



Earth
Species
Project



2025 Annual Report



Contents

- 3 Letter From Our Co-Founder
- 4 Welcome to Our New CEO
- 5 Research to Decode
- 16 Building the Field
- 19 Preparing Society for Impact
- 24 Looking Ahead
- 26 Change the World With Us

Dear Supporters,

2025 was an exciting year for Earth Species Project (ESP) and for the accelerating field of **Animal Language Processing**. Advances in AI are making it possible to study animal vocalizations and behavior at unprecedented scale, revealing patterns that were previously impossible to detect. Last year came with some important firsts for the field:

- We open-sourced **NatureLM-audio**, the first large audio language model designed to analyze animal sounds.
- We published **14 new papers**, with work accepted at leading international conferences including ICLR and NeurIPS.
- We partnered with the **Raincoast Conservation Foundation** to study the function of killer whale vocalizations in an increasingly noisy ocean.
- We co-organized the first **AI for Non-Human Animal Communication** workshop at NeurIPS — an important milestone for the field at the leading machine learning conference.
- We partnered with **The Collective Intelligence Project** to understand global attitudes on AI and animal communication, and **found strong public support** for animals having a say in democratic decision-making.
- We partnered with **Earth Law Center** to map the legal, political, and cultural implications of interspecies understanding.

- We grew our team by adding eight new full-time employees and brought on research fellows and interns to accelerate work in the field.
- We welcomed big names to our board, including **Christiana Figueres** and **Thomas Zurbuchen** to scale our science and impact.
- We were covered in 16 outlets including **Nature**, the **BBC**, and **CNN**.

What once felt like science fiction is quickly becoming reality. As we look ahead, we remain committed to breaking the interspecies understanding barrier and to building a future where that deeper understanding benefits all life on Earth.

None of this would be possible without our partners, supporters, and community. Thank you for believing in this work and helping bring it to life.



Aza, Co-Founder and Chief Explorer

Guiding the Next Phase of Interspecies Understanding

In early 2026, Earth Species Project welcomed Steven VanRoekel as our new Chief Executive Officer. Steve brings decades of leadership experience across technology, government, and mission-driven organizations, including serving as Chief Information Officer under the Obama administration and EVP for Innovation of The Rockefeller Foundation.

Throughout his career, Steve has been drawn to big, complex challenges where technology can help reshape systems and expand what's possible. That same instinct led him to ESP, where his professional experience meets a long-standing personal curiosity. Time spent outdoors early in life, along with time spent traveling and observing wildlife in Africa, sparked a sense of wonder that left a lasting impact on him.

Get to know Steve in [this Q&A](#) and hear from him directly in this [short video message](#).

[Read More →](#)



*“What excites me most about ESP is the chance to help more people experience the **profound perspective shift** I felt after spending time in nature and observing the **richness of animal communication**. If we can begin to listen more deeply to other species, we have an opportunity to **reshape our relationship** with the rest of nature.”*

Research to Decode



Overview

We conduct deep dive experiments to understand the communication systems and diverse intelligences of species across the Tree of Life. In partnership with leading biologists around the world, we're studying killer whales, carrion crows, zebra finches, and more.

In 2025, we grew our team by 6 full-time employees and extended our research program by supporting three Post-doc researchers, four interns, and one apprentice, allowing us to bring specialized expertise to different projects and accelerate our research pipeline significantly.




Fig. The “virtuous cycle” feedback loop designed to improve our general models over time.


We focused on two core areas: 1) developing foundation models for animal communication, such as NatureLM-audio, and 2) partnering with biologists to point these tools at specific Decode projects. These projects generate new datasets and insights, which improve our general models over time

Our research was recognized and accepted at major international conferences:

- [**NatureLM-audio**](#) was accepted at ICLR 2025
- [**Bidenoising: Animal Vocalization Denoising without Access to Clean Data**](#) paper was accepted at ICASSP 2025
- One paper accepted at a [**workshop**](#) and one main conference [**acceptance**](#) at NeurIPS, 2025
- Presented [**What Matters for Cross-Taxonomic Bioacoustic Models**](#) at Living Data 2025
- [**Crossing the Species Divide: Transfer Learning from Speech to Animal Sounds**](#) was accepted at DCASE 2025
- [**AVEX: What Matters for Animal Vocalization Encoding**](#) has been accepted for ICLR 2026

Decode Projects

To understand how communication works across the animal kingdom, our research spans a diverse set of species. Each one offers a different window into the structure, function, and evolution of communication. By studying systems that vary in ecology, cognition, and social organization, we can identify patterns that generalize across the Tree of Life as well as patterns that do not.



Some of the species ESP tools
were used to study in 2025

Decode: **Killer Whales**

Earth Species Project partnered with **Raincoast Conservation Foundation** (RCF) to launch a **new research initiative** exploring how killer whales communicate in an increasingly noisy ocean. By combining Raincoast's decades of field expertise with our state-of-the-art AI tools, the collaboration aims to better understand how vocal communication supports the complex social lives of killer whales and how rising underwater noise may be affecting them.

In September, our teams set out into the waters of British Columbia to conduct a pilot study, collecting synchronized data from drone footage, underwater acoustic recordings, and time-stamped behavioral observations. This multimodal approach links vocalizations to individual whales, social behaviors, and environmental conditions, creating a uniquely detailed dataset for studying how killer whales coordinate movement, forage together, and maintain the social bonds that hold their societies together. This data may also help us understand how increasing ocean noise may be disrupting these vital conversations.



Video: In 2025, we conducted a successful pilot season with the Raincoast Conservation Foundation. It is illegal to fly drones over marine mammals if not under permit. These video clips were taken under SARA permit XMMS-2-2022-Amendment 2 (at a minimum 30m altitude).

Decode: Killer Whales



Video: Synchronized drone footage and underwater acoustic recordings from the Raincoast Conservation Foundation. It is illegal to fly drones over marine mammals if not under permit. These video clips were taken under SARA permit XMMS-2-2022-Amendment 2 (at a minimum 30m altitude).

RCF helped us prototype a new ESP tool which automatically transcribes voice notes recorded during field observations and uses its predictions to populate a given ethogram. This has led to a significant reduction in manual transcription and data entry, and also opens the door for researchers to revisit and analyze years of archived voice recordings in a scalable way. We have already seen early validation of this analysis pipeline.

In 2026, the partnership is moving beyond the pilot phase to expand this work across multiple killer whale populations. Future efforts will connect findings from British Columbia with research in other regions, helping establish shared methods and datasets to understand the communication and culture of these remarkable animals.

[Read More →](#)

Decode: Carrion Crows

Using our machine learning models, our partners Prof. Vittorio Baglione and Prof. Daniela Canestrari at the University of León discovered that **wild crows** often vocalize quietly, suggesting they engage in extensive private, close-range conversations. These kinds of vocalizations have been difficult to measure with traditional microphones. By using mini-biologgers that capture close-range, low volume sounds, our collaborators revealed that the iconic loud “caw” we usually hear only scratches the surface of the crows’ communication system.



Image credit: University of León

Decode: Carrion Crows


Detection

Voxaboxen treats each vocal element as its own event, giving a more detailed description of how vocalizations are sequenced in time.



Classification

Detected vocalizations are predicted on the following labels:



In this [paper](#) published in *Animal Cognition*, Vittorio and Daniela highlight our joint work using ESP's detection model, **Voxaboxen**, to detect and classify more than 127,000 vocalizations from adults as well as many more from their chicks and even parasitic cuckoo nest-mates. This scale of analysis wouldn't have been possible using standard ethological methods, demonstrating how machine learning can unlock new levels of insight into animal communication.

They also found that the mini-biologgers had no statistically-significant impact on crow welfare, key for both ethical and scientific validity. The study sets the stage for future work on exploring what variety might exist in these quiet vocalizations and decoding the full complexity of crow language.

[Read the Paper →](#)

Fig 1. A diagram of Voxaboxen, a machine learning tool for research in bioacoustics.

Decode: Zebra Finch

Throughout 2025, we worked with McGill University to investigate how female zebra finches modulate their calls during natural social exchanges. As part of the project, we developed ZF-AIM, a generative audio language model designed to engage in real-time, two-way vocal interactions with live birds. We analyzed more than 1.5 million vocalizations from female zebra finches, an unprecedented scale of analysis that would not be possible using traditional methods.

Most playback experiments only last around 10 seconds. With our models, we were able to analyze back and forth interactions for 48 hours. This enabled us to study fine-grained patterns of turn-taking and vocal adaptation that are difficult to capture with traditional playback experiments.

This approach opens new possibilities for complex and controlled playback experiments. The [preprint](#), released in February 2026, shares the first findings from this work.

[View the Preprint →](#)




Image credit: Quang Nguyen Vinh on Pexels




Fig 2. A diagram of ZF-AIM, a generative audio language model designed to engage in real-time, two-way vocal interactions with live birds.

Foundations

NatureLM-audio: A New Class of Audio Language Models

One of the year's most significant milestones was the continued development of our flagship foundation model NatureLM-audio, the first large audio language model designed specifically for analyzing animal sounds.

Testing showed strong generalization across diverse acoustic domains, including the ability to distinguish animal sounds from human speech. In 2025, NatureLM-audio was also accepted to the 2025 International Conference on Learning Representations (ICLR), signalling growing recognition of using AI to study animal communication.

We open-sourced the model, code, and training datasets on [GitHub](#) and [Hugging Face](#). Our interactive [demo](#) lets users upload animal sound recordings (or choose from our Sample Library) and ask questions in plain English. Together, these resources make it even easier for researchers worldwide to experiment with a large-scale bioacoustic audio language model.

Thank you to the Patrick J McGovern Foundation for their support for the development and scaling of NatureLM-audio.

[View the Project Page →](#)




Fig 3. A diagram showing how NatureLM-audio works. It combines a fine-tuned audio encoder (BEATs) with a large language model ([Llama 3.1 8B Instruct](#)), enabling researchers to query bioacoustics data using natural language.

Foundations

Case Study: Listening for Frogs at Scale With NatureLM-Audio

FrogID is the largest frog-focused citizen science project in the world, led by the Australian Museum. They evaluated NatureLM-audio on five core identification tasks and found strong performance even without fine-tuning.

Despite not being trained explicitly on Australian frogs, NatureLM-audio had:

- **99%** accuracy on classifying frogs from non-frogs
- **82%** accuracy identifying focal species among FrogID's top 5 most recorded frogs
- **Strong** generalization to related tasks like detecting human speech

This highlights NatureLM's ability to generalize across tasks it was not explicitly trained for, and its flexibility points to a broader potential across new bioacoustics tasks, species, and conservation efforts. Read the full case study on [our website](#).

[Learn More →](#)




Image: FrogID Validator. Photo by Nadiyah Roslan/Australian Museum.



Foundations

What Matters for Bioacoustic Encoding

Alongside model development, a major focus in 2025 was understanding how to train bioacoustic foundation models that generalize across species and tasks.

AVEX: What Matters for Animal Vocalization

Encoding, we conducted a systematic study of model architectures, training data composition, and training strategies to identify what drives performance in real-world settings. We found that adding human generated sounds (music, speech) to the training data significantly improved the performance and especially when it comes to transfer to new species and tasks.

This work led to a practical training framework for building more robust and transferable audio encoders that can support a wide range of species and ecological contexts.




Fig 4. Our empirical study diagram, assessing (1) models, (2) training data, (3) training paradigms, and proposing an (4) extended evaluation data and methodology.

[View the Project Page →](#)

Building the Field



Building the Field

AI and Animal Communication at NeurIPS

Solving a challenge as ambitious as interspecies understanding requires an interdisciplinary approach and thriving community of researchers, institutions, and partners to advance the field of animal language processing.

In 2025, ESP co-organized the [AI for Non-Human Animal Communication Workshop](#) at NeurIPS in San Diego, alongside representatives from Google DeepMind, ENES Lab Saint-Etienne, Naturalis Biodiversity Center. The workshop marked the first time this topic was formally introduced at a major AI conference.

It brought together researchers from AI, bioacoustics, ethology and beyond to share emerging work and explore how advances in machine learning can accelerate the study of animal communication.

[Learn More →](#)



Image: Earth Species Project at the AI for Non-Human Animal Communications Workshop, co-organized with Google DeepMind, Naturalis Biodiversity Center, and ENES Lab Saint-Etienne.

By creating space for interdisciplinary exchange, the workshop helped lay the groundwork for a more connected research community and a shared agenda for advancing this rapidly developing field. We are featuring many of the researchers who presented work at NeurIPS in our monthly [Community Spotlight](#) blog series.

Building the Field

BioDCASE Challenge: Advancing Methods for Bioacoustic Research

Our team contributed to the 10th [DCASE Workshop](#) in Barcelona, presenting the following work:



[Robust Detection of Overlapping Bioacoustic Sound Events](#)



[Synthetic Data Enables Context-aware Bioacoustic Sound Event Detection](#)



[Crossing the Species Divide: Transfer Learning from Speech to Animal Sounds](#)

We also co-organized a challenge for BioDCASE focused on [Multi-Channel Alignment](#), which tackles the difficult problem of resynchronizing and aligning audio from multiple recording devices. The challenge drew strong participation from the research community. You can see the winners on the [BioDCASE GitHub page](#).

LINGUAE-AL: Animal Linguistics Seminar

In 2025, Earth Species Project co-organized the [LINGUAE-animals](#) Seminar with Phillippe Schlenker, Lucie Pavaux, [Institut Jean-Nicod](#), and [LSCP](#).

This was an interdisciplinary forum bringing together researchers working at the intersection of linguistics, cognitive science, bioacoustics and, of course, ethology. The seminar series explores how concepts and methods from human language research can inform the study of communication across species.

The seminar fosters new discussions and collaborations around the structure, meaning, and evolution of animal communication. These conversations are an important step toward building a shared scientific framework for studying communication systems across the Tree of Life.

Preparing Society for Impact



Preparing Society for Impact

Our work is not just about science, it's about reimagining our place on Earth and inspiring new systems that will bring us back into balance with the rest of nature.

Alongside our core research, we work to bring scientists, policymakers, and the public into conversation about what interspecies understanding could mean for society, and how those insights might shape future relationships between humans and other species.

In 2025, we launched two critical initiatives to help drive this work: a [global survey](#) to establish a baseline on public sentiment around AI-enabled interspecies understanding, and a landscape study of the legal, political, and cultural implications of our work, in order to help refine our theory of change and to identify new partners and initiatives to engage with.



*“Mutually beneficial change in our relationship with nonhuman animals requires **meaningful public engagement**. This foundational report exemplifies the importance and richness of **connecting the public** to AI-based efforts to understand and respond to animals, ensuring that **ethical frameworks** and **inclusive dialogue** guide technological advances.”*

–Melanie Challenger & Erin Ryan, Co-Founders of [Animals in the Room \(AiR\)](#)


Preparing Society for Impact

7 in 10 people




want to know what animals are **thinking** and **feeling**

Many respondents shared **enthusiasm** about the application of AI in animal communication




There was **strong support** for strict rules on how companies **profiting** from animals can communicate with them




94%

of respondents believe animals have **emotions**




91%

of respondents believe animals have their own **languages**



62%

of respondents believe animals have their own form of **culture**



Most respondents agree that if animals had a **voice** and **representation**, they **should** be allowed to participate in **human democratic processes**



of respondents identified **conservation** and **species protection** as the top application of AI for interspecies understanding

47%

~78%

point gaps between countries on some statements, demonstrating significant **cultural divergence** exists

What the World Thinks About AI and Animal Communication

We partnered with **The Collective Intelligence Project** to conduct the first **global survey** on AI-powered interspecies understanding, capturing perspectives from more than 1,000 people across 67 countries.

We found:

- **Global Consensus:** that animals have complex inner lives, including emotions (94%), language (91%), and culture (63%)
- **Collective Readiness:** 70% of respondents want to know what animals think and feel
- **Democratic Revolution:** 60% are open to animals having some form of participation in human decision-making
- **Proactive Governance:** 85% believe strict rules should be in place before this technology is widely deployed

The findings were presented at The MOTH Festival of Ideas 2026, bringing this conversation to a broader global audience.

Preparing Society for Impact

Mapping the Legal, Political, and Cultural Implications of Our Work

In partnership with [Earth Law Center](#), we conducted a landscape study exploring the potential of ESP technologies across political, legal, ethical, and cultural domains. The findings have helped us identify pathways, policies, practices, and cultural moments that could drive broader systemic change and reimagine legal and political participation across species lines.

Ongoing Thought Leadership

Advances in Animal Language Processing must be accompanied by thoughtful public engagement, ethical frameworks, and inclusive global dialogue. In 2025, we were involved in a number of thought leadership initiatives to advance the narrative.

- ESP's Director of Strategic Impact, Jane Lawton, joined a [panel](#) at the launch of the London School of Economics' Collier Centre for Animal Sentience, discussing how emerging technologies may reshape how we think about animals and their interests.
- We submitted a written piece to the Berggruen Institute's [Planetary Compendium](#) in partnership with Raincoast Conservation Foundation and Earth Law Center. This online multimedia collection showcases how people and institutions are addressing



Image: Jane Lawton, ESP's Director of Strategic Impact, was part of a panel launching the new Centre for Animal Sentience at London School of Economics.

planetary-scale challenges, and is set to have a large impact and reach.

- We continued collaborating with partners advancing the rights of nature movement, including co-hosting the Ecocentric Law Beach Party with [Earth Law Center](#) at the UN Ocean Conference and supporting the opening of their London office.

Preparing Society for Impact



Image credit: Cavan Images / Adobe Stock

Transforming Our Relationship With the Rest of Nature

Earth Species Project is entering a new chapter, one that not only pushes the boundaries of AI and animal behavior research, but also readies humanity to meet this moment with care, awe, and accountability.

As breakthroughs in AI accelerate our ability to listen more deeply to other species, they also give us the opportunity to rethink humanity's relationship with the natural world. Preparing society for this moment means building the conversations, governance models, and partnerships that will help guide these technologies toward outcomes that will benefit all forms of life.

Shifting Mindsets

Our work continues to capture the awe and wonder of global audiences. See our work in leading publications:



Looking Ahead

As the field continues to accelerate, our focus in the year ahead is to continue our incredible progress and partnerships in translating rapid advances in AI into real understanding across species—building the scientific, ethical, and organizational foundations needed to turn possibility into lasting impact.

We've had several papers accepted at major international conferences in 2026, including [AVEX: What Matters for Animal Vocalization Encoding](#) at ICLR and [Multi-layer attentive probing improves transfer of audio representations for bioacoustics](#) at ICASSP. We will also host a joint workshop with Raincoast Conservation Foundation at the [26th Biennial Conference on the Biology of Marine Mammals](#) and have exciting new papers currently under review.

Alongside this, we are continuing to evolve our core tools to improve performance, usability, and application across a wider range of species and tasks. In early 2026, we open-sourced AVEX, a new Python package for training and evaluating bioacoustics representation learning models. Access the [repository](#), [documentation](#), and [models](#) to get started. We're also making improvements to our flagship model, NatureLM-audio.



Image credit: Blueringmedia / Adobe Stock

Looking Ahead



Image: Steve, ESP's CEO (left) and Aza, ESP's Co-Founder (right) on stage at SXSW 2026.

We are working closely with researchers across AI, machine learning, ethology, and bioacoustics. Growing datasets and advances in machine learning have led to a shift from species-specific studies to more data-driven, cross-species, and interdisciplinary approaches. We refer to this convergence as [**Animal Language Processing**](#) (ALP). ALP is a working definition and we invite the community to critique, revise, and improve this framework together. We'd love [**your feedback!**](#)

In early 2026, ESP's Co-Founder Aza Raskin delivered a [**keynote at SXSW**](#), bringing the vision of interspecies understanding to a global audience. ESP's Director of Strategic Impact Jane Lawton also [**spoke at ChangeNOW**](#) in Paris on a panel exploring how to make nature protection more compelling and actionable. We also presented our joint research with the Collective Intelligence Project at the [**MOTH Festival of Ideas**](#).

We are also exploring new experimental pilot programs and partnerships that connect scientific discovery with clear pathways to impact, aiming to create a shift in how we relate to the rest of life on Earth.

Change the World With Us

Breakthrough science is only possible with bold supporters. Our donors make it possible for ESP to pursue one of the most ambitious scientific challenges of our time: decoding the communication of other species. Their belief in this work helps power new discoveries, new tools, and a growing global movement to better understand life on Earth.

We are deeply grateful to every donor who has helped make this work possible—and we invite others who are inspired by this vision to join us.

The supporters listed here recognize cumulative gifts of \$10,000 or more over the past five years.

[Support Our Work →](#)

\$5M+

Reid Hoffman
Waverley Street Foundation

\$1M – \$5M

Steve & Genevieve Jurvetson
Chris Larsen
Allen Family Philanthropies

\$500K – \$999K

Google Cloud
Britt Selvitelle

\$100K – \$499K

Alex Asseily
Builders Vision
Ezra & Marisa Callahan
Filecoin Foundation for the
Decentralized Web
Kate James
Brewster Kahle, Internet Archive
Mike Krieger
Kumavis
Adam & Melony Lewis
Mae Philanthropies
Matt Mullenweg

National Geographic Society
Patrick J McGovern Foundation
Courtney Roberts & Stephen Preston
Justin Rosenstein
Annika Sten Pärson, The Inner Foundation
The Payne Brodeur Giving Fund

\$10K – \$99K

Ajaz Ahmed
Amr & Shirin Charitable Giving Fund
Craig Clemens
Mike Cotton
Bill Duane
Conor Farese
Simone Giacomelli
Tom Gruber
Timothy Herzik & Garry Nolan
Andy Lam
Bethanie Maples
Will Marshall
Dave & Brit Morin
Optas Capital
Shiva Rajaraman
Milo Runkle
Schrock Family Philanthropy
Evan & Christina Sharp
Dick Simon
Halle Tecco & Jeff Hammerbacher
Telluray Foundation
Steve Vassallo

...and a number of donors who prefer to remain anonymous.

“For most of human history, we have assumed that intelligence and language belonged to us alone. The possibility that AI might let us listen, for the first time, to the communication systems of other species is a profound one. If that promise is realized, it would mark not only a scientific advance but a shift in how we understand our own place in the natural world.”

–**Patrick McGovern**, trustee of the Patrick J. McGovern Foundation



ESP is a 501c3 non-profit organization and relies on philanthropic support to make our work possible.
Please reach out to ESP Director of Impact, Jane Lawton, to discuss how you can be part of this incredible journey.
jane@earthspecies.org

