

Deep Learning for underwater acoustics



MERIDIAN

Fabio Frazao

Institute for Big Data Analytics/MERIDIAN
@ Dalhousie University

Detection and Classification in Marine Bioacoustics signals with Deep Learning

Outline



- What is deep learning?
- A few popular architectures
 - Convolutional Neural Networks
 - Sequence Models
 - Generative Models
- Level of details in annotations
- MERIDIAN's workflow vision
- Repository of public resources

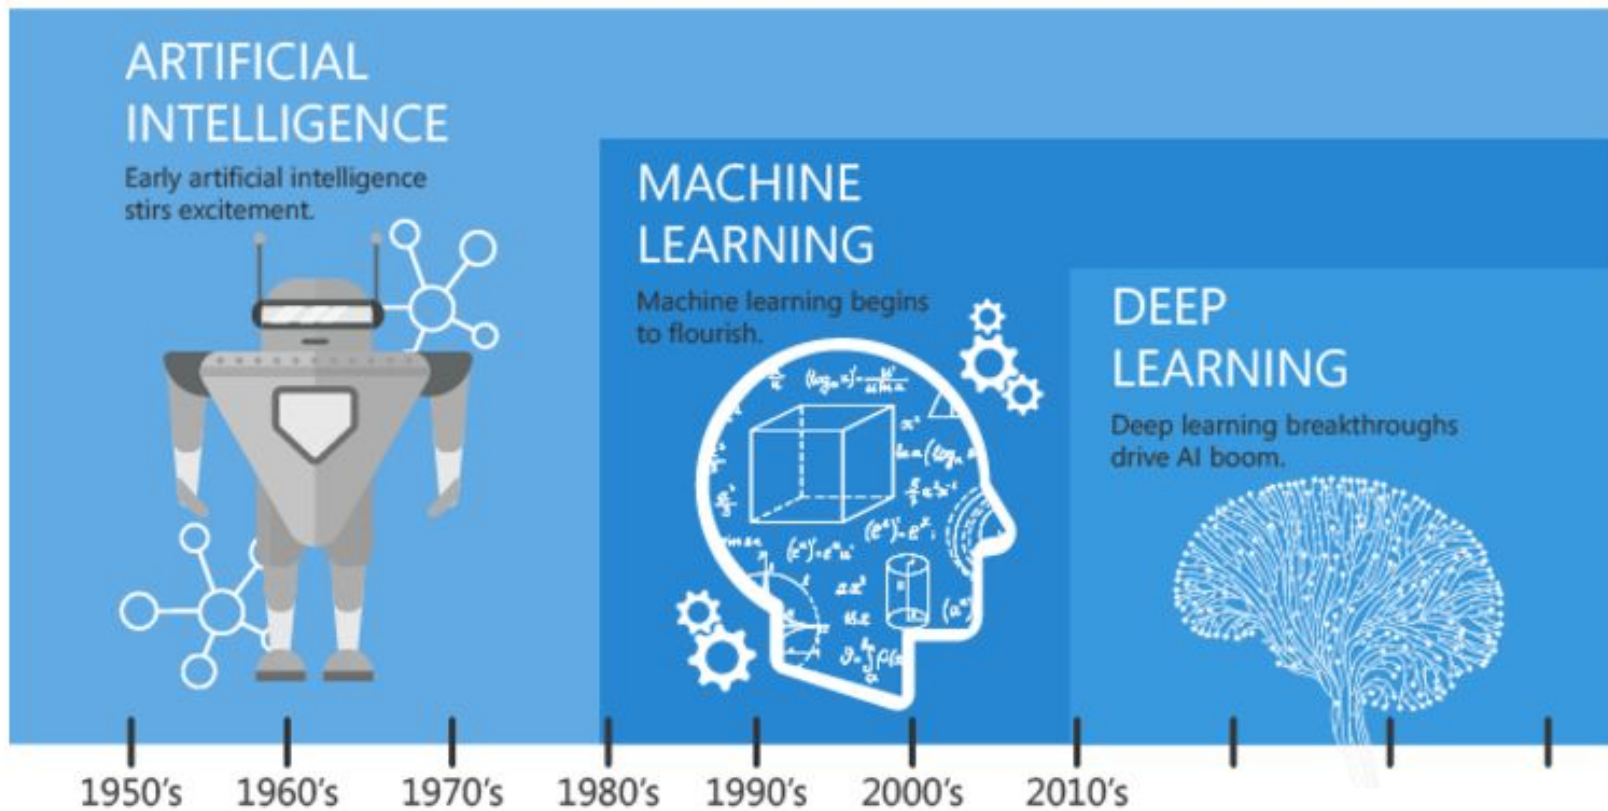
What is Deep Learning



What is Deep Learning?



A set of machine learning techniques using neural networks



What is Deep Learning?




The “learning” in machine learning refers to getting a program to do a task without having to tell it every single step of the process.

What is Deep Learning?




The “learning” in machine learning refers to getting a program to do a task without having to tell it every single step of the process.



What is Deep Learning?



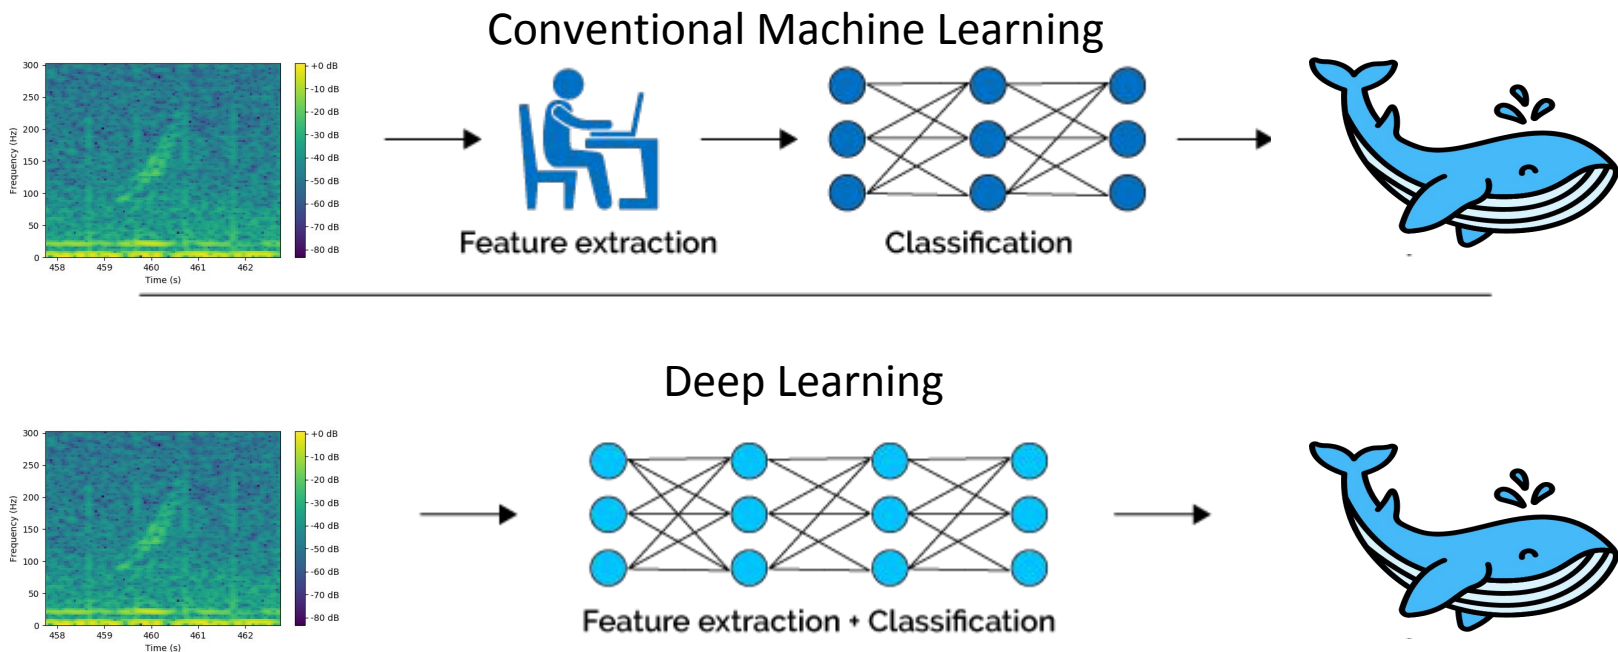
The “learning” in machine learning refers to getting a program to do a task without having to tell it every single step of the process.



What is Deep Learning?



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



What is Deep Learning?




What is a Neural Network?

What is Deep Learning?



What is a Neural Network?



Grant Sanderson, [3blue1brown](https://www.youtube.com/channel/UCBaWw44pJQ7Pn71K4u3Fm4g)

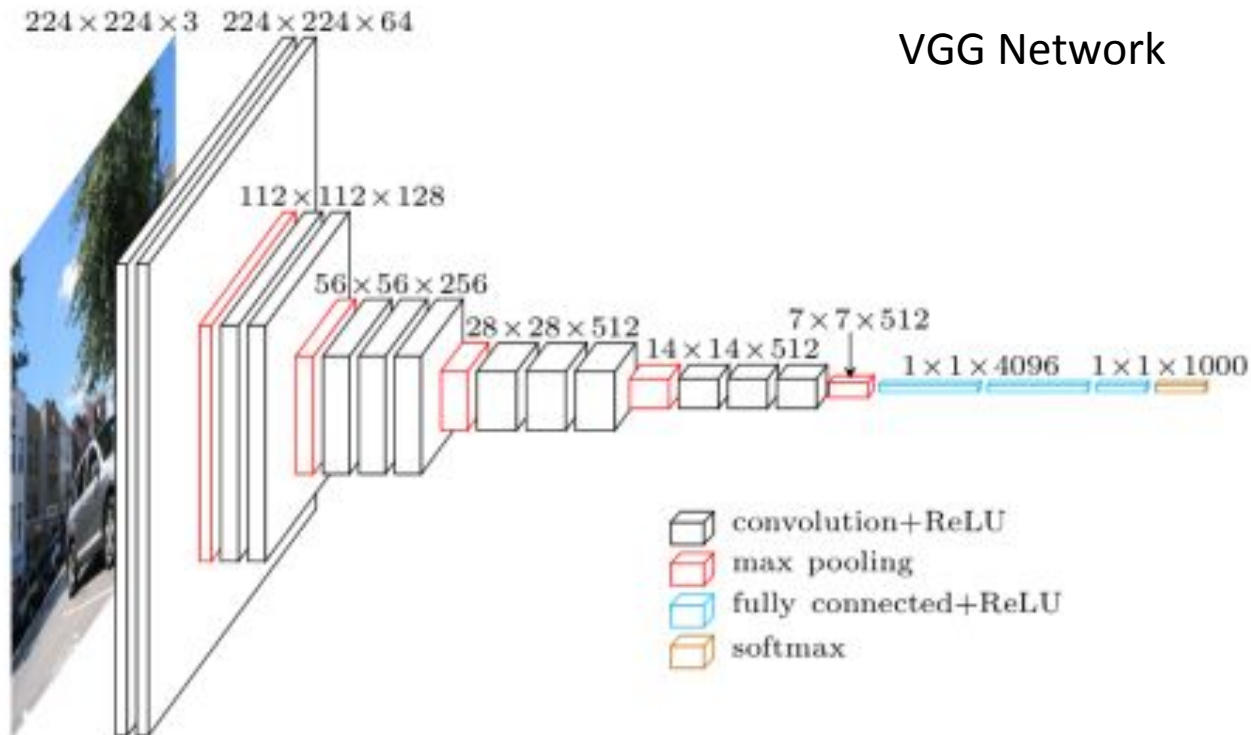
Popular Architectures



A few popular architectures



Convolutional Neural Networks (CNNs)



A few popular architectures



Convolutional Neural Networks (CNNs)



A few popular architectures

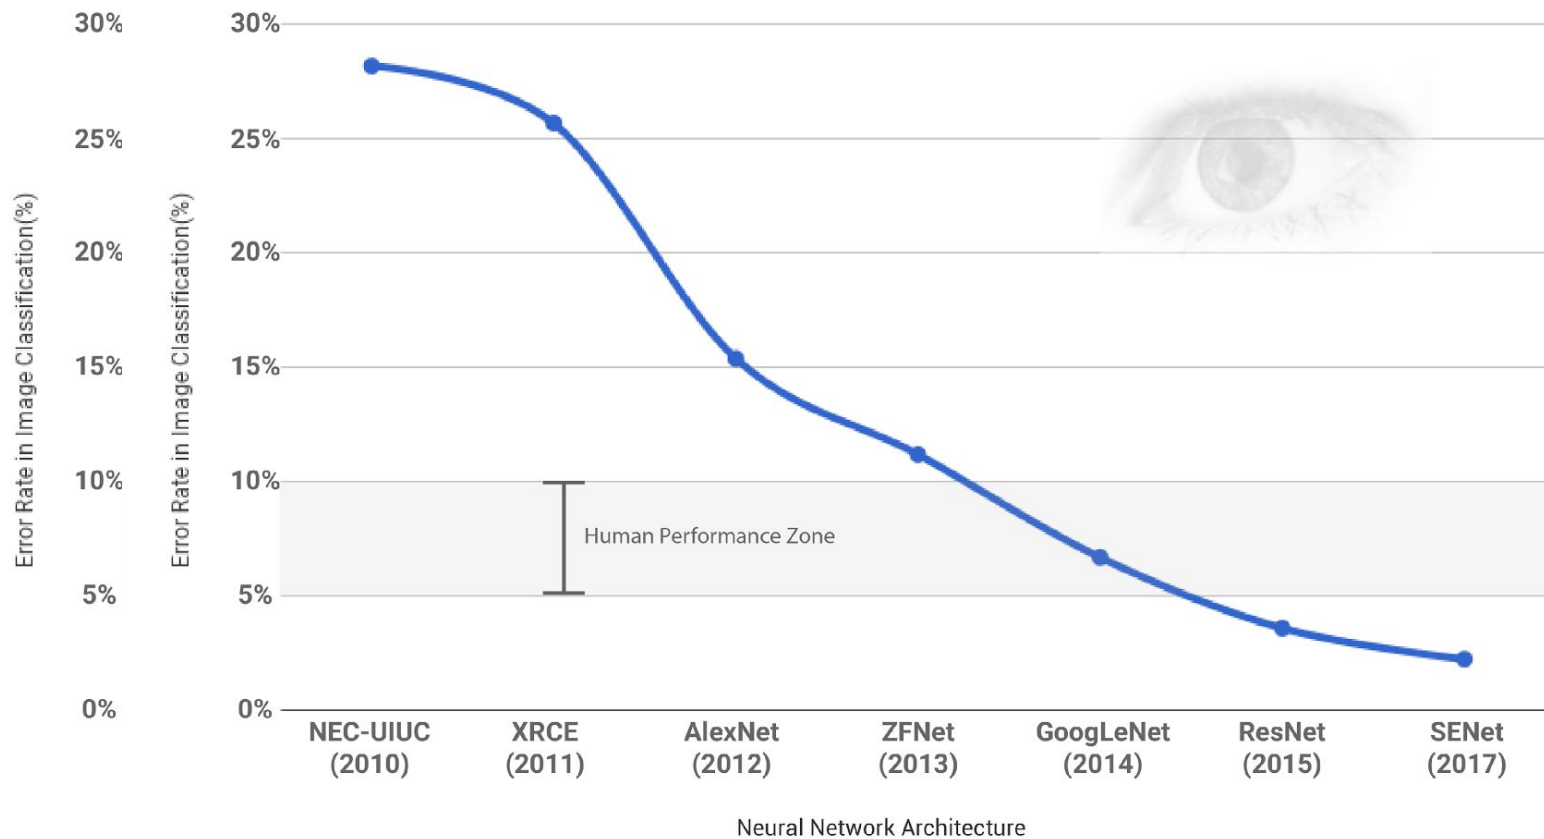


IMAGENET Large Scale Visual Recognition Challenge (ILSVRC)

Competition

The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) evaluates algorithms for object detection and image classification at large scale. One high level motivation is to allow researchers to compare progress in detection across a wider variety of objects – taking advantage of the quite expensive labeling effort. Another motivation is to measure the progress of computer vision for large scale image indexing for retrieval and annotation.


A few popular architectures



A few popular architectures



Object Detection (YOLO, R-CNN, Fast R-CNN, Faster R-CNN, etc)



A few popular architectures




Sequence models

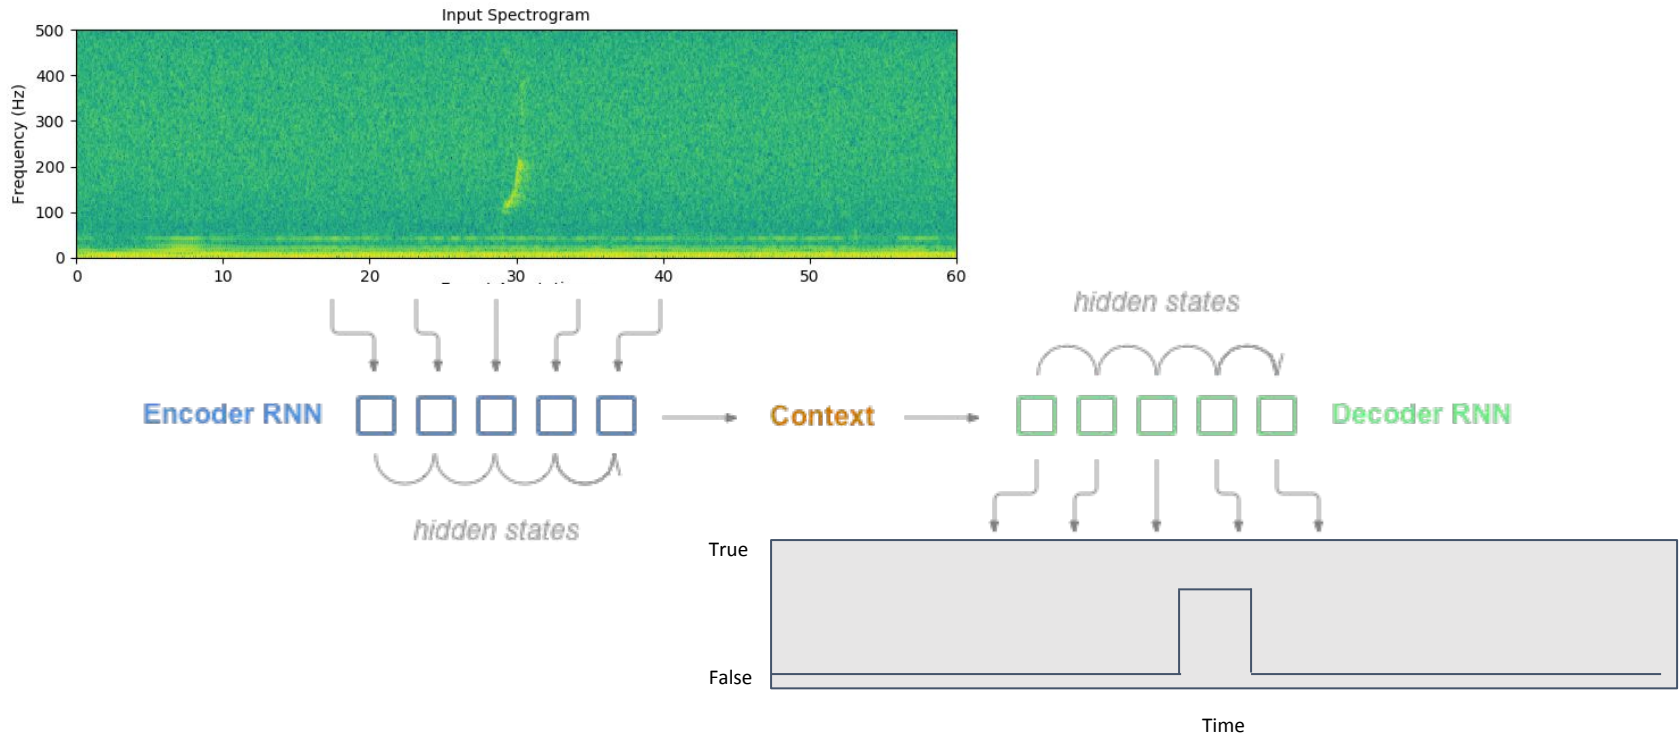
A few popular architectures



Sequence to sequence models



A few popular architectures



A few popular architectures



Generative models: see Bruno's presentation


Annotations



Detail level of annotations and the challenges of learning from another machine




Shediac_ete_2018-07-06_103047_1.wav



Detail level of annotations and the challenges of learning from another machine


Shediac_ete_2018-07-06_103047_1.wav



File annotated with lower level of details

| File | Label |
|-------------------------------------|-----------------|
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) |

Detail level of annotations and the challenges of learning from another machine




File annotated with higher level of etails

| File | Label | Timestamp (start time, end time) |
|-------------------------------------|-----------------|----------------------------------|
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (12,14) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (17,19) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (28,30) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (39,41) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (52,54) |
| | | |

Detail level of annotations and the challenges of learning from another machine

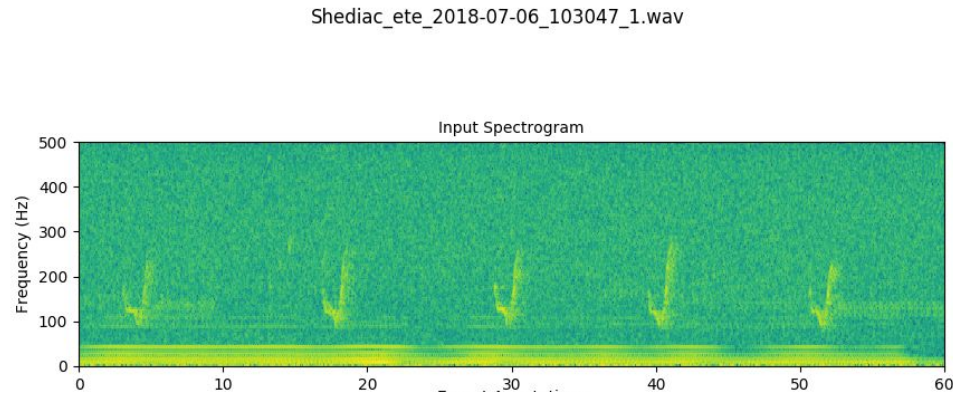
Shediac_ete_2018-07-06_103047_1.wav



File annotated with higher level of etails

| File | Label | Timestamp (start time, end time, min. freq., max freq.) |
|-------------------------------------|-----------------|---|
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (12,14,80,250) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (17,19,80,250) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (28,30,80,250) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (39,41, 80, 300) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (52,54, 75, 220) |
| | | |

Detail level of annotations and the challenges of learning from another machine




File annotated with higher level of etails

| File | Label | Timestamp (start time, end time, min. freq., max freq.) |
|-------------------------------------|-----------------|---|
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (12,14,80,250) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (17,19,80,250) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (28,30,80,250) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (39,41, 80, 300) |
| Shediac_ete-2018-07-06_103047_1.wav | 1 (NARW upcall) | (52,54, 75, 220) |
| Shediac_ete-2018-07-06_103047_1.wav | 2 (humpback) | (57, 59,80 350) |
| Shediac_ete-2018-07-06_103047_1.wav | 3 (seismic) | (53, 54,50 500) |

Detail level of annotations and the challenges of learning from another machine




Shediac_ete_2018-07-06_103047_1.wav



Detail level of annotations and the challenges of learning from another machine



Detector 1




Detail level of annotations and the challenges of learning from another machine



Expert validation

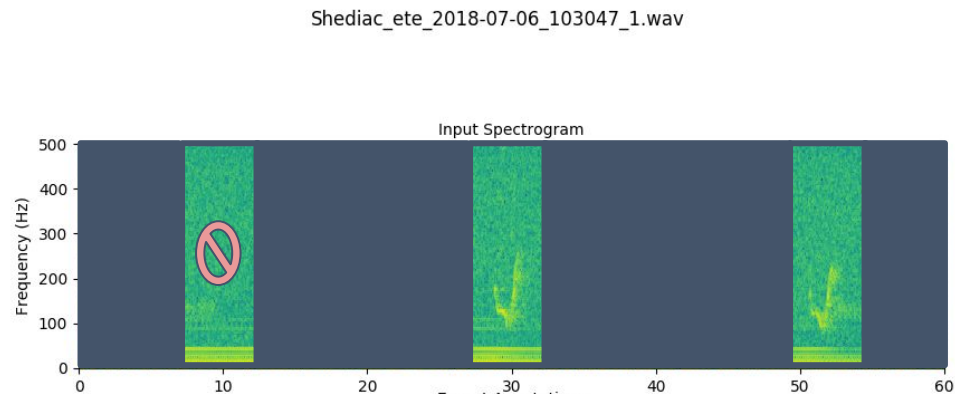
Shediac_ete_2018-07-06_103047_1.wav



Detail level of annotations and the challenges of learning from another machine



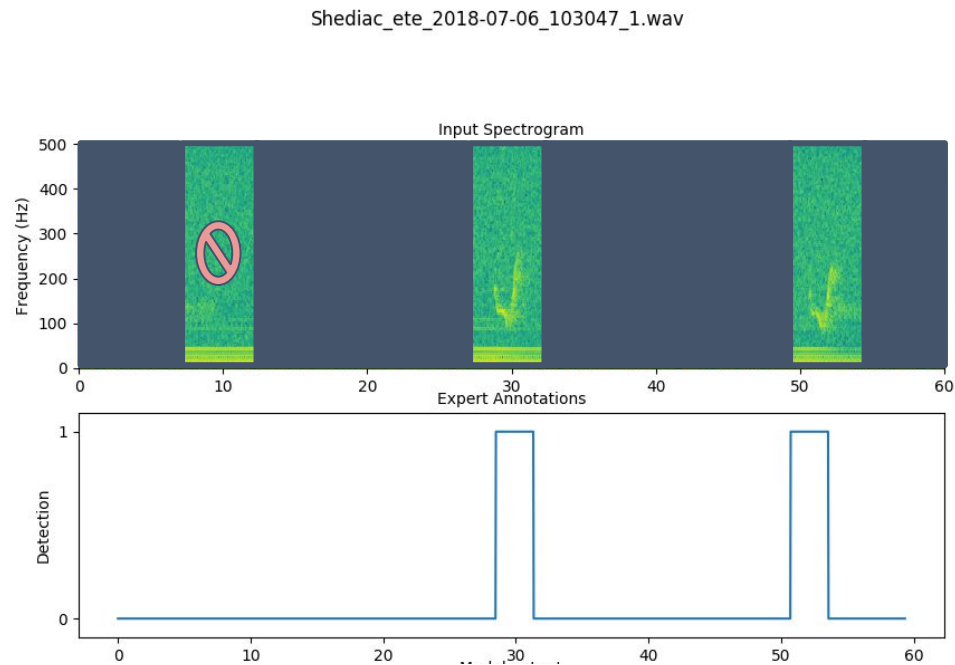
Expert validation



Detail level of annotations and the challenges of learning from another machine



Expert validation

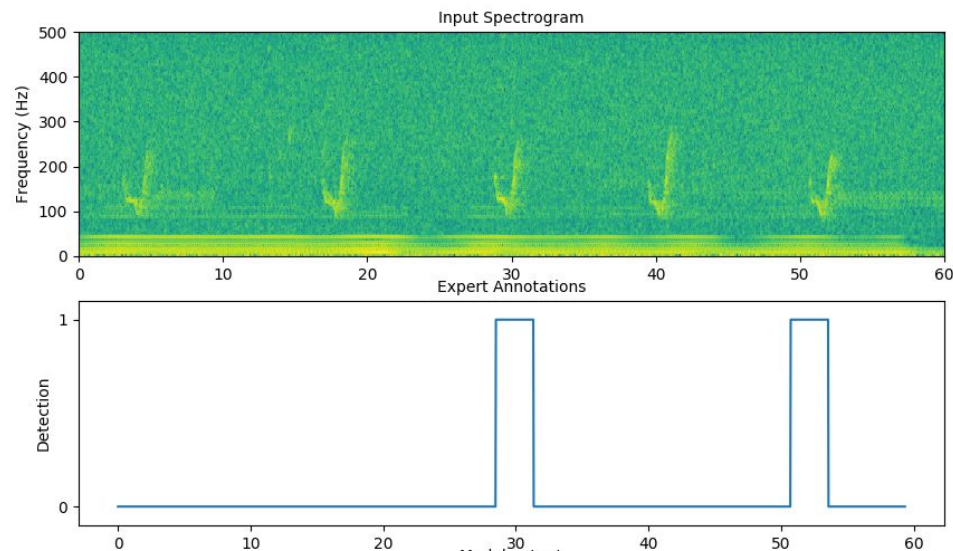


Detail level of annotations and the challenges of learning from another machine



Detector 2 training

Shediac_ete_2018-07-06_103047_1.wav




Detail level of annotations and the challenges of learning from another machine



Detector 2 training

Shediac_ete_2018-07-06_103047_1.wav





Workflow vision



Workflow visions



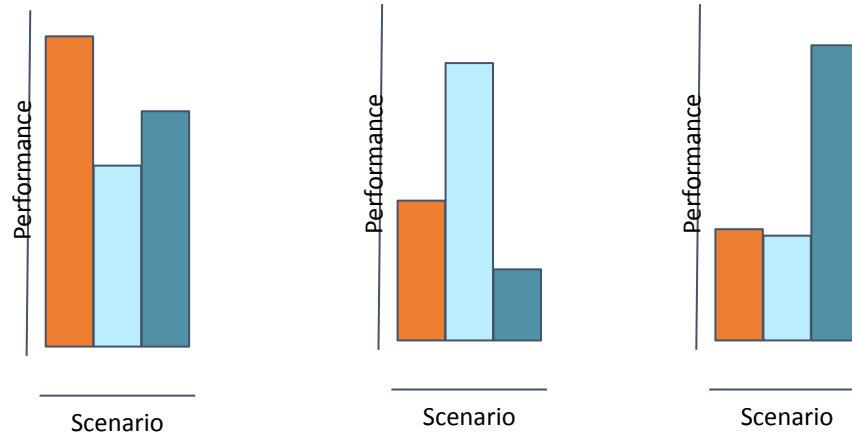

MERIDIAN's library of pre-trained models



Workflow visions



MERIDIAN's library of pre-trained models



right whales
right whales + ships
right whales + humpback whales






User's scenario

right whales + humpback whales


Workflow visions



MERIDIAN's library of pre-trained models






User's scenario



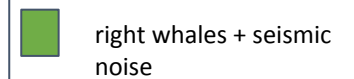
Workflow visions



MERIDIAN's library of pre-trained models






User's scenario




Workflow visions



MERIDIAN's library of pre-trained models








User's scenario



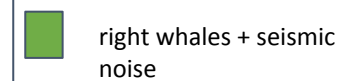
Workflow visions



MERIDIAN's library of pre-trained models




User's scenario




Workflow visions





Transfer learning



User's scenario

 right whales + seismic noise



Neural network adaptation to new data



Pre-trained network

The Deep Neural Network detects sounds and proposes classifications.



Validation

A human analyst validates the network's detections and classifications.




Improved training data

The new validated examples are added to the network's pool of training data.



Enhanced performance

The human analyst can now retrain the neural network to improve its performance.




Moving Forward as a Community




Repository of public resources



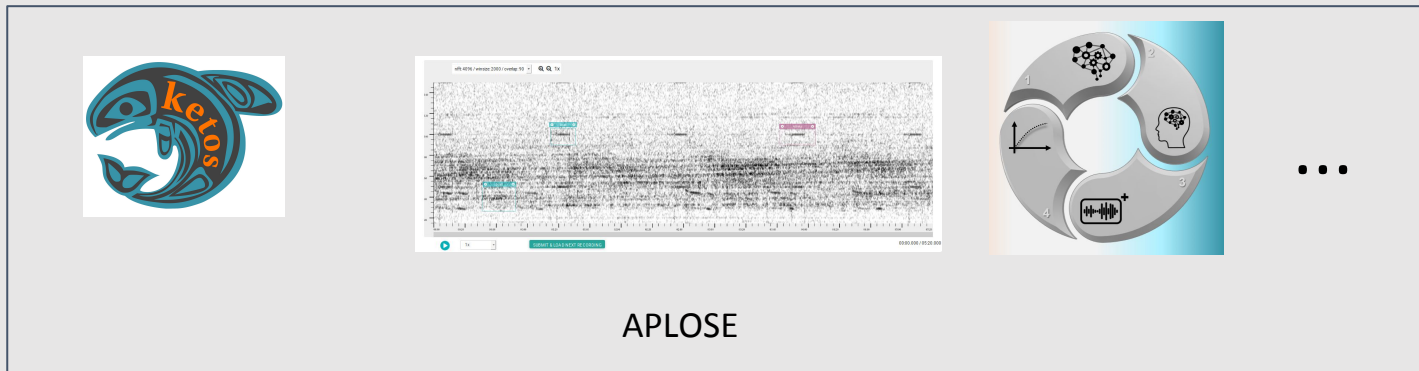
Models



Training datasets



Tools



Thank you!

