Orcasound & OrcaAL

Orcasound's open approach to orca detection
OrcaAL = Orca Active Learning

Presentation by Scott Veirs and Kunal Mehta
Meridian "Winter Webinar on Tuesday, December 15, 2020

Orcasound & OrcaAL

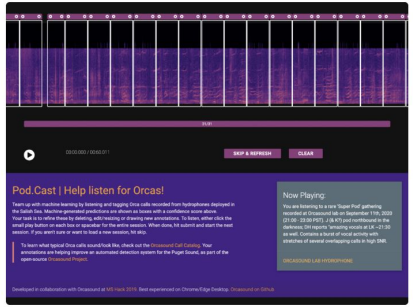

POD.CAST ANNOTATION SYSTEM

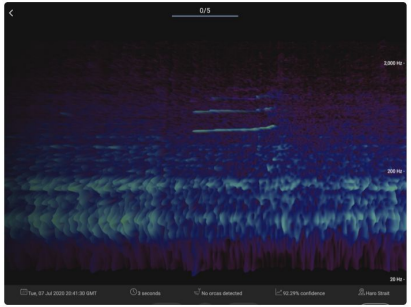

ORCA AL – SRKW

ML – ON THE EDGE

AI for orcas (#ai4orcas) -- ai4orcas.net

towards open bioacoustic data science...

ORCASOUND

A hydrophone network (WA, USA) & open-source software community

- 3 cabled nearshore sites streaming 24/7 in 2020
- Citizen scientists detect orca sounds in real-time via web app -- live.orcasound.net
- How can AI & human listeners work synergistically to conserve orcas & advance marine bioacoustics?

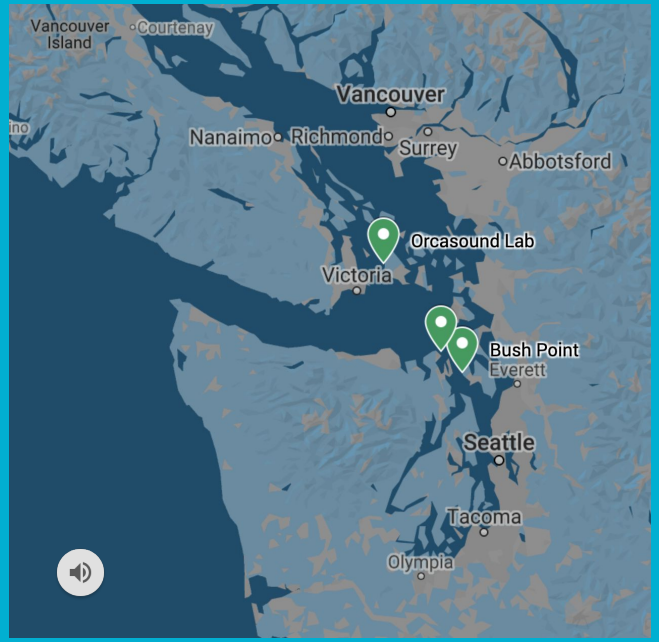

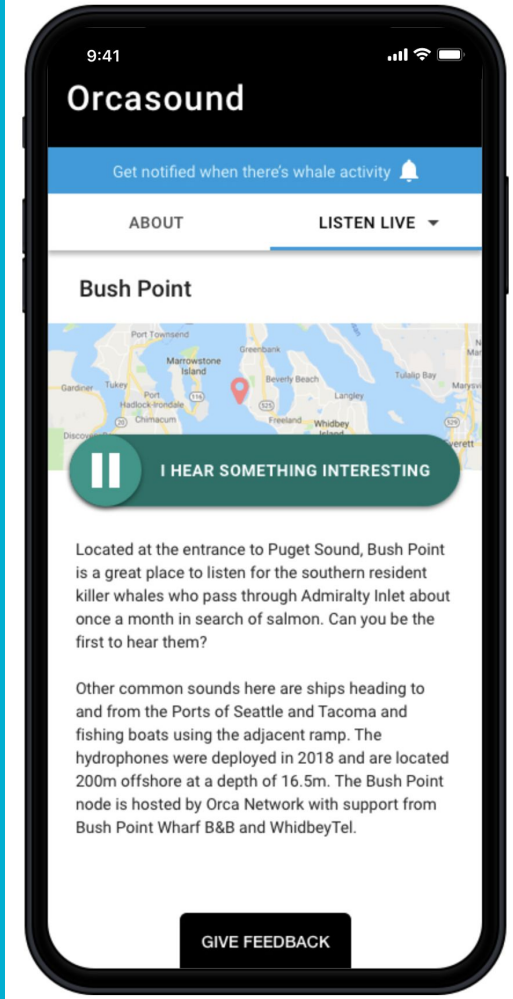

Orcasound = a hydrophone network

1970: Lime Kiln / Seasound

2002: Orcasound Lab

2008-2012 NOAA funding (expansion to 5 nodes)

2013-2015 Philanthropy only (decline to 2 nodes)

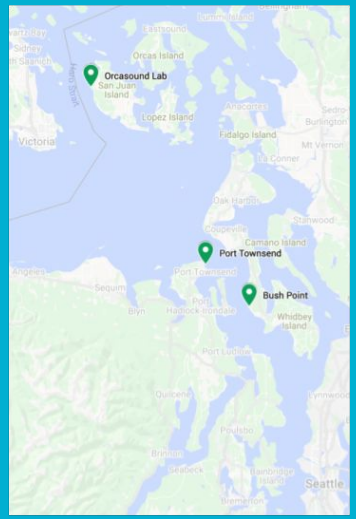

2016-2020
Cooperative agreement (now 14 NGOs)
2017 crowd-funding ($20k Kickstarter)
2018+ open-sourcing, crowd-sourcing, open data...

Orcasound = open-source soft/hard-ware

<2002-17
-- humans listening via mp3 streams
-- Val building custom software, alone
-- Scott building web static web sites, Google sheets, manual Twitter/email notifications...

2017-18: software + hardware (Kickstarter for v1 web app), live.orcasound.net launched Nov ‘18

2019-20: v2 UI beta-testing in Nov ‘19, launched in May, 2020

2021: v3 UI + a proliferation of related projects (17 Orcasound Github repositories)

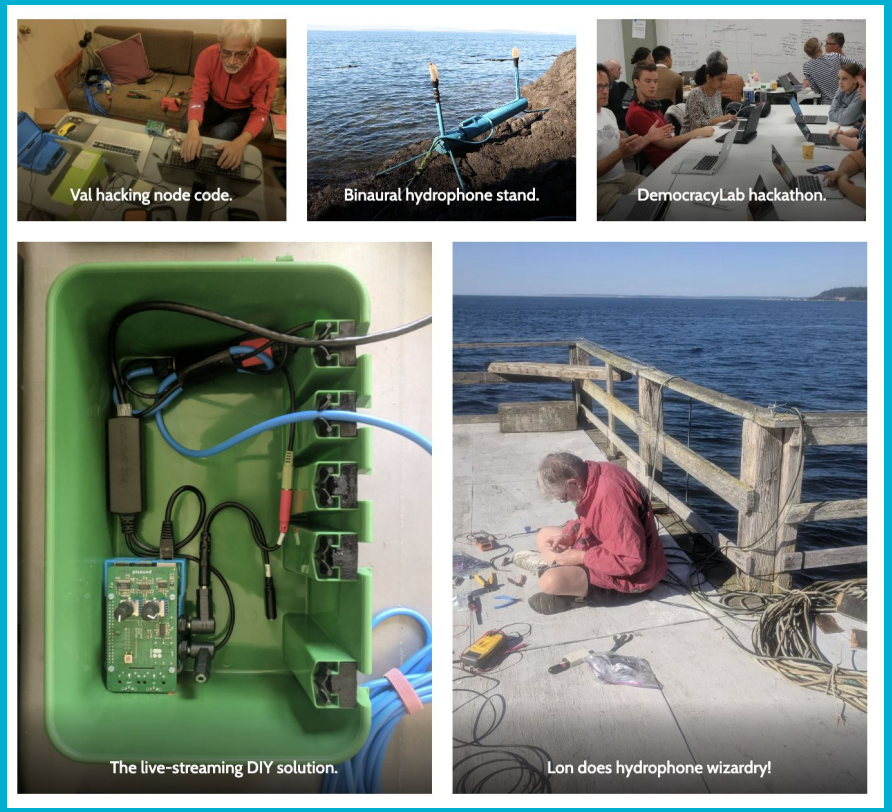

Orcasound = SRKW conservation tech

So many choices!

- Hardware, Software
- Programming language(s)
- Deployment
- Data management
- Licensing

Guiding principles:

- User-centered design
- Measuring effectiveness
- Open roadmap (Trello)

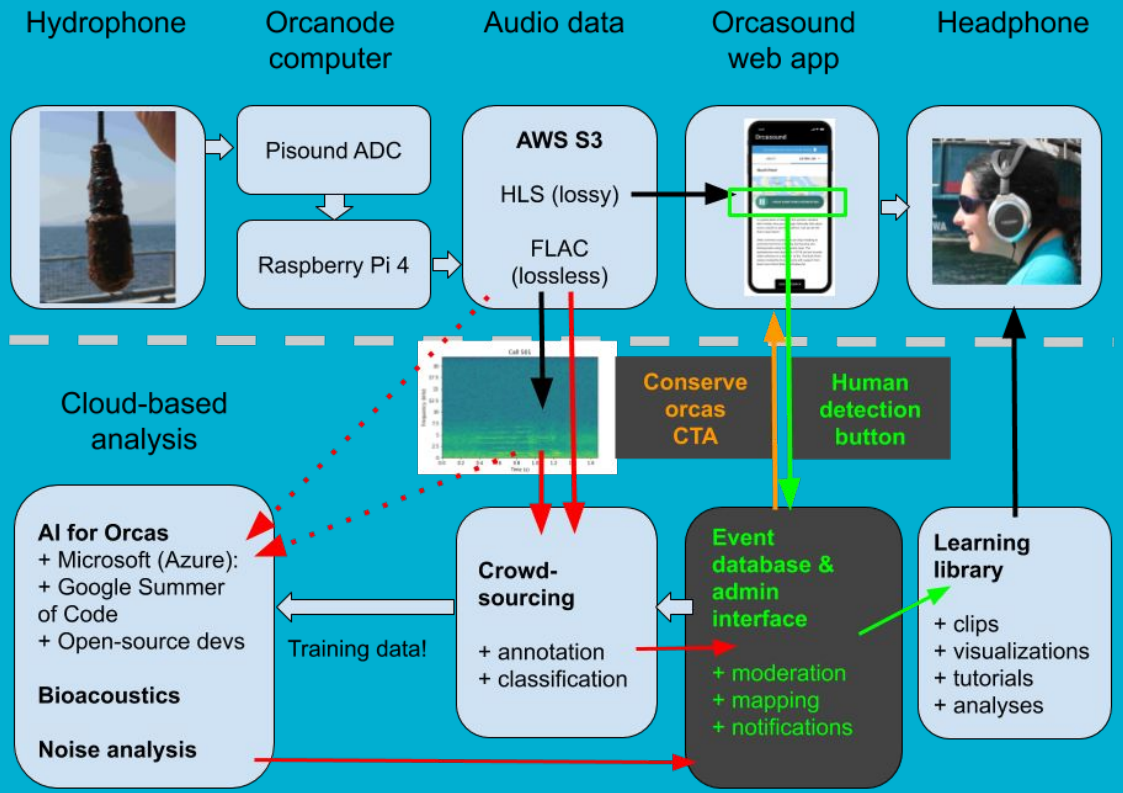

Orcasound = open raw data

```
aws --no-sign-request s3 sync streaming-orcasound-net
```

Data volume

- 2.3 TB (HLS audio only)
- 0.4 TB/yr/node
- Increasing soon! (FLAC, stereo)

Data costs

- <$100/mo w/3 nodes
- ~30$/mo/node, but scaleable
- Free with credits!

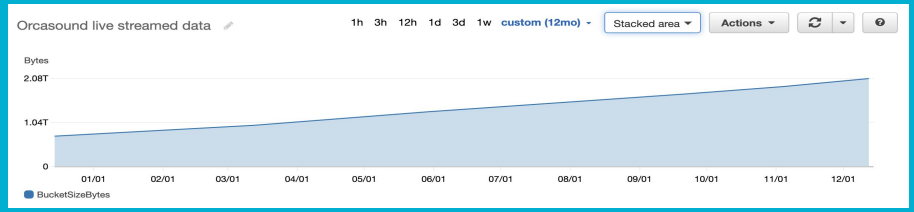

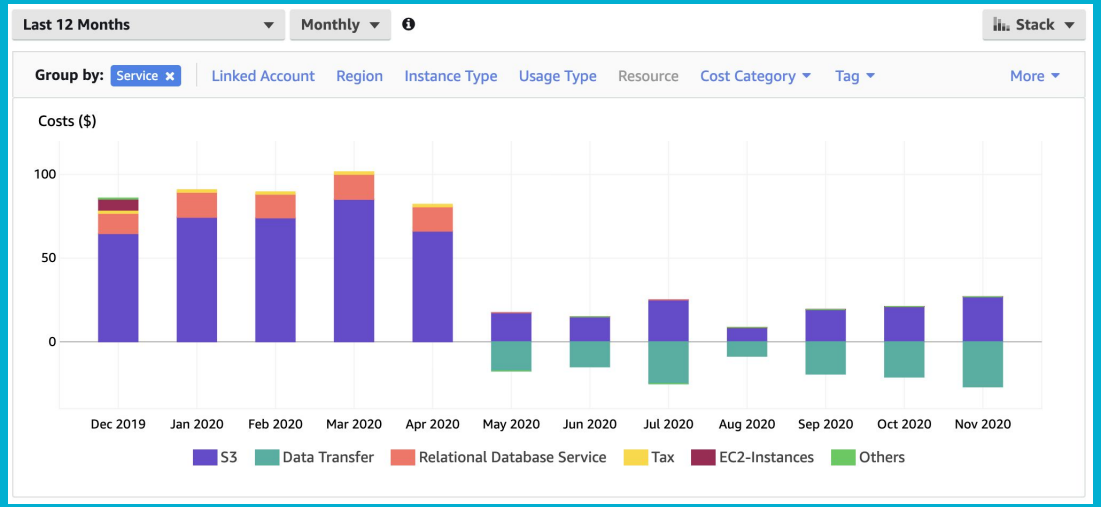

Orcasound = open labeled data

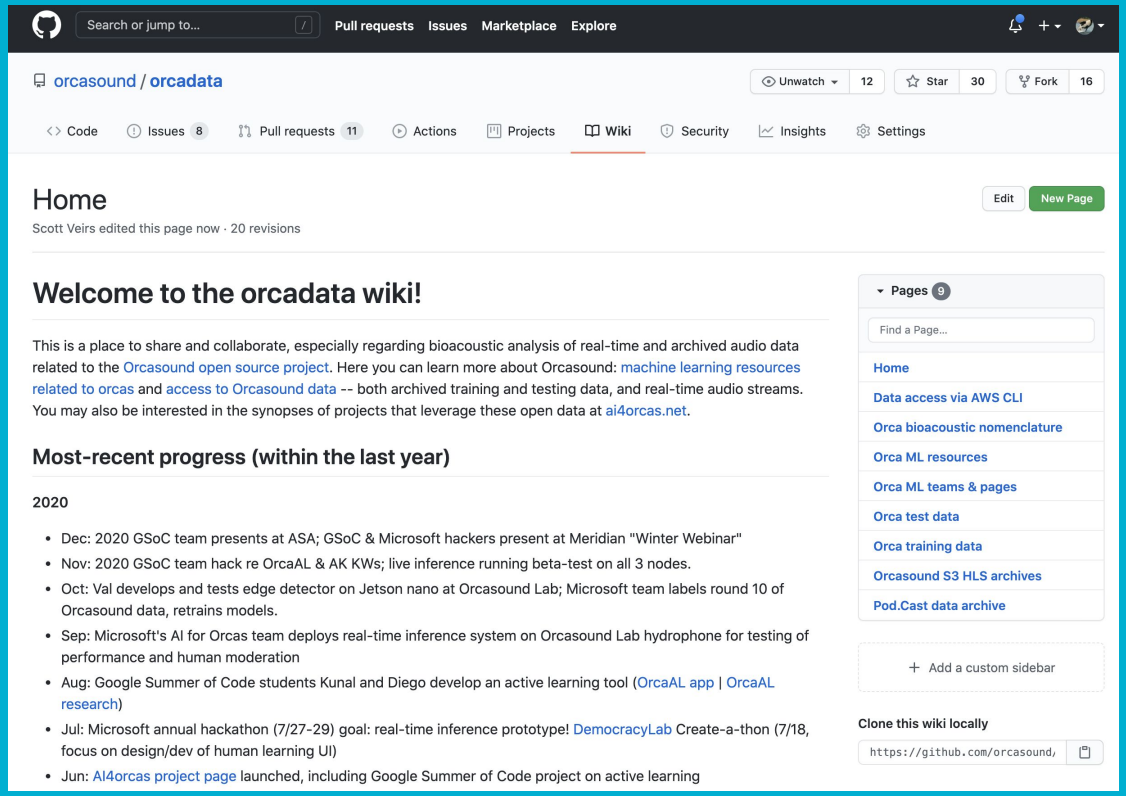

Grand challenges (human-machine synergy)

Where are they now?

Applied conservation problems

Automated (vs human-moderated) real-time notification of SRKW presence

1. Notify as soon as possible? only if you're sure it's a SRKW?
2. What is user's false positive tolerance?
3. Scott's dilemma: how to evoke empathy via a live "concert" of singing orcas without being annoying
    a. How long will it last?
    b. What will be the signal:noise?
    c. How critical is the call to action for the orcas?

What are they saying?

Basic biological problems

Bioacoustics:
-- marine mammal communication systems

-- biosonar

-- marine acoustic ecology

-- soundscape analysis

Active Listening & Learning for Orca Sound Detection

Kunal Mehta, Diego Rodriguez, Valentina Staneva, Jesse Lopez, Abhishek Singh, Val Veirs, & Scott Veirs

Presented at Acoustical Society of America

Thursday, December 10, 2020

Orcasound & Google Summer of Code

- **GSoC:** supports students & open-source software projects, like Orcasound
- **2019:** OrcaCNN for Alaskan killer whale calls
- **2020:** Build a tool to speed up the labeling of calls from the Southern Resident Killer Whales (SRKWs)

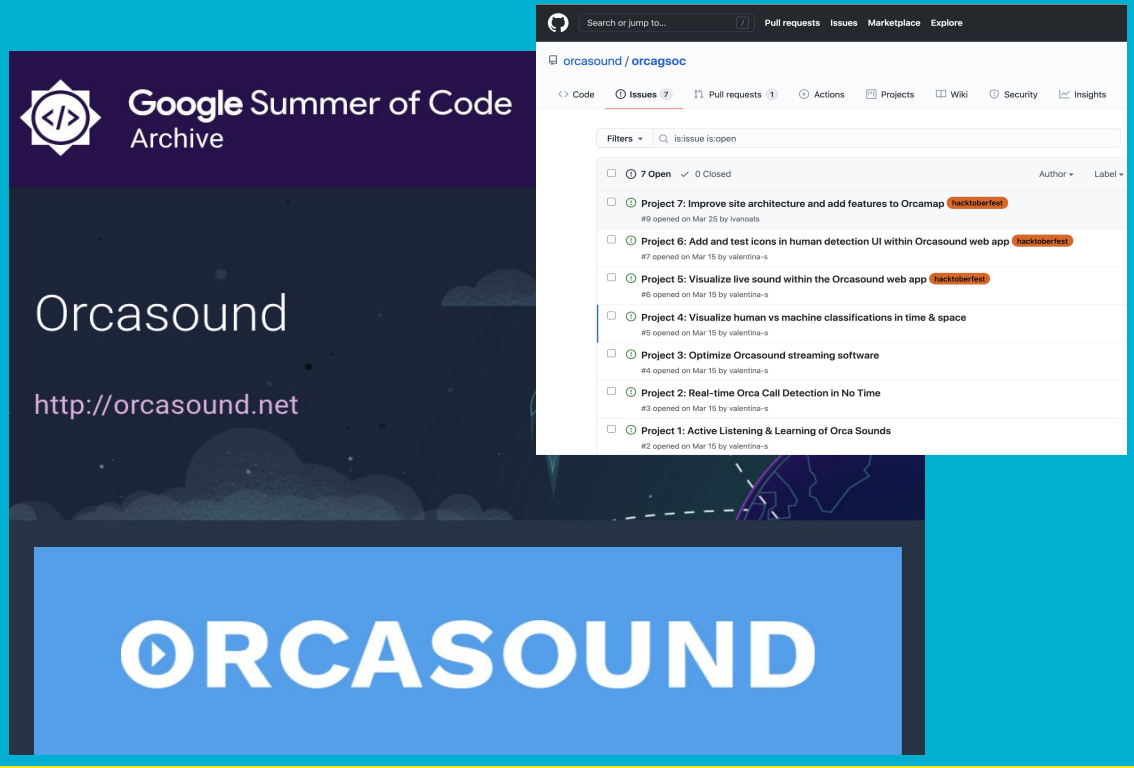

Orca call detection: open-source efforts

Open labeling tools for marine bioacoustics:

- PAMGuard
- APLOSE
- Whaledr App
- Whale-FM

Emerging open orca call machine learning models:

- Orca-SPOT (NRKWs)
- OrcaCNN (Alaskan KWs)
- Orcasound models (SRKWs)
  - University of Washington efforts
  - Microsoft hackathons
  - Other independent contributors
- More coming... (SRKWs)

Challenge: integrate efforts of human annotators & machine learning developers

Active Learning

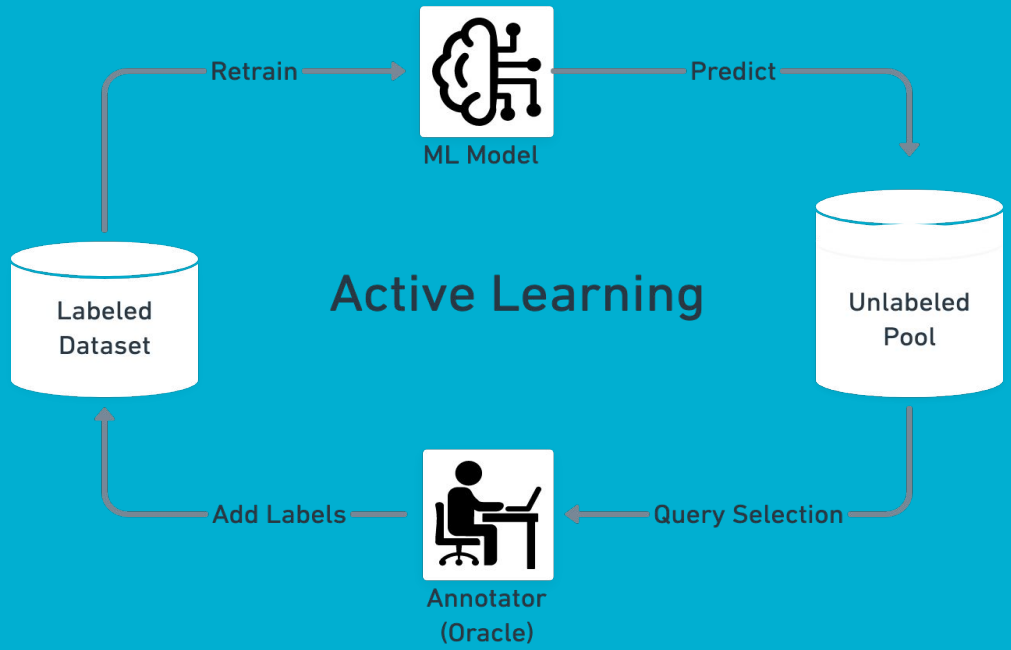

By labeling only a selective subset of samples we can save time, efforts, money, etc.

Orca active learning (OrcaAL)

- Train a **deep learning** algorithm on an initial small labeled dataset
- Use "**uncertainty sampling**" strategy to label only those samples for which the model is most uncertain (i.e. samples with confidence near 0.5)
- Integrate model training, sample selection, and annotation in an AL tool.

Motivating question: How much does labelling a subset of the samples increase accuracy?

Demo of the new citizen-science labeling tool!

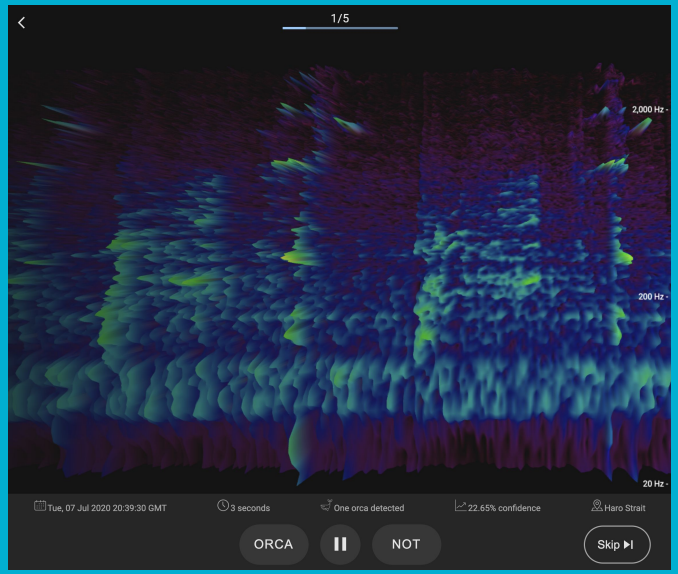

AL: uncertainty sampling

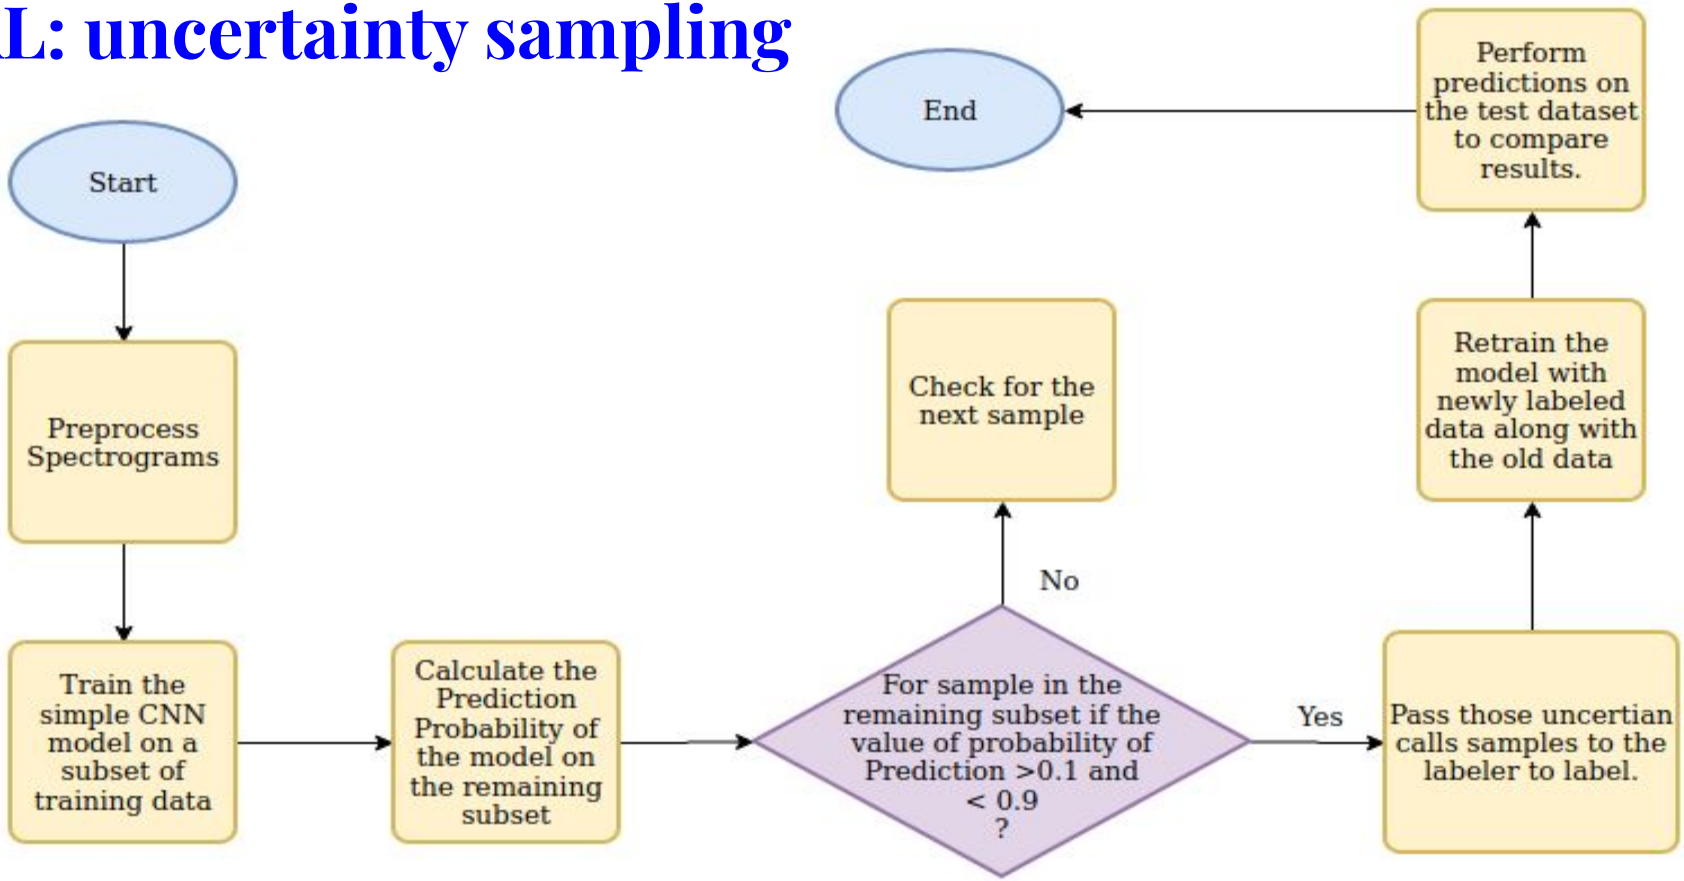

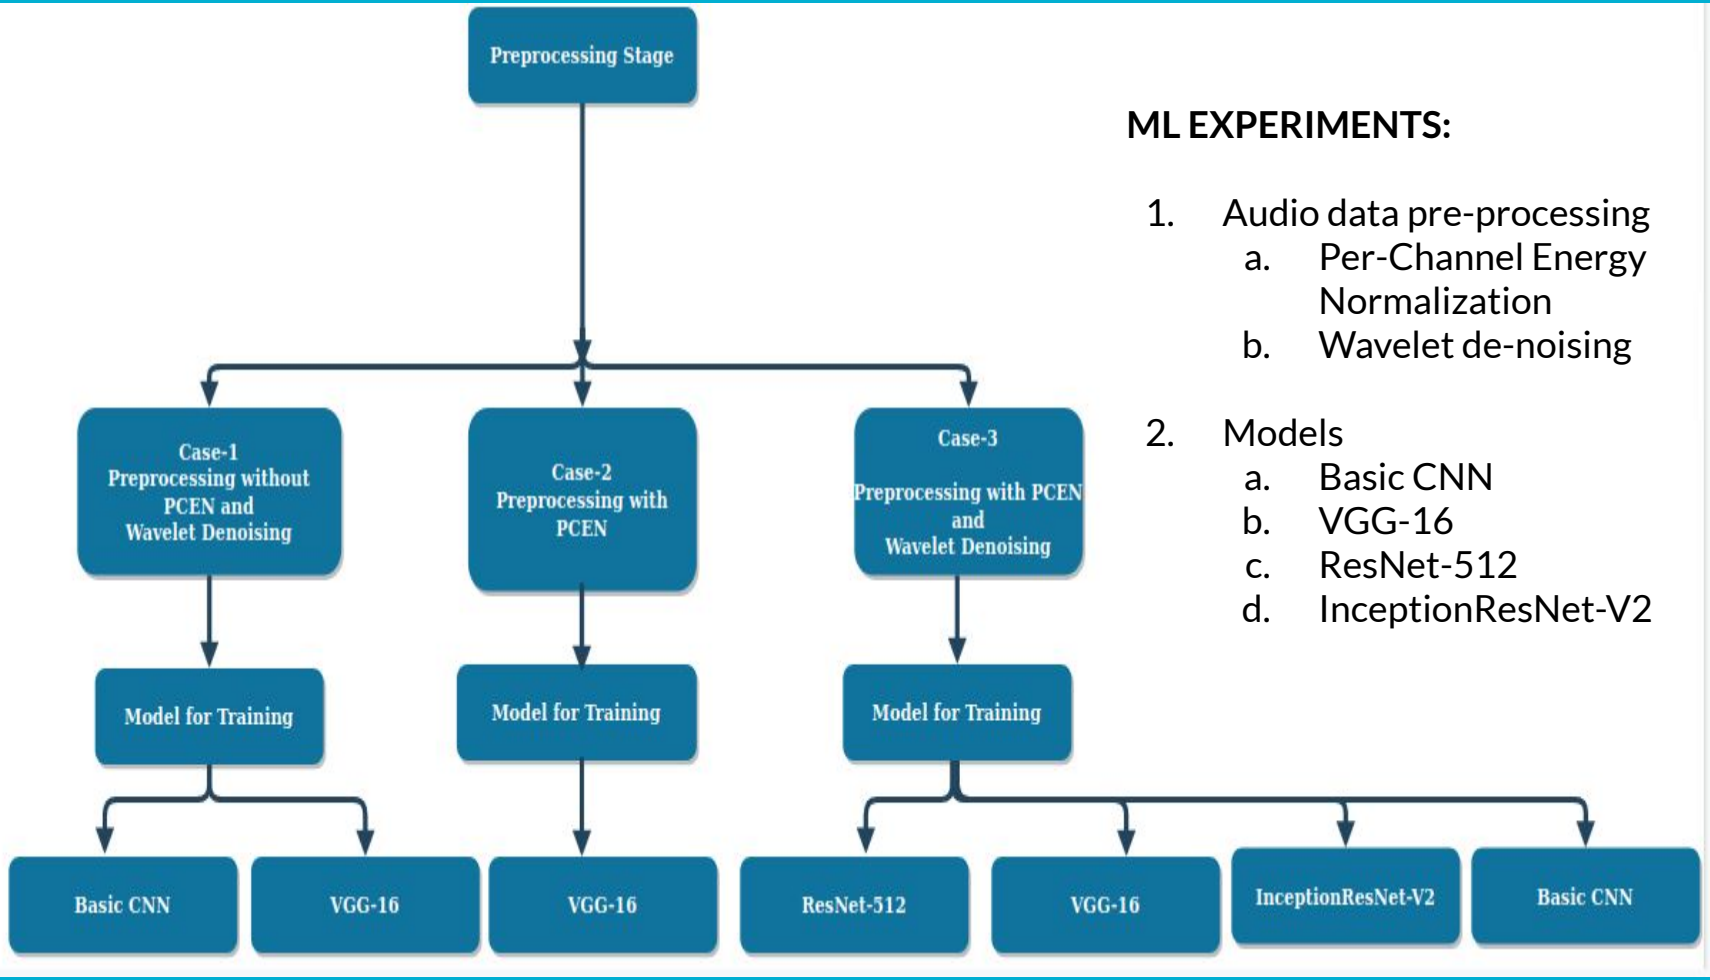

ML EXPERIMENTS:

1. Audio data pre-processing
   a. Per-Channel Energy Normalization
   b. Wavelet de-noising

2. Models
   a. Basic CNN
   b. VGG-16
   c. ResNet-512
   d. InceptionResNet-V2

Preprocessing spectrograms

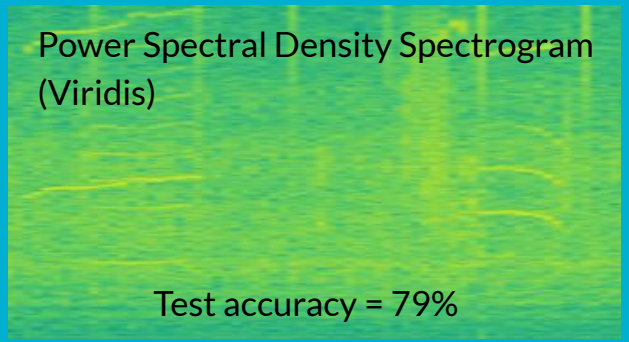

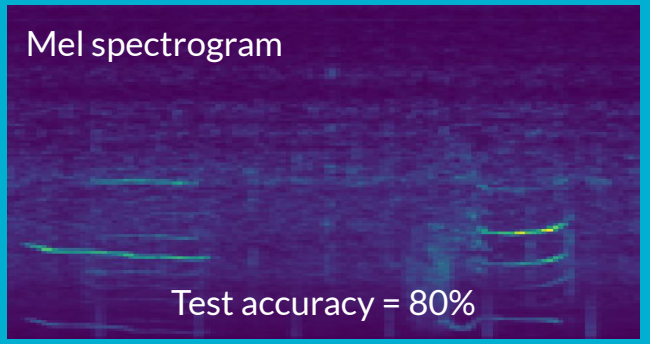

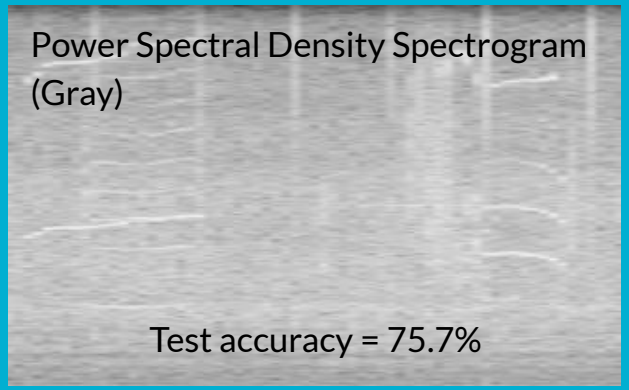

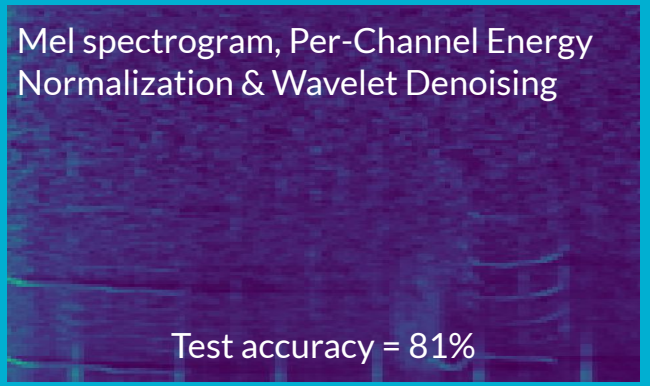

Results: model performance

AL increased accuracy by 1 %
(i.e 83.5% ---> 84.5%)

Dataset sizes:
Training: 1394 samples
Active learning: 176 samples
(37 are uncertain)
Retraining: 1570 samples
Test dataset: 201 samples

(calls & no-calls equally distributed in initial dataset)

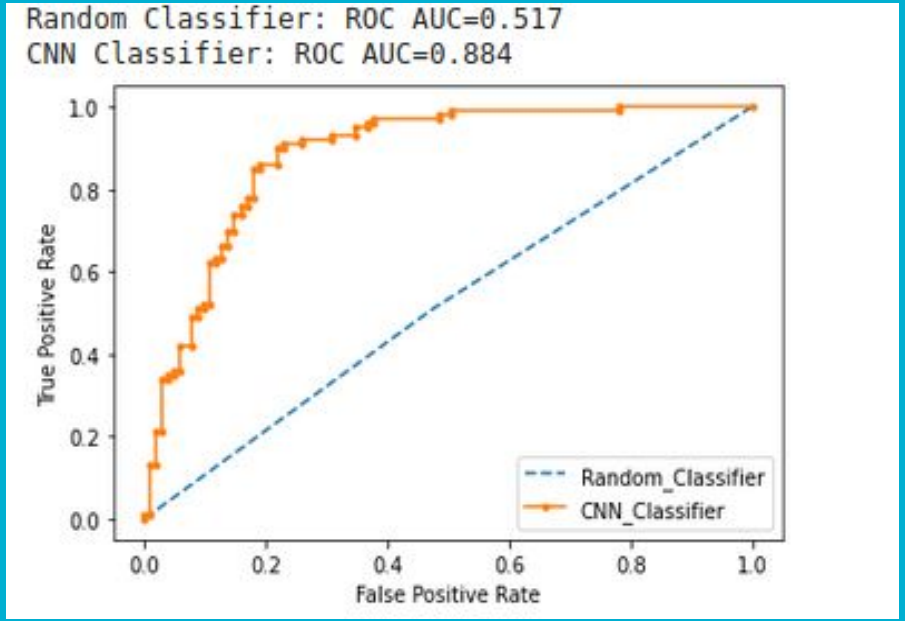

Live demo of OrcaAL

http://orcaal.ai4orcas.net/

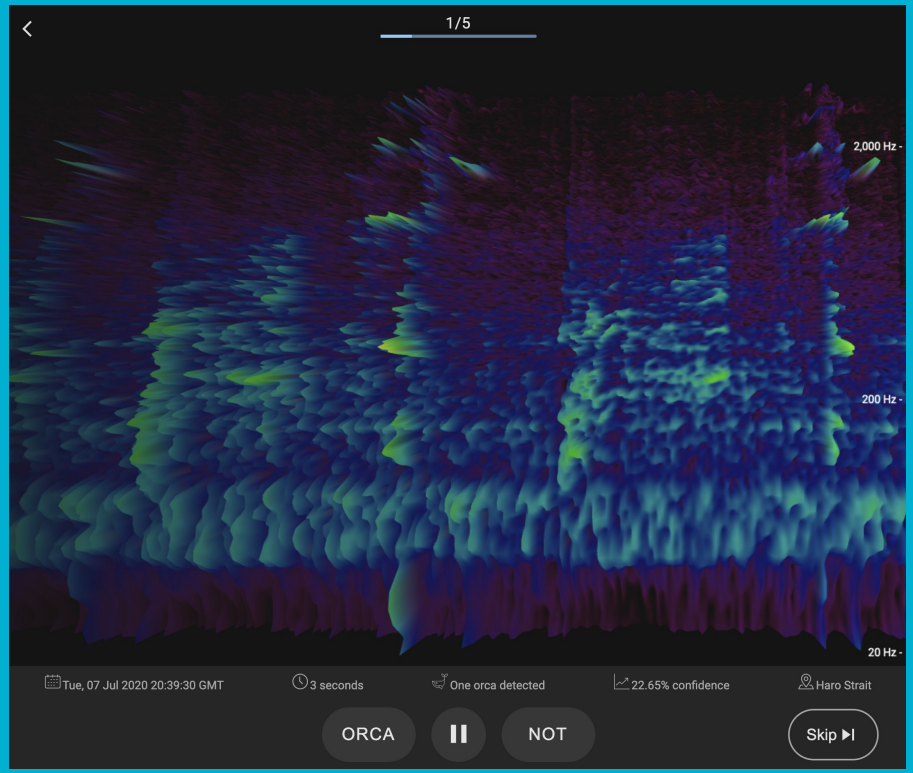

OrcaAL's architecture

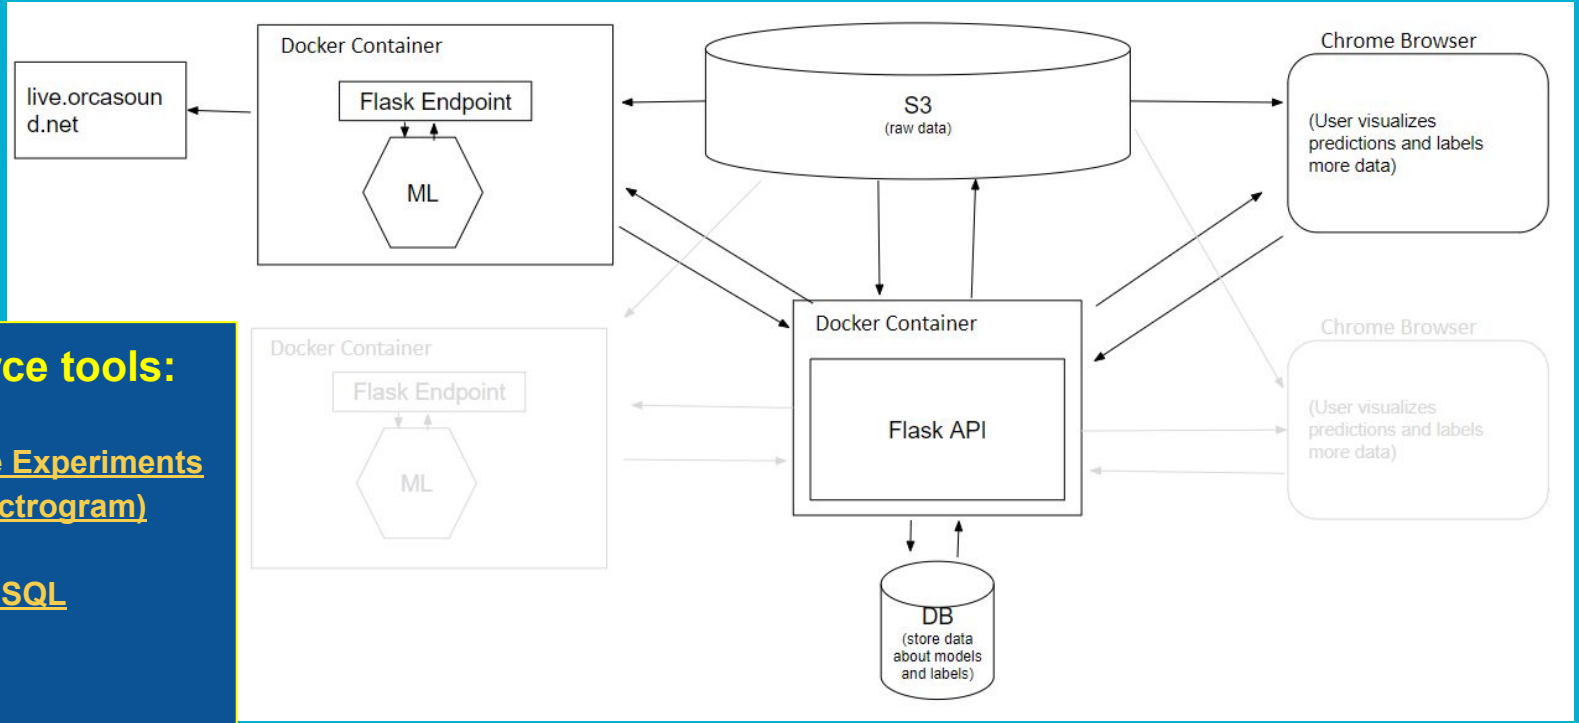

Open-source tools:

- Chrome Experiments (3D spectrogram)
- Docker
- PostgreSQL
- Flask

OrcaAL Acknowledgements & links

Thanks to all our collaborators!

- Google Summer of Code
- UW eScience Institute
- Axiom Data Science
- Beam Reach
- North Gulf Oceanic Society

Try the live demo!
orcaal.ai4orcas.net

Orcasound acknowledgements & links

Thanks to all our supporters!

- The Orcasound open-source community's volunteer hackers!
- $10k Amazon start-up credits to Beam Reach
- Microsoft AI for Earth grant of $15k Azure credits to UW eScience
- Project HALLO (2020-2022)

Join us!
ai4orcas.net