THE PEW CHARITABLE TRUSTS

2nd Global Artificial Intelligence in Fisheries Monitoring Summit Report

2024

Background

The Pew Charitable Trusts hosted its Second Global Artificial Intelligence in Fisheries Monitoring Summit (AI Summit) in Honolulu, Hawaii from January 23-25, 2024. It convened 35 global experts (from Asia, Latin America, North America, Europe, Australia, and the Pacific Islands) to discuss how artificial intelligence and machine learning (AI/ML), in conjunction with electronic monitoring (EM), can help secure greater transparency and accountability of the world's fisheries. Attendees presented their experiences developing and implementing AI technologies and associated emerging technologies and discussed real and perceived obstacles to expanding the use of AI, and ways to address and overcome those challenges. The three days were divided into seven sessions, which included presentations, panel discussions, Q&A sessions, and breakout groups. Many of the topics discussed were first raised during Pew's first AI Summit and EM Service Providers for RFMO Engagement Workshop in 2023 and Global Electronic Monitoring Symposium in 2022. The sessions, presentations, and main takeaways from each are summarized below.

Pre-Summit Survey

Before the Summit, attendees completed a short survey to gather information on their role and background, how they define AI, how they are currently using AI for EM, and where they would like to use AI by 2025. Attendees represented different sectors including software engineers, fisheries managers, fisheries scientists, Secretarial staff, and non-governmental representatives from environmental non-profits and academia. Although AI was defined in a variety of ways, common themes included concepts such as automation, simulating human intelligence, and the use of algorithms to produce desired outputs or knowledge.

When asked about their current work, many of the attendees said they were using AI or developing algorithms to detect fishing events, to identify target species, or account for and monitor bycatch. A smaller number of participants responded that they were using it to estimate size or length, track objects, estimate effort, monitor compliance, or identify social indicators (e.g. crew wellbeing). Looking to the future, attendees suggested using AI for tasks like automated video clipping, rare event detection, environmental assessments, identifying closely related species, improved size and weight estimation, and on-vessel edge computing.

When asked how AI could improve data collection and fisheries management, participants identified several areas, including providing real-time monitoring and predictive analytics to allow for dynamic, real time management, better estimation of total catch and bycatch, synthesis of data across different fisheries providing fine scale data and insight into what is happening across the ocean, expediting data analysis, transfer, storage, and transparency, and reducing time spent on repetitive manual tasks.

Stakeholder Perspectives, Motivations, & Intersections

The three previous EM-related events hosted by Pew highlighted the fact that the use of EM and AI leads to the involvement of groups that have not been traditionally engaged in fisheries management. In developing and implementing an EM or AI program, fisheries managers need to collaborate with EM vendors and AI software developers. Along with new terminology and ways of working, there is significant complexity in each of these fields that is not necessarily familiar or easily understood by the other groups.

To support and facilitate more effective collaboration between these groups, participants from the first Summit developed an AI/ML glossary geared towards fishery stakeholders in the electronic monitoring space. Before the second Summit, the Steering Committee compiled an initial list of stakeholders and their interests, perspectives, and motivations in relation to AI/ML, identifying overlaps and intersections between the groups. This initial mapping provided a foundation for a group discussion on how to better understand, map, and take advantage of the relationships between stakeholder groups. Attendees were asked a series of questions including:

- How do different stakeholders consider and approach AI/ML and why?
- What are the unique considerations and requirements needed by each party to successfully integrate automation into a project?
- Where are the opportunities for beneficial collaboration to ensure success in integrating automation?
- Are there specific issues around technical language or divergent motivations?

After a general discussion to identify any missing stakeholders, perspectives, motivations, drivers, or interests, attendees joined breakout groups to have deeper discussions on specific stakeholder interactions, goals shared between the stakeholder groups, and any strategies for overcoming competing or conflicting motivations.

Participants noted that there is significant interest from all stakeholders, principally from governments, EM and AI vendors. Participants also noted that other stakeholders, for example industry companies and crews, third party certifiers and consumers, also have a direct interest in the issues arising from the development and/or implementation of AI. The participants noted that there can be significant overlapping interests and that the development and implementation pathway is not linear. Rather there are many ways and opportunities for partnerships and AI development to occur. The biggest takeaway from the session was that it is essential for the conversations among stakeholders to occur as early as possible, and that stakeholders should go into the planning process without a pre-expectation of the desired ‘solution’, but rather hold discussions to fully understand the ‘problem’ and then develop a ‘solution’ that best achieves the desired result. The results of the discussions will be combined into a document that will be circulated in the future.

Data Sharing – Advantages, Challenges, & Pathways Forward

At the previous Pew events, availability of data has been consistently raised as a critical bottleneck in the development and expansion of EM programs that incorporate AI/ML. As such, this session was designed to explore the critical aspect of data sharing, including the advantages and opportunities of data sharing while also being realistic about the challenges. This session included three presentations that focused on why data sharing is needed, utilizing open training data libraries, and transferring AI models across fisheries.

The first presentation highlighted three main reasons why data sharing is important. One, that the availability of labeled data is the largest bottleneck affecting AI model development. The beginning phase of AI/ML development relies on significant volumes of data and having access to a sufficient amount of data to both train and test the model is critical to the development of robust and accurate algorithms. Having access to shared and labeled datasets can help provide that data. Second, shared data can increase auditability of a model, which ensures proper functioning of AI over time. As identified above, it is important for AI/ML models to be tested on datasets different to those they were developed on to ensure that the model is working as intended. Third, the use of shared data can strengthen and improve the model's functionality and transparency. Having access to a range of datasets enables models to be tested on new or unseen data thereby identifying any errors or biases in the model and allowing for continued refinement and improved functionality.

The presenter also introduced three key considerations for developers when they begin to utilize data sets:

- What type and how much data do I have (e.g., video, imagery, acoustic)?
- What type of annotations or cleaning does the data need?
- What information does my data set have and does it cover the necessary range of information (e.g., weather conditions, day/night, vessel type)?

Finally, the presenter called attention to the FAIR principles and best practices for making data accessible and shareable.

- **Findable**: Metadata and data should be easy to find for both humans and computers.
- **Accessible:** The exact conditions under which the data is accessible should be provided in such a way that humans and machines can understand them.
- **Interoperable:** The (meta)data should be based on standardized vocabularies, ontologies, thesauri etc. so that it integrates with existing applications or workflows.
- **Reusable:** Metadata and data should be well-described so that they can be replicated and/or combined in different research settings.

The second presentation highlighted a project that used multimodal datasets to develop an AI model. This project combined images and videos from open datasets with other sources of fisheries information such as logbooks, geographics, event detection. The project successfully

developed an AI model using the multiple data sets and found that it produces better informed models and requires less label data and therefore less time labeling datasets by humans, which is the most resource intensive component of the program.

The third presentation discussed the lessons learned when applying the same model across a variety of fisheries.

In relation to object detection, the presenter noted the following lessons:

- There was a need for a lot of training data;
- EM system installation can really impact the AI model's performance; and
- Different types of fishing can require very different models.

In relation to object tracking, the presenter noted that:

- Successful object tracking translates well across object types;
- Tracking depends on fishing routine, so the ability for models to be applied across fisheries is impacted by new vessels and fishing types; and
- As the pace of fishing increases, tracking performance of the model needs to be maintained.

Finally, for object classification, the lessons included:

- Non-target species or less common species are hard to classify correctly;
- AI models will make similar errors to the ones human observers make; and,
- It might be preferable to aim for high recall when tuning classification models.

Participants noted the value that has been derived from The Nature Conservancy's Fishnet.AI Open Images Database, including its use during the development of early models. Participants noted that there is likely to be significant benefits if the next generation iteration of the Fishnet.AI platform is developed, including the significant advancement of new AI models.

Evaluating AI Performance

This session started with a presentation that explored why, how, and when AI algorithm outputs should be evaluated and verified. The presenter identified *recall* and *precision* as two baseline metrics that can be useful for measuring algorithm performance. They also discussed confidence values, which when combined with a chosen threshold, can be used to determine when a model is producing results that are "correct" or "incorrect." Thresholds can be chosen by hand, based on what "looks good", automatically by optimizing the F1 score (also known as a F-score or F-measure), a metric that combines the recall and precision elements, or visually on a graph that compares precision and recall.

Discussion focused on the importance of early and ongoing conversations between fisheries managers and AI developers to ensure there is a clear understanding of the objectives and motivations of the project, as well as the available resources, including time, data, and money.

Participants noted that the specific 'fisheries problem' combined with the available resources will heavily influence which metrics and thresholds are realistic and achievable for a given project, and how the developer will create the most cost-effective solution. Importantly, participants noted that decisions surrounding metrics is likely a policy-focused decision but should be guided by input on how the data will be used and the associated risks and consequences of choosing that metric. These early conversations may impact the way the model is developed, refined, and ultimately how the data generated from the AI model is used in decision making.

AI Contracting and Procurement

This session focused on 1) highlighting the complexity associated with developing and responding to AI procurement requests and 2) an in-depth discussion of the procurement process, including examination of the different perspectives of the involved parties.

The session begun with a presentation on some general principles when considering procuring the development of algorithms for use in EM programs:

- As complexity of the desired task and the level of desired performance goes up, costs increase, and the robustness and generalization of the model go down.
- Understanding program needs, data availability, timing, and cost are important to determine early in the procurement process. Clear communication on these issues will help calibrate expectations for all parties.
- Model development generally reaches a point of diminishing returns – 20% of the effort is spent getting 80% of the way to the end result, while the remaining 80% of the effort is spent getting to the final 20%.

The second presentation focused on the experiences of an EM provider who incorporates the use of AI and ML into their EM services. They have noticed that the use of AI is becoming an increasingly common request from those groups looking to procure EM and have noted several important considerations that can help facilitate efficient procurement of AI models.

- The importance of clear specifications in the request, which can vary and will depend on if the client is seeking to develop an AI algorithm or if they are looking to incorporate existing algorithms into an EM program.
- A good understanding of the market and EM provider business models, which can generally fall into two categories – licensing an AI algorithm and using it to analyze data or sending the data to a third-party AI company to be analyzed.
- Realistic expectations when it comes to costs and capabilities of AI and building a recognition that AI is a cost savings tool, not a cost elimination solution, and that sometimes the marginal costs of utilizing an AI model are going to outweigh its benefits.

Finally, the presenter provided six recommendations for clients seeking to develop a request to incorporate AI into their EM program:

1. Understand AI is a tool for enhancing and supporting the work of reviewers, not replacing them.
2. Evaluate the trade-offs in AI implementations across technologies and fisheries, and consider if human reviewers make more sense, either logistical or economically.
3. Understand the limitations of AI in terms of computational resources, available image libraries, inferencing costs, and accuracy levels.
4. Determine the best place and environment to run the AI model (e.g., onboard, locally, in the cloud).
5. Dedicate appropriate long-term resources, as AI will require continuing costs and resources, even with the best AI.
6. Be realistic about the capabilities and costs of AI when setting performance standards and specific data requests, (e.g. not every data field of a logbook or observer book need to be validated).

After the presentations, the attendees reviewed an example Request for Proposal (RFP). The activity was designed to identify opportunities and takeaways that could be used to strengthen future RFPs by exposing participants to the perspectives of those writing RFPs (i.e. government officials) and those responding to the RFP (i.e. technology providers). The discussion generated a list of suggested best practices for future RFPs:

- Including clear objectives and metric(s) for measuring success.
- Providing enough information to allow for the Provider to calculate effort or margin.
- Focusing on the desired outcomes of the proposal, rather than being prescriptive about the tasks or inputs to complete the project.
- Developing a list of minimum standards that bidders must meet.
- Linking to examples of desired elements.
- Including a list of available data from the beginning to allow realistic estimates of cost and time needed.
- Allowing for modular contracting, which allows for multiple potential providers to revise their bids based on a review of initial datasets.
- Ensuring that the RFP process is truly open and that a broad number of companies could meet the project requirements.
- Providing clarity on who will own the intellectual property.
- Having conversations with potential providers before developing the RFP to ensure requests are realistic.
- Involving data scientists in creating and writing the RFP.

AI Ethics & Addressing Privacy Concerns

The two presentations in this session sought to dive deeper into previous discussions surrounding privacy concerns related to the development and use of AI algorithms in conjunction with electronic monitoring. The first presentation focused on ethical considerations

surrounding AI-assisted monitoring including risks of intrusive surveillance, lack of consent, data exploitation, and over reliance on technology. EM providers and other stakeholders will need to be aware of emerging government regulations, IGO and NGO standards, and ethical best practices when utilizing AI as part of an EM program. They should also consider adopting mitigation measures, including increased transparency, impact assessments, privacy-driven camera setups, pseudonymization to protect individual privacy, and audit and compliance checks.

The second presentation introduced an ongoing pilot project that aims to monitor working conditions onboard fishing vessels using EM. While AI is not currently being used as part of the project, it could be incorporated in the future to help protect crew privacy by automatically blurring faces and identifying information; detecting when crew are on deck to help quantify working hours; detecting man overboard events, potential abuse, and proper use of safety equipment; and providing real time alerts of serious event through the use of onboard edge computing.

Privacy and confidentiality concerns will remain but there are opportunities currently available to manage and mitigate much of the risks.

Optimizing AI and Reviewer Partnerships

This session included two presentations on the use of AI in regional EM programs and was followed by general discussion about lessons learned, pitfalls and expectation management. Both presenters emphasized the importance of continuing to employ human observers, both to provide a direct source of employment, but to also serve as a pool for future fisheries managers or scientists. The repetitive nature of the work was also raised, with discussion about opportunities to alternate time observing onboard fishing vessels with time as a “dry observer,” reviewing video footage with the assistance of AI.

Areas of Future Discussion

For the final session of the Summit, participants were asked to propose topics that could be covered in future gatherings. Suggestions included:

- How to best integrate AI into existing EM systems and programs, including ensuring interoperability
- The use of AI in legal, enforcement, and compliance proceedings
- Increasing accessibility to requests for information (RFI) and RFPs
- Data sharing logistics, including databases, warehouses, and frameworks
- Incentivizing data sharing, including cost recovery and licensing models
- Evaluating performance and the use of metrics
- Privacy and labor issues
- Broadening participation by including representation from fishing companies and cloud vendors

Acknowledgements

Planning and guidance for this event was provided by the expert steering committee:

- Amanda Barney (Teem Fish, Inc)
- Jamie Gibbon (The Pew Charitable Trusts)
- Raiana McKinney (The Pew Charitable Trusts)
- Claire van der Geest (Seven Seas Consulting)
- Benjamin Woodward (CVision AI)