The Ketos Package

Fabio Frazao

MERIDIAN, Institute for Big Data Analytics,
Dalhousie University, Halifax, Canada

Introduction

Deep learning can ...

- help us create better acoustic detection and classification (DC) tools
- change the way we develop and use these tools

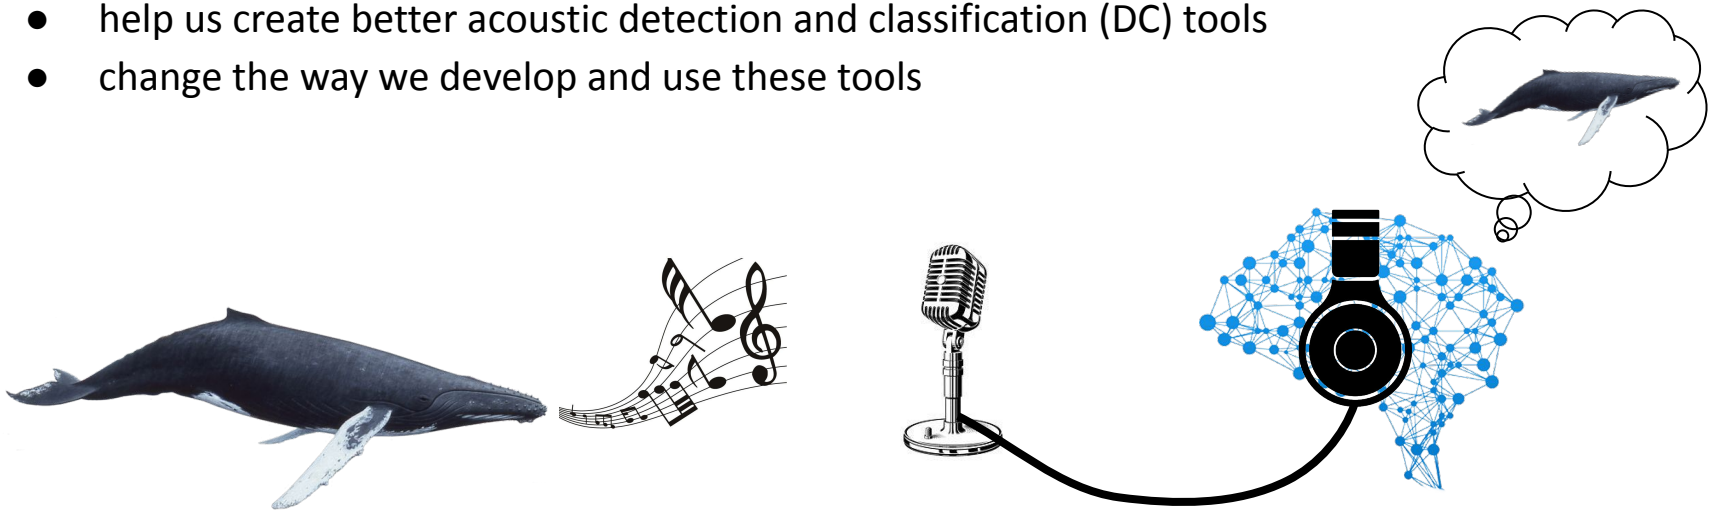

Deep learning in marine bioacoustics

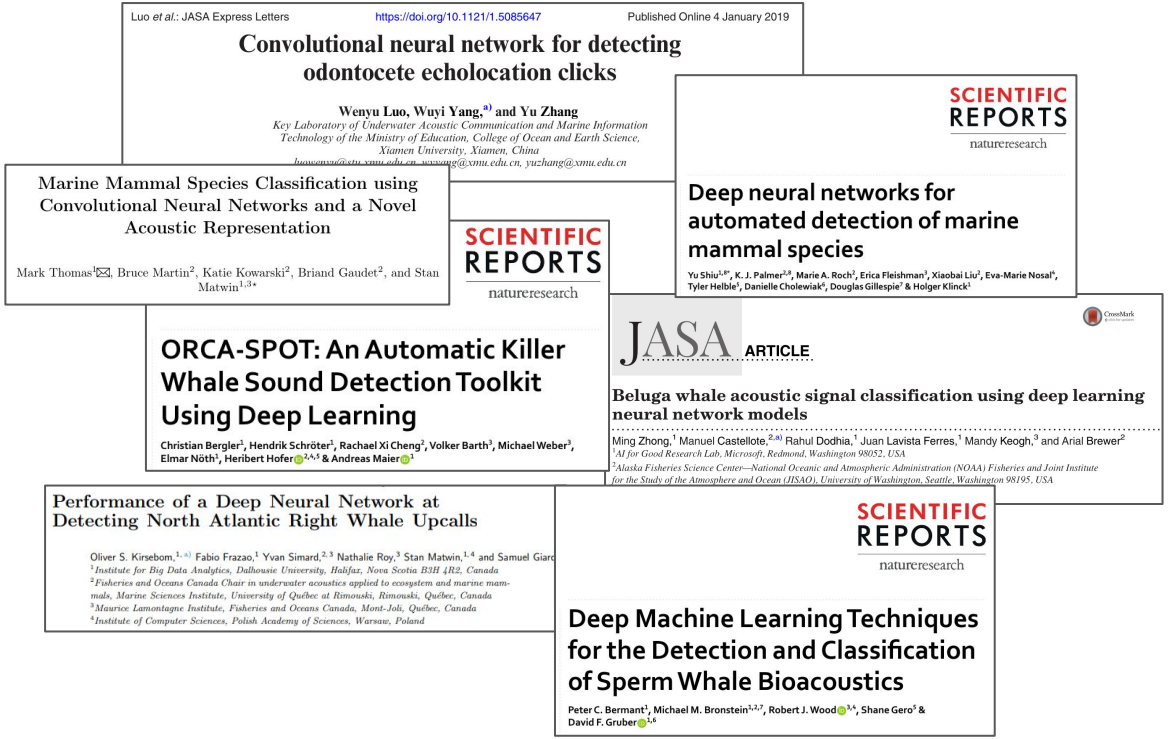

- Deep learning works, is now a mature technique
- Often outperforms more conventional DC algorithms
- However, these algorithms are not always accessible to marine acousticians (or application developers supporting marine scientists)

Ketos - at a glance

Neural network architectures

Audio processing

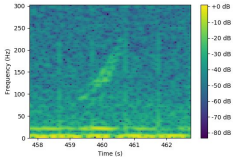

	A	B	C	D
1	filename	label	start	end
2	IML-BA_Sample_2019-06-02_025209.wav	1	0	3
3	IML-BA_Sample_2019-06-02_065235.wav	1	0	3
4	IML-BA_Sample_2019-06-08_012119.wav	1	0	3
5	IML-BA_Sample_2019-06-21_201729.wav	1	0	3
6	IML-BA_Sample_2019-06-29_124902.wav	1	0	3
7	VAS_Sample_2019-06-06_052058.wav	1	0	3
8	VAS_Sample_2019-06-06_052147.wav	1	0	3
9	VAS_Sample_2019-06-06_052545.wav	1	0	3

Annotation tables

Training Database

Training methods

Trained DC model

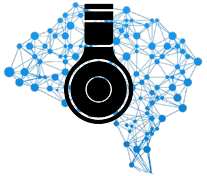

Ketos - at a glance

Neural network architectures

- Standard Convolutional networks (AlexNet, VGG,etc)
- ResNets (including 1D/temporal version options)
- DenseNets
- Inception networks
- More coming soon (sequential and segmentation architectures)

Audio processing

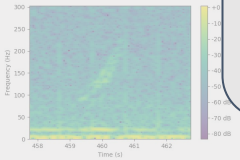

	A	B	C	D
1	filename	label	start	end
2	IML-BA_Sample_2019-06-02_025209.wav	1	0	3
3	IML-BA_Sample_2019-06-02_065235.wav	1	0	3
4	IML-BA_Sample_2019-06-08_012119.wav	1	0	3
5	IML-BA_Sample_2019-06-21_201729.wav	1	0	3
6	IML-BA_Sample_2019-06-29_124902.wav	1	0	3
7	VAS_Sample_2019-06-06_052058.wav	1	0	3
8	VAS_Sample_2019-06-06_052147.wav	1	0	3
9	VAS_Sample_2019-06-06_052545.wav	1	0	3

Annotation tables

Training Database

Training methods

Trained DC model

Ketos - at a glance

copyleft license

- Written in Python
- GNU GPLv3 license - freely available to use and modify
- Hosted on GitLab: https://gitlab.meridian.cs.dal.ca/public_projects/ketos
- Documented code, including examples: https://docs.meridian.cs.dal.ca/ketos/
- Tutorials, version history, and more ...
- Available on the Python Package Index (PyPi) - the official third-party software repository for Python

Built on top of open-source libraries

1. Audio data manipulation & processing

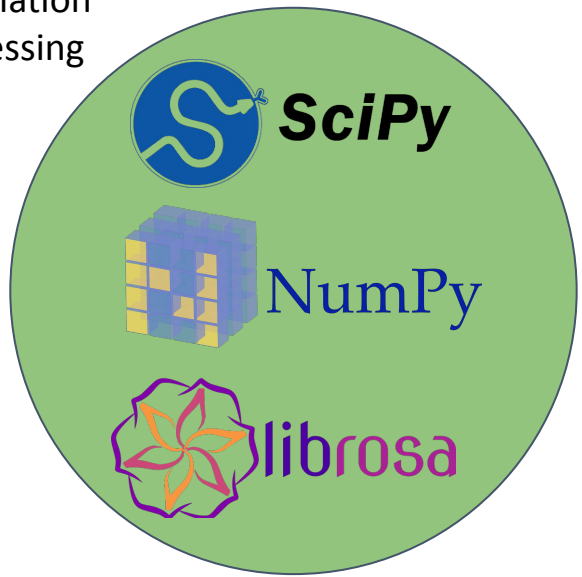

2. Data handling & storage

3. Machine learning, neural networks

Documentation - https://docs.meridian.cs.dal.ca/ketos/

ketos

Underwater acoustic detection and classification with deep neural networks

2.2

Search

- Introduction
- Installation
- Tutorials
  - Creating a training database (Basic)
  - Creating a training database (Extended)
  - Training a ResNet classifier
  - Creating a detector
- Modules
- How to contribute
- Versions

Docs » Tutorials » Tutorial: Training a Binary ResNet Classifier

View page source

Tutorial: Training a Binary ResNet Classifier

Note

You can download an executable version (Jupyter Notebook) of this tutorial and the data required to follow along here.

North Atlantic Right Whale detector-part 1

This is the first of a two parts tutorial illustrating how to build a deep learning acoustic detector with ketos.

We'll use the database built in the Creating a training database tutorial, in which we converted raw audio files to spectrograms of the North Atlantic Right Whale's stereotypical upcall. If you didn't follow that tutorial, you can find the resulting database in the .zip file linked at the top of this page. There you will also find an executable version of this jupyter notebook, in case you want to follow along.

Our final goal is to have a detector that can take a long .wav file (e.g.: 30 min) and tell us where within that file are the right whales upcalls.

The core part of such detector will be a binary classifer that takes 3-s long spectrograms and classifies them into "contains an upcall" or "does not contain an upcall". We will treat these two classes as "1" and "0". This is what we'll cover in this tutorial.

The second part will take this binary classifier and turn it into a detector.

Contents:

1. Importing the packages
2. Creating the data feed
3. Creating and training the Neural Network

The lines below define the random seeds used in the tutorial. This is necessary to ensure that you get the precisely the same results every time you run the code.

Typical Ketos users

Acoustic data analyst

Application developer

Deep learning developer

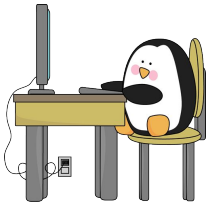

Deep learning researcher

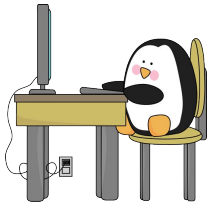

Typical Ketos users

Acoustic data analyst

Application developer

Deep learning developer

Deep learning researcher

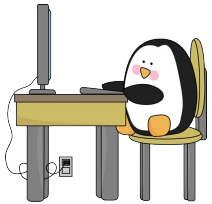

- Has domain expertise
- Has some basic programming experience (if using ketos directly)
- Interested the application of detectors/classifiers to their data
- Mostly follows the default routes/uses pre-trained models

Typical Ketos users

Acoustic data analyst

Application developer

Deep learning developer

Deep learning researcher

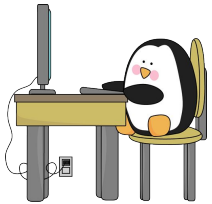

- Experienced software developer
- Not (necessarily) experienced in machine learning/data analysis
- Interested in developing applications (web, desktop,etc) around trained detectors/classifiers.
- The product of their work is used by data analysts (allowing them to benefit from ketos indirectly)

Typical Ketos users

Acoustic data analyst

Application developer

Deep learning developer

Deep learning researcher

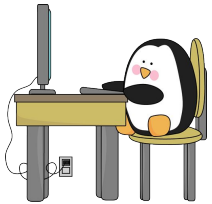

- Has experience with data analysis and machine learning
- Is interested in applied deep learning: wants to build models that work for a given acoustic application
- Not interested in developing new machine learning methods

Typical Ketos users

Acoustic data analyst

Application developer

Deep learning developer

Deep learning researcher

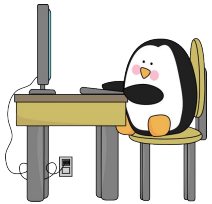

- Has expertise with data analysis and machine learning
- Is interested in developing new machine learning methods (for acoustics)

Interfaces

Ketos can be used with different interfaces:

- Scripts
- Command-line interfaces
- Jupyter notebooks
- In the backend of web/desktop applications
- Through other applications/frameworks compatible with the exported models

Interfaces

Command-line interface for headless operation onboard a floating data collection station

Raspberry Pi 3B

+

DFO/ Maurice Lamontagne Institute (MLI)

Interfaces

A web application running Ketos in the backend

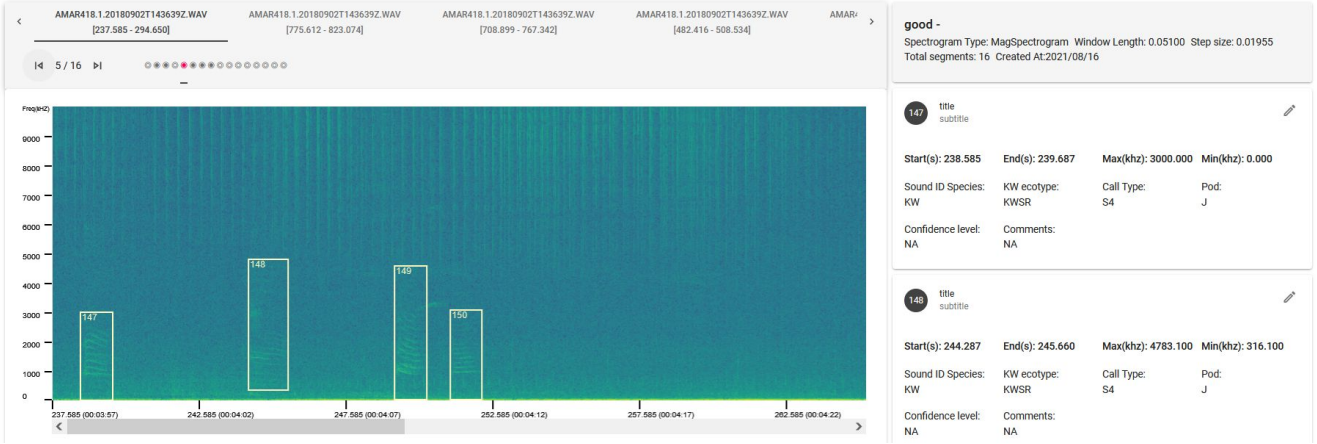

Interfaces

Exported Ketos model being used by PAMGuard

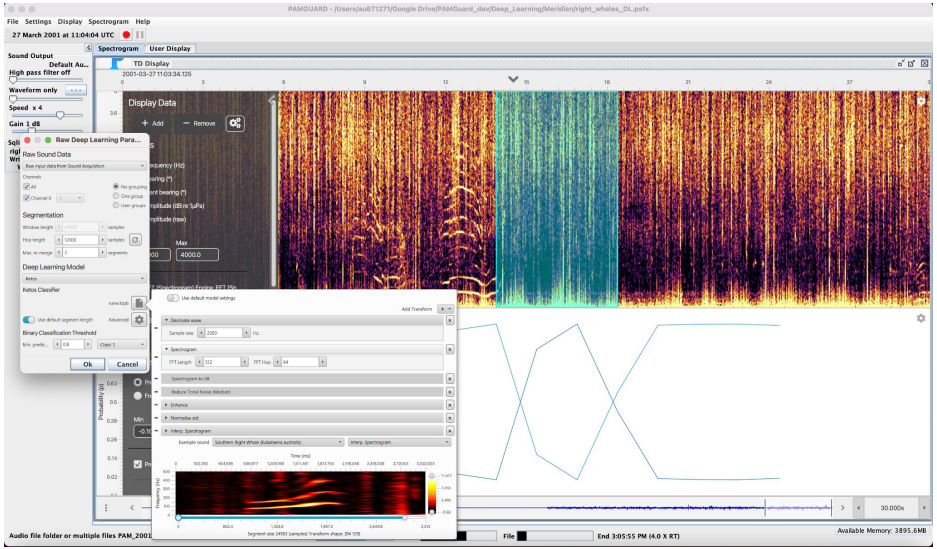

Read more about the new PAMGuard deep learning module here: https://conservationcoding.com/2021/06/07/deep-learning-in-pamguard/

Summary

Deep learning can ...

- help us create better acoustic detection and classification (DC) models
- change the way we develop and use these models

Ketos ...

- is an open-source Python package for developing deep learning based acoustic detectors and classifiers
- provides neural network architectures, transfer learning capabilities, tools for dealing with larger-than-memory datasets, audio processing, saving and sharing of models, and more
- is accessible to different kinds of users
- Can be used with different interfaces
- offers documentation and step-by-step tutorials
- check it out at https://docs.meridian.cs.dal.ca/ketos/

Thank you!