Discovery of Sound in the Sea (DOSITS) Webinar
How are Passive Acoustics Data Used to Inform the Decision-Making Process?
Dr. Elizabeth Henderson, U.S. Naval Information Warfare Center (NIWC), Environmental Readiness Branch
October 22, 2020
Webinar archive: https://dosits.org/decision-makers/webinar-series/2020-webinar-series/passive-acoustics-data/

Question and Answer Summary

(questions asked/answered in real time during the webinar or in the Q&A box)

Questions asked/answered live by Elizabeth Henderson during the webinar

How do you define the radius of interest around a recorder? Could we define the distance between the whale and recorder?

That's a great question. There are two things to look at. Your hydrophone, that instrument with the microphone that's recording the sound, it's going to have some limitations. It's going to be set up to record a certain DB sensitivity (intensity). You can have a hydrophone we would say more sensitive, it's going to detect things further away, less sensitive means that it's going to be able to get louder stuff without the clipping. You also have to think about how the sound that you're interested in is propagating through the water. High frequency sounds propagate much shorter distances and low frequency sounds, things like blue whales, you know, you can pick up their calls hundreds of kilometers away, whereas something like an echolocation click, you maybe you're only going to get within a few kilometers. Your radius of interest is going to depend on your instrumentation itself, but then also what you're interested in in listening to.

Do people work with hydrophone arrays that are attached to recorders that have different types of hydrophones so they can look at it, or are people generally looking at specific types of calls or specific species when they deploy these?

I am not familiar with any recorders like that it's a single recorder with different sensitivities of hydrophones. But I know of studies…we did a study actually up in Alaska, where we used multiple recorders that were set at different sensitivities, so that we could capture loud sounds, but we could also try and capture quiet sounds.

When you hear a sound of a particular species, I think the cause of the right whales are great example…Do you want talk for a little bit about the algorithms and how accurate they are at detecting a particular call. For example, with the right whale monitoring off of Boston, I know those [detections] trigger an alert and there's a human that reviews it but how well does the automation work, and particularly for those recorders, are they set to turn on, to record only

when it detects the call. What is the reliability of this? (Are there algorithms available that can accurately detect mammal sounds (and thus start recording when those specific sounds are heard?)

Let me start at the end of that question and work backwards. I know that there are some recorders that are duty cycled for detection files. They have the detection algorithms in them and they turn on and they listen and they record whatever detection they're getting. But also, if a detection occurs, they are triggered to record a WAV file for a period of time. That does exist. But you can have something that is a compromise between both. The reliability, I think it depends on a number of things. You need to look at how regular that signal is. The right whale up call is a really good example of this signal that's very regular and very similar every time it's produced. That's going to be something that you can have a pretty good detector for, and that's going to be pretty reliable. And yes, you still want that human review factor, but those detections are going to be pretty reliable. The CPods that I mentioned are click detectors, and those are just looking for clicks and then you can set those [CPods] for different bandwidths. We know porpoises produce only very high frequency clicks. So you can say, I just care about porpoises…I only want clicks that in this range versus I want to know all clicks that happened, and then there might be some dolphin sounds in there as well. Once you start to get a little bit more complex and you're looking at really variable stuff then it's going to be a lot more human analysis. I'm just trying think of any examples of recorders that have detection algorithms for more complicated vocalizations, and off the top of my head, I'm not thinking of any. I think in most cases when you're starting to think about things like whistles or pulsed calls or, you know, humpback song, or things like that, you know, certainly you can say, yes, you know it's humpback song, we know what that is. But if you're actually looking at the notes and things like that, that needs to come back to a human analyst, generally. But there are a lot of really great people working on these questions, and I am sure that there will come a point where some of this stuff will be will be automated.

How do you identify the source? How can one be sure of the sound source? [How do you identify the source…how do you decide, these are dolphins, those are porpoises these are whatever? To non-experts the species' sounds, sound familiar]

That's definitely true and that's one of the things that is actually being worked on right now. Researchers are trying to distinguish between very similar sounding species, a lot of the delphinid species sound very similar. They have whistles in a similar frequency band and their echolocation clicks are in a similar frequency band. And there is some fantastic work being done right now, again aggregating all of that information and then automating the decision making to say, well, if it's this whistle and this click type, then it's most likely this species. But all of that, it doesn't matter what species you're talking about, all of that has to start with having a recorder out and eyes on the water. You have to start by seeing the animal and recording it at the same time and being confident that nothing else is in the area that could be making that call. And that's where we build one hundred percent of our information. For example, there are some species of beaked whales right now where we know where the species are distributed based on where they've been identified visually, and then we have some acoustic detections that are very reliable and we know those detections, but we have never put them together. We've never seen, say Gervais' beaked whales, and heard that click at the same time. So even

though we suspect this is their click, we don't know, and we can't say, until we actually confirm it. It always starts with the visual identification.

Combining two questions about looking at the data. A lot of the acoustic data when it's being imported, is it band pass filtered so that you are getting to your species of interest? How often or how much possibility do you think there is to go back and look at old data because it's only the band pass filters part that is saved or is there other information there that could be processed as we learn more, maybe about a particular species? [Is most acoustic data bandpass filtered for the species of interest or is it possible to go back and look for something new in old data?]

That's a great question. And again, I think this is one of those things that you actually really have to think about ahead of time when you're planning your study. If you, in this particular study, are only interested in one very specific species, you can absolutely band pass filter your recorder, whiten out all the other noise. For those of you that don't know what that means, band pass filtering is just putting limitations, so you're basically whitening out, or cutting off…if you do a high pass filter you're cutting off low frequency noise, if you do a low pass filter you're cutting off the high frequency stuff. This is commonly done, for example, let's say you're looking at dolphins that are higher frequency, but you're in a busy shipping area. If you don't care about that shipping noise, and that really is truly just noise and you're never going to look at that, you can high pass filter that out. It makes your data much cleaner, makes a higher frequency part of the data much cleaner. You absolutely can do that at the recorder. What is more typically done is full band recording, bring it into your lab, and you can band pass there, but then you still have that original raw data that you can go back to and look at. And that, to me, is really critical. Having that data that you can go back to. Something that we do in our lab- we're constantly developing new detectors and classifiers and going back to old data and reanalyzing it. So having that full band where you can really ask different questions down the road that you maybe didn't think about at the time is critical. I prefer to do it that way but you absolutely can do it at your recorder as well if that's what you want to do.

So there's a second part of this, which is which we sort of started down the path. And that is, you know, are you aware of any studies where people have sort of gone back and mined hydrophone data to sort of expand our estimates of say species' populations or ranges? [Has anyone mined the data to find: (1) out how many species you can identify? (2) rough population size for each of those species? (3) rough location(s) of high populations? (lobster, crab, sea urchins, finfish, seals, turtles, ...]

For sure. Like I said, you have to have that data, you have to have that raw data stored. While we think of passive acoustics as a relatively new science, it really is, you know, a few decades old at this point. And there are some great labs that have been doing it from the early days…Cornell's lab and the Scripps lab, and there's a number of groups that have decade's worth of data and people definitely go back and look at it. You know, things like knowing how the oceans are getting noisier now, that necessarily wasn't something that you were looking at in the early days of doing passive acoustic monitoring. But if you have the data, you can go back and you can say, well how low was that ambient background level, you know, in the early 90s, or whenever you started versus today. There are definitely people doing that.

Can you can you determine population density using only whistles, so not clicks…you cited a couple of studies that used acoustics and/or visual or different acoustic methods to look at population densities. Could you talk more about, you know, using only acoustic methods to determine population densities. [Is it possible to determine densities using only whistles (not clicks) and only for acoustic methods (without the visual confirmation)?]

Let me preface this by saying this is not my area of expertise, and again, there are other really amazing people out there doing this work. There's a woman, Caitlin Frazier, who is pioneering acoustic density research and Tiago Marquez at University of St. Andrews. I mean, there are some fantastic people that do this work. That is not me. But, as I said with acoustic density, you still have to have some information and that information, again, typically goes back to some visual data. For group counting methods for example, you have to know how many animals are typically in your group, and you're only going to get that originally visually. But yes, you can use any cue. Your cue can be anything you want. The two examples I gave of the Blaineville's, these whales don't whistle. They only echolocate. Then they have these pulse calls, that are these tonal calls that happen at depth, but they don't whistle like a like a dolphin. So they had to rely on their clicks for that method. Whistles are actually far easier to use. Clicks are really difficult, especially if you have multiple recorders when you're trying to match up where an animal is. So if you're trying to if you're getting a number of clicks and you want to say is that one group or is that maybe two different groups and you have multiple sensors, in order to localize where a group is you have to be able to align those calls on your different sensors. And clicks are really, really difficult to do that with. Whistles are fantastic to do that with because they have a very distinct shape. And so you can look at your data, and you can say, oh, yeah, that's the same whistle that showed up here, here, and here. So we know that that group is here, and no, there are no other groups in the area, so we know that that's one group that we can count. So you can definitely use whistles as your acoustic cue, but, again, you have to know information about how often your species whistles, does that change with their behavior. For common dolphins their whistle rates change drastically, with different behavioral state, and group size, as well. So there is information that you have to know to use these cues to do the acoustic densities. But once you have that, it's a really powerful tool.

We talked about putting out a recorder. We talked about…there's a hydrophone it's attached some hardware that that then save that data, a hard drive, whatever. What's a typical deployment? It's not a single hydrophone in most cases, is that true? Are people putting out arrays or there's also floating ones out and all kinds of things? Could you talk about what a typical deployment might be?

I don't even know if there's such a thing [as a "typical deployment"]. Certainly, things have evolved as well. Let's take the high frequency accorded recording package that Scripps has. That's the thing that sits on the seafloor, gets deployed from a ship, and then it gets recovered. That originated as a single hydrophone that floats above, or there might be two in there, but it's just very simplistic. And it was a big heavy frame, something that was robust that could really handle sitting on the seafloor for months at a time. That has really evolved. They now have ones that have four elements in a tetrahedron, so they can do localization. They have them

now as moorings, so they have multiple hydrophones on it, but then they also have other sensors like a ADCP to look at wave movement, and echosounders to look at prey and just all kinds of things. That's just, that's an example of one type of recorder that has evolved into many different forms. So again, it's your recorder… what you're going to put out is really dependent on the questions that you're asking. You know you can you can use a single hydrophone on the single recorder and a lot of really good studies do just that, just to get a sense of what's going on. And there's a lot of great data that you can find just from a single sensor. In Hawaii, they discovered that the vertical migrators, the mesopelagics, when they're doing their dial migration at night, in Hawaii on the islands, not only are they moving up and down in the water, they're moving in and out. And so, if you have a single recorder sitting there, you really can get almost the soundscape as that happens throughout the course of like a 24-hour cycle and you, you see what that how marine mammals are responding to that. You see how they're coming in and out as that happens. There's fantastic things you can do with a single sensor. But, but if you want to localize if you want to try and get some of these other questions, you know, then you need to start adding in additional sensors. So, it just depends, although I will say, Danielle Harris at St. Andrews was able to estimate density with single sensor data, which is like mind blowing how you do that, but she was able to do it. So yeah, I don't, I don't think there are any limits, except for yourself, or how you can do things.

"Big data", it's a buzzword. Is big data analysis being employed to look at the data from some of these larger arrays and then things like the decision trees used by the plankton identifying cameras. Is that kind of stuff going on? And I don't know if you've been involved in any of that. [How much is big data analysis being employed to analyze the data? Things like decision trees used by the plankton identifying camera from McLane in Falmouth, MA?]
Absolutely. That's definitely a field in which, you know, where passive acoustics is headed and is actively happening. Things like machine learning is really a huge area of growth right now. It's not necessarily the large ranges…I mean, we can do things like machine learning as well because we have you know a ton of sensors and a lot of data. A temporal record can be just as useful. And that's what people are doing where they're training these automated methodologies and they're feeding in the data and then they're getting out some really cool stuff. Ann Alan actually coordinated with Google to use some of their machine learning algorithms to try to do humpback whale song identification, and that was just a really cool, fun project. That was, you know, an example of machine learning application. A lot of the newer classifiers, especially with those odontocete species that we were talking about that are so similar, when you start putting in all of the you know the whistles and the burst pulses and the clicks and everything else, machine learning is really the place where those classifiers are getting good. They're getting really amazing. And it's the big data. And it's just kind of putting it all in there and doing some training and saying, this is, this and this is this now go. And it's phenomenal what's coming out.

For real time networks (e.g. Whale Alert), is there any communication back to the sensors? Can you adjust you know the gain or the or the way its operating? If you are suddenly interested in a different species or something else…Is there a two-way communication with

the sensors, or are they just sending out the satellite signal and not get anything back? [Can the gain be adjusted dynamically (automatically) to capture a wider range of sounds?]
To my knowledge, it's one way. I am not familiar with anything where you can actively change settings on something that's deployed. Now, I might be wrong, there's a lot of work going on out there…stuff is changing all the time. So, there could be. But to my knowledge, no. To my knowledge, it's all one way. You have to, you have to preset everything before you put it out.

If you put out a couple of hydrophones recorders and they're out there for a year…but, you get them back and you put it through the algorithms, give it to the machine learning, you know, what's the turnaround time on that? I ask that only because in a decision-making situation, you know, looking at in Rhode Island off the coast, here we have the Block Island Wind Farm and they want to build, you know, hundreds more turbines, and that's great, but, there's a there's a decision making process that happens, and data gets collected…And so what kind of timeline do you think people are thinking about, when can I make a decision, what is my process?
Again, I think that's a question based on if you're going to do something like presence/absence. That's going to be really quick. Or, if you just want to say I want to know if my species is there, this time of year, or this whatever, you can just really quickly run through the data. You can use your automated processes, you can just say, oh, yes, there was a detection. So yes, they're present and you don't really need to go any more into it than that, it's going to be pretty quick. But for a lot of what we do you, want to do a little bit more than that. You want to get a little bit deeper than that. Even though automation is becoming wonderful and we're getting better and better at it, we really still rely on the human analysts to do the to do the bulk of the work, or at the very least check. In our lab we do a lot of a lot of automated things and then we have to do some checking, just to make sure…we don't want to be saying things are one way and then you know our algorithms got messed up, and we were totally wrong. So, it does come back to the human analyst and how much the human analyst has to do depends on the questions that you're asking. Things like propagation modeling, and if you're trying to get received levels, and you don't have something like an acoustic tag or you have the level at your instrument that you want to, you know, propagate out to somewhere the animals or something like that, you know, that kind of thing takes a good amount of time. There's not an easy answer to this question. It really just depends on what information you need for your decision making.

Could these risk functions change as animals become more and more habituated/sensitized to the noise? How to control for that?
I can't speak for any other organization, but the Navy does update their risk functions every 5 years - the permits that the Navy gets from NMFS are for a 5-year period, and so the analysis methods can't be changed more frequently. However, we are continuously reviewing and assessing new research and information, and all new data is integrated into the next round of criteria. So, if some kind of sensitization or habituation was observed and was quantifiable, it would be incorporated into the next round of criteria.

I wanted to know more about the actual decision-making process of changing stakeholders' minds. Do all of them grasp this kind of information? What can we as scientists do to make the message come across more easily?
Who do you mean by the stakeholders? Are you talking about the regulatory organizations such as NMFS, BOEM, or the Navy? To my knowledge all of these organizations are continuously reviewing new science as it emerges. As mentioned above, criteria are typically developed over some time period, and these criteria go through public review and so can't just be constantly changed to incorporate small bits of new information. But on the cycle of each organization, new data is incorporated as possible. My recommendation for scientists is to make sure data is published; it is very difficult to use reports or presentations in our criteria, but published data is usable and defensible. Also, do participate in the public review process; your feedback is always read and incorporated where possible.

Questions asked/answered in the Q&A box during the webinar

(answers provided by members of the DOSITS Team)

Are whales always passive acoustic? Or, are some vocalizations active acoustic? If some are active, are Navies considering using active whale sounds for active acoustic sensing?
Dolphins use active acoustics when they are doing echolocation to find food or avoid obstacles. Scientists have used the characteristics of those signals to develop similar signals, but I am not aware of Navies using dolphin signals themselves.

Are you aware of any research monitoring fish behavior in the wild using passive acoustic devices?
There has been some work with fishes during reproductive stages when they are vocalizing (males trying to attract females to their nests). There is also a lot of work where fish are tagged with active tags and VEMCO recorders listen for the tag pings to document animal movements and migrations. Check out this DOSITS page: https://dosits.org/people-and-sound/investigate-marine-animals/how-is-sound-used-to-study-the-distribution-of-marine-fishes/

Are acoustic tags glues on or stabbed into the individual?
Acoustic tags can be attached via a variety of methods, including attached with a long pole (and affixed via suction cups), shot from a gun or crossbow to be embedded in the skin or blubber layer of the target animal, or they can be glued directly onto an animal's body. For more info, please visit this page on DOSITS: https://dosits.org/galleries/technology-gallery/observing-and-monitoring-marine-animals/passive-acoustic-recording-tags/

What kinds of seafloor (sediment types) host the material that whales eat when they bottom-feed?
When marine mammologists talk about bottom feeding, they mean deep in the water column, not necessarily right on the seafloor. However, there are some species, like gray whales, that

scrape the seafloor to eat amphipods and other critters, typically in softer sediments.

Are tools like the whale watch or apps that are available to mariners used in a regulatory way. For example, if right whales are known to be in an area is there an alert that vessel speed must be less than a certain speed in a particular area?
Yes, the exact purpose of the apps is to alert mariners when whales are in the area and to caution them to slow their vessel speed to less than 10 knots within a specific area that is defined by the vocalizations that have been recorded.

Does the Navy's hydrophone range include visual (camera) sensors?
The hydrophone ranges are pretty far offshore and spread across large areas. I don't believe there are camera sensors; the hydrophones are powered by cables on the seafloor.