

Active Listening & Learning for Efficient and Interactive Marine Sound Annotation and Classification

Kunal Mehta, Jose Giraldo, Diego Rodriguez,
Jesse Lopez, Val Veirs, Scott Veirs,
Valentina Staneva

vms16@uw.edu

9th DCLDE Workshop,

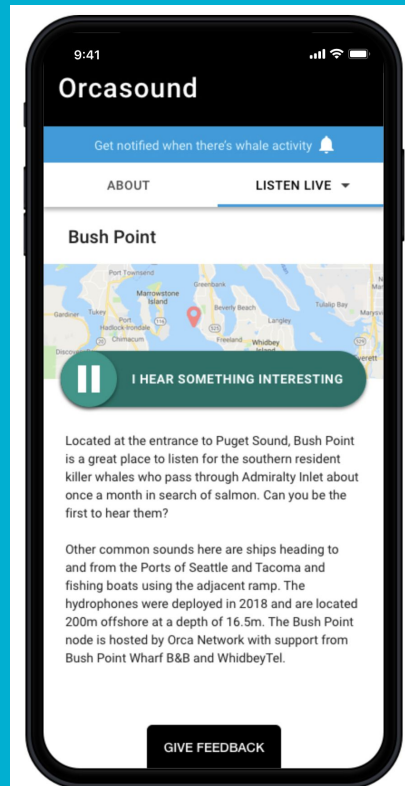

March 8, 2022



ORCASOUND

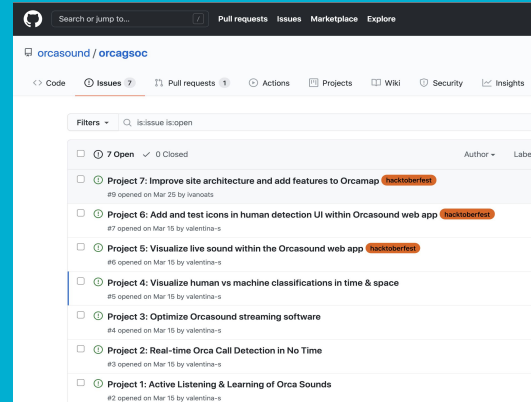
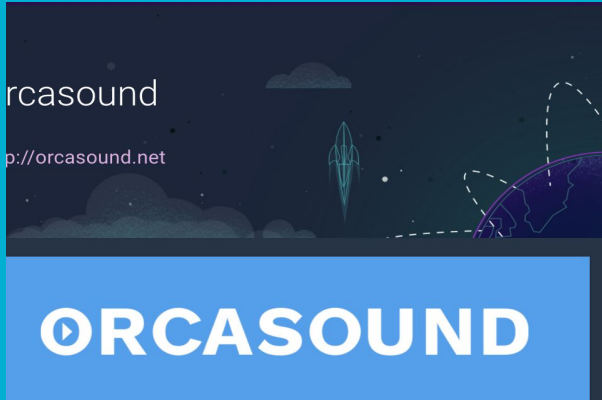
A hydrophone network (WA, USA) & open-source software community

- 3 cabled nearshore sites streaming 24/7
- 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 & Google Summer of Code

- GSoC: supports students & open-source software projects, like Orcasound
- 2020: Build a tool to speed up the labeling of calls from the Southern Resident Killer Whales (SRKWs)
- 2021: Sampling diverse sounds through open embeddings



Orca call detection: open-source efforts

Open labeling tools for marine bioacoustics:

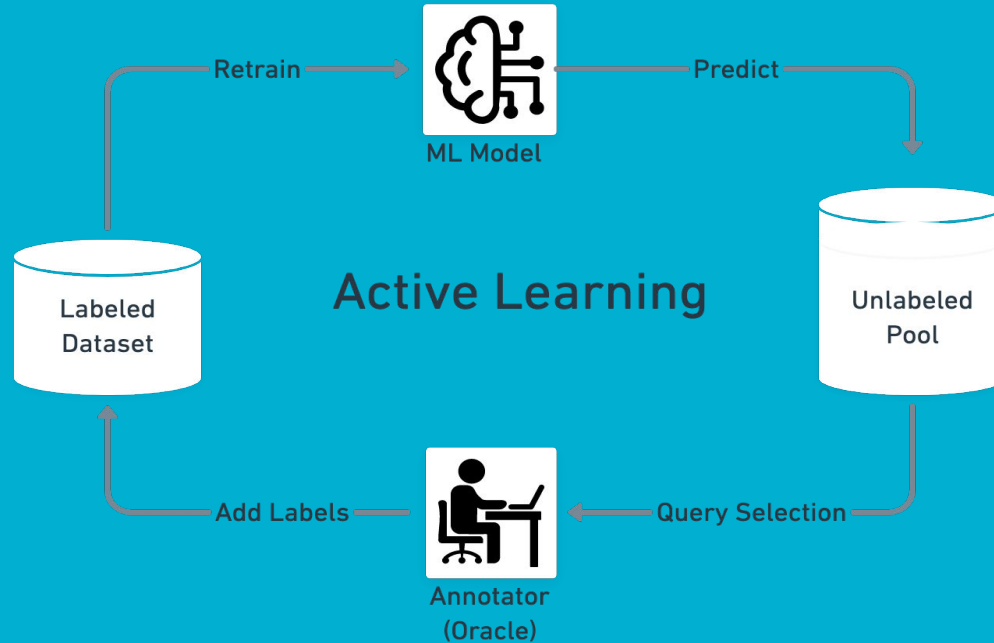
- [PAMGuard](#)
- [APLOSE](#)
- [Whaledr App](#)
- [Whale-FM](#)
- ...

Emerging open orca call machine learning models:

- [Orca-SPOT](#) (NRKWs)
- [OrcaCNN](#) (Alaskan KWs)
- [Orcasound models](#) (SRKWs)
- More coming... (SRKWs)

Goal: integrate efforts of human annotators & machine learning developers

Active Learning

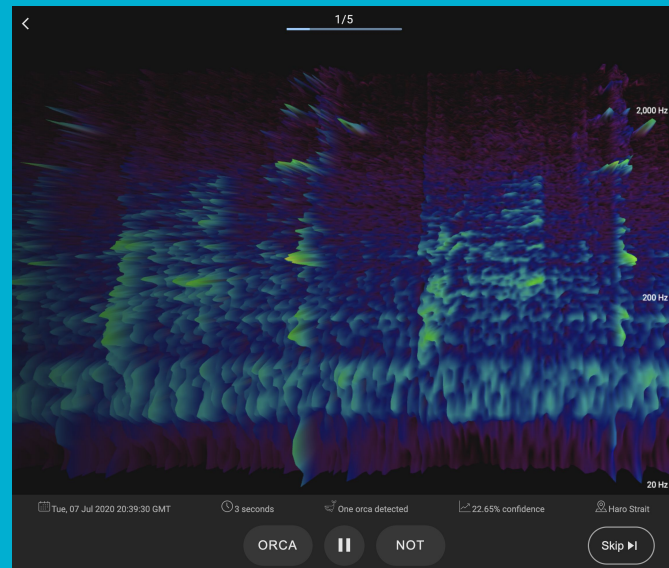


By labeling only a selective subset of samples we can save time, efforts, money, etc.

Orca active learning (OrcaAL)

- Train a **Convolutional Neural Network** algorithm on an initial small labeled dataset
- Select a new sample to label for which the algorithm is **uncertain**
- Integrate model training, sample selection, and annotation in an AL tool

Demo in the end!



Uncertainty Sampling

Least Confidence Sampling:

- apply current algorithm to unlabeled observations
- calculate confidence score for each prediction
- select to label n observations with least confidence




Image from “Human-In-The-Loop Machine Learning”, Monarch, R.

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)

Random Classifier: ROC AUC=0.517

CNN Classifier: ROC AUC=0.884



Diversity Sampling

Samples with same confidence score may look similar, and may not enrich the training set.

Select samples which are most different from the labeled set.

$$diversity(x) = 1 - \frac{1}{L} \sum_{l=1}^L sim(x, x_l)$$

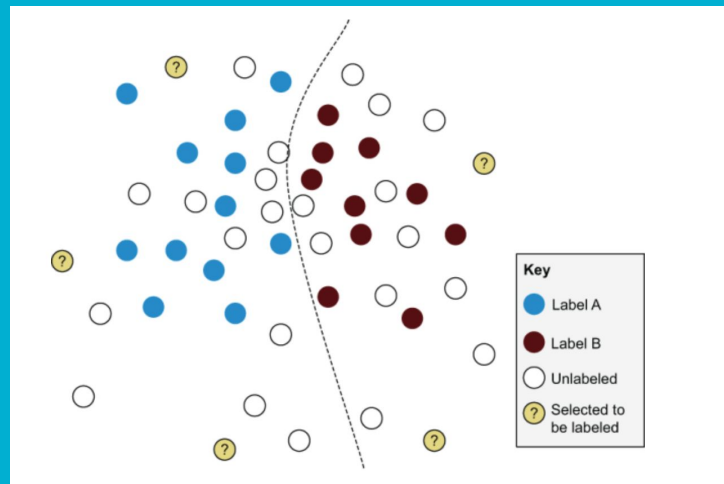


Image from “Human-In-The-Loop Machine Learning”, Monarch, R.

Lonely Outliers!

Density Sampling

Select observations not far from all other observations in the **unlabeled** set.

$$density(x) = 1/U \sum_{u=1}^U sim(x, x_u)$$

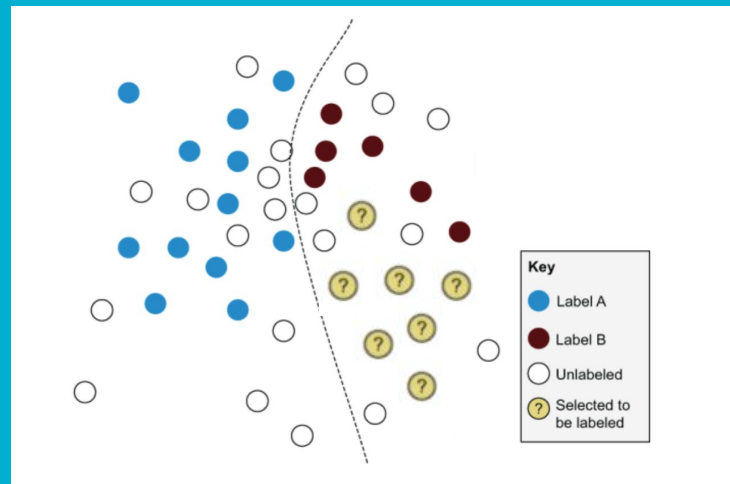
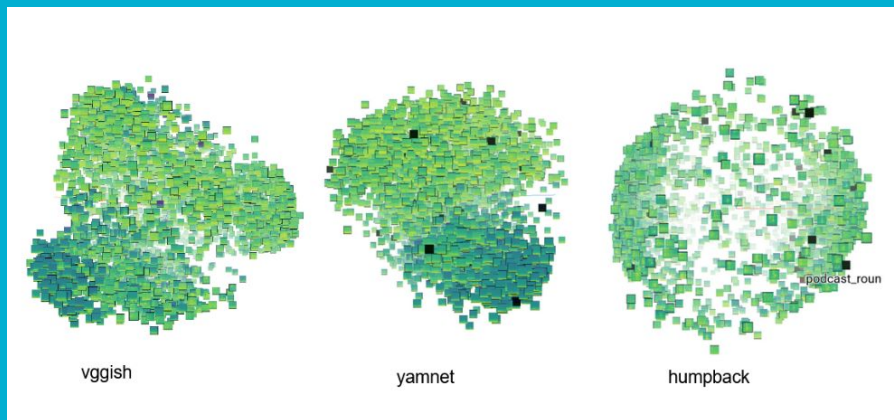


Image adapted from “Human-In-The-Loop Machine Learning”, Monarch, R.

Sound Embeddings

Apply existing open pretrained models (on [TensorFlow Hub](#)) to reduce the dimensionality of the orca sounds:


- [VGGish](#): trained on YouTube videos
- [YAMnet](#): trained on YouTube videos
- [Humpback whale](#): trained on Humpback songs



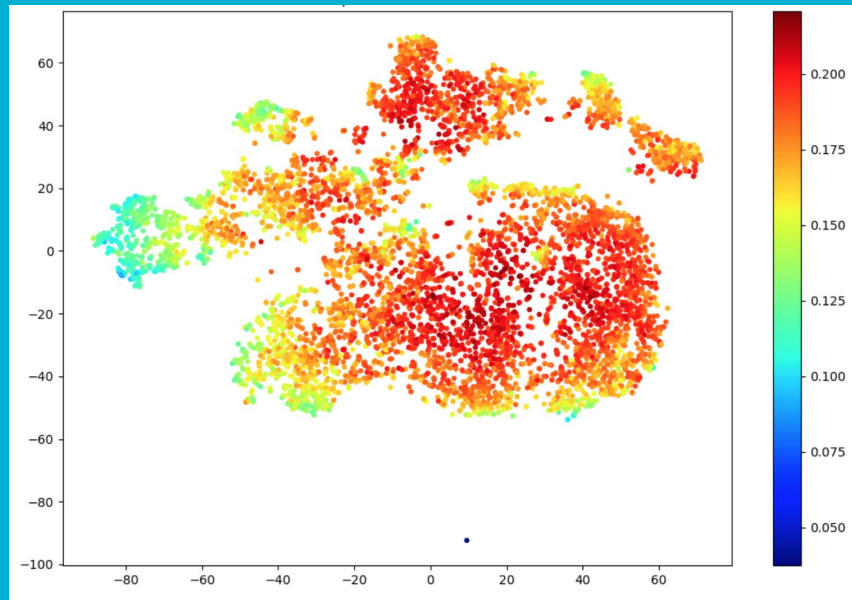
➤ Similar sounds get grouped together!

Uncertainty vs Density Sampling

Uncertainty




Density



Demo of OrcaAL

<http://orcaal.ai4orcas.net/>



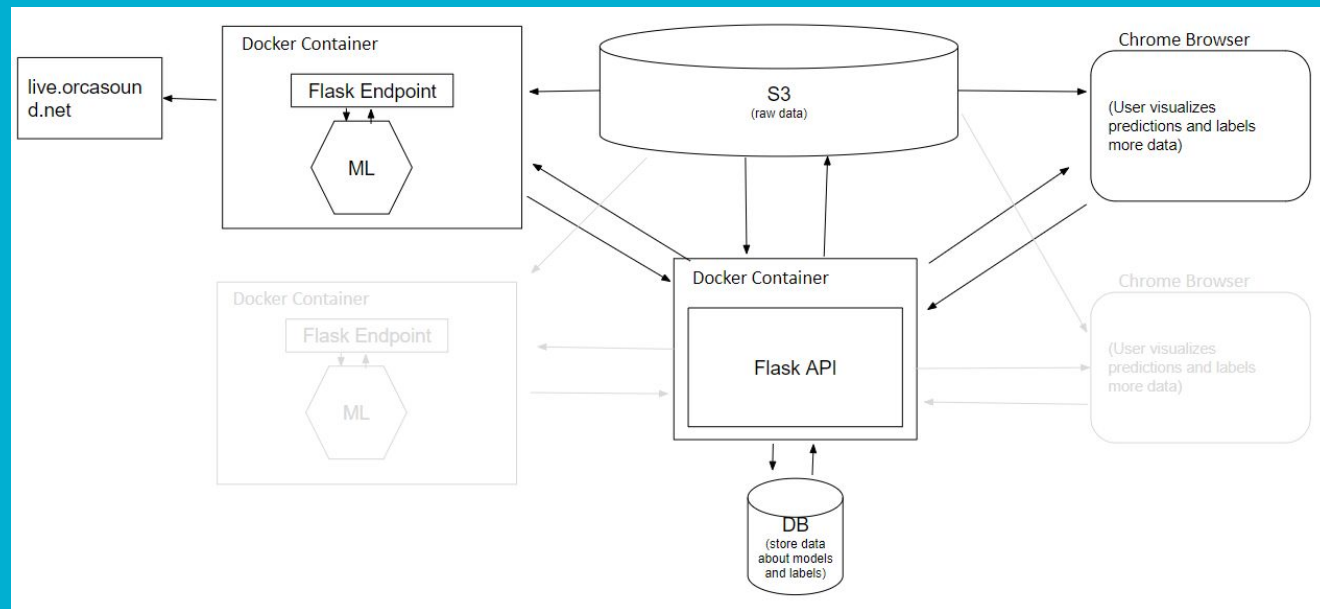
OrcaAL's architecture

Built with modularity in mind.



Open-source tools:

- [Chrome Experiments \(3D spectrogram\)](#)
- [Docker](#)
- [PostgreSQL](#)
- [Flask](#)



Acknowledgements & links

Thanks to all our collaborators!

- [Google Summer of Code](#)
- [UW eScience Institute](#)
- [Axiom Data Science](#)
- [Beam Reach](#)
- [North Gulf Oceanic Society](#)
- The Orcasound open-source community's volunteer hackers!
- \$10k [Amazon start-up credits](#) to Beam Reach

More info:

- [orcasound.net](#)
- [ai4orcas.net](#)
- [github.com/orcasound](#)

Try the live demo!

orcaal.ai4orcas.net

