My name is Jill Munger, I’m a Fisheries and Wildlife undergraduate student working with the CIMRS acoustics lab as a research assistant. I've spent a lot of time underwater as a diver, and I'm fascinated with how fish communicate with each other. Advances in acoustical science give us the ability to study this in detail, and I've spent the last 9 months doing just that.

Fish have developed the largest diversity of sound-generating mechanisms among vertebrates. For example, these black drum fish will drum on their air bladder with specialized muscles <clip>, and there are catfish that rub their bones together to create sound <clip>.

We can calculate a visual representation of sound using spectrograms. Spectrograms have frequency on the y-axis which can be thought of as the pitch of the sound, time on the x-axis representing the length of the sound and energy or the loudness of sound represented by the brightness of color.

Damselfish create sound by snapping their pharyngeal teeth together. The resulting sound is resonated by the swimbladder <play clip>. They can create pops and clicks which are generally associated with aggressive behavior and nest defence, but the sound you heard in the clip is called a pulse train and is associated with the courtship display of the males. The various sounds are not only correlated with different behaviors, but also with different species, and even different individuals.

My research focuses on identifying temporal patterns of fish vocalizations, which is critical to understanding how environmental factors and conditions are changing over time in response to climate change and shifting fishing practices. Unfortunately, the ocean creates all kinds of challenges for underwater research: saltwater is corrosive to electronics, communication between researchers is difficult, and the deeper you go, the more pressure you and your equipment have to deal with. More traditional field methods like towed-diver or snorkel surveys are labor intensive and can be prone to mistakes.

Hydrophones, however, are minimally invasive, can record for long-periods of time, and are comparatively cost effective. My research escalated quickly when I received more than 2 years of continuous recordings from a 4-year period of time in American Samoa. Noise Reference Station #10 was deployed at 33 meters within a marine protected area managed by the National Park Service in proximity to fringing coral reefs. The challenge now was to see the ecological stories in the data.

First, I identified the calls as damselfish by listening to recordings and comparing the acoustic attributes with the literature. I tried manually sampling the data by visually and aurally scrolling through spectrograms, but because damselfish vocalize so frequently, I only got through an hour of data in about 45 minutes real time. I needed a more efficient way to look at 19,000 hours of sound and I wondered if I could utilize machine learning to automate the process?

Machine learning is a way of defining a computer program by learning from data instead of writing it ourselves. We can have the computer perform a task without explicitly programming each step. Deep learning is a subset of machine learning that uses neural networks with multiple layers. For complex recognition tasks like detecting fish sounds in a very large and noisy ocean, deep neural networks have proven extremely effective.

My colleague, a machine learning engineer, became a big part of this research and together, we trained a deep neural network to detect damselfish vocalizations. We leveraged state-of-the-art image recognition models by training them with our spectrograms. We then used our model for inference over the entire dataset, broken up into 2-second segments, detecting the presence of calls in each segment. Our data reflects the percentage of time in which vocalizations were detected by the model.

Of course, it’s incredibly fun to listen to the underwater soundscape, but there are also some pretty interesting ecological applications of bioacoustics. Vocalizations of damselfish are presumed to change in occurrence over time, so understanding the variation will give us a baseline for comparison as conditions shift. As human influence causes changes in the environment, we presume the vocal behaviors of the fish will also change.

One of the things I’m investigating is if damselfish sound can be used as an indicator for coral reef health. Some species can actually be detrimental to the reef by fueling the overgrowth of macroalgae. Another interesting aspect is the effect of overfishing on more commercially important fishes, such as snapper and grouper who feed on damselfish. But in order to understand how these interactions are changing through time, I first needed to know if machine learning techniques are capable of detecting fish vocalizations at all.

We were able to train the model to a self-reported accuracy of 94%! Empirical testing shows an accuracy rate of 79%.
The hydrophone was deployed at 3 distinct time periods with breaks in between. In the first period of data, we see that vocalizations were detected 5.5% of the time, on average. In the

second time period, this dropped to 3.3% and 3.6% in the third. If we ran the model over an even longer period of time, is it possible that we would detect a trend indicating a decline in the vocal behaviors of the fish?

Looking at this segment of data, you can see significant seasonality, potentially related to the lunar cycle or courtship activities of the fish. Research has shown that damselfish reproduction is correlated with lunar and tidal cycles, presumably to synchronize with water conditions that are favorable for pelagic larval survival. Ideal timing happens if the larvae are swept out to sea by tidal forcing where there are fewer predators looking for an easy meal.

Zooming into just a week's worth of data shows diel variation in the vocalizations, with the mean percentage of calls being higher in the afternoon than in the morning or at night. Damselfish are algal mat gardeners, and the diel feeding hypothesis suggests that this pattern is related to the nutritional quality of the algae. The algae spends the morning hours photosynthesizing and developing nutrients and starches which is what the fish wants to take advantage of. It would make sense that there are territorial vocal displays during peak feeding hours. Damselfish also have a diel spawning cycle.

Of course this is all fascinating to me as a fisheries and wildlife student and as a person who loves to explore the ocean. But it should also matter to you. Seafood serves as the main source of protein for more than 3 billion people worldwide. Coral reef structures help to protect our coastal cities from storms, and the ocean does us the not-so-small favor of slowing down climate change by absorbing carbon dioxide from the atmosphere. In order to protect the species that provide these services, we have to understand the role of communication and how human activities are influencing the vocal behaviors of fishes. Machine learning is a technology that was developed for an entirely different purpose, but it has enormous implications on how we study the ocean.

Thank you to my supportive mentors, Joe Haxel and Samara Haver, who encouraged me to pick myself up, dust myself off and keep pushing forward and thanks to Dan Herrera, an incredibly talented programmer and very patient teacher. And thanks to all of you for spending a part of your afternoon learning about my research.