

Deep Learning in the Wild


Embedded PyTorch for Real-World Conservation Bioacoustics




GUESS THAT SOUND!



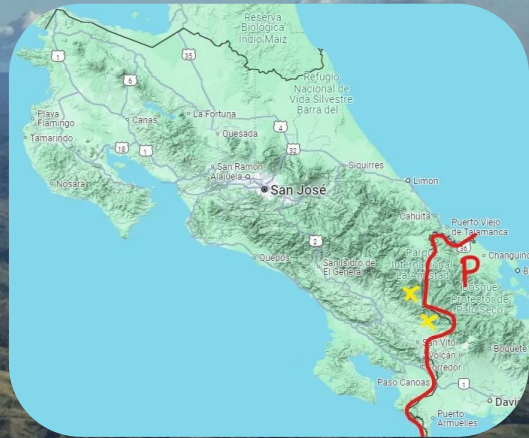
**Chestnut-Mandibled
Toucan**



Me




**Three-Wattled
Bellbird**



La Amistad is a Remote Area

- No power grids
- No cell towers
- Satellites are expensive and power hungry


THREE BODY PROBLEM



**Acoustic Machine
Learning Model**




**Battery
Optimization**



**Low Power
Connectivity**

OWL IOT – MESH NETWORK




PHASE I

Oct 2024 - December 2024



PHASE I – MACHINE LEARNING MODEL



K. Lisa Yang Center for
Conservation Bioacoustics

ARCHITECTURE

How BirdNET works

BirdNET processes raw acoustic data through a multi-stage pipeline designed for ecological accuracy. By transforming varying soundscapes into standardized feature representations, the system can isolate and identify subtle vocal signatures across thousands of species while accounting for biogeographical variations in species distribution.



1. Capture



2. Spectrogram




3. Neural Net




4. Result




Audio is captured at 48 kHz and divided into 3-second segments, optimizing the balance between model input size and the natural duration of avian vocalizations.



Signals are processed into two log-scaled Mel-spectrograms, visualizing frequency patterns from 0 to 3 kHz and 150 Hz to 15 kHz for detailed analysis.



A Convolutional Neural Network (CNN) scans these visuals, utilizing millions of trained weights to detect species-specific patterns.



Initial predictions are cross-referenced with local metadata (location and date) to produce high-confidence species probabilities.


PHASE I – TECH STACK

Machine Learning


OWL IoT



Microphone




RaspberryPi

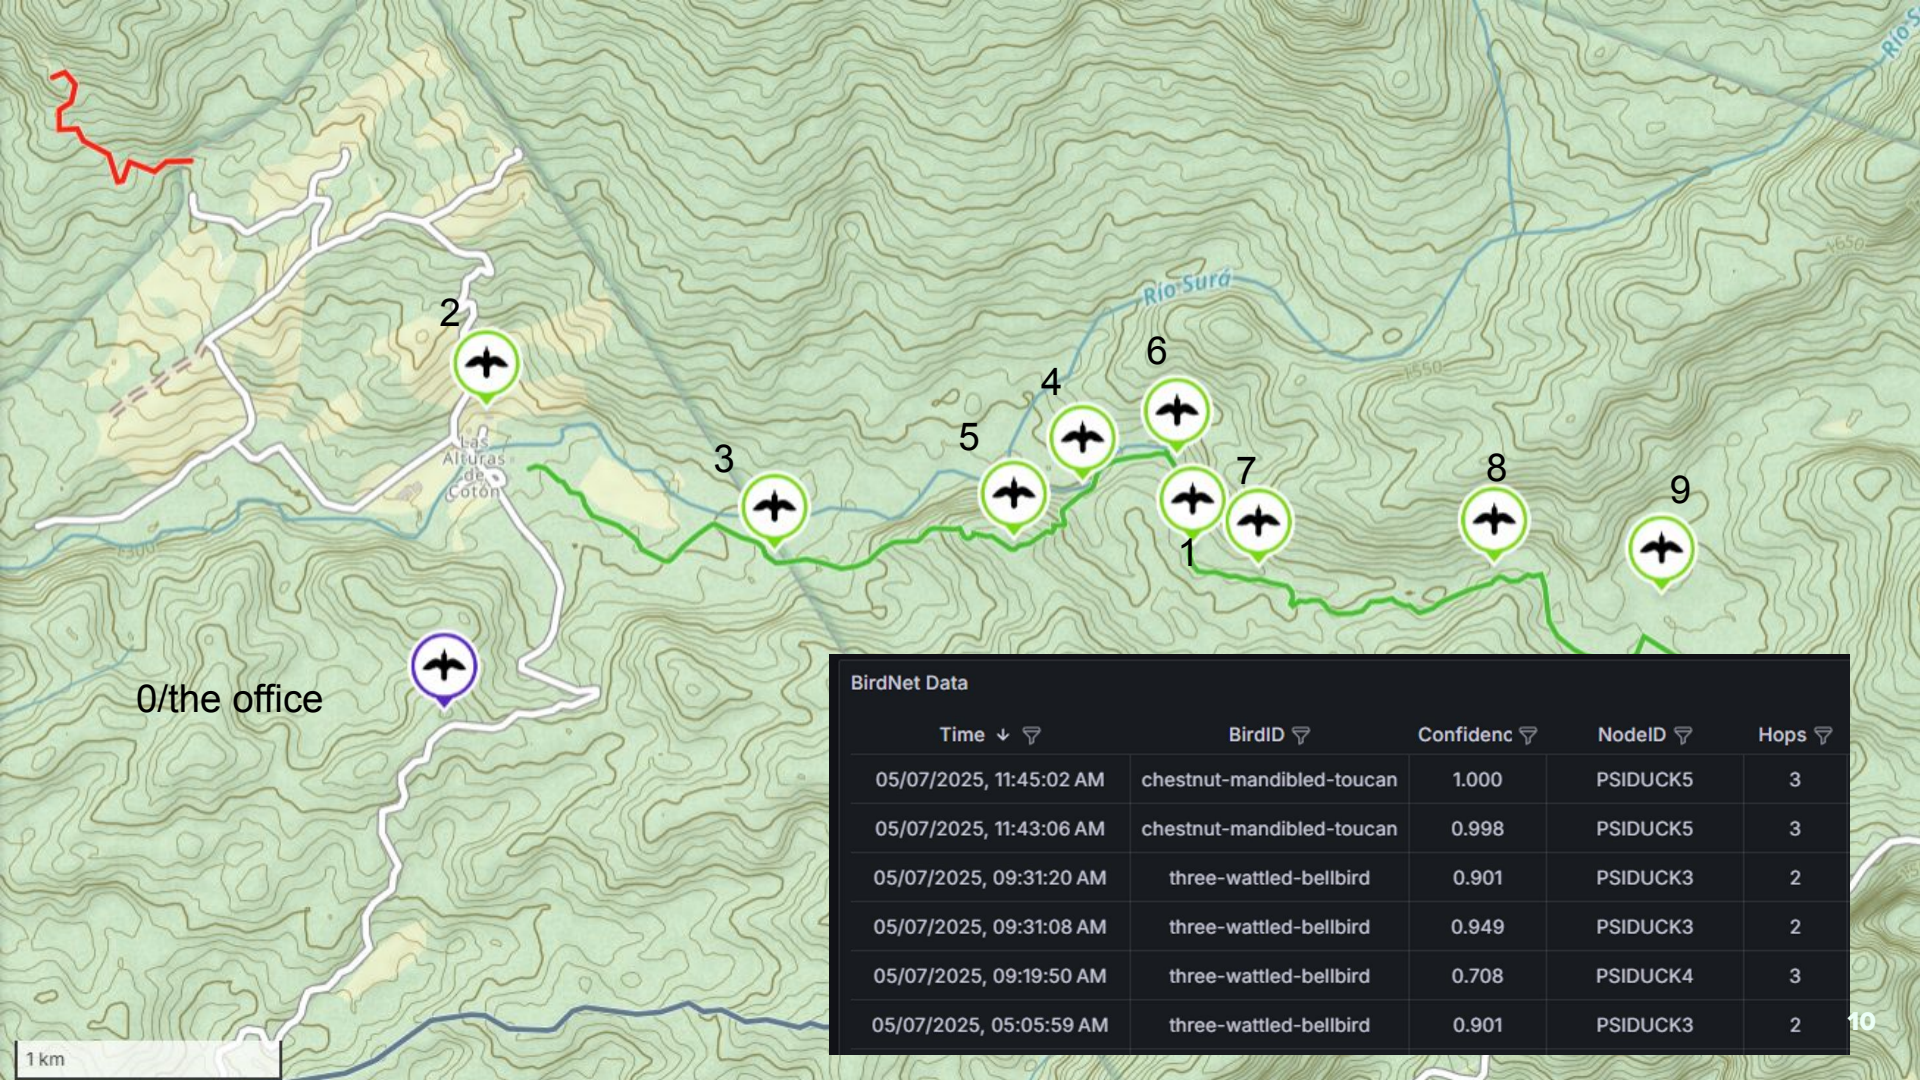


TBeam

Power: 20,000 mAh Battery &
20 W Solar Panel

PHASE I – DEPLOYMENT JAN 2025





BirdNet Data

Time ↓	BirdID ↓	Confidenc ↓	NodeID ↓	Hops ↓
05/07/2025, 11:45:02 AM	chestnut-mandibled-toucan	1.000	PSIDUCK5	3
05/07/2025, 11:43:06 AM	chestnut-mandibled-toucan	0.998	PSIDUCK5	3
05/07/2025, 09:31:20 AM	three-wattled-bellbird	0.901	PSIDUCK3	2
05/07/2025, 09:31:08 AM	three-wattled-bellbird	0.949	PSIDUCK3	2
05/07/2025, 09:19:50 AM	three-wattled-bellbird	0.708	PSIDUCK4	3
05/07/2025, 05:05:59 AM	three-wattled-bellbird	0.901	PSIDUCK3	2

PHASE I – SUMMARY

Wins

- ML model works on the edge
- Relay detections through OWL IoT

Improvements

- Reduce the false positives
- Improve battery life

PHASE II

January 2026 - July 2026



PHASE II – DRAWING BOARD


How do we improve power?

- Raspberry Pi is the most power consuming component.
 - However we needed it to run CNN and FFT on the edge.

How do we improve ML?

- Do we need to run Birdnet? Can we train our own lightweight model?

DIGITAL SIGNAL PROCESSING (DSP)



PHASE II – PIPELINE

DSP STAGE

Microphone Input
Harmonic Filtering



ML STAGE

Machine Learning
Send Results with OWL IoT



PHASE II – MODEL ARCHITECTURE

- Using DSP, we gathered about 170 features
- Used Feature Selection to select top twenty features
- Random Forest
 - Fast and can works on edge hardware.
- Lightweight Neural Networks
 - Fast and more accurate.
 - Need to convert from PyTorch to TFLite for Microcontrollers.
 - Can shrink size of model.

PHASE II – TECH STACK

ML and DSP

OWL IoT



Microphone

Teensy 4.1

T-Beam

Power: 20,000 mAh Battery &
20 W Solar Panel

PHASE II – MODEL PERFORMANCE

	BirdNet	Random Forest	MLP
Inference Time			
TWB (F1 Score)			
CMT (F1 Score)			
Size (MB)			
Estimated Lifetime			


PHASE II – IMPROVEMENTS & Next Steps

Model Optimization

- DSP Filter
- Lightweight Model
- Collect More Data
- Continue Tuning Parameters

Battery Optimization

- DSP runs most of time.
- Replaced Raspberry Pi 0 2W with Teensy
- Underclocking Teensy
- Sleep Mode



*We can now stray from the
beaten path of boring,
predictable gains from throwing
compute at the problem.*

Sara Hooker,
“On the Slow Death of Scaling”

THANK YOU!



Taraqur Rahman
taraqur@owlintegrations.com
www.owlintegrations.com



Owen O'Donnell
owen@owlintegrations.com
www.owlintegrations.com

