The Ketos Package

Oliver Kirsebom
Fabio Frazao
Bruno Padovese

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

Introduction

Deep learning can ...

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

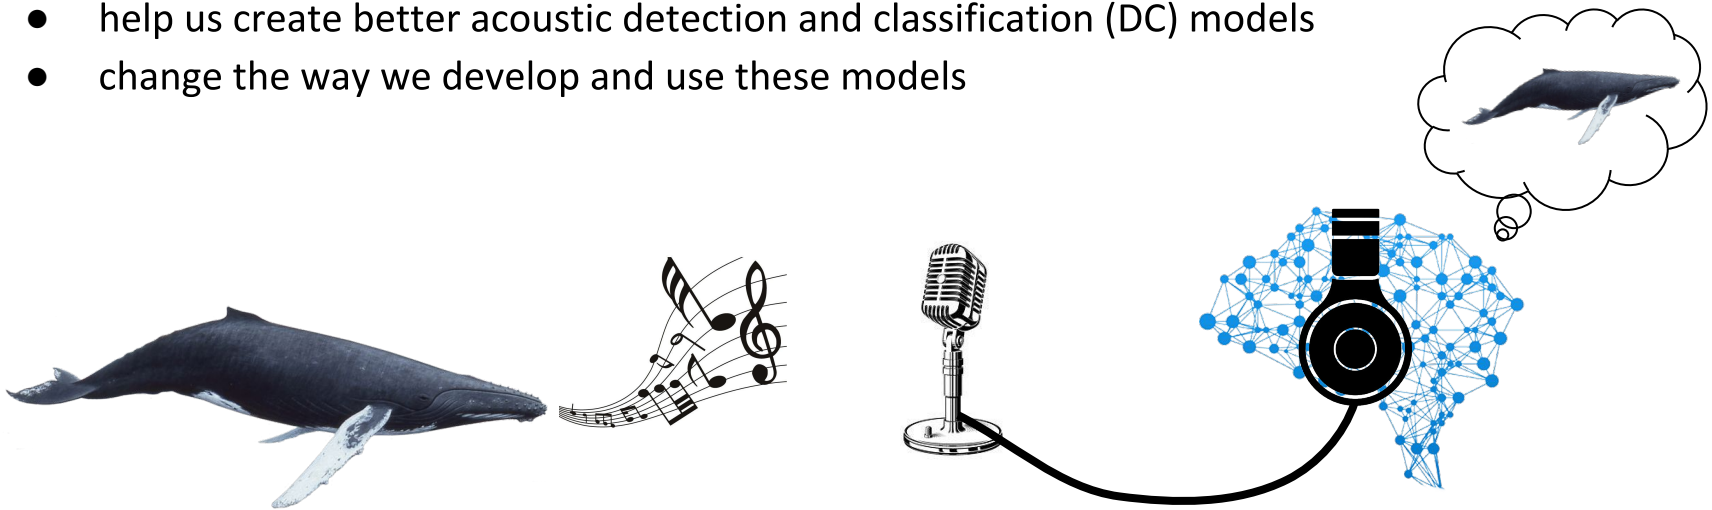

Deep learning - everyday examples

Deep Learning,
Deep Neural Networks,
Artificial Intelligence (AI) ...

Speech recognition/ synthesis

Facial Recognition

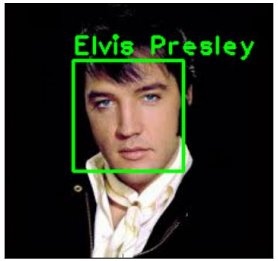

Translation

Object detection

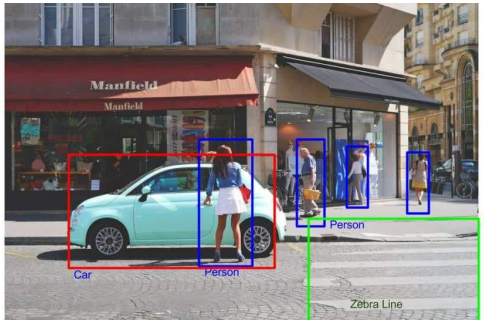

Deep learning in marine bioacoustics

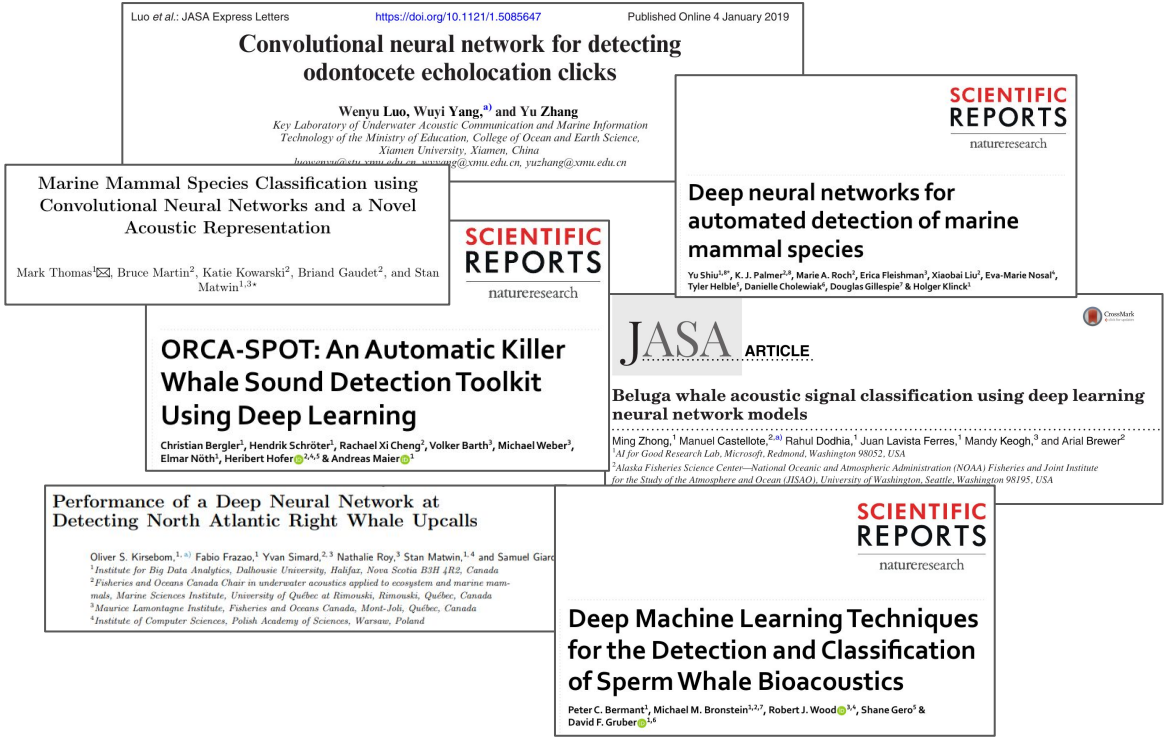

- Deep learning works, is now a mature technique
- Outperforms conventional DC algorithms
- It's time to build tools and applications that make these algorithms accessible to marine bioacousticians

Deep learning in a nutshell

Deep Learning aims to be end-to-end (but usually isn't)

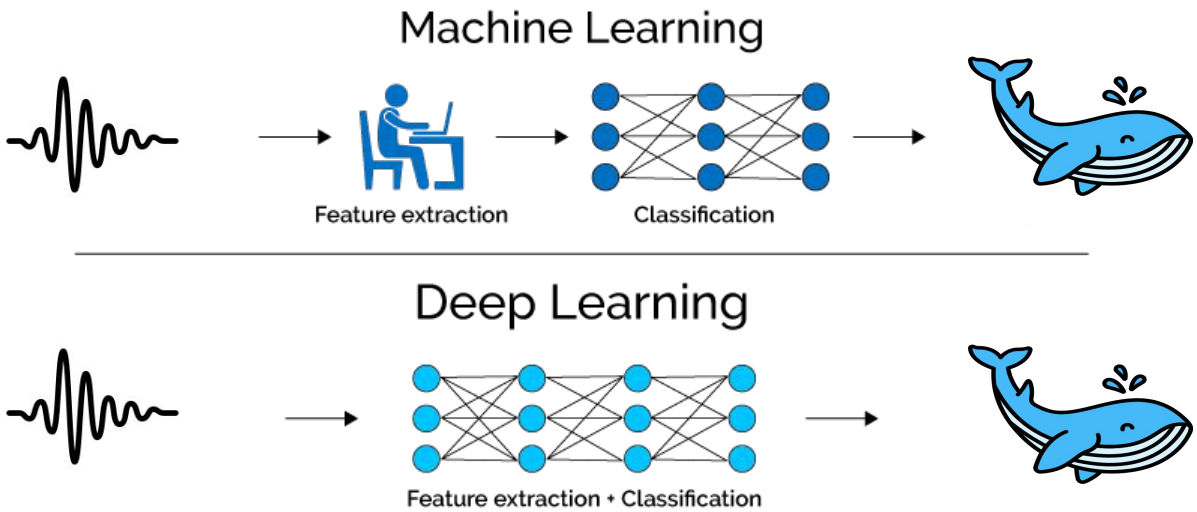

Deep learning in a nutshell

Deep Learning aims to be end-to-end (but usually isn't)

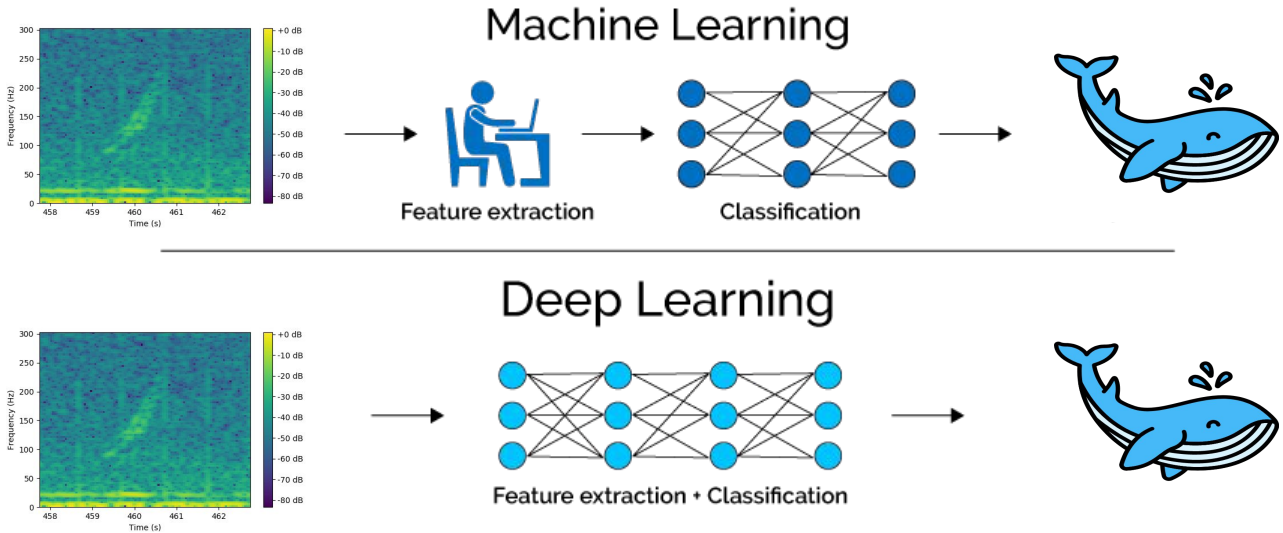

Transfer learning (model adaptation)

- Transfer learning can ...
  - drastically reduce amount of training data and training time
  - make models more adaptable and reusable

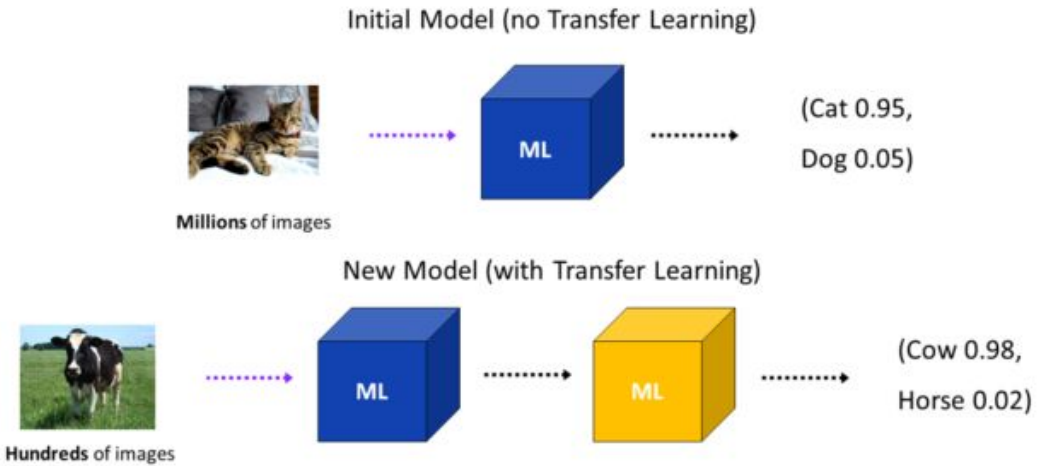

Workflow

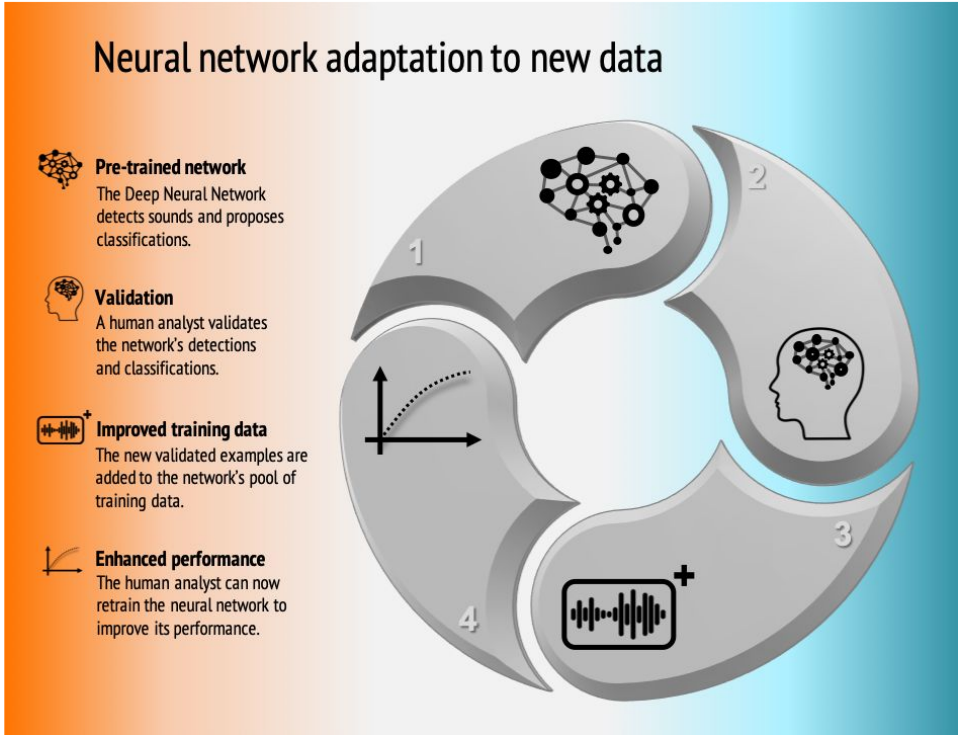

p. 144-145:
"Towards Versatile and Adaptive Detection Algorithms in Underwater Acoustics"

Making deep learning accessible

TensorFlow → 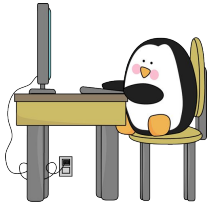 Deep learning developer

TensorFlow →  → 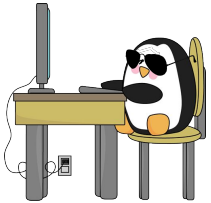 Data scientist / marine biologist

Etymology

- In Ancient Greek, *ketos* denotes a large fish, whale, shark, or sea monster.
- *ketos* is also the origin of the scientific term for whales, cetacean.

Ketos - at a glance

Neural network architectures

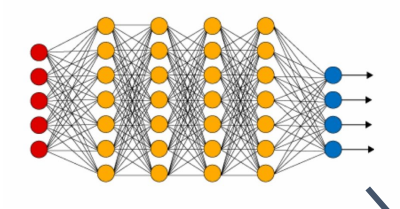

Audio processing

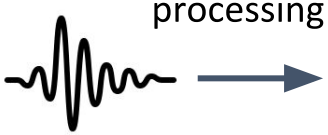

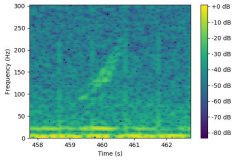

	A	B	C	D
1	filename	label	start	end
2	IML-BA_Sample_2019-06-02_025209.wav	1	0	3
3	IML-BA_Sample_2019-06-02_065235.wav	1	0	3
4	IML-BA_Sample_2019-06-08_012119.wav	1	0	3
5	IML-BA_Sample_2019-06-21_201729.wav	1	0	3
6	IML-BA_Sample_2019-06-29_124902.wav	1	0	3
7	VAS_Sample_2019-06-06_052058.wav	1	0	3
8	VAS_Sample_2019-06-06_052147.wav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

Annotation tables

Training Database

Training methods

Trained DC model

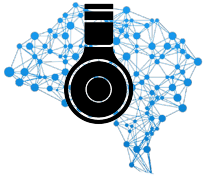

Ketos - at a glance

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
- Fully 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

Open-source libraries

1. Audio data manipulation & processing

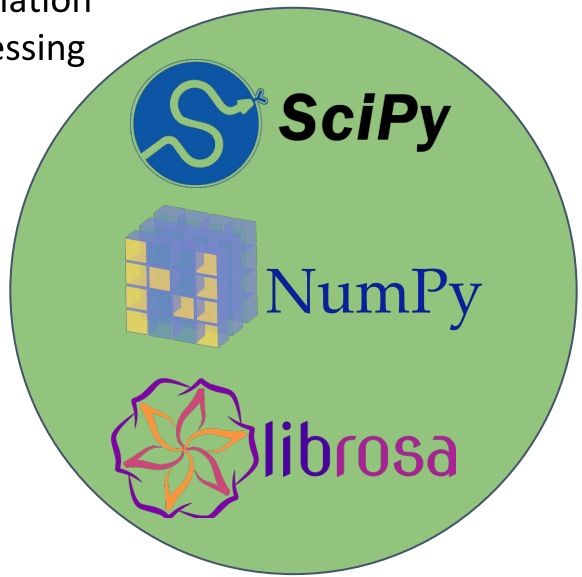

2. Data handling & storage

3. Machine learning, neural networks

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

ketos

Underwater acoustic detection and classification with deep neural networks

2.1

Search

- Introduction
- Installation
- ⊞ Tutorials
- ⊞ Modules
- ⊞ How to contribute
- ⊞ Versions

Docs » Welcome to Ketos's documentation! View page source

Welcome to Ketos's documentation!

Introduction

Ketos provides a unified, high-level interface for working with acoustic data and deep neural networks. Its main purpose is to support the development of deep learning models for solving challenging detection and classification problems in underwater acoustics.

Ketos is written in Python and utilizes a number of powerful software packages including NumPy, HDF5, and Tensorflow. It is licensed under the GNU GPLv3 license and hence freely available for anyone to use and modify. The project is hosted on GitLab at https://gitlab.meridian.cs.dal.ca/public_projects/ketos .

Ketos was developed by the MERIDIAN Data Analytics Team at the Institute for Big Data Analytics at Dalhousie University. We are greatful to Amalis Riera and Francis Juanes at the University of Victoria, Kim Davies and Chris Taggart at Dalhousie University, and Kristen Kanes at Ocean Networks Canada for providing us with annotated acoustic data sets, which played a key role in the early phases of the project. The first version of Ketos was released in April 2019.

The intended users of Ketos are primarily researchers and data scientists working with (underwater) acoustics data. While Ketos comes with complete documentation and comprehensive step-by-step tutorials, some familiarity with Python and especially the NumPy package would be beneficial. A basic understanding of the fundamentals of machine learning and neural networks would also be an advantage.

To get started with Ketos, follow the Installation instructions and then proceed to the Tutorials section. For an example application of Ketos, see Kirsebom, Frazao, et al., Performance of a deep neural network at detecting North Atlantic right whale upcalls, JASA, 147, 2636 (2020) (preprint).

The name Ketos was chosen to highlight the package's main intended application, underwater acoustics. In Ancient Greek, the word ketos denotes a large fish, whale, shark, or sea monster. The word ketos is also the origin of the scientific term for whales, cetacean.

Indices and tables

- Index
- Module Index
- Search Page

Next

Ketos example: North Atlantic right whale

Training and test data

- Gulf of St Lawrence & Gulf of Maine
- Surface boys and bottom moorings
- Over 5,000 NARW upcalls
- Data and annotations available at FRDR: doi.org/10.20383/101.0241

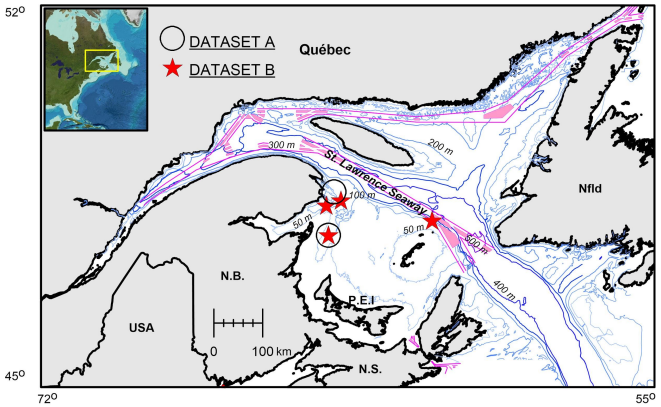

Kirsebom, Frazao, *et al.* (2020) JASA **147**, 2636

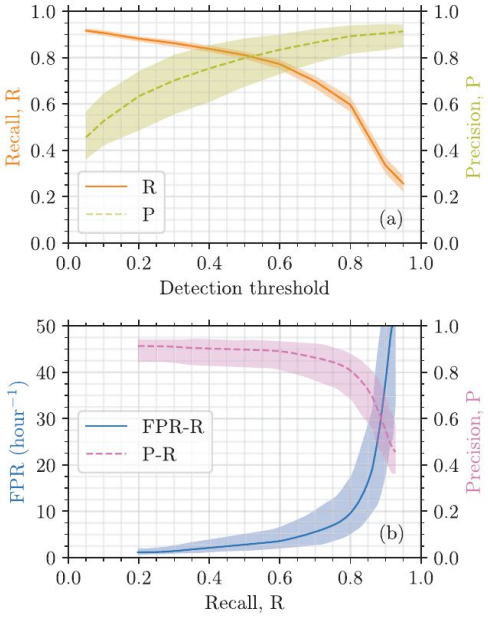

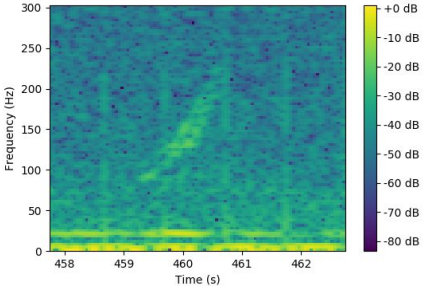

Example of NARW upcall

Computer code available at Zenodo:
doi.org/10.5281/zenodo.3736625

Step-by-step tutorials at:
docs.meridian.cs.dal.ca/ketos/

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
- offers extensive documentation and step-by-step tutorials
- check it out at https://docs.meridian.cs.dal.ca/ketos/