NOAA Technical Memorandum NMFS-AFSC-408

Report of the Image Processing Workshop

September 16–20
Seattle, WA

SEPTEMBER 2020

U.S. DEPARTMENT OF COMMERCE

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

The National Marine Fisheries Service's Alaska Fisheries Science Center uses the NOAA Technical Memorandum series to issue informal scientific and technical publications when complete formal review and editorial processing are not appropriate or feasible. Documents within this series reflect sound professional work and may be referenced in the formal scientific and technical literature.

The NMFS-AFSC Technical Memorandum series of the Alaska Fisheries Science Center continues the NMFS-F/NWC series established in 1970 by the Northwest Fisheries Center. The NMFS-NWFSC series is currently used by the Northwest Fisheries Science Center.

This document should be cited as follows:

Angliss, R., K. Sweeney, E. Moreland, B. Hou, E. Richmond, C. Khan, B. Sanderson, S. Robinson, M. Lynn, and A. Martinez. 2020. Report of the Image Processing Workshop. U.S. Dep. Commer., NOAA Tech. Memo. NMFS-AFSC-408, 77 p.

This document is available online at:

Document available: https://repository.library.noaa.gov/welcome

This document is available to the public through:
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161

www.ntis.gov

NOAA FISHERIES

Report of the Image Processing Workshop

September 16–20 Seattle, WA

Hosted by the National Marine Fisheries Service's Alaska Fisheries Science Center
and sponsored by the Office of Atmospheric Research's UAS Program Office Report

R. Angliss[1], K. Sweeney[1], E. Moreland[1], B. Hou[1], E. Richmond[1,2],
C. Khan[3], B. Sanderson[4], S. Robinson[5], M. Lynn[6], and A. Martinez[7]

[1] Marine Mammal Laboratory, Alaska Fisheries Science Center, NOAA, National Marine Fisheries Service, 7600 Sand Point Way NE, Seattle, WA 98115

[2] Cooperative Institute for Climate, Ocean, and Ecosystem Studies, University of Washington, 7600 Sand Point Way NE, Seattle, WA 98115

[3] Protected Species Branch, Northeast Fisheries Science Center, NOAA, National Marine Fisheries Service, 166 Water Street, Woods Hole, MA, 02543

[4] Ecosystem Analysis Program, Northwest Fisheries Science Center, NOAA, National Marine Fisheries Service, 2725 Montlake Blvd E, Seattle, WA 98112

[5] Hawaiian Monk Seal Research Program, Pacific Island Fisheries Science Center, NOAA, National Marine Fisheries Service, 1845 Wasp Blvd, Honolulu, HI 96818

[6] Marine Mammal and Turtle Division, Southwest Fisheries Science Center, NOAA, National Marine Fisheries Service, 8901 La Jolla Shores Dr, La Jolla, CA 92037

[7] Protected Resources and Biodiversity Division, Southeast Fisheries Science Center, NOAA, National Marine Fisheries Service, 75 Virginia Beach Dr, Key Biscayne, FL 33149

U.S. DEPARTMENT OF COMMERCE

National Oceanic and Atmospheric Administration

National Marine Fisheries Service

Alaska Fisheries Science Center NOAA Techncial Memorandum NOAA-TM-AFSC-408

September 2020

CONTENTS

EXECUTIVE SUMMARY

Imagery and associated data collected from aircraft and vessel surveys have been critical for estimating marine mammal abundance, stock structure, behavior, health, and other information needed to meet the National Marine Fisheries Service's (NOAA Fisheries) mission for decades. With the advent of digital photography, and most recently, advances in image resolution, unmanned aerial systems (UAS), autonomous underwater vehicles (AUV), and remote cameras, the total volume of electronic imagery collected by NOAA Fisheries researchers has increased rapidly. Biologists experience common challenges at regional NOAA Fisheries science centers, including the collection, storage, processing, and analysis of terabytes (TB) of data. For unmanned systems (UxS) to fully transition to operations in NOAA Fisheries, we must have efficient solutions for processing this unprecedented amount of image data. There is a strong need to increase efficiency by using artificial intelligence (AI) techniques, such as machine learning and computer vision, to automate image processing and analysis.

The goals of the 2019 Image Processing Workshop were to bring NOAA Fisheries scientists, private industry, and external scientific collaborators together in order to understand how AI is being used to process marine mammal data in NOAA Fisheries, to discuss what is needed to harness this technology more effectively, to develop a community of NOAA Fisheries AI experts, and to train participants in available tools for implementing AI.

Workshop participants identified key priorities identified:

- Streamline access to crowdsourcing and data science competitions.
- Researchers should increase engagement with AI experts.
- Engagement with industry is critical.
- Invest in dedicated AI staff, training, and collaborations.
- Streamline the process to use UxS for mission needs.

Recommended future agency investments include:

- Dramatically increase data storage (e.g., cloud service providers, larger network storage systems)
- Increase processing power (e.g., cloud computing resources, computers with GPUs).
- Provide a separate source of funds to enable innovations in AI simultaneously with advancing survey efforts through traditional (manned surveys) or new approaches (e.g., UxS, satellites).
- Establish an agreement with NASA's Center of Excellence for Collaborative Innovation (CoECI) to streamline access to data science competitions and crowdsourcing; in the long term, develop ready access to competitions directly through NOAA.
- Establish a team with proficiency in AI projects at each regional NOAA Fisheries science center with the responsibility to coordinate and collaborate with researchers interested in AI.
- Provide training for biologists and staff so they can excel in the new roles required by AI projects (e.g., annotation, programming, project management, etc.).

Workshop appendices provide a crosswalk between workshop recommendations and NOAA AI Strategic Goals, a list of known researchers incorporating AI into their projects at NOAA Fisheries science centers, a "beginner's guide" to starting an AI project, specific sources of training for various roles in an AI project, project presentation summaries, and suggested recent literature on AI relevant to marine mammal researchers.

INTRODUCTION

Image data collected from manned aircraft and vessel-based surveys – and more recently, unmanned systems (UxS) and remote cameras -- have been critical for estimating abundance, trends in abundance, stock structure, behavior, health, and other information needed to meet the National Marine Fisheries Service's (NOAA Fisheries) mission (Fritz et al. 2017, Pace et al. 2017, Hinke et al. 2018, Khan et al. 2018). In the past decade, NOAA Fisheries has made significant advances in remote collection of image data using equipment in manned aircraft, UxS, and remote cameras, resulting in an enormous increases in the volume of images to be processed. Image processing and analysis has been primarily manual, and automated processing of marine mammal data sets has been rare until recent years, with the exception of a few trailblazers in the NOAA Fisheries marine mammal community (Melancon et al. 2011, Conn et al. 2014, Bogucki et al. 2018). The NOAA Fisheries marine mammal community's transition to partially or fully automated processing and analysis has been delayed due to the need to seek external funding to develop software, hardware, and workflows that could improve the processing speed.

Researchers at NOAA Fisheries recognize that supporting just the collection of image data is no longer an option: significant resources must be directed at new hardware and software approaches to advance our ability to process and analyze these data on a timeline that is relevant to NOAA Fisheries decision-makers. Thus, in September 2019, NOAA Fisheries convened a week-long workshop in Seattle, Washington, to meet the following goals:

- Ensure that NOAA Fisheries scientists who collect large image datasets for assessing marine mammals are aware of the "state of the science" for collection and automated processing throughout NOAA Fisheries and in the broader scientific community.
- Engage NOAA Fisheries collectors of large image datasets in a discussion about what is needed to improve data collection, processing, and analysis, and frame recommendations to NOAA leadership that can guide future research investments.
- Provide hands-on training about use of NOAA-funded machine learning software, VIAME, and other types of software available to automate image processing.

The workshop included invited presentations about a variety of successful, ongoing, and new projects using various approaches to automate processing of large volumes of imagery data (Appendix A). Break-out sessions focused on specific types of image data collected (e.g., images collected for animal population counts and trends, for understanding underwater foraging, etc.) provided opportunities to get data-type specific recommendations from researchers. Presentations from technical companies (Kitware, Microsoft, West Inc., WildMe) provided information on company programs, technical processes, and software available to enable automation. Industry representatives provided insights into their perspectives on the various roles and responsibilities required to complete an AI project, and provided advice about how NOAA Fisheries might be more successful in future AI efforts. NOAA Fisheries leadership (William Michaels, NOAA Fisheries Ocean Technology Program) provided important agency-level context to our workshop.

The steering committee, which planned the workshop, consisted of members throughout the regional NOAA Fisheries science centers. Attendees and participants included a variety of individuals from NOAA, other government agencies, universities, and private industry (Appendix B).

This workshop was designed to build on the recent NOAA Fisheries Strategic Initiative on Automated Image Analysis, which was a multi-year initiative from 2014 to 2019 (Richards 2015, Richards et al. 2019) that resulted in two main products: the Video and Image Analytics for a Marine Environment (VIAME) open-source software toolkit and the CoralNet web-based solution for benthic image analysis. Although VIAME was focused on analysis of underwater imagery, the basic infrastructure can be adapted to address many other image processing needs in NOAA Fisheries, with additional support. Some adaptations to VIAME are in progress for aerial marine mammal applications in NOAA Fisheries. Some NOAA Fisheries biologists also have AI projects developed outside of the VIAME framework (Science Center AI project descriptions are in Appendix C).

This workshop also provides significant information that may be useful for NOAA's efforts to increase the agency's adoption of artificial intelligence. In particular, in 2019, an Executive Order on Maintaining American Leadership in Artificial Intelligence was issued. In response, NOAA developed the NOAA AI Executive Committee, which includes senior members from all NOAA line offices. This Committee is charged with developing a NOAA AI Strategy, with the following goals:

- Goal 1: Establish an efficient organizational structure to advance AI across NOAA.
- Goal 2: Advance AI research and innovation in support of the NOAA mission.
- Goal 3: Accelerate the transition of AI research into operational efficiencies.
- Goal 4: Strengthen and expand AI partnerships.
- Goal 5: Promote AI proficiency in the NOAA workforce.

In 2020, NOAA will be developing an internal AI Implementation Plan which will be based on the NOAA AI Strategy and will set the direction for future NOAA work with AI; the Implementation Plan is projected to be completed by the fall of 2020. The development of Cooperative Research and Development Agreements (CRADAs), Memoranda of Opportunity (MOUs), and other agreements were identified by NOAA Fisheries leadership as ways to make progress with AI before the completion of the Implementation Plan.

This workshop report provides information for consideration by the NOAA AI Executive Committee that may aid their work in addressing Goals 1-5; crosswalks between these goals and the results of this workshop are provided in Appendix D.

Table 1. -- Definitions of key terms and acronyms identified in the workshop and used throughout this workshop report.

KEY TERMS AND ACRONYMS	
Term	**Definition**
Algorithms	Machine learning algorithms build a mathematical model based on sample data, known as "training data", in order to make predictions without being explicitly programmed to perform the task.
Annotation	Data annotation is the task of labeling data (text, audio, images). Annotation is crucial because clean, annotated data are necessary to train AI models.
Artificial Intelligence (AI)	Artificial Intelligence is the theory and development of computer systems able to perform tasks that normally require human intelligence.
Convolutional Neural Network (CNN)	Convolutional neural networks (CNN) are a specific type of artificial neural network that uses perceptrons, a machine learning unit algorithm, for supervised learning to analyze data. They have become the model of choice for image processing applications and have been effective in image classification, segmentation, object detection, face recognition, and microscopy.
Crowdsourcing	The practice of obtaining services, ideas, or content by soliciting contributions from a large group of people, and especially from the online community. Crowdsourcing events may be paid or unpaid, and are often set up as a contest with a prize offered for the most successful contribution. Crowdsourcing is one effective method that is used to develop new AI algorithms.
Deep learning	Deep learning is a type of convolutional neural network where the word "deep" refers to the number of layers through which the data are transformed.
Graphical User Interface (GUI)	A front end display that may include icons, cursors, and buttons which is more user-friendly than a text based command line for running code. Algorithms developed using AI are frequently controlled through a GUI for end users.
Graphics Processing Unit (GPU)	A GPU (or Graphics Processing Unit) is a specialized electronic circuit designed to rapidly perform parallelized calculations to accelerate the creation of images in a frame buffer intended for output to a display device. GPUs are used in embedded systems, mobile phones, personal computers, workstations, and game consoles. GPUs are well suited for many tasks related to AI, machine learning, and computer vision as their architecture lends itself well to performing identical operations on a large amount of data simultaneously. This parallelization delivers tremendous performance increases when processing imagery data.

KEY TERMS AND ACRONYMS	
Term	**Definition**
Hackathon	A hackathon event is similar to crowdsourcing, where individuals spanning multiple disciplines collaborate closely on a project or task, with the goal of producing some functional product or solution to a problem by the end of the event. Hackathons can be a good way to engage external partners. Breaking the overarching project into sub-projects that can be completed during a hackathon can be helpful.
Machine Learning	The term "machine learning" was coined in 1959 by Arthur Samuel, and is a branch of artificial intelligence based on the idea that systems can learn from data, identify patterns and make decisions with minimal human intervention. In other words, machine learning is a method to achieve AI.
Model	A machine learning model is a mathematical model that generates predictions by finding patterns in your data. *What is the difference between an algorithm and a model?* An algorithm is the general approach you will take. The model is what you get when you run the algorithm over your training data and what you use to make predictions on new data. You can generate a new model with the same algorithm using different data, or you can create a different model from the same data with a different algorithm.
Neural networks	Neural networks are a type of machine learning modeled after the human brain where nodes or elements of the model are interconnected, sending and receiving inputs, and are capable of learning. Networks are trained by modifying weight values assigned to various input nodes. When new data are presented to the network, these weight values are used for inference.
Structure from motion	Structure from motion is a photogrammetric technique by which three-dimensional structures can be derived from a sequence of two-dimensional images. Careful planning on how the imagery is collected improves the quality and accuracy of the derived structures.
Supervised learning	In supervised machine learning, the algorithm learns on a labeled dataset (for example, classification and regression).
Testing data	Testing data are simply data that are withheld from the training data and then used to evaluate the performance of the algorithm at a later stage.
Training data	Training data are data used to train the algorithm. Usually, the training data are labeled or annotated.

KEY TERMS AND ACRONYMS	
Term	**Definition**
UAS	Unmanned aircraft system, or drone.
Unsupervised learning	In unsupervised machine learning, the algorithm tries to make sense of unlabeled data by extracting features and patterns on its own.
UxS	Unmanned system operated by air, land, or water.
VIAME	Video and Image Analytics for Marine Environments, an open-source system for analysis of video and imagery for fisheries stock assessment and a variety of other applications (viametoolkit.org).

PROJECT PRESENTATIONS

To bring awareness of the current "state of science" for data collection and automated processing, presentations were invited from NOAA Fisheries scientists, the broader marine mammal community involved in AI, and industry specialists. The presentations covered a broad range of themes including approaches to image processing, AI for enumeration, automating image processing for photo-identification and photogrammetry, statistical considerations, and hardware, processing, and storage. Presentation summaries can be found in Appendix E.

Table 2. -- List of presentations and authors presented on day one to show the current "state of science" for AI-applications.

LIST OF PROJECT PRESENTATIONS	
Approaches to Image Processing	
Citizen scientists help researchers investigate Steller sea lions	Sweeney (NOAA Fisheries)
Computer vision: practice and pitfalls	Redmon (University of Washington)
Machine learning and deep learning differ in effort: Use satellites to inform drone work – Tip & Cue Frameworks	Johnston (Duke Univeristy)
AI/ML for wildlife conservation and myth busting for biologists	Morris (Microsoft)
UAS applications with bottlenose dolphins, birds and bats	Thompson / Erickson (WEST)
Using AI to identify whales from visually represented acoustic data	Allen (NOAA Fisheries)
AI for Enumeration	
Applications of Machine Learning Algorithms to Automate Data Extraction from Images	Altukhov (North Pacific Wildlife Consulting)
Automated surveys for ice-associated seals in the Arctic	Moreland (NOAA Fisheries)
Collection and analysis of imagery and video by the SWFSC Antarctic Ecosystem Research Division	Hinke (NOAA Fisheries)
Use of Images for Photo-Identification & Photogrammetry	
Computer vision for conservation: Automating right whale photo ID	Khan (NOAA Fisheries)

LIST OF PROJECT PRESENTATIONS	
Flukebook: Multi-modal, multi-stage machine learning for marine mammal research with citizen science	Holmberg / Parham (WildMe)
Species identification and stereo measurements	Lauffenburger (NOAA Fisheries)
Morphometrics and volumetrics of pinnipeds from imagery	Shero (Woods Hole Oceanographic Institution)
Statistical Considerations	
Accounting for species misclassification	McClintock (NOAA Fisheries)
Estimating abundance with automated detection systems	Conn (NOAA Fisheries)
Notes on statistical considerations for photoID mark-recapture	Conn (NOAA Fisheries)
Hardware, Processing, & Storage	
Network effects: storage, processing, and connections	Hou (NOAA Fisheries)
Moving to the cloud	O'Neil (NOAA Fisheries)
Matching UAS platforms to your data objectives	Seymour (United States Geological Survey / Cherokee Nation)

KEY PRIORITIES IDENTIFIED

NOAA Fisheries should streamline access to crowdsourcing platforms

There are many platforms (Zooniverse, Kaggle, TopCoder, etc.) that allow crowdsourcing of an image processing project. Crowdsourcing allows researchers to harness a vast talent pool with broad expertise that is often able to solve a complicated problem at a minimal cost and faster speed relative to traditional contracting for services or products. Agency workshop participants noted that biologists and contracting specialists seldom have the expertise needed to set up and manage a crowdsourcing acquisition. The NOAA AI strategy, once implemented, may provide easy access to crowdsourcing, but the expected timeline for this is months to years. Fortunately, National Aeronautics and Space Administration (NASA) has an existing, streamlined acquisition process for setting up crowdsourcing projects that NOAA is welcome to use (for an overhead fee). Workshop participants strongly recommended that NOAA Fisheries pursue an agreement with NASA so that NOAA scientists can use their acquisition process. Industry participants noted that the agency could attract good crowdsourcing contest participation by ensuring that project data are clean, easy to access, ready to process, and that the challenge posed is interesting from a data science standpoint.

NOAA Fisheries should increase engagement with AI experts

Most NOAA Fisheries researchers are still in the early days of learning how to use AI, and direct engagement with AI experts will improve our utilization of AI in accomplishing NOAA's mission. The incorporation of crowdsourcing and AI into image analysis projects is expected to save time over the long term, but not all datasets will benefit from AI or crowdsourcing. For instance, large datasets of annotated data are needed for training and testing an algorithm; if a project has a small dataset, it is likely not a good candidate for deep learning as thousands of ground truth values are required for more sophisticated approaches. Similarly, if it is challenging for humans to accurately categorize data, then the project may not be a good fit for an AI algorithm to make the classification.

It is also clear that incorporating crowdsourcing and AI into data processing will initially make the process more time-intensive and complicated. Industry representatives and NOAA Fisheries staff with experience in these processes noted that it will be important for knowledgeable project managers to evaluate two key planning points: 1) at what processing step will crowdsourcing or AI be most helpful (e.g., annotation, detection, rotation of image components, etc.), and 2) what is an appropriate threshold for success. For instance, the winning algorithm for matching North Atlantic right whales was 87% accurate, which was considered sufficient for the stated purpose of the project (Bogucki et al. 2018). In contrast, a similar percent accuracy was not considered sufficient for identifying the presence/absence of whales in images collected in the U.S. Arctic (Ferguson et al. 2019). An algorithm has not been needed to find Steller sea lions in images collected from remote cameras; human volunteers participating in StellerWatch on Zooniverse process the dataset sufficiently to reduce NOAA Fisheries' researcher time. Researchers may have to choose between very good algorithms or approaches that are very specific to a particular dataset, and more broadly applicable algorithms with lower classification rates. Modeling acceptable precision and bias in advance of the project will guide these decisions.

Engagement with industry is critical

Industry participants from Microsoft, Vulcan, West Inc., Kitware, and WildMe contributed to the success of the workshop by providing outside technical advice that prompted novel discussions about required expertise to work effectively on AI, advice about contracting language and how to evaluate potential AI service vendors, and crowdsourcing outcomes. Industry participants highlighted which programming languages should be learned by key agency staff involved in AI projects (Python, C/C++, R), discussed under what circumstances processing should be based in the cloud or on a local GPU system, how to attract good contest participation, what types of conferences and workshops would most benefit NOAA Fisheries biologists, and why image file type matters for some AI projects.

NOAA Fisheries should invest in dedicated AI staff, training, and new collaborations

Currently, NOAA Fisheries biologists must perform several roles outside of their domain of expertise to incorporate AI into a project. Investment in dedicated staff to perform vital tasks such as project management, data science, systems engineering, algorithm development, IT, and software engineering is necessary to support various project teams. Collaborations can also play a vital role in advancing the implementation of AI. Larger data sets provide both more training data and test data, and there is an opportunity for research teams to pool data so different types of scientific questions can be addressed using AI. Working out collaborations between research teams will take considerable time, outstanding communication skills, and a willingness to embrace data sharing across groups internal and external to the agency.

Biostatisticians will have an ongoing role in this quickly evolving field. Understanding bias and precision in an AI context, and how errors may be propagated through the AI process in a way that affects an important output (e.g., an abundance estimate) will be critical. Industry should work with biostatisticians so they understand how advances in algorithm development impact the intended products of the analyses.

Streamline the approval process for UxS platforms that meet agency mission needs

Many workshop participants were interested in crowdsourcing and AI primarily because of the large number of images generated when using UxS platforms. In some cases, images from UxS platforms approved by the agency cannot be used in an AI context due to the platform's inability to effectively and reliably collect usable data critical for post-processing (e.g., accurate time, location, altitude) and imagery during flight. Some platforms don't support the integration of custom payloads, or the manufacturer won't provide information on the integrated, proprietary camera. The lack of UxS platform options stifles innovation, and may have serious downstream impacts to researcher ability to collect and process data to address mission needs. The process to introduce platforms for investigation and approval should be streamlined and include the researchers who use these platforms as stakeholders.

IDENTIFICATION OF FUTURE AGENCY INVESTMENTS

Workshop participants in NOAA Fisheries currently integrating AI into project workflows are doing so independently and predominantly by finding funding sources outside NOAA Fisheries. Specific areas of agency support were identified to move the Agency forward in adopting AI techniques to improve efficiency. The following recommended investments were identified by NOAA researchers focused on abundance estimation, video processing, photo-identification, and photogrammetry.

Significantly increase data storage

Researchers need access to larger internal servers and cloud storage providers for access to large image training sets. Use of high resolution imagery allows aircraft to fly at higher altitudes, which increases the survey area and reduces disturbance of animals, but also results in larger image file sizes. The use of UxS has also increased the overall number of images collected by researchers. This increase in quantity and file size of imagery has exceeded the internal data storage and management capacity of at some local research groups. In addition, imagery needs to be moved easily between NOAA and partners working to implement AI approaches to image analysis. Moving imagery through the existing IT framework is cumbersome and inefficient that researchers resort to physical transfer using external hard drives. Improvements to IT infrastructure to increase data storage and overall capacity are needed as well as easy access to cloud storage providers to improve collaborative efforts with external partners.

Increase processing power

Training and running AI models on large image sets requires access to new kinds of powerful computing resources. Currently, researchers with access to computers with the required graphics processing units (GPUs) have to work on these machines off-line. They have not been integrated into the Center IT network, often due to limitations of the local IT support group. Most researchers doing AI work do not have access to GPU machines or cloud-computing services with adequate processing capabilities and memory to work with large imagery sets. NOAA Fisheries should evaluate the pros/cons of increasing internal computing resources or using cloud service providers, ensure that the NOAA Fisheries research community knows about various options, and IT teams should facilitate access so that resource limitations are not an obstacle to progress.

Financial resources

All participants identified a need for project funding. Traditional, established surveys and analytical approaches must be continued while researchers simultaneously initiate new AI projects, run pilot projects, test the outcomes, and analyze any differences between the established approaches and the new AI approach. Marine mammal researchers in NOAA Fisheries have successfully found funds through competitive processes to support AI projects, but a stable source of funding for research and development would speed up progress. Financial resources are needed to increase data storage and efficiency, add technical staff, conduct training, and to develop systems, survey protocols, and statistical approaches. Without dedicated funds to support the transition from traditional image analysis approaches to AI, adoption of AI will be slow and cost-prohibitive for many projects. AI Investments are important for significantly reduce time, effort, and cost of data processing relative to traditional approaches (i.e., human analysis).

Improve access to AI developers

Researchers need easy access to the expertise required to develop AI models specific to the image set. This includes access to companies that harness large international communities of developers through competition. Established routes to work with companies such as Kaggle and TopCoder are required to solve the individual challenges of developing new AI models. There are countless approaches to model development and these platforms are the most effective way to identify the appropriate approach for a new problem. Marine mammal researchers with experience in AI identified the lack of easy access to data science competitions as a major hurdle to the development of new AI approaches to processing marine mammal data.

Access to AI expertise should be improved across NOAA line offices by creating opportunities for better communication between researchers who are geographically or bureaucratically separated. Opportunities and incentives should be created for offices to work together and provide expertise and guidance to projects in their infancy and to share information to improve AI efficiency and implementation across the agency. Access to AI expertise within the agency could be effectively facilitated by sponsoring annual NOAA or NOAA Fisheries workshops for researchers involved in AI projects. Providing support to researchers to participate in AI conferences and workshops designed for private industry would expand researchers' professional networks to include AI expertise outside of the agency. Links to outside expertise will allow researchers to continually take advantage of the latest technology far more efficiently than having to hire or train enough AI expertise to stay current.

NOAA Fisheries has made a major investment in VIAME and it currently provides easy-to-implement solutions to certain types of imagery problems. However, its origin is in processing video of fishery resources and has only recently been considered for use by NOAA biologists using other sensors. Modifications to the program to accommodate aerial survey imagery have been funded at the program level but are still in their infancy. Greater investment in the front end user interface of VIAME and support for the web-based application is needed to radically increase access to researchers who lack programming skills. To ensure VIAME flexibility, availability, and usefulness for NOAA Fisheries researchers, there needs to be a broad effort to incorporate the specific needs of researchers studying protected species.

Establish an agreement with NASA's Center of Excellence for Collaborative Innovations (CoECI) to streamline access to data science competitions and crowdsourcing

One way to improve access to AI developers is through collaborations with groups doing innovative work. Data science competitions and other forms of crowdsourcing have become a leading approach to innovation. A single competition can result in thousands of approaches tested and hundreds of solutions submitted to address a single problem. Competition participants benefit by gaining experience with solving novel, real-world problems, and the prestige gained when they develop a winning solution can be used on a resume to demonstrate proficiency. The types of data science challenges presented by NOAA Fisheries are of particular interest to competition participants because they can make a direct contribution to science and conservation efforts and are often more interesting than the more common competition topics from the insurance and finance industries. AI experts caution that one potential issue with one-time solutions on a particular dataset are generalization and operationalization: the limited time involvement those generating solutions could become a road block in the event the code needs to be modified or maintained (e.g., inadequate or unclear documentation of the solution).

The Center of Excellence for Collaborative Innovation (CoECI) at NASA provides researchers access to innovative solutions through internal and external avenues. Challenges are contracted through an IDIQ with companies with proven success in delivering innovative results. These companies provide novel solutions by harnessing the expertise of international communities of innovators who compete to find solutions. Rather than having a handful of solutions from a small team of experts through traditional vendors, this approach results in hundreds or thousands of creative approaches. Setting up a challenge through the NASA process typically takes approximately a month, and each challenge runs for one to three months depending on project needs. For comparison, with other avenues for creating challenges, the process can take upwards of 6 months.

CoECI also works to connect other government agencies to crowdsourcing providers and NOAA Fisheries is in the process of establishing an interagency agreement to support researchers in need of crowdsourced deep learning solutions to image processing. This type of approach is used by NASA for more than AI. They crowdsource innovation through international competition to solve challenges with materials development, satellites, and even hull design and logos. If NOAA developed the same type of arrangement with crowdsourcing vendors, it would establish a clear path for innovation accessible to all line offices.

Internally, NASA also runs small competitions within the agency to leverage the talent of its own staff. These are limited to smaller problems to avoid disruption of expected tasks and result in a number of benefits to the agency including increased morale, creative collaboration, and connecting staff across NASA. NOAA has extensive expertise throughout the agency, however distinct researcher groups are effectively isolated, causing projects to contract out work that could be done internally. Setting up a similar program could help researchers leverage the skills of staff from other programs, reduce costs, build connections, and strengthen morale across the agency.

Establish a team with proficiency in AI projects at each regional NOAA Fisheries science center

To enable diverse applications of AI throughout NOAA, each science center needs an in-house team comprised of specific skill sets that can launch new projects and enable biologists to incorporate AI into research. This includes staff proficient in machine learning, data annotation, algorithm development, software development, AI project management, and AI quality control. This team would be needed whether AI model development was conducted internally or through external partnerships. Biologists involved in AI also need sufficient expertise in AI to establish appropriate and strong contract language, evaluate potential vendors, and evaluate vendor performance.

Provide training to biologists so they can excel in the new roles required by AI projects

The incorporation of AI into a project adds a new level of both technical and organizational complexity. Workshop participants identified training as a need for biologists to learn the language of AI while taking on new roles to usher a project into operations. Some biologists may have the coding proficiency to learn to train models in a familiar language (i.e., R or Python), or may have the untapped ability to learn to code and develop much needed skills in house. Others may require a higher level working understanding of the process to be effective project managers. Training is readily available to help biologists transition to appropriate roles to successfully implement AI techniques. A list of online training courses and resources for AI can be found in Appendix F and suggested literature in Appendix G.

Research, development and operational use of new UxS technology

UxS technology has increased the use of imagery as data and pushed approaches that used to be cost prohibitive to the forefront of current methodology. Harnessing the benefits of this technology requires dedicated research and development to optimize integration into operations. Greater support for development is needed to create--or adapt and adopt--the tools needed to develop and control custom payloads and integrate AI into the data management and processing workflow. Increased access to new types of UxS are needed, which will require streamlined approvals of systems for agency use.

INSIGHTS FROM INDUSTRY

Workshop participants invited industry experts to identify ways that marine mammal researchers could more quickly see improvements in our use of AI. In an effort to better design and implement AI projects, the AI industry participants recommended the following: standardization of file and attribute naming conventions across NOAA projects, effective annotation, building better relationships with internal IT staff, increasing capacity for network storage, shifting to the cloud for data computing and sharing, and (if shifting to cloud storage) staying within one cloud computing ecosystem (e.g., AWS, Azure, etc.).

AI industry experts at the workshop identified barriers to developing successful agency partnerships in the past, including challenges with data sharing, poor management of time and expectations, poor communication between developers and biologists because of difficulty in understanding and communicating domain-specific biology or computer-science terms (i.e., not speaking the same language) and change of scope over the duration of the project. In addition, in part because agency staff are in the early stages of understanding AI, industry experts note that there have been problems with lack of clarity in requests for proposals and statements of work, incomplete documentation about datasets, lack of agreed-upon checkpoints throughout the process, and lack of reasonable expectations about what "success" looks like at the end of an AI project. These problems cause frustration and can lead to incomplete projects.

In order to write a successful AI statement of work with sufficient details and metrics, it is important to remember that AI is only a small part of the process. Metrics of success need to be specified without having unreasonable expectations of the AI component of the project and benchmarks need to be set on its performance. A clear vision of what the need is, or how the deliverable(s) will be used must be clearly communicated (i.e., GUI development, algorithm development, ongoing cloud servers for running algorithms, etc.) to ensure that the partner has a clear understanding of the expected outcome: many concepts that are clear to biologists can be foreign to software engineers and experts in other fields. An example of this is how uncertainty or error in AI outputs is qualified (i.e., often the same terms can have different meanings between biologists and computer scientists). Proper scoping and definition is critical when working with external collaborators. Project stages need to be broken down into manageable milestones with clear, fully documented deliverables. A final consideration is that survey biologists often have a requirement of stability in time-series of analyses. That is to say once a method is implemented, constant incremental improvements in methodology may not be the correct approach as they can bias annual trends, but more infrequent full re-analyses of the datasets may on occasion be required and a more optimal approach for updates.

To effectively evaluate a potential AI vendor, it's important for biologists and others setting up a contract or partnership to check professional references, GitHub contributions, and scientific publications. It may also be beneficial to identify someone within the agency with the background and expertise to assess people and relevant talent. Requesting assistance from colleagues in the AI industry may help with evaluating credentials and capabilities of a potential partner.

There are a variety of GUIs that make AI algorithms usable by the broader research community - WildBook and VIAME are two examples. Creators of the GUIs provided a variety of cautions to workshop participants. Some of these GUI platforms are hosted centrally and used remotely, and computing issues may occur when multiple users are accessing the systems simultaneously. In addition, some amount of financial support for the GUI developers will likely be required for the long term because of the need for ongoing software engineering support for users. Old algorithms and models will need to be retired as new techniques and methodologies are developed, and there are software engineering and statistical implications of these changes. If researchers plan to use an existing GUI after a crowdsourcing effort, it's important to realize that models produced by competition may not readily work within an existing GUI or may encounter problems during the transition from testing- to production-level processing.

DATA COLLECTION CHALLENGES

Novel approaches are becoming more widely used in the field of wildlife sciences including unoccupied aircraft systems, camera tags, camera traps, satellite imagery, and more. Some of these new approaches collect new types of data, or ameliorate the burdens associated with traditional data collection methods. Adoption of these methods improve survey efficiency, productivity, accessibility, and often decrease cost (Christie et al. 2016). However, in many of these cases, adopting these methods has transferred the burden to data processing and analysis because of the enormous amount of data products collected (e.g., imagery, video, audio, etc.; Angliss et al. 2018, Ferguson et al. 2018, Moreland et al. 2015). These increases in data collected have led to advancing and investing in using AI to automate image data processing. Although the primary goal of this workshop was to discuss challenges and solutions of data processing and analysis using AI, it also highlighted some issues in the data collection process and capabilities.

Many researchers are now using UxS to collect various types of data on wildlife that NOAA Fisheries requires to meet our mission. Alex Seymour (USGS/Cherokee Nation) provided recommendations about what UxS were best able to collect various types of data (Appendix E). He specifically noted the importance of having access to manual camera controls, especially shutter speed.

Throughout this workshop and during break-out group discussions, common issues emerged as almost ‘universal challenges’ pertaining to various data collection themes (e.g., for photogrammetric, enumeration, or individual identification studies). Overall, there is no one-size-fits-all solution for study design and it is important that each new technology or novel approach be assessed for the specific study site, species, habitat, etc. There are some Science Center teams who are already successfully implementing innovative approaches for data collection using UxS. As is commonly the case with developing novel approaches in natural resource areas, lack of funding for development and implementation causes setbacks and delays. Most programs or divisions do not have funding to support research and development of data collection systems, a better coordinated approach of posting an up-to-date list of external funding options that can be shared and circulated regularly would invaluable.

One key data collection challenge is that NOAA employees are only permitted to use UAS platforms that have been approved as airworthy by NOAA Aircraft Operations Center (AOC). The list of approved platforms is limited, and approved systems do not always have the capabilities required for certain projects. If there isn’t already a platform that has been approved, this could extend the research and development of a project, in order to research and adequate platform to propose to AOC to go through the airworthiness process. If there were a more streamlined approach and guidelines for biologists to follow in order to initiate an evaluation of new platforms, this could greatly expand the agency’s ability to meet our mission needs.

Finally, issues associated with transferring, storing, summarizing, and organizing data from UxS missions was another common issue reported at the Workshop. This relates to data collected from sensors, as well as UxS platform data and performance information. Manually transferring data collection products from one storage medium to another is time consuming, and, if issues arise, can result in data loss, and requires verification that all the data has been completely and accurately transferred. A centralized understanding of various data products and methods for collection, management, and storage would be beneficial. This would allow biologists to the opportunity to share various techniques, storage technologies, and computer code to automatically retrieve, sort, and validate data. This would also be helpful for sharing techniques for evaluating UxS platform data (e.g., UAS flight log processing to produce summaries for Situation Reports, etc.). This is another example where increased communication, coordination, and collaboration between researchers would greatly improve efficiency and reliability. Using a cloud storage provider also streamlines access to running data science competitions to develop algorithms.

GUIDANCE FOR LAUNCHING AN AI PROJECT

Evaluating whether AI is the right tool

Some tasks that seem simple to a human may be quite challenging for a computer, but, as a general rule, if a human can classify the images, it's likely that AI could as well. For example, automatically counting animals in an image can be broken down into two different computer vision problems: semantic segmentation (separation of animals from one another or from the background), and classification ("sorting" identified animals into species or age-sex classes). A clear understanding of the AI task at hand is vital, especially when looking to work with external partners to develop computer vision algorithms.

While full automation is highly desirable, it may take more time, effort, and resources than available. Instead, implementation of AI into some steps of a research project may be more achievable and can still provide tremendous value to a research project. Focusing AI development efforts on solving repetitive tasks on a large volume of data may deliver great benefits to the project. This sort of semi-automated approach could be ideal for those instances when full-automation is not yet achievable. Semi-automated approaches allow AI to work for you to make the initial cut of the data (with a measure of uncertainty), and then enable humans to review a fraction of the data, for difficult cases).

Though it may be tempting to use AI for all projects and problems, there may be more cost effective and efficient methods to process the data. In some cases, it may be possible to get equally accurate results without large investments of time and resources for the development of AI. Consulting with experts who have used AI or other methods to address similar problems and reviewing literature is the best way to make this determination. Always keep the overall goal of the project in mind.

Major stages of an AI project

Workshop participants identified five main stages of an AI project: *scoping*, *data preparation and annotation, model selection-training-testing*, *model evaluation and re-training*, and *deployment and integration*. A full table of major milestones and tasks is provided in Appendix H. Throughout the course of the project, maintaining good documentation is key. A clear understanding of objectives, requirements, system parameters, techniques, and responsibilities will help reduce confusion and waste.

In the *scoping* phase, team members should discuss the problem at hand, whether AI is the right tool for the job, budgeting, resources (e.g., time, labor, expertise, computing hardware), and software necessary to be successful. Consulting with experts and reviewing what has been done by other teams will help this happen smoothly. These conversations should occur while determining whether AI is the right tool. In addition, consulting a statistician on the required accuracy rate of the AI system is important as some degree of error can be accounted for with statistical measures.

The *data preparation and annotation* phases are focused on creating a clean dataset of annotated images that will be used to train the AI model. In essence, an annotation is a human-created classification that gives the AI information that it can learn from: by identifying and providing information on what class an object is, we can train the AI. Creating quality annotations can be costly in terms of both labor and time, yet is extremely important: poor quality annotations will likely result in poor model performance. The size of a training set also varies depending on the imagery, objective, and method: some models may only need a training set of a few hundred examples, while others may require thousands. Many machine learning approaches require very large datasets of thousands of annotated images, especially for the more sophisticated deep learning approaches. It is important that a percentage of the annotated images is withheld for algorithm evaluation (commonly 25%) and this testing data set should not be given to the algorithm developers until the final evaluation stage.

The *model selection-training-testing* and *model evaluation-re-training* stages focus on testing different approaches and reviewing metrics on accuracy/error rates. If the error rate is too high, it may be necessary to make changes to the annotation process or change models and approaches. Once the AI algorithm has achieved the required accuracy rate, the AI system can be *deployed and integrated* into research projects.

Competitions as a source of solutions

Competitions can be a great way to engage AI experts, software engineers, and expertise from a diverse field of disciplines. Often, individuals participating in competitions do so for bragging rights, learning opportunities, rankings, and credentialing. As such, the end result of a competition could exceed the quality of a product procured by traditional means. Having a large number of participants from a wide field of disciplines trying different approaches can help identify ones that may be successful in solving the problem.

Participation is key in a competition, and there are several things that can be done to help attract competitors and yield the desired results. A clear, well defined, and interesting problem should be posed. Clean data that is easily accessible and ready to process helps projects run smoothly and reduces the need for competitors to focus on data cleaning tasks instead of the main challenge. AI development is only part of the overall process: delivery of a great algorithm may still need additional investment for successful integration and implementation. It may be necessary to work with a software engineer and/or a machine learning data scientist to make the newly developed solution usable by biologists.

Often, competitions will require the project team to release a significant amount of data. Carefully consider the source of the imagery and any special permissions that are associated with that data.

Limitations of AI models and training sets

Any time a model is trained on one set of conditions and then applied to another, lower accuracy and greater bias are likely. For example, a computer vision model that is able to accurately detect and classify animals on rocky surfaces will likely suffer when presented with animals on ice. Changes to the scale of an image induced by changes to sensors (cameras) or survey altitude may also have detrimental effects. Depending on the model and underlying algorithms, these changes may have different impacts on the accuracy of the AI. Changes to the environment and methodologies may also necessitate the need to retrain the AI with new data, or trigger creation of a whole new algorithm and workflow. As the world changes, so must the AI; scientists must consider what kinds of changes would trigger the need for modification to the AI after implementation. New AI algorithms may need to be phased in, replacing old ones. Triggers may not always be obvious: new advances in AI development may provide new opportunities for greater automation. The field is rapidly advancing and solutions only a year or two old may be considered outdated. It is important to point out that constant incremental improvements in methodology may not be the correct approach as they can bias data outputs for certain scenarios (e.g., abundance surveys where incremental changes could bias trend), but more infrequent full re-analyses of the datasets may on occasion be required and a more optimal approach for updates.

There are some important considerations in determining the ideal trade-off in model specificity and generalizability. Detector models, such as those performing presence/absence are typically easily generalizable. Specific classifiers (distinguishing one species or individual from another) are less generalizable and likely require retraining for individual projects. Algorithms from contests, academic partners, or publicly available repositories may not be ready for use for a specific purpose. Issues may arise in accuracy and scalability: some models may be trained for one type of data and will struggle with novel data, and some proofs of concept may not perform well on larger datasets.

Roles and responsibilities of an AI project team, and recommended training

The incorporation of AI into a workflow introduces a new level of complexity into existing projects already struggling to manage large quantities of data. Depending on the project, team members with new types of expertise may be needed, and biologists will require training in new skills to meet the challenges of this new approach. The following describes the different roles of a hypothetical "team" with the responsibility for deciding and implementing a new AI approach for a project, and the types of training recommended for each role. Note that on many teams, single individuals will have several roles. Appendix I provides a list of online courses and resources.

Principal investigator:

Leads the project. This individual should have a clear understanding of the required outcome of the project and who the end-user of the project will be. Oversight of a project that involves AI will require a broad understanding of how to use AI and what it can (and can't) contribute to a project. The individual in this role should have training that provides a high-level introduction to AI (e.g., Coursera Machine Learning course, or a course designed for executives or project managers). Ideally, the course should include an introduction to basic Python scripting, which is commonly used to write machine learning algorithms.

Advocate:
The individual with this role has to be able to “sell” the project in order to secure monetary, staffing, and technical resources needed to do the work. The advocate should understand AI sufficiently to explain the needs of the project to a moderately-technical audience, either internal or external to the agency.

Domain expert/biologist:
Expert in the scientific question being addressed; should have biological knowledge of the species of interest, relevant ecosystems, and background/context about the scientific question that needs to be addressed. The individual with this role should have more in-depth training in AI to ease communication between AI experts and other team members; courses such as Coursera Deep Learning and AI workshops at conferences and meetings (e.g., American Geophysical Union Fall Meeting or Conference on Computer Vision and Pattern Recognition) are good places to gain this level of expertise.

Project manager:
The addition of AI to a project will significantly increase the project’s complexity, and it is recommended that someone be designated the role of project manager. This role involves identifying milestones and ensuring that milestones and objectives are met on time and within budget, and making sure that the biologists, IT resources and staff, the AI project lead, systems engineers, etc. are working together to meet the goals of the project. This role requires substantial attention to detail, ability to communicate effectively with both team members and managers, and time/budget management skills. The individual will be expected to manage project scope creep, ensure that data agreements and/or contracts are in place, and understand significant technical aspects of the project (e.g., data volume involved, where it needs to be transferred, etc.). In addition to basic training in AI, the individual in this role should have professional project management training, which is available through multiple vendors.

AI coordinator:
The individual in this role would advise the project team about at what stage in the workflow AI would be most useful, and what type of AI approach should be pursued. This individual should have considerable AI technical expertise gained through training, relevant experience, or both. One AI coordinator could be shared between multiple AI projects.

Systems engineer:
This role ensures and verifies that various data collection instruments, computer hardware, and software are working and providing expected results throughout the project. The individual should have a broad skill set, an understanding of the project goals, underlying biology and principles, and the ability to troubleshoot equipment being used during the project. This individual should have training in programming, basic computer vision and AI, computer repair and troubleshooting, and networked computing resources. In addition, experience with the survey platform, deep domain-specific knowledge, and general engineering and software development principles.

Data annotator:

The individual with this role must review and annotate image data to make it useful for AI by drawing bounding boxes, labeling key points/pixels, or otherwise marking the image data. Annotation is a critical step because it may impact model results, but may be tedious due to the number of images needed for a successful AI project. The process requires considerable patience, ability to work with others looking at the same images to agree on an annotation approach, and ability to mark image attributes at different levels (e.g., object level and pixel level). Biological domain knowledge may or may not be necessary depending on the problem to be solved (e.g., if the goal requires differentiating multiple similar species, expertise in biology may be necessary; if the goal requires assessing presence/absence of an object expertise in biology may not be required). In some cases, training for this task is best provided by project staff familiar with the images and the goal of the project. If the project does not require scientific expertise, then it may be more efficient to outsource the annotation work. Alternatively, external annotators can do a first pass, which can then be verified by NOAA biologists.

Algorithm developer:

Algorithms will typically be developed by data scientists outside the agency. These individuals have extensive training and experience in writing code to solve problems and are typically found at universities or in private industry. Expertise is gained through university training and individuals will have a range of credentials, up to and including a Ph.D. in a technical field. Individuals with this role may write code over a short period of time and not expect to support the project over the long term, thus requiring software engineer support. Project managers should be clear about licensing the product as open source in early stages of the project. Data scientists with the skillset to write machine learning algorithms are highly sought after and very expensive to hire.

Software engineer:

The individual or team with this role is responsible for designing the user interface, engineering software, and developing backend infrastructure to support the algorithm deployment. This role may be internal or external to the agency. The individual or team needs to be familiar with how to build software applications for both the cloud and for a desktop, manage servers for application hosting, storage and transfer of images, and connect algorithms running on GPU systems to the workflow pipeline. In addition, this individual or team must perform quality assurance and quality control testing to identify, resolve, and validate fixes for software bugs, gaps in functionality or features, and work closely with project members to ensure that the product delivered meets project needs.

Data scientist:

The individual with this role curates data, trains models, conducts data quality control, and may also develop algorithms. This field is rapidly becoming an independent professional specialty. Expertise required includes a background in computer science, database design and manipulation, applied statistics, and programming in R and Python. Training is offered through degree programs, intensive bootcamps, which are offered at major universities and private companies, such as General Assembly and many others. A basic understanding of the field can be gained through various online systems available (e.g., Lynda.com or Coursera).

Data manager/Data librarian:

The individual with this role tracks data, develops/manages metadata, oversees people conducting data annotation prior to AI work, manages retention of data, and data access.

Quality Assurance/Quality Control:

The individual with this role should evaluate each step in the data collection, annotation, and AI process and ask basic questions about whether the system is doing what is expected from end to end. Beyond the traditional role of a Quality Assurance/Quality Control (QA/QC) engineer, this position would be performing checks during the complete project life cycle. This QA/QC role is critical; when working with external partners or contractors, there is the opportunity for confusion or misinterpretation of terminology, or objective, especially when multiple disciplines are involved. This role may be the same as the data scientist and is needed for any system that is being developed (not AI-specific).

The individual tasked with quality control and assurance should evaluate the inputs and outputs of each step of any process, data-collection system, or software, and ensure that they are performing as expected. Additionally, there is a risk of training bias into an AI model. This potential risk should be identified by QA/QC and assessed by the biologist to identify and quantify the impact of this bias. Uncertainty and error should also be quantified and provided to the statistician. Whenever a cross-disciplinary or cross-organizational juncture exists, there is the possibility for confusion or misinterpretation of intent, terminology, and expectations. This individual within the project team is responsible for checking the quality of produced products to ensure that they meet research requirements.

For example, this individual would be responsible for ensuring that a solution built by software engineers performs as required by scientific researchers and delivers data products correctly in the expected output format. During the annotation phase, they would need to perform checks on the quality and accuracy of annotations, regardless of whether the annotations are produced within the agency or not. If issues arise, they can identify and escalate the issue and ensure that it is resolved correctly.

Statistician:

The individual with this role should be involved in the project from its inception to advise on data collection, QA/QC of datasets, and what data should be used and withheld for algorithm training and testing. It is critical that the statistician understand the model sufficiently to be able to quantify the uncertainty in the final product (e.g., the CV of a population abundance estimate) and be involved in model testing. As algorithms are developed, retired, and improved over time, statisticians will need to understand the implications of these changes to the level of uncertainty in the final products (e.g., a marine mammal abundance estimate) and be able to communicate this effectively to agency managers. The individual should have a basic understanding of AI similar to that of the PI or project manager.

Contracting/acquisitions:
Many AI projects will involve some type of acquisition. We recommend that a small number of acquisitions professionals be designated to process AI contracts so they develop a familiarity with new processes, terminology, qualified vendors, etc. and are better able to work with the team to procure services. In the short term, a portion of NOAA Fisheries' need for a crowdsourcing platform may be filled by a new interagency agreement with NASA. However, NOAA acquisitions should begin developing the expertise needed to effectively process and communicate about the types of contracts needed for AI.

Data architect:
The individual with this role provides the big-picture design of the software and database system used for the project. The individual must have database and domain skills and a big-picture understanding of the project.

IT support:
The individual with this role must know how to store, share, and provide secure access to data, and have expertise in managing large volumes of data needed by various team members and understand how to move large volumes of data to external partners.

EXISTING SOFTWARE

VIAME

VIAME (**V**ideo and **I**mage **A**nalytics for **M**arine **E**nvironments) is an open-source software framework designed to reduce the barriers to development and implementation of AI models. Developed by Kitware Inc., under NOAA's Automated Image Analysis Strategic Initiative (AIASI; Richards 2015, Richards et al. 2019), VIAME enables rapid, low-cost integration of new algorithmic modules and datasets, enabling adaptability to new workflows. VIAME is comprised of a collection of GUI and command-line tools that offers capabilities for object detection, imagery annotation, training and running AI models, inter-frame object tracking (primarily for video analysis), and stereo measurement. In addition, relatively recent developments include iterative query refinement (IQR), which allows for rapid development of models, as well as new capabilities for server-based image processing.

Within NOAA Fisheries, several research programs have successfully utilized these capabilities: IQR allowed researchers at the SWFSC to rapidly develop a full-frame classifier to identify frames of interest collected by an animal-borne video camera, greatly reducing the amount of time and effort required to identify these frames (Hinke, SWFSC). Similarly, researchers at the AFSC were able to integrate and test various detection and classification models developed by external partners (Moreland, AFSC). In addition, another AFSC group is working with Kitware to develop an end-to-end program to automate image processing and counts of Steller sea lions from aerial imagery (Sweeney, AFSC).

A key benefit of VIAME is that it provides a GUI interface: many common AI model development tools, such as Tensorflow, Keras, and SciKit Learn are primarily controlled via programming languages and have a high learning curve. The modular framework of VIAME allows for software engineers to delve deeply into model development while providing a friendly interface for researchers and analysts. Participants of a VIAME training session reported that developing models was quite straightforward.

VIAME can significantly reduce the barriers to implementing AI for research activities but challenges still remain. Advanced techniques like fusion of multi-spectral imagery may require engineering time to add capabilities. In addition, models obtained from other collaborators, contractors, or contests may need to be modified in order to work correctly within the VIAME framework. While specifications can be provided, software engineering efforts may be required on behalf of the GUI designer to modify a third-party model to accept proper inputs and provide outputs in the correct format.

At this time, VIAME is minimally supported by NOAA Fisheries now that the strategic initiative has ended. VIAME is easy to use when the AI project closely matches an existing module already included in the system, but VIAME does not have many modules that are relevant to marine mammal researchers (e.g., counting animals on a beach or on ice). VIAME is being modified to be of more use to marine mammal researchers, but progress is slow because funds must be found from outside NOAA Fisheries through competitive requests for proposals; substantial additional support would be needed for VIAME to develop a suite of modules that would be relevant for common marine mammal applications.

WildBook

Wildbook is an open source software framework designed to support collaborative, crowd-sourced wildlife and ecological studies. Developed by WildMe, the WildBook framework is designed to be run on web platforms (e.g., web servers, cloud computing services) to allow for easy collation of large volumes of data from multiple participants. WildBooks aims to be broadly accessible by scientists, research and industry partners, government agencies, and the public. A primary objective of the platform is to assist data curation, freeing up resources to focus on analysis of individuals and populations. Capabilities for filtering and aggregating data allow for extraction of data for specific analyses.

The platform uses AI for both imagery processing and integration of new sources of data. Flukebook, an implementation of a WildBook, is a prime example of these capabilities. Multimodal hierarchical algorithms identify species and employ different algorithms to make a determination of an individual where possible. “Tweet a Whale” provides a convenient way for the public to submit imagery for analysis via Twitter: if an individual is successfully identified, additional spatial-temporal data are also made available via a link. Access to datasets can be restricted via opt-in or peer-approval mechanisms. In addition, data are stored in a standard format, and specific queries and results can be shared with web links, or extracted for more advanced analyses.

Jason Holmberg (WildMe) identified several challenges with the WildBook platform and with AI in general. Current work on WildBooks are focused on increasing the performance of the system, especially when multiple collaborators are using the system simultaneously. Resource constraints, especially for running algorithms on GPU servers, also pose challenges when processing a large amount of data.

As is the case with any AI platform, there are challenges for developing and integrating new models and retiring old models. AI, models, and algorithms sourced from external sources will likely require additional software engineering time and resources to properly integrate into an implementation of a WildBook. In addition, adding project-specific functionality and features requires modification to the system and additional investment for software engineering. Development of GUIs and features would require a great deal of communication between software engineering teams and project teams to ensure that the delivered product is as expected. While it is designed for use in cloud environments, there are no serious technical barriers to running "local" instances of WildBook within an agency.

Conceptually, the WildBook platform differs substantially from VIAME in that WildBooks is an end-to-end solution for projects while VIAME is more of a "standalone" solution for image analysis. The magnitude of change required to incorporate a WildBook solution for a research program is likely quite high relative to the investment required for VIAME. The collaborative capabilities of the platform presents significant benefits, though they need be weighed against elevated customizability costs.

FinFindR

FinFinR is a program created by the WEST group, led by Jaime Thompson. FinFindR uses a multiphase approach to identify bottlenose dolphins using natural markings on their dorsal fins. During the first phase, the program identifies and crops images to isolate the dorsal fins. In phase two, a CNN highlights key features of a fin and traces the dorsal fin to isolate the fin edges by tracing the edges based on contrast between the fin and the background. In phase three, an algorithm matches a given fin to others like it in the catalogue. The WEST group has also used AI for correcting counts of bats using thermal video data.

CONCLUSION AND NEXT STEPS

Participants in the workshop agreed that meeting others in the agency involved in using AI to improve the processing and analysis of large volumes of data on marine mammals was very helpful. Feedback from participants indicated particular appreciation for the presentations about how marine mammal researchers internal and external to NOAA Fisheries were using AI, involvement of AI experts from private industry in the discussion, the discussion about the steps required to get started on an AI project, demonstrations and training in AI tools (VIAME, WildBooks) and software needed to process UAS imagery (Pix4D), and the internal and external networking opportunities (Appendix I).

However, participants acknowledged that much progress is needed to effectively use AI to process marine mammal image data. Feedback from participants noted that we need a much better understanding of the AI approach in general to know when it will create efficiencies (and when it won't), and what cloud services are available within NOAA and how to access those services.

The need for considerable future discussion, training, and in-depth exploration of various topics was clear throughout the workshop. Participants identified the following topics for future workshops:

- Methods-oriented discussions that involve both the fisheries and marine mammal research communities using AI in NOAA Fisheries.
- How to manage “big data”, specifically data annotation for AI, AI for individual identifications from photo ID databases.
- How to write a successful contract for AI in detail and how to evaluate products.
- A “how to” workshop on computer vision for biologists with a session on results validation and what types of information to present in publications.
- One-on-one sessions with experts in private industry involved in AI (WildMe, Vulcan, Microsoft, Pix4D Structure from Motion, others) to understand how they use AI to process data and learn how the approach can be applied to data in NOAA Fisheries.
- Advice on how to pick a cloud service provider, data types/size considerations, costs, and which options are supported by NOAA.

In addition, workshop participants identified specific training that would benefit individuals with various roles in AI projects in the agency. Training will be needed in technical topics, such as basic and advanced coding in R, Python, and other programming languages that are commonly used for AI applications and development, as well as business topics, such as formal project management training, which is recommended for AI projects because of the complexity of working with multiple teams with different expertise, deadlines and many project dependencies (see *Gudiance for Launching an AI Project* section for additional training recommendations for various roles required for an AI project).

Most participants who provided feedback on the workshop indicated that they now see new ways to approach a research problem and developed new connections within and outside NOAA Fisheries that are likely to lead to future collaborations.

Further, the workshop resulted in many specific, actionable recommendations that we expect will improve the speed and effectiveness with which NOAA Fisheries researchers can explore and adopt new AI approaches to data processing (Appendix D). These recommendations may be implemented at the Center, Line Office, or NOAA level. Many recommendations are focused on improving connectivity between researchers in NOAA Fisheries implementing AI projects, which participants expect will significantly improve the speed at which AI is incorporated into agency image processing workflows.

CITATIONS

Angliss, R. P., M. C. Ferguson, P. G. Hall, V. T. Helker, A. S. Kennedy, and T. Sformo. 2018. Comparing manned to unmanned aerial surveys for cetacean monitoring in the Arctic: Methods and operational results. J. Unmanned Veh. Syst. doi: 10.1139/juvs-2018-0001.

Bogucki, R., M. Cygan, C. B. Khan, M. Klimek, J. K. Milczek, and M. Mucha. 2018. Applying deep learning to right whale photo identification. Conserv. Biol. doi: 10.1111/cobi.13226.

Christie, K. S., S. L. Gilbert, C. L. Brown, M. Hatfield, and L. Hanson. 2016. Unmanned aircraft systems in wildlife research: Current and future applications of transformative technology. Ecol. Soc. Am. doi: 10.1002/fee.1281.

Conn, P.B., J. M. Ver Hoef, B. T. McClintock, E. E. Moreland, J. M. London, M. F. Cameron, S. P. Dahle, and P. L. Boveng. 2014. Estimating multispecies abundance using automated detection systems: Ice-associated seals in the Bering Sea. Methods Ecol. Evol. 5(12): 1280-1293. doi: 10.1111/2041-210X.12127.

Ferguson, M.C., R. P. Angliss, A. Kennedy, B. Lynch, A. Willoughby, V. Helker, A. A. Brower, and J. T. Clarke. 2018. Performance of manned and unmanned aerial surveys for estimating arctic cetacean density and associated uncertainty. J. Unmanned Veh. Syst. doi: 10.1139/juvs-2018-0002.

Fritz, L., K. L. Sweeney, D. Johnson, M. Lynn, T. Gelatt, and J. Gilpatrick. 2016. Aerial and ship-based surveys of Steller sea lions (*Eumetopias jubatus*) conducted in Alaska in June-July 2013 through 2015, and an update on the status and trend of the western distinct population segment in Alaska. U.S. Dep. Commer., NOAA Tech. Memo, NMFS-AFSC-321, 72 p

Hinke, J.T., A. Barbosa, L. M. Emmerson, T. Hart, M. A. Juáres, M. Korczak-Abshire, G. Milinevsky, M. Santos, P. N. Trathan, G. M. Watters, and C. Southwell. 2018. Estimating nest-level phenology and reproductive success of colonial seabirds using time-lapse cameras. Methods Ecol. Evol. 9:1853–1863.

Khan, C., A. Henry, P. Duley, J. Gatzke, L. Crowe, and T. Cole. 2018. North Atlantic right whale sighting survey (NARWSS) and right whale sighting advisory system (RWSAS) 2016 results summary. U.S. Dep. Commer., Northeast Fish. Sci. Cent Ref. Doc. 18-01, 13 p.

Melancon, R. A. S., S. Lane, T. Speakman, L. B. Hart, C. Sinclair, J. Adams, P. E. Rosel, and L. Schwacke. 2011. Photo-Identification field and laboratory protocols utilizing FinBase version 2. U.S. Dep. Commer., NOAA Tech. Memo, NMFS-SEFSC-627, 46 p.

Moreland, E.E., M. F. Cameron, R. P. Angliss, and P. L. Boveng. 2015. Evaluation of a ship-based unoccupied aircraft system (UAS) for surveys of spotted and ribbon seals in the Bering Sea pack ice. J. Unmanned Veh. Syst. doi: 10.1139/juvs-2015-0012.

Pace, R. M., P. J. Corkeron, and S. D. Kraus. 2017. State–space mark–recapture estimates reveal a recent decline in abundance of North Atlantic right whales. Ecol. Evol. 7(21): 8730-8741. https://doi.org/10.1002/ece3.3406.

Richards, B. L, A. B. Albu, E. Clarke, G. R. Cutter, E. Edginton, D. Hart, A. Hoogs, D. Kreigman, C. Kunz, M. Piacentino, L. Prasad, C. Thompson, K. Williams, and W. Michaels. 2019. NOAA Fisheries strategic initiative on automated image analysis. [accessed 14 Jan. 2020]. https://www.st.nmfs.noaa.gov/aiasi/Home.html .

Richards, B. L. 2015. NOAA Fisheries Strategic Initiative on Automated Image Analysis, p. 4-20. In M. Mellody (ed.), Robust Methods for the Analysis of Images and Videos for Fisheries Stock Assessment: Summary of a Workshop. The National Academies Press. Washington, DC. doi.org/10.17226/18986.

Sweeney, K. L., B. Birkemeier, K. Luxa, and T. Gelatt. 2019. Results of Steller Sea Lion Surveys in Alaska, June-July 2019. U.S. Dep. Commer., NOAA, NMFS, Alaska Fisheries Science Center, Marine Mammal Laboratory Memorandum to The Record. 6 December 2019. https://www.fisheries.noaa.gov/resource/data/2019-results-steller-sea-lion-surveys-alaska .

APPENDICES

APPENDIX A. Workshop Agenda

Monday, September 16: Symposium - Building 3; Oceanographer Room

8:00 – 8:30	*Coffee / Tea*	
8:30 – 8:45	Welcome and introduction to challenges	Angliss (NOAA Fisheries)
8:45 – 9:00	Key definitions	Khan (NOAA Fisheries)
9:00 – 9:45	**Approaches to image processing**	**Moderator: Sweeney**
	• Citizen scientists help researchers investigate Steller sea lions • Computer Vision - practice and pitfalls • Machine learning and deep learning differ in effort Use satellites to inform drone work	Sweeney (NOAA Fisheries) Redmon (UW) Johnston (Duke, remote)
9:45 – 10:00	*Break*	
10:00 – 11:00	**Approaches to image processing (continued)**	**Moderator: Moreland**
	• AI/ML for wildlife conservation and mythbusting for biologists • UAS applications with bottlenose dolphins, birds and bats • Using AI to identify whales from visually represented acoustic data *15-minute session discussion*	Morris (Microsoft) Thompson / Erickson (WEST) Allen (NOAA Fisheries)
11:00 - 12:00	**Use of AI/ML/etc. for counting animals in images**	**Moderator: Khan**
	• Applications of Machine Learning Algorithms to Automate Data Extraction from Images • Automated surveys for ice-associated seals in the Arctic • Collection and analysis of imagery and video by the SWFSC Antarctic Ecosystem Research Division *15-minute session discussion*	Altukhov (NPW) Moreland (NOAA Fisheries) Hinke (NOAA Fisheries)
12:00 – 13:15	*Lunch*	

***Continued*—Monday, September 16: Symposium - Building 3; Oceanographer Room**

13:15 – 14:30	**Use of images for photo-identification and photogrammetry**	**Moderator: Khan**
	• Computer vision for conservation: Automating right whale photo ID • Flukebook: Multi-modal, multi-stage machine learning for marine mammal research with citizen science • Species identification and stereo measurements • Morphometrics and volumetrics of pinnipeds from imagery *15-minute session discussion*	Khan (NOAA Fisheries) Holmberg/Parham (WildMe) Lauffenberger (NOAA Fisheries) Shero (WHOI, remote)
14:30 – 14:45	*Break*	
14:45 – 15:45	**Statistical considerations**	**Moderator: Angliss**
	• Accounting for species misclassification • Estimating abundance with automated detection systems • Notes on statistical considerations for photoID mark-recapture *15-minute session discussion*	McClintock (NOAA Fisheries) Conn (NOAA Fisheries) Conn (NOAA Fisheries)
15:45 – 17:00	**Hardware, processing, and storage**	**Moderator: Angliss**
	• Network effects: storage, processing, and connections • Moving to the cloud • Matching UAS platforms to your data objectives *15-minute session discussion*	Hou (NOAA Fisheries) O'Neil (NOAA Fisheries) Seymour (USGS/Cher. Nation)
Optional evening event - Magnuson Café & Brewery (5:30 - 7:30pm)		

Tuesday, September 17: Introduction to VIAME and discussion - Building 3 - Oceanographer Room

8:00 – 8:30	*Coffee*
8:30 – 9:45	Introduction to VIAME (Dawkins/Kitware; 1hr with a 15min discussion section)
9:45 – 10:00	*Break*
10:00 – 11:00	Discussion about presentations - what approaches are working and why?
11:00 – 12:00	How to evaluate sensors, platforms, and post-processing techniques
12:00 – 13:15	*Lunch*
13:15 – 14:15	Introduction to Flukebook: Basic concepts (Holmberg/WildMe)
14:15 – 14:30	*Break*
14:30 – 16:00	**Break-out group discussions:** ● Photogrammetry (moderator: Richmond) ● Enumeration/abundance (moderator: Krause) ● Photo-ID (moderator: Khan) ● Video analysis (moderator: Sanderson)
16:00-16:30	Report-out and discussion

Wednesday, September 18, morning: Discussion - Building 3 - Oceanographer Room

8:00 – 8:30	*Coffee*
8:30 – 9:45	Recommendations for future research investments
9:45 – 10:00	*Break*
10:00 – 11:00	Roles/responsibilities of a team designing/implementing an AI/ML project involving imagery
11:00 – 12:00	NASA approach to contracting for competitions (Rader/NASA)
12:00 – 13:15	*Lunch*
13:15 – 14:30 Parallel sessions	**Internal NOAA Fisheries/NOAA strategy session (CLOSED)**
	Internal NOAA Fisheries initiative process & where NOAA is heading with AI/ML (special guest, Bill Michaels, NOAA Fisheries F/ST)
	How can NOAA strengthen external partnerships?
	The NASA approach to contracting AI/ML (special guest, Steve Rader, NASA)
	What automation/AI/ML issues (legal, organizational, other) need to be addressed at a higher level in the agency?
	Develop/prioritize recommendations for future investments in AI/ML
	Industry break-out session: Insights that may improve future agency success with AI/ML (OPEN)
	Given the examples of NOAA Fisheries' projects discussed at this meeting, what types of additional expertise should we seek?
	Do external experts see anything "cringeworthy" about how we are designing and implementing our projects?
	How can public/private partnerships be strengthened? What barriers to developing successful partnerships have you encountered in the past?
	What degree of AI/ML proficiency is required at various levels in the agency, and where can training be found?
	How can biologists write successful statements of work with sufficient detail/metrics to get a great AI partner?
14:30 – 14:45	*Break*

14:45 – 15:45	What do you need to get started? Beginners guide to successful AI/ML development.
15:45 – 16:45	Report from industry break-out session
16:45 – 17:00	Action items and next steps

Thursday, September 19 - Building 3 - Oceanographer Room

Time	Description	Location
8:30 – 5pm	Introductory VIAME training	Bldg 3, Oceanographer Rm

Friday, September 20 - Building 4 - Rooms 2011 & 2039

Time	Description	Location
8:30 – 5:00pm	One-on-one consultations with Kitware about VIAME/Matt Dawkins	Bldg 4, room 2011
8:00 – 12:00	Structure-from-motion training/Alex Seymour	Bldg 4, room 2039
1:00 – 5:00	Machine learning using R/Alexey Altukov	Bldg 4, room 2039

APPENDIX B. List of Participants

* Workshop organizers (*see footnote for affiliation acronym definitions in order as they appear*)

Name	Affiliation	Email
National Oceanic & Atmospheric Administration		
Alexey Altukhov	NPW & AFSC, NOAA Fisheries	alexey.altukhov@noaa.gov
Robyn Angliss*	AFSC, NOAA Fisheries	robyn.angliss@noaa.gov
John Bengtson	AFSC, NOAA Fisheries	john.bengtson@noaa.gov
Burlyn Birkemeier	CICOES & AFSC & NOAA Fisheries	burlyn.birkemeier@noaa.gov
Peter Boveng	AFSC, NOAA Fisheries	peter.boveng@noaa.gov
Michael Cameron	AFSC, NOAA Fisheries	michael.cameron@noaa.gov
Cynthia Christman	CICOES & AFSC & NOAA Fisheries	cynthia.christman@noaa.gov
Paul Conn	AFSC, NOAA Fisheries	paul.conn@noaa.gov
Shawn Dahle	AFSC, NOAA Fisheries	shawn.dahle@noaa.gov
Brian Fadely	AFSC, NOAA Fisheries	brian.fadely@noaa.gov
Nancy Friday	AFSC, NOAA Fisheries	nancy.friday@noaa.gov
Tom Gelatt	AFSC, NOAA Fisheries	tom.gelatt@noaa.gov
Stacie Hardy	AFSC, NOAA Fisheries	stacie.hardy@noaa.gov
Diana Haring	CICOES & AFSC & NOAA Fisheries	diana.haring@noaa.gov
Benjamin Hou*	AFSC, NOAA Fisheries	ben.hou@noaa.gov
John Jansen	AFSC, NOAA Fisheries	john.jansen@noaa.gov
Devin Johnson	AFSC, NOAA Fisheries	devin.johnson@noaa.gov
Carey Kuhn	AFSC, NOAA Fisheries	carey.kuhn@noaa.gov
Brett McClintock	AFSC, NOAA Fisheries	brett.mcclintock@noaa.gov
Molly McCormley	CICOES & AFSC & NOAA Fisheries	molly.mccormley@noaa.gov
Sharon Melin	AFSC, NOAA Fisheries	sharon.melin@noaa.gov

Name	Affiliation	Email
Erin Moreland*	AFSC, NOAA Fisheries	erin.moreland@noaa.gov
Erin Richmond*	CICOES & AFSC & NOAA Fisheries	erin.richmond@noaa.gov
Stacie Robinson*	PIFSC, NOAA Fisheries	stacie.robinson@noaa.gov
Christy Sims	CICOES & AFSC, NOAA Fisheries	christy.sims@noaa.gov
Katie Sweeney*	AFSC, NOAA Fisheries	katie.sweeney@noaa.gov
Rick Towler	AFSC, NOAA Fisheries	rick.towler@noaa.gov
Paul Wade	AFSC, NOAA Fisheries	paul.wade@noaa.gov
Janice Waite	AFSC, NOAA Fisheries	janice.waite@noaa.gov
Lisa Conger	NEFSC, NOAA Fisheries	lisa.conger@noaa.gov
Elizabeth Josephson	NEFSC, NOAA Fisheries	elizabeth.josephson@noaa.gov
Christin Khan*	NEFSC, NOAA Fisheries	christin.khan@noaa.gov
Elizabeth Clarke	NWFSC, NOAA Fisheries	elizabeth.clarke@noaa.gov
Abi Powell	Lynker & NWFSC, NOAA Fisheries	abigail.powell@noaa.gov
Beth Sanderson*	NWFSC, NOAA Fisheries	beth.sanderson@noaa.gov
William Michaels	Office of Science & Technology, NOAA Fisheries	william.michaels@noaa.gov
Zac Schakner	Office of Science & Technology, NOAA Fisheries	zachary.schakner@noaa.gov
Ann Allen	PIFSC, NOAA Fisheries	ann.allen@noaa.gov
Kym Yano	JIMAR & PIFSC, NOAA Fisheries	kym.yano@noaa.gov
Nathan Lauffenburger	PIFSC, NOAA Fisheries	nathan.lauffenburger@noaa.gov
Kym Yano	PIFSC, NOAA Fisheries	kym.yano@noaa.gov
Anthony Martinez	SEFSC, NOAA Fisheries	anthony.martinez@noaa.gov
Jefferson Hinke	SWFSC, NOAA Fisheries	jefferson.hinke@noaa.gov
Beth Jaime	OAI & SWFSC, NOAA Fisheries	beth.jaime@noaa.gov

Name	Affiliation	Email
Trevor Joyce	SWFSC, NOAA Fisheries	trevor.joyce@noaa.gov
Douglas Krause	SWFSC, NOAA Fisheries	douglas.krause@noaa.gov
Morgan Lynn*	SWFSC, NOAA Fisheries	morgan.lynn@noaa.gov
Jon O'Neil	Systems Management, NOAA Fisheries	jon.oneil@noaa.gov
Philip Hall	UAS Program Office, OAR/OMAO	philip.g.hall@noaa.gov
Industry		
Matt Dawkins	Kitware	matt.dawkins@kitware.com
Dan Morris	Microsoft	dan@microsoft.com
Paul Albee	Vulcan	palbee@vulcan.com
Gracie Ermi	Vulcan	gracieer@vulcan.com
Sam McKennoch	Vulcan	samm@vulcan.com
Wally Erickson	WEST, Inc	werickson@west-inc.com
Jaime Thompson	WEST, Inc	jwthompson@west-inc.com
Jason Holmberg	WildMe	jason@wildme.org
Jason Parham	WildMe	parham@wildme.org
External Scientific Partners		
Kelly Hastings	ADFG	kelly.hastings@alaska.gov
Lauri Jemison	ADFG	lauri.jemison@alaska.gov
David Johnston	Duke University	david.johnston@duke.edu
Steven Rader	NASA	steven.n.rader@nasa.gov
Jamie Womble	National Park Service	jamie_womble@nps.gov
Jessie Lindsay	University of Washington	jmlinds@uw.edu
Alexander Seymour	USGS & Cherokee Nation	alexander.c.seymour@gmail.com

NPW—North Pacific Wildlife Consulting
AFSC—Alaska Fisheries Science Center

CICOES—Climate, Ocean and Ecosystem Studies
PIFSC—Pacific Island Fisheries Science Center
NEFSC—Northeast Fisheries Science Center
NWFSC—Northwest Fisheries Science Center
JIMAR—Joint Institute for Marine and Atmospheric Research
SEFSC—Southeast Fisheries Science Center
SWFSC—Southwest Fisheries Science Center
OAI—Ocean Associates, Inc.
OAR/OMAO—Oceanic and Atmospheric Research/Office of Marine and Aviation Operations
ADFG—Alaska Department of Fish and Game
NASA—National Aeronautics and Space Administration
USGS—United States Geological Survey

APPENDIX C. NOAA Fisheries Science Center Projects Using AI

Alaska Fisheries Science Center

Marine Mammals:

Erin Moreland - The Marine Mammal Laboratory has been investigating automated collection and processing of imagery for several years, and has had significant success (Conn et al. 2014). Researchers from the Polar Ecosystem Program have been developing AI models that are compatible with VIAME for processing and analyzing visual and thermal imagery collected from occupied aircraft to identify ice seals and polar bears. Future goals are to incorporate ultraviolet imagery to improve detection, expand use to terrestrial and glacial harbor seal surveys, miniaturize the system to enable deployment on UAS, and optimize the data collection and data processing workflow.

Katie Sweeney - Researchers from the Alaska Ecosystem Program are working with Kitware to develop a Steller sea lion Automated Count Program using VIAME to count individuals by age-sex class from aerial imagery collected by sUAS and occupied aircraft. This count program will include image processing to avoid double counting sea lions in overlapping imagery. Earlier efforts using Kaggle to produce successful algorithms showed the potential of creating this program. There is also an effort to develop a UAS-based approach for surveying northern fur seals, which will involve using AI for image processing and analysis. Researchers are also utilizing citizen science to process imagery from remote camera images with the Steller Watch project hosted by Zooniverse.org.

Vladimir Burkanov - Researchers from North Pacific Wildlife Consulting are working with the Alaska Ecosystem Program to automate processing of remote camera images to isolate tiles of marked animals in images. This method uses R with Keras and U-Net.

Manuel Castellote - The Cetacean Assessment and Ecology Program has been working closely with Microsoft on automating the classification of beluga whale acoustics. This approach uses visual pattern recognition on extracted spectrograms. The plan is to ultimately transition to a signal-based approach.

Paul Wade - The Cetacean Assessment and Ecology Program is interested in automating the photo-identification of Cook Inlet beluga whales and plans to initiate a project in FY20.

Zooplankton:

David Kimmel, Calvin Mordy, Piotr Margonski, and Eugene Berger - AFSC FOCI is pursuing a collaboration with PMEL researchers in the Innovative Technology for Arctic Exploration (ITAE), the Poland Plankton Sorting and Identification Center, and Google to develop machine learning algorithms for identification of zooplankton samples. The eventual goal is to refine the machine learning algorithm for deployment with a glider for real-time processing of usable images onboard.

Northwest Fisheries Science Center

Marine Mammals:

Blake Feist and Jameal Sahmouri - Researchers have manually processed images downloaded from Flickr® to consider how rates of whale watching have changed over time in Monterey Bay. They used data from the social media platform Flickr®, an online image and video hosting service where users upload photographs with tags for time, keywords, user identification code and location. They developed a Python script that used the Flickr API to query their servers for photographs from 2009 through 2017 associated with whale watching captured within the vicinity of Monterey Bay, CA.

Fish, Shellfish, & Wetlands:

Beth Sanderson - Researchers have been collecting underwater videos to examine species using several types of nearshore habitats such as shellfish aquaculture, eelgrass, and bare sediment. In the past several years, thousands of hours of video have been collected and processed manually using freeware (Boris). In the past 1.5 years, researchers have collaborated with Dr. Jenq-Neng Hwang (UW Engineering) and Microsoft AI for Good/Earth programs to explore ways to 1) note when fish and invertebrates are present in the video and 2) identify the species present. Those efforts are ongoing.

Curtis Roegner - Researchers at the Pt. Adam Research Station and colleagues from OSU have been utilizing underwater video from benthic landers (GoPro cameras) and video sleds (Canon Vixia HF R20 camcorder) to study the effects of dredged sediment disposal on mobile epifauna. Video processing was conducted with Windows Media Player 2010, Adobe Premiere Pro CC, or CyberLink PowerDirector 13.

Curtis Roegner - Researchers from Pacific Northwest National Laboratory and colleagues from RykaUAS are remote sensing wetland habitat by fusing data from an imaging spectrometer and high resolution RBG cameras with LiDAR. All data are acquired using unmanned aerial vehicles. RBG imagery is stitched and used to create a structure from motion (SfM) digital model using Pix4D.

Elizabeth Clarke - The NWFSC has been collecting still images of benthic habitats off the west coast since 2005. The standard analysis protocol is carried out by experts who identify selected fish and invertebrates and measure fish lengths using the software program OneTwoRedBlue. Currently investigations include the potential to use machine learning based analysis using VIAME. In 2017, there was an unusual pyrosome bloom in the NE Pacific and large numbers of pyrosomes were observed on the seafloor in AUV imagery collected off the Olympic Peninsula in Washington. Researchers have been trialing the rapid model generation tools in VIAME to detect pyrosomes in the AUV imagery in order to estimate densities during this event. Researchers are also in the early stages of collaborating with Lynker Analytics in New Zealand to develop a rockfish detection model.

Southwest Fisheries Science Center

Marine Mammals:

Trevor Joyce - Investigating using fixed wing UAS for line transect photo survey and AI to detect animals, primarily cetaceans, in images.

Dave Weller - Using FLIR video imagery and AI to detect whale blows and count animals. This effort has been put on hold due to inability to get it to work with current techniques.

Jim Gilpatrick - Investigating using AI software to auto-match individual gray whales from still images taken vertically from UAS and oblique from hand held cameras.

Doug Krause and Jefferson Hinke - Researchers from the Antarctic Ecosystem Research Division are currently managing two research projects which apply VIAME approaches to: 1) processing imagery of snow, ice, and rocky beaches and detecting Euphausiid swarms in animal-borne video, and 2) filtering animal-borne video to resolve different life stages, and to identify key foraging and behavior events for seals from UAS thermal imaging and seabirds.

Northeast Fisheries Science Center

Marine Mammals:

Christin Khan - Automating the individual identification of North Atlantic and Southern right whales with machine learning through the use of a Kaggle competition and collaboration with WildMe.

Elizabeth Josephson - Exploring the use of VIAME to detect and count harbor seals and gray seals of different age classes from still images collected from manned aircraft and drones.

Genevieve Davis & Sofie Van Parijs - Working to create an automated North Atlantic right whale upcall detector with the Google AI team that built the humpback whale detector through collaboration with PIFSC.

Fish & Shellfish:

Dvora Hart - Using VIAME and convolutional neural networks to detect sea scallops, crabs, skates and other fish in images collected during HabCam towed camera surveys.

Nichole Rossi - High quality fish images were collected by the NOAA bottom trawl survey to cultivate an image library and develop algorithms in support of machine learning applications to enhance the efficiency and accuracy of automated species identification from Electronic Monitoring. Researchers are also leveraging observer data collected in support of species identification for machine learning applications to enhance the efficiency and accuracy of the Observer Species Verification Program.

Julie Rose and Renee Mercado-Allen - Collecting underwater video in shellfish aquaculture habitats and manually processing using EthoVision software.

Southeast Fisheries Science Center

Marine Mammals:

Jenny Litz - The SEFSC conducts bottlenose dolphin surveys in multiple embayments each year using identification of dorsal fin photos as a tool to estimate dolphin abundance. Each survey results in thousands of photos that need to be processed and matched to a catalog of individual dolphins. Researchers at the SEFSC are using FinFindR to automate dolphin dorsal fin matching following photo-identification surveys. FinFindR is an R package that uses a machine learning algorithm to identify and rank potential matches and automate a portion of the photo processing that previously would have been done manually, saving valuable staff time.

Fish, Shellfish, & Corals:

Christian Jones - Leading an effort to detect manta rays from satellite imagery utilizing machine learning products to scan images for classification and further processing**.**

Beverly Barnett - Researchers are engineering a revolutionary approach for improving age determination efficiency in fish using Fourier transform near-infrared spectroscopy. The roadmap to operationalizing the FT-NIRS ageing technology across science centers is envisioned as three major related tasks that are staged at varying, overlapping time frames. These tasks include: (1) application development, (2) application implementation, and (3) stock assessment integration. Our project team, consisting of individuals from AFSC, NWFSC, SEFSC and NEFSC, will be split into west coast (AFSC, NWFSC) and east coast (NEFSC, SEFSC) application development centers. AFSC and SEFSC will serve as the hub for application development. To make this work relevant to already established age data production processes and age data use in stock assessments, the work will focus on three different managed fish stocks from each region with data covering a 5-year time frame. This will provide an opportunity to evaluate FT-NIRS performance across species with differing life history characteristics (e.g., short-lived vs. long-lived), and the inter-annual stability of calibration model parameters.

Phil Caldwell - Facilitating Essential Fish Habitat (EFH) consultation in the Southeast Region by clarifying EFH Boundaries and reducing the need for consultation. The primary goal of this project is to generate a publicly accessible geographic information system (GIS) to improve and update the efficiency of the EFH Mapper in the Gulf of Mexico. This will be accomplished by combining several methods including a probabilistic analysis of the most current NAIP imagery using random forest or softmax function algorithms.

Matthew Campbell - The SEAMAP Reef Fish Video project is using VIAME software (developed from NOAA AIASI) to detect and identify species of reef fish observed in the Gulf of Mexico. These methods will be applied in the RESTORE funded GFISHER survey to begin in 2020. This semi-automated system is still in the testing/supervised phase and has resulted in moderate success. It is probably 5 years from full implementation.

Farron Wallace and David Gloeckner - The Observer Program at SEFSC will be beginning a new AI project this year to develop detectors for our Fisheries Observer Electronic Monitoring program. They will begin annotating image datasets from previous EM research projects soon and will be deploying electronic monitoring into the shrimp and menhaden fisheries late in the year. Based on past experience I am guessing our dataset will be in the 20 - 40 terabyte range and will be stored on the local server. Initially we plan to retrain species and length algorithms previously developed at the University of Washington and AFSC for our fisheries in the Gulf of Mexico.

Turtles:

Chris Sasso - Researchers plan to combine manned and unmanned aerial survey imaging with machine learning in order to count turtles and estimate sizes from aerial video.

Pacific Islands Fisheries Science Center

Marine Mammals:

Amanda Bradford and F. Vivier - Point of contacts on an ongoing collaborative project led by Fabien Vivier (with University of Hawaii) that is looking at age structure and body condition of small odontocetes using UAS-based vertical photogrammetry. Although the current process is still done manually, the aim is to ultimately be able to fly over groups of free-ranging small cetaceans and make length and width measurements on as many individuals as possible. There could be several or tens of individuals (or more) in a frame or series of frames. An AI approach would be extremely useful and make the inference needed to drive management decisions available much more quickly.

Charles Littnan - Currently collecting photos from remote cameras at some remote haul-out locations for monk seals. Cameras are set to take photos every 10-30 minutes, producing datasets of thousands of photos each year. AI has not yet been employed on this project, but it would add great utility.

Stacie Robinson - We collect video using animal-borne cameras. The primary analytical activity with the video is identifying seal behaviors (particularly foraging behavior). AI has not yet been employed on this project yet, but it would add great utility. This could be a fun challenge discerning patterns of movement to categorize dynamic behaviors rather than still images.

Marie Hill - The PIFSC Cetacean Research Program (CRP) collects photo-identification images during cetacean surveys within the Pacific Islands Region. The photographs are organized, analyzed, and matched against catalogs of previously encountered individuals. The processing steps are currently done manually and require many hours/days to analyze. The CRP has proposed a project to improve features in FinFindR, an R package developed by the WEST Group. The FinFindR package uses a machine learning algorithm to identify and rank potential matches using the trailing edge of dorsal fins and was trained using bottlenose dolphins. The CRP uses dorsal fin photo-identification for a variety of delphinid species. Some species have few identifiable features on the trailing edge of the dorsal fin (e.g., spinner dolphins), so we also use features on the leading edge of the dorsal fin for identifying an individual and features on the peduncle of the individual as secondary identifiers. Other species, such as short-finned pilot whales, have dorsal fin shapes that are very different from bottlenose dolphins. We hope to increase the functionality of FinFindR by training the algorithm to recognize dorsal fins of other cetacean species, as well as expand the area that FinFindR traces for identifiable features to include the leading edge of the dorsal fin and peduncle.

Fish, Shellfish, & Corals:

Ben Richards - The Fisheries Monitoring and Research Division has made great strides in using AI to detect fish and identify species in underwater camera footage. Much of this work has been done under Strategic Initiative on Automated Image Analysis (initiated by S&T), with the goal of creating an open-source software toolkit (VIAME) allowing for automated analysis of optical data streams to provide fishery-independent abundance estimates for use in stock assessment.

Bernardo Vargas-Angel - Funding has been received to build out CoralNet's capacity to annotate point clouds from large area, photomosaic dense point cloud datasets, allowing researchers to scale up annotation from plots roughly 1 m2 to over 100 m2. A cloud-based annotation platform for point annotation of benthic photo quadrats was used by employing convolutional neural networks (CNNs) to provide human-in-the-loop machine annotation.

Matt Carnes - The Fisheries Monitoring and Research Division of PIFSC has committed to fostering the growth of machine learning in fisheries electronic monitoring. FRMD is annotating data from pelagic longline fishing vessels and training models to significantly cut down the amount of human review required for high quality fishery dependent data collection.

APPENDIX D. Crosswalk Between This Workshop and the NOAA AI Strategic Goals

The following section and related summary table are provided to facilitate identification of specific recommendations resulting from the IPW that meet the NOAA AI Strategic Goals. Additional discussion about certain recommendations are included below; discussion about each recommendation can be found in the main body of the workshop report.

Goal 1: Establish an efficient organizational structure to advance AI across NOAA

1.1 Establish a team with proficiency in AI projects at each NOAA Fisheries Science Center to advise, coordinate, and collaborate on projects.

- The team should include biologists from both the fish and protected resources research communities who have implemented successful AI projects, programmers, and other experts. Participating in the team should be part of each members' official duties.

1.2 Improve requests for proposals by improving agency staff understanding of AI principles and by improving private industry understanding of the realities and constraints of biological ecosystems

1.3 Develop an approach for providing long-term support for software developers and software engineering either internal or external to the agency

1.4 The transition to AI will require hiring or training staff to fill new roles, including data scientist, systems engineer, data annotator/classifier, and project manager

Goal 2: Advance AI research and innovation in support of the NOAA mission

2.1 Provide financial support for new AI projects to be started concurrent with traditional data collection and processing approaches

2.2 Provide for training needed for new roles and responsibilities

2.3 Support travel to internal NOAA meetings of researchers using AI and to external workshops and conferences involving AI and the broad types of wildlife and environmental data used by NOAA Fisheries

- There are topical and geographic barriers between communities of AI researchers with NOAA Fisheries. Creating opportunities for staff to meet and exchange information about their projects will improve the speed at which new projects can be envisioned, developed, and implemented. In addition, travel to external workshops and conferences involving AI will provide researchers with new professional contacts in private industry who can provide advice about projects.

2.4 Conduct R&D on equipment/sensors/survey protocols/statistics needed to collect and analyze data that involve an AI processing component

- The need for AI assistance with processing data has resulted, in part, from advances in equipment that is used to collect data. However, many data collection systems used by the marine mammal research community are still in a prototype or testing stage, and would need to be improved substantially to take full advantage of advances in AI. Further, there is a feedback loop between equipment and AI, and in some situations, transition to using AI for processing has led directly to a need for instrument upgrades to make the AI easier to implement. Resources for AI, equipment R&D, and capital purchases of new operational systems are needed.

2.5 Expand the types of UxS that can be used to collect data on wildlife.

- Many participants' interest in AI was sparked in part because our ability to collect data using UxS has greatly outpaced their ability to process data on timelines relevant to decision makers. As NOAA Fisheries researchers have explored new types of UxS, it has become clear that the approved UxS models do not meet NOAA Fisheries' needs, and the process for gaining approval for new models should become more transparent and streamlined.

2.6 Convene specific AI-related workshops for researchers in NOAA Fisheries that address key AI topics

Recommended topics include: methods discussion that includes entire NOAA Fisheries research community, how to manage "big data" used for AI, how to integrate AI with photo-id databases, how to write a successful AI contract and evaluate products, computer vision for biologists with a focus on results validation, how to pick a cloud service or GPU system.

- When possible, use existing opportunities for internal NOAA workshops to convene key groups of people involved in AI and task them with discussing how AI workflow will affect how they develop products for managers. For example, hold a joint meeting of the NOAA Fisheries Stock Assessment and Protected Resources Stock Assessment groups, add AI to the agenda, and ask that participants advise management about targeted questions (for instance, how will statisticians measure and depict uncertainty in a new workflow that includes AI in the development of an abundance estimate of a species of interest).

2.7 Develop ability to effectively annotate data

2.8 Standardize file and attribute naming conventions when it's likely that projects may have a common AI approach

2.9 Develop internal competitions within the agency for solutions to leverage the talent of our own agency to solve AI (or other technological) problems

- NOAA has extensive AI and technical expertise throughout the agency, however distinct researcher groups are often isolated, causing projects to contract out work that could be done internally. Setting up a program similar to that at NASA could help researchers leverage the skills of staff from other programs, reduce costs, build connections, and strengthen morale across the agency.

Goal 3: Accelerate the transition of AI research into operational efficiencies

3.1 Develop an interagency agreement with the NASA Center of Excellence for Collaborative Innovation so NOAA Fisheries researchers can use this resource to contract competitions (short term).

- NASA has existing agreements with multiple companies that organize competitions. Partnering will greatly improve the speed at which NOAA Fisheries can access crowdsourcing resources for AI projects.

3.2 NOAA should develop the internal capacity to quickly access a range of options for crowdsourcing competitions (long term).

- This will require flexible contracts with various vendors and technical training for involved acquisitions staff so they understand the language of competitions and AI and can work effectively with both vendors and researchers.

3.3 Dramatically increase data storage (e.g., cloud, larger networked storage systems)

3.4 Increase processing power (e.g., cloud, GPU computing systems)

3.5 Develop a statistical understanding of bias and precision in an AI context, and how uncertainty may be promulgated through the process to impact an important scientific product (e.g., a population abundance estimate) needed by NOAA Fisheries managers.

- Statisticians and modelers developing assessments of fisheries and marine mammals have developed ways to describe assessment uncertainty to decision makers so they can understand the scientific confidence in the assessment result. Statisticians and modelers already familiar with the agency frameworks for assessment will need to learn about AI so they can innovate new ways to describe uncertainty and bias in projects involving AI. Staff with this expertise will also be critical in the development of AI models, workflows, and methodologies.

Goal 4: Strengthen and expand AI partnerships

4.1 Expand NOAA researcher professional networks to include software developers and AI practitioners by supporting researchers to attend external workshops and conferences on AI.

- We expect cutting edge research on AI to occur in the private sector, but researchers can benefit by improved professional linkages with private industry. Industry involvement in our workshop led to specific examples about what the agency should change, including what programming languages are critical, under what circumstances processing should be based in the cloud vs a local GPU system, how to attract good contest participation, how to write a better request for proposals, upcoming workshops and conferences that are biologist-friendly, and why image file type matters for AI applications. The AI field is moving so quickly that industry recommendations about these issues (and many others) will change rapidly over time, and maintaining multiple professional contacts with external private industry experts will help NOAA Fisheries researchers understand advances and how they can be applied internally.

4.2 Provide ongoing support for VIAME so modules can be built to address high priority marine mammal data processing needs.

Goal 5: Promote AI proficiency in the NOAA workforce

5.1 Support travel to conferences and meetings for NOAA Fisheries researchers integrating AI into their research workflows.

- NOAA Fisheries employees pursuing AI are topically stove-piped by research community and are geographically dispersed. Developing a network of experts in the agency will be accomplished most efficiently by periodic co-location at AI-related conferences and workshops, and sharing research experiences in person.

5.2 Provide training to employees who will be involved in AI.

- Incorporation of AI into data processing will change the types of expertise needed for designing, implementing, contracting, and completing some types of projects. Supervisors, researchers, and acquisition professionals will need to understand new terminology, processes, how to set realistic milestones, and be able to manage projects that are substantially more complex.

Crosswalk Between Workshop Recommendations and Goals in the NOAA AI Strategic Plan

XX = primarily supports this NOAA Strategic Plan Goal X = supports additional NOAA Strategic Plan Goals

	NOAA AI Strategic Plan Goals				
Recommendation	**Goal 1**	**Goal 2**	**Goal 3**	**Goal 4**	**Goal 5**
1.1 Establish a team with proficiency in AI projects at each NOAA Fisheries Science Center to advise, coordinate, and collaborate on projects	**XX**	X	X	X	X
1.2 Improve requests for proposals by improving agency staff understanding of AI principles and by improving private industry understanding of the realities and constraints of biological ecosystems	**XX**	X	X	X	X
1.3 Develop an approach for providing long-term support for software developers or software engineering either internal or external to the agency	**XX**	X	X		X
1.4 Hire, partner with, and/or train staff to fill new roles, including data scientist, systems engineer, data annotator/classifier, and project manager	**XX**	X	X		X
2.1 Provide financial support for new AI projects to be started concurrent with traditional data collection and processing approaches		**XX**	X		X
2.2 Provide training needed for new roles and responsibilities that accompany the transition to using AI		**XX**	X	X	X
2.3 Support travel to internal NOAA meetings of researchers using AI and to external workshops and conferences involving AI and the broad types of wildlife and environmental data used by NOAA Fisheries		**XX**	X	X	X
2.4 Conduct R&D on equipment/sensors/survey protocols/statistics needed to collect and analyze data that involve an AI processing component		**XX**	X		
2.5 Expand the types of UAS that can be used to collect data on marine mammals		**XX**	X		
2.6 Convene specific AI-related workshops for researchers in NOAA Fisheries that address key AI topics (recommended topics include: methods discussion that includes entire NOAA Fisheries research community, how to manage "big data" used for AI, how to do AI for photo-id databases, how to write a successful AI contract and evaluate products, computer vision for biologists with a focus on results validation, how to pick a cloud service or GPU system)		**XX**	X		X
2.7 Develop ability to effectively annotate data		**XX**	X		X

Recommendation	NOAA AI Strategic Plan Goals Goal 1	Goal 2	Goal 3	Goal 4	Goal 5
2.8 Standardize file and attribute naming conventions when it's likely that projects may have a common AI approach		**XX**	X		X
2.9 Develop internal competitions within the agency for solutions to leverage the talent of our own agency to solve AI (or other technological) problems		**XX**	X		X
3.1 Develop an interagency agreement with the NASA Center of Excellence for Collaborative Innovation so NOAA Fisheries researchers can use this resource to contract competitions (ASAP)		X	**XX**	X	
3.2 NOAA should develop the internal capacity to quickly access a range of options for crowdsourcing competitions (long-term)			**XX**		
3.3 Dramatically increase data storage (e.g., cloud, larger networked storage systems)		X	**XX**		
3.4 Increase processing power (e.g., cloud, GPU computing systems)		X	**XX**		
3.5 Develop a statistical understanding of bias and precision in an AI context, and how uncertainty may be promulgated through the process to impact an important scientific product needed by NOAA Fisheries managers		X	**XX**		X
4.1 Expand NOAA researcher professional networks to include software developers and AI practitioners by supporting researchers to attend workshops and conferences on AI.		X	X	**XX**	X
4.2 Provide ongoing support for VIAME so modules can be built to address marine mammal data processing needs.		X	X	**XX**	X
5.1 Support travel to conferences and meetings for NOAA Fisheries researchers integrating AI into their research workflows.					**XX**
5.2 Provide training to employees (managers, contract specialists, biologists, etc.) who will be involved in AI (see Appendix F for recommendations from industry representatives)		X	X		**XX**

APPENDIX E. Project Presentation Summaries

Approaches to Image Processing

Citizen scientists help researchers investigate Steller sea lions - Sweeney (NOAA Fisheries)

Katie Sweeney (AFSC) summarized a project where citizen scientists helped researchers track endangered Steller sea lions in remote locations. This project aims to collect vital rate data on branded Steller sea lions with the goal of understanding why the recovery of this species differs among regions. Images are captured by remote cameras, with the potential for the collection of over a million new images per year. In addition, variation in camera angle, distance, and image resolution resulted in the need for human interpretation of imagery. Katie described how her team worked with Zooniverse, a crowdsourcing citizen-scientist platform, to process vast amounts of imagery. The two-step image processing tasks asks citizen-scientists to determine whether Steller sea lions were present within an image, and if so, whether there were any marked animals. The platform tracked the answers submitted by the crowd and aggregated their responses to arrive at a final determination for each question posed to the crowd. For example, determining whether an animal was present or not required 5-7 responses, while determining the presence of a marked animal required 14 responses. A multifaceted approach was employed to engage audiences in different ways, including the development of a software application, creation of a project blog, press releases, and other methods of engagement like the "Sea Lion of the month feature." Combined, these approaches have provided great opportunities for public engagement. Overall, the project has been successful in helping determine which images need to be reviewed by experts. Challenges with the time required for crowd-sourcing to tackle the imagery, as well as the need for pre/post processing of data still exist; these challenges may be addressed by promising new developments in using AI resources.

Computer Vision - practice and pitfalls - Redmon (UW)

Joseph Redmon provided an overview on the history and evolution of computer vision and provided insight on the differences and relationships between machine learning, deep learning/AI, and computer vision. Computer vision developed independently from machine learning and was heavily influenced by research conducted at MIT during the 90s. Since then, many tools and techniques have developed and continue to advance. Redmon introduced several key concepts and definitions to help participants understand new concepts. Imagery can be considered arrays of numeric values upon which functions (such as convolutions) can be performed to identify features like edges or areas of rapid change. These operations allow for new images/arrays to be derived from the original. Often times, deep learning can use these convolutional processes and outputs to extract features and learn from the data. Deep learning is best suited to solve problems that revolve around pattern recognition but may not be the best choice when multiple steps (i.e., reasoning) is required. Deep learning also relies heavily on training data, and is best suited to process new imagery that is similar to the training set; when presented with new data outside of the scope of the training data, it is likely to underperform. Other problems within the field of computer vision include classification, tagging, object detection, semantic segmentation, and instant segmentation. Each problem requires different processing times and has different requirements.

Machine learning and deep learning differ in effort – Use satellites to inform drone work – Tip & Cue Frameworks - Johnston (Duke)
Dave Johnston leads the Marine Robotics and Remote Sensing Lab at Duke University. Prior to his current position at Duke, he led the Cetacean program at the Pacific Islands Fisheries Science Center. His presentation covered some topics on machine learning and deep learning, but focused on the use of satellite imagery to "tip and cue" subsequent UAS survey methods. Tips are assessed via a management framework from which priority targets can be identified. Subsequent assessment cues took the form of using UAS to obtain a high resolution sample. An automated change-detection system allowed the framework to identify priorities. The use of satellite data in conjunction with UAS surveys allows for greater environmental intelligence. Johnston highlighted some work on Otter Island where pinniped surveys were informed by satellite based thermal data.

AI/ML for wildlife conservation and myth busting for biologists - Morris (Microsoft)
AI for Earth is Microsoft's initiative to stimulate innovation in machine learning for sustainability. With a dedicated staff of data scientists and computer engineers, AI for Earth is able to provide assistance and cloud-computing credits to various research projects. Morris highlighted the *Snapshot Serengeti* project, in which AI technologies were able to process millions of images captured by wildlife camera traps in the Serengeti. The project generated great hope, but also identified some key areas that were in need of improvement. While the AI was successful at classifying wildlife imagery, it was hard to apply the generated model to another project due to how it was trained and validated. The AI had difficulty classifying animals when multiple species were present. Despite the high accuracy rate of classifications, rare events or species could be missed as the model was not able to identify such occurrences as being significant. Last, some features that may be important to biologists may be difficult to "express" to an algorithm, which makes it challenging to develop an AI model and algorithm. These challenges are being investigated actively. Whenever an AI model is trained on one set of conditions, such as some population, background, or habitat, and a different set of circumstances are applied, it's likely that bias and inaccuracy will increase.

Morris offered some advice on model development and flexibility. Any time an AI is trained on one set of conditions (e.g., one population distribution, background habitat) and then is applied to another, lower accuracy and greater bias are likely. There are some important considerations in determining the ideal trade off in model specificity versus generalizability. Detector models performing presence-absence of a desired object are typically easily generalizable, while specific classifiers (e.g., distinguishing one species or individual from another) are less generalizable and will likely require training for individual projects.

Often AI may be best suited for intermediate level problems – full automation of a given problem could take more time to create a model than it will save the scientist, but there may be simpler processing tasks where AI can be most efficient and save steps for the scientist. Carefully consider the project needs, resource availability, and project end goals: the fanciest, most advanced solution isn't necessarily the best choice.

UAS applications with bottlenose dolphins, birds and bats - Thompson / Erickson (WEST)
WEST Environmental and Statistical Consultants employ AI on a variety of projects, including identification of bottlenose dolphins, detecting eagle carcasses at wind farms, fusing acoustic and thermal data to detect bats, and detecting terns with minimal disturbance. Their methods focus on a hierarchical

application for AI. The *FinFindR* project employs multiple stages and techniques to identify individual bottlenose dolphins. First, an image is processed and cropped to the region of interest (dorsal fin). With the dorsal fin identified, another algorithm traces the edge of the fin then attempts to match it with existing entries in a catalog. In a bat study, AI was used to detect bats in thermal imagery to correct counts obtained from call detectors; often, these detectors misrepresent the number of animals present. Work with video data underscored the importance of accounting for spatial-temporal factors. These factors were accounted for by use of a long short-term memory network.

Thompson and Erickson also provided a few suggestions. When training a model, they had success with cropping images to contain the specific animal or part of the animal they were training upon rather than using large images. Second, when true positives are categorized as negatives (e.g., not present), this can harm AI accuracy: being careful about what data are presented to an AI during training is key.

Using AI to identify whales from visually represented acoustic data - Allen (NOAA Fisheries)
Ann Allen (PIFSC) has been collaborating with Google to detect humpback whale songs from spectrograms derived from passive acoustic recorders. Years of recording at various sites have resulted in over 170,000 hours (10+TB) of data. After reaching out to Google, Allen was able to recruit assistance after the request was distributed through an internal Google communication channel. With the assistance of a full time programmer, AI experts and resources, an AI model built on the ResNet50 convolutional neural network was developed using the collected data. Allen noted that communication between PIFSC and Google was excellent, with plenty of back-and-forth communications to ensure that the outputs were as expected. The AI was able to identify 100% of all known whales in the dataset with a 3% false-positive rate. Allen and the PIFSC will continue to work with the algorithm to adapt it for different species (e.g., fin whales, blue whales) and different types of equipment.

AI for Enumeration

Applications of Machine Learning Algorithms to Automate Data Extraction from Images - Altukhov (NPWC)
A long-term monitoring program has been capturing data on Steller sea lions (herein abbreviated to SSL) in the Western Aleutian (US), Commander and Kuril islands (Russia) with the objective of collecting marked/branded SSL for life history, demography, and distribution studies. Beginning in 2010, these efforts were augmented by the deployment of 55 remote cameras at various rookery and haulout sites in Russia and Alaska. Collectively, these cameras collected over 15 million images over an eight year period. This volume of imagery presents a serious data-processing challenge with some sites requiring hundreds of hours of manual review; often, it would take an entire year to analyze a single site.

To improve the timeliness of deriving data from this imagery, Altukhov utilized Keras, an open source neural-network library, to develop an AI model that was able to find brands on animals. With the use of UNet CNN, it was possible to locate animals within the image; the VGG16 CNN was then used to perform classification into branded and unbranded animals. Compared with manual processing, the model performed well overall, detecting 25% more brand locations than human observers while missing only 5% of brand locations. Animals with marks, such as scars and spots, did confuse the model: 16% of brands were incorrectly classified. This new approach only required 5 days to process 200,000 images; a significant time savings compared to manual methods.

Utilizing an image catalog of branded animals and specifying brand locations, additional refinement of the model was possible. Further development of the AI model revolved around prediction of the location of a branded animal within an image, then attempting to identify the brand itself. While this approach had trouble identifying 15% of brands (classified as "Unknown"), it allowed for the processing of brand ID imagery within a month.

AI was also utilized for the enumeration of multiple species (Steller sea lions, Northern fur seals) in imagery captured from a UAS platform. After stitching the imagery together with Agisoft Photoscan software, a UNet CNN based AI was developed to classify Steller sea lions, Northern fur seals and their harems. The accuracy of this model varied between 70 and 99%, with weather and animal density being the primary driver of this variance.

Automated surveys for ice-associated seals in the Arctic - Moreland (NOAA Fisheries)
Ice-associated seals are broadly distributed and strongly dependent on sea ice as habitat for pup rearing, molting and resting. They are important resources to indigenous coastal communities and their population abundance and trends are poorly understood.

Past surveys were conducted in the Bering Sea (2012 and 2013), as well as the Chukchi Sea (2016) with the use of three pairs of infrared and color cameras mounted in a fixed-wing manned aircraft. These survey efforts collected approximately 41TB of imagery (~4 million images), requiring approximately a year to process the Bering Sea imagery and 6 months to process the Chukchi Sea imagery.

Two thermal detection approaches were implemented: a manual review method where an analyst looked for thermal signal "peaks," and semi-automated software running a modified outlier based algorithm on thermal imagery. The manual review method analyzed thermal data for relative thermal signal peaks to identify frames containing animals; this worked for most, but not all species. Often, the thermal signal of ringed seals would be lost in the noise. The modified outlier algorithm provided a semi-automated approach but still required a visual review of hotspots as many detections were false positives. This approach was effective for detecting ringed and bearded seals.

Current objectives for automated, instrument-based surveys include development of open-source general detection algorithms for animals on sea ice, elimination of anomalous thermal signals, and an in-flight software system to run algorithms upon. Ultimately, these algorithms would be used to control data acquisition for manned and unmanned surveys. Seal surveys will be conducted in the spring of 2020 in the Beaufort Sea with the new in-flight camera software, with the goal of testing detectors.

Kitware's VIAME software is being used to test detection and classification models developed by Microsoft AI for Earth and Xnor.ai. Some models look at one channel only (i.e., color OR thermal); future work involves refinement of models that utilize both sources of imagery to identify and classify animals. Past survey efforts have highlighted the need to capture imagery simultaneously: investments in a higher end inertial measurement unit, better handling of timestamps and triggering allow for pixel alignment of multi-spectral imagery.

Collection and analysis of imagery and video by the SWFSC Antarctic Ecosystem Research Division - Hinke (NOAA Fisheries)
The Antarctic Ecosystems Research Division (AERD) is conducts long-term monitoring of Antarctic fisheries, and studying the Antarctic ecosystem to provide management advice. The AERD employs multiple cameras systems for data collection. Time lapse cameras are used to estimate the timing of egg-laying and hatching events as well as chick production in penguin colonies. The APH-28 UAS platform is used to capture imagery for census studies of seals and penguins (Goebel et al. 2015), health assessments of seals using photogrammetry (Krause et al. 2017), and photo ID. Video loggers attached to seals, chinstrap and gentoo penguins, as well as Antarctic fur and leopard seals, provide information on krill density, prey encounter rates, and other key behavioral characteristics.

The collection of this data presents a significant processing challenge. Presently, it takes 2-4 analysts approximately 170 hours to analyze imagery from 34 time lapse cameras. Analysis of 30 minutes of footage captured by chinstrap penguins requires more than an hour. AERD hopes to use VIAME to automate counts of pinnipeds and to utilize thermal imagery to determine age class. There is also the goal of developing AI models to remove frames from videos that are not of interest (e.g., open water, surface) and identifying frames where the animal is exhibiting certain behavior (predator-prey interaction, interactions with other conspecifics), and the presence of krill swarms. A new model built by a graduate student at the SWFSC was able to achieve accurate results from video with the use of VIAME's iterative query refinement (IQR) tool. This process only took a few days to complete and provided high success rates (> 97%) for identifying frames with select, simple characteristics.

Use of Images for Photo-Identification & Photogrammetry

Computer vision for conservation: Automating right whale photo ID - Khan (NOAA Fisheries)
Entanglements and vessel strikes are significant threats for North Atlantic right whales. To monitor abundance, aerial surveys are conducted to capture images of individuals. These images are then compared with the New England Aquarium photo identification catalogue to identify individuals. Christin Khan (NEFSC) coordinated with Kaggle to host a competition to produce an algorithm that could identify individual animals. Using 7,000 images, a total of 470 competitors in 364 teams competed. The winning algorithm was developed by Deepsense and utilized several algorithms chained in a pipeline to locate the animal's head, crop and rotate the image into a standard "passport photo" of the whale, and then matched this image to a specific individual. This approach yielded an accuracy rate of 87%. Next steps include development and integration into Wild Me's Flukebook platform, and retraining of the algorithms on Southern right whales, as well as updated images provided by the North Atlantic Right Whale Consortium. In addition, there is interest in integrating oblique imagery captured from vessels. Despite being a "side project" and not directly funded, persistence paid off.

Flukebook: Multi-modal, multi-stage machine learning for marine mammal research with citizen science - Holmberg / Parham (WildMe)
Jason Holmberg (WildMe) provided an overview of Flukebook: an online platform that allows scientists, citizen-scientists, industry partners, academia, government agencies and others to collaborate. Wildbook, the platform Flukebook is based upon, seeks to help researchers focus on analysis of data instead of curation, engaging the public and connecting data sources, and performing mark-recapture type analyses

with non-invasive AI techniques. Despite being a highly collaborative platform, security and data access-control are built in as well, with the ability to restrict access to data.

Artificial intelligence is built into multiple portions of the Flukebook/Wildbook platform. Data from multiple sources (e.g., public social media posts, citizen scientists, biologists) are fed into the data management and image analysis server. Fully automated pipelines are able to run multiple multistage algorithms and display the output to the user, providing a visualization of what the AI sees and serves as a tool for engaging the public. Flukebook agents (software tools) analyze public social media content and use AI technology to determine what metadata are missing from a sighting report and request this from the user. In addition, Flukebook has tools and services that give the public an easy way to interact with it, such as “Tweet a Whale,” which strives to identify a specific whale within the provided image and returns information about the individual to the requester. Tools to communicate how AI helps bridge the gap between computer vision scientists and the public.

Different species have different detection models and algorithms: the Whaleshark Wildbook is based on NASA’s algorithm for identifying star patterns in the sky while other Wildbooks (e.g., Zebra Book, Giraffe Spotter) utilize background subtraction algorithms as part of a sequence to determine species and individual.

In discussions, Holmberg identified several challenges with the WildBook platform and AI in general. These challenges included the need to retire algorithms over time as new ones are developed, computing issues when multiple collaborators are attempting to use Wildbook AI models concurrently, and the need to make software engineering improvements for performance. In addition, Holmberg also mentioned that models produced by contests or in academia may not be able to deliver accurate results using data collected from new projects or data collected under different circumstances than what the model was developed on. In addition, these models may encounter issues related to performance and scalability when integrated into a research project.

Species identification and stereo measurements - Lauffenburger (NOAA Fisheries)
Nathan Lauffenburger presented an overview of the CamTrawl system developed by the University of Washington and NOAA. A stereo camera system mounted at the end of the trawl captures 2 megapixel images at 4 hertz as fish enter the codend. These images are downloaded processed at sea: C++ and Matlab algorithms are used to subtract the background from the imagery, collect length measurements, and identify the species. This method allows for collection of data on fish who are not retained in the codend, but remain acoustically relevant. A 20-minute trawl takes about an hour of processing to complete at sea. Results can be manually validated in about 15 minutes.

The species identification algorithm is jointly developed with the University of Washington. Thumbnails of fish are created from the imagery and cropped/oriented to a standard format in preparation for feature identification using the SIFT algorithm and codebook learning. Using 200-500 images per species, the algorithm achieved an accuracy rate of 98% for five species/classes of fish and 100% accuracy for Pollock. Lauffenberger also noted that CamTrawl faced several challenges, including turbidity obscuring animals, high density occlusion of fish (overlapping animals), measurement of curved fish, and fish in orientations that are suboptimal for measurement and classification.

In the future, it may be possible to process catches solely with the CamTrawl and eliminate the need for extracting fish from the environment. Currently, catch data are being used to ground truth the output from the CamTrawl system. It may also be possible to compare echogram and CamTrawl results.

Morphometrics and volumetrics of pinnipeds from imagery - Shero (WHOI)

In January of 2019, the Woods Hole Oceanographic Institution (WHOI), Duke University, and Canadian Department of Fisheries & Oceans (DFO) conducted pinniped research on changes in animal mass over an 18 day lactation period in Nova Scotia, CA. Using a Freefly Alta 6 hexacopter and Sony Alpha a5100 (24 megapixel, 30mm lens), oblique imagery of Gray seals was captured while orbiting the target at 20-40 m. Using python code to geotag the images, 3D models, orth-mosaic, and digital surface models were developed using Pix4D's structure-from-motion (SFM) processing. Compared to the true mass of 17 Gray seals handled, the body mass estimate from SFM had an error of 11kg (5.8%) for adult females and 4 kg (11%) for pups. The higher percent error for pups is related to their size.

Shero found that an orbit of 360 degrees around the target was sufficient for this analysis, and that capturing imagery at two altitudes would work as well assuming there was vertical overlap between the two orbits. A point in the scene was captured in five or more photos before being added to the 3D point cloud to reduce error due to animal movement, while tall objects near the target made modeling challenging.

This technique was also used on Weddell seals in the Antarctic, and in collaboration with New Zealand in the Ross Sea. Several problems were encountered. Compass issues caused by close proximity to the poles were mitigated by turning the platform compass off. Cold weather limited drone operations, and the homogenous ice substrate resulted in artifacts appearing in the 3D models. To address the 3D artifacts, a spiral flight path around sedated animals was flown to collect both oblique and nadir images and the generated model had fewer artifacts.

This approach is a completely non-invasive means for collection of length and mass estimates and is applicable for many species, habitats, and group sizes.

Statistical Considerations

Accounting for species misclassification - McClintock (NOAA Fisheries)

Even experts don't always agree with a species ID from non-invasive, image-based, aerial surveys. If not properly accounted for, severe bias can occur in estimators of species distribution and abundance. One approach to quantify and incorporate this uncertainty is exemplified using the aerial imagery of ice-associated seals. The approach involved multiple observers reviewing the same image set, determining species, and recording their confidence in the species determination (100% = "positive", < 99% = "likely", or "guess"). A misclassification model builds a categorical distribution based on observed species and confidence, rooted in "positive" identifications as truth, providing the probability of a given species being misclassified by the observer. It's also possible to assign a probability to unknowns based on this analysis. Experts don't always agree and they're not always right; it's not a problem as long as it's measured and incorporated into the analysis. Species classification doesn't have to be 100% correct.

Estimating abundance with automated detection systems - Conn (NOAA Fisheries)
Using AI for detection in abundance estimation is structurally similar to traditional models for human observations in that they all include detection and classification processes. Using ice seal surveys as an example, we'll start with a conceptual model for abundance and detection with aerial survey data, with the goal of describing how abundance varies over space. We need to measure availability (e.g., proportion of seals out of water), detection probability of the automated detector (in human observers, this would be detection probability of the human observer), disturbance, and species misclassification. An important question to consider is: how do improvements in detection probability using AI (versus human observers, or different types of AI) impact the overall abundance estimate? Ultimately, management is based on population estimates and relative measures of uncertainty, so the relevant question is how adopting AI procedures will affect bias and precision of resultant estimators. In this light, it is important to weigh the amount of work required to improve detection to the overall variance estimate of the final abundance estimate. This consideration is also helpful to determine what AI detection rate is "good enough."

If detection is entirely automated (i.e., no manual removal of false positives) then it's necessary to account for the spatial distribution of those false positives. Out-of-sample data are crucial for estimating detection and species classification probabilities. In other words, do not use all the data in the course of development of a detection/species classification model. A large set of imagery for training and testing set is needed but a portion of this imagery must also be withheld for the measurement of detection probability. This information is necessary for incorporation into the abundance estimation model.

In discussion, there was significant interest in better understanding how AI error is, or should be incorporated into abundance estimation models. A question on whether AI development also modeled and considered error arose. With regard to required model accuracy, Conn recommended a focus on precision and bias to determine where to focus efforts. Prioritizing resources and effort should be done with the end goal in mind.

Notes on statistical considerations for photoID mark-recapture - Conn (NOAA Fisheries)
The primary statistical consideration with photo-id analyses is with matching error and the effect on parameter estimates (e.g., survival, abundance) from capture-recapture of encounter histories. Matching errors tend to cause "ghost" encounter histories, which are false histories where the animal is only observed once. Analysis with such records tend to lead to overestimates of abundance and underestimates of survival. To minimize matching error, it is recommended to limit datasets to only high-quality imagery and to limit inference to individuals with highly distinctive marks. A number of researchers are working to deal with images captured on left and right flanks of animals. Most approaches assume misidentification is constant across individuals; in reality, it's more likely a function of mark distinctiveness. Conn also suggested retaining of out-of-sample data to measure error rates as a function of distinctiveness and to conduct exercises to examine precision-bias tradeoffs using different distinctiveness categories and approaches.

Hardware, Processing, & Storage

Network effects: storage, processing, and connections - Hou (NOAA Fisheries)
Traditional image based projects include off the shelf cameras, copying imagery from memory cards, collecting a GPS track, and using a companion application for data entry. Some projects require many hard drives due to the high volume of data collected. Often, these projects encounter challenges with

providing multi-user access to imagery, tracking changes and versions, and higher costs related to capturing, storing, and required providing public access.

New tools being incorporated into large-scale surveys include a shift to machine vision cameras, allowing for cleaner management and real-time processing to integrate intelligence into the image capture and storage process, reducing some of the traditional challenges. All cameras are controlled by on-board computers and data are stored on a network attached storage (NAS) setup with redundancy (RAID). This approach reduces migration friction and provides high redundancy with little effort. Networked cameras, computers, and storage allows for real-time processing and storage of images. Colocating imagery on one large volume helps with automation and provides multiple users access to work on a single dataset. Efforts are underway to develop detection and classification models to incorporate into on-board processing.

Future directions include intelligent systems allowing targeted image collection and moving to incorporate unoccupied aircraft into survey efforts. This work helps reduce the need for storing all images and increase endurance of the platform

In discussion, the topic of staffing arose. The AFSC has a couple of specialized IT staff who work closely with biologists to provide critical assistance with advanced technologies. These staff are particularly ambitious, interested in the work, and expanded their skills to meet research needs and were not hired directly for this role. Other centers do not have the same in-house support but would like IT Specialists with passion for the projects, unique technical skills to provide specialized support, and communication skills to bridge the gap between biologists and technology.

Moving to the cloud - O'Neil (NOAA Fisheries)
NOAA's Big Data Project provides public access to NOAA's large datasets. Agreements have been established with organizations better equipped to provide open data, such as Amazon, Microsoft, Google, IBM, and the Open Cloud Consortium (non-profit). This project includes five separate 4-yr agreements to understand NOAA data and offset cost for assessing needs. An example is how NOAA worked with partners to develop services to make satellite data available to customers. Industry partners provided access but NOAA experts were able to ensure quality of the data.

Publishing data on the cloud improves access, facilitates the use of data, improves security posture, allows development of authentic tools, and enables new economic and research opportunities. It also reduces requests for data and the internal data management load—requests for data went down by 80% and cloud platforms were able to deliver 1.2 petabytes of data over a 4-month period, 30-100 times more than what NOAA was providing previously. The Big Data Project brought an awareness of the true demand for NOAA data. Developing agreements with industry to host data on the cloud has allowed increased public access and usage of NOAA data with reduced costs.

In discussion, Participants expressed that it was often unclear who was allowed to use cloud agreements, the process to obtaining access, or how to comply with regulations. While cloud access is free under the CRADA, no guidance has been provided to NMFS researchers.

Matching UAS platforms to your data objectives - Seymour (USGS / Cher. Nation)
Alex Seymour (USGS/ Cherokee Nation) provided insight on a project Duke University's remote sensing laboratory undertook. Using a fixed-wing eBee UAS platform, the project team was able to capture imagery for multiple purposes, including colony mapping. Flying overlapping nadir transects, the platform "scanned" the area for seals and other animals. This data was usable for enumeration and classification of animals, and could also be processed to produce georeferences mosaics, index maps, and for volumetrics.

Seymour was able to provide advice on sensors, platforms, and techniques:

- Use of low-distortion fixed focal length lenses and global shutter helps reduce distortion at the edge of an image.
- When planning to map a colony or collect imagery for enumeration, images should be collected with a minimum of 60-70% overlap with nadir transect lines. These transects should be planned to run on the shorter axis of the survey area to reduce movement of animals in overlapping imagery.

- The use of structure-from-motion software (e.g., Pix4D, Agisoft, Drone2Map) are helpful for mosaic generation.
- Dual-sensors (thermal, color) may be helpful for specific projects but it is necessary to ensure that the imagery is aligned. Some platforms and manufacturers provide imaging systems that perform this alignment automatically.
- Morphometric and volumetric surveys require high resolution imagery to reduce error in measurements. Rotorwing aircraft are better suited for these missions; the ability to reduce blur from wind and platform motion helps reduce error in measurements.
    - Most missions can be accomplished with off-the-shelf systems; single-image nadir morphometric projects are an important exception.
- Camera and lens aperture, shutter speed, and ISO should be optimized for the highest image quality. Imagery capturing imagery closer to the target is desirable (i.e., captured at lower-elevation). If disturbance is a factor, a telephoto lens may be helpful.
- The use of a ground-control-station with live video feed is helpful for target identification.
- Gimbal mounts are helpful for capturing both volumetrics and morphometrics. For morphometrics, the gimbal should be set to ensure that nadir imagery is collected.

- Fixed-wing UAS with vertical take-off and landing capabilities may struggle with landing in windy and gusty conditions.
- The reliability and accuracy of compass and GPS may suffer in polar regions. In addition, some platforms may not allow for control of GPS settings.
- Avoid UAS systems that are at the start of their product lifecycle: more mature systems have had more time to address problems that arise.

Appendix Table E-1-- Summary of recommended UAS platforms for each marine mammal mission type.

Mission Type	Survey Characteristics	UAS Platform Characteristics
Colony Mapping & Counting	● Overlapping nadir images ● Transect lines perpendicular to the long axis (reduce animal movement in overlapping images) ● Images may be scaled to a measurement system ● Georeferenced mosaic or index map is a primary data output	● Long endurance at cruising speed make fixed wings optimal to map large areas ● Duel sensors to capture thermal and visual imagery ● Global rather than a rolling shutter and prime lens to reduce distortion **Software / compatibility considerations** ● Need SFM software with mosaic editing ability(e.g., Pix4d, AgiSoft) **Example platform:** SenseFly eBee
Morphometrics (2D, single image)	● Single nadir images ● Seek, acquire and follow animal ● Capture individual specific details ● Some uses include health indices and energy transfer ● Manually flown missions	● Multi-copter required, larger craft for better stability in winds ● Ultra low distortion lens ● Digital RGB sensor ● 3-axis gimbal sensor ● Precise altitude geotagging (laser altimeter) ● Robust telemetry for signal ● Live feed with large screen **Example platform:** Freefly ALTA 6
Volumetrics (3D, multi-image)	● 2D images are not suitable, must construct 3D mosaic ● Low altitude orbits around target ● Conduct orbits quickly to reduce animal movement ● Images capture at oblique angles ● Flown closer to animals for higher resolution	● Multi-copter required, larger craft for better stability in winds ● Ultra low distortion lens ● Consider smaller / quieter copters to reduce disturbance or telephoto lens for greater distance buffer ● 2 or 3 axis gimbal sensor ● Second controller for gimbal/payload control ● Live feed with large screen **Software / compatibility considerations** ● Orbit mode in flight management software **Example platform:** Freefly ALTA

APPENDIX F. Online Courses & Resources for AI

Program	Description
Online Courses	
Cognilytica	AI & ML Project Management Training & Certification
Coursera	Online courses in various topics, including AI
DataCamp	Learn R, Python, & data science online
Kaggle	Free data science online micro-courses
MIT Courses	Machine learning online course
University of Georgia	Practical machine learning and data science for executives online course
Online Resources	
AI TED talks	Forbes article citing seven AI & ML TED talks
AI Today	A Podcast focusing on relevant information about what's going on today in the world of AI
Machine Learning for Everyone	A blog about ML for non-AI experts

APPENDIX G. List of Suggested Literature

The following is a list of helpful publications as a reference for AI and detection, photo-identification, enumeration, photogrammetry and health (e.g., body condition), satellite imagery, and data collection.

Title	Author	Year	Journal
AI & Detection			
Towards automated annotation of benthic survey images: variability of human experts and operational modes of automation	Beijbom et al.	2015	PLoS ONE
Robust methods for the analysis of images and videos for fisheries stock assessments	National Research Council	2014	Proceedings of the NRC Workshop
Computer-automated bird detection and counts in high-resolution aerial images: a review	Chabot and Francis	2016	Journal of Field Ornithology
Deep learning for small object detection in satellite infrared imagery	Crall et al.	2018	Proceedings of the MSS National Symposium on Sensor and Data Fusion
Scallop detection in multiple maritime environments	Dawkins and Stewart	2011	Rensselaer Polytechnic Institute
Quantifying variation in killer whale (*Orcinus orca*) morphology using elliptical Fourier analysis	Emmons et al.	2018	Marine Mammal Science
Two-stage detection of north Atlantic right whale upcalls using local binary patterns and machine learning algorithms	Esfahanian et al.	2017	Applied Acoustics
Working group on machine learning in marine science (WGMLEARN)	ICES	2019	ICES Scientific Reports
Deep learning	LeCun, et al.	2015	Nature: Review
Multi-view object-based classification of wetland land covers using unmanned aircraft system images	Liu and Abd-Elrahman	2018	Remote Sensing of Environment
Robust methods for the analysis of images and videos for fisheries stock assessment: summary of a workshop	National Research Council	2014	National Academies Press

Title	Author	Year	Journal
A Trainable system for object detection	Papageorgiou and Poggio	2000	International Journal of Computer Vision
Right whale recognition using convolutional neural networks	Polzounov et al.	2016	Arxiv
Machine learning to classify animal species in camera trap images: Applications in ecology	Tabak et al.	2018	Methods in Ecology and Evolution
Report of the National Marine Fisheries Service Automated Image Processing Workshop	Williams et al.	2012	NOAA Tech Memo NMFS-F/SPO-121
Photo-Identification			
An astronomical pattern-matching algorithm for computer-aided identification of whale sharks Rhincodon typus	Arzoumanian et al.	2005	Journal of Applied Ecology
Applying deep learning to right whale photo identification	Bogucki et al.	2018	Conservation Biology
HotSpotter - patterned species instance recognition	Crall, et al.	2013	Proceedings of the IEEE Workshop on Applications of Computer Vision
Recognition of juvenile hawksbills *Eretmochelys imbricata* through face scale digitization and automated searching	Dunbar et al.	2014	Endangered Species Research
A note on an automated system for matching the callosity patterns on aerial photographs of southern right whales1	Hiby and Lovell	2001	Journal of Cetacean Resource Management
Photo-identification software for bowhead whale images	Hillman et al.	2008	Final Report of OCS Study MMS2008-002. SC/60/BRG24, Appendix II
Drones and convolutional neural networks facilitate automated and accurate cetacean species identification and photogrammetry	Moskvyak et al.	2019	Cornell University

Title	Author	Year	Journal
Robust re-identification of manta rays from natural markings by learning pose invariant embeddings	Moskvyak et al.	2019	Cornell University
EM and component-wise boosting for Hidden Markov models: a machine-learning approach to capture-recapture	Rankin	2016	bioRxiv
Manta Matcher: automated photographic identification of manta rays using keypoint features	Town et al.	2013	Ecology and Evolution
Enumeration			
Noninvasive unmanned aerial vehicle provides estimates of the energetic cost of reproduction in humpback whales	Christiansen et al.	2016	Ecosphere
Automated detection and enumeration of marine wildlife using unmanned aircraft systems (UAS) and thermal imagery	Seymour et al.	2017	Nature: Scientific Reports
Photogrammetry & Health (e.g., Body Condition)			
A rapid UAV method for assessing body condition in fur seals	Allan et al.	2019	Drones
Estimating body mass of free-living whales using aerial photogrammetry and 3D volumetrics	Christiansen et al.	2019	Methods in Ecology and Evolution
Inexpensive aerial photogrammetry for studies of whales and large marine animals	Dawson et al.	2017	Frontiers in Marine Science
Drones and convolutional neural networks facilitate automated and accurate cetacean species identification and photogrammetry	Gray et al.	2019	Methods in Ecology and Evolution
An accurate and adaptable photogrammetric approach for estimating the mass and body condition of pinnipeds using an unmanned aerial system.	Krause et al.	2017	PLoS ONE
Satellite Imagery			
Use of high resolution space imagery to monitor the abundance, distribution, and migration patterns of marine mammal populations	Abileah	2001	IEEE

Title	Author	Year	Journal
Aerial-trained deep learning networks for surveying cetaceans from satellite imagery	Borowicz et al.	2019	PLoS ONE
Whales from space: Four mysticete species described using new VHR satellite imagery	Cubaynes et al.	2018	Marine Mammal Science
Penguins from space: faecal stains reveal the location of emperor penguin colonies	Fretwell et al.	2009	Global Ecology and Biogeography
Whales from space: Counting southern right whales by satellite	Fretwell et al.	2014	PLoS ONE
Convolutional neural networks for detecting great whales from orbit in multispectral satellite imagery.	Gray and Johnston	2018	AI for Wildlife Conservation Workshop
Automatic whale counting in satellite images with deep learning	Guirado et al.	2018	bioRxiv
Novel Data Collection Methods			
Noise levels of multi-rotor unmanned aerial vehicles with implications for potential underwater impacts on marine mammals	Christiansen et al.	2016	Frontiers in Marine Science
Unoccupied aircraft systems in marine science and conservation	Johnston	2019	Annual Review of Marine Science
Studying cetacean behaviour: new technological approaches and conservation applications	Nowacek et al.	2016	Animal Behavior
A small unmanned aerial system for estimating abundance and size of Antarctic predators.	Goebel et al.	2015	Polar Biology

APPENDIX H. AI MILESTONES

This section describes major milestones and tasks associated with a machine learning project and may not encapsulate all steps or phases. While every project is unique, this guide serves as a basic outline. Table 2 lists online coursework and other resources.

Milestone	Common Tasks
1 - Scoping	
Problem definition	● What problem am I trying to solve? ● What is this problem reframed in computer vision? ● Is this problem best solved with AI? Consult literature and AI experts to understand whether AI is the right tool for the task, and how likely machine learning approaches are to the desired level of accuracy. In general, if a human can perform some image analysis, it's likely that an AI model could do the same. Bear in mind, some tasks that seem "easy" to a human may be quite challenging for a computer. Understanding the problem at hand from a computer-vision perspective will be very helpful. For example, automatically counting marked animals in an image can be broken down into three different computer vision problems: segmentation (separation of animals from one another), classification (determination of whether the object is an animal or substrate), character recognition (for a given animal, determine whether there is a man-made marking on the animal). Be mindful of objectives and don't let the computer scientists choose the hardest method possible: not all problems require the use of AI. Some other computer vision techniques may be equally effective and more cost efficient.
Budgeting and resourcing	● Who will annotate the data? ● Where will the data be stored and accessed? ● What computing resources are needed? How do I get access to it? ● Is there sufficient funding to get to a viable model? Development of a viable model can be an expensive process; if funding is insufficient, it may be necessary to re-evaluate project objectives. Data annotation and model training are highly resource intensive tasks. Annotation of a large dataset may require significant human capital to complete in a meaningful time frame. Without access to GPU computation hardware, development of a model could take an unfeasible amount of time. In addition, effective storage and distribution of large datasets can be a significant challenge, especially when collaborating with external partners. The use of network-storage systems is highly recommended.
Statistical considerations	What level of accuracy is sufficient for study needs? Absolute accuracy is likely unattainable, and may even be unnecessary as some degree of error can be

Milestone	Common Tasks
	incorporated into a statistical model. Consulting a statistician will help with determining when a sufficiently accurate model has been developed--this will help reduce resource expenditure.
Identification of imagery that is appropriate for training a model	The more similar training data are to "new imagery," the better the results will be. Use of historic imagery may be tempting, especially for rarely sighted species, but the inclusion of this data may have adverse effects on a model's predictive power. Some models can be trained on just a few hundred images while others may require thousands. Data augmentation (e.g., creation of synthetic data) may be necessary to develop a robust training set.
2 - Data preparation and annotation	
Data preparation	Once a training dataset has been identified and is appropriate for use in training an AI model, some data preparation, standardization, and documentation may be required. This may include aggregating imagery in one location (e.g., one server share), setting up server, software, and processes; or even finding a way to transfer the imagery to an external partner. If applicable, be sure to understand who has access to the data and that all parties have agreed to contribute their data to the project. Be sure to reach out experts and project-team members as needed. If there are plans to collaborate with external partners or to host a contest, carefully consider the source and availability of the contest data. If the dataset slated for use in a contest dataset is publicly and freely available, there is the possibility of contestants artificially enhancing the performance of their entry. Contest administrators can provide guidance and advice.
Annotations	To train the AI model, it is often necessary to provide the "correct answer" in the form of an annotation. For imagery, annotations are commonly created by setting bounding boxes around the object of interest. The computer is then able to mathematically describe what is within the box and can inform the AI model. Depending on the AI model and methodology, different annotation methods may be necessary, or may deliver better results. Though annotating data are extremely tedious and time consuming, high quality annotations are critical for the development of an accurate model: garbage in, garbage out.
Quality control	Regardless of whether the annotation is being performed in-house, by contractors, or obtained via crowdsourcing, it is necessary to perform spot-checking of annotations. This will help ensure that the annotations are usable.
Partitioning the training set	To avoid overfitting the AI model, the dataset needs to be split into a training and testing dataset. The training dataset is used directly to train the model, and the testing set is used to validate the accuracy of the model's predictions.

Milestone	Common Tasks
	Many common AI/machine learning libraries (e.g., Tensorflow, sci-kit learn) have this functionality built in. If the framework used to develop a model does not have this capability, this task may need to be manually performed.
3 - Model selection, training, and testing	
Selection	There are many models and networks to choose from, such as VGG, ResNet, DenseNet, and YOLO, to name a few. It is recommended to consult an expert to see which approaches may work best for your project. Each model has specificities that may better suit them for a problem.
Training	With a few potential approaches selected, it is time to pass the training imagery and annotation data to the model. During this process, the computer uses a series of mathematical operations on the specific portion of the imagery identified in the annotation. The model training process is computationally intensive and will require access to GPU hardware in order to be completed in a practical amount of time.
Testing	After a model has been developed, the testing set can be processed to derive metrics on model performance. The accuracy of predictions is the primary metric to consider.
4 - Model evaluation and retraining	
Evaluation	Having developed and tested multiple models, it is possible to evaluate and decide on a viable approach(es). In some cases, an ensemble-approach of combining multiple model outputs may deliver the required accuracy. It is important to keep in mind that the results of an AI model are likely to be fed into a statistical model that can account for errors. Consulting a statistician to determine what level of accuracy is sufficient is recommended.
Retraining	In some cases, it may be necessary to make some changes to training data annotation methods, model choices, and parameters to improve the accuracy of a model.
5 -Deployment and integration	
Integration into a workflow	The final major stage is to integrate the model into the project workflow. Considerations may include developing a pipeline for collecting and organizing images, validating model outputs, and statistical considerations.
Determining triggers for modification	Changes to the environment, sensors, and methodologies may necessitate retrain the model with new training data or even adoption of a different technique overall. For example, increasing the elevation at which imagery is collected may impact the relative size of an animal within an image, resulting in lower positive detection rates.

APPENDIX I. Post-Workshop Feedback Summary

The IPW Steering Committee sent out a request for feedback survey after the workshop and received 17 responses. Respondents answered various questions about workshop logistics and the impacts the workshop had on their current perceptions and future direction, in regards to AI. All responses were from biologists or natural resource managers (these results combined), or industry. Industry responses were reported separately with “Industry, #,” with # indicating the number replied, when applicable. Written responses are summarized.

Feedback Responses		**17**
	Biologist	13
	Industry	3
	Manager	1

New to AI:		**11**
	Biologist	10
	Manager	1

Experience with AI:		**6**
	Biologist	3
	Industry	3

How would you rate workshop logistics and organization? Scale of 1 (poor) to 5 (excellent)

Level 4	5	*(Industry, 1)*
Level 5	12	*(Industry, 2)*

What do you think went well? Why?
Conference logistics (i.e., coffee, snacks, lunch, agenda, schedule, hotel shuttle to campus, evening social, and moderating) were well thought out, organized and executed. Communication about the agenda and flexibility in changes in the agenda to reflect organic flow of workshop was beneficial. The workshop provided a good overview of what is currently being done and future efforts to pursue. The diverse group of attendees (i.e., government scientists and managers and industry), especially those from industry, was beneficial to have in attendance.

What do you think could be improved? How?
Logistically, there could have been more people with foreign national escort training to ease this burden of non-US citizens to gain access. Communicate changes in the agenda more clearly throughout the workshop. While we disseminated information on the computer programs and instructions required for the trainings, this could have been shared earlier to allow more time for people to troubleshoot issues ahead of time. The length of the conference was a bit long to have traveling required on both weekends before/after the workshop week. There still seems to be a need for increasing collaborations with NOAA Fisheries. Have a nominal workshop fee to cover costs for food and beverages so the Steering Committee didn’t have to cover costs not covered by donations.

How would you rate workshop content and focus? ***Scale level from 1 (poor) to 5 (excellent)***

Level 4	8	*(Industry, 1)*
Level 5	9	*(Industry, 2)*

What do you think went well? Why?

The workshop was well organized with the first day providing a summary and background of a variety of projects covering a wide-range of data products, statistical considerations, and methods. The second day offered a good background of basics about AI methods (especially the steps to consider getting started and what you'll need to consider moving forward) and their applications to analyze data followed by break-out sessions for a free-flow of discussion for various topics of interest. Additionally, the various demonstrations were helpful for conceptualizing AI practices. The presence of industry and NOAA scientists implementing AI provided a good mix of expertise present (including guiding questions to help facilitate conversation). These experts provided good insights for how to approach using AI methods. Great networking opportunities with industry leaders and other NOAA scientists implementing AI. The flexibility of the workshop to adapt to specific needs and topics that came up as important for attendees to concentrate on was beneficial.

What do you think could be improved? How?

The VIAME overview/demonstration was beneficial; however, it could have been better tailored to the audience, which was a majority not computer scientists. Using too much jargon can make VIAME seem intimidating or more challenging and there is no reason NOAA scientists cannot implement and utilize VIAME themselves. However, some also mentioned that the VIAME overview was too many steps and didn't need to be click-by-click demonstration of steps. Allow time for prepping computers for training (and provide two IT staff for troubleshooting support. Break-out sessions were so helpful it would have been better to have more of those, as well as more industry experts in attendance. Shifting the agenda and not communicating changes well was problematic for scheduling and those who were 'walk-ins.' That statistical considerations was helpful and it would have been helpful to have more talks in this regard, especially summarizing results of AI.

What are one or two themes you would like to see highlighted at future workshops?

- Specific asks of the AI community: what specifically can interested AI experts do to help biologists incorporate AI?
- How to manage big data.
- Building successful partnerships (outside of NOAA) for AI / computer science.
- How to write a successful vendor contract in detail, evaluate your products, and how to secure funding.
- Incentivizing collaborations between fish and marine mammal groups.
- Cetacean dorsal fin ID using scars and coloration patterns.
- Resource sharing (NOAA specific and otherwise) and lessons learned.

What training sessions would you like to see at future workshops?

- Workshop on the nuts and bolts of computer vision for biologists with session on results validation and what to present in publications.
- Data management for big data, data annotation for AI, AI for individual ID from photo databases.

- Some basic AI coding / getting started / using typical software in use by industry.
- One-on-one sessions with more industry teams (WildMe, Vulcan, NASA, Microsoft, Pix4D Structure from Motion).
- How to pick a cloud service, data type/size considerations, costs, which are approved by NOAA, etc.
- Additional R or Python training for integrating machine learning.

If you participated in training sessions, do you feel your skill level improved from these opportunities?

Yes	9	
No	1	
Maybe	2	
No response	5	*(Industry, 3)*

What benefits did you gain from attending this workshop?
This workshop illuminated problems within NOAA, including gaps in current software that needs remedies. A better understanding of CNNs and what challenges they're best suited for solving. A better foundational understanding of the current state of technology, on VIAME, and capabilities for AI approaches, especially a realistic grasp on these elements. Networking with fellow scientists using AI and industry experts. Hands-on experience as well as demonstrations were helpful, especially opportunities for one-on-one discussions with some industry experts.

Did this workshop allow you to network with other NOAA scientists and experts in the field?
All respondents agreed this workshop was a successful networking event.

Do you think these new connections will lead to future collaborations?

Yes	8	*(Industry, 2)*
No	1	
Maybe	8	*(Industry, 1)*

Please describe any new contact, project, or potential collaboration that has come out of this effort.

- Industry experts expect NOAA biologists to follow-up on potential collaborations.
- Potential USGS and NOAA collaboration on advances in single track aerial imagery for structure from motion application.
- Connection with folks using VIAME at AFSC.
- Steve Rader (NASA) collaboration was very beneficial.
- Potentially collaborate on a joint grant proposal for future AI analysis.
- Further learning and collaboration with Polar Ecosystems Program (AFSC) on harbor seal imagery.

Having seen what is possible, do you see a new way to approach a research question or problem that you or another colleague have?

Yes	11	
No	1	*(Industry, 1)*
Maybe	4	*(Industry, 1)*
No Response	1	*(Industry, 1)*

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

Acting Under Secretary of Commerce
for Oceans and Atmosphere
Dr. Neil Jacobs

Assistant Administrator for Fisheries
Chris Oliver

September 2019

www.nmfs.noaa.gov

OFFICIAL BUSINESS

National Marine
Fisheries Service
Alaska Fisheries Science Center
7600 Sand Point Way N.E.
Seattle, WA 98115-6349