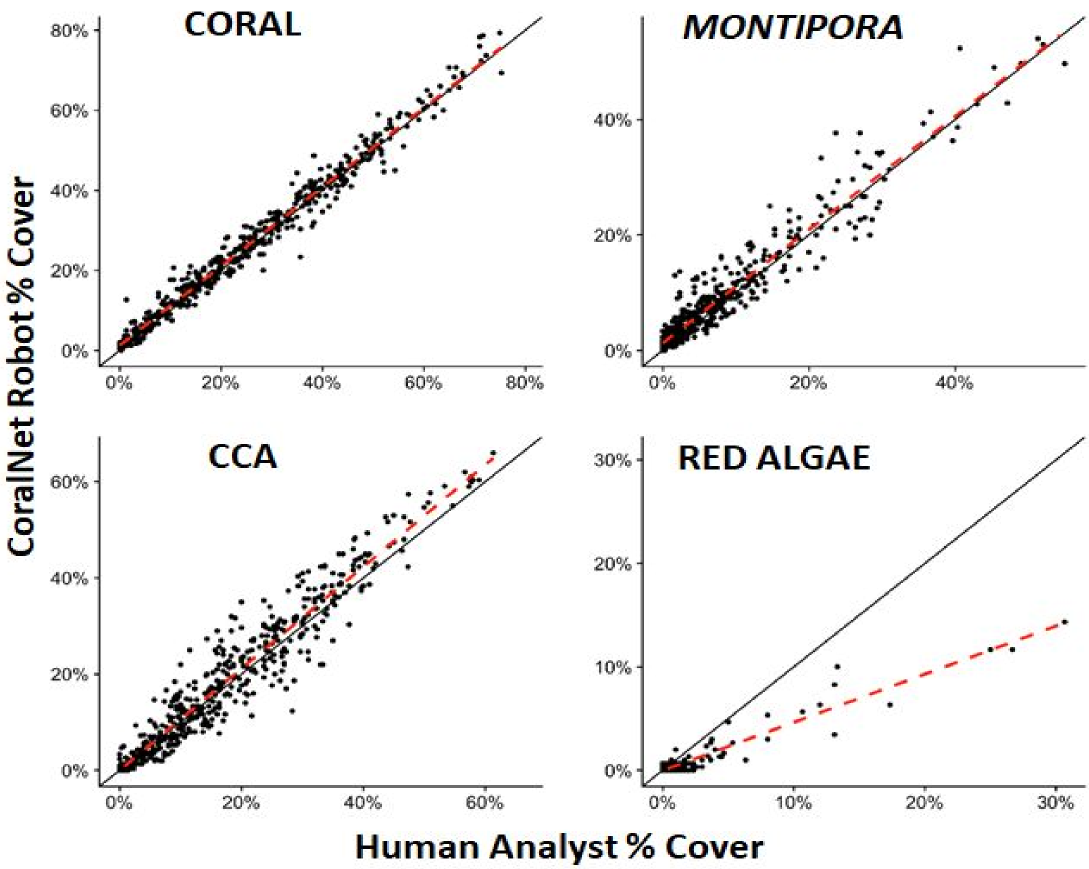

Figure 4-8. Percent cover of benthic features generated by the fully automated CoralNet classifier compared to percent cover from human annotations of the same imagery collected from American Samoa reefs in 2015.

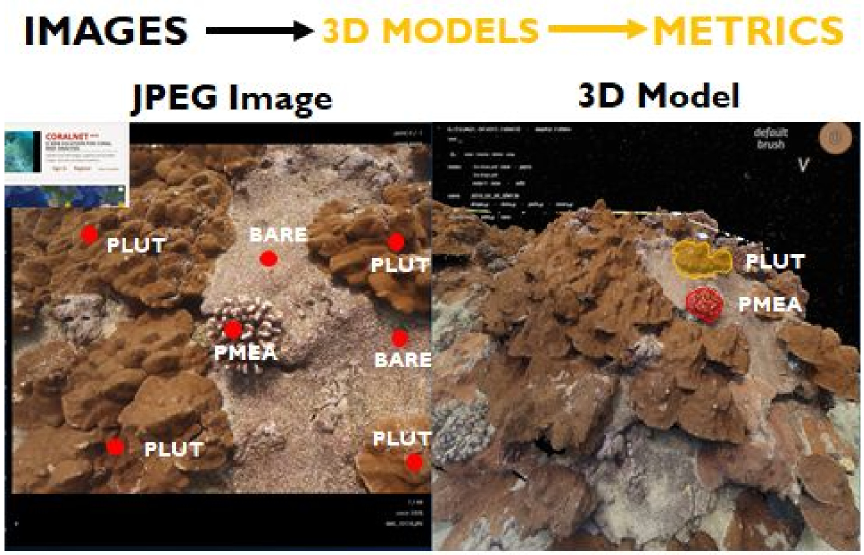

Figure 4-9. A future goal is to overlay CoralNet point annotations on top of 3D models of coral reefs then merge with coral colony segmentations that are generated through a human-assisted machine.

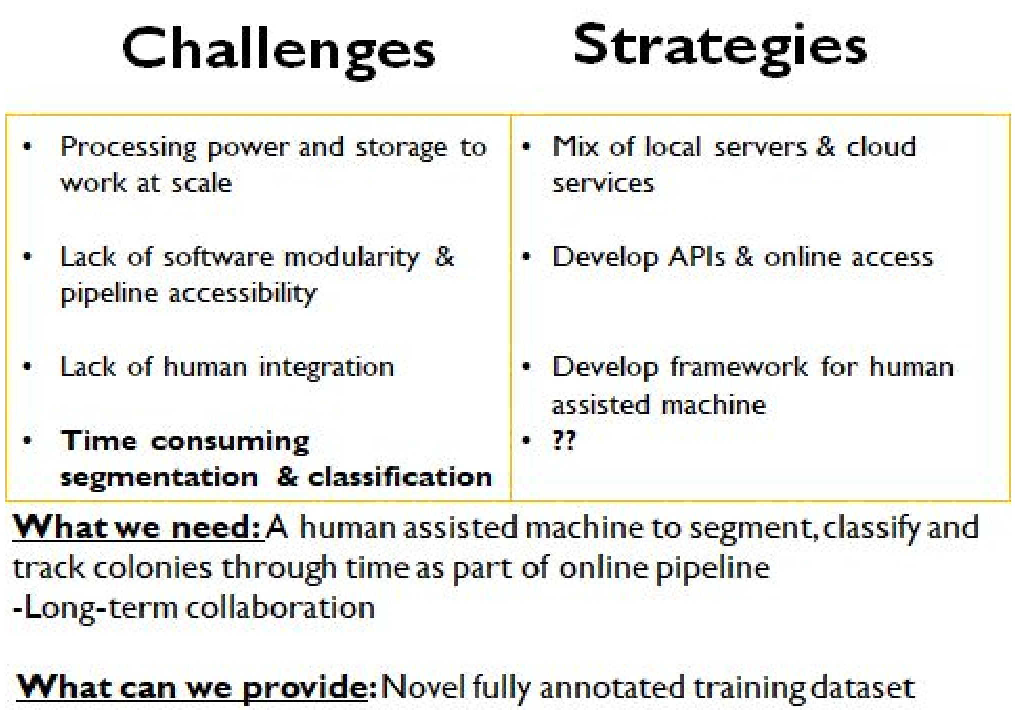

Figure 4-10. Challenges and strategies for scaling up coral reef imagery and extracting timely data from 3D models of the reef.

Discussions: CoralNet was released in 2015, and its use has expanded to about 400 international users. NOAA has invested in CoralNet, and has processed about 20% of its benthic images with this web-based software. This human assisted ML tool has reduced processing time; however, improvements are still needed with processing power, APIs, accurate geo-location coordinates, and improved pipeline (possibly cloud-based architecture) for accessibility to the world community.

The wish list for automated processing of underwater images for the coastal survey programs.

Kim Halvorsen, Tonje K. Sørdalen, Nils Olav Handegard

Institute of Marine Research, Bergen, Norway

Extended Abstract

Underwater stereo video is increasingly being used to monitor coastal ecosystems in Norway. Most surveys use baited remote underwater video systems (BRUVs) to increase observations and hence time needed for manual video analysis, but this method imposes strong selectivity and induces unnatural behavior. On the other hand, unbaited cameras can capture the natural temporal dynamics, and large-scale deployment would produce unprecedented volumes of observations. However, only a fraction of these images can be analyzed manually, greatly limiting the potential advances that can be made from these data streams. Further, the procedures for extracting information may not be standardized between surveys. If we could develop a consistent computer vision analysis framework for identifying, counting, sexing, and sizing the key coastal species, large datasets could be created by interchanging data across surveys, making it possible to address new questions, such as how environmental factors (e.g. season, temperature, depth, habitat, geography) shape fish communities. Moreover, with computer vision, the cost and time used for post-processing and reviewing videos will no longer be a limiting factor – providing a stronger incentive for expanding sampling effort within current projects, as well as establishing new video-based surveys and coastal observatories. Access to a large amount of labelled data is probably the largest bottleneck to reach this, but citizen science is a promising tool to assist us in building large datasets – we present a pilot project that is under establishment where users can capture and label images from a live video stream from a cabled observatory in Lindesnes (Fig. 4-11).

Figure 4-11. Example images of different species in the coastal zone. Automated analysis of these enables large scale deployment of camera-based survey systems.

Automatic interpretation of otoliths using deep learning.

Endre Moen, Nils Olav Handegard, Vaneeda Allken, Ole Thomas Albert, Alf Harbitz, Ketil Malde

Institute of Marine Research, Bergen, Norway

Extended Abstract

The age structure of a fish population has important implications for recruitment processes and population fluctuations, and is a key input to fisheries-assessment models. The current method of determining age structure relies on manually reading age from otoliths, and the process is labor intensive and dependent on specialist expertise. Recent advances in machine learning have provided methods that have been remarkably successful in a variety of settings, with potential to automate analysis that previously required manual curation. Machine learning models have previously been successfully applied to object recognition and similar image analysis tasks. Here we investigate whether deep learning models can also be used for estimating the age of otoliths from images. We adapt a pre-trained convolutional neural network designed for object recognition to estimate the age of fish from otolith images (Fig. 4-12).

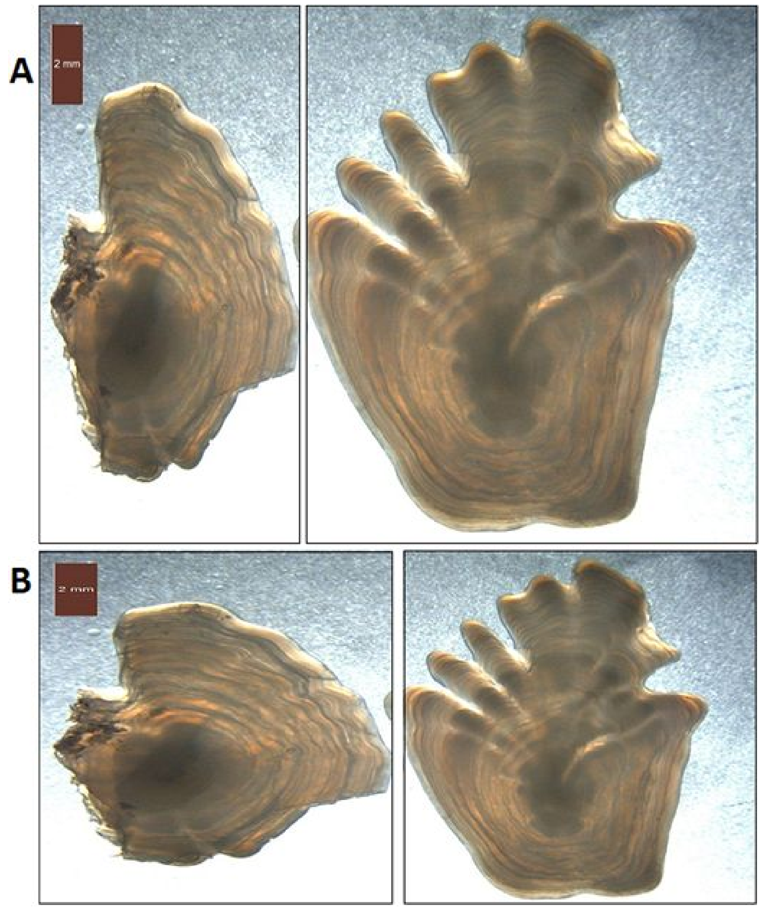

Figure 4-12. A pair of otoliths from 2014 with an estimated age of 13 years. Due to the size difference between the otoliths, the image was split with a substantial offset from the middle (A). There was also a small horizontal overlap causing a fragment of the right otolith to remain in the left image. Resizing causes stretching of the images (B), which is particularly evident in the image of the left otolith.

The model is trained and validated on a large collection of images of Greenland halibut otoliths. We show that the model works well, and that its precision is comparable to documented precision obtained by human experts. Automating this analysis may help to improve consistency, lower costs, and increase the extent of age estimation (Fig. 4-13). Given that adequate data are available, this method could also be used to estimate age of other species using images of otoliths or fish scales. Automatic interpretation of otoliths using deep learning is licensed under CC-BY 4.0 (Moen et al. 2018).

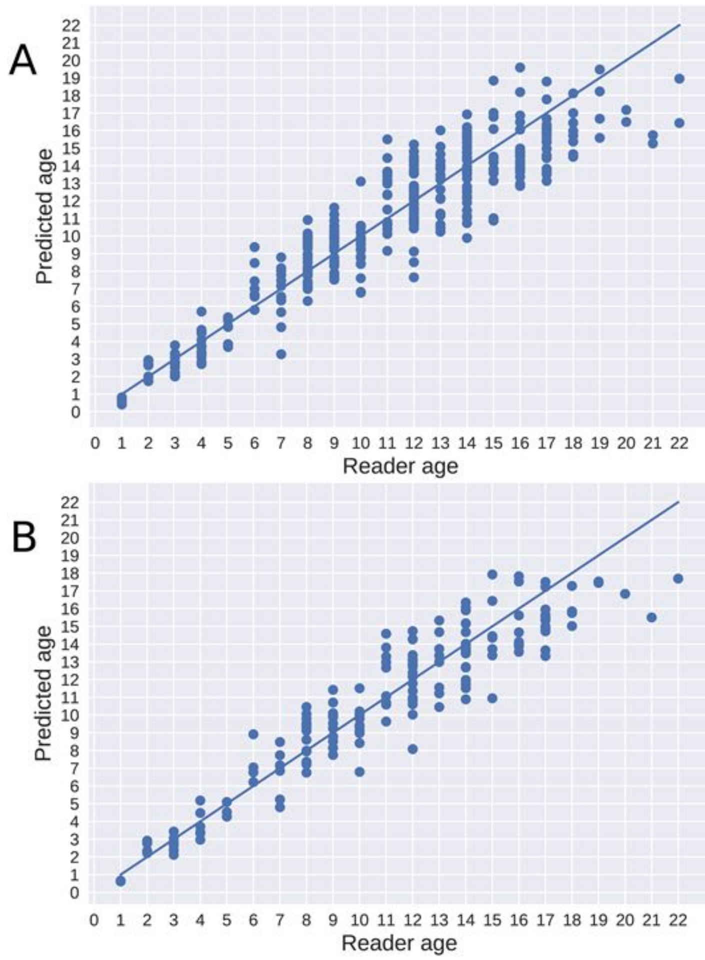

Figure 4-13. Age prediction. Predictions are shown using single otoliths (A) and using the average prediction of each pair (B), compared to the age estimated by a human reader.

Big data sets from unmanned surface vehicles as critical enablers of machine learning insights at scale.

Sebastien de Halleux

Saildrone Inc., Alameda, CA, USA

Extended Abstract

The world's oceans are experiencing significant change, which will have a profound impact on ecosystems, fish stocks, and climate. Furthermore, the areas where some of the biggest changes are occurring are also some of the least measured and understood. This is largely due to their remote location and/or harsh environment, where the cost of deploying sensors via traditional ship-based methods is very high. In response to these factors, new technologies are required to supplement ships and mooring data to meet the demand for longer, more economical deployments with the ability for real-time data, adaptive sampling and shore-based assessment.

Saildrone[15] Unmanned Surface Vehicles (USVs) were designed to meet this need, providing the ability to reach almost any part of the world's oceans, without requiring additional ship time. Deployed from a dock, the Saildrone USVs navigate autonomously to the area of interest, where they operate for extended periods in open seas, before returning to shore for servicing and subsequent re-deployment. A sophisticated suite of on-board science sensors collects climate-quality data at a much greater scale compared than conventional technologies like ships and buoys, enabling these assets to be refocused on high added value activities. As the number of these USVs rises, ultimately numbering in the hundreds, so does the size of datasets they generate, from basic oceanographic to complex acoustic or even eDNA data (Fig.4-14).

Figure 4-14. Saildrone USV on its way to a data collection mission in the Pacific.

[15] Saildrone Inc., https://www.saildrone.com/

This is where advances in automated processing, specifically ML, has the promise to help, transforming big data into new insights, without requiring large increases in headcount, shifting instead the burden to computer resources. However, the challenge in the ML application to ocean data from USVs resides mostly in the need for large amount of training data, i.e. labelled data sets assembled in such a way as to be machine readable.

In the realm of fisheries acoustics for example, new USV-derived big data sets could play a critical role in enabling large scale ML based processing to derive enhanced insights. In the case of USVs, there are three distinct requirements when developing a pathway for public-private partnerships:

Step 1: Survey design - Scientists define the required survey design based on their accumulated domain expertise, including survey tracks, acoustic frequencies, and other methodologies.

Step 2: Data collection - Public and private assets are mobilized to collect acoustic data at scale conforming to the specified survey design. These assets include government ships, charter ships, and USVs. Ships with onboard scientists can tightly control data quality while USVs specialize in cost-efficient raw data collection at scale.

Step 3: Data labeling – Public-private partnerships (such as NOAA's collaborative R&D agreements) can combine scientists' expertise with private sector software tools to label these large acoustic data sets at scale, in such a fashion that the resultant labels are machine readable.

These machine-readable labeled data sets are the crucial enablers of the development of ML algorithms by academic, research institutions, and/or private sectors, which, as they evolve, can both enable processing ever increasing data sets efficiently while reducing human bias and deliver enhanced societal benefits.

5 Case Studies: Machine Learning for Electronic Monitoring

The role of machine learning (ML) in fisheries monitoring: Examples from Alaska.

Farron Wallace[1], Jenq-Neng Hwang[2], Craig Rose[1] and Suzanne Romain[1]

[1] NOAA Fisheries Alaska Fisheries Science Center, Seattle Washington, USA

[2] University of Washington, Seattle Washington, USA

Extended Abstract

In the face of climate change and the consequential impacts to our natural resources, NOAA Fisheries has invested significantly in the development of automated "eMonitoring" tools. Research focuses on development of camera-based systems, sensors, and methods that incorporate the latest developments in machine learning (ML). These systems continuously monitor catch as it comes aboard the vessel or is offloaded at plant, thereby greatly increasing the efficiency and scope of the monitoring process.

Alaska fisheries account for 60% of all landed catch in the US and are monitored by one of the largest Observer programs in the world. Observers gather information from numerous fisheries and plants to support catch accounting, compliance, and stock assessment. Development of remote monitoring alternatives include a chute system to estimate halibut discards by trawl gear, a rail system to estimate discards by longline gear, and a conveyer belt system in the plants to ensure compliance (Fig. 5-1, 5-2, 5-3). For each monitoring requirement, ML algorithms are being developed through use of training datasets collected from the fishery to provide real-time automated image analysis that identify catch by species, measures individual fish, and counts catch events (Fig. 5-4, 5-5).

Our research is leading to development of an improved alternative to currently available remote monitoring systems by providing length/weight measurements, a key data element to estimate total discarded catch and for use in stock assessments. Through integration of machine learning, real-time automated image analyses greatly improve timeliness for extracting data from EM and solve storage capacity issues of storing voluminous video data onboard vessels. Storage costs related to video will also be greatly reduced, since most imagery would not be retained.

These advances will benefit all remote monitoring programs world-wide as the technology is transferable to nearly all fisheries and the machine learning algorithms can be re-trained for any new image data stream (Huang et al. 2016; Wang et al. 2019).

Figure 5-1. Camera chute system deployed on a Bering Sea trawler.

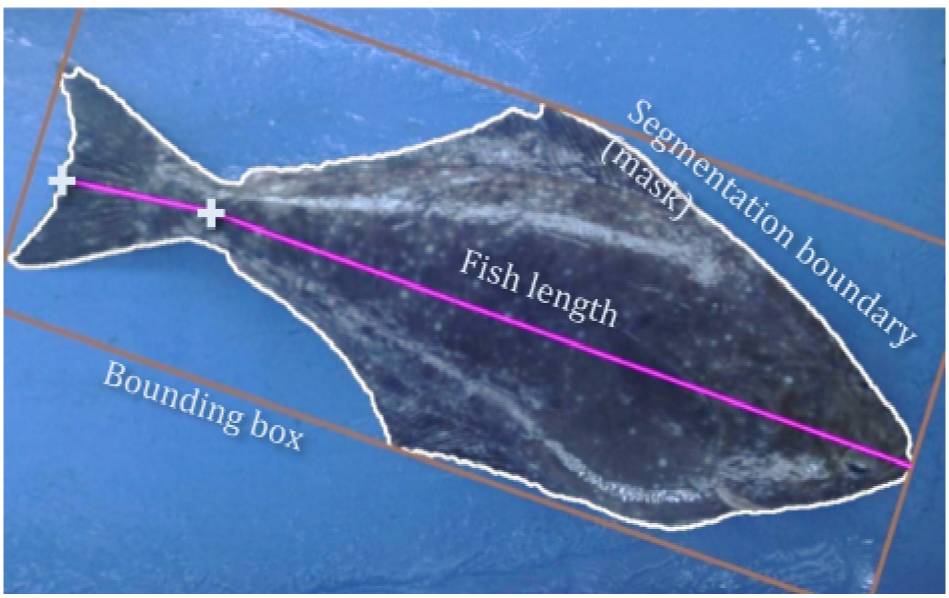

Figure 5-2. Automated segmentation mid-line measurement of halibut in camera chute system.

Figure 5-3. Stereo camera placement (vertical camera system) on an Alaskan longline vessel.

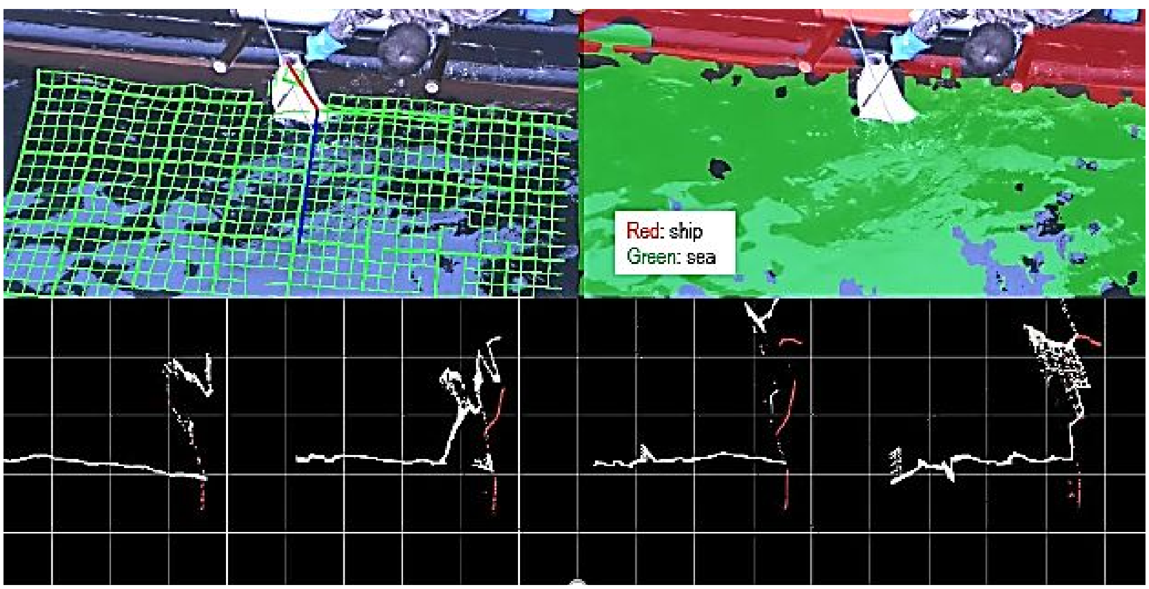

Figure 5-4. Extracting catch events from complex variable back grounds using depth information derived from stereo cameras.

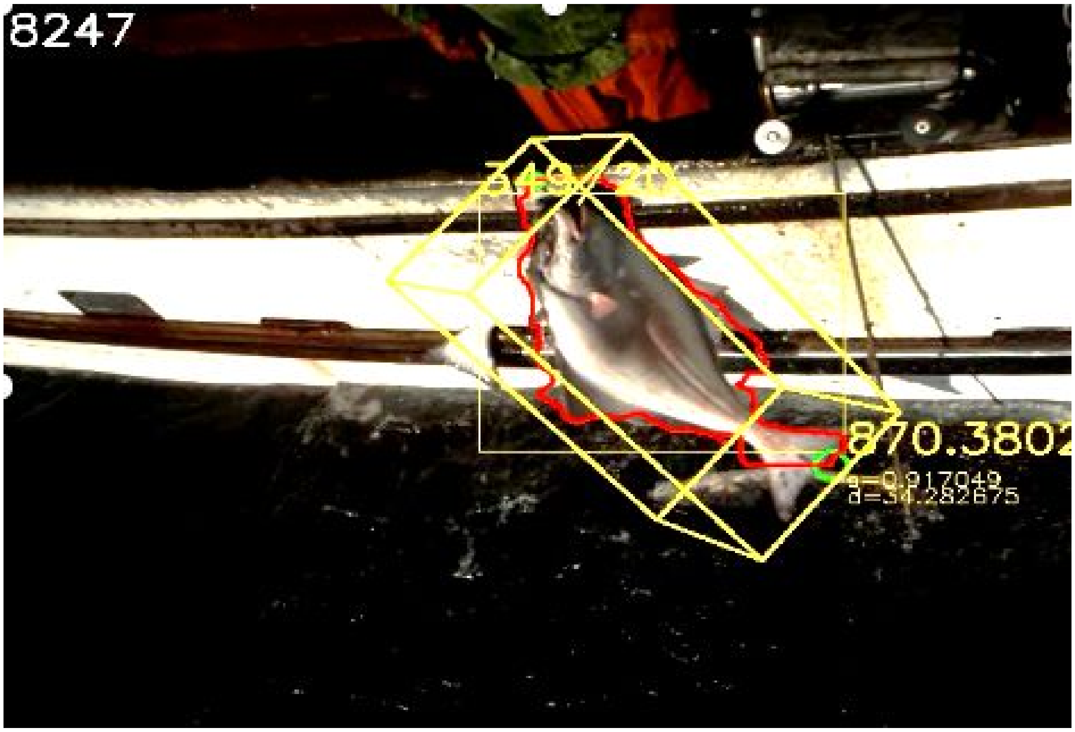

Figure 5-5. Three-dimensional bounding box and depth map to estimate length measurement.

Discussions

Lack of data availability and access is a significant challenge to leverage a wide variety of recent machine learning developments from various agencies, NGO, and commercial entities. We support advancement of a framework to foster development of machine learning algorithms through creation of an annotated data access model that can be used to share non-confidential data. This dataset coupled with development of an image analysis tool box of machine learning algorithms will provide researchers significant leverage from the considerable investments in current (and future) image analysis using machine learning algorithms including likely contributions from academic and non-governmental partners.

Electronic monitoring from fishing vessel operations:

The need for and progress in automating the video review for fisheries electronic monitoring footage.

Nichole Rossi[1] and Ben Woodward[2]

[1] NOAA Fisheries, Northeast Fisheries Science Center, Woods Hole, MA, USA

[2] CVision AI, Medford, MA, USA

Extended Abstract

Electronic monitoring (EM) is being used to revolutionize the way that fisheries dependent data is collected. EM systems, which incorporate multiple cameras and sensors to produce a record of fishing activity, are being used as a monitoring tool in a wide variety of fisheries across the globe. In the Northeast US, investigations of the potential utility of EM systems by the agency have been ongoing since 2010. These have included efforts to use EM for both catch estimation and compliance applications. In these pre-implementation programs technology from a variety of EM companies has been tested and information on the efficiency of these systems has been researched. One salient point found in each application of EM has been that human video review is labor intensive and thus expensive and program expansion inhibited.

As the region moves towards industry-funded monitoring and increased coverage across multiple fisheries, there is a need to address programmatic concerns such as video review. As a result, the agency is investigating incorporating efficiencies such as machine learning advancements to automate the processing of EM video data, thus making the video review process more efficient and reducing related program costs. To that end, a set of projects spearheaded by CVision AI[16] have leveraged existing EM video data to train algorithms for automated fish identification, fishing activity detection, and length/weight estimation (Fig. 5-6). Additionally, data science competitions have been organized to help develop creative algorithms with which to process video data. A number of other efforts in this vein are also underway including the development of open source video review software and data collection in support of machine learning. We briefly detail these efforts and discuss the path forward in machine learning based EM video processing.

Discussion: Electronic monitoring is expanding rapidly and the human manual processing of imagery data will not meet the requirements for immediate turn-around for policy decisions. In addition to the big data bottleneck, there is need for improving the data collection protocols to ensure standardized metadata with annotations and activity recognition.

[16] CVision AI, http://cvisionai.com/

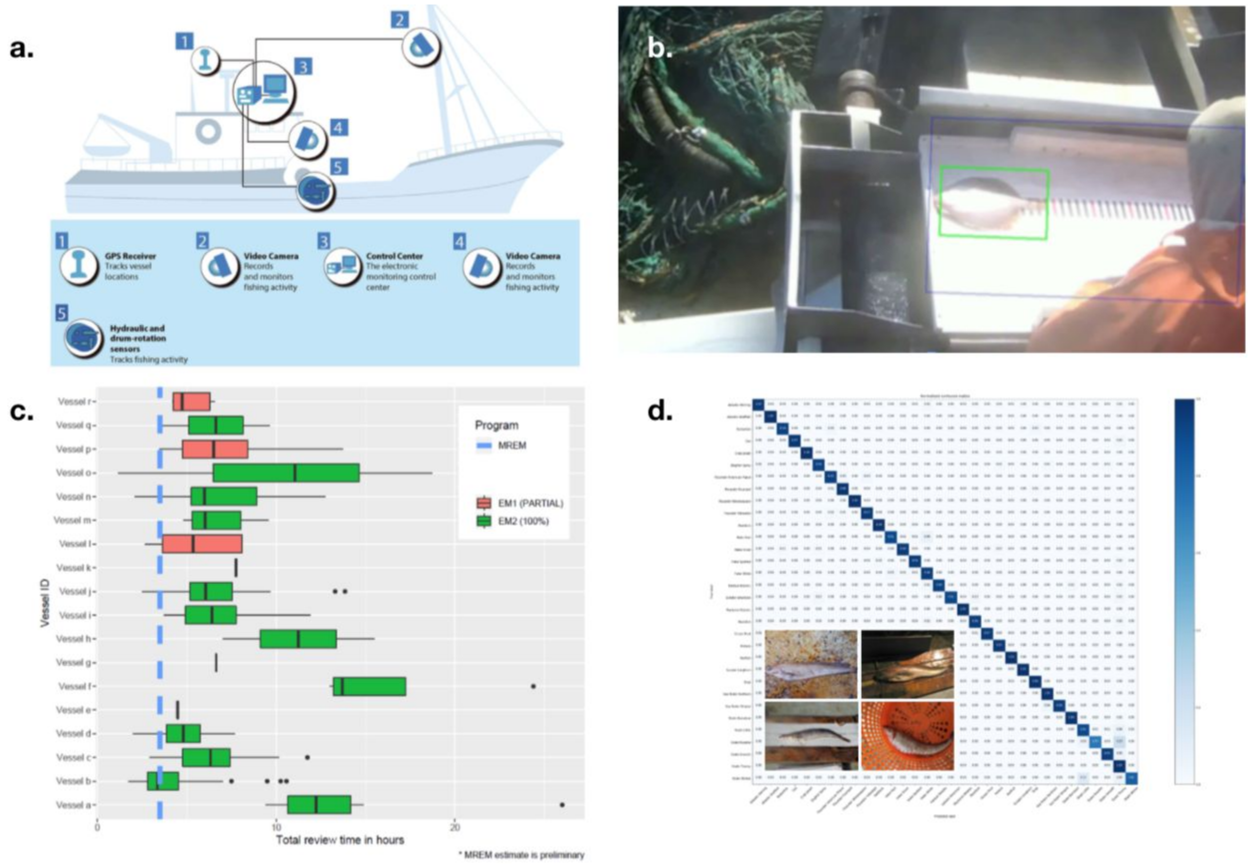

Figure 5-6. Applying Machine Learning to Fisheries Monitoring. Electronic monitoring is expanding in New England and other regions of the U.S. With this expansion comes the need to build more efficient methods for processing EM data. Modern machine learning methods that leverage cutting edge deep learning models show great promise and are actively being explored in New England's Multispecies Groundfish Fishery. **a.** A depiction of an electronic monitoring system. **b.** An image from CVision AI's Open EM software. **c**. A figure showing the estimated review times for a set of EM vessels. Generally, review time is lengthy, between five and ten hours. **d.** Results of a classification model showing that ~30 species were well handled by modern machine learning methods. Subset shows representative image from the Northeast Fisheries Science Center Observer Program.

Automated fish analysis aboard research vessel operations: Building a library for image processing and machine learning to support electronic monitoring programs.

Andy Jones[1] and Ben Woodward[2]

[1] NOAA Fisheries, Northeast Fisheries Science Center, Woods Hole, MA, USA

[2] CVision AI, Medford, MA, USA

Extended Abstract

One consistent challenge for those interested in developing computer vision methods is the initial difficulty in collecting sufficient training data. To begin to develop machine learning methods in the Northeast US we have sought to overcome this hurdle by leveraging existing scientific collections of biological data to build training data sets (Fig. 5-7). We recently initiated a project with CVision AI[(17)] to collect images from the existing biannual bottom trawl survey. This survey, which has been conducted by the Northeast Fisheries Science Center for > 50 years, aims to collect fish using short (~20 min) bottom trawl tows from > 375 stations that stretch from Cape Hatteras to the Gulf of Maine. Catch includes all of the commercially important species from the region. From this catch a wide variety of biological data including the length, weight, sex, and stomach contents are collected from the catch, and stored in NEFSC databases. Therefore, this survey represents an untapped resource for developing training data for machine learning algorithms (collections of paired biological data and images).

To begin to develop these paired collections, CVision AI outfitted the NOAA fisheries survey vessel Henry B. Bigelow with a camera system and began collecting images in the fall of 2018. Following the completion of the cruise legs and auditing of specimen data, video annotations were made for a list of priority species. Here we discuss our progress on this project and the potential application of this type of a program to other regions with ongoing survey collections.

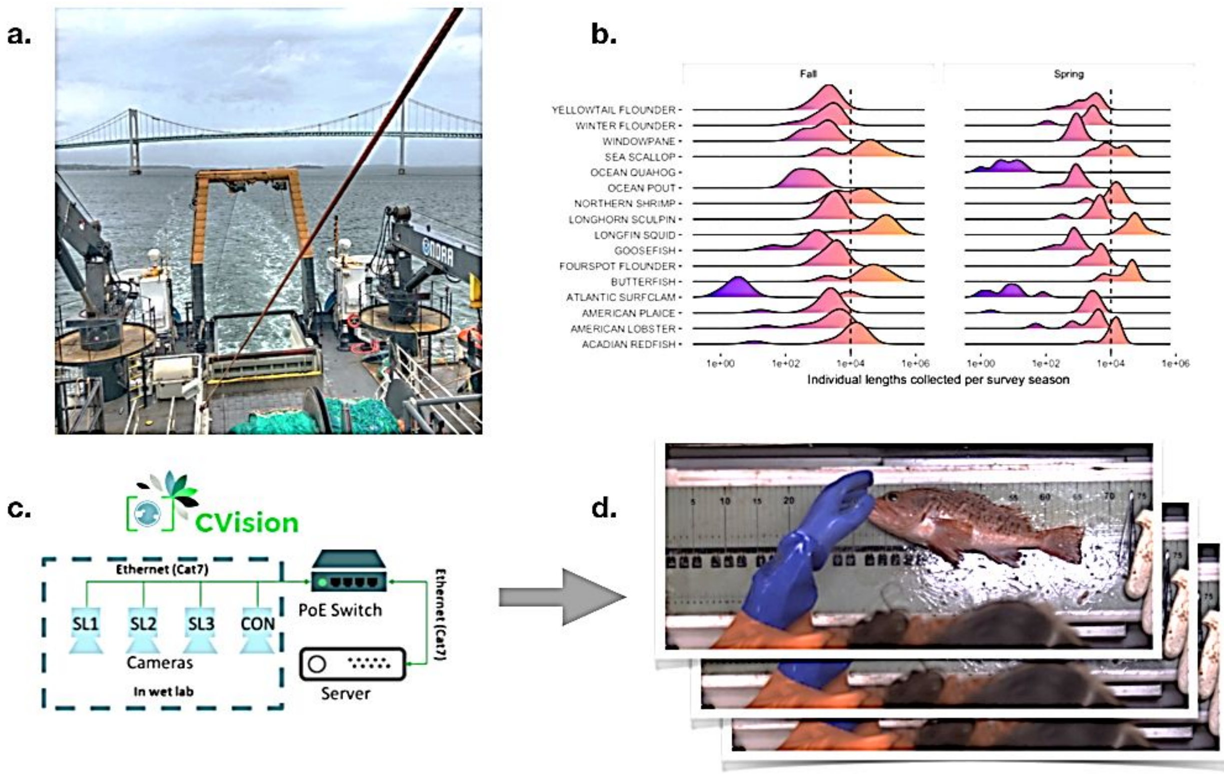

Figure 5.7. Building a Library by Leveraging Survey Operations. Developing machine learning solutions for fisheries problems demands large training data sets. Leveraging existing survey operations may allow for the rapid development of a high volume of annotated image training data. This concept is being trialed in the Northeast Bottom Trawl Survey which processes a large number of individuals each year. **a.** A picture of the stern of the R/V *Henry B. Bigelow*, the research vessel for the bottom trawl survey. **b.** An example of the number individuals processed per year by the fall and spring bottom trawl surveys. Data are derived from surveys 2010 - 2017. A subset of species is shown. **c.** A schematic showing the components of CVision's system for image capture system. **d.** An example of images captured by the system.

Catch identification and analysis – challenges at the Norwegian Directorate of Fisheries.

Roger Fosse, Kine Iversen, and Atle A. Øinas

Norwegian Directorate of Fisheries, Norway

Extended Abstract

The Directorate of Fisheries shall promote profitable economic activity through sustainable and user-oriented management of marine resources and the marine environment. Monitoring, control, and surveillance are essential in achieving this objective. We have jurisdiction over a large area and have limited control resources. Our goal is to develop automatic risk-based systems to monitor and control the fishing industry at sea and on land. We collect data through the value chain in the fishing industry and gain information from a number of data sources, such as Vessel Monitoring Systems, Electronic Reporting Systems, and sales notes (Fig. 5-8). These data can be used in data analysis to perform more targeted controls. The challenges we face are related to consistency, quality and coupling of our data sources, as well as to develop good training data sets for the potential use of machine learning.

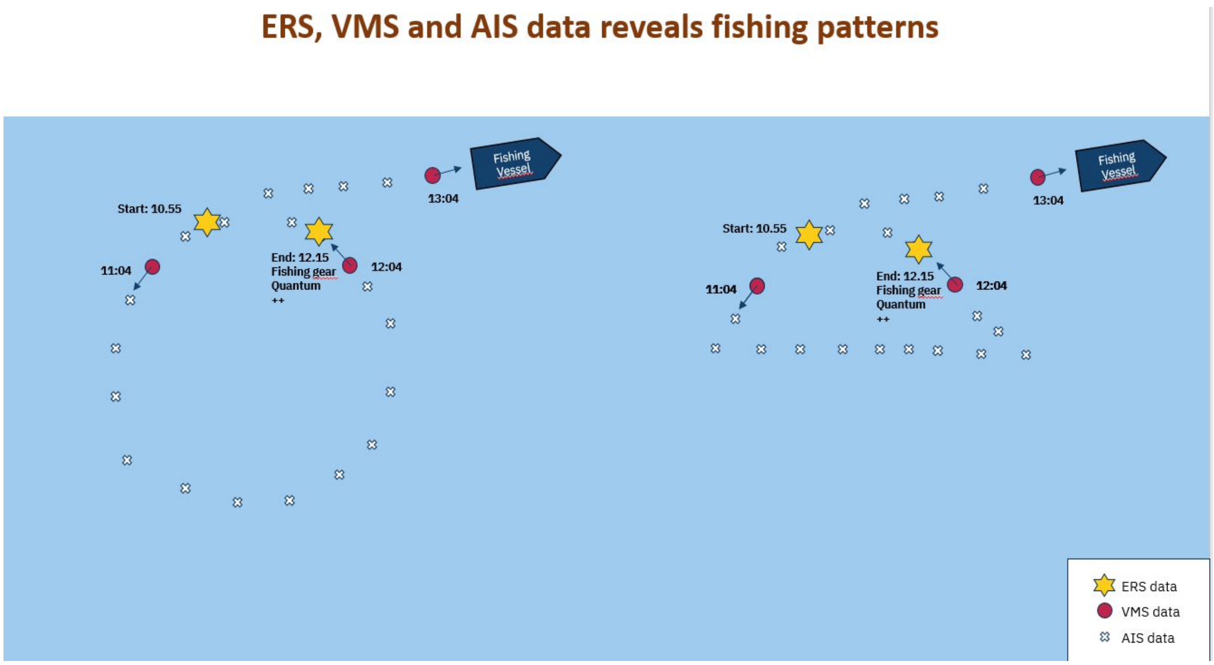

Figure 5-8. Vessel Monitoring Systems and Electronic Reporting Systems are used to gain information from a number of data sources through the value chain in the fishing industry.

Discussion: The monitoring of fishing vessel activity can be automated using ML; however, the challenge is to establish more consistent standards when collecting information from our data sources.

6 Case Studies: Machine Learning Methods

Collaborative machine learning through expert knowledge – Implementation.

Håvard Vågstøl

Scantrol Deep Vision AS, Bergen, Norway

Extended Abstract

While textbook examples of classification models typically begin with a fixed size training set and end with prediction results on a test set, the Deep Vision Bio Base needs to be a dynamic system. How strongly should the influx of new data influence the model? Which is the best way to allow for new species to be classified by the model? How do we best ensure that we do not break the model accuracy on older data? In this panel discussion we will discuss goals, pitfalls and solutions for a collaborative, online solution with the right adaptability and accuracy to meet the requirements of Deep Vision Bio Base.

Challenges to be solved with the implementation of the Deep Vision Bio Base system:

Data storage and organization:

- Data format for annotations, masks, metadata
- Data storage
- Data access for training
- User access to models

Data ownership:

- Who can do what with the data?
- Who might have an interest in keeping their data secret?
- Who might have an interest in using the data and annotations?

During cruises:

- Each user input should improve future predictions, without extensive retraining of a classifier.
- New classes found during research cruises will only have a small number of examples.
- Could we combine a backbone network from the DVBB with a local network based on online learning?

We can retrain a classifier in the Deep Vision Bio Base, but we need to be careful not to reduce the prediction accuracy for existing classes. Could we make an architecture that enables us to proceed with training on the basis of the existing weight, while also enabling us to introduce new classes? What machine learning architecture would be the most appropriate?

Discussion: Data annotation logging software is a critical requirement for developing training data sets for ML classifiers, but the question was presented on whether the annotated data should be open or owned by the private sector. The private sector would be concerned with how to maintain the data and/or classifier, while scientists would be concerned with avoiding the 'black box' classifiers. This brings to question the ML requirements for scientists conducting research and development, in comparison to the requirements for those conducting more routine survey operations.

Underwater Imaging: from color calibration to machine learning.

Rafael Garcia[1,2]

[1] University of Girona, Spain

[2] Girona Vision Research, Spain

Extended Abstract

Underwater imaging needs to deal with several limitations imposed by the medium. The interaction between the light and the aquatic environment includes basically two processes: absorption – where light is gradually attenuated and eventually disappears from the image-forming process, and scattering – a change in the direction of individual photons, mainly due to the various particles suspended in water. These transmission particularities of this medium result in additional challenges in underwater imaging, such as blurring, limited range, color shift, clutter, non-uniform illumination, and "marine snow" due to suspended particles. In this talk we will describe the work carried out by the Underwater Vision Laboratory(18) of the University of Girona in Spain to understand the underwater image formation process. We will address the peculiarities of underwater imaging in the scenarios of seafloor mapping and fish detection, and we summarize the results of using machine learning on those underwater datasets (Fig. 6-1).

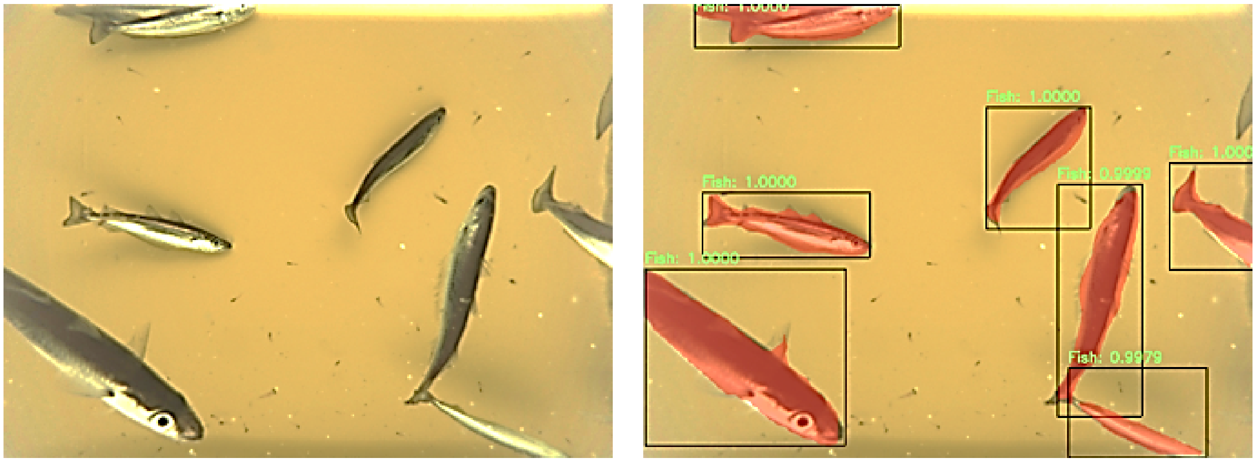

Figure 6-1. Fish detection and semantic segmentation using a Mask-RCNN architecture for the images acquired by the Scantrol Deep Vision system. The network has been trained using 1,564 images from a different trawl.

Discussion: Color correction for underwater video and image data is a physics problem in regard to the light attenuation and scattering through the water, and is not an ML problem. ML is important for object detection and classification, and can only help with color correction if you have training data sets for conditions that distort the light, such as changes in turbidity of the area and season, and the daylight cycle.

A data management platform for autonomous marine measurements.

Jeremy Cook

NORCE Norwegian Research Center, Bergen, Norway

Extended Abstract

We have developed a layered data management platform for data acquisition from autonomous marine measurement platforms. Real-time data is streamed, through satellite communications, to ground stations where it is processed. Bulk data (hydrophone and echo sounder) cannot be transferred by satellite due to bandwidth and data quota limitations, so it is uploaded when the vehicles are recovered (Fig. 6-2). The database of time-series and bulk data can be analyzed with interactive visual analysis tools.

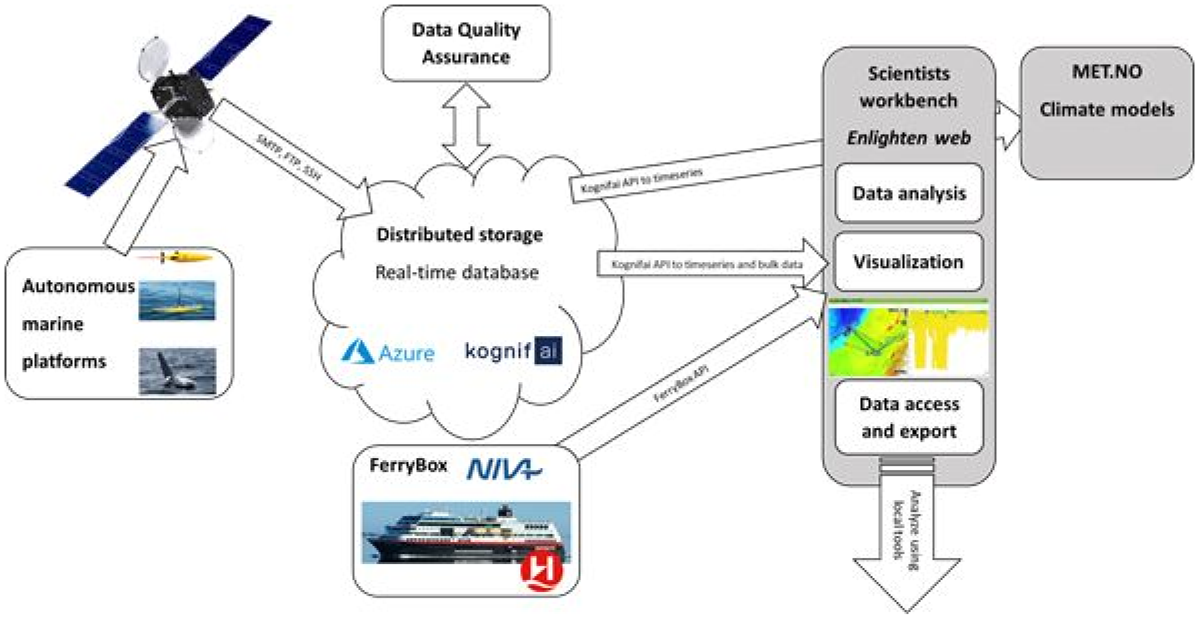

Figure 6-2. The system provides convenient access to survey data in real-time as well as providing scientists with advanced tools for analysis. The user accessible front end is written using W3C web standards and will run on many platforms including a mobile platform solution. Key parameters for survey operations include: longitude, latitude, depth, current, salinity, temperature, depth, O_2, fluorescence, audio, CO_2, wind speed, wind direction, and pressure.

The portal has been on-line for 2018s survey and continues with 2019 surveys (started March 2019. The Enlighten Web visualization solution handles large data volumes – adaptive data visualization based on load (Fig. 6-3). NORCE[5] partnerships (Table 6-1) enable the system to conduct data acquisition from multiple platforms and data analysis of multiple parameters and platforms simultaneously.

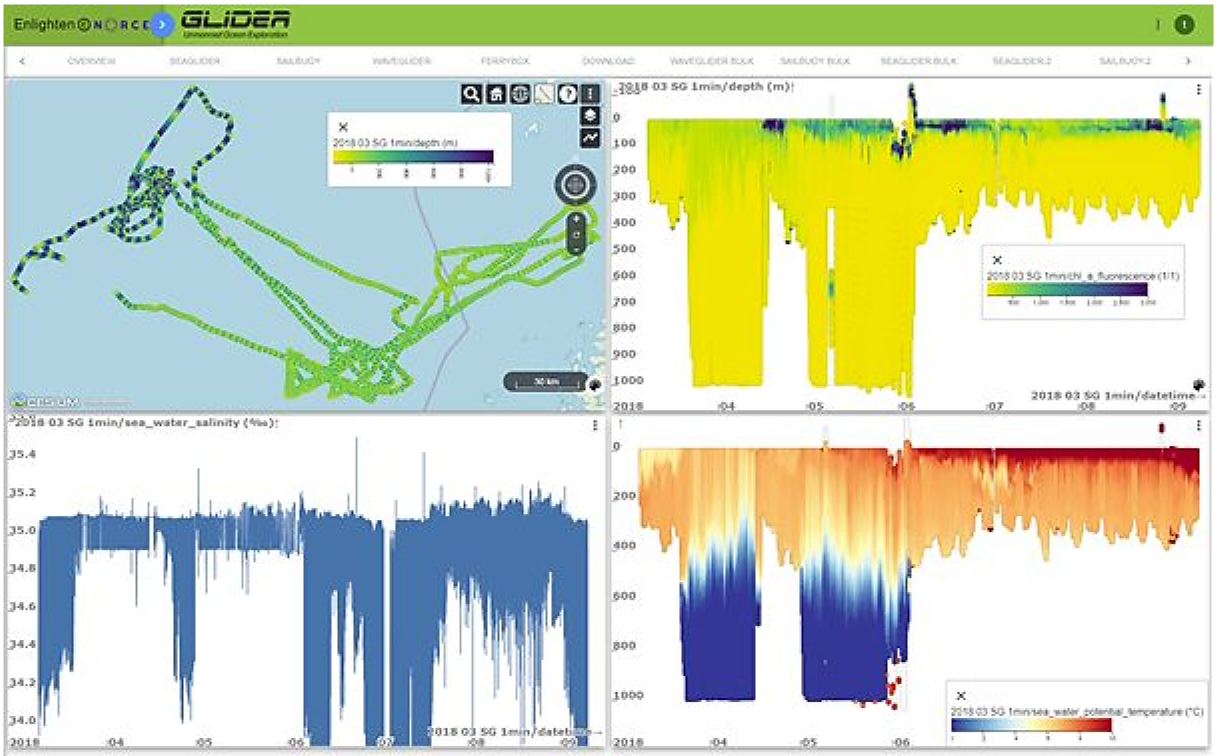

Figure 6-3. The portal is developed using Enlighten Web, developed by NORCE over several years. Data visualization for millions of data points is possible. The figure below illustrates visualization of survey parameters for the entire 6-month survey period.

Table 6-1. NORCE's partnerships enable the system to carry out data acquisition from multiple platforms and data analysis of multiple parameters and platforms simultaneously.

Kongsberg Maritime / Digital	Subsea development glider commissioning, operation & piloting, echo sounders, sensors, data storage services, K-Lander, business development
Offshore Sensing	Sailbuoy commissioning and operations, piloting
NIVA - Norwegian Institute of Water Research	Sensor selection, implementation and data transfer
NORCE Technology	Marine acoustics, data integration and analysis, visualization
Maritime Robotics	Wave glider operations, piloting

Aanderaa Data Instruments	Sensor selection and field coordination
UiT The Arctic University of Norway	Ecological data interpretation
Met.no - Norwegian Meteorological Institute	Data assimilation into ocean models

7 Training Session: Tools for Machine Learning

Tutorial: GitHub and Continuous Integration.

Ibrahim Umar

Institute of Marine Research, Bergen, Norway

Extended Abstract

Version Control System (VCS) are one of the most useful tools for computer programmers. VCS enables users to safely records a file's or set of files' "history" (or their changes over time) and allow users to retrieve a specific version of a file or set of files within their recorded history. VCS also enables easier comparison between sets of changes and is far more convenient than having to keep several different versions of the files. Today VCS software comprises more than just history keeping, but has evolved to support distributed "social coding" collaboration, continuous integration (CI), and continuous deployment (CD) pipeline.

In the tutorial (F1) we will learn and practice the basics, advantages, and use-case examples of VCS followed by hands-on in using GitHub (www.github.com), which is one of the most prominent and feature-rich VCS software applications available. The tutorial will cover several topics from the Git basics to the usage of GitHub for CI and CD pipeline, where we are going to try making a multi-platform and automated testing/deployment pipeline for our own program (Fig. 7-1).

For the training session (F1), please bring your own laptop and register with these web services before the tutorial session (if you haven't done this already):

- GitHub (https://github.com/join, do this first)
- Travis-CI (https://travis-ci.org/, you have to use your GitHub account to sign up)
- Appveyor (https://ci.appveyor.com/signup, you have to use your GitHub account to sign up and select "FREE -..." on the "plan" drop down menu)

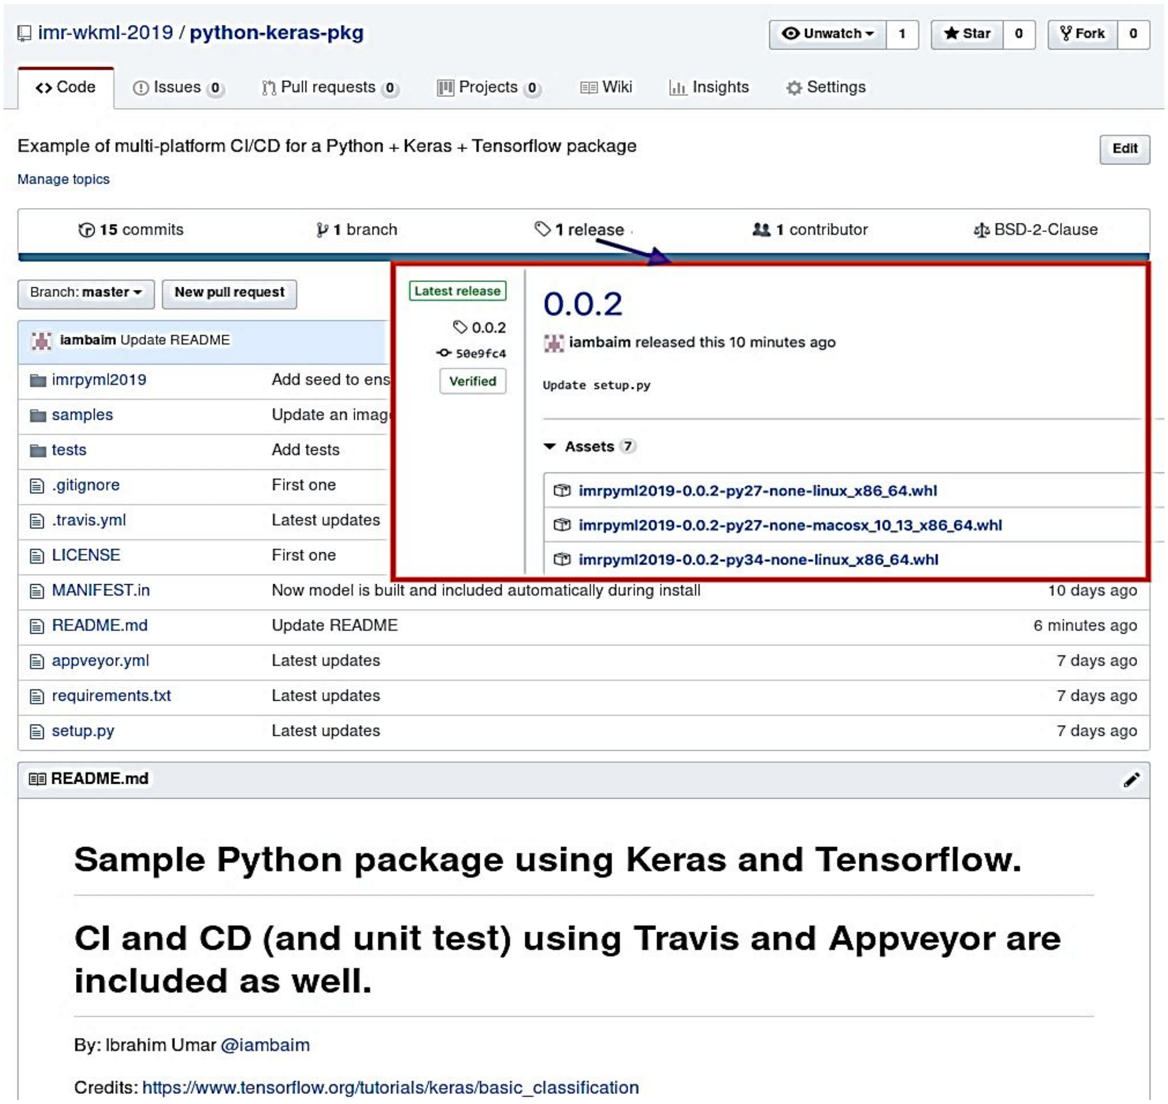

Figure 7-1. Screen capture of a sample Python package using Keras and Tensorflow.

Installation and hands-on training with the open source Video and Image Analytics for Marine Environments (VIAME) software.

Matt Dawkins

Kitware Computer Vision, Inc., Clifton Park, NY, USA

Extended Abstract

This is a hands-on training session with the Video and Image Analytics for Marine Environments (VIAME) toolkit including software installation, running on novel data, detector model refinement, software integration, and other topics as desired. Our preliminary training session is as follows:

- System Overview and Software Installation
- Running Existing Detectors and GUI basics
- Training New Detection Models via Manual Annotation
- Rapid Model Generation for Detectors or Frame-Level Classifiers
- Comparing Detectors, Archive Summarization, and Aggregate Statistics
- Auxiliary features: Stereo Measurement, Image Registration
- Planned Future Work and GUI Development
- Experimentation, Questions, Discussion

Bringing a laptop and video or image data specific to your problem (if available) is helpful, otherwise surrogate data will be available. Installers will be provided on the workshop day, though it can be beneficial to download them ahead of time from https://github.com/Kitware/VIAME to confirm functionality on your laptop. A new software release will be made a week before the workshop. Feel free to contact viame.developers@gmail.com with any questions.

8 Conclusions

8.1 Applications in Fisheries and Ecosystem-Based Marine Science

The marine science community is eager to apply the rapidly evolving ML tools to a wide range of cross-functional missions relevant to the sustainability of living ocean resources. The goal of this workshop was to examine progress with case studies to understand the requirements for transitioning ML research into operations. While most of the case studies strive to increase the accessibility of data for discovery and forecast assimilations by the wider scientific community, other case studies have specific operational objectives with confidentiality constraints on the data accessibility. For this reason, the case studies discussed in this report are divided into two ML application categories:

Surveys and ocean observations:

Most of the case studies described in the report apply ML analytics to improve the processing and analysis of various data types collected from aerial and underwater surveys and ocean observation operations. These data tend to be the largest in volume and are often used for abundance estimates, stock and habitat assessments, and analysis of uncertainty in the estimations. In this case, data accessibility for ML will be open to the wider scientific community for discovery, time series analysis, and predictive forecasting. For this category, optimizing the data storage and accessibility with cloud resources platforms for the ML applications by the public provides added value to the data enterprise. The case studies herein examine ML applications for various data types such as:

- **Imagery:** Most of the ML case studies provided examples of successful automated object detection and classification for imagery data. Plankton images with well-defined features from video recorders provide an opportunity to conduct ML research to refine ML methods, while ML applications for underwater and aerial visual surveys for fish and marine mammals tend to be more challenged with the quality of imagery collected from environments impacted by light, turbidity, and other environmental variability. Electronic monitoring is also impacted by environmental factors. Object detection and classification is easier when the targets occur in less complex backgrounds (e.g., fish in water column); therefore, it is important to train with targets in more complex backgrounds (e.g. fish along the seafloor bottom). Another example of a more complex application of ML for imagery is the pixel-based benthic habitat characterization (refer to CoralNet example in Section 4).
- **Active acoustics:** Large volumes of acoustic backscatter data are collected from fishery surveys and ocean observation systems, and there is considerable interest and benefit with the application of ML to reduce manual processing time. Unlike imagery data with distinct features that can be used to classify to the species taxonomic level, acoustic surveys rely on

concurrent other technologies and sampling to interpret the acoustic backscatter by species. Therefore, the immediate benefits of applying ML are to reduce the time and human bias associated with the laborious post-processing of active acoustic data.

- **Passive acoustics:** An example was presented with the successful application of ML using cloud resources for the classification of marine mammal songs from passive acoustic data. The success can be attributed to the development of training datasets for the species-specific songs as well as the false positives, and the partnership with the cloud expertise.
- **Other data types:** There is recognition that the ML applications for various environmental data are beneficial to providing higher quality and more timely scientific products for the sustainability of living ocean resources. For example, ML can be applied to 'shotgun sequencing' (e.g., meta-genome assembly or meta-transcriptome assembly) or 'amplicon sequencing' (PCR meta-barcoding) of environmental DNA (eDNA) used to detect species in the marine environment. Additionally, ML applications for data fusion and assimilation from multiple observation systems holds considerable promise for improving forecast modeling in marine ecosystem and earth science.

As these data become more readily accessible with training datasets for ML applications and discovery by the wider scientific community, economic benefits will be realized through significant cost savings in data processing, improved estimates and forecast modeling, and more timely scientific products for ecosystem-based management and policy decisions for healthy oceans.

Electronic monitoring:

Electronic monitoring herein is in reference to data collection aboard commercial and recreational fishing boats to provide scientific information, such as catch per unit effort and bycatch estimates for a specific region that can be used in harvest control regulations (Section 5). Vessel monitoring systems (VMS) and imagery data from camera systems are two examples of data collection from fishery-dependent sampling. These data have confidentiality restrictions that prevent open access, and will not be used by the broader community for discovery or forecast modeling using ML. Furthermore, the development of training datasets will likely be specific for the regional fisheries and catch. Therefore, data accessibility, training datasets, and ML computations will likely be best accommodated with regional on-premise servers. Overall, ML applications for electronic monitoring data will result in cost savings from reduced processing time and more timely scientific information for regional fisheries harvest control regulations.

8.2 Data Quality and Accessibility

The case studies presented during the workshop emphasized a common theme on the need to improve data accessibility with enriched metadata and enhanced data enterprise architecture to utilize machine learning analytics. As the working group addressed the following question:

What should the data enterprise consider for increasing the accessibility of enriched annotated metadata for ML analytics?

This question focused discussions on two priorities; the need for enriched metadata, and enhanced storage/access to the data for ML computations.

Metadata and annotations: The application of ML typically requires large volumes of high-quality data with appropriate metadata and annotations. Data logging routines to annotate event activities and other features are important for the descriptive and structural metadata (Margolis et. al. 2019). Data collection formats can vary between regions, such as video data formats; therefore, awareness on recommended formats can be improved. There were some that suggested some degree of flexibility is needed to allow the user to define the metadata and logging requirements. Overly prescriptive mandatory fields with the intent standardizing formats can impose an administrative overhead that effectively hinders data submission, and at times cause difficulties in the implementation of ML tools. Data and metadata standards should therefore be designed in close collaboration between data producers and users, to ensure all needs are adequately met. Some participants recommended to keep things simple based on the necessary metadata standards to prevent too much overhead on the data management side. Administrative metadata is necessary for data sharing policies that consider security, ethics, legal, and societal implications. Ongoing efforts to modernize the data enterprise with metadata and annotations should consider not only the original operational and research objectives of the data collections, but also consider the scientific products derived from research and discovery by the broader community using ML.

Storage and accessibility: In most situations where larger databases are collected from surveys and ocean observation infrastructure, much of these data reside on on-premise servers, and there were discussions about the need for more storage capacity and the benefits of improving data accessibility using cloud resources. Cloud platforms are increasingly used as cost-effective solutions for storage, accessibility, and ML tools. The cloud platform provides effective data access to the wider scientific community for discovery, and in this situation, creating libraries of training datasets for ML is also recommended. A recommendation from the workshop is when large volumes of data (e.g., imagery, active and passive acoustics, and other environmental data) are stored, special attention regarding storage and accessibility requirements is needed. There was also recognition that in some cases, hybrid architectures can optimize the use of on-premise and cloud resources to improve storage and accessibility of data, access to libraries of training datasets, and end-to-end workflows using cloud ML computing. As previously mentioned, some types of data (e.g., electronic monitoring) have confidentiality restrictions and produce regional-based products. In this situation, an access control system is needed with appropriate security levels and regional-based servers for data storage, restricted accessibility, and regional-based training data for ML applications.

8.3 Training Data Sets and Model Performance

ML provides cost-saving benefits by significantly reducing data processing time and delivery of timely scientific products; however, the integrity, trustworthiness, and quality of the scientific output from ML is the overriding goal. The end users, such as assessment scientists, require a rigorous evaluation of the model performance with quantifiable uncertainty for the scientific products used in policy decisions. Deep learning algorithms such as neural networks can approximate any function; however, performance is highly dependent on the quality and size of available training sets (Hornik et al. 1989). The limited interpretability of the models may also limit the usefulness in some cases. Fitting a model to a training set does not assure good prediction performance. In situations where the model is complex with a number of parameters or the datasets are small, overfitting the model to a particular training set can result in low predictive performance for new data. There are bias-variance trade-offs where a simpler model is less prone to overfitting, but also less likely to capture complex relationships between input and targets.

The goal is to have good predictive performance of the ML on the input data, which is called generalization. To determine the generalization error, a portion of the training set should be reserved for independent testing of the model performance. A good overview on training deep learning algorithms is provided by Aggarwal (2018). To ensure that the model generalizes well to new data, it is common to set aside a separate validation data set. Models are trained using the bulk of the available data, and then the final model is selected based on performance on the validation set. Since the validation set is used in the selection, another previously unused test data set is used to estimate the performance of the final model (Fig. 8-1).

.

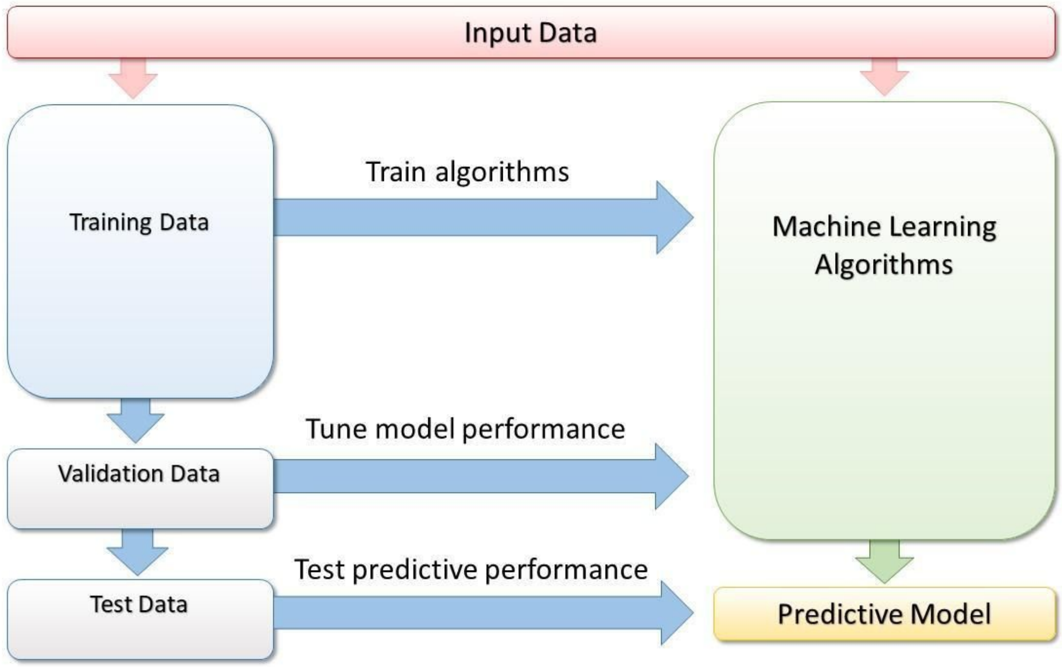

Figure 8-1. The success of applying machine learning is highly dependent on the quality of the data and labeled training data. A portion of the training data in used for independent validation and testing to determine the predictive performance of the model.

There was consensus among the workshop participants that more and better labeled training sets to train, validate, and test the predictive performance of ML deep learning algorithms is highly needed. While classifying and labeling datasets is time consuming, the need to produce trustworthy output for reliable and credible science is unquestionable the highest priority. As the data enterprise architecture is modernized for effective data accessibility for ML computations, pooled annotative libraries with labeled training data must also be made available. It is also important to note that the training data collected from a specific time and location might provide different results from another time and location. Therefore, routine updates to training and test datasets are important. Researchers should always have a thorough understanding of the data, training data, ML algorithms, and statistical evaluation methods to ensure the integrity of the model performance.

8.4 Open Source Tools and Cloud Resources

In recent years, the availability of open source tools and libraries of deep learning algorithms have and will continue to advance ML applications for marine science and other disciplines. For the purpose of this report, "open source tools" is a phrase used for programs or tools that perform specified tasks, in which the source code is openly published and available. In the open source environment, collaborative efforts can make improvements in the source code which is shared with the wider community. All the case studies presented during the workshop used some form of open source tool for machine learning. For example, some used open source libraries downloading ML algorithms to build, train, and test models. Others used open source tools, such as VIAME, that provide end-to-end pipeline framework for integrated computer visualization to build, train, and test ML models. There are a wide range of open source tools that have become increasingly available for ML, and some examples are provided Table 8-1. It can be challenging for researchers to stay updated and knowledgeable as these tools continue to rapidly evolve. For this reason, scientific collaboration and exchange is necessary as we strive to apply ML for high quality and timely scientific products.

Table 8-1. Examples of open source tools and cloud computation resources for machine learning (ML).

Name	Description	Website
Blocks	ML library; Theano	https://github.com/mila-iqia/blocks
Caffe	Deep learning framework by Berkeley AI Research	http://caffe.berkeleyvision.org/
Chainer	Deep learning framework; supports CUDA computation	https://chainer.org/
CNTK	Microsoft Cognitive Toolkit for deep learning	https://github.com/Microsoft/cntk
Fuel	Data pipeline framework for ML; supports Pylearn2	https://fuel.readthedocs.io/en/latest/
Keras	High-level neural networks API; Python; runs TensorFlow, CNTK, or Theano	https://keras.io/
Neon	ML library; Python	https://neon.nervanasys.com
Orange3	Data visualization and ML models	https://orange.biolab.si/
PyBrain	ML Library; Python	http://pybrain.org/
Pylearn2	ML Library; built on Theano	http://deeplearning.net/software/pylearn2/
PyTorch	ML Library; supports Amazon AWS and Microsoft Azure	https://pytorch.org/
Scikit-learn	ML library, data mining and data analysis; Python	https://scikit-learn.org/stable/
Shogun	ML library; community based	https://www.shogun-toolbox.org/
TensorFlow	ML library; multiple programming languages including Python; originally by Google and now widely available	https://www.tensorflow.org/
Theano	ML library; supports NumPy, Keras, Pylearn2	http://deeplearning.net/software/theano/

VIAME	Open source ML toolkit; developed by Kitware; GitHub user friendly	http://www.viametoolkit.org/ https://github.com/VIAME/VIAME

Another discussion point was the need to improve the data enterprise architecture to enhance ML applications. While some endorsed on-premise architecture, others desired more access to cloud computing resources. Cloud computing uses a network of remote servers on the Internet for on-demand services to store, manage, process data, and use applications. Cloud computing relies on shared computing resources over the Internet for a cloud service fee. Cloud environments provide on-demand self-service, broad network access, and rapid elasticity (resources can be scaled up or down depending in the need). Cloud platforms are considerably less costly than on-premise platforms by eliminating the infrastructure overhead of hardware, software, and IT personnel (Moreno et al. 2019). Cloud solutions for ML applications include enhanced efficiencies in data storage and accessibility, data mining and knowledge discovery, and computational resources. For ML applications, the cloud also provides elastic services for scalable performance and scalable storage of big data repositories (Talai 2019). A key challenge in moving enterprise software systems is the migration of their database, and this involves careful strategic planning, cost estimation, and migration duration considerations (Ellison et al. 2018). Some examples of cloud computing resources are provided in Table 8.2.

Table 8-2. Examples of open source tools and cloud computation resources for machine learning (ML).

Name	**Description**	**Website**
Amazon AWS Sagemaker	Cloud resources for building, training, and deploying ML models	https://aws.amazon.com/sagemaker/
Google Cloud Machine Learning	Cloud resources for building, training, and deploying ML models	https://cloud.google.com/ml-engine/
IBM Watson Machine Learning	Cloud resources for building, training, and deploying ML models	https://www.ibm.com/cloud/machine-learning

Scientists that are conducting specific research projects tend to download ML software and algorithms from open source libraries to on-premise servers and computers where their data are stored. On-premise refers herein to servers and software that are locally maintained within the organization's firewall. In this situation, concern was expressed that investigators should eliminate the duplication of datasets and use only the centralized datasets to eliminate potential inconsistencies. Other preferences for on-premise servers include situations where there are regional specific data, objectives, products, and confidentiality constraints. Some argue on-premise servers provide more control for workflow in data preparation, transformation, and analysis when scientific products must be delivered quickly. For IMR, the situation is somewhat different from that for NOAA. IMR has one large institution with a more centralized data center, whereas NOAA has several science centers with a distributed data storage architecture. Some participants indicated that IMR may put a higher priority on increasing data storage, while NOAA may put a higher priority on improving data accessibility.

Confidentiality constraints such as the case with electronic monitoring from fishing vessels would prevent data access for discovery by the wider community, and this situation may be better serviced with on-premise computing resources. On the other hand, the scientists that use larger and multiple historical datasets from routine survey operations, data discovery, and forecast modeling might prefer the enhanced data accessibility and integrated computing capabilities that are available with cloud platforms.

There is also recognition that the solution would likely be a hybrid cloud architecture that optimizes both on-premise and cloud resources, and this could include a combination of private and public cloud infrastructures. Although the organizational culture of the data enterprise may be slow to adopt cloud solutions, any transitional changes to the data enterprise must minimize disruptions and be positive to the workflow and delivery of scientific products. Furthermore, modernization of the data enterprise and its digital policies and applications are most successful when multiple organizations and sectors work collaboratively towards shared objectives.

8.5 Building Partnerships

Participants from academic, governmental, and private sectors were invited to the workshop to provide perspectives on the application of ML, and this provided the opportunity to discuss the importance of embracing partnerships. While data, computing power, and open source tools have become more readily available for ML applications, there is a shortage of scientists with ML expertise. Each sector has different value-based drivers for developing and applying ML applications (Fig. 8-2). During the workshop, the following questions were asked to encourage discussion on building partnerships:

How do we build partnerships and build workforce competence to implement ML?

How to accelerate the contributions of private sector to advance AI capabilities?

Academic institutions bring strengths in research, training, and discovery for developing ML. The need to train our next generation of ML experts is critical to resolve the workforce shortage and to ensure the integrity and reliability of scientific products derived from ML. Cooperative agreements between governments and academic institutions are good partnership mechanisms to build ML capacity through training, recruitment, and collaborative research.

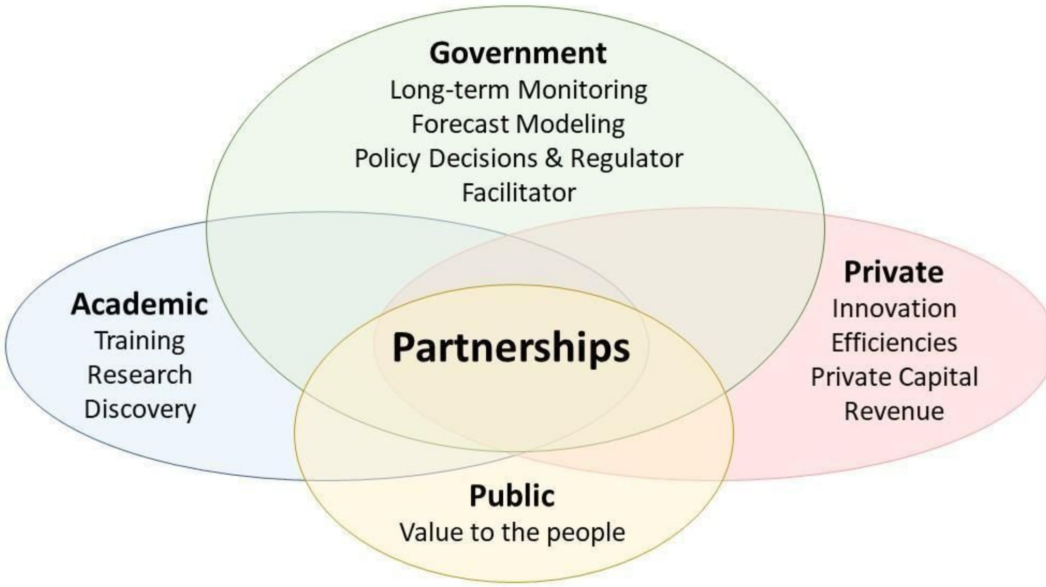

Figure 8-2. The value-based drivers for the academic, governmental, and private sectors that should be considered when developing partnerships.

Science-based government agencies are mandated to provide scientific information for policy decisions, and this typically involves routine monitoring programs resulting in large databases. Due to the dramatic increase in data collections from emerging technologies, the application of ML is critical to reduce the processing costs to deliver high quality and timely scientific products. In addition to building partnerships with academic institutions to train staff and conduct research, the government must work with the private sector to improve its organizational efficiencies. Optimizing their data enterprise with cloud solutions is one example to improve their data processing and workflow with ML, and to make their data more accessible to the wider community for discovery which brings added value to their data enterprise. Another example is to build partnerships with the private sector that provide the innovative technologies for their data collections, and further efficiencies can be realized with the integration of ML into the data collection, post-processing, and analysis relevant to these sampling technologies. A good mixture of government, academic, and private partnerships is necessary in formulating the best practices for integrating ML into organizational efficiencies.

Partnership with the private sector brings the benefits of innovations and efficiencies. The private sector is continuously conducting market research to determine what gaps exist before investment, and for this reason, improved communications and cooperative agreements between government and private sectors help direct the required developments. The industry needs to develop systems that are user friendly and stable. It is only through the use of the system that they will improve, and ensuring good user feedback is essential in further developing the system.

During the workshop, two companies were asked about the key opportunities and challenges between building private-public partnerships.

From the point of view of the Deep Vision project which involves computer vision and ML technologies for automated underwater image analysis (Sections 3 and 4), key opportunities and challenges in building partnerships between private and public sectors include:

Opportunities:

- Work with a unique and diverse pool of knowledge,
- Develop technology in an innovation-friendly environment,
- Test technology in operational environments such as on research vessels, and
- Gain access to end customers to ensure that the technology is customized to their needs.

Challenges:

- Understand the data needs and requirements of the marine researcher,
- Align diverse input and requirements from individuals within the research organization,
- Communication from different departments may lack direction by management,
- Gain an understanding of the fact that the project needs to be commercialized and gain profit in order to be continually supported and developed by the company.

The participants discussed the challenge of private enterprise's proprietary interests for commercialization and profit gain, while scientists often need open data and systems to meet their mandates to provide trustworthy scientific products for policy decisions. This requires a rigorous peer review of the data and methods to confirm the integrity, reliability, and credibility of the science.

Saildrone Inc. provided their perspective (refer to Section 4) on the importance of building complimentary partnerships. Academic institutions train our next generation of experts, and conduct research and development to advance AI. The private sector provides the engineering and feasibility studies for deploying technology platforms. Government agencies require the scientific information for policy decisions, and the efforts of their scientists tend to be devoted to calibrations, standardized survey operations, and statistical models to provide time series indices. These roles for each sector are interconnected and should be coordinated as a collective and complementary effort.

Saildrone's expertise in engineering has produced effective autonomous surface platforms that collect large volumes of acoustic and environmental data. Their immediate concern during deployments is how useful the data collections are for the end-users (e.g., assessment scientists). This is a function of three critical factors:

Survey design: The scientists that utilize these data typically require the appropriate statistical survey design that address the assumptions and sources of uncertainty in assessments. While the private sector provides the expertise to engineer and deliver the technology platforms, government institutions provide guidance on the calibration and survey design requirements for the platform deployments. Utilizing the AI technology for timely processing of the big data collected from these technologies is closely linked with these governmental considerations for delivery of scientific products for policy decisions.

Data collection: Private platforms are deployed to gather large volumes of data, and these data collections should be scaled for the specified survey design by government agencies or field study designs by academic institutions. Ships with onboard scientists can tightly control data collection quality with appropriate data logging metadata, including annotations. While the private sector develops platforms that may specialize in cost-efficient raw data collection at scale, the governmental or academic end-users must determine if raw data is required.

Data labeling: The private sector benefits from combining scientific expertise of government and academic institutions to assure the data collection software provides the appropriate metadata and annotated labeling of the big data collections at scale, so the resultant labels are machine readable.

Overall, the community is eager to engage in ML and understands partnership is a force multiplier to bring together a diversity of assets. Leadership, careful planning, and commitment of the parties are important elements in building partnerships. Furthermore, when there is a mutual understanding of the value-based drivers of each sector (Fig. 8-2), the foundation of trust and sustainability can be established for successful partnerships to advance ML applications.

8.6 Scientific Exchange and Training

Resources should be allocated to build competency in the use and application of ML tools. This can be accomplished by academic training programs, recruitment of expertise, national and international pools of experts, scientific exchange, and collaborations. Building competency is critical to prevent a 'black box' or poorly annotated training datasets that could result in poor quality classifiers and results, thereby harming the integrity of science used for policy decisions. To promote the application of ML, one of the key questions presented to the participants was:

What resources are needed to enable the application/deployment of ML analytics?

In addition to the need to improve the data enterprise with cloud computing and to build partnerships, another key priority was the need for resources to train staff, recruit expertise, and promote scientific collaborations. Similar to any technology, the research and deployment of ML requires the best practices and rigorous scientific evaluation by experts for scientific quality assurance. This is especially true when scientific products are used for policy decisions on the sustainability of our ocean resources, and the reliability and credibility of the science is paramount for stakeholders. There was consensus among the participants that scientific exchange and training should be a high priority for any institution that applies ML. In addition to the integrity of the development, validation, and independent testing of ML models, there must be a rigorous statistical evaluation of the model performance that satisfies the accuracy and precision requirements of the end users.

There is recognition that there are different categories of users that will require varying degrees of training and expertise. Again, the value-based drivers of each sector will drive the training requirements (Fig. 8-2). Researchers will need a strong foundation in statistics and ML computations, while those engaged in data assimilations and forecast modeling will require additional expertise such as time series analysis. There will also be scientists devoted to survey operations and post-processing large volumes of data collection. In this case, the training requirements might be more focused on the application of user-friendly ML software to rapidly process and analyze data with established classifier models. Managers may also need to understand the fundamentals of ML for effective decision making on building ML capacity.

Funding can be used to connect the government, academic institutes, and industry to ensure the collaborative developments are aligned with the operational objectives in support of their value-based goals. Funding should also be allocated to support graduate programs and train staff in ML applications relevant to the goals. The appropriate balance is needed in recruiting and developing technical proficiencies, and it is important that the different user groups have a very clear idea of how this balance can be achieved. In any case, managing expectations and maintaining competence is extremely important. This requires capacity building throughout the partnerships, and the continuity of proficiency in ML is important. This involves strategic partnership and collaboration across institutions and with industry. Considerations must also address how to optimize the expertise across sectors, while minimizing redundancy and optimize the collective efforts to implement ML analytics.

8.7 Recommendations

There has been considerably more interest in ML application in marine science, which can be attributed in part to the rapid development and availability of open source tools in ML analytics. As we proceed with supporting the implementation of ML, there is recognition that capabilities of ML analytics will continue to rapidly evolve in the coming years. Therefore, the working group was presented with the following questions:

How to develop a roadmap for implementation of the rapidly developing ML analytics?

Although the answer to this question involves more strategic planning, the case studies presented during the workshop provided considerable insight on the key components for advancing ML research and innovations. Key recommendations from the various perspectives of government, academic, and private participants included:

Develop and enable access to high-quality datasets: High quality data with enriched metadata is critical for developing reliable and accurate ML models for streamlined data processing, data assimilation, and forecast predictions. There was consensus on the importance of improving data accessibility for ML applications, including access to the wider community for research and knowledge discovery that provides added value to the data. Administrative policies are an important consideration, and some data may have confidentiality restrictions.

Model performance: The performance, error rates, and accuracy of the ML models must be documented and acceptable for the end user, and this requires labeled training datasets for training, validation, and independent testing of the model. The reliability and credibility of the scientific products derived from ML are of paramount importance.

Data enterprise modernization: Enhance data storage, accessibility, processing, and workflow capacity using open source tools and cloud computing when efficiencies can be gained. Hybrid solutions that integrate on-premise and cloud resources may well be the vision for future improvements in data architecture. For successful migration to cloud computing, careful strategic planning must consider cost estimations, migration duration, administrative policies, need to minimize disruptions to workflow, and delivery of scientific products. Partnerships must be built toward shared objectives with an understanding of how ML will complement and augment human capabilities.

Address the big data bottleneck: Emerging technologies have resulted in a dramatic increase in the volume of data collected which exceeds manual processing capacity; therefore, user-friendly ML toolkits are needed to reduce processing time and cost with automated detection and classification capabilities.

Promote collaborative ML research: In addition to improving ML algorithms and methods, open source tools that integrate ML and computer vision technologies into more user-friendly end-to-end pipeline workflow are needed. The community also needs to be engaged in the development of best practices, technical standards, and benchmarks to maintain the integrity of ML science.

Partnerships: Trusting and sustained partnerships are built on three requirements: understanding the value-based drivers of each sector, leadership and careful planning, and commitment based on the significance of the collective goals.

Scientific exchange: Scientific exchange is critical to remain knowledgeable and to make improvements in the rapidly evolving ML discipline. Redundancy and duplication of effort is inevitable in a rapidly developing field, but should nevertheless be minimized. Improvements in organization, communication, and collaboration would be a way to address this. The workshop is a step towards better exchange between the U.S. and Norway within this field, and similar efforts should be encouraged.

Training: Building and maintaining proficiency is a critical requirement for building ML capacity, and the best investment is to train your dedicated workforce. Scientists engaged in ML research will need more rigorous ML training including a strong aptitude in statistical computations. Online introductory ML training would also be helpful to managers. Cooperative agreements and academic support to develop ML focused training programs are also necessary to resolve the present shortage of ML experts and build our next generation of ML experts.

In conclusion, ML has already revolutionized how we will process and analyze scientific data. We are experiencing a turning point in data enterprise culture with how ML complements and augments how humans process and analyze data. The intent with these recommendations is to highlight some of the key requirements for advancing human-ML collaborations. The importance of partnerships is interconnected with each of these recommendations, and we hope this report helps the organizational culture shift to effectively utilize ML tools to deliver higher quality and more timely scientific products for policy decisions on the sustainability of ocean resources.

9 Literature Cited

Aggarwal. C.C. 2018. Neural networks and deep learning. Springer. 497 pp. ISBN 978-3-319-94463-0.

Allken, V., N.O. Handegard, S. Rosen, T. Schreyeck, T. Mahiout and K. Malde. 2019. Fish species identification using a convolutional neural network trained on synthetic data. ICES J. Mar. Sci. 76: 342-349.

Bengio, Y., A. Courville, P. Vincent. 2013. Representation learning: A review and new perspectives. IEEE Transactions on Pattern Analysis and Machine Intelligence. 35(8): 1798-1828.

Breiman, L. 2001. Random forests. Machine Learning. 45(1): 5-32. ISSN: 1573-0565. https://doi.org/10.1023/A:1010933404324

Cai, Z., D. Xu, Q. Zhang, J. Zhang, S.M. Ngai, J. Shao. 2015. Classification of lung cancer using ensemble-based feature selection and machine learning methods. Molecular Biosystems, 11(3): 791-800.

Chang, J.H., D.R. Hart, B.V. Shank, S.M. Gallager, P. Honig and A.D. York. 2016. Combining imperfect automated annotations of underwater images with human annotations to obtain precise and unbiased population estimates. Methods in Oceanography, 17:169-186.

Dawkins, M., L. Sherrill, K. Fieldhouse, A. Hoogs, B. Richards, D. Zhang, L. Prasid, et. al. 2017. An open source platform for underwater image and video analytics. In 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 898-906. IEEE.

Ellison, M., R. Calinescu, and R.F. Paige. 2018. Evaluating cloud database migration options using workload models. J. Cloud Computing Advances, 7(6): https://doi.org/10.1186/s13677-018-0108-5

Goodfellow, I., Y. Bengio, A. Courville. 2016. Deep Learning. Massachusetts Institute of Technology. 775 pp. ISBN 9780262035613.

Hornik, K., M. Stinchcombe, and H. White. 1989. Multilayer feedforward networks are universal approximators. Neural Networks, 2(5): 359-366.

Huang, T. W., J. N. Hwang and C. S. Rose. 2016. Chute based automated fish length measurement and water drop detection. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).

Kong, Y., T. Yu. 2018. A deep neural network model using random forest to extract feature representation for gene expression data classification. Nature Scientific Reports. 8(4):16477. ISSN 2045-2322.

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

LeCun, Y., B. Boser, J.S. Denker, R.E. Henderson, W. Hubbard, L.D. Jackel. 1989. Backpropagation applied to handwritten ZIP code recognition. Neural Computation, 1:541-551.

Lozada-Misa P., B. D. Schumacher, and B. Vargas-Ángel. 2017. Analysis of benthic survey images via CoralNet: a summary of standard operating procedures and guidelines. Pacific Islands Fish. Sci. Cent., Natl. Mar. Fish. Serv., NOAA, Honolulu, HI 96818-5007. Pacific Islands Fish. Sci. Cent. Admin. Rep. H-17-02, 175 p. https://doi.org/V5/10.7289/V5/AR-PIFSC-H-17-02.

Malde, K., N.O. Handegard, L. Eikvil and A.B. Salberg. 2019. Machine intelligence and the data-driven future of marine science, ICES Journal of Marine Science, fsz057, https://doi.org/10.1093/icesjms/fsz057

Margolis, S., W.L. Michaels, B. Alger, C. Beaverson, M.D. Campbell, E.J. Kearns, M. Malik, C.H. Thompson, B.L. Richards, C.C. Wall and F. Wallace. 2019. Accessibility of Big Data Imagery for Next Generation Machine Learning Applications. NOAA Tech. Memo. NMFS-F/SPO-194, 68 p.

Moen, E., N.O. Handegard, V. Allen, O.T. Albert, A. Harbitz and K. Malde. 2018. Automatic interpretation of otoliths using deep learning. PLoS ONE 13(12): e0204713. https:/doi.org/10.1371/journal.pone.0204713

Moreno-Vozmediano, R., R.S. Montero, E. Huedo, and I.M. Llorente. 2019. Efficient resource provisioning for elastic Cloud services based on machine learning techniques. J. Cloud Computing Advances, 8(5): https://doi.org/10.1186/s13677-019-0128-9

Richards, B.L., O. Beijbom, M.D. Campbell, E.M. Clarke, G. Cutter, M. Dawkins, D. Edgington, D.R. Hart, M.C. Hill, A. Hoogs, D. Kriegman, E.E. Moreland, T.A. Oliver, W.L. Michaels, M. Piacentino, A.K. Rollo, C. Thompson, F. Wallace, I.D. Williams, K. Williams. 2019. Automated Analysis of Underwater Imagery: Accomplishments, Products, and Vision. NOAA Technical Memorandum NOAA-TM-NMFS-PIFSC-83. 58 p. https://doi.org/10.25923/0cwf-4714.

Ronneberger, O., P. Fischer, T. Brox. 2015. U-Net: Convolutional networks for biomedical image segmentation. https://arxiv.org/pdf/1505.04597.pdf

Russakovsky, O., J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, et al. 2015. Imagenet large scale visual recognition challenge. International Journal of Computer Vision, 115: 211–252.

Russel, S.J., P. Norvig. 2010. Artificial intelligence: A modern approach. Third Edition. Prentice Hall. ISBN 9781036042594.

Samuel, A. 1959. Some studies in machine learning using the game of checkers. IBM Journal of Research and Development, Vol. 3, Issue 3.

Schneier, B. 2016. The hidden battles to collect your data and control your world. W.W. Norton and Co., First Edition.

Shelhamer, E., J. Long, and T. Darrell. 2017. Fully Convolutional Networks for Semantic Segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39, 640-651. https://doi.org/10.1109/TPAMI.2016.2572683

Stone, V.J. 2019. Artificial intelligence engines. Sebtel Press. 201 pp. ISBN 97809563728819.

Talai, D. 2019. A view of programming scalable data analysis: from cloud to exascale. J. Cloud Computing Advances, Systems and Applications, 8:4, https://doi.org/10.1186/s13677-019-0127-x

Wang, G., J. N. Hwang, C. Rose, and F. Wallace. 2019. Uncertainty-Based Active Learning via Sparse Modeling for Image Classification. *IEEE Transactions on Image Processing*, *28*(1), 316-329.

Williams, I.D., C. Couch, O. Beijbom, T. Oliver, B. Vargas-Angel, B. Schumacher, and R. Brainard. 2019. Leveraging automated image analysis tools to transform our capacity to assess status and trends on coral reefs. *Frontiers in Marine Science*, *6*: 222.

Appendix A. Workshop Terms of Reference and Agenda

Norway – U.S. Workshop on Machine Learning to Improve Science for the Sustainability of Living Ocean Resources

Institute of Marine Research (IMR) in Bergen, Norway

April 23-25, 2019

Terms of Reference: As big data and analytical tools become more commonplace for research and scientific operations in the marine environment, there is an increasing need for fisheries and marine scientists to build collaborative partnerships with experts in machine learning, computer vision, and artificial intelligence. Sampling technologies, such as acoustic and optical remote sensing, have in recent years become more readily available and deployed in the ocean environment. This has resulted in a dramatic increase in data collections that have exceeded our ability to process in a timely manner using conventional manual processing methods. Therefore, the urgency to utilize artificial intelligence methods that incorporate computer vision and machine learning analytics is a priority to reduce the cost of big data processing for more accurate and timely scientific products for ocean policy decisions. In recent years, consensus has been developing on the use and benefits of deep learning methods, yet collaborative work is still required to provide guidance on enhancing the data accessibility, workflow, and metrics to fully utilize the analytical tools that are becoming more readily available.

The goal of the workshop is to exchange information on the current state of development, progress, and applications of computer vision and machine learning analytics. The 3-day workshop will be held at the Institute of Marine Research in Bergen, Norway. The first day will provide an overview of ongoing research efforts and progress pertinent to the various applications of machine learning for fisheries and marine science. Day 2 of the workshop will provide a more technical focus on the research for applying the analytical methods of machine learning and to define the priorities for the collaborative roadmap to effectively advance the use of machine learning. It will also include a training session on standard tools widely used in machine learning. Day 3 of the workshop will be devoted to hands-on training with the newly released open source Video and Image Analytics for Marine Environments (VIAME) toolbox that utilizes computer vision and machine learning algorithms.

Location of workshop: Institute of Marine Research (IMR) in Bergen, Norway (see agenda).

Co-chairs:

Hege Hammersland-White, Scantrol Deep Vision AS, Bergen, Norway; Email: hege@deepvision.no

Nils Olav Handegard, Institute of Marine Research, Bergen, Norway; Email: nilsolav@hi.no

Ketil Malde, Institute of Marine Research, Bergen, Norway; Email: ketil.malde@hi.no

William Michaels, NOAA Fisheries, Office of Science and Technology, Silver Spring, MD, USA;

Email: william.michaels@noaa.gov

Workshop Agenda

Day 1 (Tuesday, April 23): IMR Pynten meeting room, Nordnesparken 38, Bergen

Session A. Introductions and overview of machine learning projects

0815:	Coffee	
0830:	Introductions and review of workshop's terms of reference.	Hege Hammersland, Scantrol Deep Vision AS
0840:	A1. Machine intelligence and the data-driven future of marine science.	Ketil Malde, IMR
0900:	A2. Utilizing machine learning analytics to enhance the value and timeliness of scientific products from big data collected in the marine environment.	William Michaels, NOAA
0920:	A3. The COGMAR project.	Arnt Børre Salberg, Norwegian Computing Center
0940:	A4. NORCE and smart oceans.	Annette Fagerhaug, NORCE
1000:	Break	

Session B. Deep vision and image analysis

1020:	B1. State of Deep Vision image analysis - Deep Vision Project.	Helge Hammersland, Scantrol Deep Vision AS
1040:	B2. Deep Vision Bio Base.	Håvard Vågstøl, Scantrol Deep Vision AS
1100:	B3. Deep Vision – From concept to production model.	Kristoffer Løvall, Scantrol Deep Vision AS
1120:	B4. Segmentation and measurements.	Ricard Prados Gutiérrez, Girona Vision Research
1140:	B5. VIAME toolkit and potential deep learning applications.	Matt Dawkins, Kitware Computer Vision Inc.
1200:	Lunch	

Session C. Machine learning applied to surveys

1300:	C1. Automated species recognition using CNNs.	Vaneeda Allken, IMR
1320:	C2. The IMR deep vision pipeline for the Herring survey, DV + CNN + LSSS.	Nils Olav Handegard, IMR
1340:	C3. Automated acoustic data processing.	Olav Brautaset, NR
1400:	C4. Population estimation by combining automated and manual image annotations.	Devora Hart, NOAA

1420:	C5. Scaling up coral reef monitoring through imagery and machine learning: Advances and next steps.	Courtney Couch, NOAA
1440:	C6. The wish list for automated processing of underwater images for the coastal survey programs.	Kim Halvorsen, IMR
1500:	C7. Automatic interpretation of otoliths using deep learning.	Endre Moen, IMR
1520:	Break	

Session D. Machine learning applied to fisheries monitoring

1540:	D1. The role of machine learning in fisheries monitoring; Examples from Alaska.	Farron Wallace, NOAA
1600:	D2. Automated fish analysis aboard research vessel operations: Building a library for image processing and machine learning to support electronic monitoring programs.	Andy Jones, NOAA; Ben Woodward, CVision AI
1620:	D3. Electronic monitoring from fishing vessel operations: The need for and progress in automating the video review for fisheries electronic monitoring footage.	Niki Rossi, NOAA; Ben Woodward, CVision AI
1640:	D4. Capture identification and analysis – challenges at the Norwegian directorate of fisheries.	Roger Fosse, Kine Iversen og Atle Øinas, IMR
1700:	Adjourn	

Social Events:

1700:	Social at Altona followed by a guided walking tour by Bergen By Experts along the waterfront of Bergen.	Scantrol Deep Vision AS
1900:	Reception at Scantrol offices. Pre-registration required.	Scantrol Deep Vision AS

Day 2 (Wednesday, April 24): IMR Sildetønnen meeting room, Nordnesboder 4, 4.etg., Bergen

Session E. Machine learning tools and methods

0815:	Coffee	
0830:	Opening remarks and new introductions.	Ketil Malde, IMR
0840:	E1. Collaborative machine learning through expert knowledge – Implementation.	Håvard Vågstøl, Scantrol Deep Vision AS
0920:	E2. Underwater imaging: From color calibration to machine learning.	Rafael Garcia, Girona Vision Research
1000:	Break	
1020:	E3. A data management platform for autonomous marine measurements	Jeremy Cook, NORCE
1040:	E5. KERAS and TensorFlow.	Vaneedas Allken, IMR

1100: Panel discussions (Co-chairs)

1200: Lunch

Session F. Training session on GitHub and continuous integration

Pre-registration required.

1330: F1. Tutorial: GitHub and continuous integration. — Ibrahim Umar, IMR

1530: Break

Session G. Training for Video and Image Analytics for Marine Environments (VIAME)

Pre-registration required.

Matt Dawkins – Kitware Computer Vision, Inc.

1600: G1. System overview and software installation.

1630: G2. Running Existing Detectors and GUI basics.

1700: Adjourn (VIAME training will continue on Day 3)

Social event

1930: Dinner at Brasilia. Pre-registration required.

Day 3 (Thursday, April 25): IMR Sildetønnen meeting room, Nordnesboder 4, 4.etg., Bergen

Session G. Training for Video and Image Analytics for Marine Environments (VIAME)

Pre-registration required.

Matt Dawkins – Kitware Computer Vision, Inc.

0815: Coffee

0830: G3. Training New Detection Models via Manual Annotation

0930: G4. Rapid model generation.

0830: G5. Comparing detectors, archive summarization and aggregate statistics.

0915: G6. Auxiliary features: Stereo measurement, image registration.

1000: Break

1030: G7. Planned future work and GUI development.

1130: G8. Experimentation, questions, closing discussions.

1230: Adjourn

APPENDIX B. PARTICIPANT LIST

Participants of the Norway–U.S. Workshop on Machine Learning to Improve Science for the Sustainability of Living Ocean Resources. The conveners are designated with an asterisk (*).

Name (last, first)	Affiliation	Email
Aase, Øystein	Scantrol Deep Vision AS, Bergen, Norway	oystein@scantrol.com
Allken, Vaneeda	Institute of Marine Research, Bergen, Norway	vaneeda@hi.no
Alger, Brett	NOAA Fisheries, Office of Science and Technology, Silver Spring, MD, USA	brett.alger@noaa.gov
Brautaset, Olav	Norwegian Computing Center, Bergen, Norway	olav.brautaset@nr.no
Couch, Courtney S.	NOAA Fisheries, Pacific Islands Fisheries Science Center, Honolulu, Hawaii, USA	courtney.s.couch@noaa.gov
Cook, Jeremy	NORCE Norwegian Research Center, Bergen, Norway	jeco@norceresearch.no
Dawkins, Matt	Kitware Computer Vision, Inc., Clifton Park, NY, USA	matt.dawkins@kitware.com
De Halleux, Sebastien	Saildrone, Alameda, CA, USA	sebastien@saildrone.com
Garcia, Rafael	Girona Vision Research, Girona, Spain	rafael.garcia100@gmail.com
Fosse, Roger	Norwegian Directorate of Fisheries, Bergen, Norway	roger.fosse@fiskeridir.no
Hammersland, Helge	Scantrol Deep Vision AS, Bergen, Norway	helge@scantrol.no
*Hammersland-White, Hege	Scantrol Deep Vision AS, Bergen, Norway	hege@deepvision.no
*Handegard, Nils Olav	Institute of Marine Research, Bergen, Norway	nilsolav@hi.no
Hart, Deborah R.	NOAA Fisheries, Northeast Fisheries Science Center, Woods Hole, MA, USA	deborah.hart@noaa.gov
Iversen, Kine	Norwegian Directorate of Fisheries, Bergen, Norway	kine.iversen@fiskeridir.no

Jones, Andrew	NOAA Fisheries, Northeast Fisheries Science Center, Woods Hole, MA, USA	andrew.jones@noaa.gov
Kim, Hyeongki	Institute of Marine Research, Bergen, Norway	1hkk1002@gmail.com
Løvall, Kristoffer	Scantrol Deep Vision AS, Bergen, Norway	kristoffer@deepvision.no
*Malde, Ketil	Institute of Marine Research, Bergen, Norway	ketil.malde@hi.no ketil@malde.org
*Michaels, William L.	NOAA Fisheries, Office of Science and Technology, Silver Spring, MD, USA	william.michaels@noaa.gov
Minnelsen, Gitle	Scantrol Deep Vision AS, Bergen, Norway	gitlem@gmail.com
Moen, Endre	Institute of Marine Research, Bergen, Norway	endre.moen@hi.no
Myers, Savannah	IMBRSea-IMR, Belgium	savannah.myers@imbrsea.eu
Prados Gutiérrez, Ricard	Girona Vision Research, Girona, Spain	ricardpg@gmail.com
Rossi, Nichole	NOAA Fisheries, Northeast Fisheries Science Center, Woods Hole, MA, USA	nichole.rossi@noaa.gov
Salberg, Arnt-Børre	Norwegian Computing Center, Bergen, Norway	salberg@nr.no
Stephansen, Annette Fagerhaug	NORCE Norwegian Research Center, Bergen, Norway	anst@norceresearch.no
Umar, Ibrahim	Institute of Marine Research, Bergen, Norway	ibrahim.umar@hi.no
Vågstøl, Håvard	Scantrol Deep Vision AS, Bergen, Norway	havard@deepvision.no
Wallace, Farron	NOAA Fisheries, Alaska Fisheries Science Center, Seattle, WA, USA	farron.wallace@noaa.gov
Woodward, Benjamin	CVision AI, Medford, MA USA	benjamin.woodward@cvisionai.com
You, Junyong	NORCE Norwegian Research Center, Bergen, Norway	you@norceresearch.no
Yu, Liangju	Institute of Marine Research, Bergen, Norway	liangjuyu@gmail.com

Appendix C. Glossary

Artificial intelligence (AI): Artificial intelligence is a branch of computer science dealing with the simulation of intelligent behavior in computers.

Computer vision: Computer vision is an interdisciplinary scientific field that deals with how computers can be made to gain high-level understanding from digital images or videos. As a scientific discipline, computer vision is concerned with the theory behind artificial systems that extract information from images.

COGMAR: Ubiquitous cognitive computer vision for marine services (COGMAR) project of the Norwegian Computing Center.

CoralNet: CoralNet is a web-based repository for coral reef images, to organize, annotate images, and view annotation statistics. Computer vision algorithms provide 50-100% automation in annotations and image analysis. Further information is available at https://coralnet.ucsd.edu/

CNNs: Convolutional neural networks (CNNs) are a deep learning method to classify, localize, and segment objects in non-standard, sparsely labelled, image data. CNNs can use unlabeled data in deep learning based on three main concepts: (i) cross-domain transfer learning, (ii) semi-supervised learning, and (iii) data augmentation and simulation.

Deep learning: Deep learning is an artificial intelligence function that imitates the workings of the human brain in processing data and creating patterns for use in decision making. Deep learning is a subset of machine learning (ML) in artificial intelligence (AI) that has networks capable of learning unsupervised from data that is unstructured or unlabeled. Also known as deep neural learning or deep neural network.

GitHub: Prominent and feature-rich VCS software available at www.github.com.

HabCam: The towed optical platform, referred to as HabCam, is used by NOAA Fisheries' Northeast Fisheries Science Center to conduct annual scallop surveys for stock assessments.

Machine learning (ML): Machine learning is the scientific study of algorithms and statistical models that computer systems use to effectively perform a specific task without using explicit instructions, relying on patterns and inference instead. It is seen as a subset of artificial intelligence. Machine learning algorithms build a mathematical model based on sample data, known as training data, in order to make predictions or decisions without being explicitly programmed to perform the task. Machine learning algorithms are used in a wide variety of applications, such as email filtering, and computer vision, where it is infeasible to develop an algorithm of specific instructions for performing the task. Machine learning is closely related to computational statistics, which focuses on making predictions using computers.

NOAA: National Oceanic and Atmospheric Administration (NOAA)

NOAA Fisheries: NOAA's National Marine Fisheries Service is referred to as NOAA Fisheries

UNet architecture: The UNet architecture is built upon the Fully Convolutional Network and modified in a way that yields better segmentation.

VIAME: Video and Image Analytics for Marine Environments (VIAME) is an open-source system for analysis of underwater video and imagery that enables rapid, low-cost integration of new machine learning algorithmic modules, datasets and workflows. Further information on VIAME is available at http://www.viametoolkit.org/

Version Control System (VCS): Tools for computer programmers to enable users to safely record the history of a file or set of files, support distributed social coding collaboration, and continuous integration.

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

Acting Under Secretary of Commerce
for Oceans and Atmosphere
and NOAA Administrator
Neil A. Jacobs, PhD

National Marine Fisheries Service
Assistant Administrator for Fisheries
Christopher W. Oliver

www.fisheries.noaa.gov

OFFICIAL BUSINESS

National Marine Fisheries Service
1315 East West Highway
SSMC3, F/ST
Silver Spring, MD 20910