Q&A WITH AZA RASKIN AND KATIE ZACARIAN

Want to know what animals are saying to each other? ***Aza Raskin*** *and* ***Katie Zacarian*** *are finding out.*

AS TOLD TO
Brian Williams

In 2017, there was a significant breakthrough that transformed language translation: Machine learning techniques were developed, allowing computers to learn a geometric representation of an entire language, like a galaxy where each star is a word and the distance and direction between stars encodes relational meaning. While the technology has moved on significantly since then, at the time, it was revolutionary. And it inspired new thinking about whether these new techniques could be applied to help understand non-human communication.

"You could say it's ironic to use technology to study the world's natural communication, but I think it's poetic that we're using the very newest tech to remember something we used to know and forgot."

That is what Aza Raskin, the co-founder of Earth Species Project (ESP), had to say about using artificial intelligence to study how other species communicate. ESP, which Raskin launched along with Katie Zacarian and Britt Selvitelle, is a nonprofit focused on using artificial intelligence to decode non-human communication, with the ultimate goal of changing the way we relate to the rest of nature. In partnership with leading biologists and researchers, the organization's work is also supporting conservation on the ground today.

In an exclusive interview, Raskin and Zacarian share with *Lifestyles Magazine* their current projects, the potential ethical pitfalls involved with using AI, and how their work may be able to support conservation.

PHOTO COURTESY OF AZA RASKIN

LIFESTYLES MAGAZINE: What inspired you to found ESP?

AZA RASKIN: I'd been thinking about ESP since 2013 when I heard an NPR piece about gelada monkeys. The researcher said these animals have one of the largest vocabularies of any primate except humans. They sound like women and children babbling and the researchers swore the animals talked about them behind their backs (which is probably true). They were out there with a recorder, hand transcribing what the monkeys were saying, trying to decode their sounds.

It really struck me: Why aren't we out there with microphone arrays and using machine learning? I looked into it at that point, and the more we looked, the more it looked like there was really something there. But at that time, machine learning couldn't do something like translate between two languages that you didn't already know the translation for.

So when the technology came out that would allow for translation without a Rosetta Stone, that inspired the genesis of ESP. Essentially the impossible just became possible. What was previously science fiction became science.

From those original 2017 days, which is like the Stone Age in machine learning, a lot has happened that's made translations without examples even better. The language map shape is no longer the only technique available, and we are currently using many other tools and approaches as part of our roadmap toward decoding non-human communication.

LM: What drew you to Earth Species Project?

KATIE ZACARIAN: Every once in a while, you meet extraordinary people doing extraordinary work. Connecting with Britt, and later Aza, felt like opening a clamshell and finding pearls. Here were two very big minds thinking about how they could harness their expertise in the development of new technologies to address an extremely big challenge, possibly one of the biggest challenges we can imagine: interspecies communication.

The "Earthshot" we're aiming for at ESP is to use AI to decode non-human communication, and then to use the new knowledge and understanding that results from that to reset our relationship with the rest of nature. That was incredibly compelling to me as a potential unlock in addressing the biodiversity and climate crises, and helping us to find ways to coexist on the planet with other species.

Perhaps most exciting for me was that it was clear the models and techniques we would develop along the way would dramatically increase efficiencies in scientific discovery and provide more tools for those on the front lines of conservation today.

LM: How do you identify your areas of research, and what are some of the projects you're working on?

AR: We have recently published a technical roadmap that describes how ESP's research projects contribute to our goal of unlocking the potential for two-way communication with other species. The roadmap builds on the recent advances in the field of machine

RASKIN PHOTO BY ALICE LIU; ZACARIAN PHOTO BY XYNN TII

learning, including the rise of foundation models trained on unlabeled data at scale, which have transformed the ability of machines to process and translate human language.

We're building the fundamentals that will drive more rapid progress in the animal domain. As an example, benchmarks are very important in advancing the development of new methods and incentivizing the progress in a field of research, but for the most part, they don't yet exist for animal communication research, meaning researchers are working in silos without a way of comparing results. Recently, we worked with partners to create the first comprehensive benchmark in animal communication—the Benchmark of Animal Sounds (BEANS)—so that when anyone comes up with a new algorithm or model or AI, you can test to see if it's better. Without the field's ability to see itself, it doesn't actually improve.

We're also working on the field's first multimodal foundation models. These models are trained on large amounts of data, typically in a self-supervised manner, and can perform difficult predictive and generative tasks. We're really excited and proud to have just published the Animal Vocalization Encoder based on Self-Supervision (AVES), the first large foundation model developed in the animal communication space. This is a foundation model that uses self-supervised learning to greatly increase the efficiency of tasks such as detection and classification. And of course, various partners and other people are starting to reach out to use those models.

KZ: We rely heavily on the many partners we work with who are engaged in front-line research to help determine our projects. Many of our partners are looking to us to help develop tools that can solve basic challenges such as automated detection of signals from animals. For example, when you have hundreds of hours of recordings of beluga whales or orangutans, you might need to sit and listen to them just to know where the vocalizations are in the recordings. A machine learning model can detect vocalizations and go further to classify them into different types of vocalizations, which saves an immense amount of time on the part of researchers.

Machine learning can also help with with denoising of recordings and source separation, which is critical when there are many overlapping vocalizations in a recording. One of the biologists we're working with, Dr. Valeria Vergara, has indicated that upward of 90 percent of some of her recordings of belugas are unusable because they are so noisy, or because she can't separate out the vocalizations of different animals.

And, of course, it can be effectively applied to discover patterns in data that might not otherwise be seen. For example, ESP is developing models to predict an individual's motion conditioned on that individual's vocal behavior (and vice versa), potentially revealing when an animal's vocalizations are predictive of its behavior. This kind of understanding can help to quantify the impacts of human activities and climate change, and potentially help to warn us of things like stranding events with cetaceans before they occur.

LM: What do you say to people who dismiss technology in conservation efforts? How will the field of conservation change as AI advances?

KZ: I would push back on technology as an easy solution to any of the major global issues we face. Whatever issue you're dealing with, taking a systemic approach and collaborating to engage all key stakeholders will always be as important as a great technological advance. It's the same with conservation, where we see that support for the nontechnical, such as training for enforcement officers, or investment in community-based conservation, is a critical part of the equation.

Having said that, there is no question that technology has made a hugely positive contribution

PHOTO BY KERI WILK

to conservation over the past two decades, and often this is because the data gathered through technology can make a case for investment in the nontechnical. Advances in geospatial imagery, data storage, the decreasing cost of low-Earth orbit satellites, advancements in machine learning—all of these offer more tools for practitioners to gather data, analyze at scale, and make decisions more quickly, at a lower cost.

We are seeing huge amounts of data now being gathered through animal-borne tags. In fact, these tags have developed so rapidly that they can be deployed on species as small as bats, and can gather data on movement, vocalizations, environmental context, even emotions. AI is playing a critical role in analyzing that data, helping practitioners in the field with myriad tasks such as identifying species and individuals, monitoring populations, and understanding human impacts.

LM: What about the ethics surrounding AI?

KZ: We are very aware of the myriad concerns surrounding AI technology, from power imbalances and the environmental impacts of training large models to the fears of AI being actively used to drive misinformation and erode human rights. In the human domain, there are questions about intellectual property ownership and whether you have the right to know who's speaking to you—is it another human or has it been generated by a machine?

As we move toward the possibility of two-way communication with other species, using signals generated by machines, the same ethical concerns

arise. We see this already with playback experiments, a research framework used by scientists for decades to test their hypotheses about animal communication. However, the scale and pace of technological advancement we are seeing with machine learning will make it unlikely that the current ethical and regulatory frameworks will be sufficient to take account of these rapidly advancing technologies.

It's a Pandora's box situation, and we don't yet have clear solutions. But we do know that AI is here to stay and that it will be critical to put in place the right safeguards to govern its use. Over the coming year, we will be exploring with our partners what kinds of governance principles, or safeguards, need to be developed to ensure that the technology is used for the benefit of other species, and doesn't have unintended consequences.

LM: What kinds of advancements would you like to see in the field?

KZ: We want to see more data that's openly available to researchers. Access to data is often at the heart of technological progress. We saw how critical open data sharing was in the rapid development of COVID-19 vaccines, and the incredible advances we have seen in AI have been underpinned by access to great datasets and open sharing of code.

Unlike human communication where computers can trawl the internet for images, text, etc., it's time consuming and expensive to gather data in the field from other species. To support this, we should be rewarding the incredibly time-intensive work of data collection and annotation in the same way that we reward analysis and publication—so that researchers build on and recognize each other's datasets.

At ESP, we are actively promoting the principles of open access and data sharing in animal communication research. At the same time, we're aware of the potential risks and unintended outcomes of making datasets openly available and accessible, especially large ones.

LM: What do you see as being your legacy for future conservationists?

AR: The biggest revolutions in science rarely come from conceptual understanding. Many come from a new tool that opens up new data, like using a microscope for the discovery of germs. I think that's the paradigm shift that's about to happen here with the creation of these new tools. That's going to change our perspective and the way we interact with the rest of nature.

If you're to name the one fundamental problem, it's disconnection. We humans are disconnected from ourselves, from our own bodies, our emotions, our relationships, the world we live in. And we're going to continue to mutilate the tree of life until we cut the branch that we depend on. A shift in perspective and the ability to listen can change us not just at the individual level, but also at the collective level, the cultural and the institutional levels, and affect law and markets.

My big hope for ESP is that we can help, within a movement, to create a moment that superpowers everyone else. You can imagine, as we get a little further, that everyone who is working on rights for nonhumans, rights for nature, or creation for marine and terrestrial protected zones, how much their work could be boosted by this moment we can create. It's a "looking back at ourselves and seeing ourselves from space" moment.

LM: Aza, a lesson you've taken from your father is "Our ability to understand is limited by our ability to perceive." Why is this a message you want the public to know?

AR: Because we are human, we are blind to so much of the natural world. If we could see at the speed at which plants grow, we would see how they bend toward the light, how the roots of pea plants seek the sound of running water. We would be aware of so much more. That's what I mean by "Our ability to understand is limited by our ability to perceive." If we could open up the aperture of imagination, of human perception, then suddenly we'd understand a bunch more. And if we would understand differently, we would act differently. We would see the beauty around us and we would change the way we behave. I don't think the way for human beings to change is through shame or guilt, but through wonder and awe. **LM**

PHOTO BY KERI WILK