Bridges-2 Webinar

Scaling Up Ecological Monitoring with AI: How Supercomputing Is Unlocking the Value of Autonomous Acoustic Sensing

Sam Lapp
University of Pittsburgh

August 23, 2024

Pittsburgh Supercomputing Center

enabling discovery since 1986

The **Pittsburgh Supercomputing Center (PSC)** provides advanced research computing capability, education, and expertise to the national research community.

Since 1986, PSC has provided university, government, and industry researchers with access to some of the most powerful systems available for high-performance computing, enabling discovery across all fields of science.

OUR AREAS OF EXPERTISE

- high-performance and data-intensive computing
- data management technologies
- software architecture, implementation, and optimization
- enabling ground-breaking science, computer science, and engineering
- user support for all phases of research and education
- STEM outreach in data science, bioinformatics, and coding

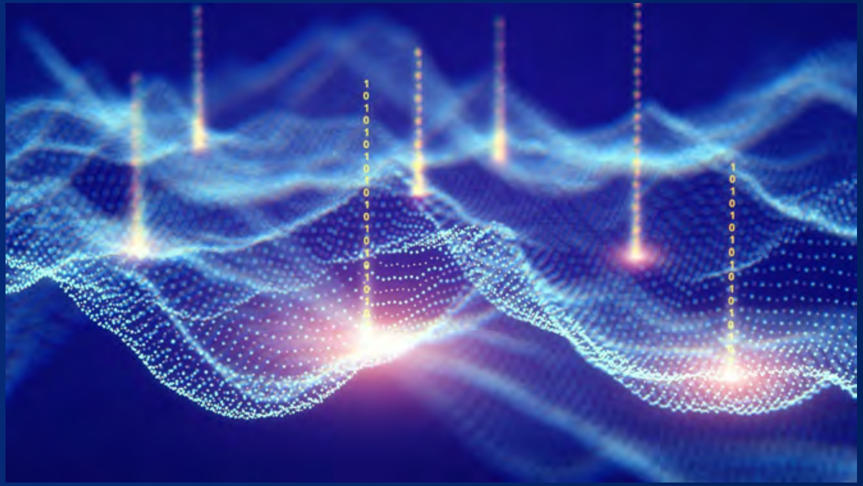

Welcome!

Hewlett Packard Enterprise is delivering *Bridges-2*

Bridges-2 Leadership Team

Sergiu Sanielevici
PI & Dir. Support for Sci. Apps.

Robin Scibek
Dir. Comms.
co-PI

Paola Buitrago
Dir. AI & Big Data
co-PI

Edward Hanna
Dir. Systems & Ops.
co-PI

Tom Maiden
User Services Mgr.
co-PI

Stephen Deems
Project Manager

Andrew Adams
Information Security Officer

Bridges-2 Webinars

- A forum for the Bridges-2 community to learn and share ideas and achievements: Bridges-2 Webinar series | PSC
- Topics and speakers of interest to work that is being done, or that may be done in future.
- Please suggest future speakers (including from your own team) and/or topics (including your own)!

Just email: sergiu@psc.edu

Introducing today’s presenter: Sam Lapp

Sam Lapp is a machine learning and sound geek with a passion for biodiversity conservation. His current research focuses on developing and applying AI methods in bioacoustics, the strategy of studying ecology through the sounds produced by living things. His projects include developing the open-source Python package *OpenSoundscape* and applying machine learning methods to the conservation of birds, frogs, and insects. Sam studied engineering, music, and sound design at Penn State and is currently working on his Ph.D. in the Kitzes Lab at the University of Pittsburgh.

Q&A Logistics

- **We abide by https://support.access-ci.org/code-of-conduct**
- All of us except Sam will be muted during his presentation.
- Please type your questions into the Zoom chat.
- We may be able to address some questions in the chat while Sam is presenting.
- When Sam finishes his presentation, he will answer questions live during the final ~10 minutes of this webinar.

peeps, pops, and teraflops

discovering hidden patterns in nature by applying supercomputing to audio recordings

Sam Lapp – Kitzes Lab @ University of Pittsburgh

sam.lapp@pitt.edu

Pittsburgh Supercomputing Center Bridges2 Seminar Series, August 2024

Sound carries rich information

Song is a fundamental aspect of bird behavior

Acoustic sensors: systematically sample biological sounds in the wild

- Large spatiotemporal scales
- Remote areas
- Hard-to-survey species
- Acoustic “specimens”


RSH-33
94V-0
PROG
CUSTOM
USB/OFF
DEFAULT
AudioMoth
1.1.0
0500


Acoustic
Monitoring
Machine
Learning

How can supercomputing crack the code?

Three stories for today:

1. intelligent sound recognition
   powered by GPU Nodes

2. probing the unknown with one-shot acoustic querying
   enabled by Extreme Memory nodes

3. high-resolution spatial tracking
   with super-parallelized CPU

Part 1: It takes a landscape to raise an owl

GPU nodes fuel automated recognition of rare wildlife sounds

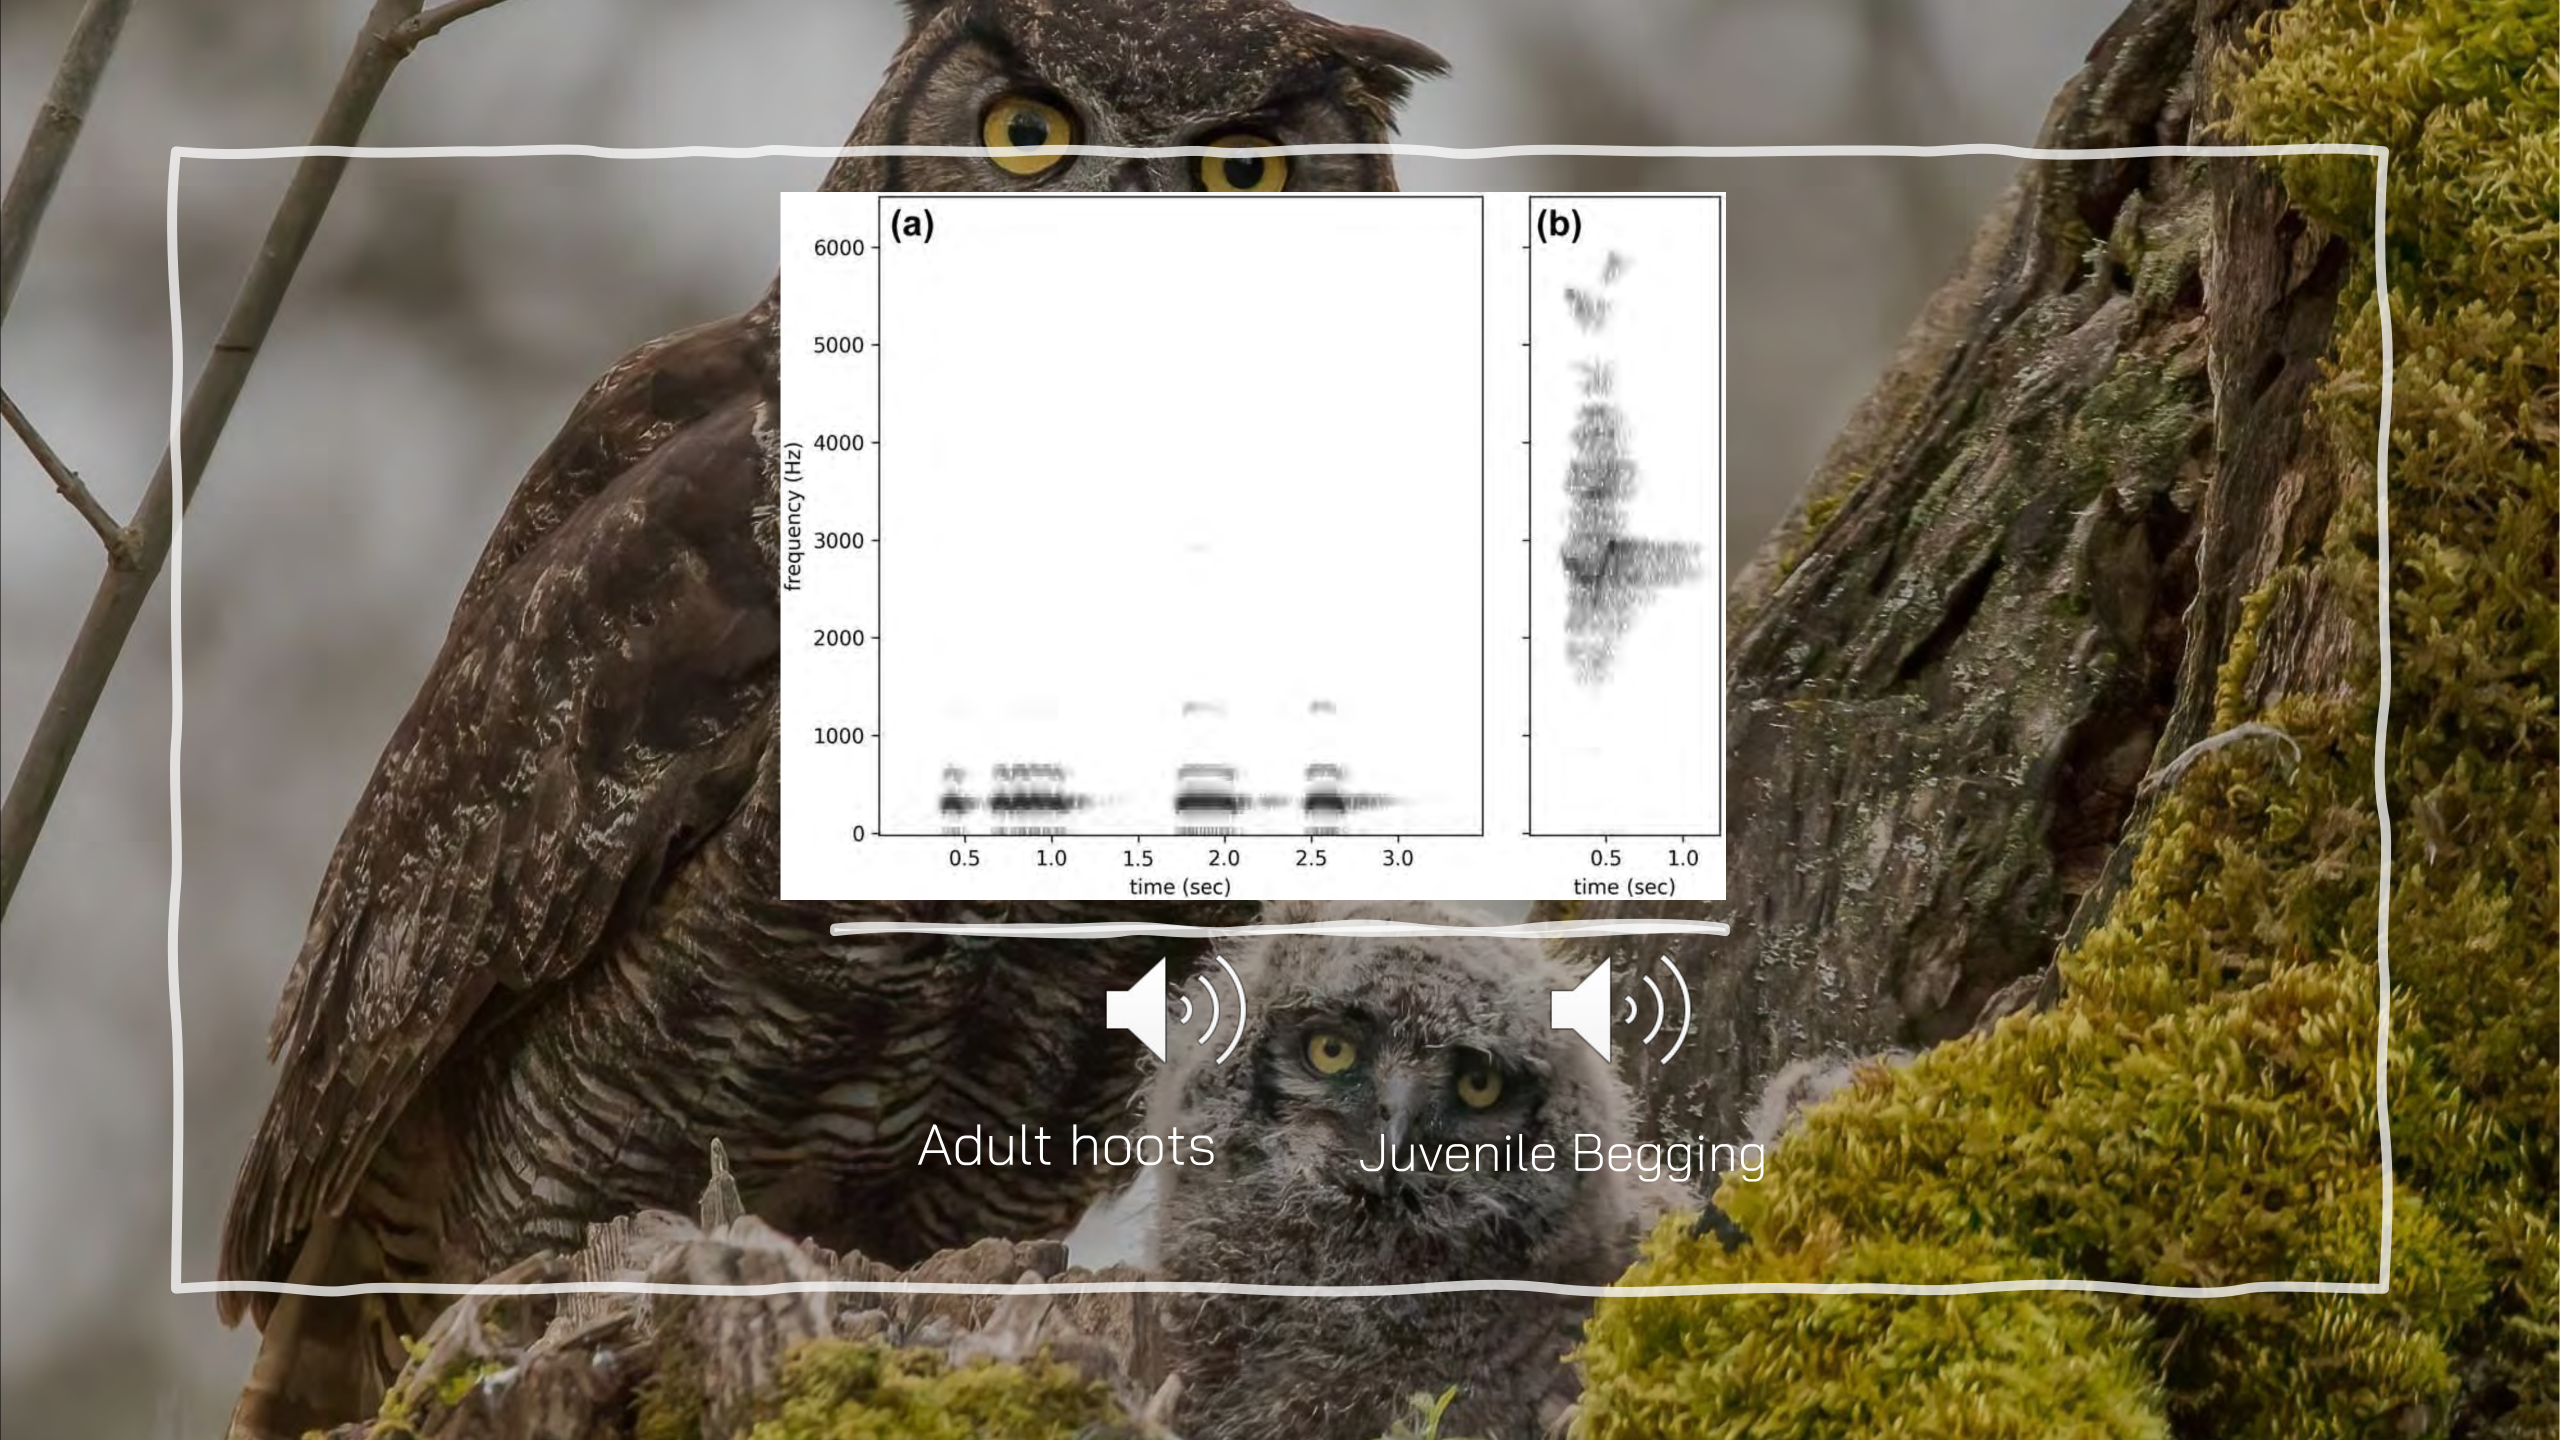

(a)
(b)
frequency (Hz)
6000
5000
4000
3000
2000
1000
0
0.5
1.0
1.5
2.0
2.5
3.0
time (sec)
0.5
1.0
time (sec)
Adult hoots
Juvenile Begging

Through ecological succession, habitats mature from fields to old growth forests

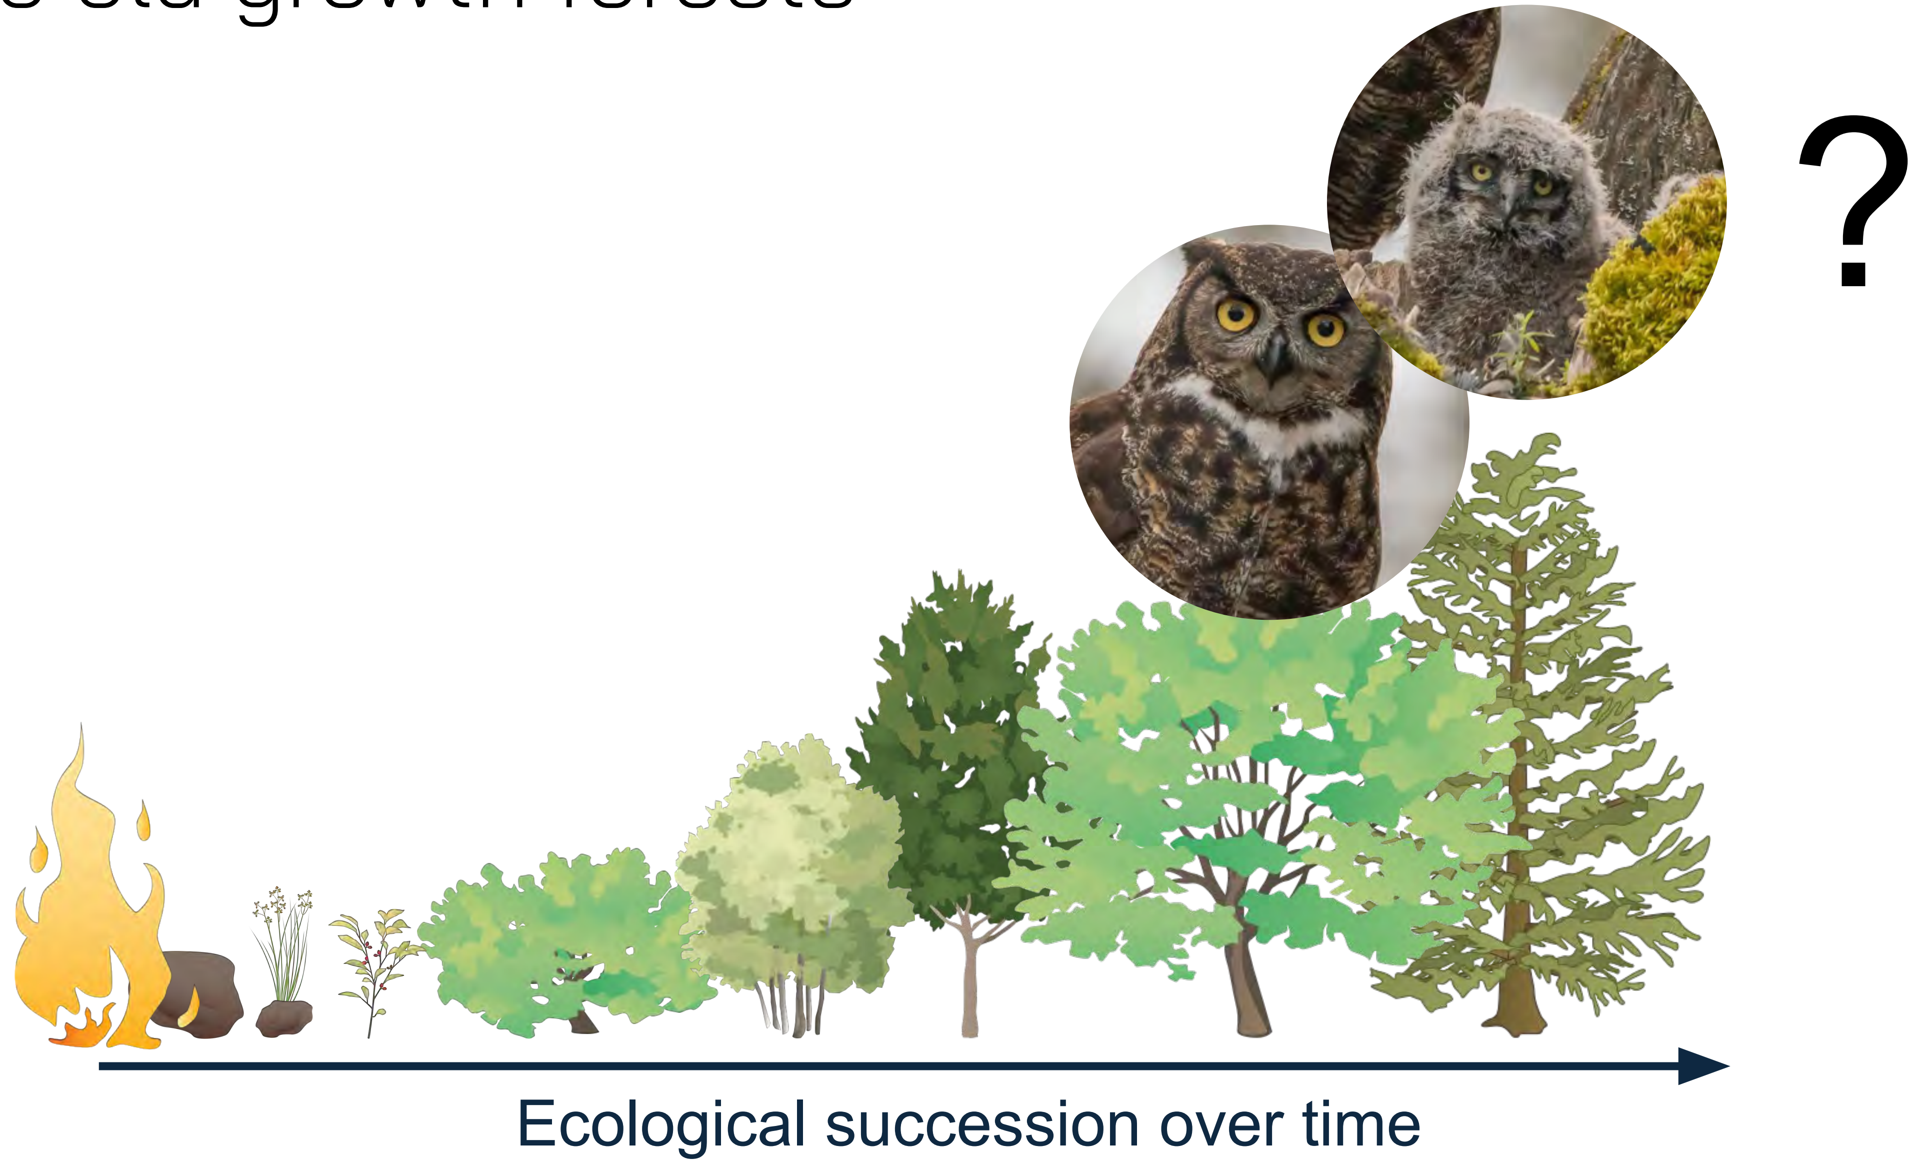

data collection by JT Larkin

Season-long acoustic monitoring across the Northeast

- recorded over 100,000 hours of audio
- covered diverse landscapes
- produced 20 Terabytes of data

If Lauren had tried to listen to it all, she'd have spent

133 years

listening for 40 hrs/week

a spectrogram represents sound as an image

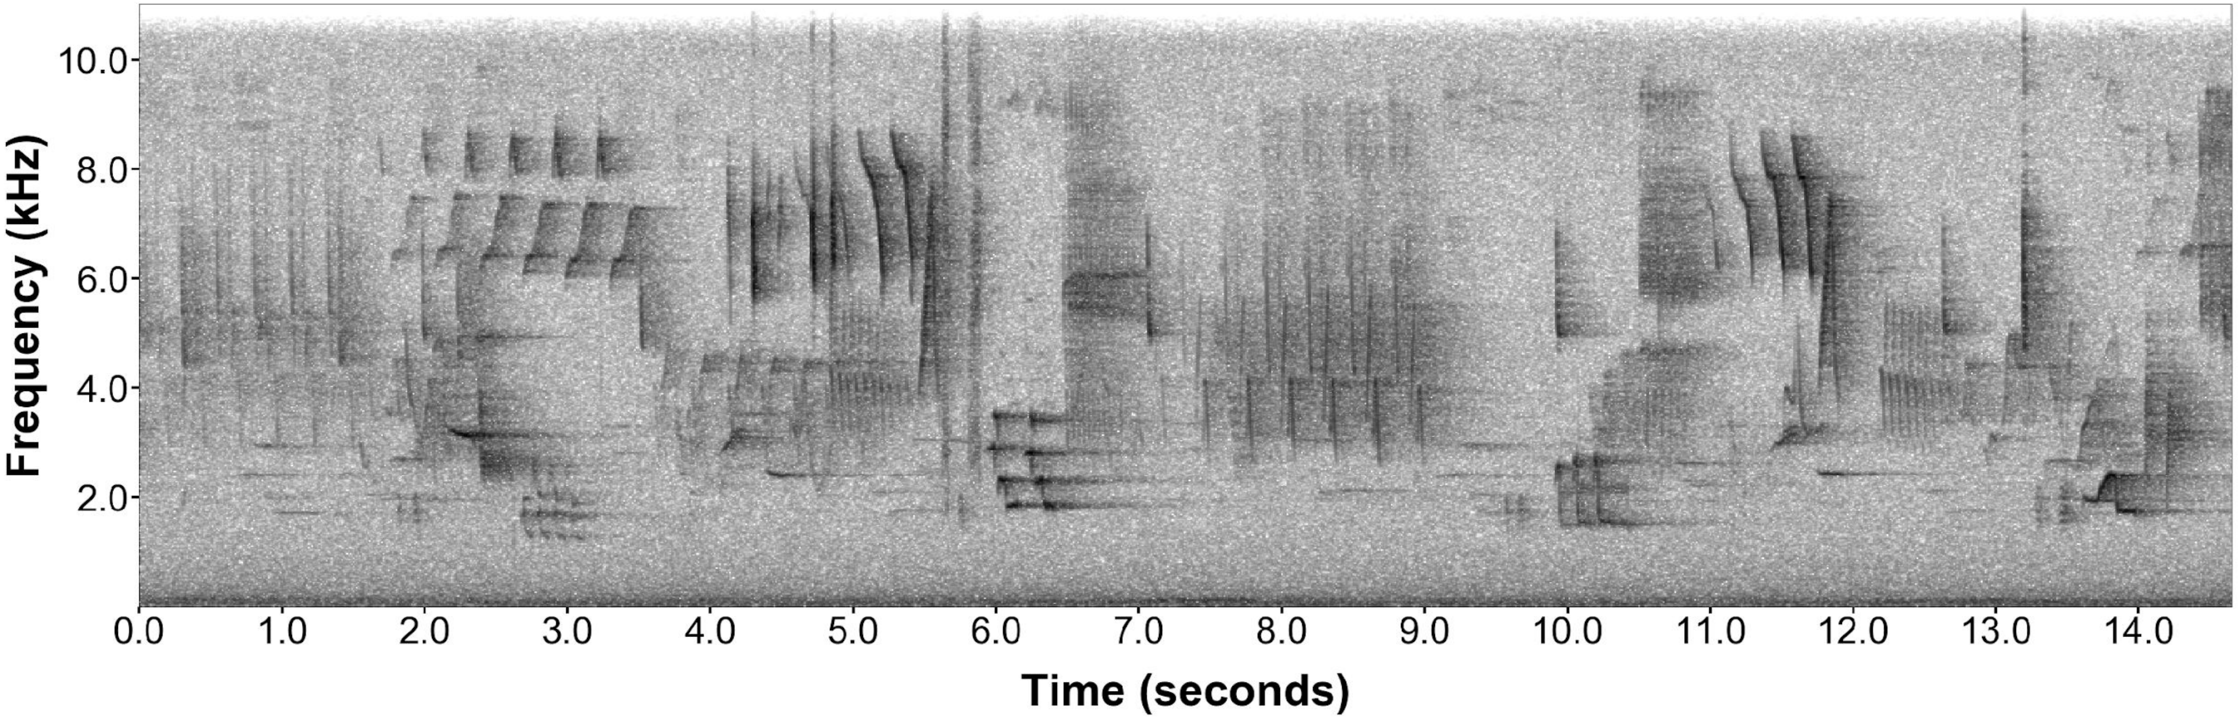

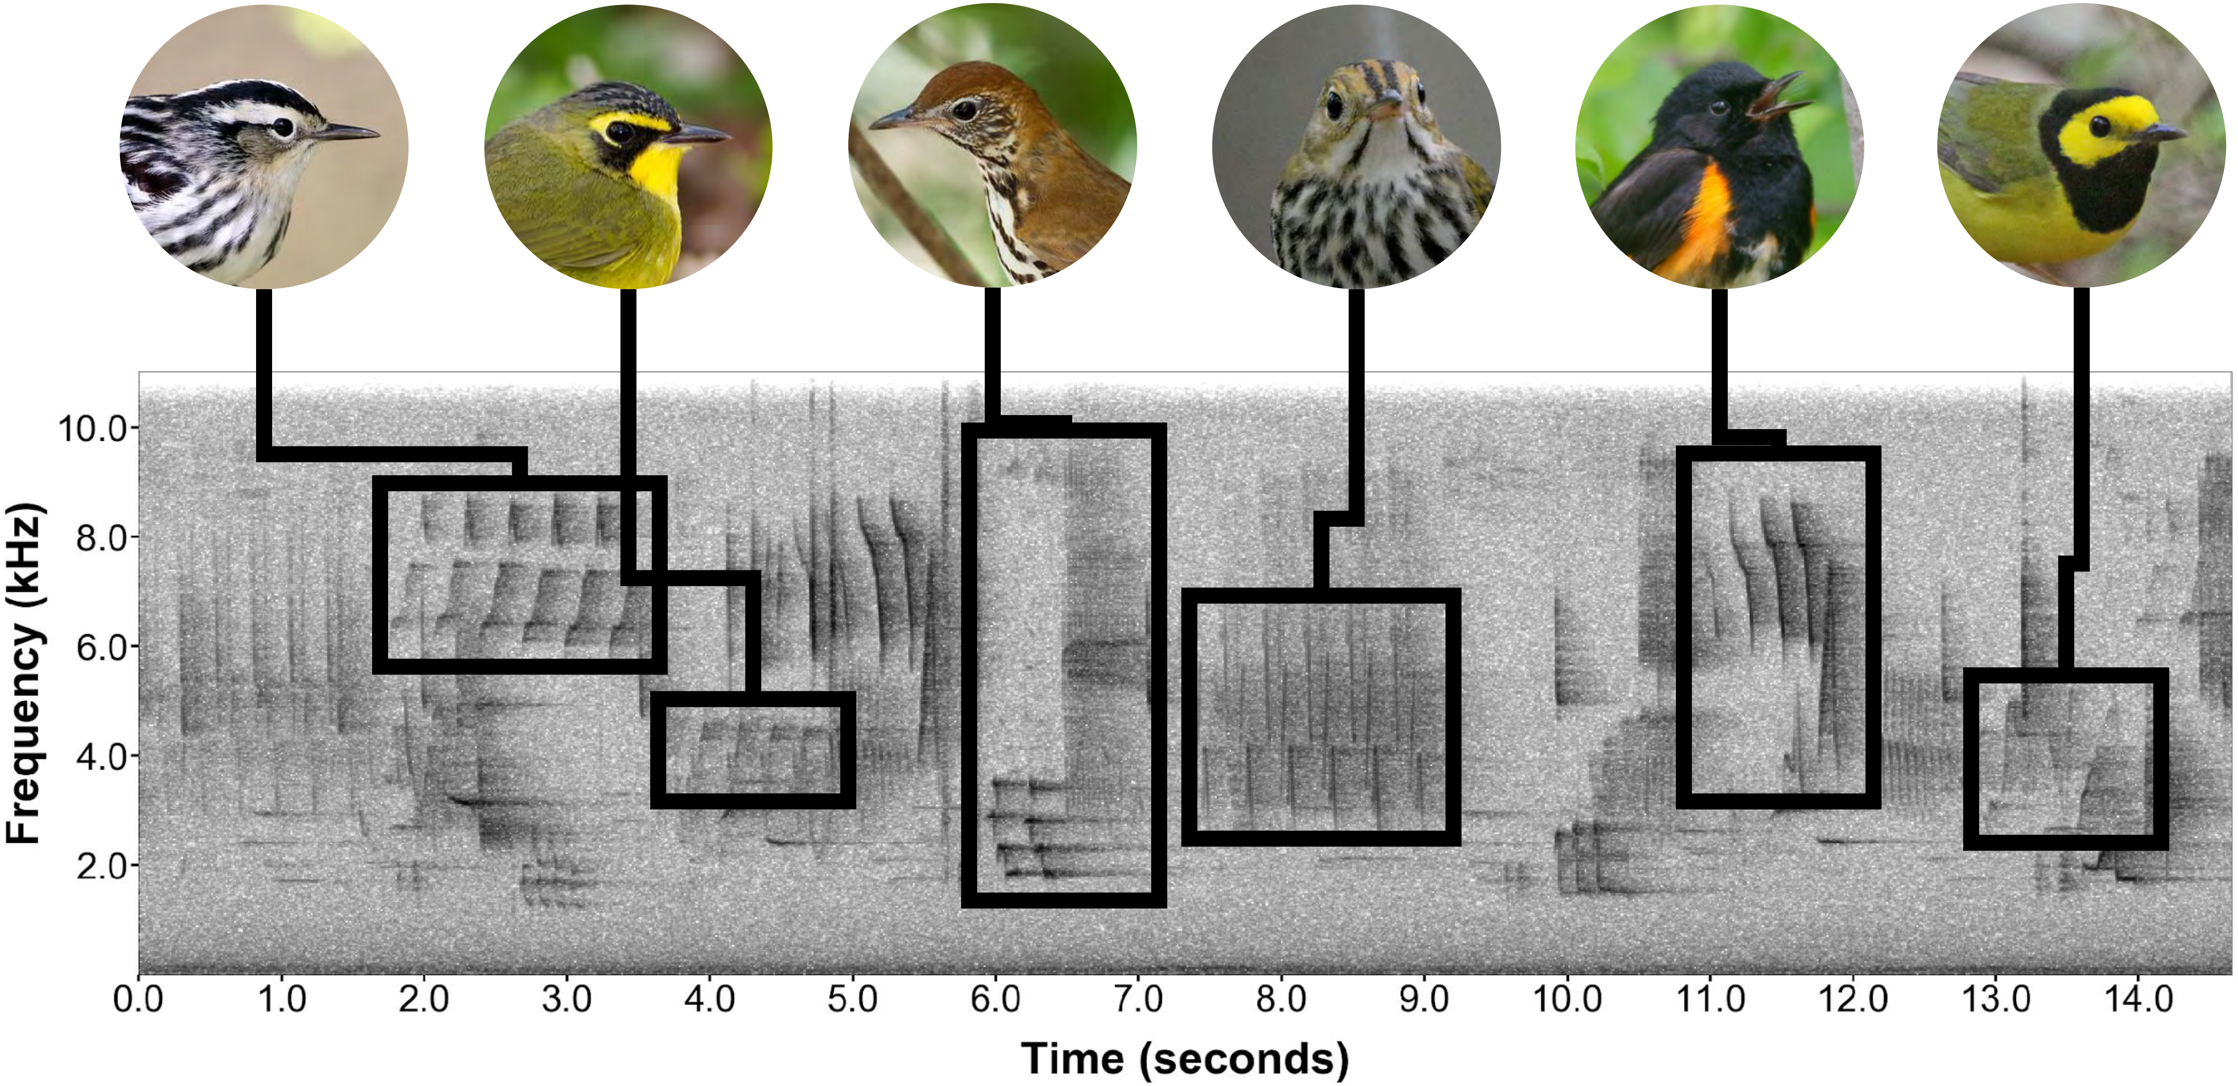

Frequency (kHz)
10.0
8.0
6.0
4.0
2.0
0.0
1.0
2.0
3.0
4.0
5.0
6.0
7.0
8.0
9.0
10.0
11.0
12.0
13.0
14.0
Time (seconds)

GPU computing accelerates deep learning sound recognition

Adult hoots

Juvenile Begging

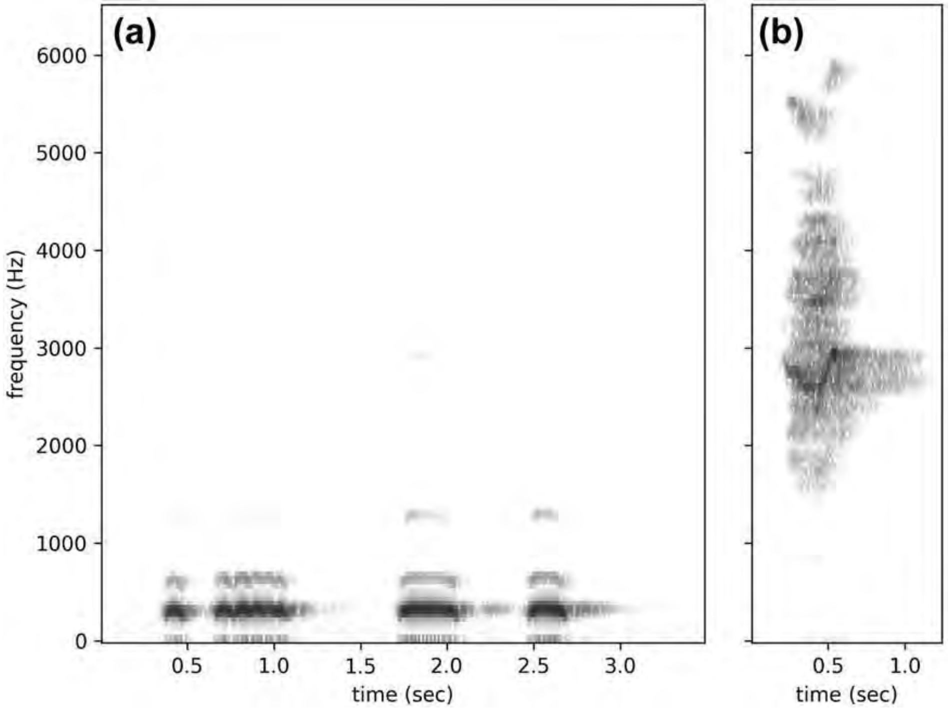

compared to CPU, GPU provides a **40x** speed-up for training and applying deep learning models

for instance, training a model could take 1 hour instead of 40

Supercomputing-powered analysis of the massive dataset resulted in the **detection** of **rare juvenile begging events**

allowing Lauren to study **habitat needs** of juvenile owls

Findings: juvenile owls require diverse and distinct habitat types, including open unforested patches

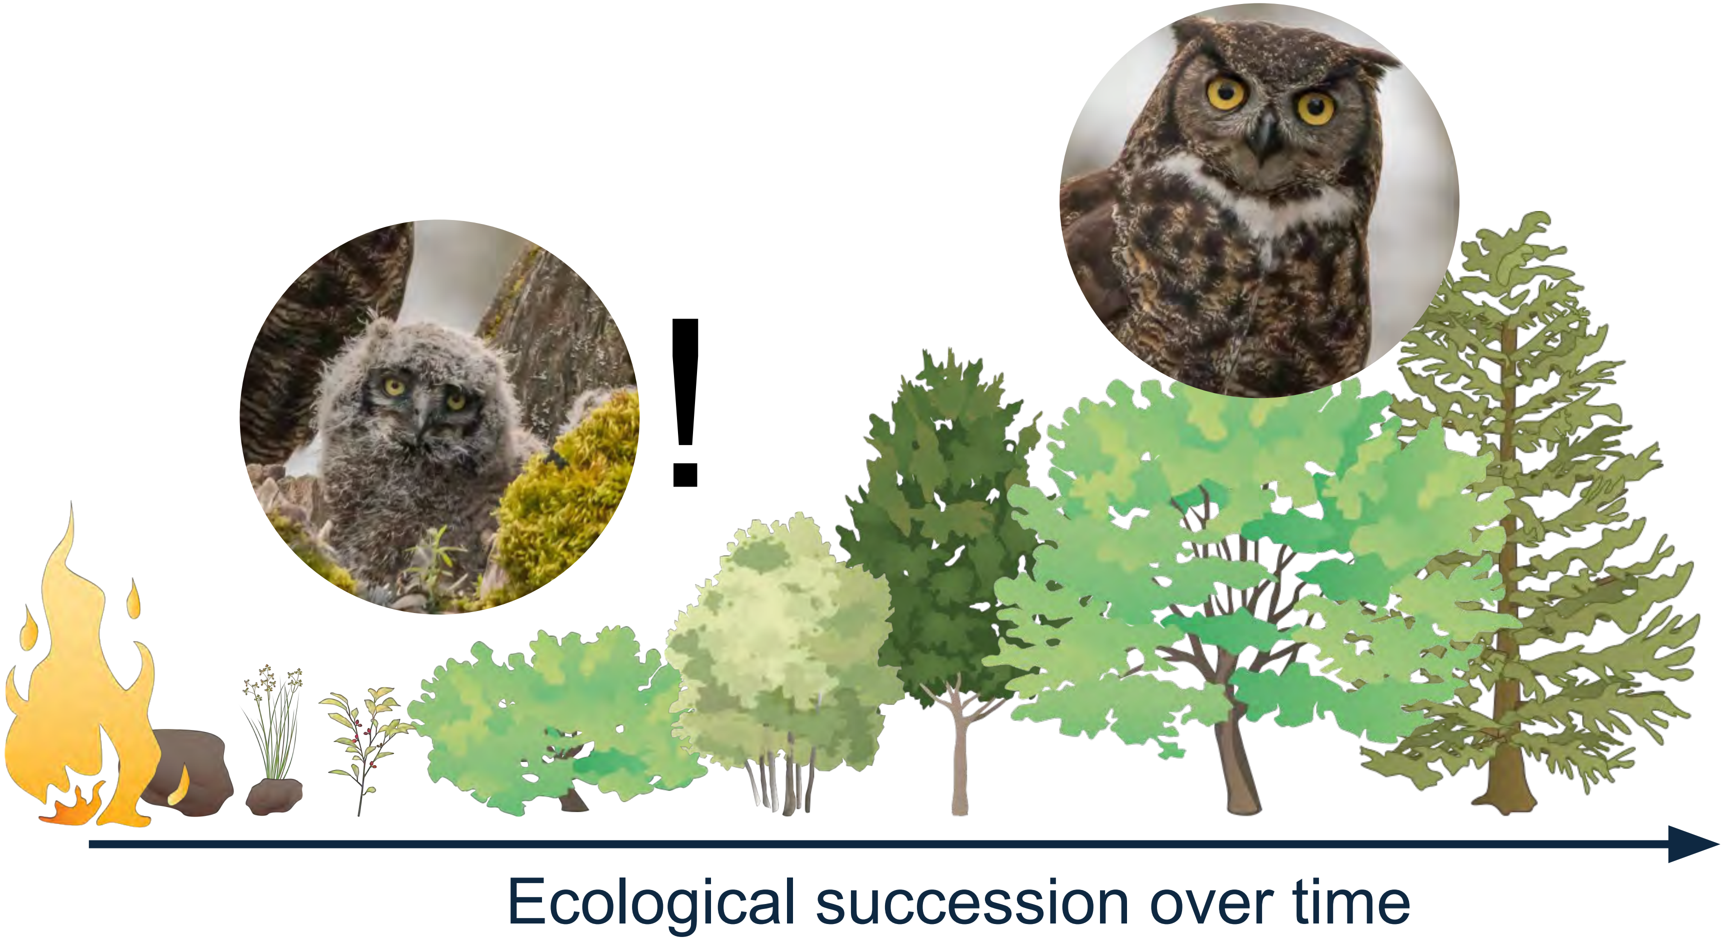

It takes a landscape to raise an owl

Part 2: searching in the dark

rediscovering lost frogs in Panama with Extreme Memory embedding querying

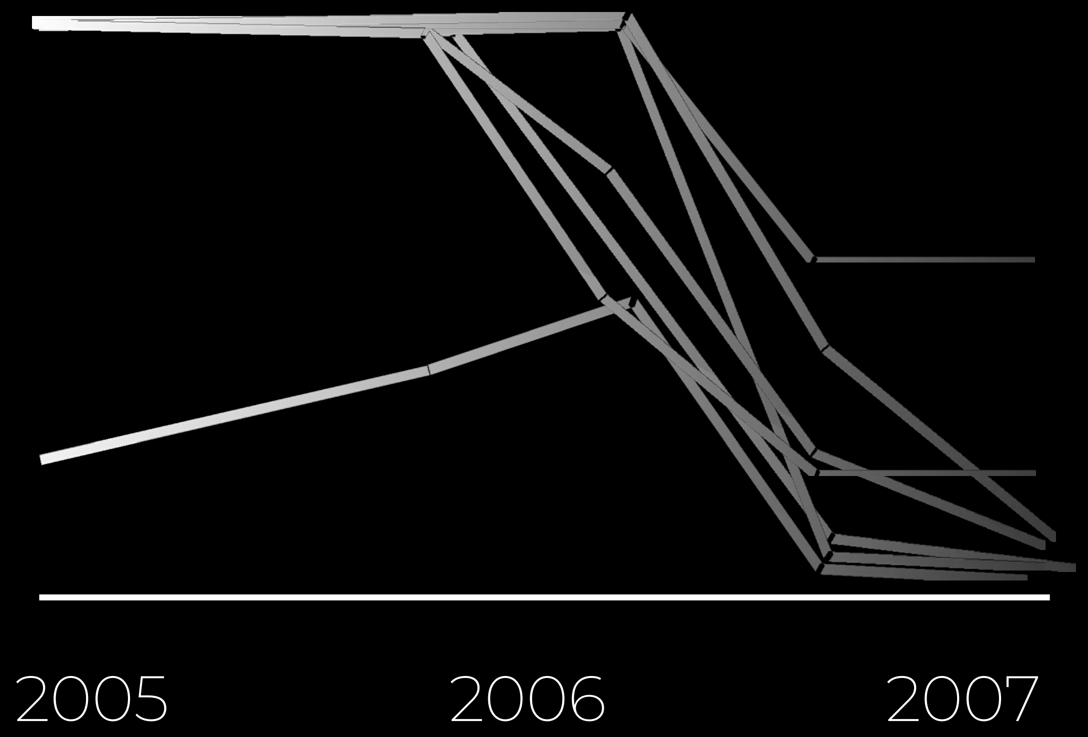

Chytridiomycosis and Amphibian Population Declines Continue to Spread Eastward in Panama
Woodhams et al, 2008. EcoHealth

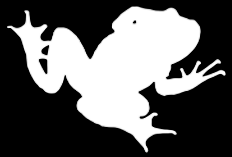
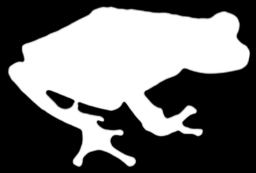


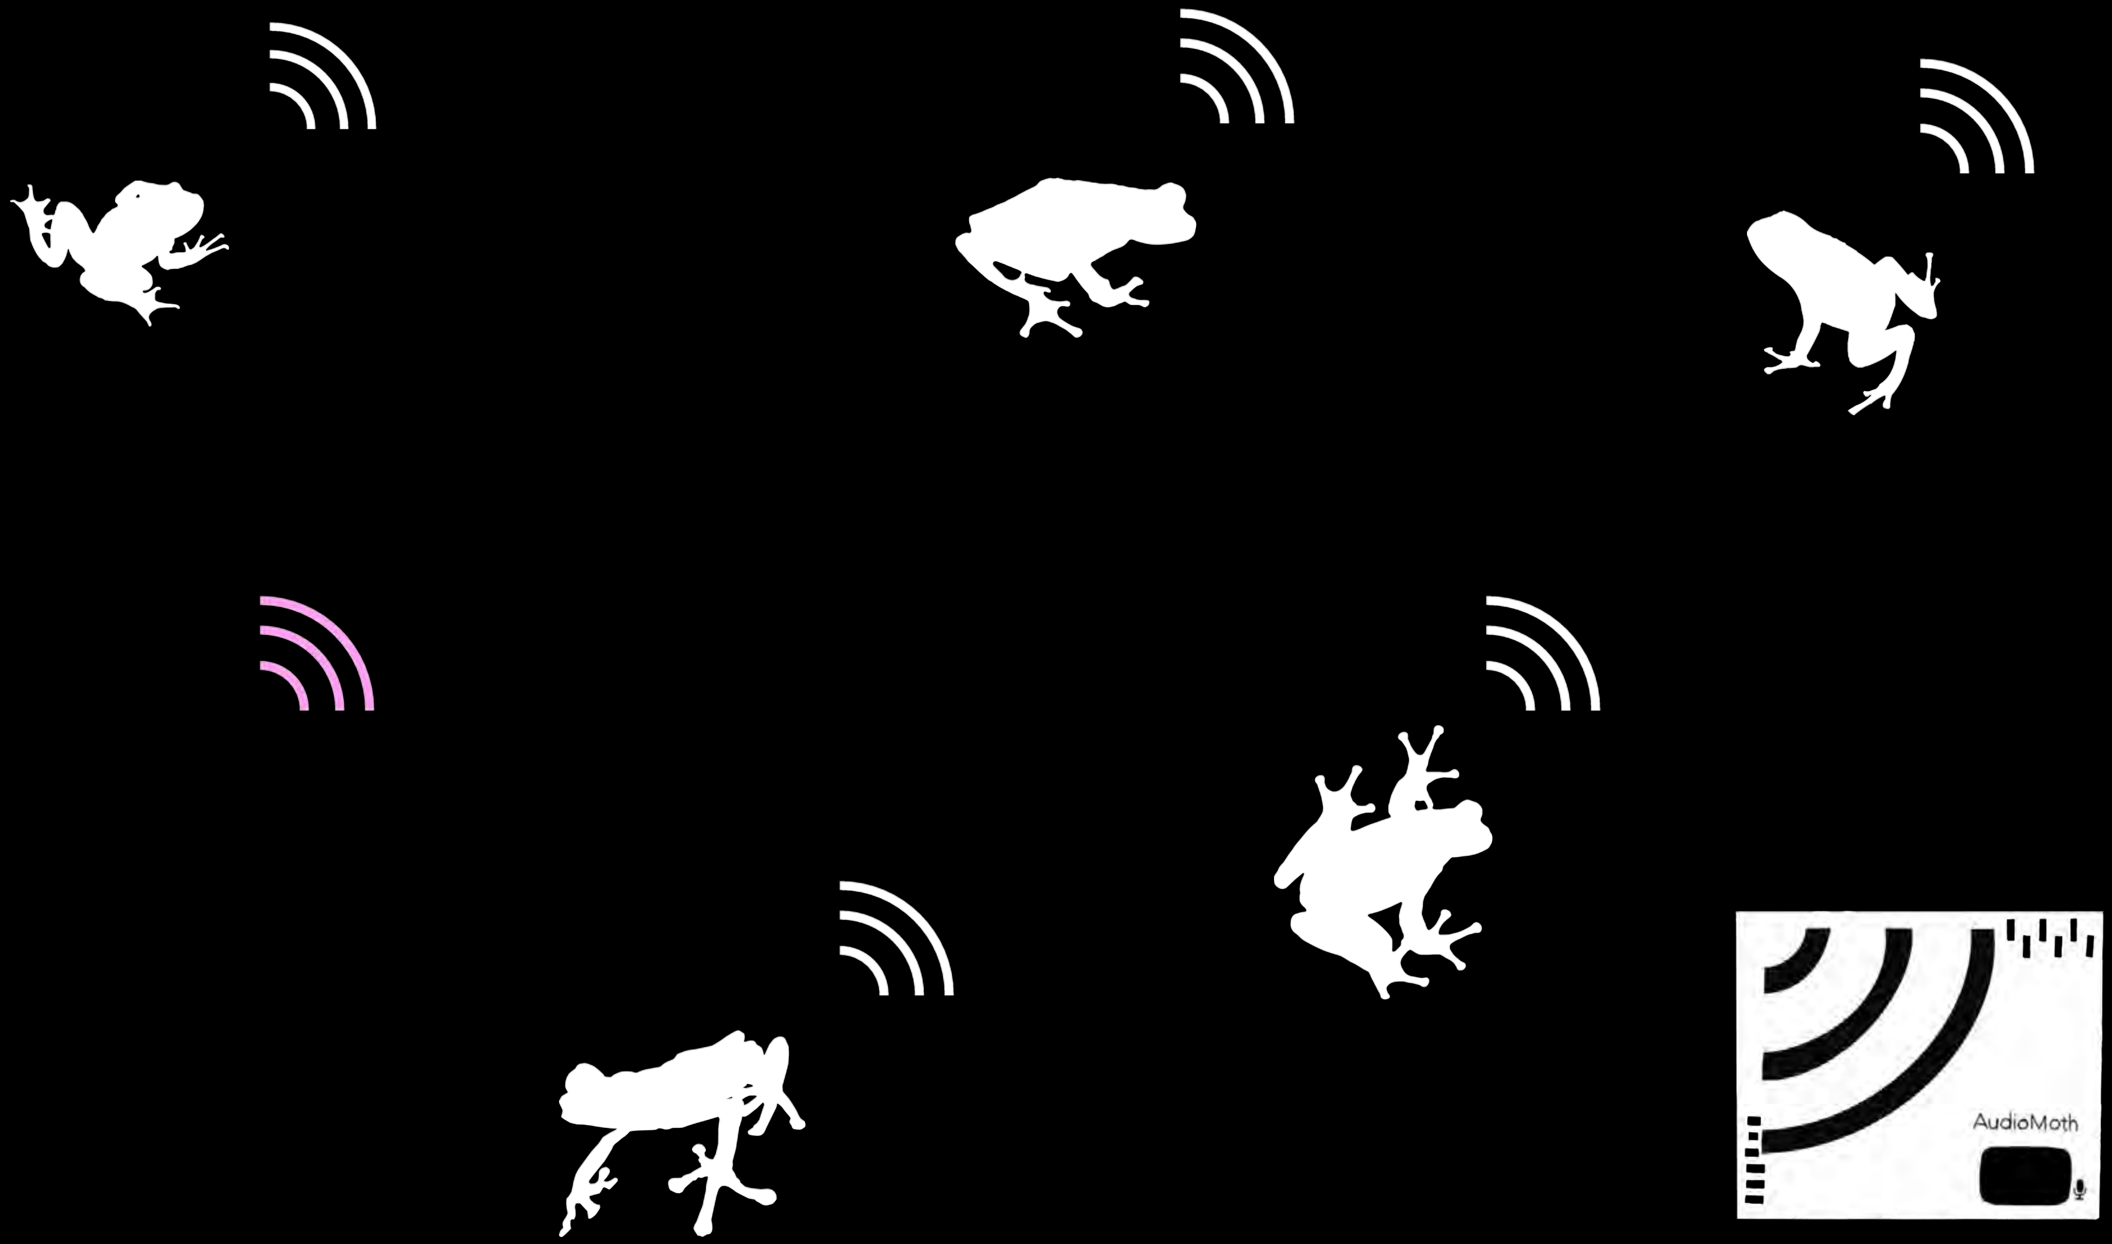
AudioMoth

Feature embeddings of a model trained to recognize bird vocalizations cluster by species

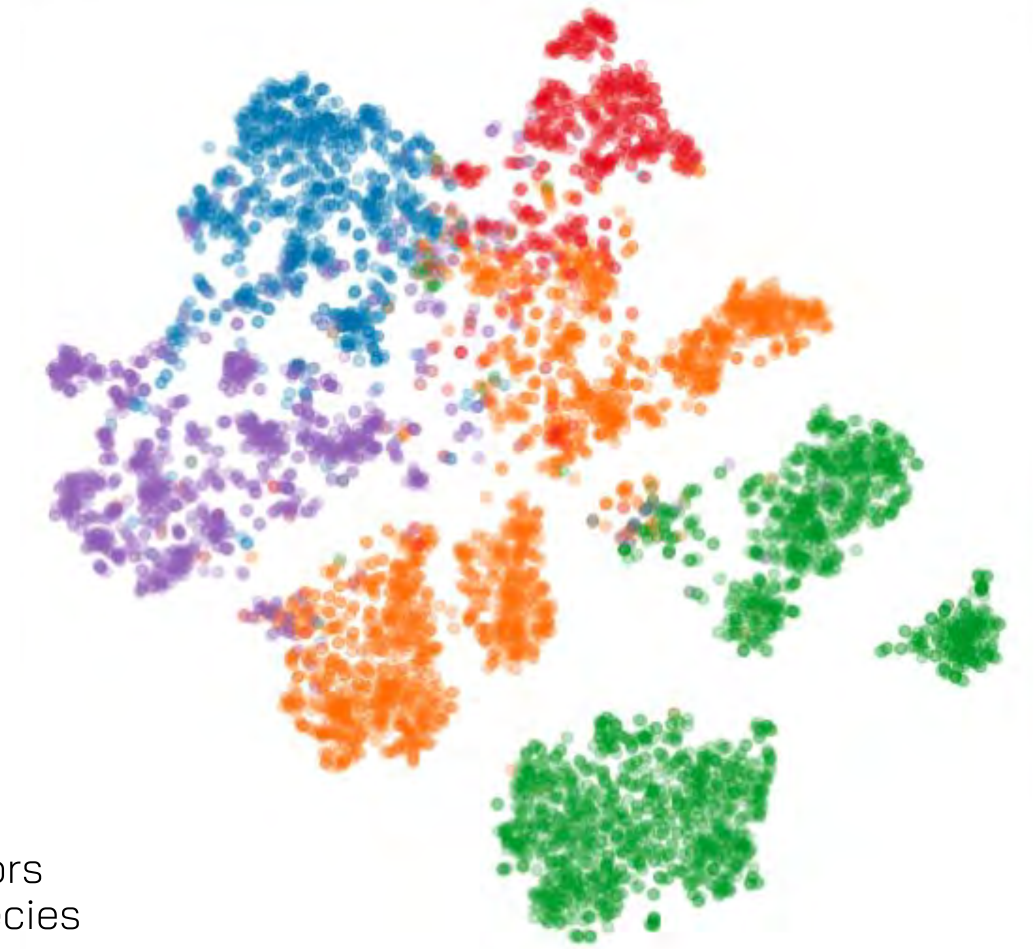

tSNE of feature embedding vectors with samples colored by bird species

Ghani e al., 2023. Feature embeddings from large-scale acoustic bird classifiers enable few-shot transfer learning.

Using embedding model trained on birds to search unlabeled data for a specific frog sound

1. Generate embedding vectors for the new dataset
   (10 million samples)
2. Generate embedding vector for a sample of the sound you want to detect in the new dataset
3. Inspect the samples in the raw data with the most similar embedding vectors

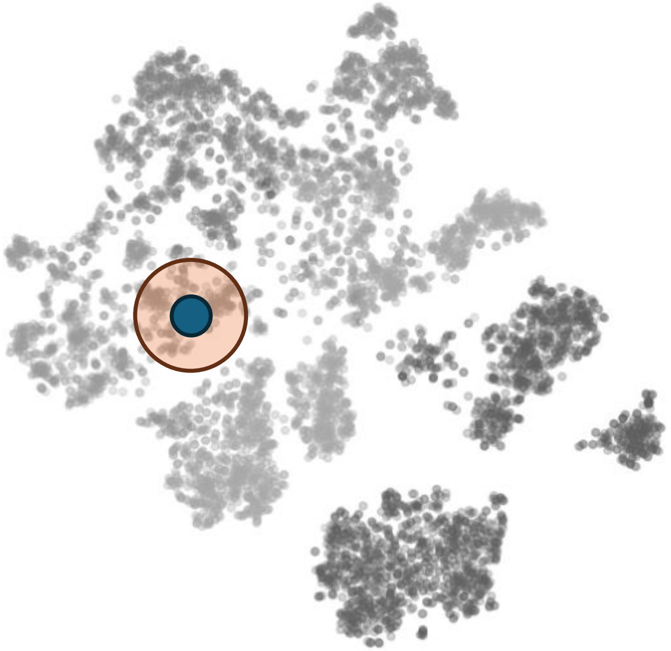

querying an entire dataset requires **extreme memory** nodes

Using embedding model trained on birds to search Panama data for a frog's sound

Query: *C. panamensis* call

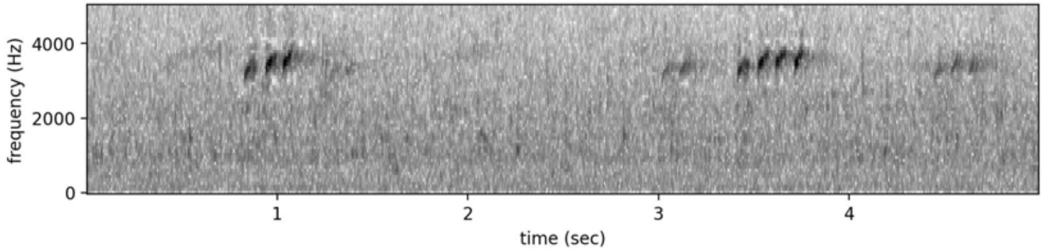

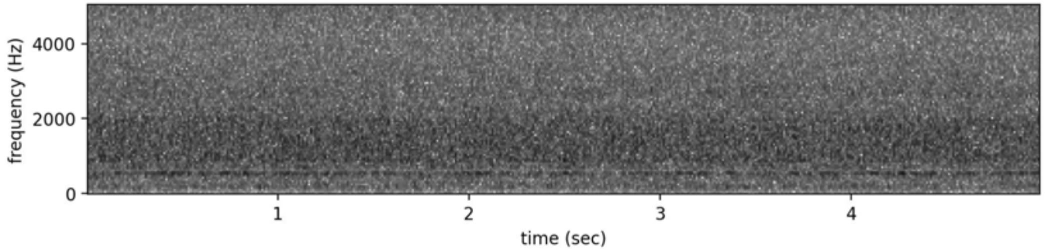

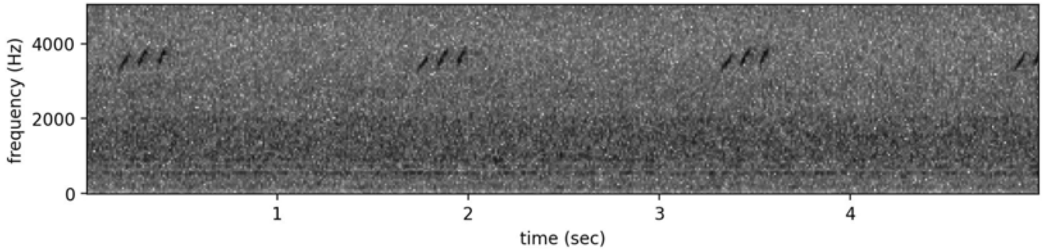

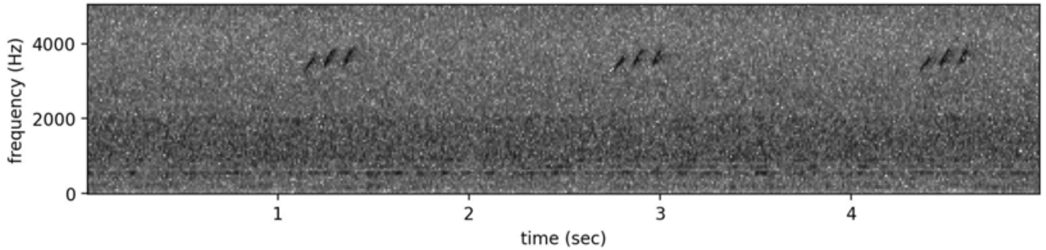

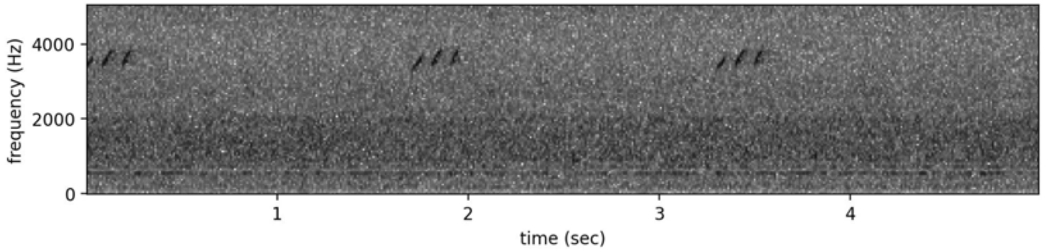

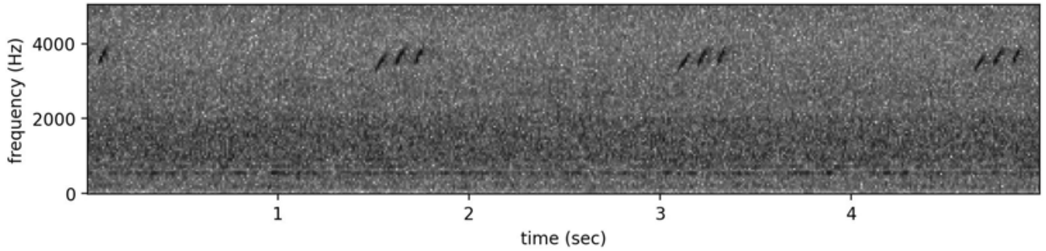

5 samples from field data with most similar embedding vectors

Atelopus varius

critically endangered

Atelopus varius

Query: *A. varius* call

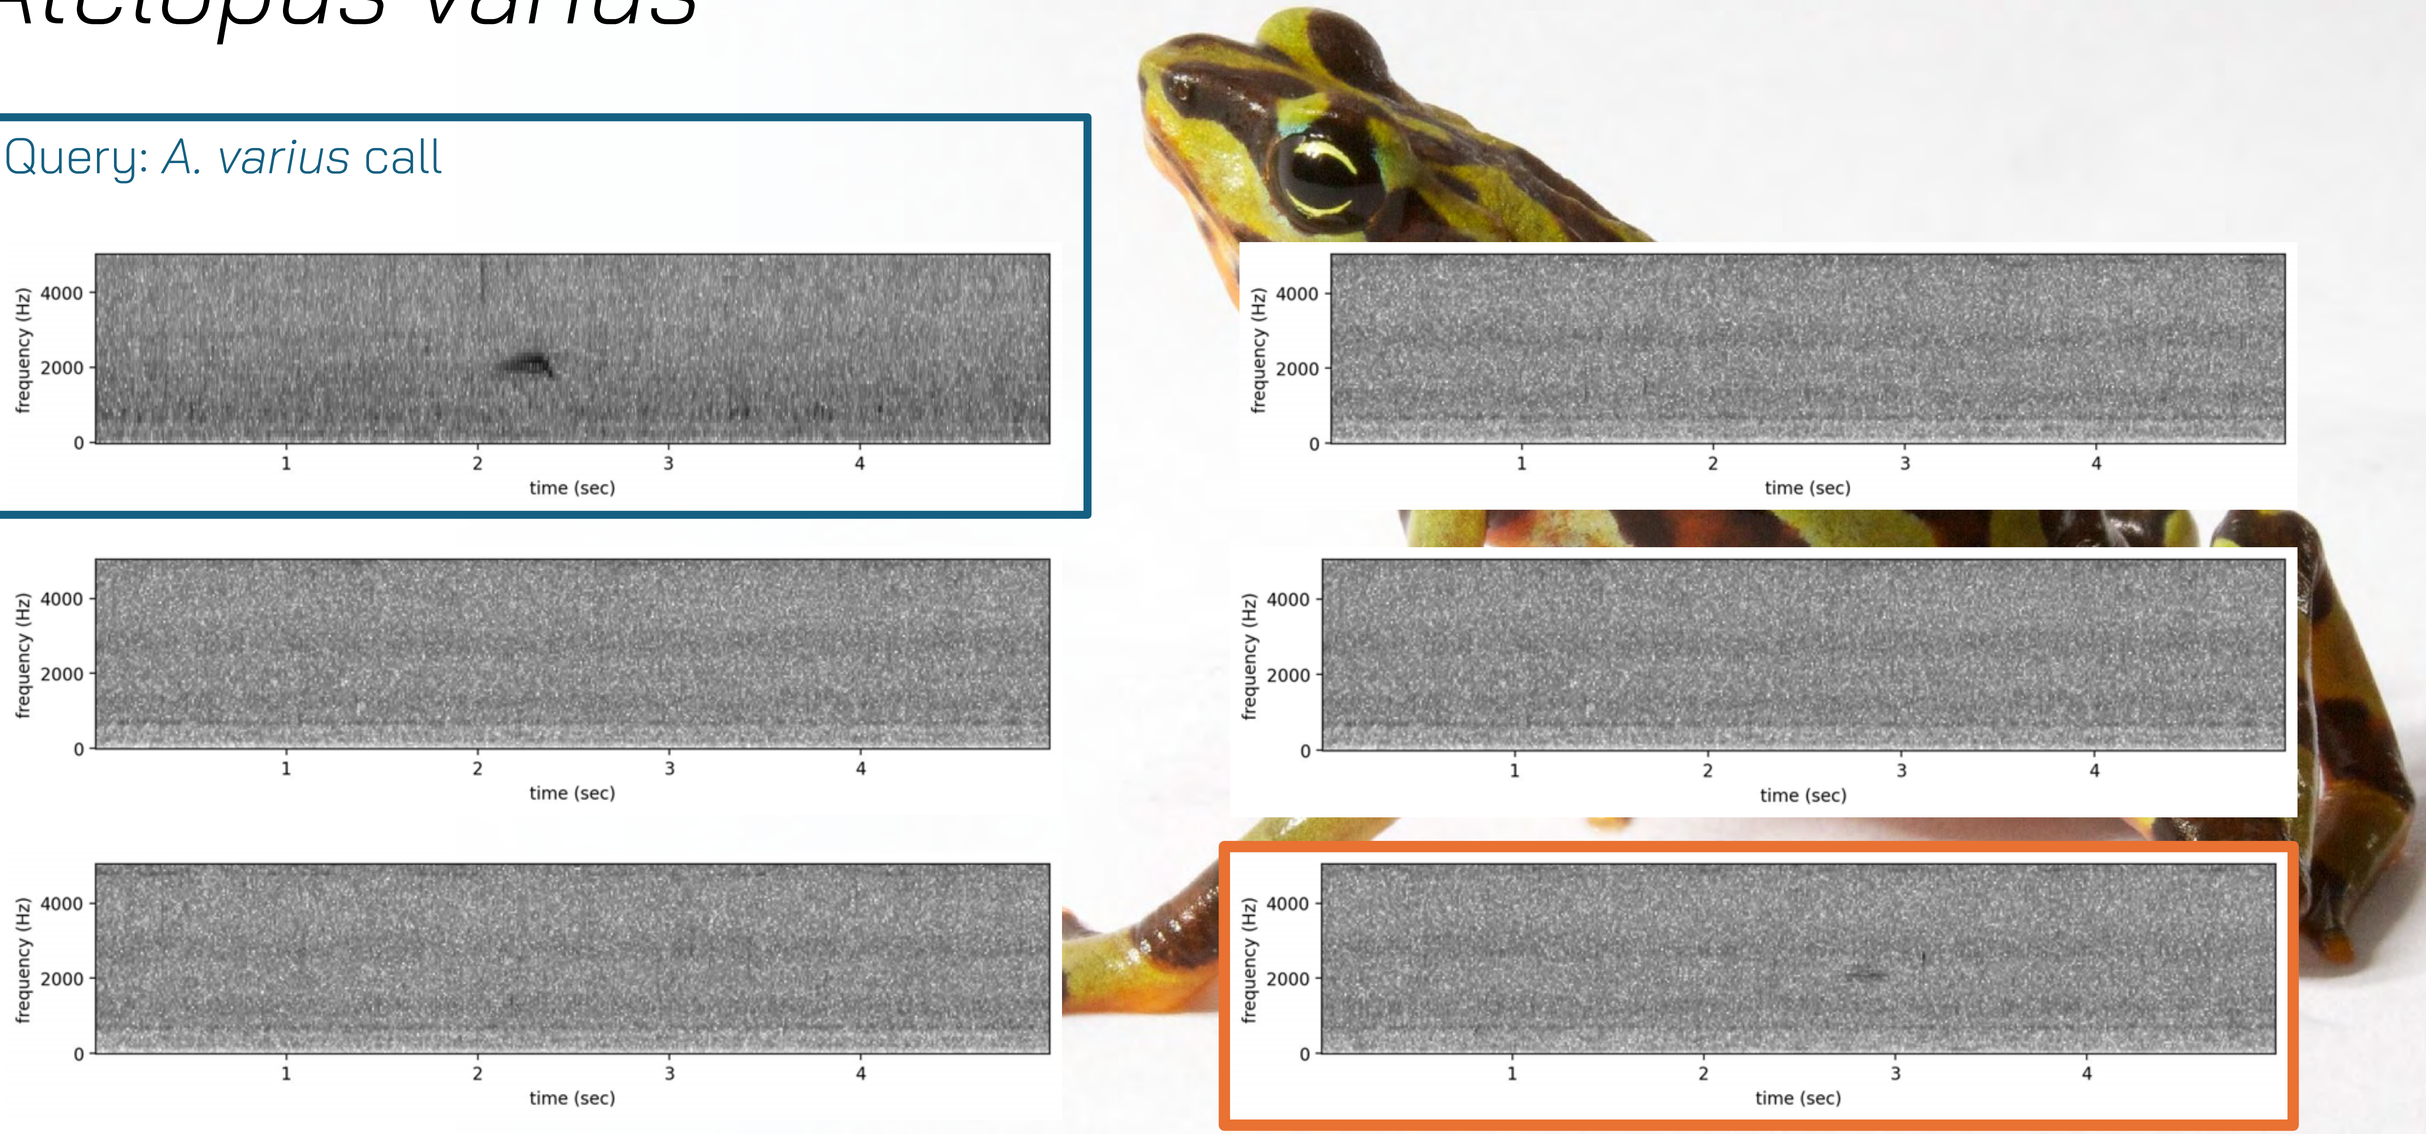

5 samples from field data with most similar embedding vectors

searching in the dark

rediscovering lost frogs in Panama with Extreme Memory embedding querying

Part 3: A day in the life of a warbler

Automated Acoustic Localization

Acoustic localization:
spatially triangulating sound positions

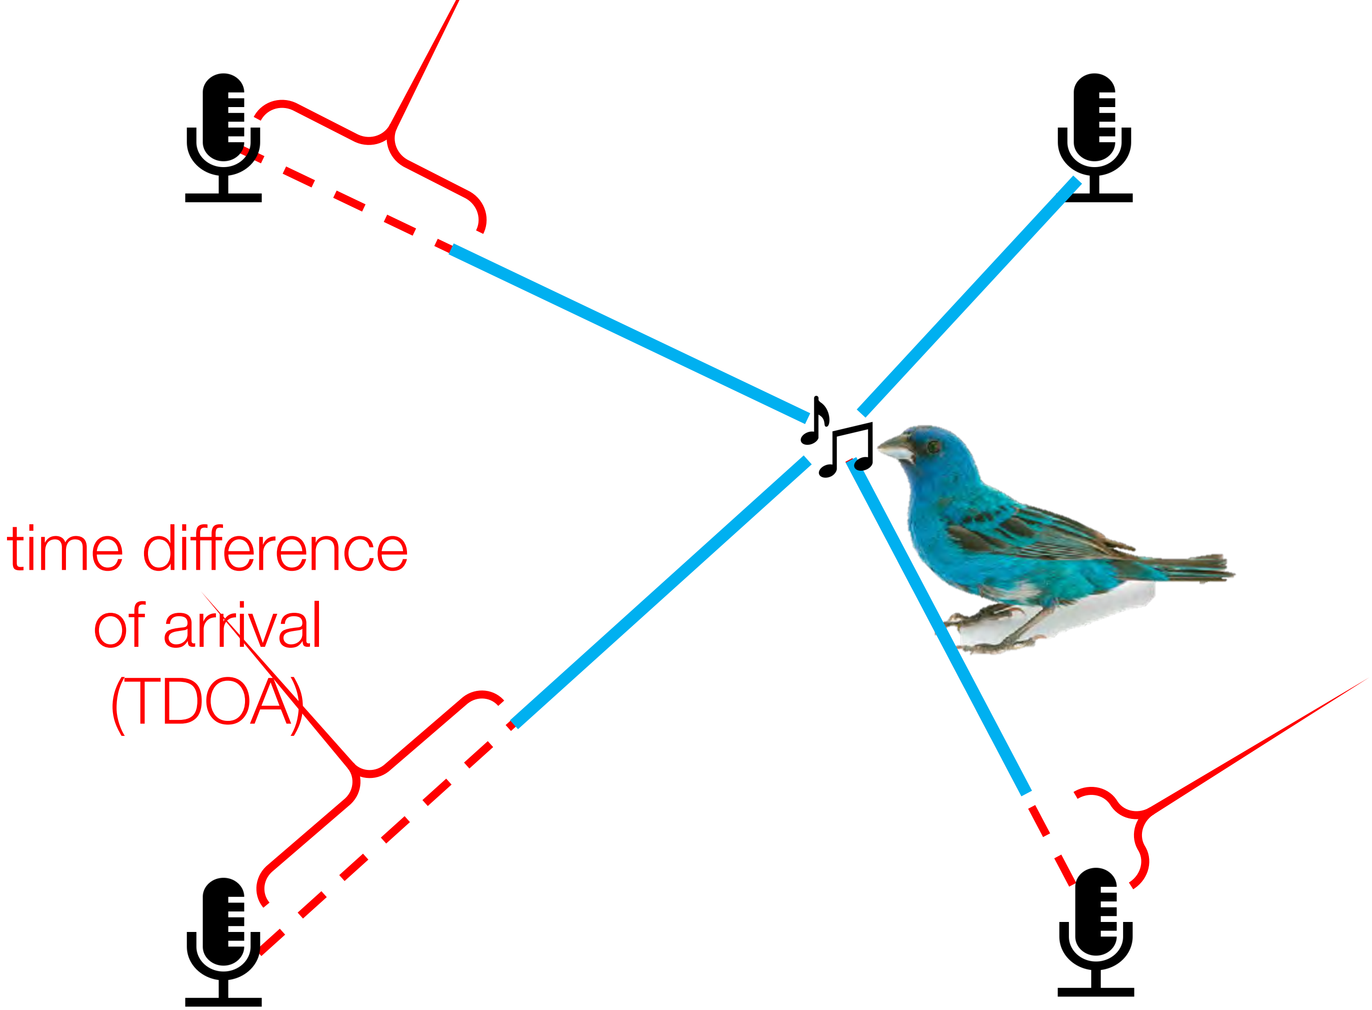

Array of time-synchronized automated recorders (ARUs)

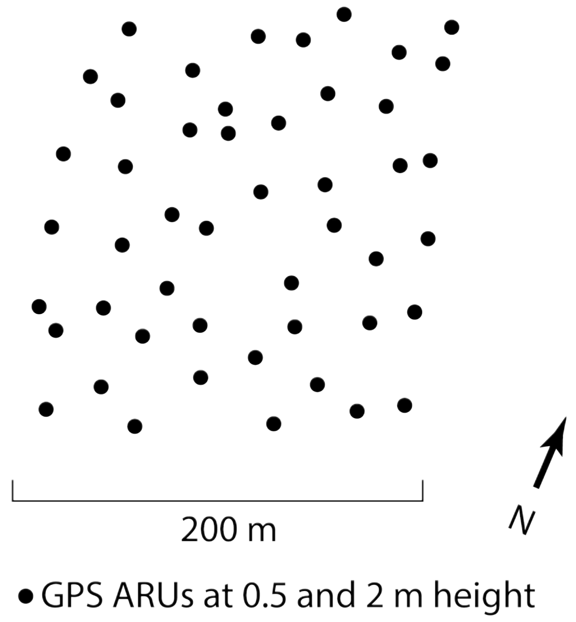

Automated detection and localization of bird song

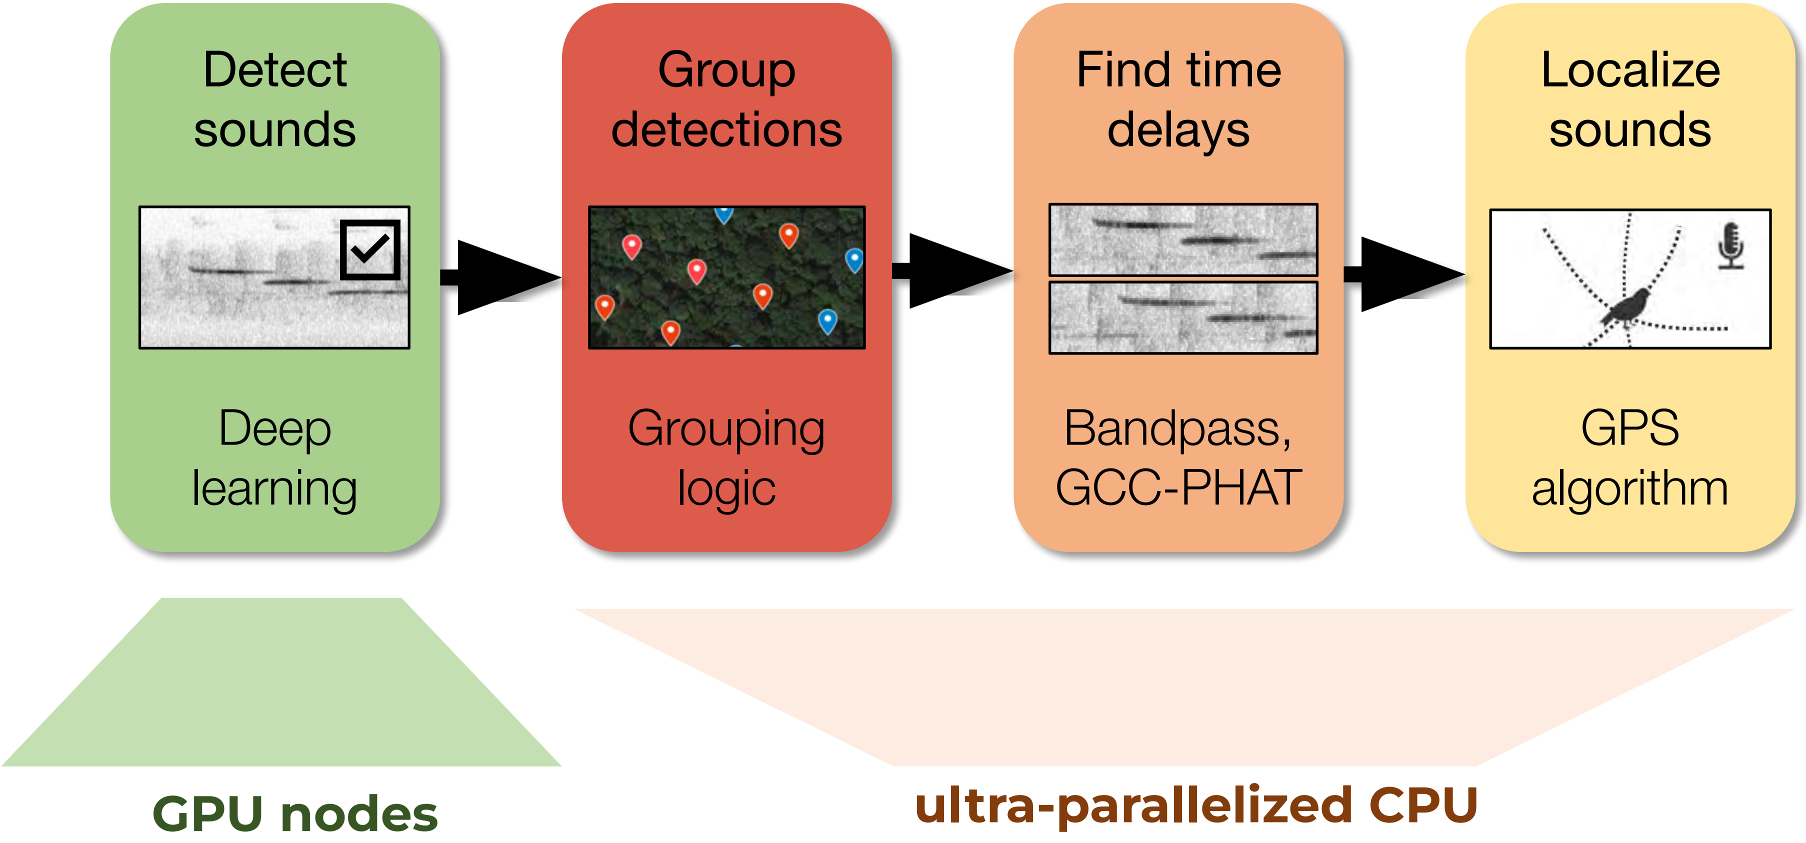

A day of singing activity for six species

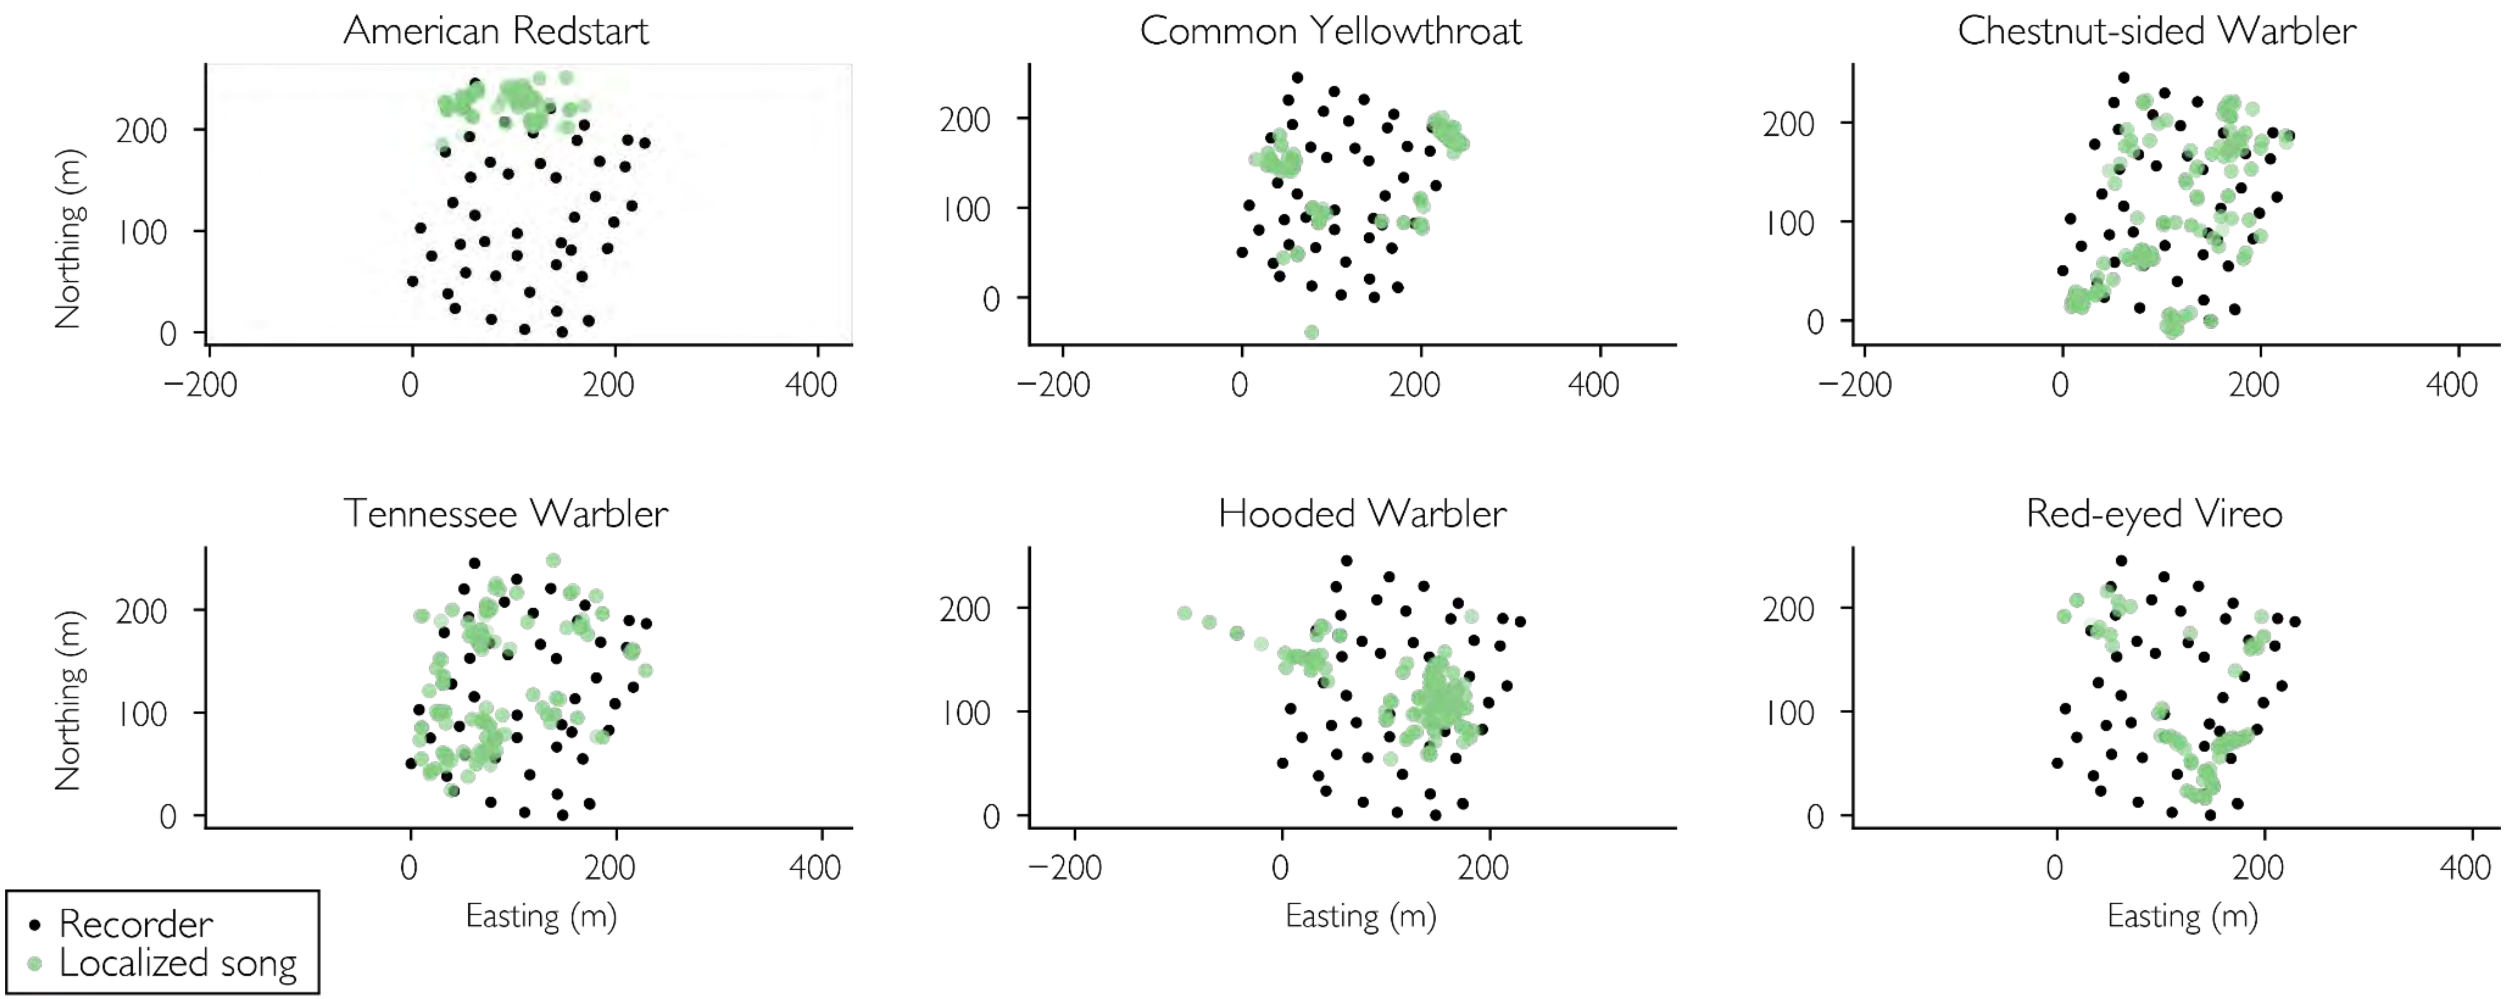

by Tessa Rhinehart

A season’s worth of localization

By day

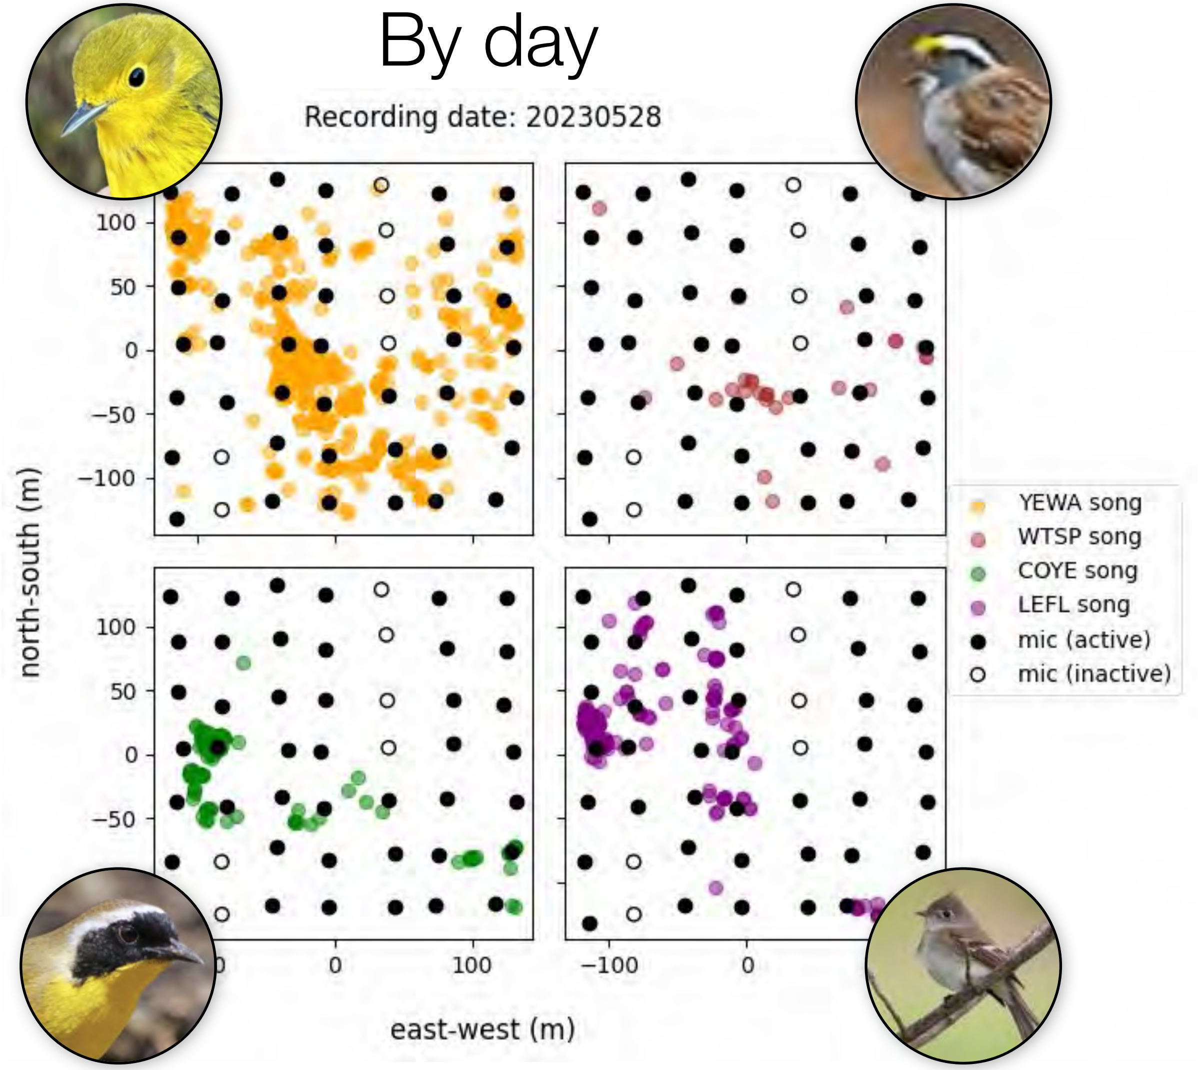

Heatmap for entire season

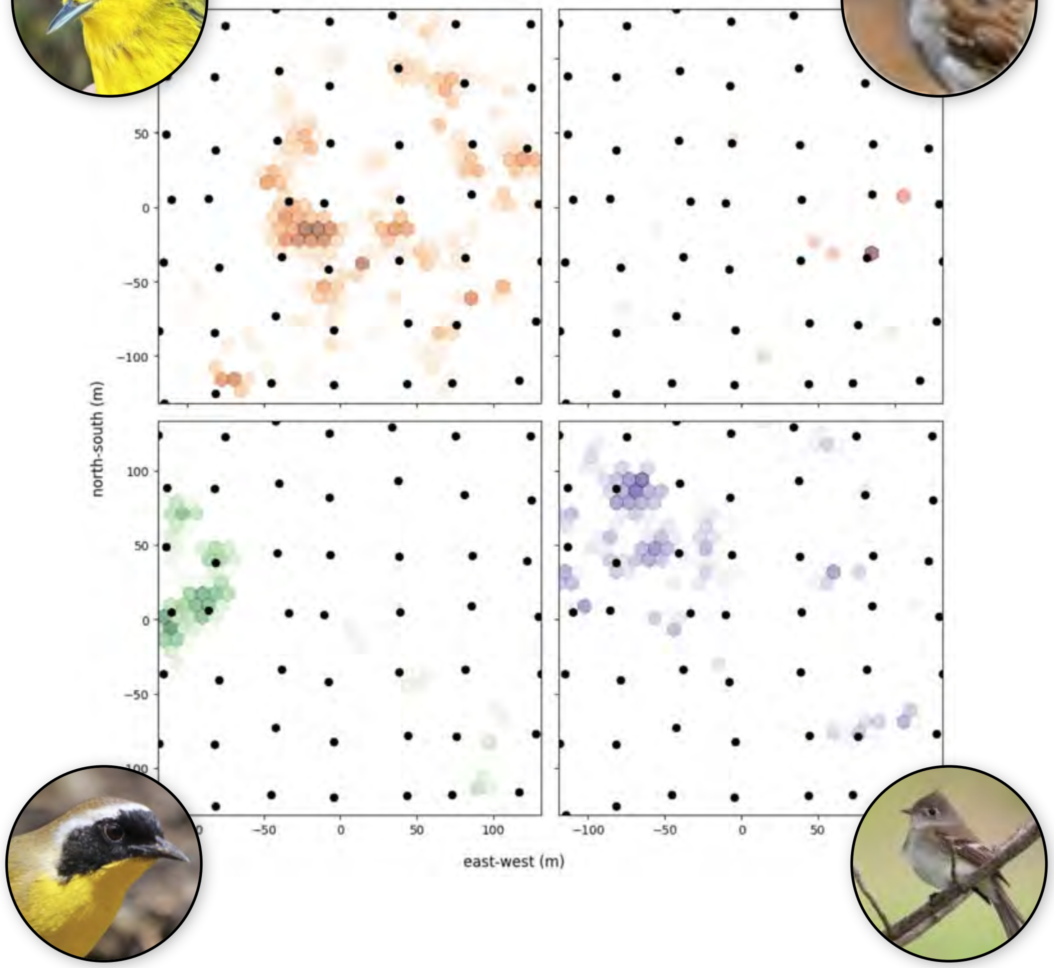

recognizing individual Song Sparrows by song

Example spectrograms of sampled song sparrow songs, each representing a different song variant.

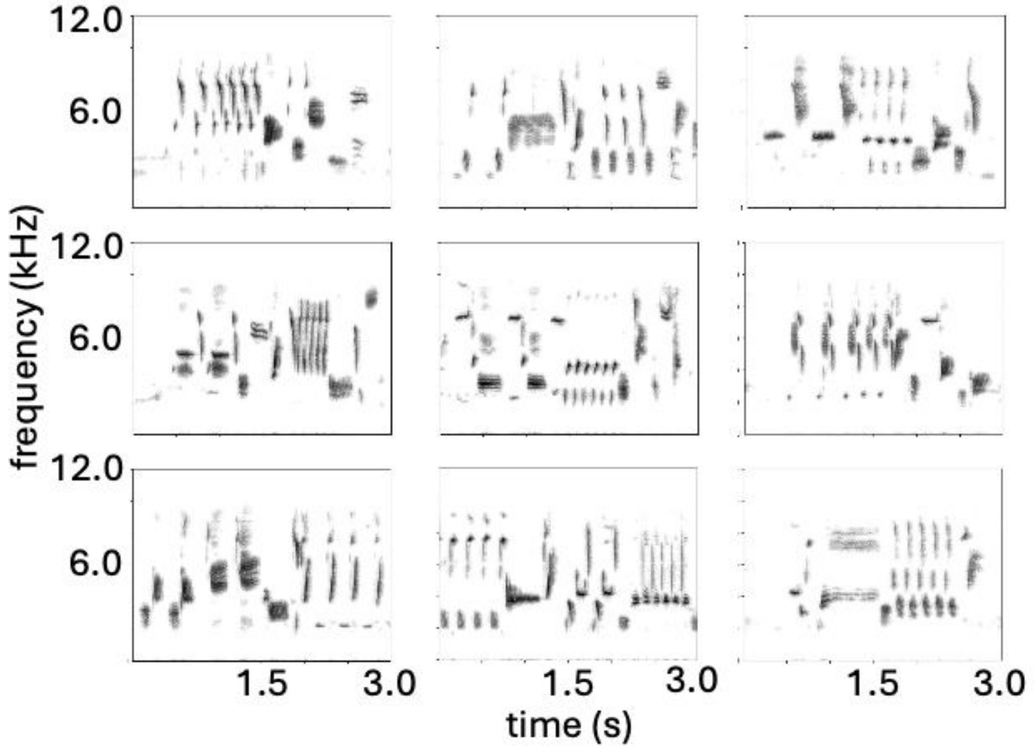

Automatically separated versus hand separated song sparrow song variants. Points are assigned by automated separation and colors indicate each variant determined by a trained observer.

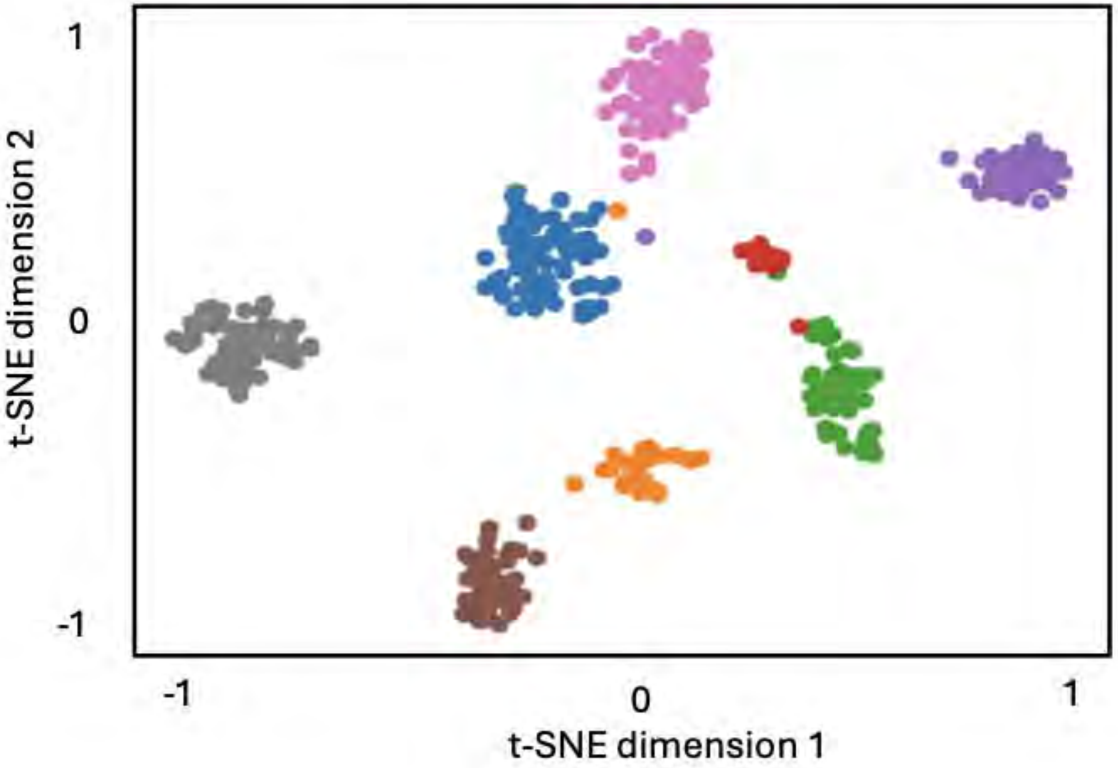

by Lauren Chronister

Acknowledgements

Advisor: Justin Kitzes

Compute: PSC Bridges2 & Pitt CRC

Funding sources:
NSF, Gordon and Betty Moore Foundation

Kitzes Lab & collaborators, especially:
Lauren Chronister, Sasha Syunkova, Santiago Ruiz-Guzman, Louis Freeland-Haynes, Tessa Rhinehart, Jeffrey T. Larkin

Thank you

Sam Lapp
sam.lapp@pitt.edu