SciTinyML - ICTP workshop

Scientific Use of Machine Learning on Low Power Devices

Motion Classification – Anomaly Detection

Prof. Marcelo José Rovai
UNIFEI - Universidade Federal de Itajubá, Brazil
Web: https://github.com/Mjrovai
Email: rovai@unifei.edu.br

Who I am

- Brazilian from São Paulo, **Data Science Master's degree by UDD, Chile**, and MBA by IBMEC (INSPER), Brazil.
- Graduated in 1982 as an **Engineer from UNIFEI** with Specialization from Poli/USP, both in Brazil.
- Worked as a **teacher, engineer, and executive** in several technology companies such as CDT/ETEP, AVIBRAS Aeroespacial, SID Informática, ATT-GIS, NCR, DELL, COMPAQ (HP), and more recently at IGT as a Regional VP, where continue as a Senior Advisor for Latin America.
- **Write about electronics**, publishing in sites as MJRoBot.org (Editor/Writer), Hackster.io (#1 Contributor), Instructables.com, and Medium.com (TDS – Towards Data Science).
- **Volunteer Professor** at UNIFEI Engineering Institute, teaching "Machine Learning applied to Embedded Devices" course (IESTI01).
- Active member of the **TinyML4D group**, an initiative to bring TinyML education to developing countries.

Marcelo Rovai

Sound

Vibration

Vision

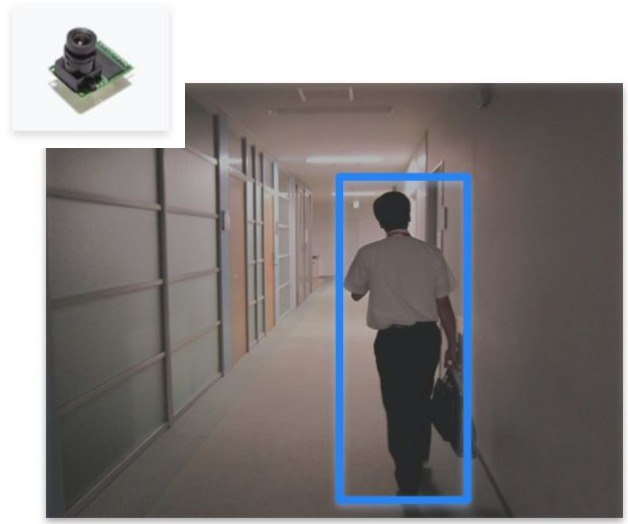

Sound

Vibration

Vision

Motion Classification

Case Study: Mechanical Stresses in Transport

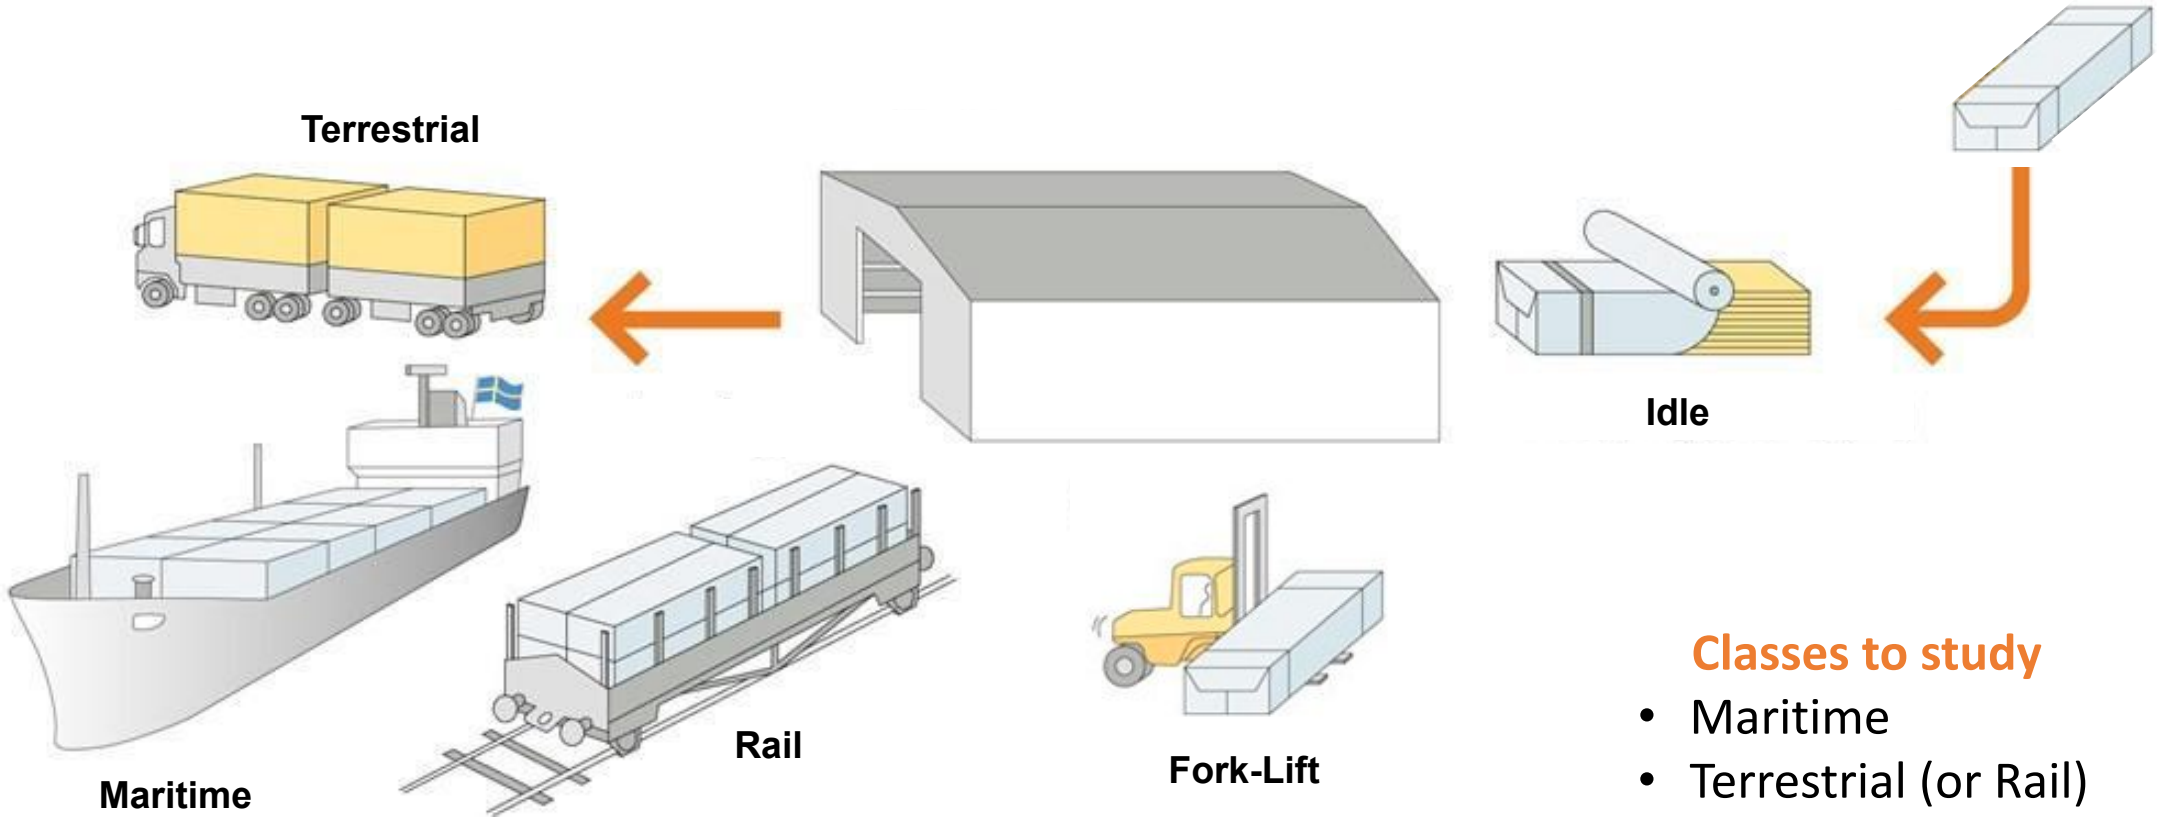

Classes to study

- Maritime
- Terrestrial (or Rail)
- Lift
- Idle

Machine Learning Workflow

Collect Data | Preprocess Data | Design a Model | Train a Model | Evaluate Optimize | Convert Model | Deploy Model | Make Inferences

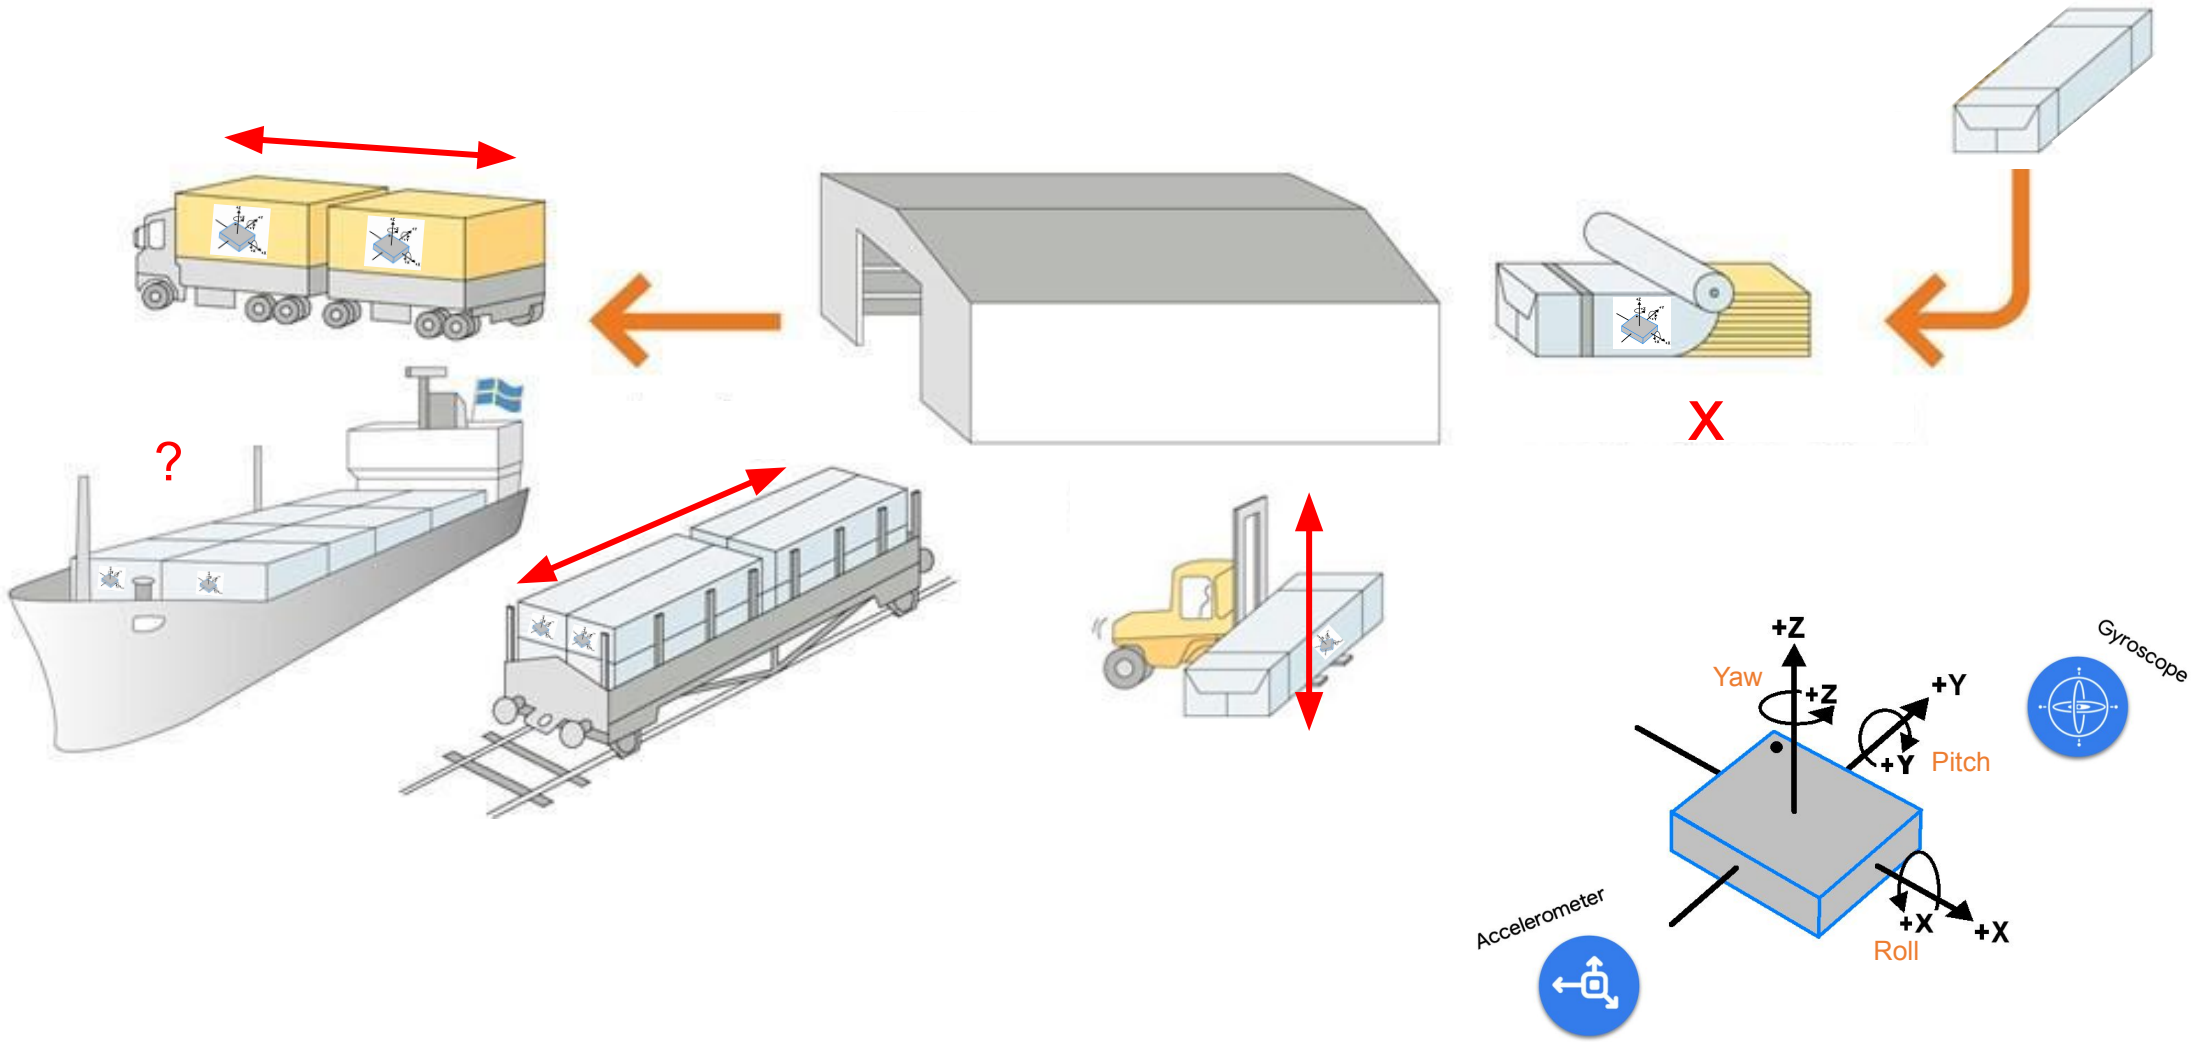

Mechanical Stresses in Maritime Transport

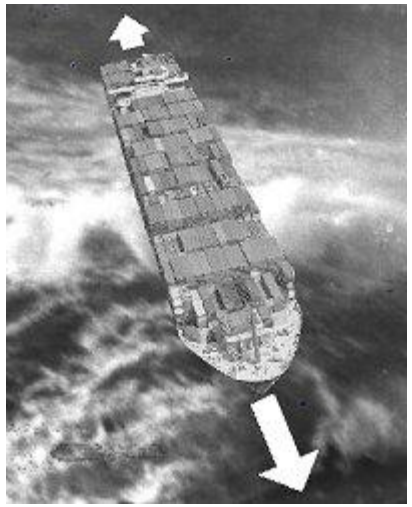

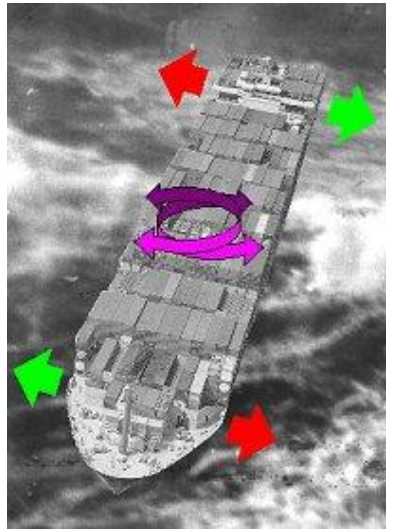

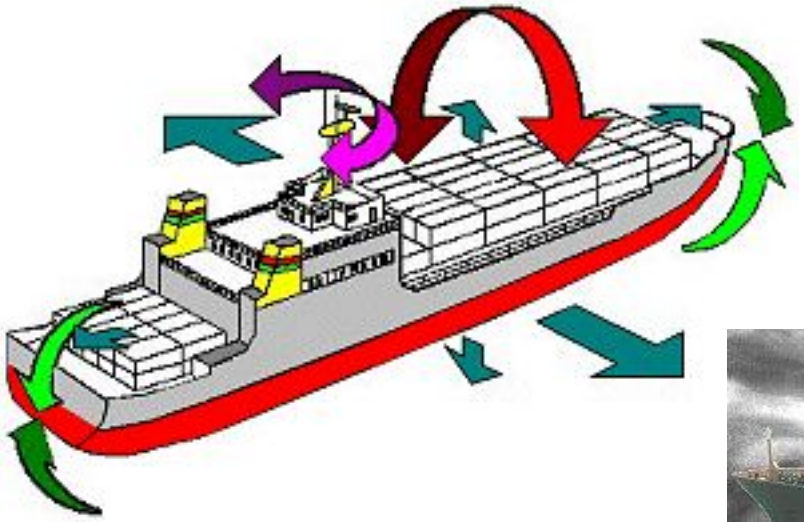

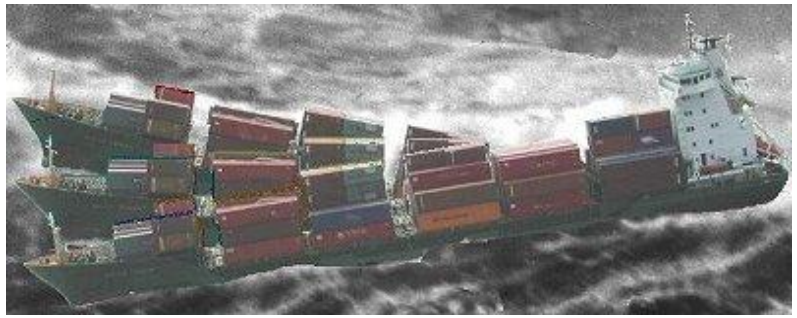

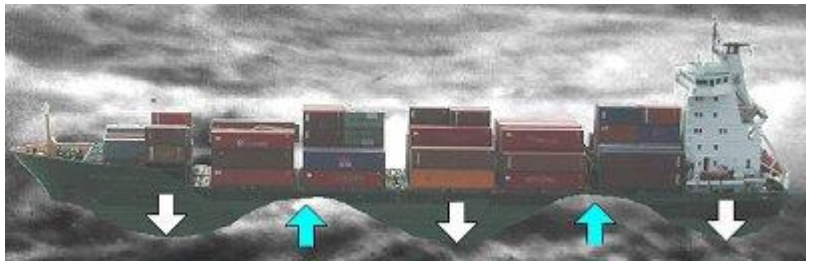

https://www.containerhandbuch.de/chb_e/stra/index.html?/chb_e/stra/stra_02_03_03.html

Mechanical Stresses in Maritime Transport

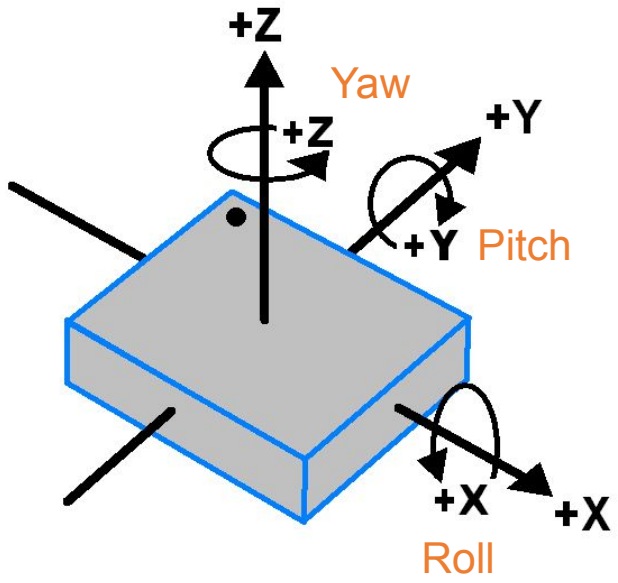

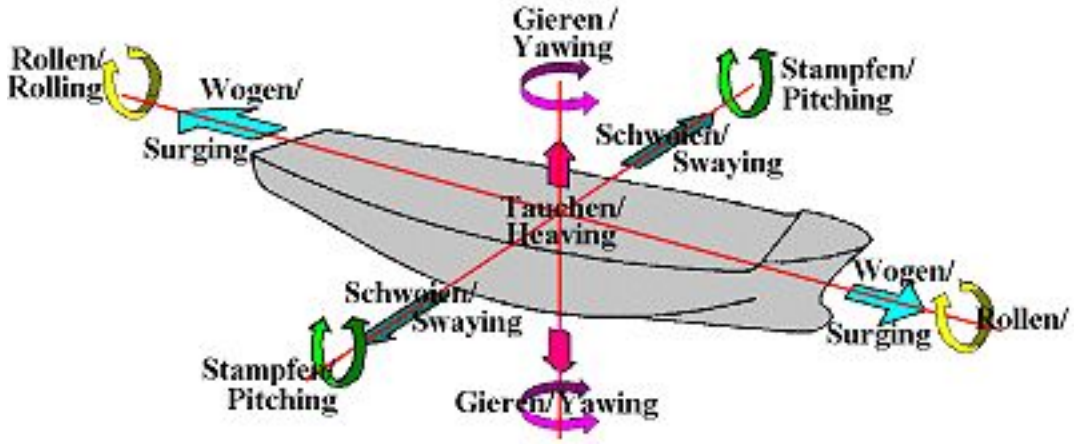

Accelerometer

Gyroscope

	Linear motion	Rotational motion	
X	Surging is motion along the longitudinal axis.	Rolling is motion around the longitudinal axis.	Roll
Y	Swaying is motion along the transverse axis.	Pitching is motion around the transverse axis.	Pitch
Z	Heaving is motion along the vertical axis.	Yawing is motion around the vertical axis.	Yaw

Summary of ship movement

Example: 10 seconds of accelerometer data, captured with a sample rate: 62.5 Hz

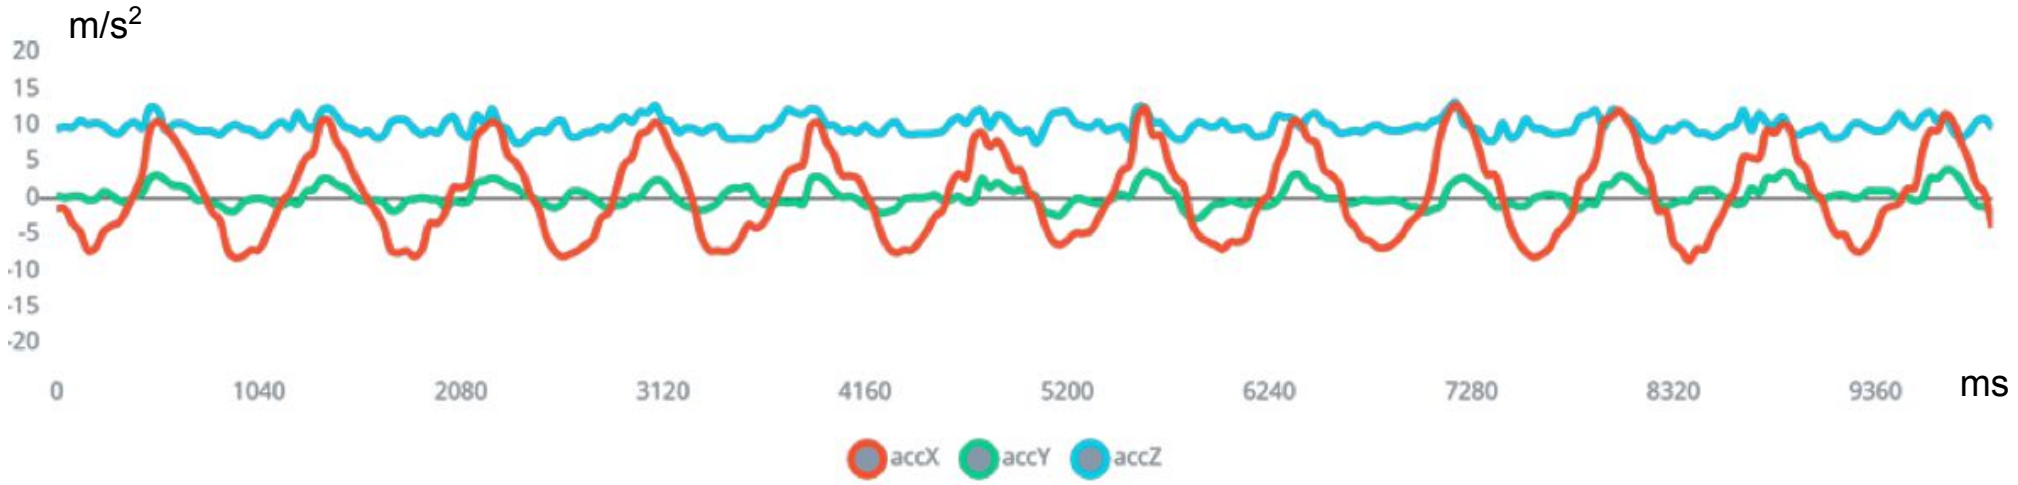

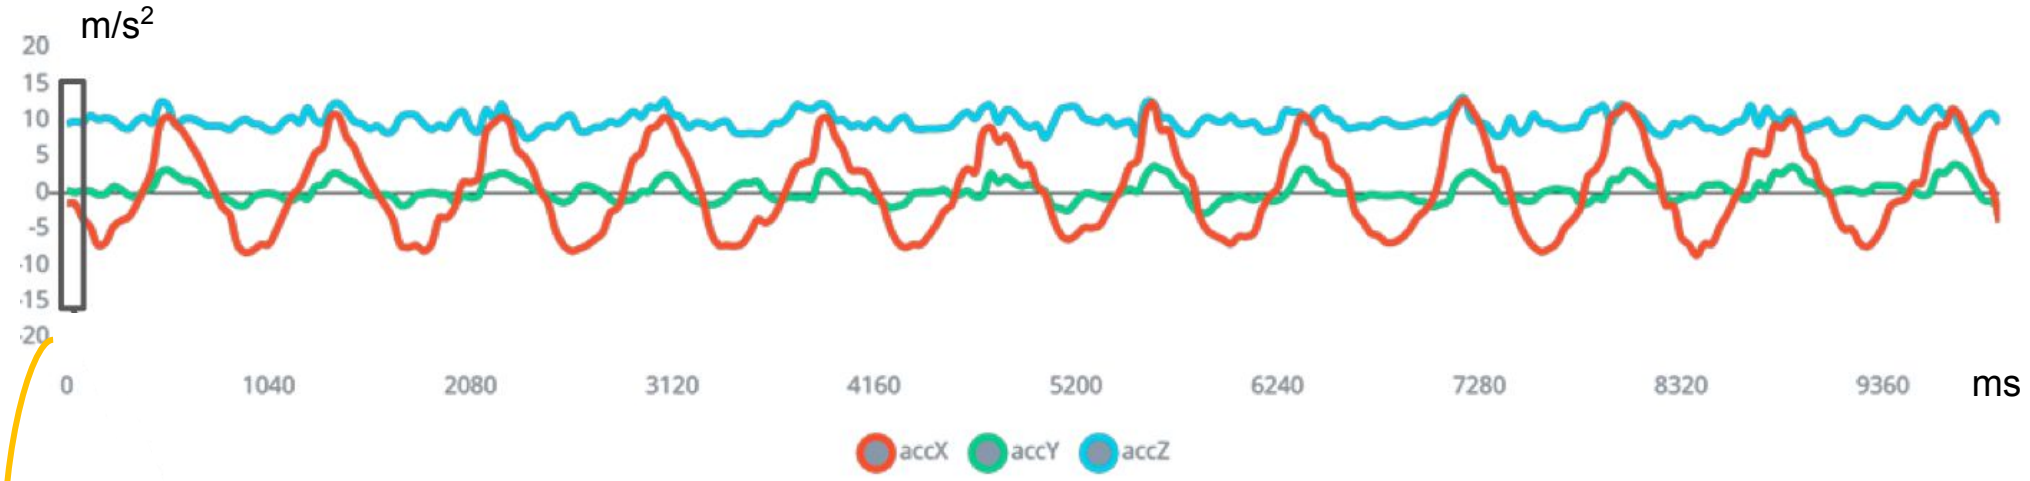

Raw Features

- accX
- accY
- accZ

Model

Classes

- Lift
- Terrestrial
- Maritime
- Idle

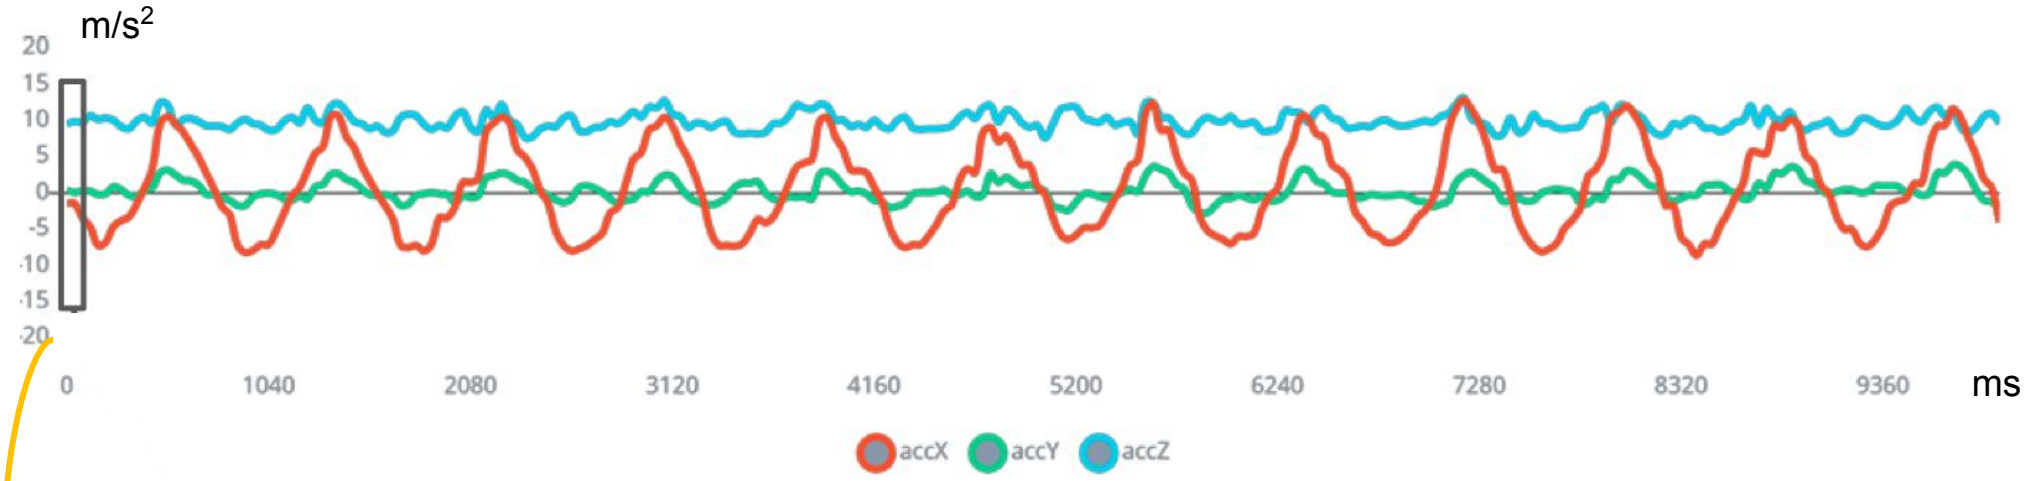

Raw Features

[-0.22, 0.01, 9.81]

Model

Classes

[1 0 0 0]

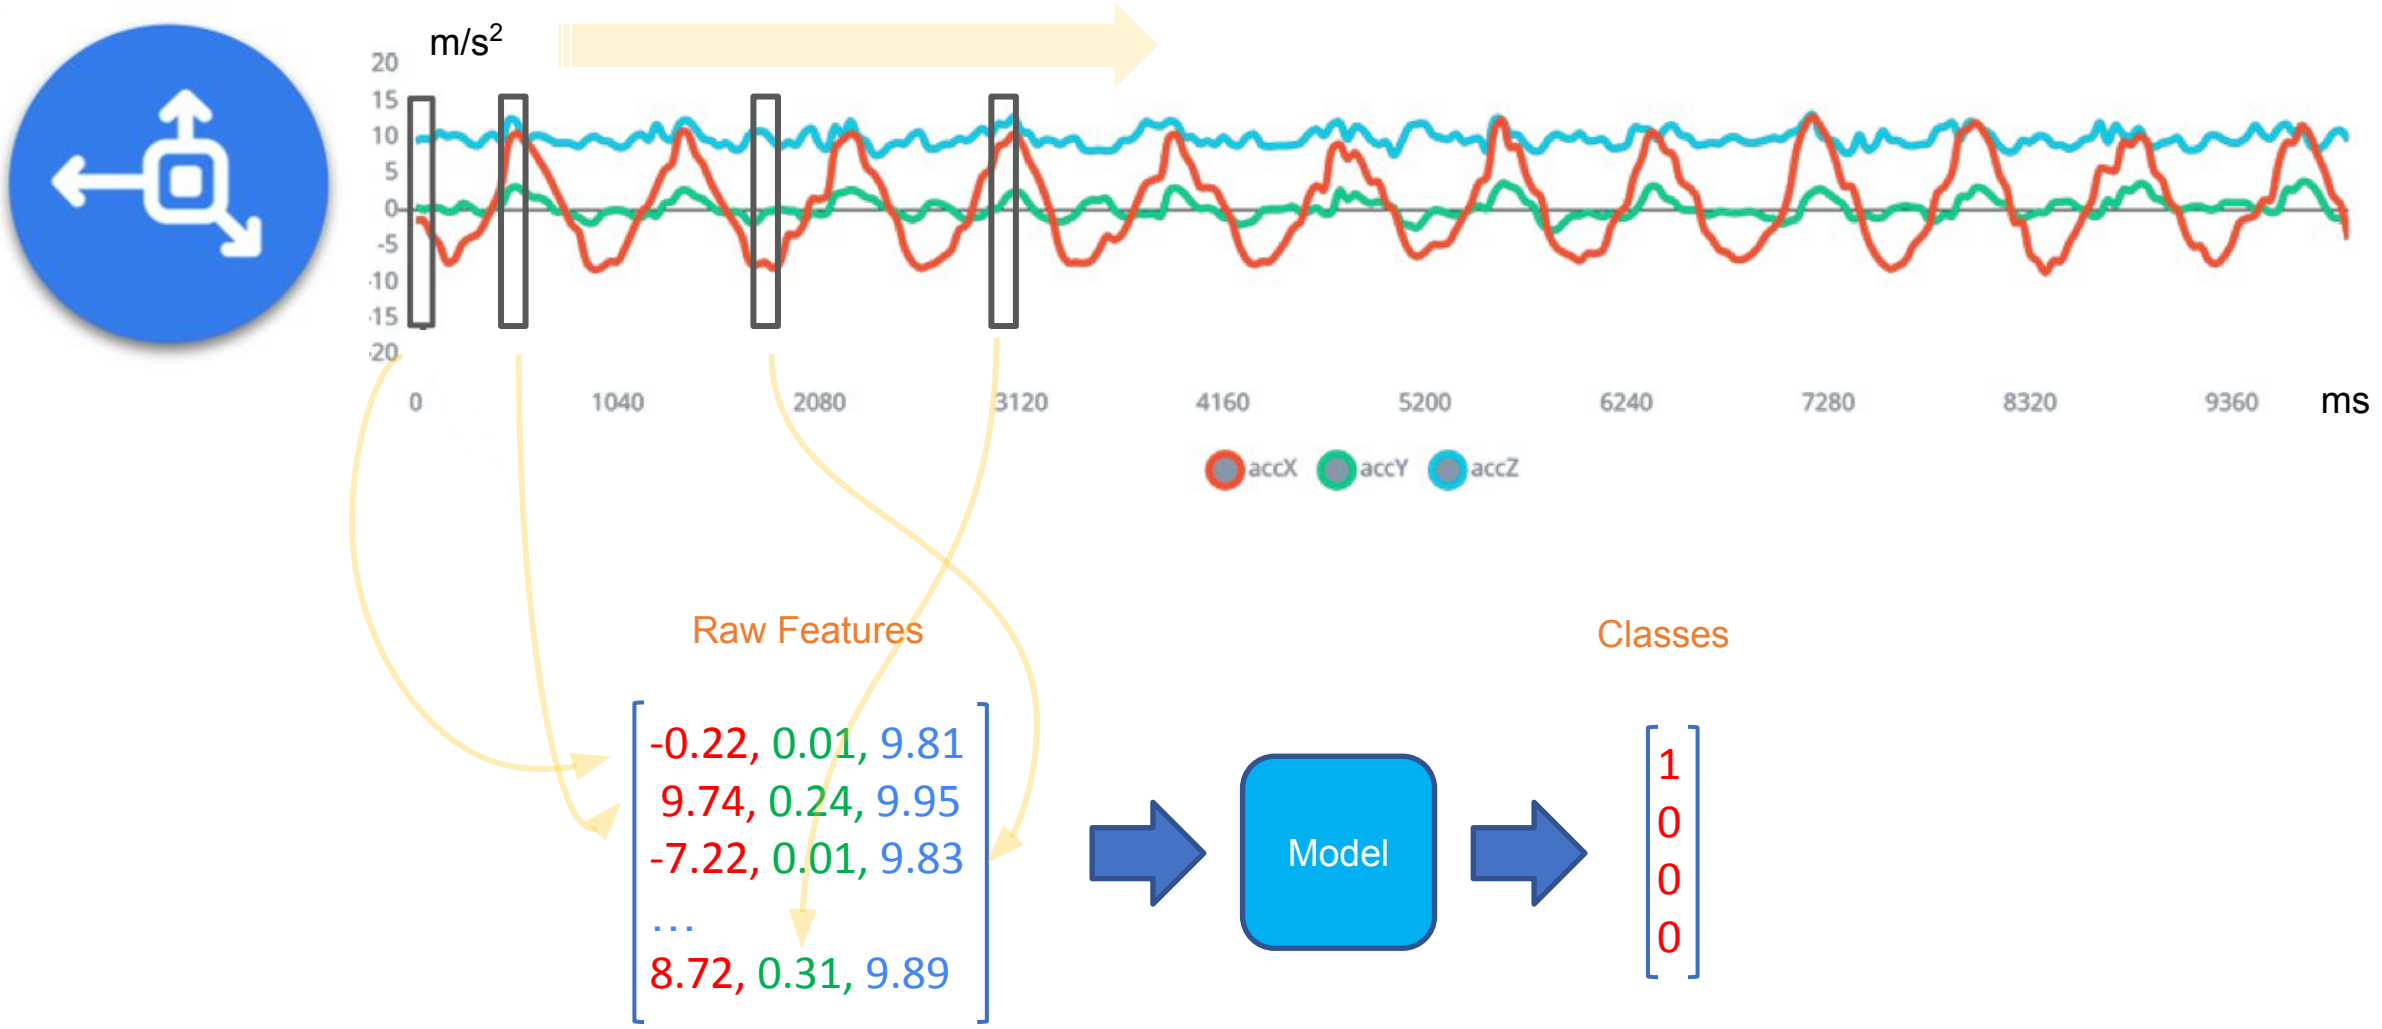
m/s2
20
15
10
5
0
-5
-10
-15
-20
0
1040
2080
3120
4160
5200
6240
7280
8320
9360
ms
accX
accY
accZ
Raw Features
-0.22, 0.01, 9.81
9.74, 0.24, 9.95
-7.22, 0.01, 9.83
…
8.72, 0.31, 9.89
Model
Classes
1
0
0
0

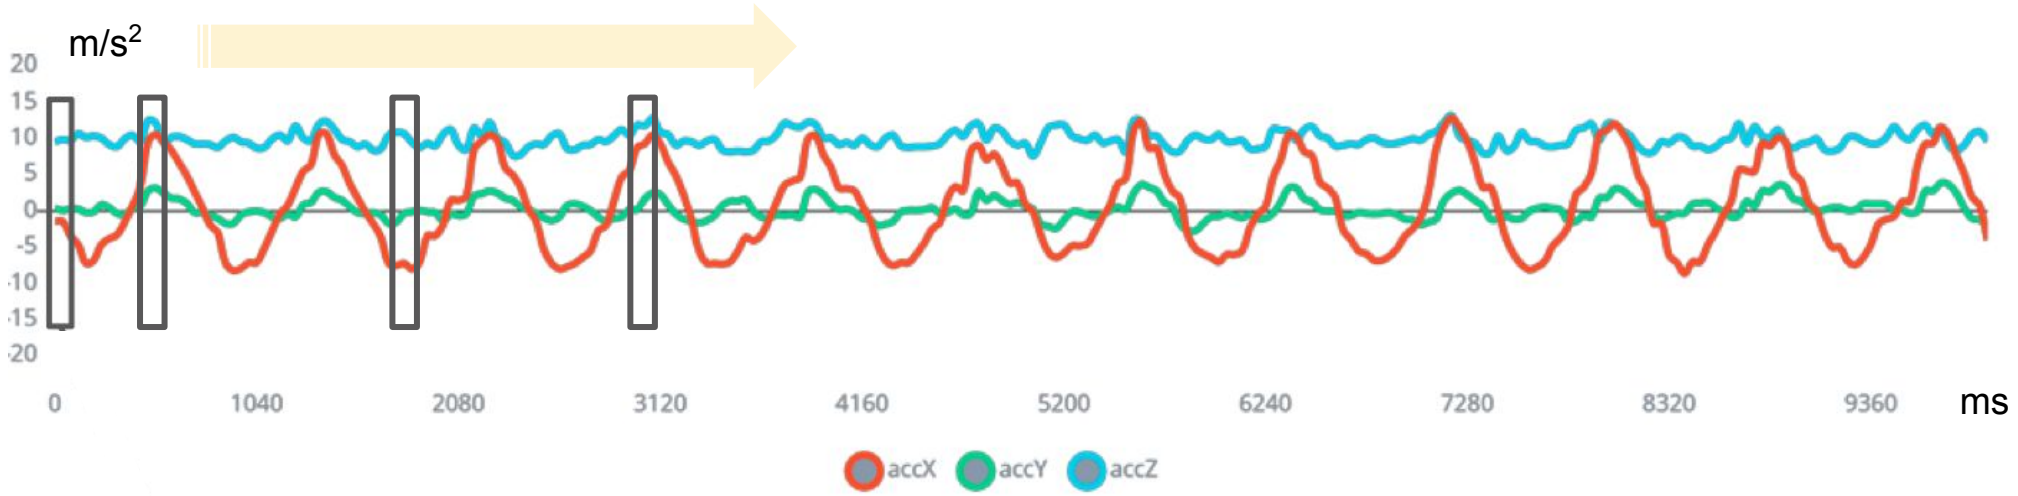

Raw Features

-0.22, 0.01, 9.81
9.74, 0.24, 9.95
-7.22, 0.01, 9.83
…
8.72, 0.31, 9.89

Model

Classes

?
?
?
?

Problem!

✔ Time-Series Data

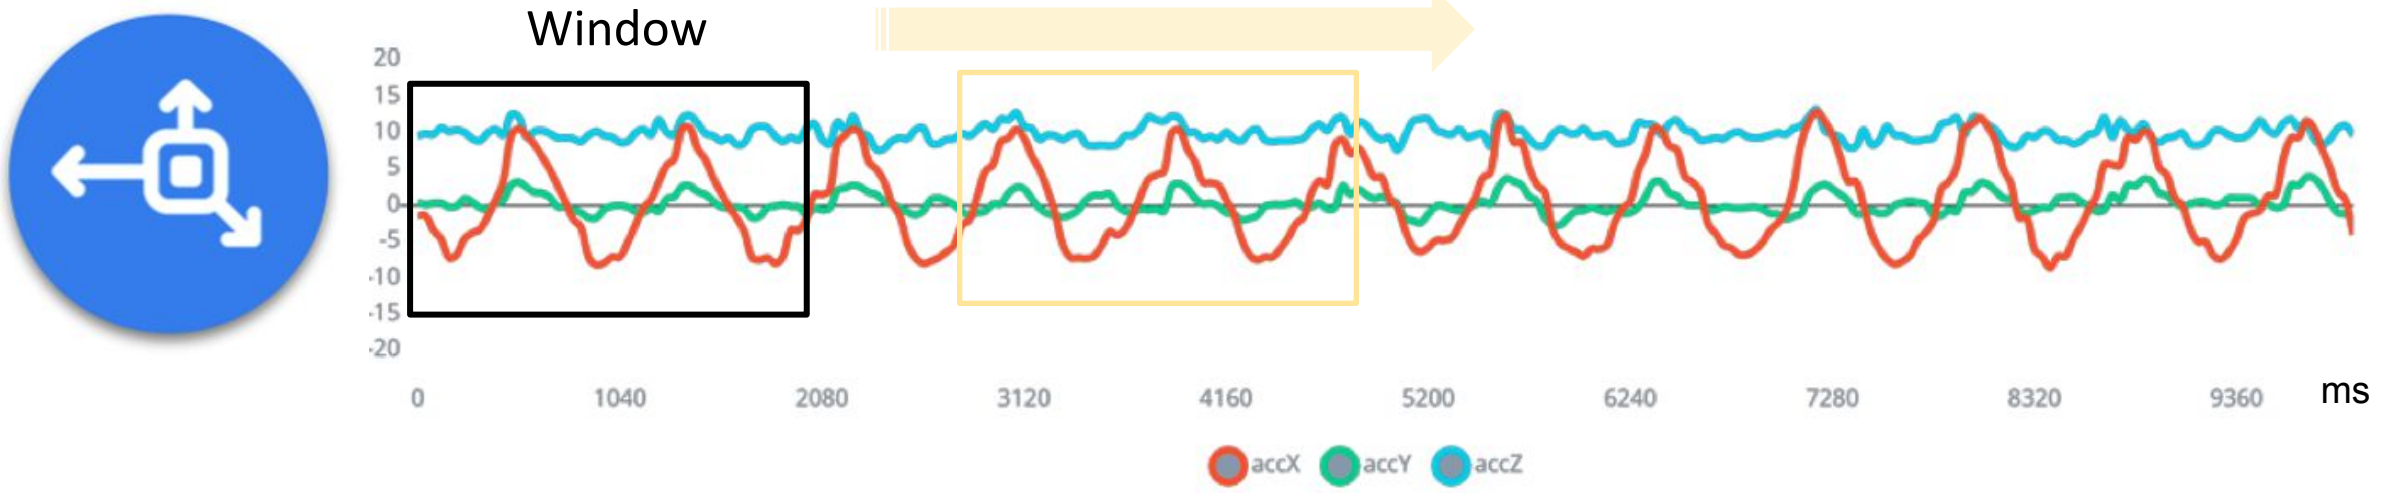
Window
20
15
10
5
0
-5
-10
-15
-20
0
1040
2080
3120
4160
5200
6240
7280
8320
9360
ms
accX
accY
accZ

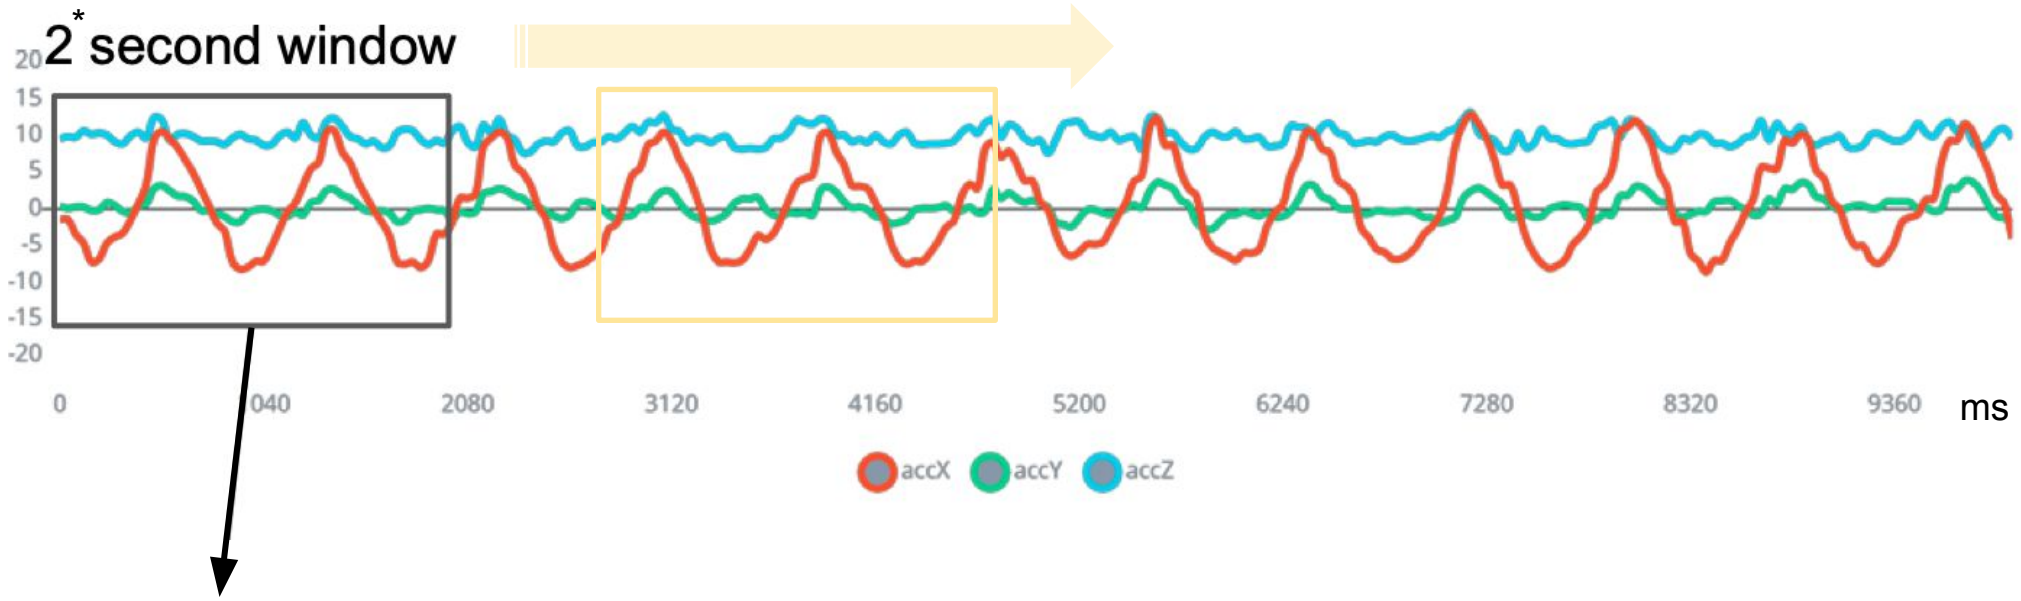

Raw Features as a window

- 125** samples for each axis (62.5Hz x 2s)
- 375 total features (125 x 3 axis)

* 2 seconds is needed to capture 1 or 2 cycles of movement

** 2 seconds at sample rate of 62.5 Hz -> 125 samples

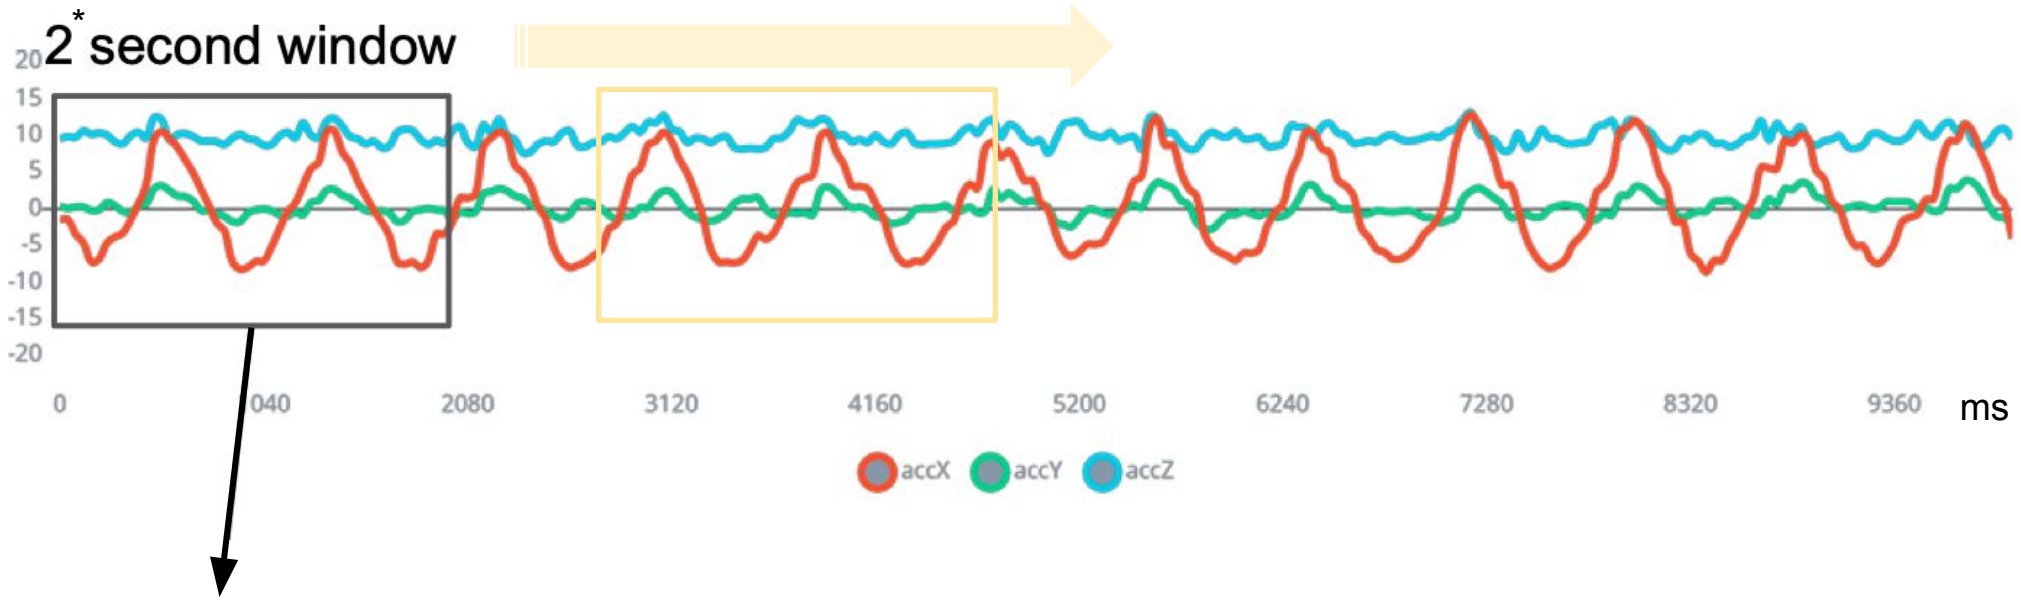

Raw Features as a window

- 125** samples for each axis
- 375 total features

Automatic Feature Extraction using DL

- Computational complexity
- Lots of training data

Problem!

✔ More Memory

* 2 seconds is needed to capture 1 or 2 cycles of movement

** 2 seconds at sample rate of 62.5 Hz -> 125 samples

Manual Feature Extraction

Data Pre-Processing

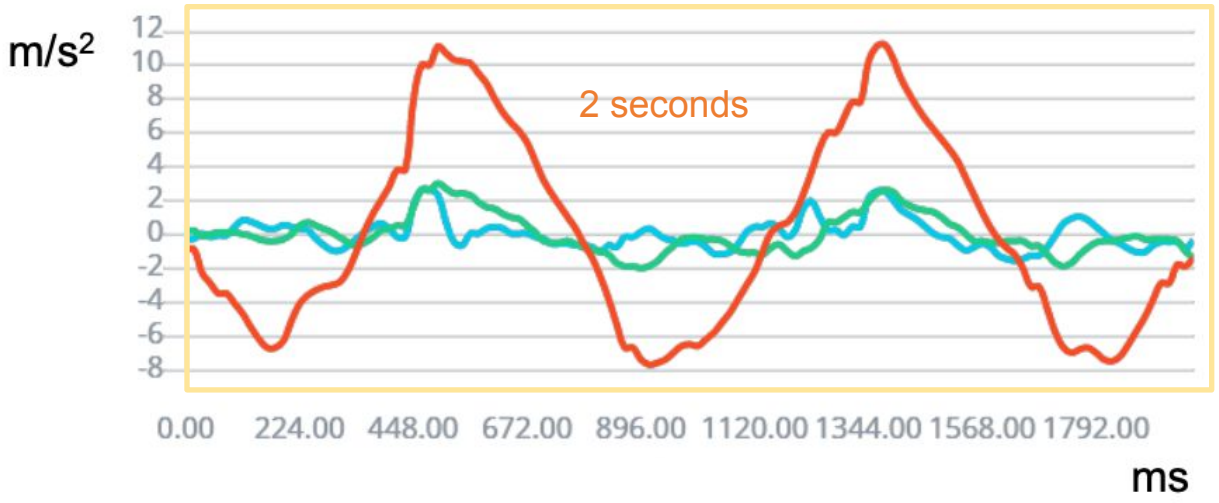

Manual Feature Extraction

3 RMS (Root Mean Square) values

- one for each axis (x, y, z)

$$x_{\mathrm{RMS}} = \sqrt{\frac{1}{n}\left(x_1^2 + x_2^2 + \cdots + x_n^2\right)}.$$

125

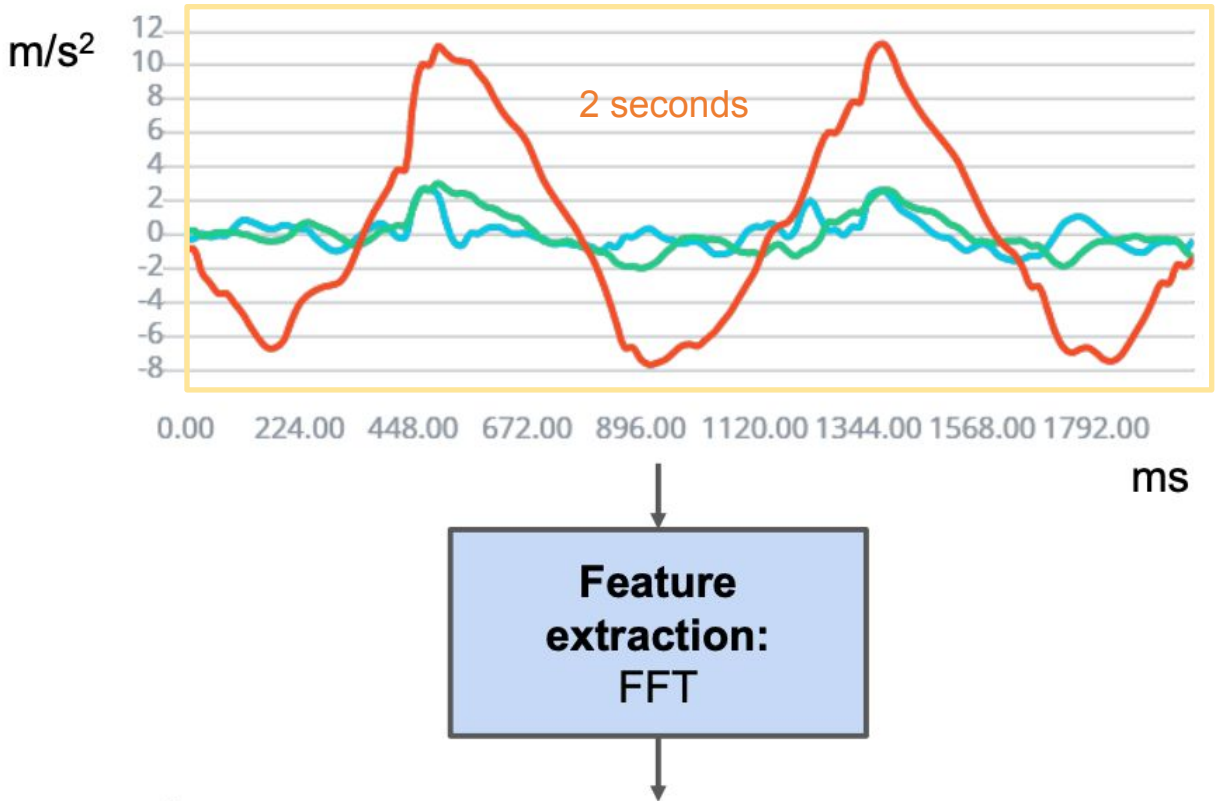

Manual Feature Extraction

3 RMS

Fast Fourier Transformer (FFT)

Raw Signal

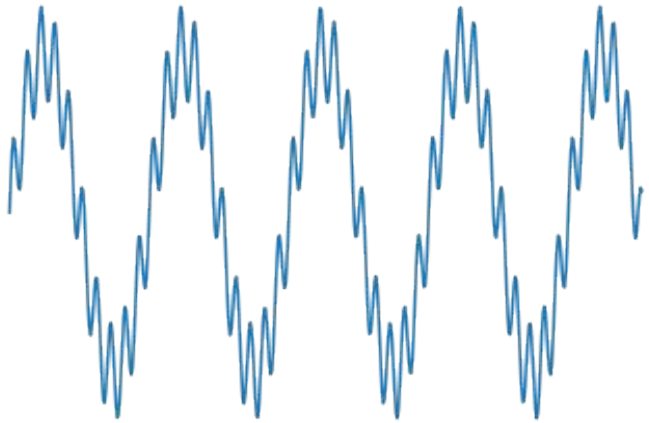

Time Domain

Fast Fourier Transformer (FFT)

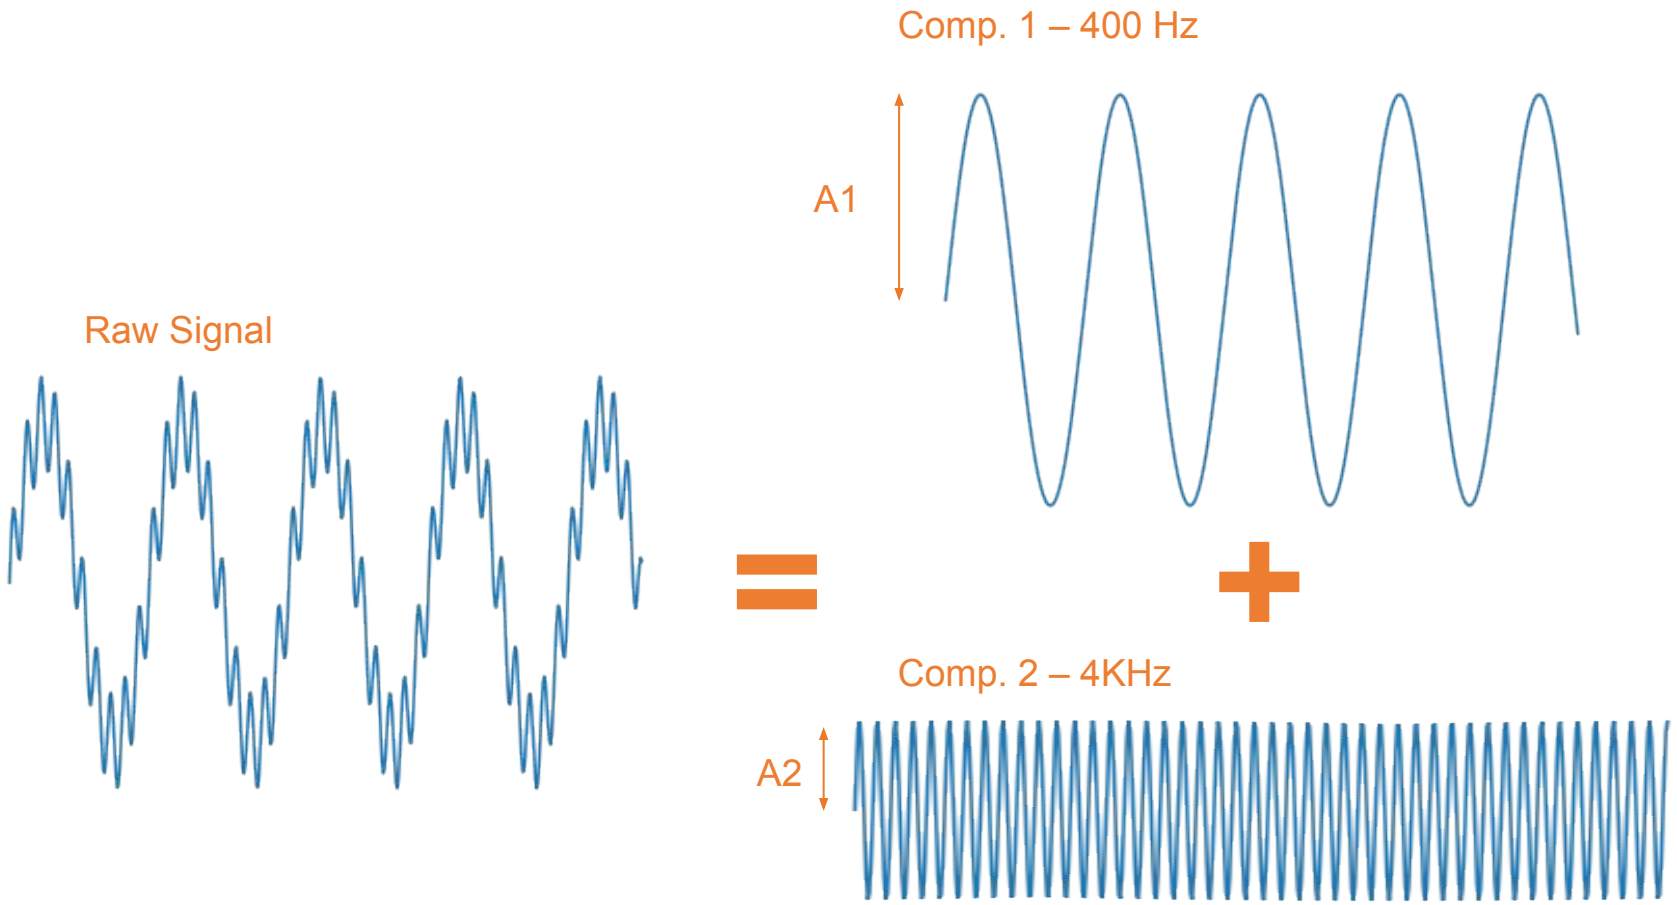

Fast Fourier Transformer (FFT)

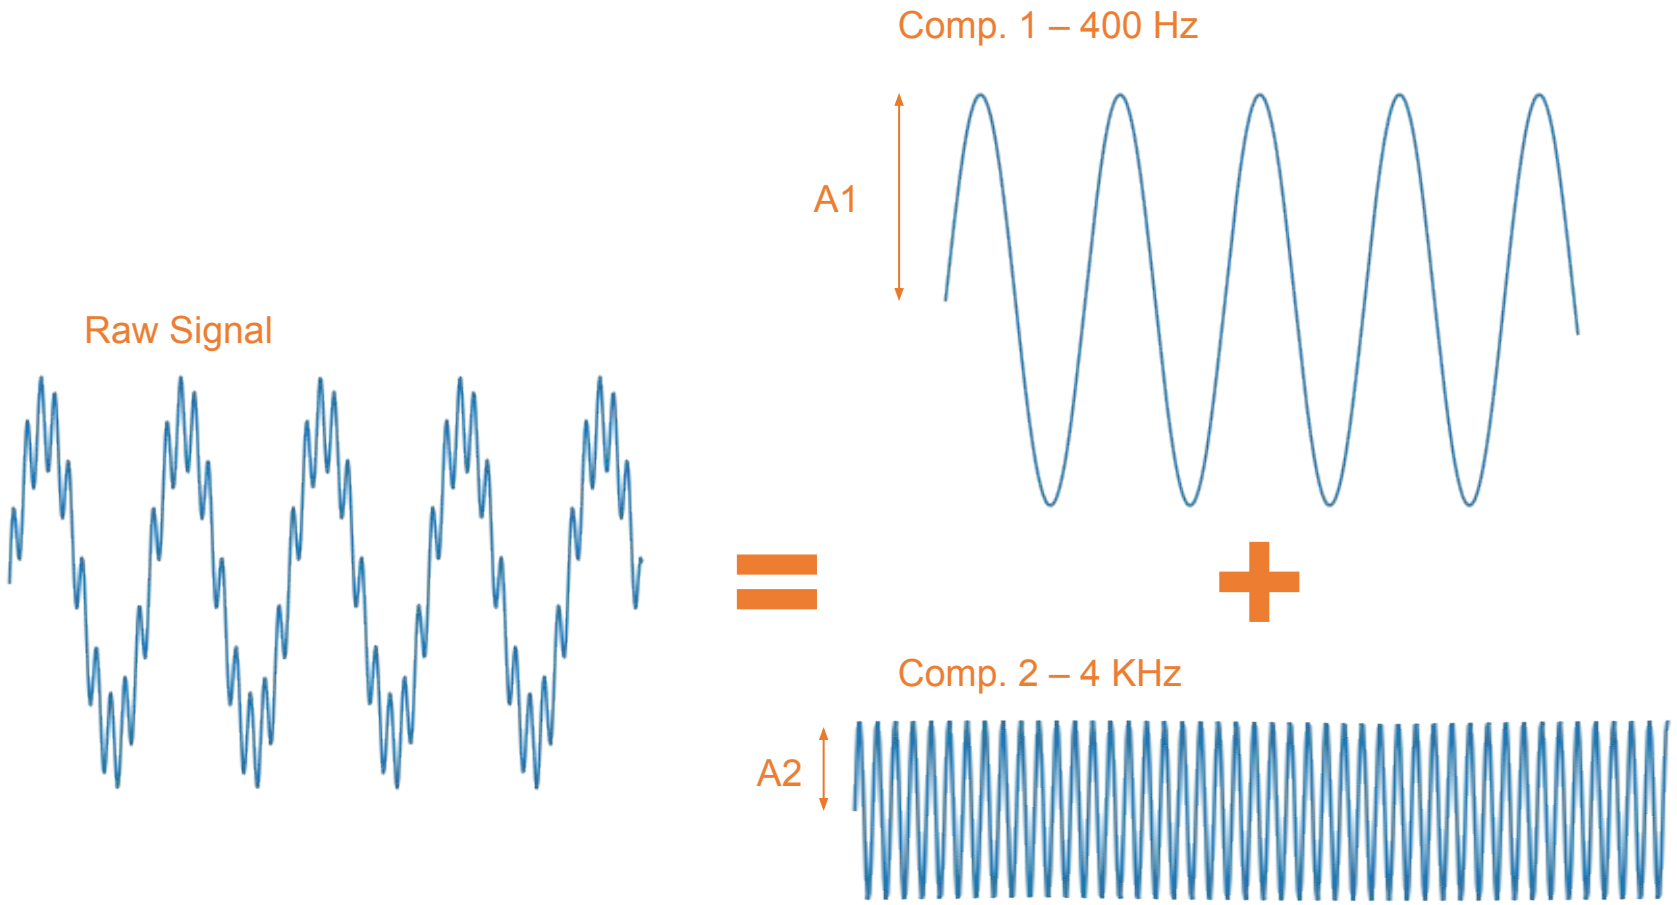

Time Domain

```
from scipy.fft import fft
yf = fft(raw signal)
plt.plot(xf, np.abs(yf));
```

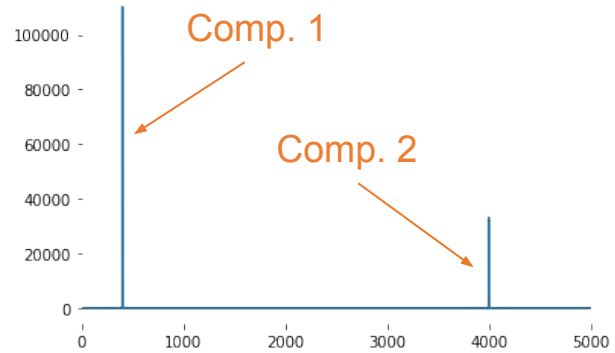

Frequency Domain

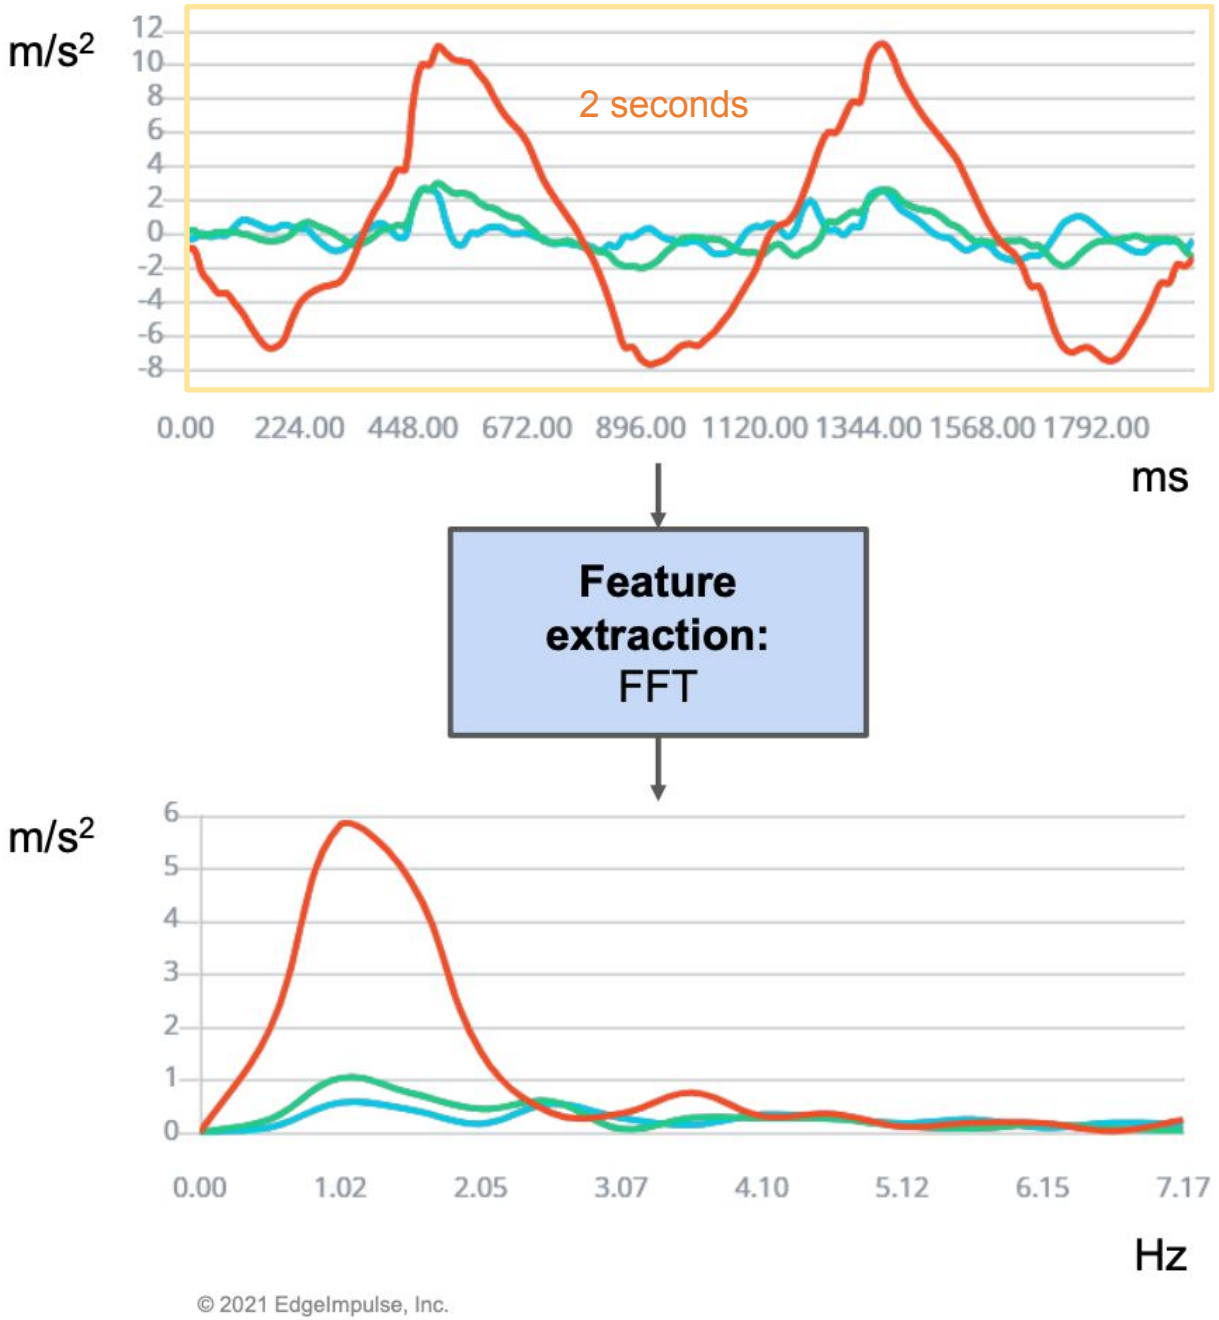

Manual Feature Extraction

3 RMS

Manual Feature Extraction

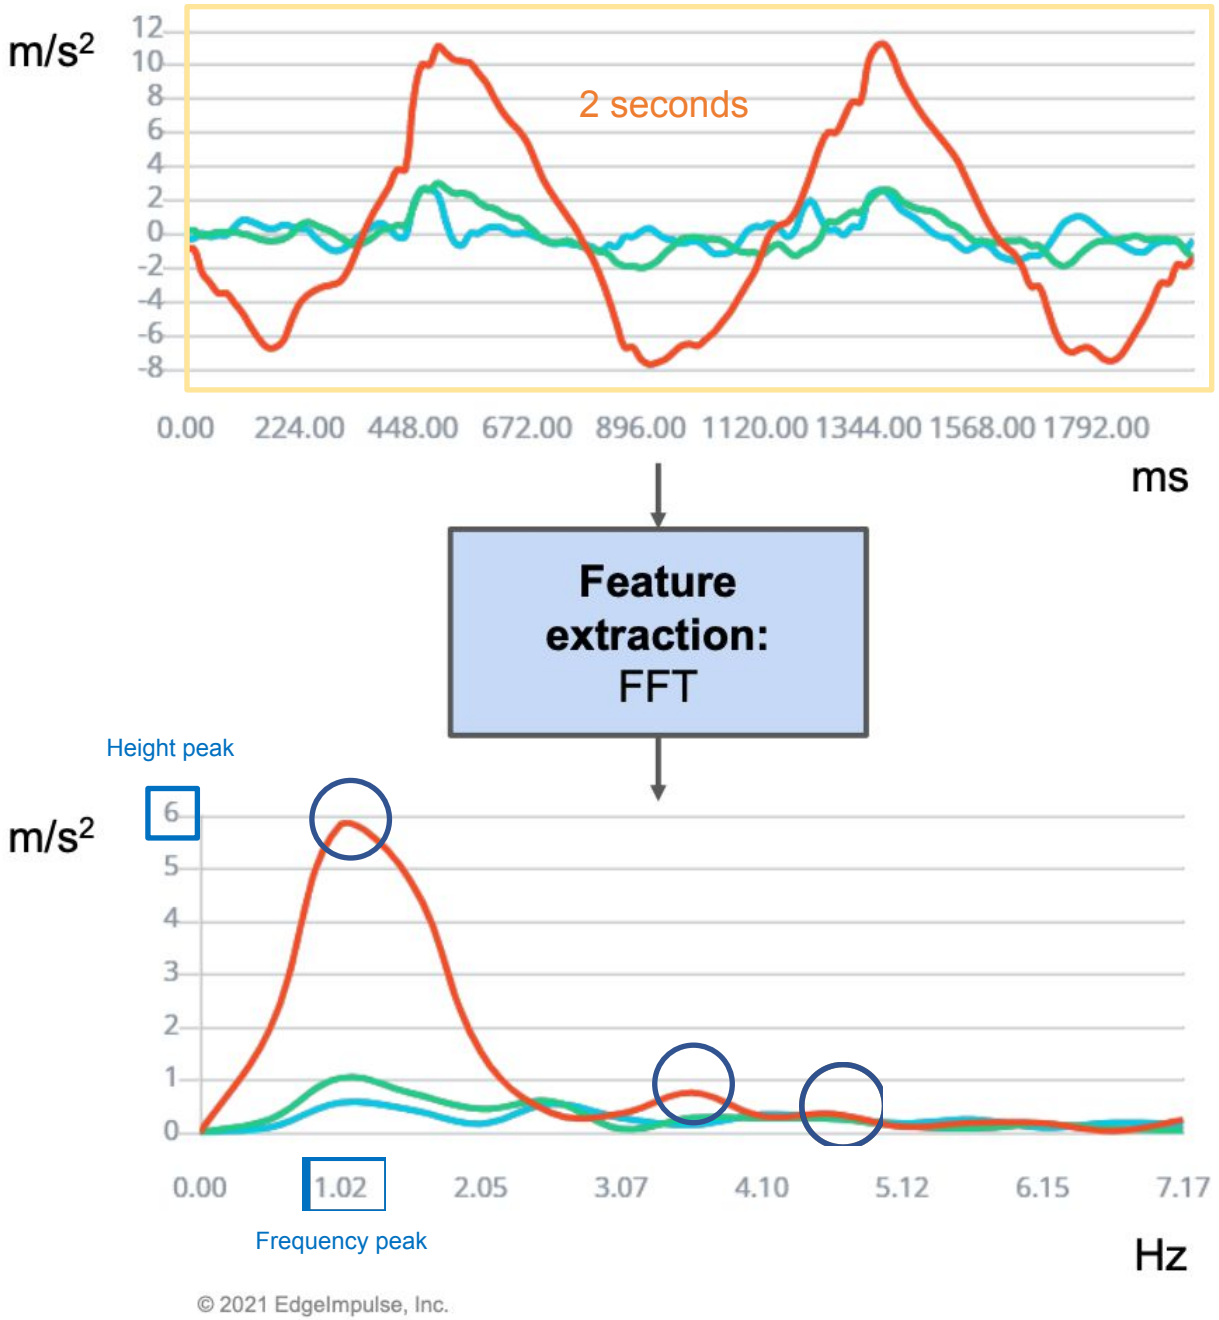

3 RMS

9 Height + **9** Freq. peak values

Manual Feature Extraction

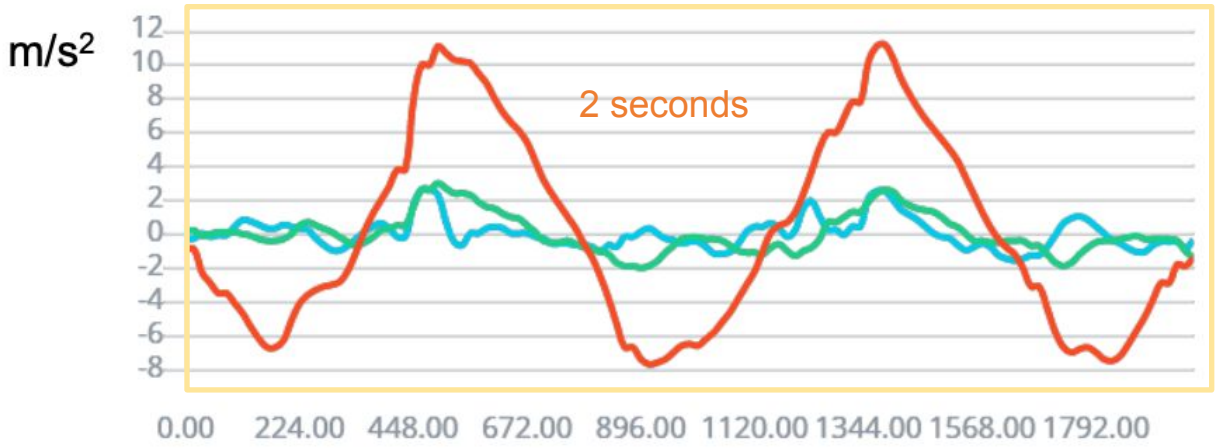

Feature extraction: FFT

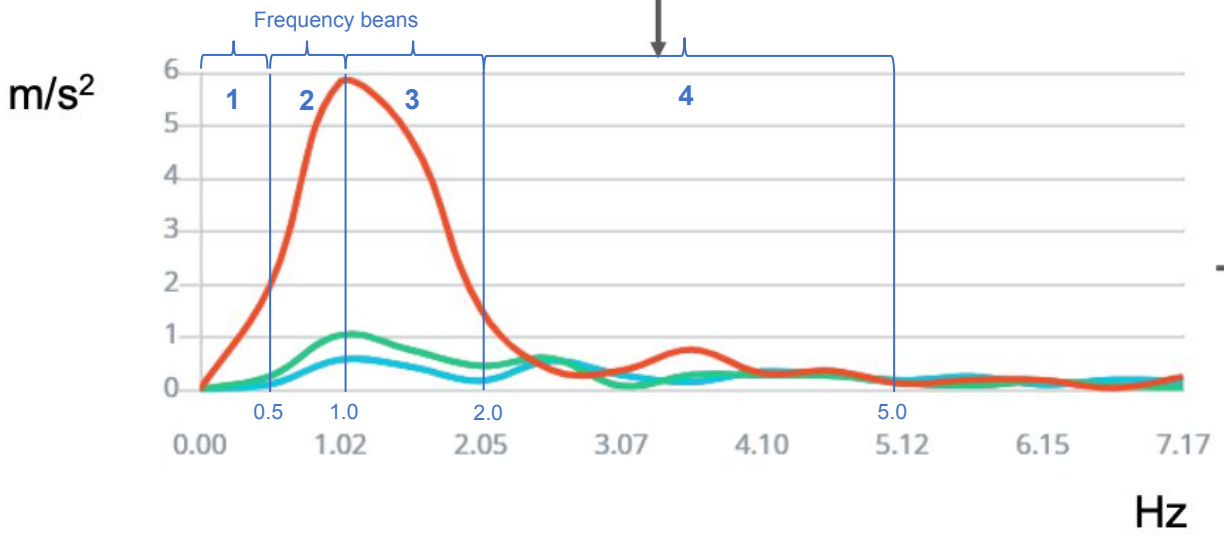

3 RMS + **9** HP + **9** FP + **12** PSD values

4 Frequency bins per axis

Feature extraction: PSD

https://blog.endaq.com/why-the-power-spectral-density-psd-is-the-gold-standard-of-vibration-analysis

Sliding Window: 80 ms

375 Raw Features

- Raw Data from sensor

Manual Feature Extraction

Spectral Analysis

11 Features

11 Features

11 Features

accX RMS
accX Peak
accX Peak
accX Peak
accX Peak
accX Peak
accX Spec
accX Spec
accX Spec
accX Spec

accY RMS
accY Peak
accY Peak
accY Peak
accY Peak
accY Peak
accY Spec
accY Spec
accY Spec
accY Spec

accZ RMS
accZ Peak 1 Freq
accZ Peak 1 Height
accZ Peak 2 Freq
accZ Peak 2 Height
accZ Peak 3 Freq
accZ Peak 3 Height
accZ Spectral Power 0.1 - 0.5
accZ Spectral Power 0.5 - 1.0
accZ Spectral Power 1.0 - 2.0
accZ Spectral Power 2.0 - 5.0

Sliding Window: 80 ms

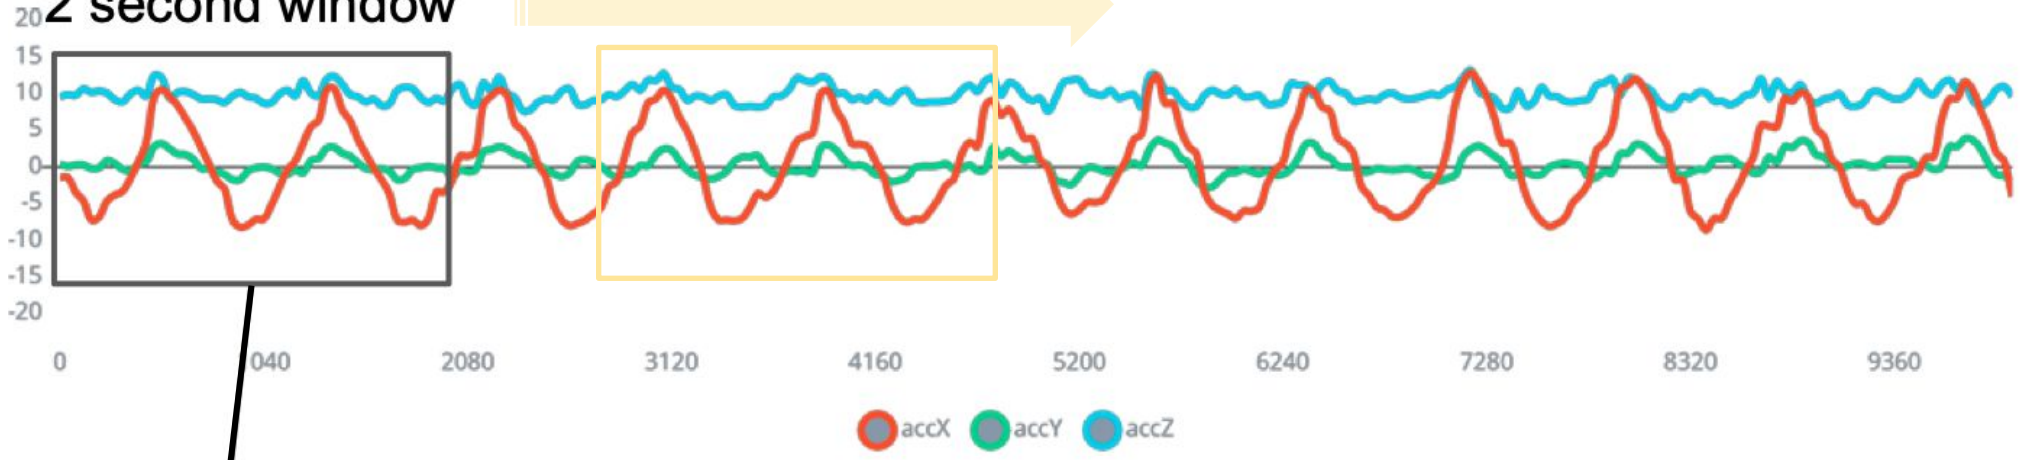

375 Raw Features

- Raw Data from sensor

Manual Feature Extraction

Spectral Analysis

33 Features

- RMS
- FFT
- PSD

NN Classifier

Classes

- Lift
- Terrestrial
- Maritime
- Idle

Model Design (NN Classifier)

Model Design (NN Classifier)

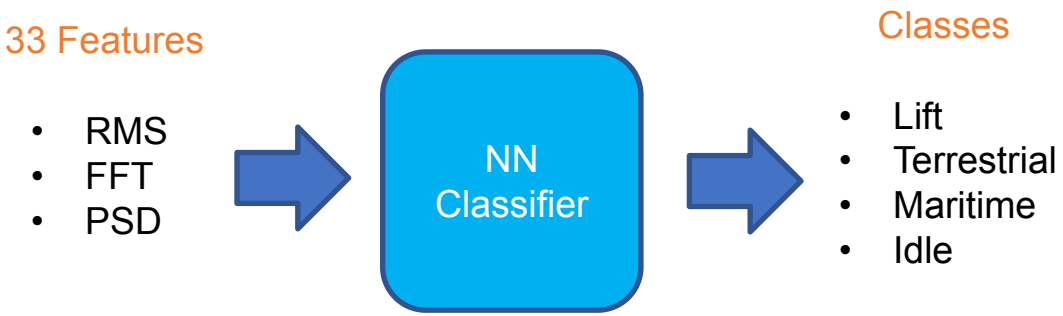

Model Design (DNN Classifier)

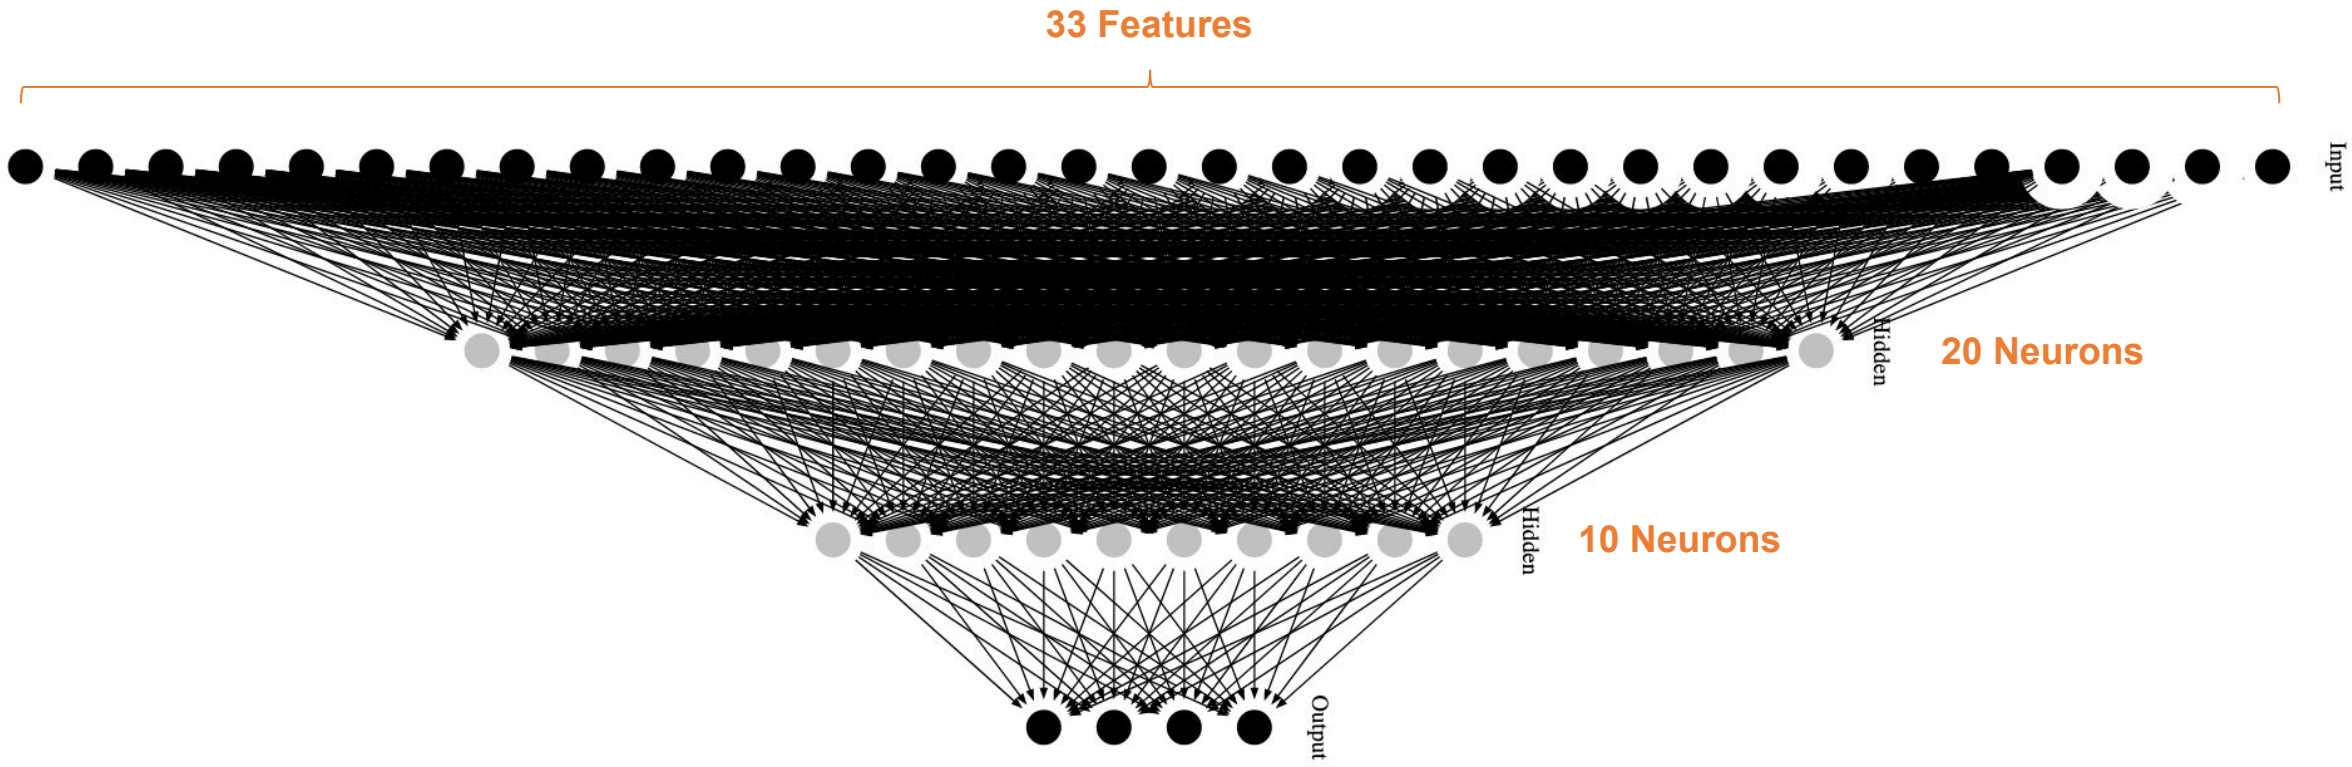

Model Design (DNN Classifier)

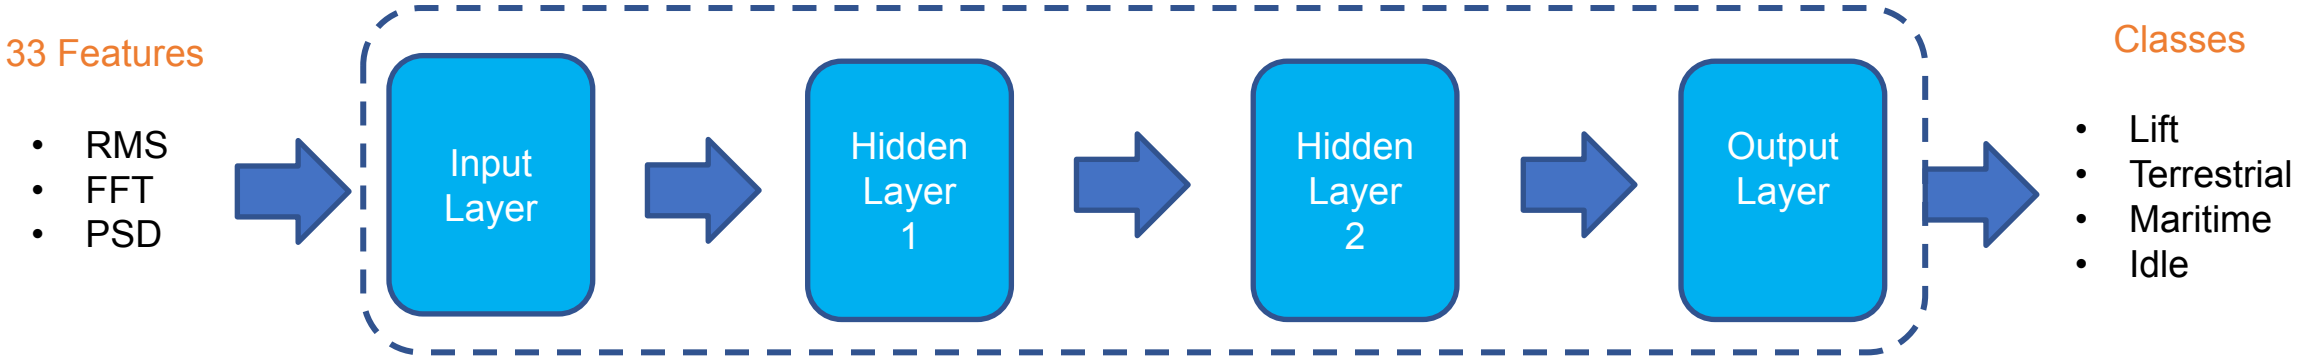

Model Design (DNN Classifier)

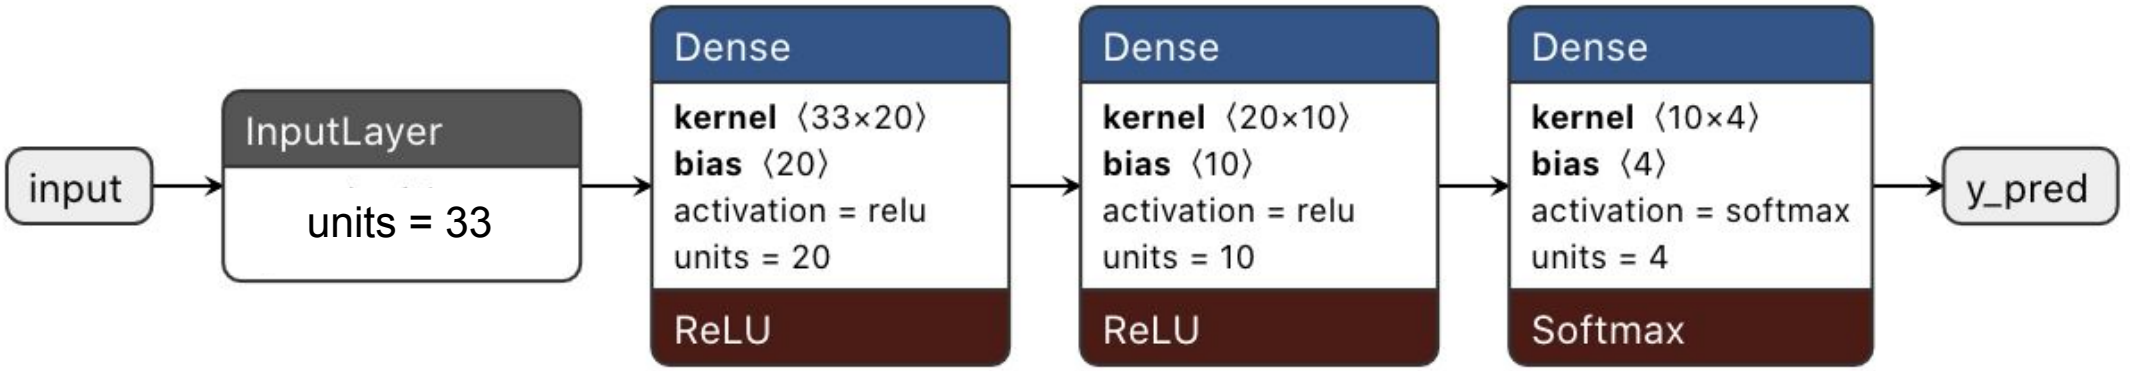

33 Features

- RMS
- FFT
- PSD

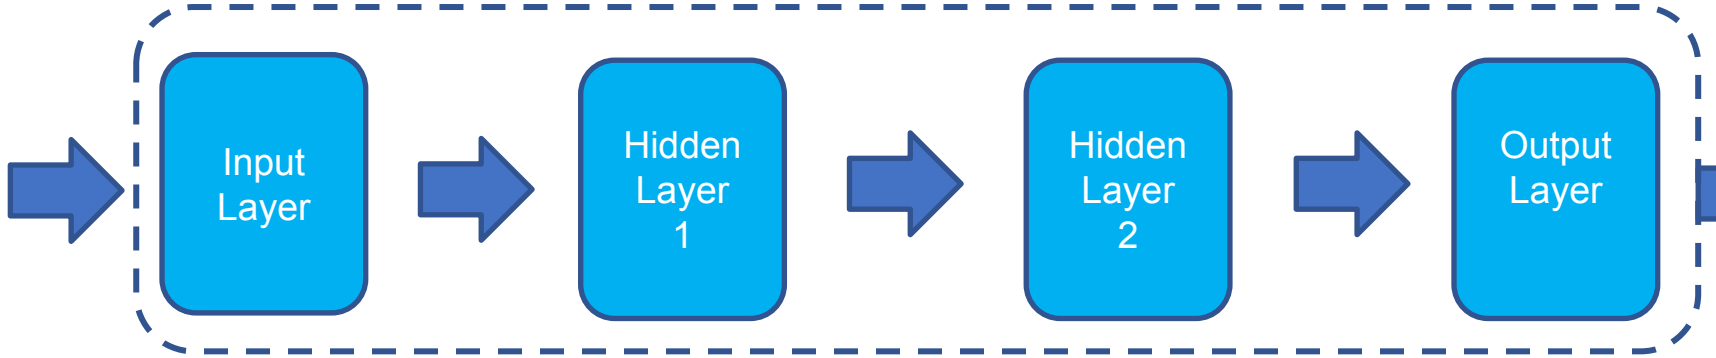

Classes

- Lift
- Terrestrial
- Maritime
- Idle

Train, Evaluate, Convert, Deploy the Model

Train, Evaluate, Convert, Deploy the Model

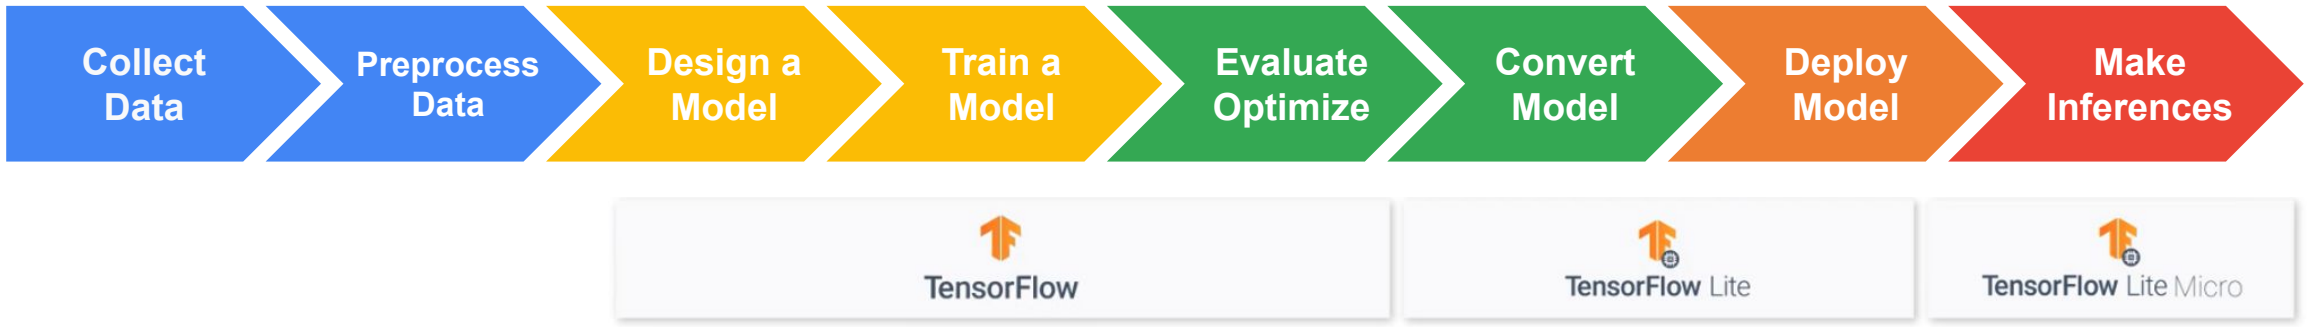

Machine Learning Workflow

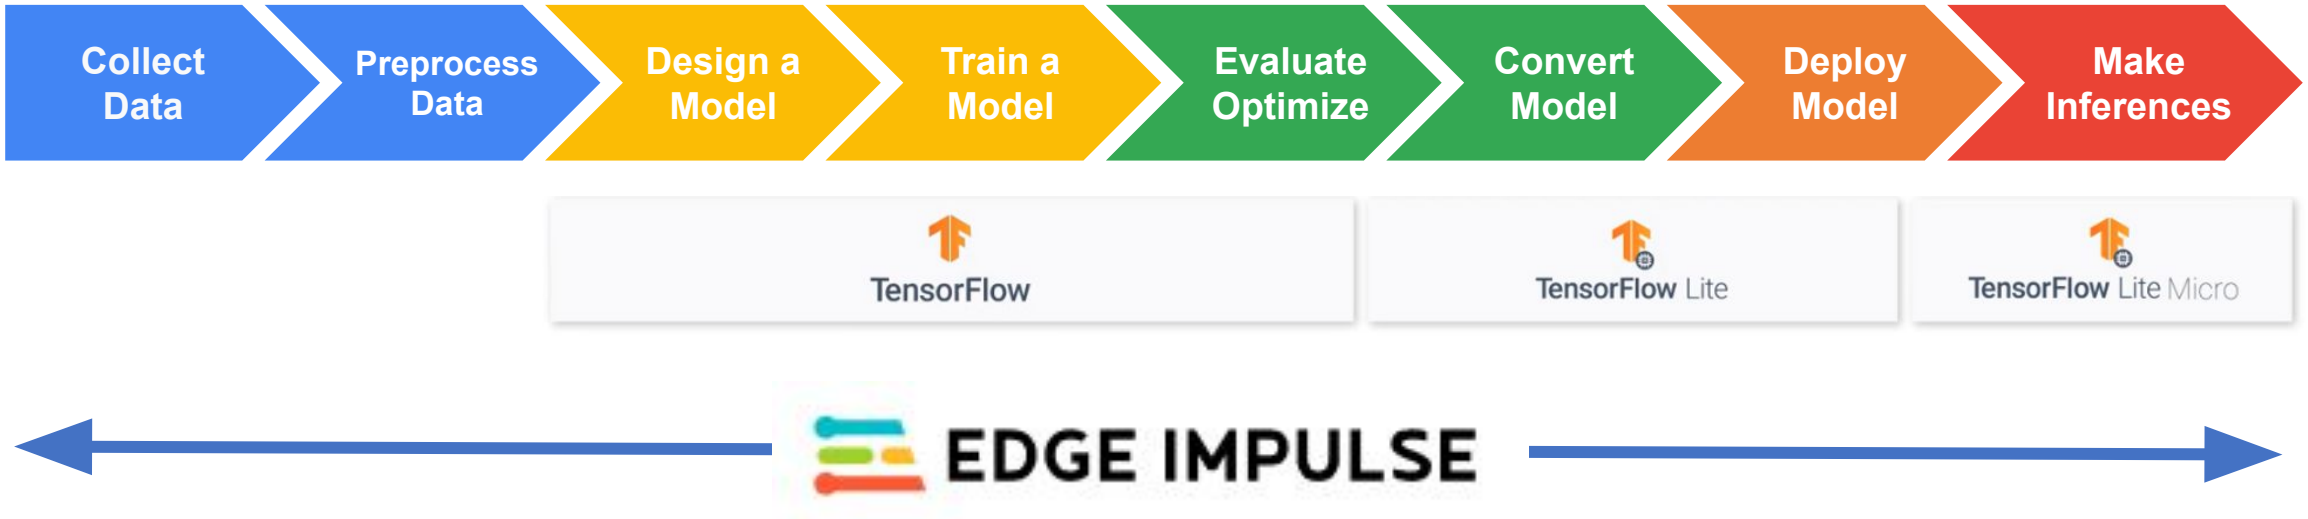

EI Studio - Embedded ML platform

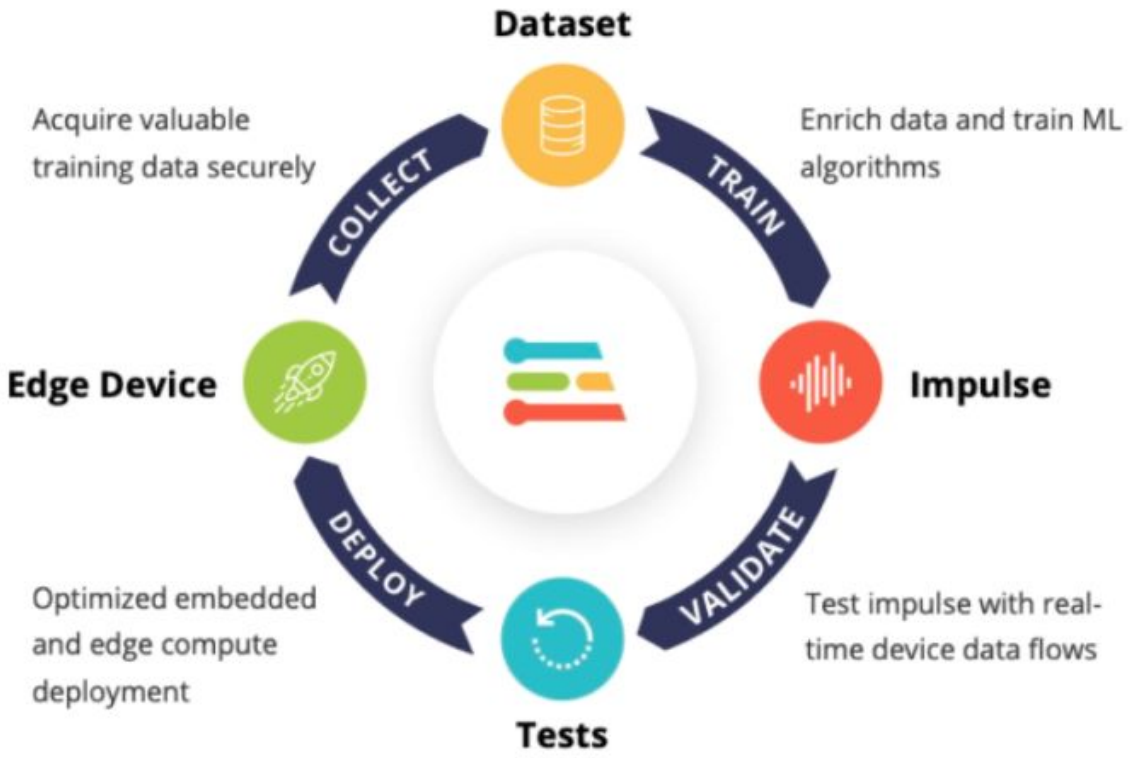

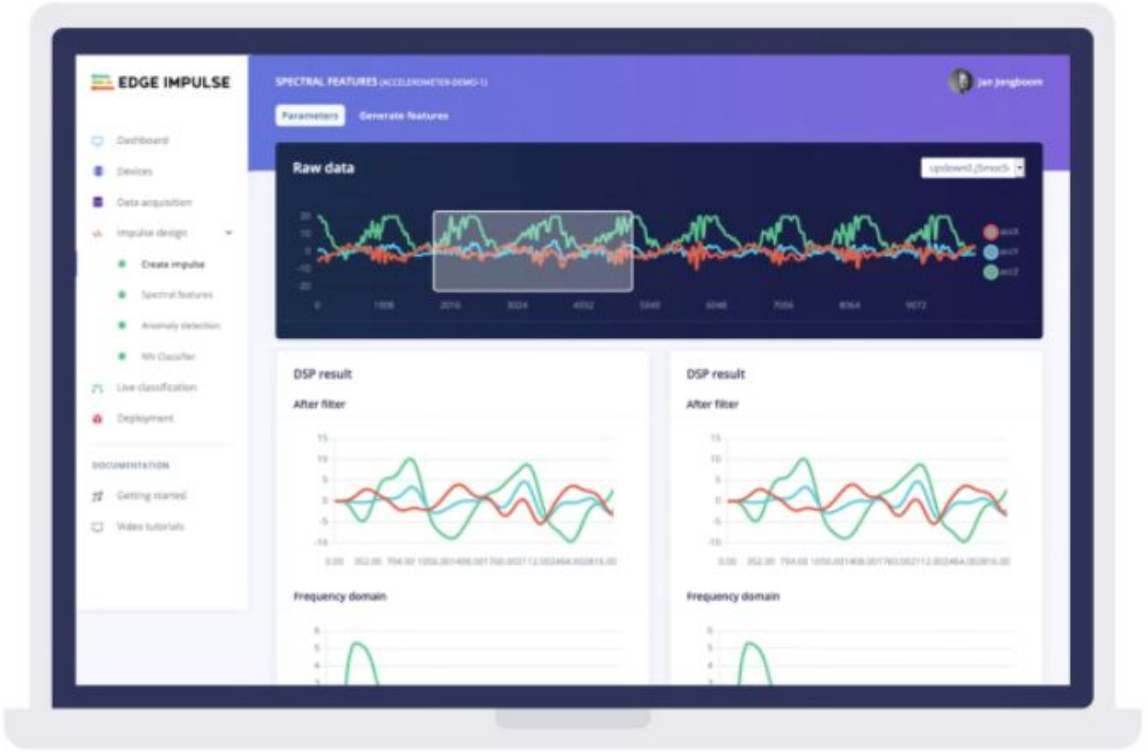

Learn more at http://edgeimpulse.com

Motion Classification

Transportation Classes:

- **lift** (up-down)
- **terrestrial** (left-right)
- **maritime** (zig-zag, etc.)
- **idle**

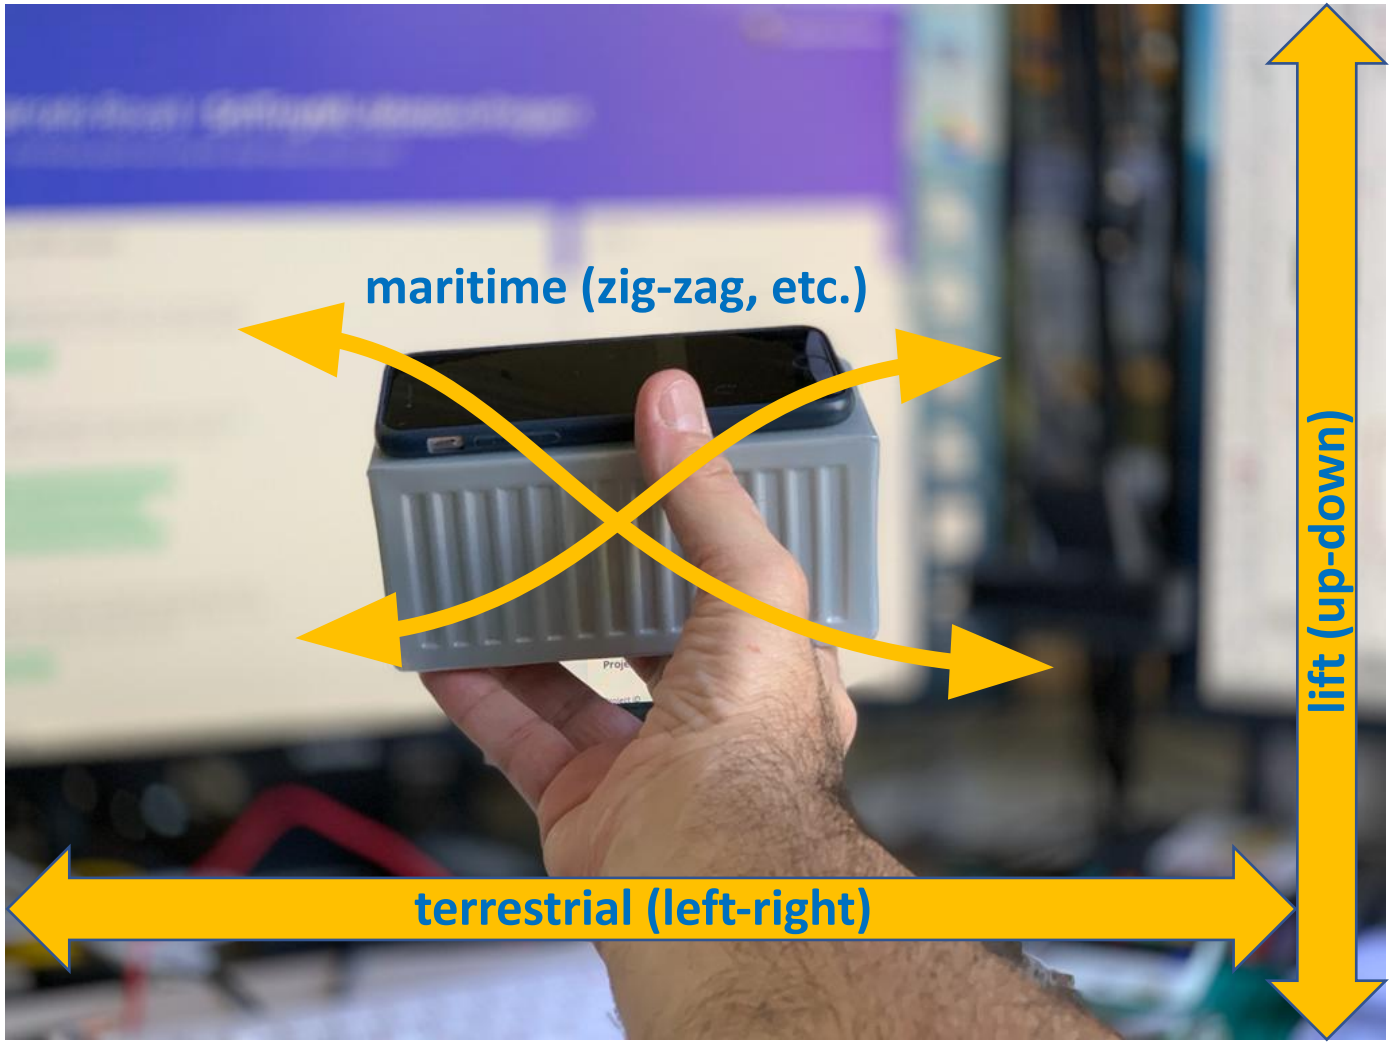

Motion Classification

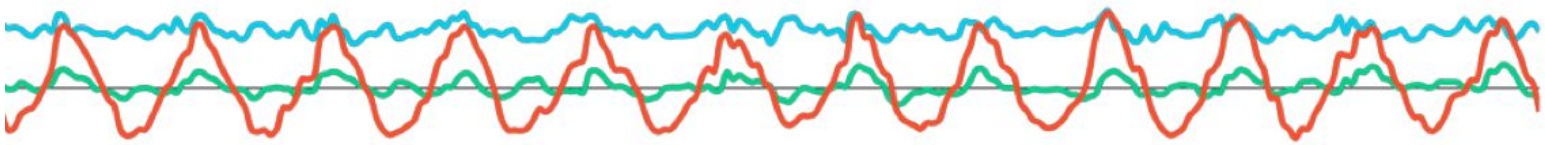

Transportation Classes

- **lift** (up-down)
- **terrestrial** (left-right)
- **maritime** (zig-zag, etc.)
- **idle**

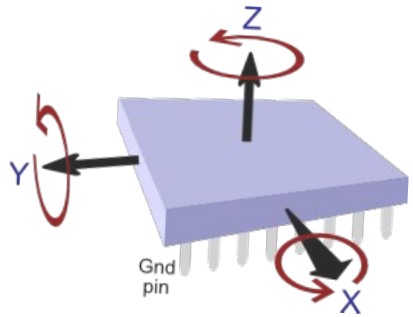

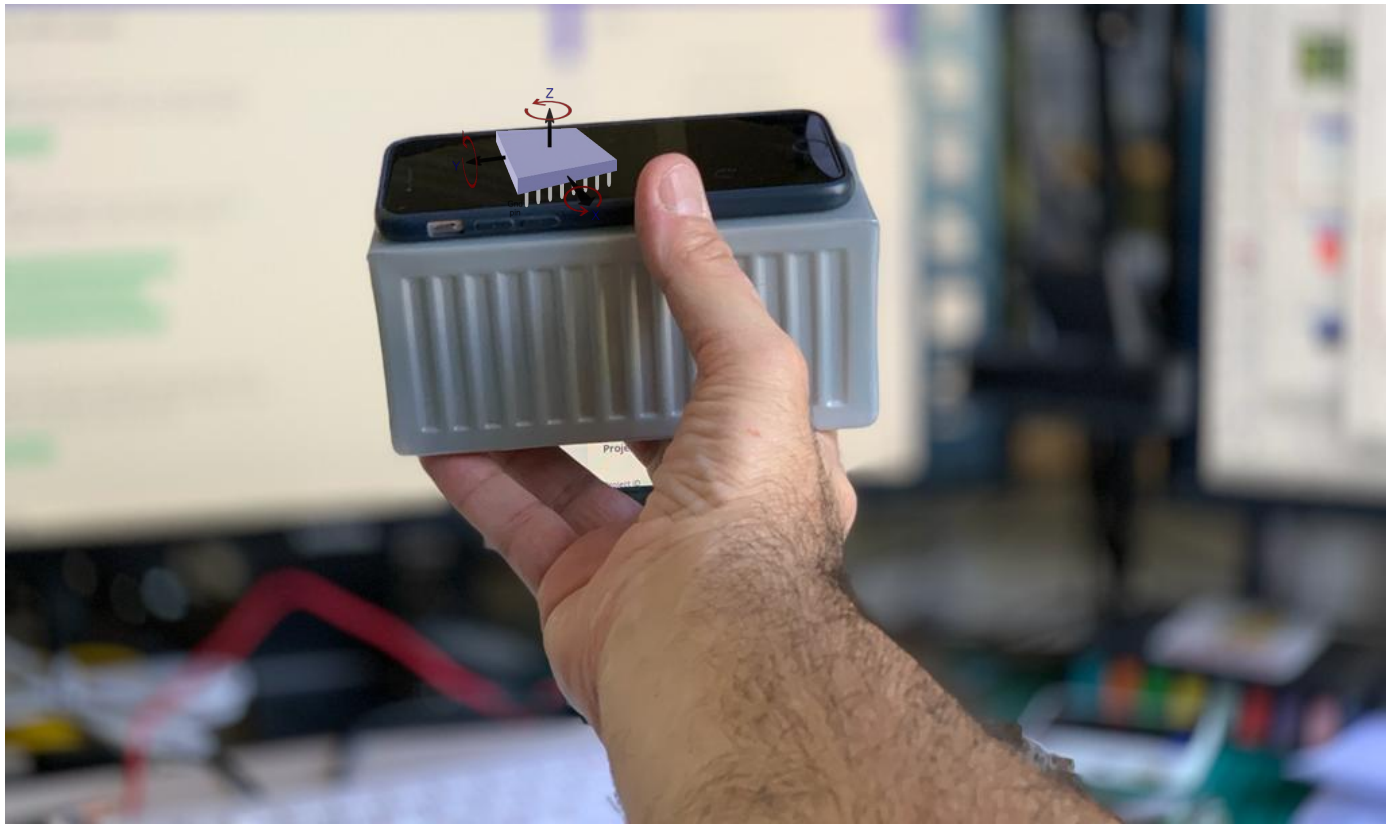

Data: collect & test using accelerometer as sensor

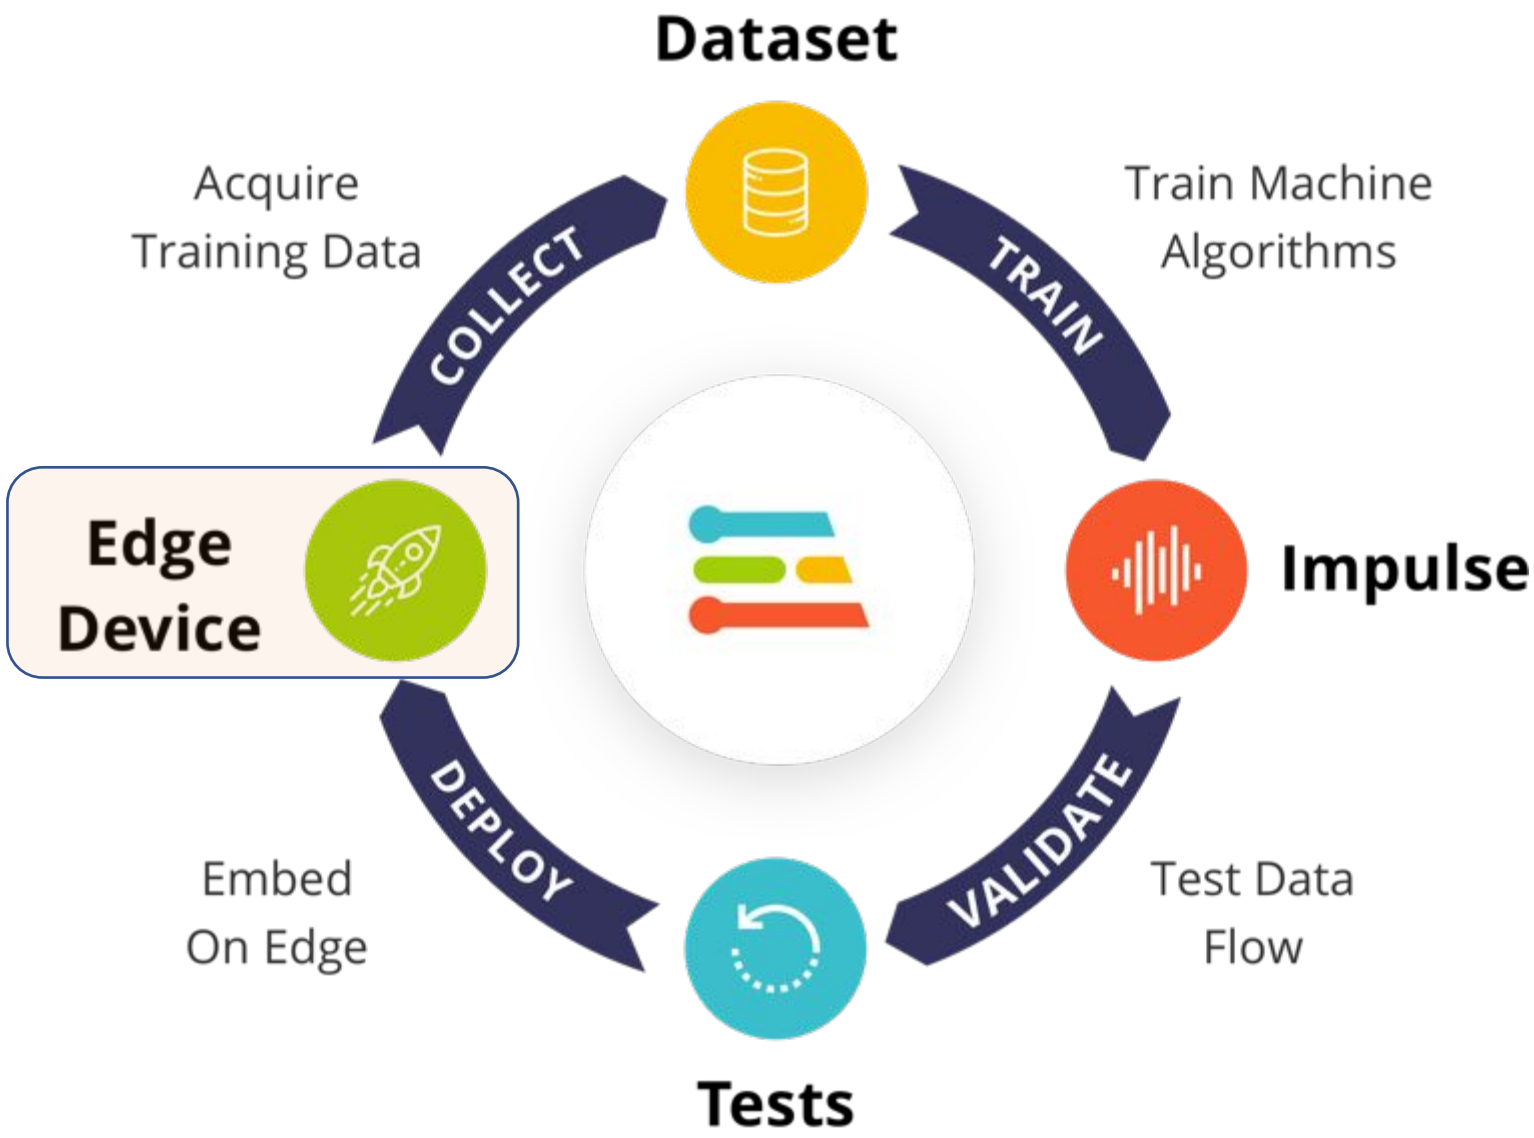

- Pre-Processing Data
- Design a Model
- Train a Model

Sensor - IMU (Inertial Measurement Unit)

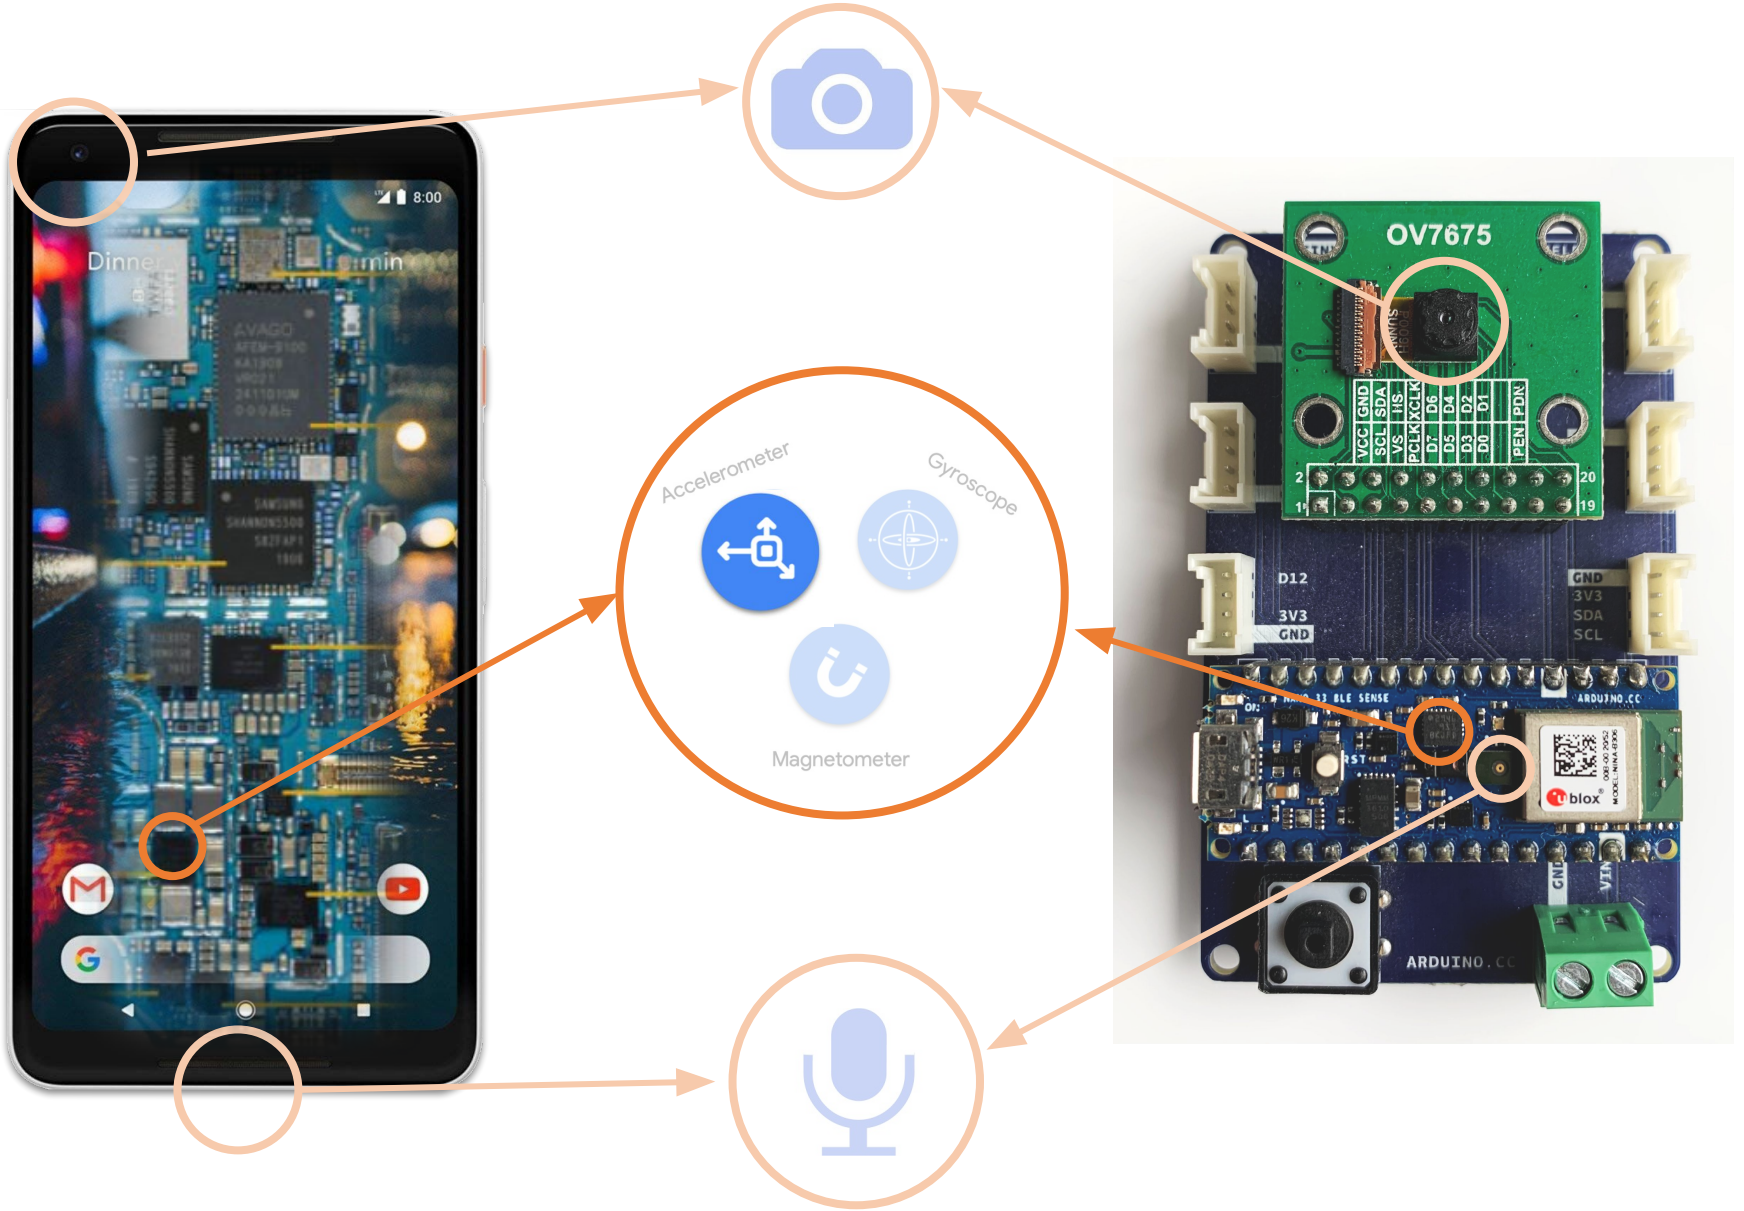

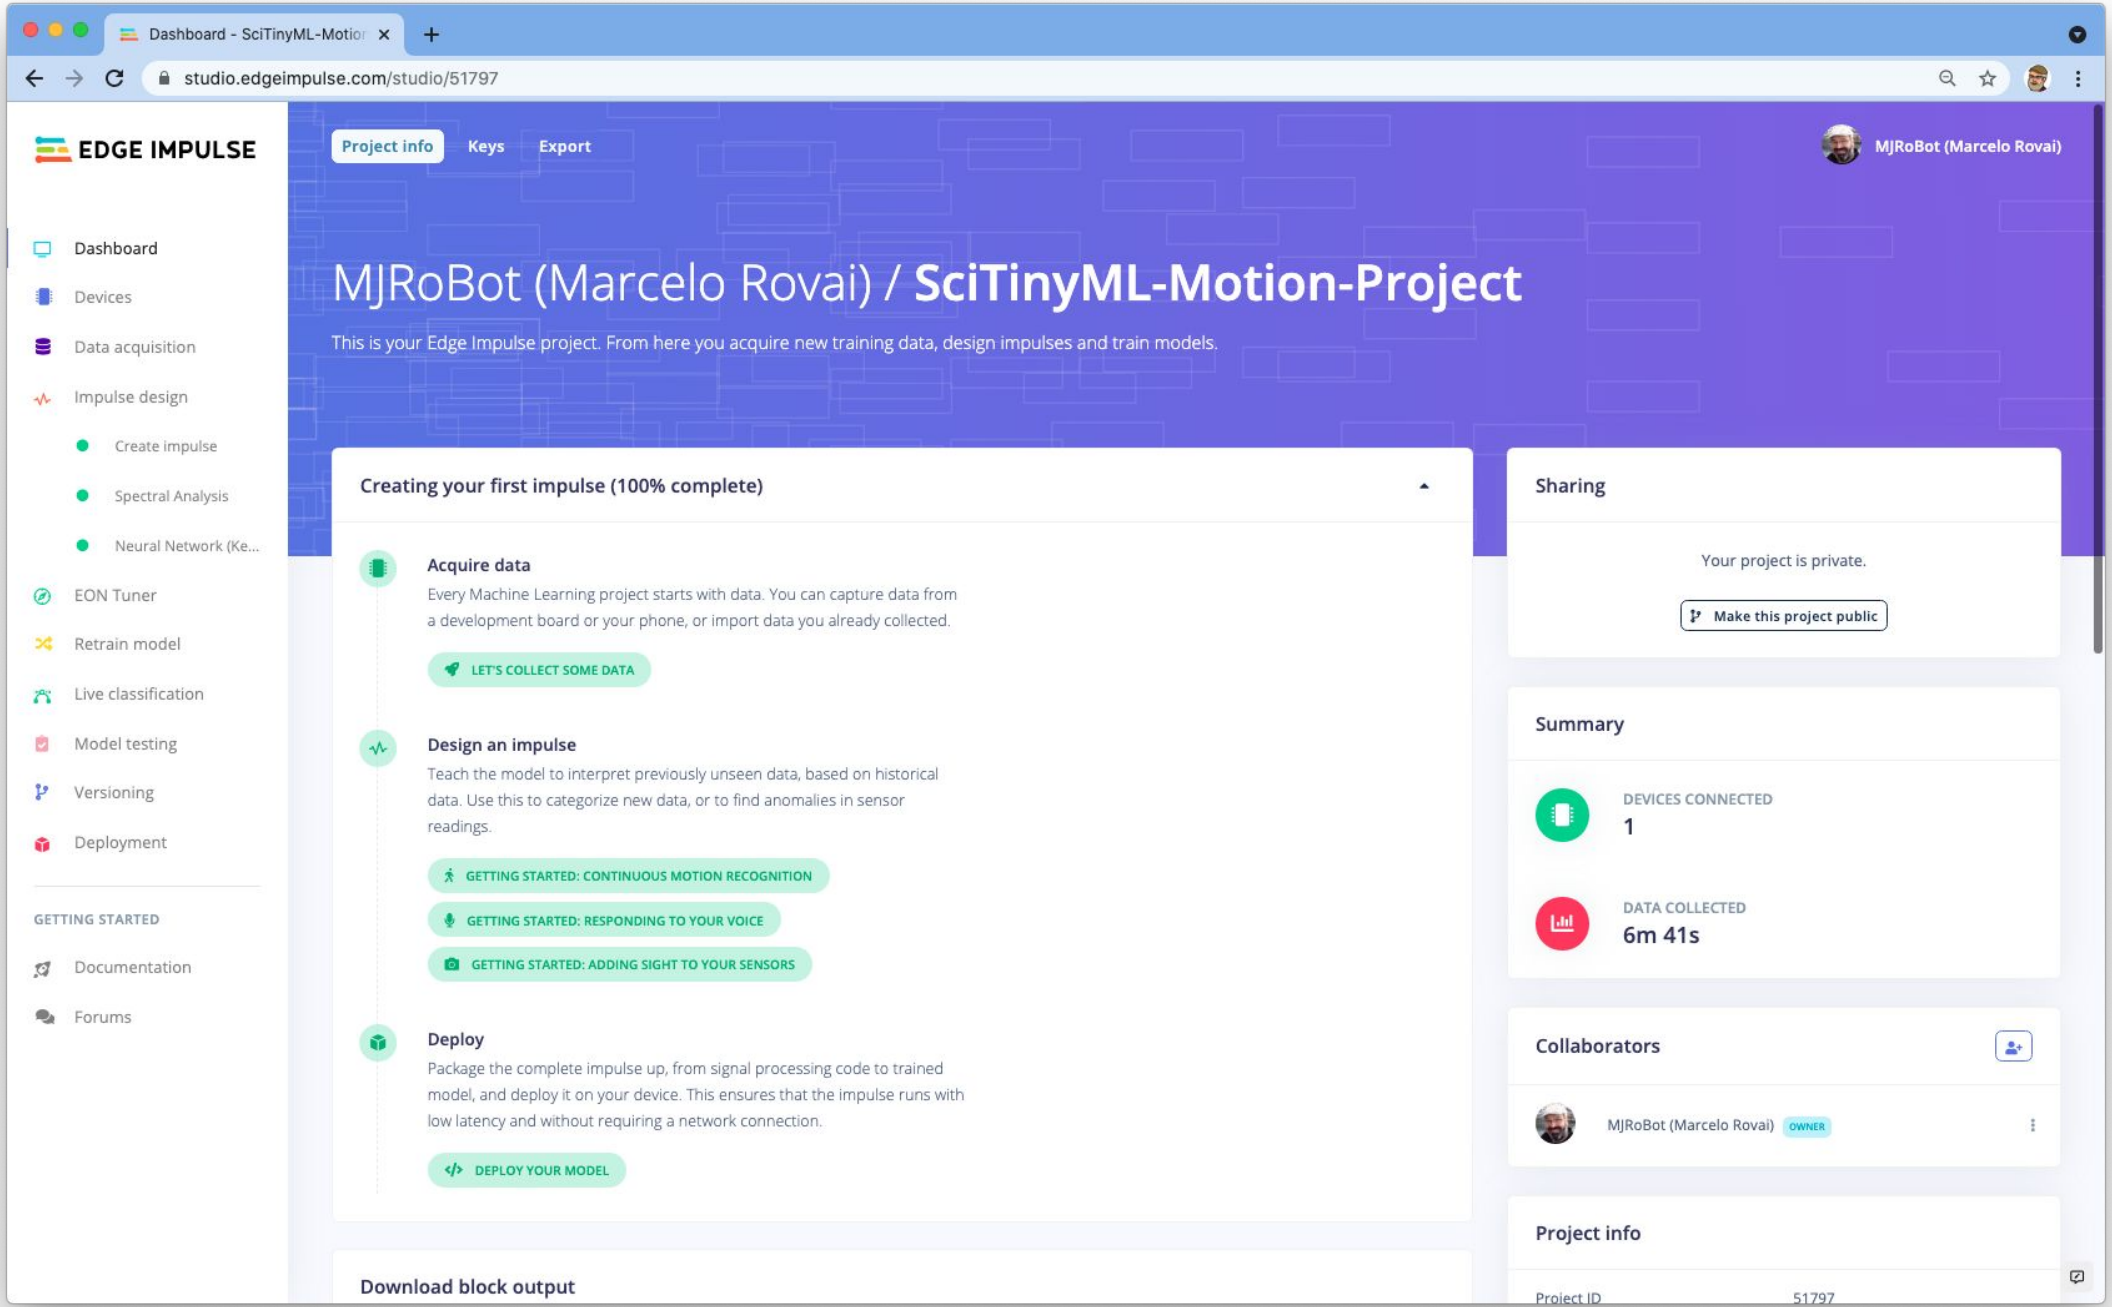
Dashboard - SciTinyML-Motion
studio.edgeimpulse.com/studio/51797
EDGE IMPULSE
Project info
Keys
Export
MJRoBot (Marcelo Rovai)
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Spectral Analysis
Neural Network (Ke...
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
MJRoBot (Marcelo Rovai) / SciTinyML-Motion-Project
This is your Edge Impulse project. From here you acquire new training data, design impulses and train models.
Creating your first impulse (100% complete)
Acquire data
Every Machine Learning project starts with data. You can capture data from a development board or your phone, or import data you already collected.
LET'S COLLECT SOME DATA
Design an impulse
Teach the model to interpret previously unseen data, based on historical data. Use this to categorize new data, or to find anomalies in sensor readings.
GETTING STARTED: CONTINUOUS MOTION RECOGNITION
GETTING STARTED: RESPONDING TO YOUR VOICE
GETTING STARTED: ADDING SIGHT TO YOUR SENSORS
Deploy
Package the complete impulse up, from signal processing code to trained model, and deploy it on your device. This ensures that the impulse runs with low latency and without requiring a network connection.
DEPLOY YOUR MODEL
Download block output
Sharing
Your project is private.
Make this project public
Summary
DEVICES CONNECTED
1
DATA COLLECTED
6m 41s
Collaborators
MJRoBot (Marcelo Rovai) OWNER
Project info
Project ID
51797

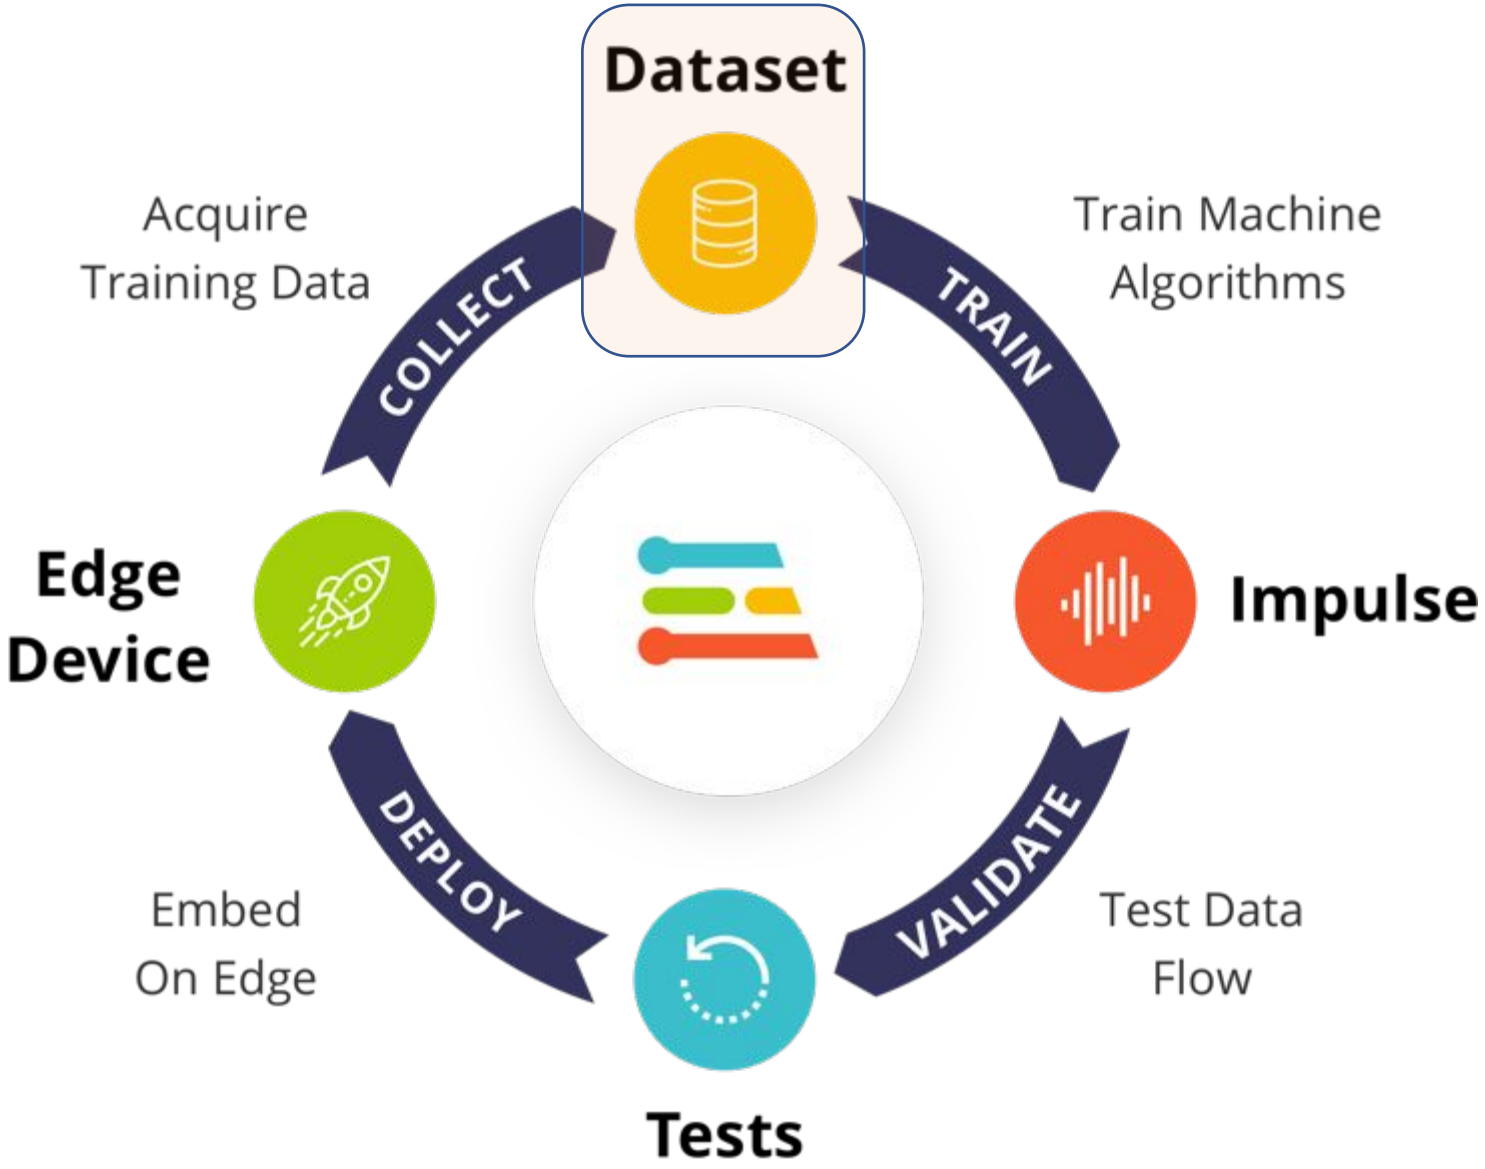

- Pre-Processing Data
- Design a Model
- Train a Model

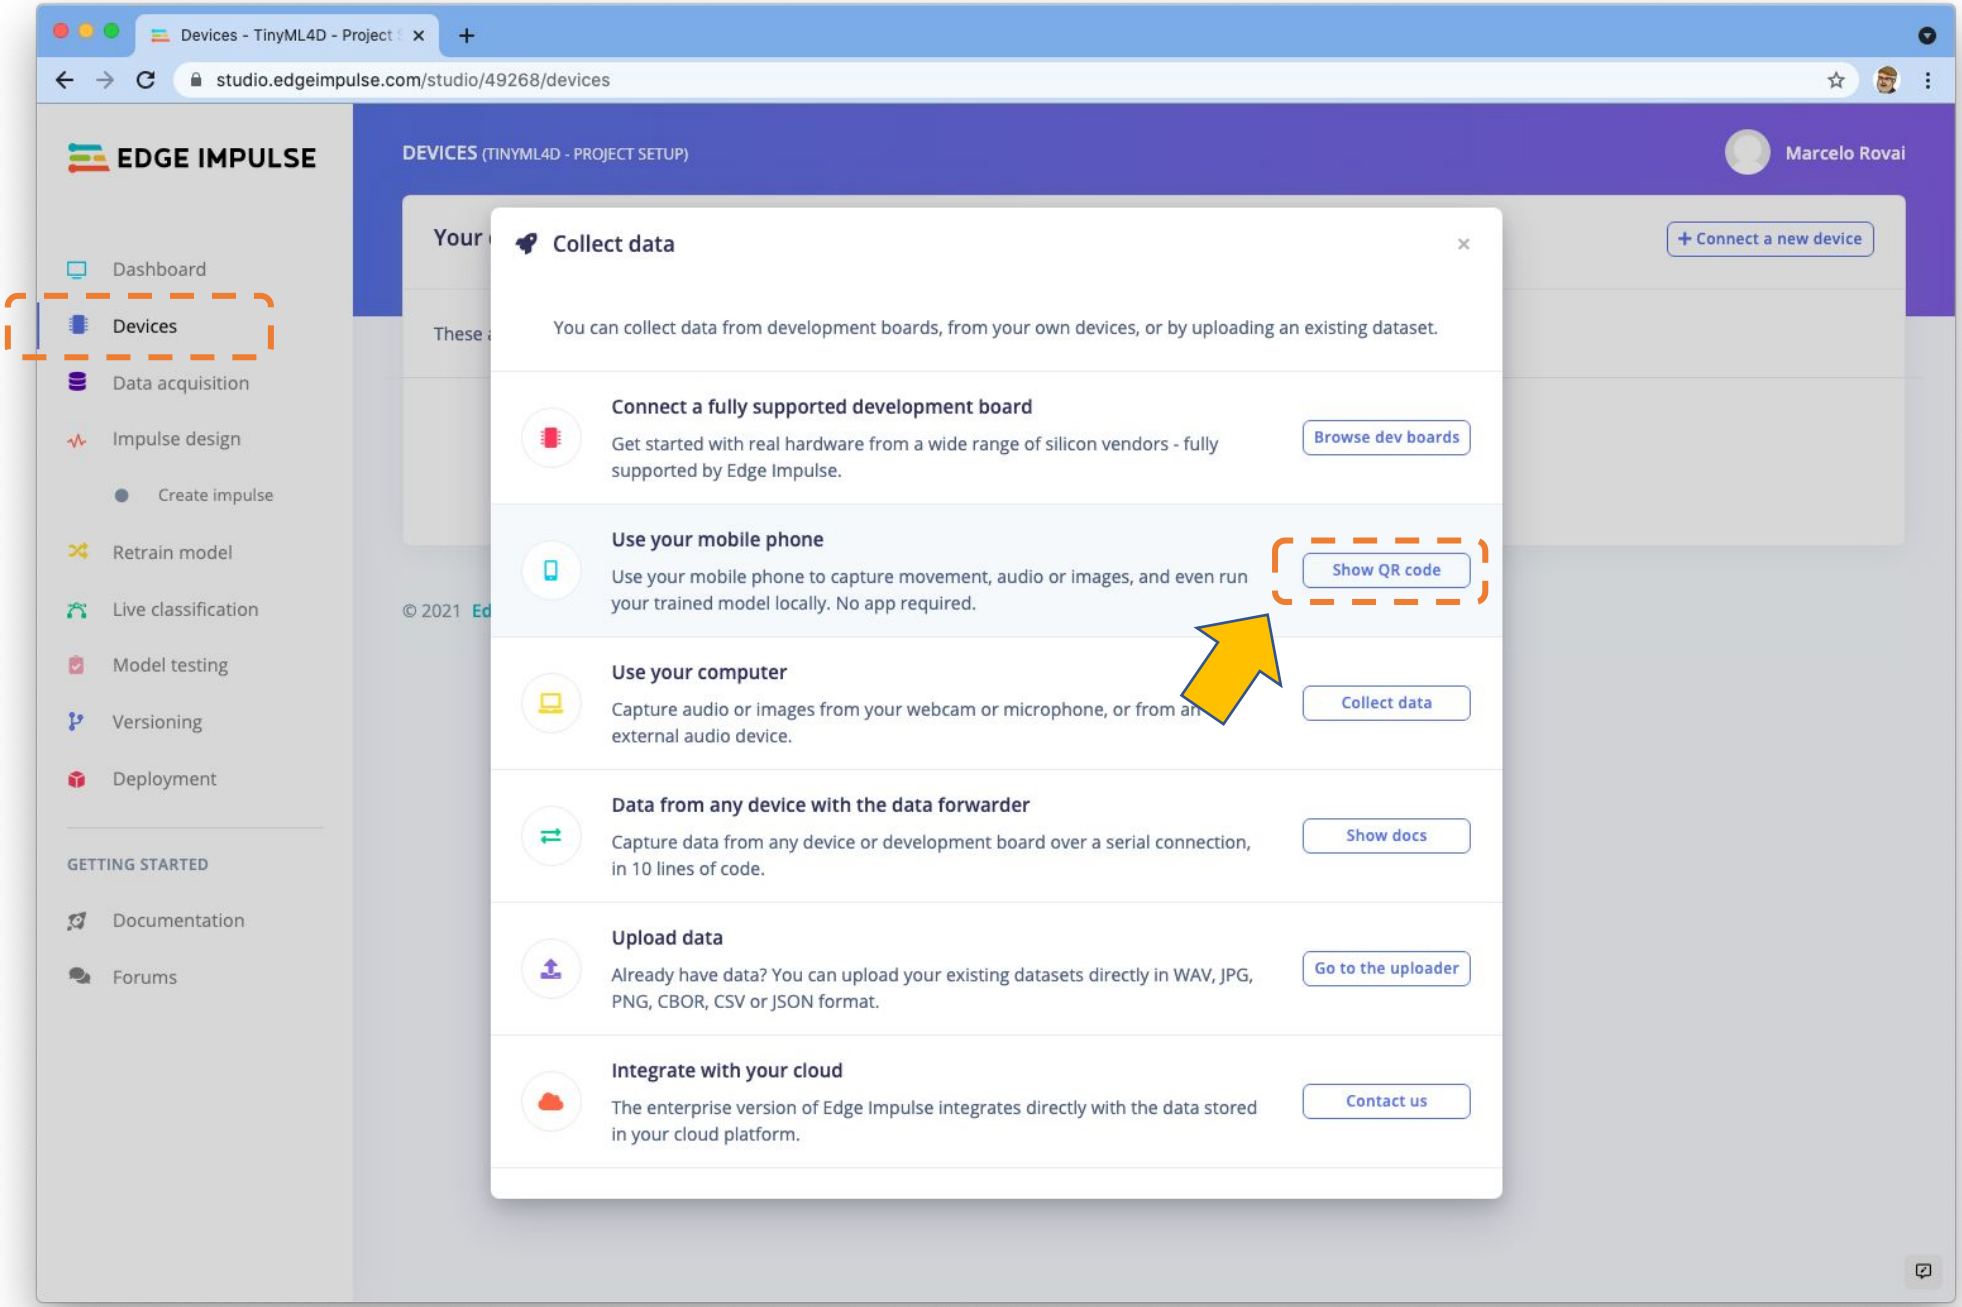
Devices - TinyML4D - Project
studio.edgeimpulse.com/studio/49268/devices
EDGE IMPULSE
DEVICES (TINYML4D - PROJECT SETUP)
Marcelo Rovai
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
+ Connect a new device
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
Use your mobile phone
Use your mobile phone to capture movement, audio or images, and even run your trained model locally. No app required.
Show QR code
Use your computer
Capture audio or images from your webcam or microphone, or from an external audio device.
Collect data
Data from any device with the data forwarder
Capture data from any device or development board over a serial connection, in 10 lines of code.
Show docs
Upload data
Already have data? You can upload your existing datasets directly in WAV, JPG, PNG, CBOR, CSV or JSON format.
Go to the uploader
Integrate with your cloud
The enterprise version of Edge Impulse integrates directly with the data stored in your cloud platform.
Contact us

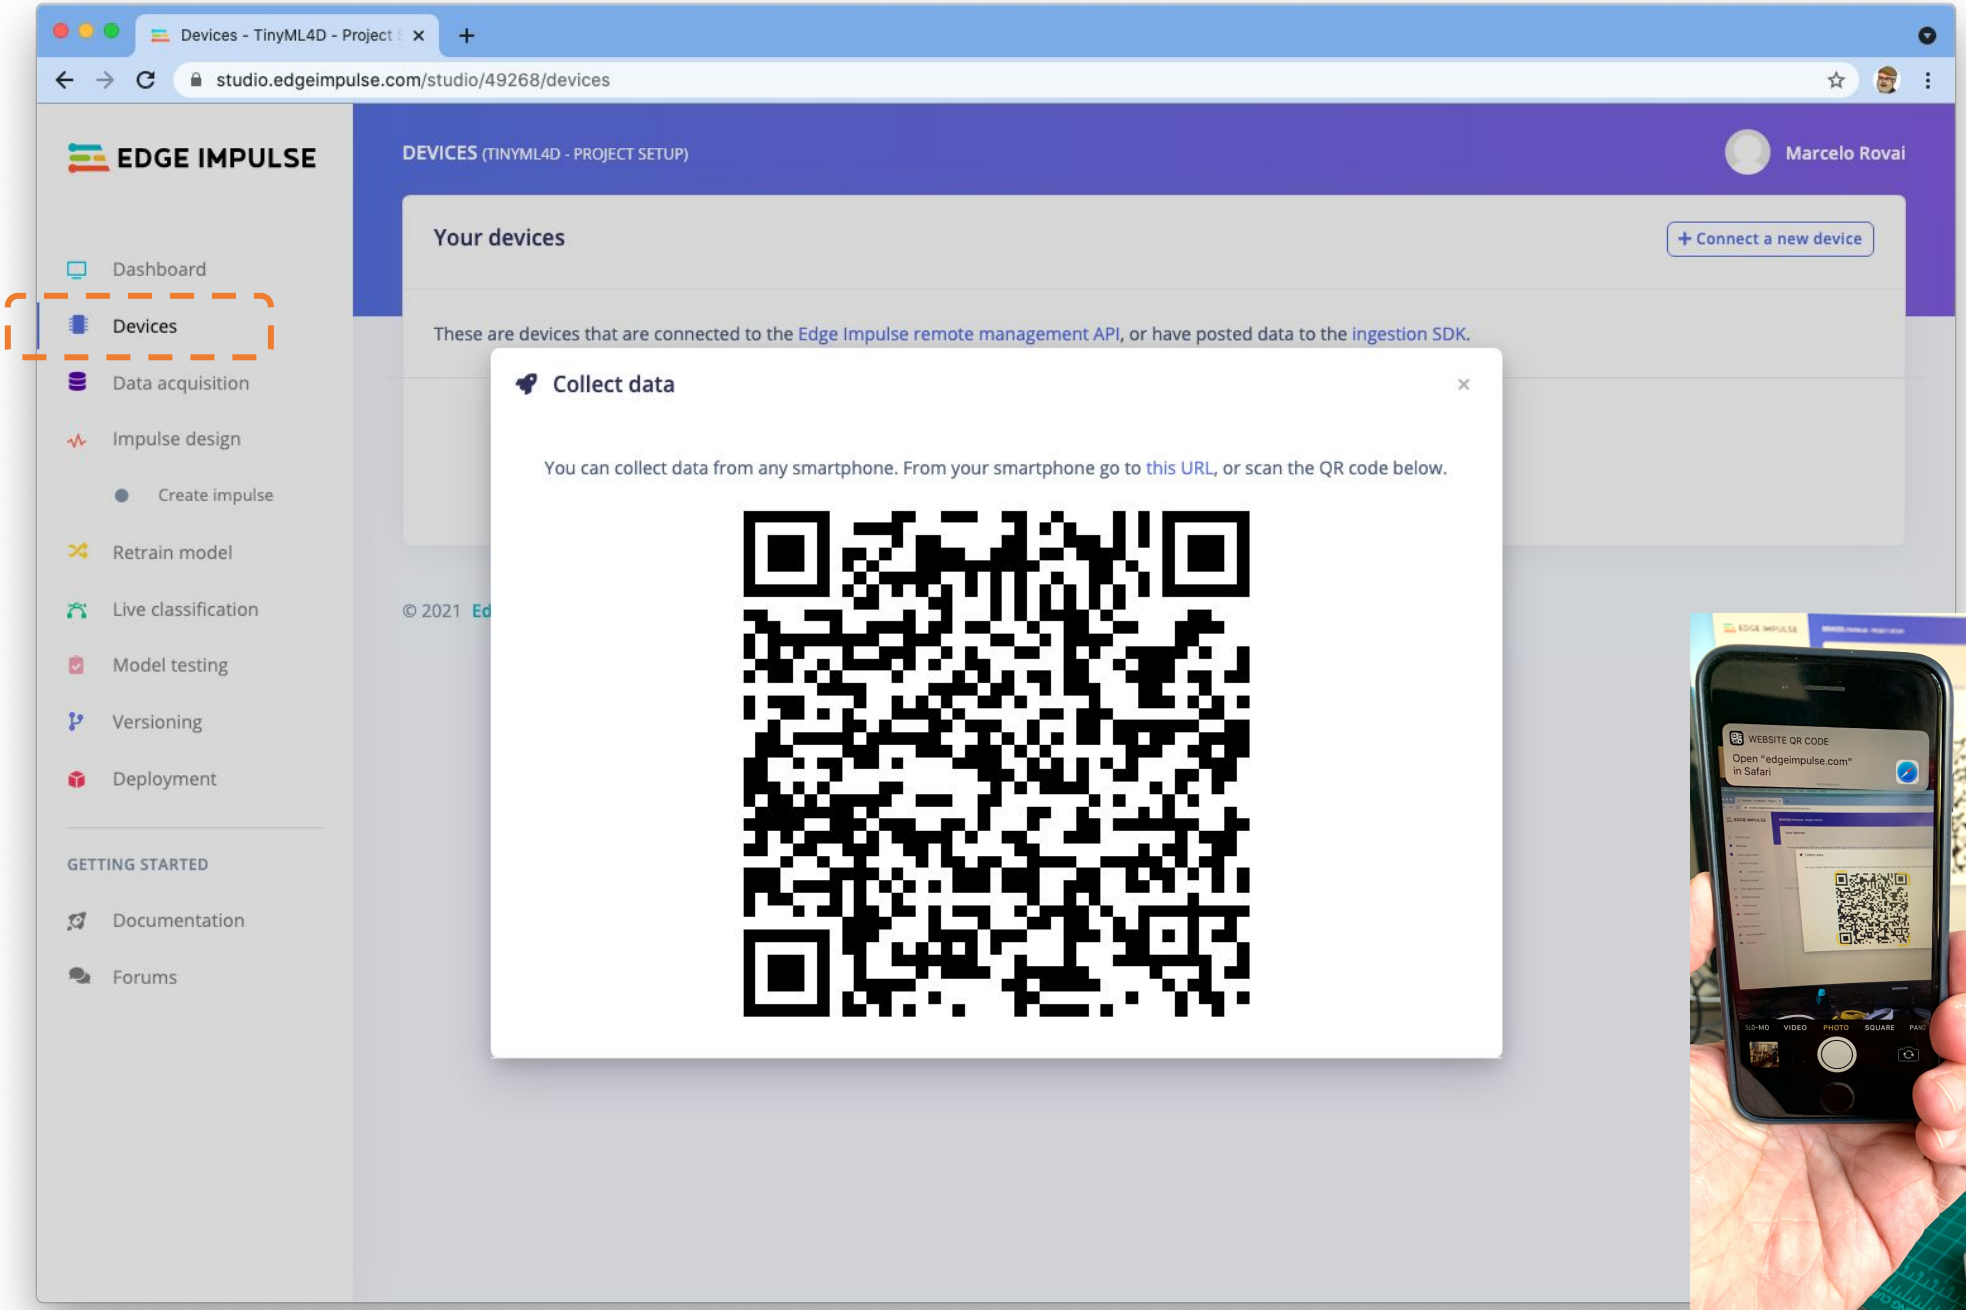
Devices - TinyML4D - Project
studio.edgeimpulse.com/studio/49268/devices
EDGE IMPULSE
DEVICES (TINYML4D - PROJECT SETUP)
Marcelo Rovai
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Your devices
+ Connect a new device
These are devices that are connected to the Edge Impulse remote management API, or have posted data to the ingestion SDK.
Collect data
You can collect data from any smartphone. From your smartphone go to this URL, or scan the QR code below.
© 2021 Ed


WEBSITE QR CODE
Open "edgeimpulse.com" in Safari

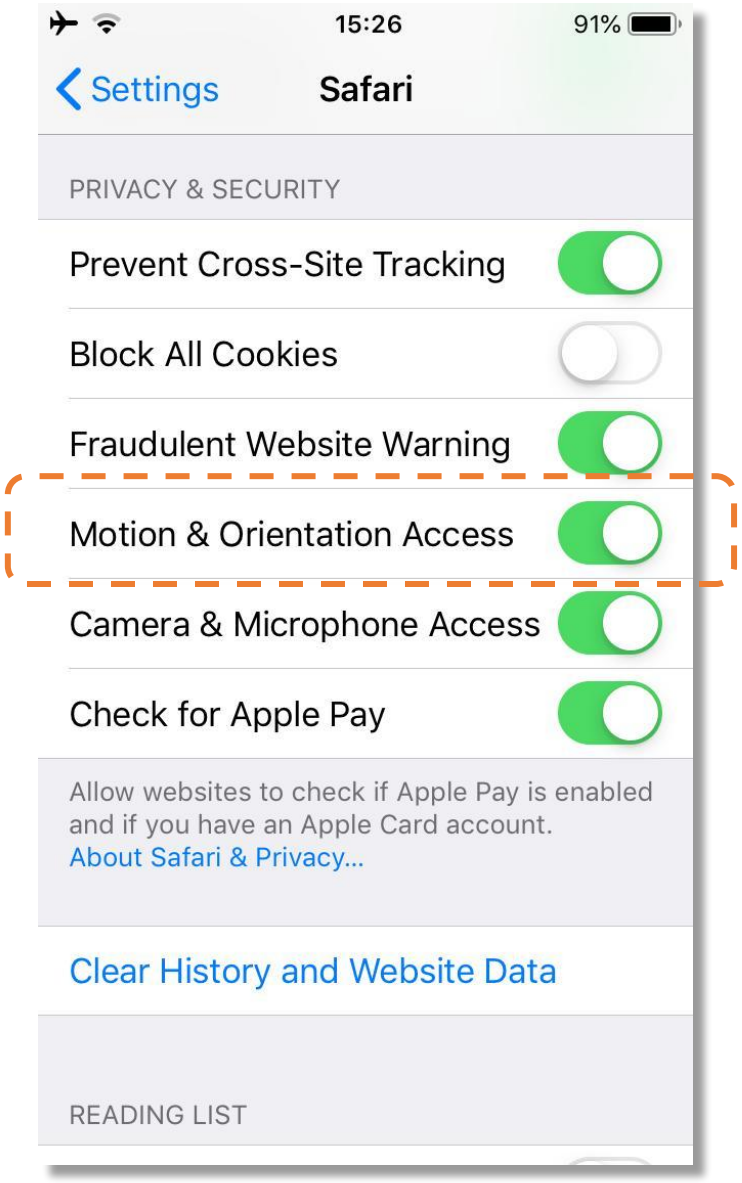

Depending on your mobile OS version and/or your Browser, could be necessary to give access to device IMU sensor

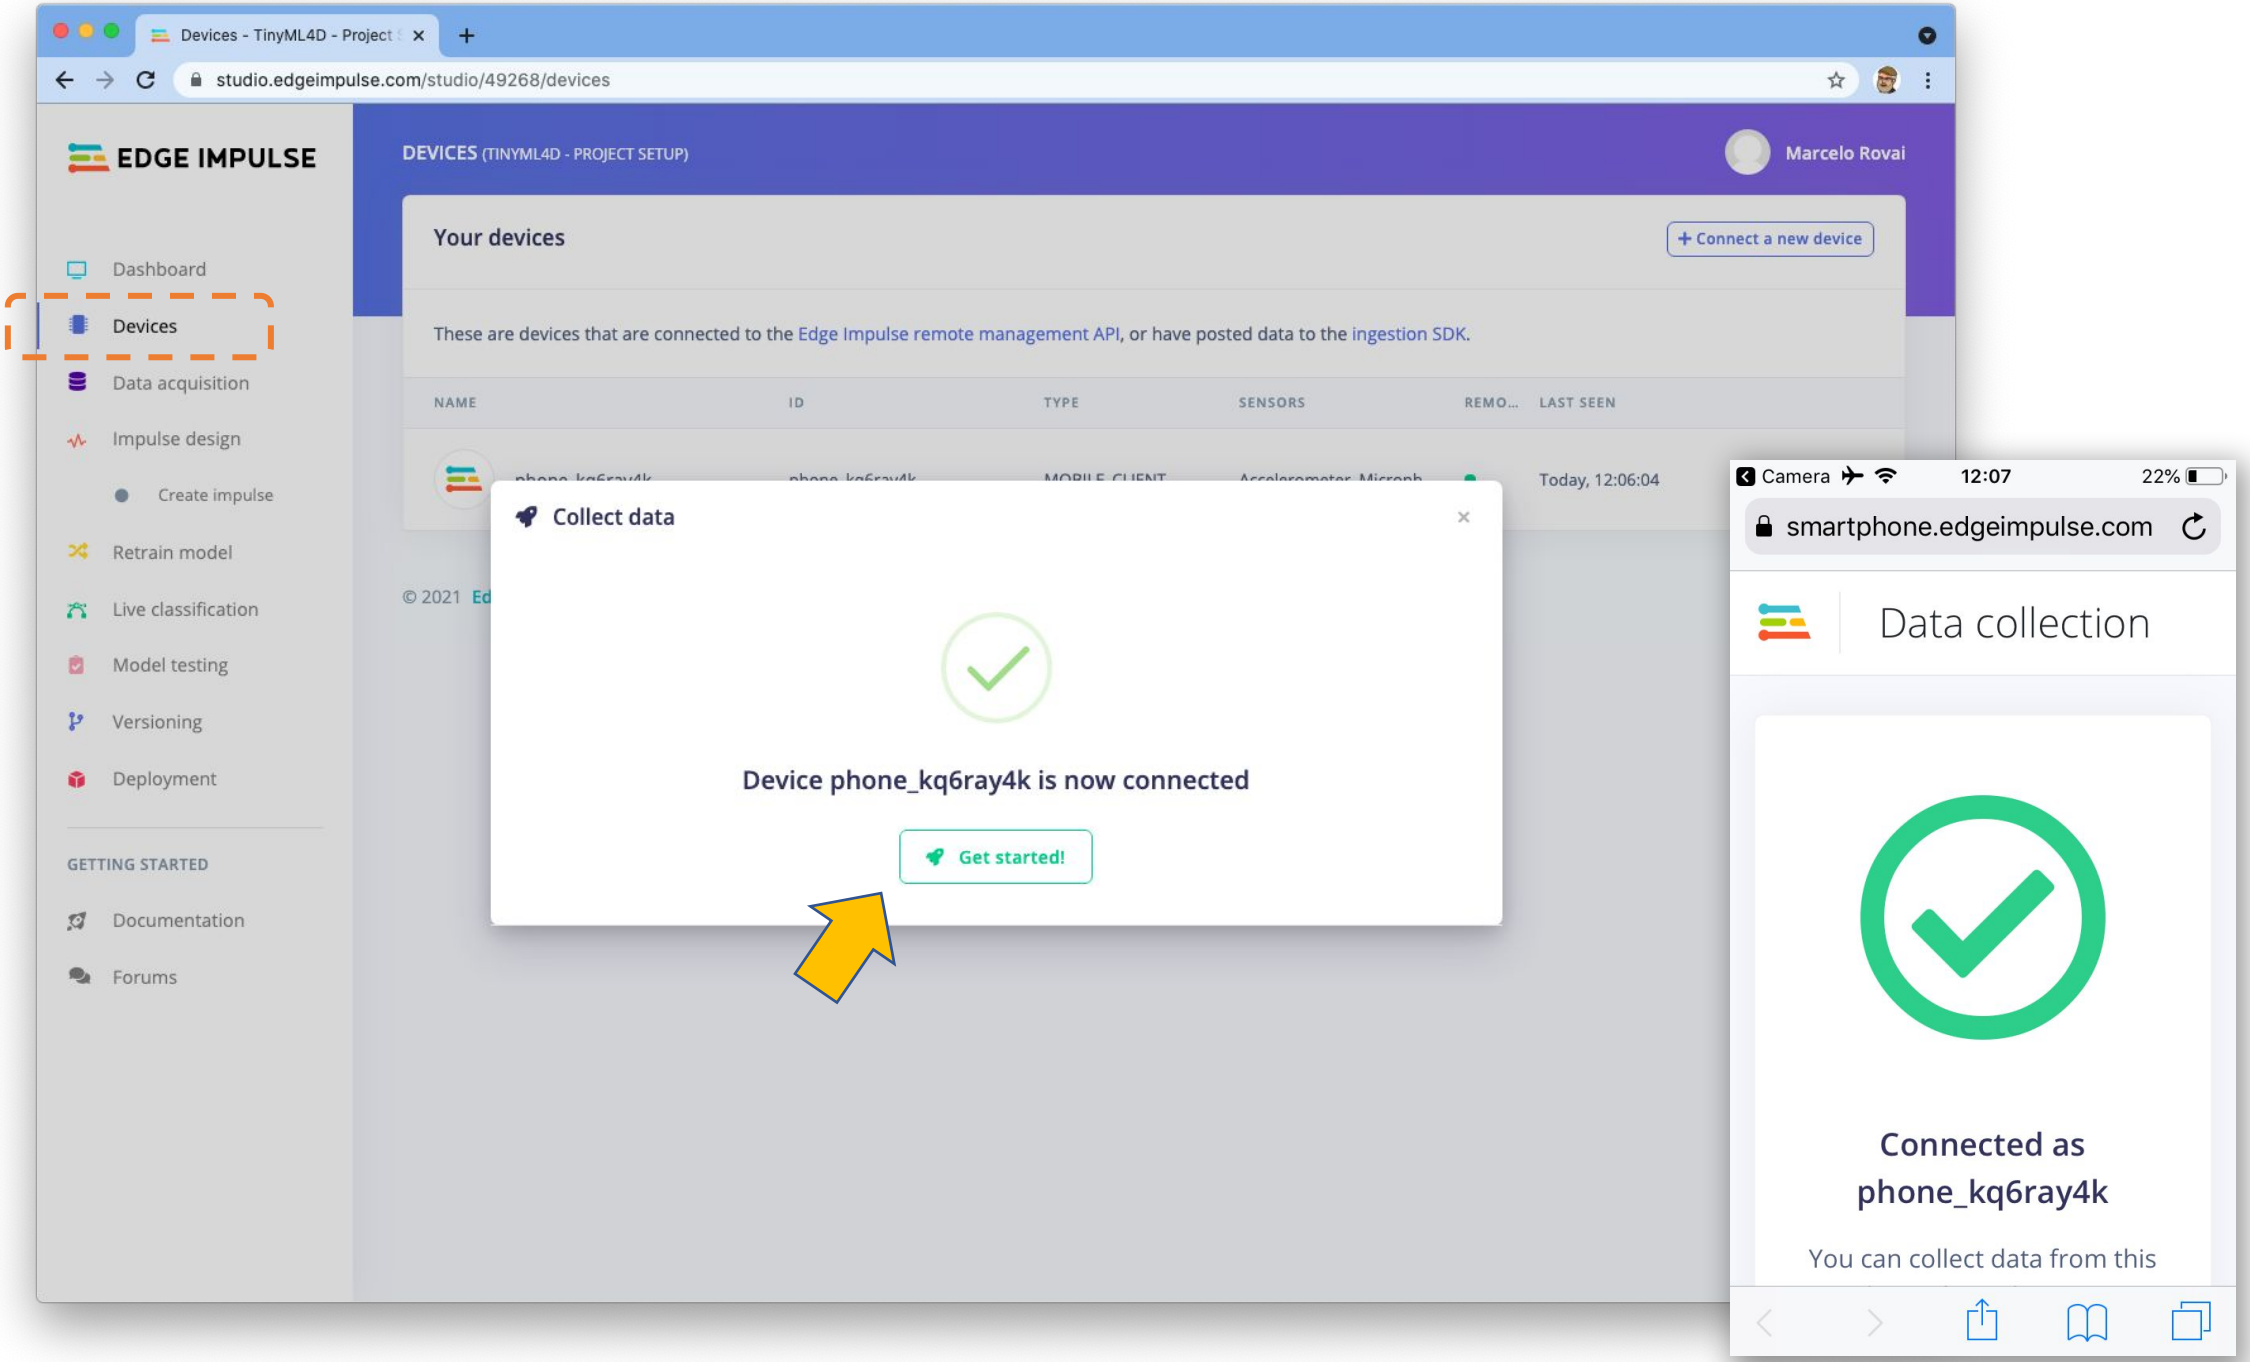
Devices - TinyML4D - Project
studio.edgeimpulse.com/studio/49268/devices
EDGE IMPULSE
DEVICES (TINYML4D - PROJECT SETUP)
Marcelo Rovai
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Your devices
+ Connect a new device
These are devices that are connected to the Edge Impulse remote management API, or have posted data to the ingestion SDK.
NAME
ID
TYPE
SENSORS
REMO...
LAST SEEN
Today, 12:06:04
© 2021
Collect data
Device phone_kq6ray4k is now connected
Get started!
Camera
12:07
22%
smartphone.edgeimpulse.com
Data collection
Connected as
phone_kq6ray4k
You can collect data from this

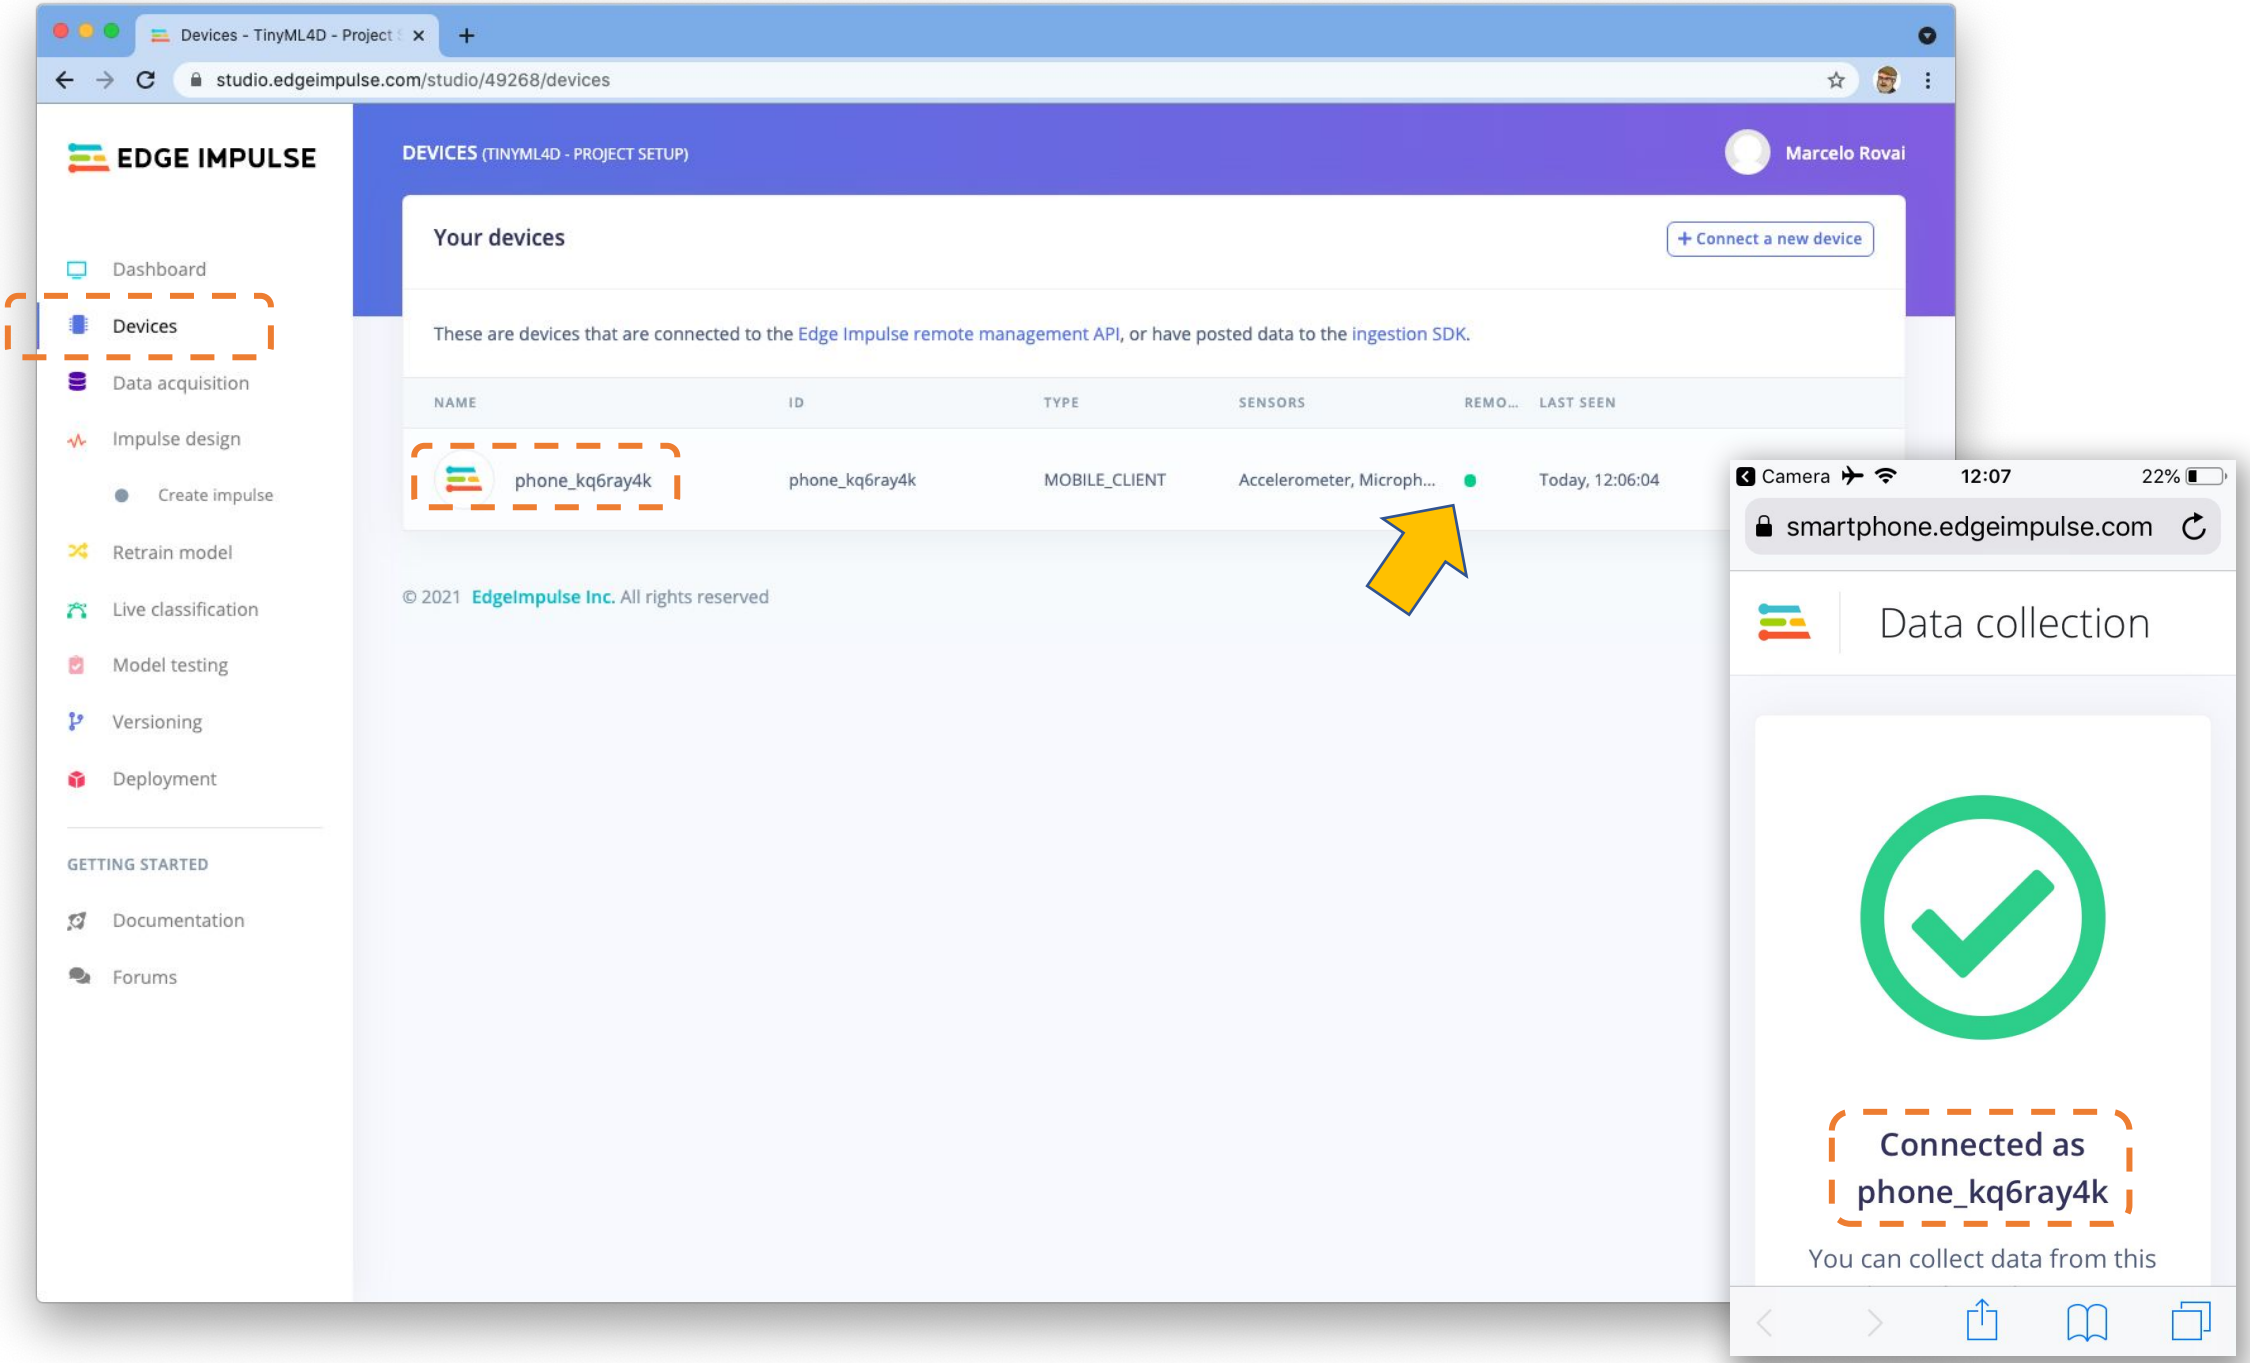
Devices - TinyML4D - Project
studio.edgeimpulse.com/studio/49268/devices
EDGE IMPULSE
DEVICES (TINYML4D - PROJECT SETUP)
Marcelo Rovai
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Your devices
+ Connect a new device
These are devices that are connected to the Edge Impulse remote management API, or have posted data to the ingestion SDK.
NAME
ID
TYPE
SENSORS
REMO...
LAST SEEN
phone_kq6ray4k
phone_kq6ray4k
MOBILE_CLIENT
Accelerometer, Microph...
Today, 12:06:04
© 2021 EdgeImpulse Inc. All rights reserved
Camera
12:07
22%
smartphone.edgeimpulse.com
Data collection
Connected as
phone_kq6ray4k
You can collect data from this

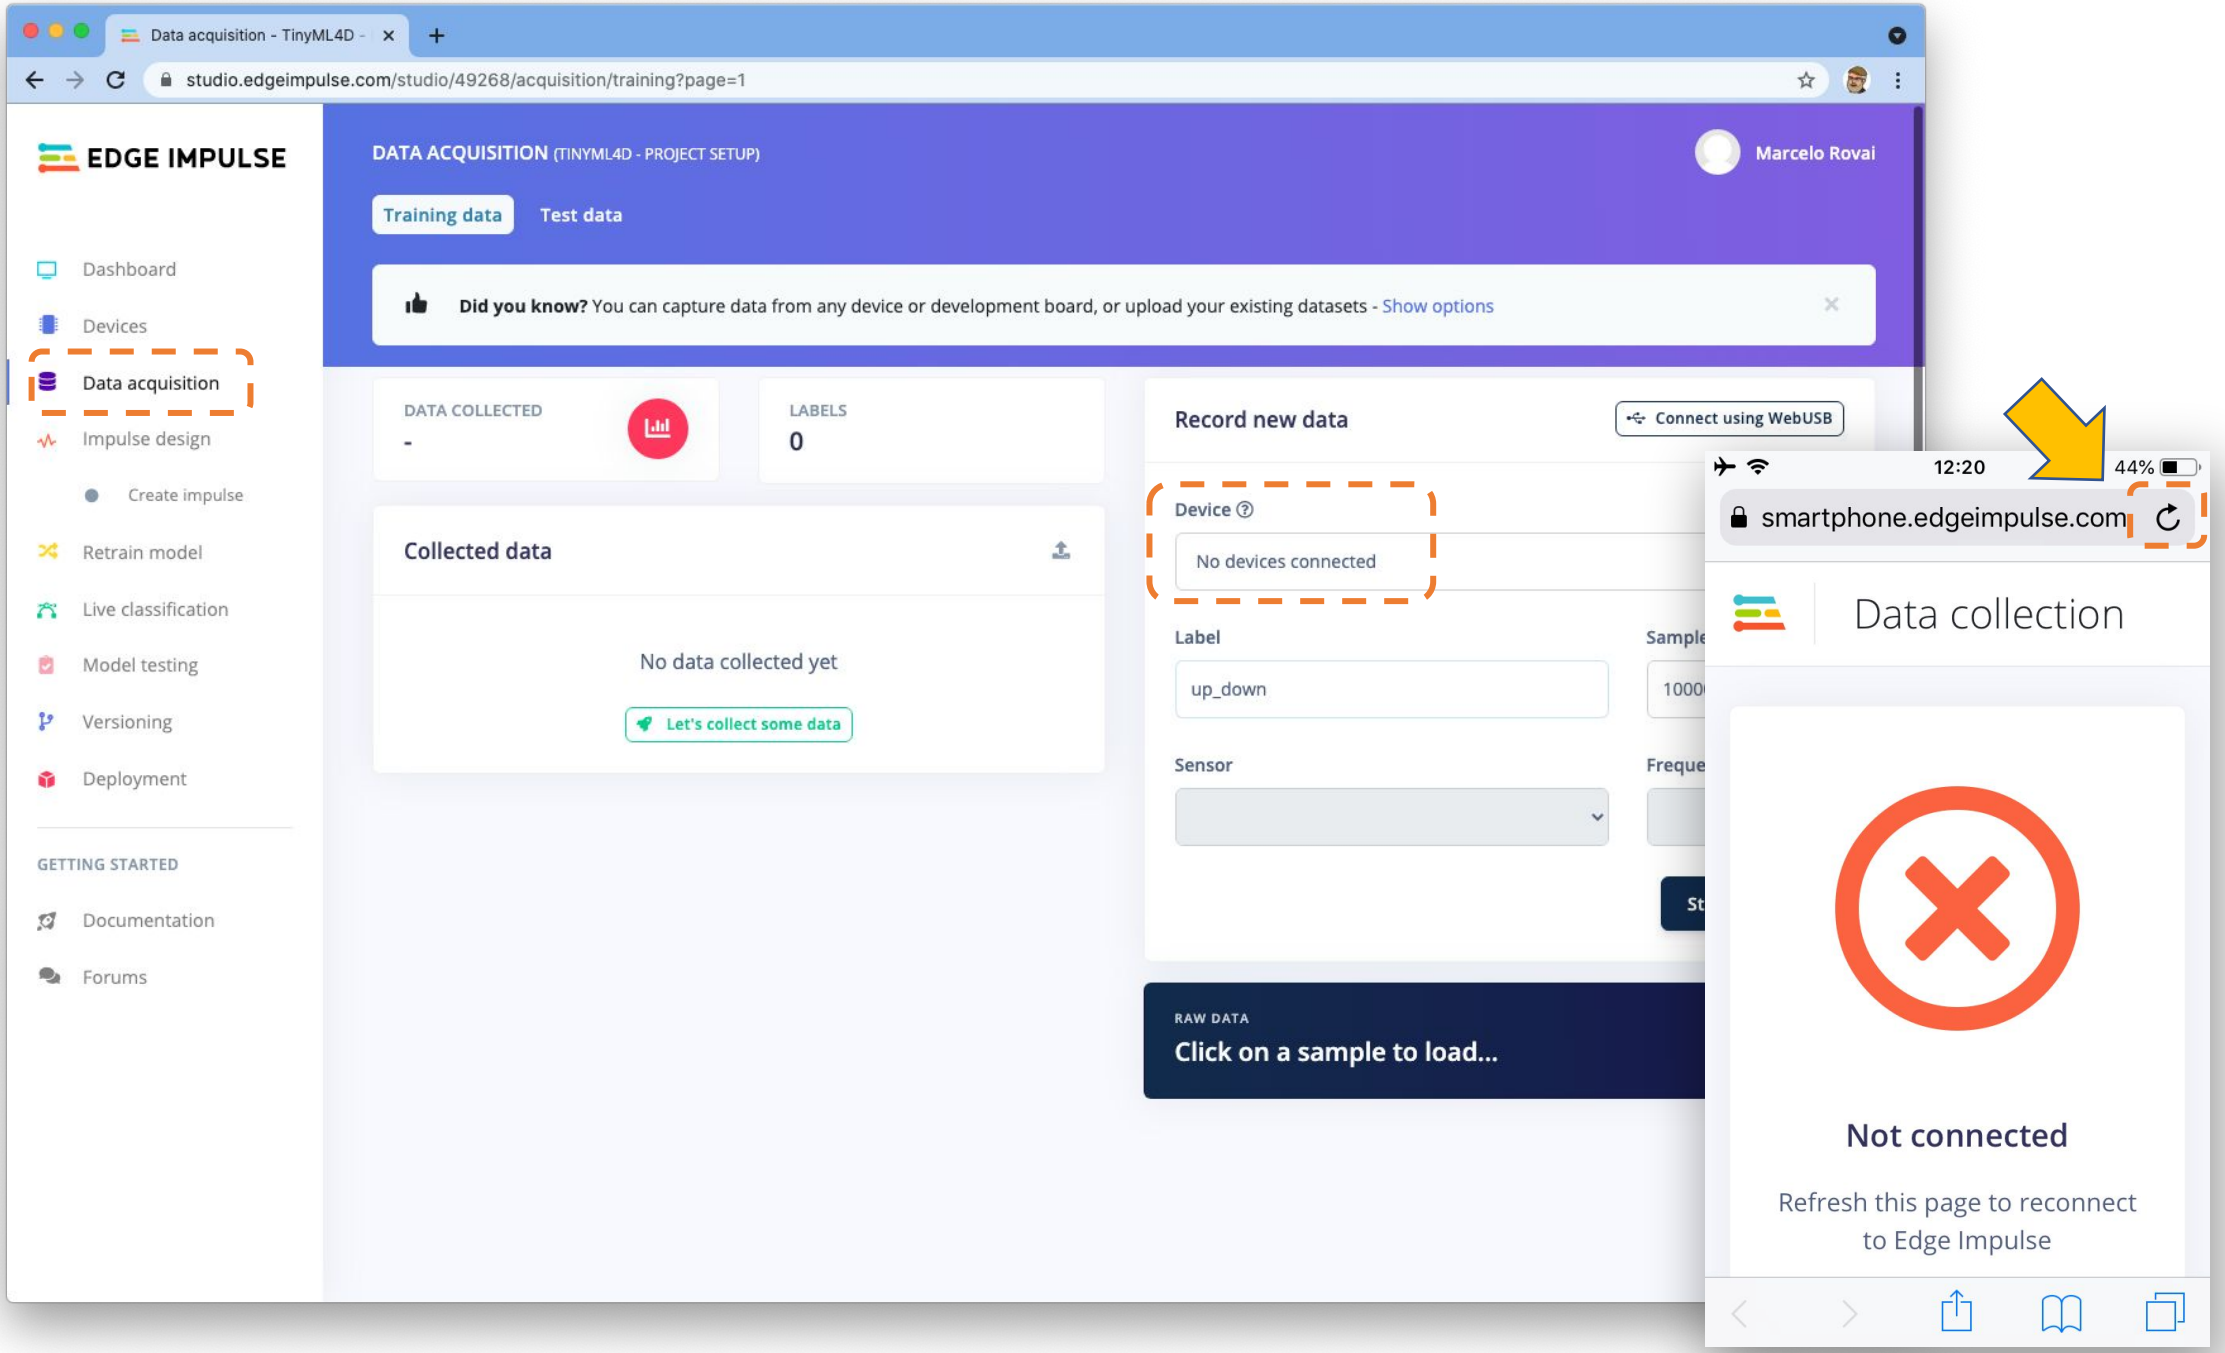
Data acquisition - TinyML4D -
studio.edgeimpulse.com/studio/49268/acquisition/training?page=1
EDGE IMPULSE
DATA ACQUISITION (TINYML4D - PROJECT SETUP)
Marcelo Rovai
Training data
Test data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
DATA COLLECTED
-
LABELS
0
Collected data
No data collected yet
Let's collect some data
Record new data
Connect using WebUSB
Device
No devices connected
Label
up_down
Sensor
RAW DATA
Click on a sample to load...
12:20
44%
smartphone.edgeimpulse.com
Data collection
Not connected
Refresh this page to reconnect to Edge Impulse

Collect Data

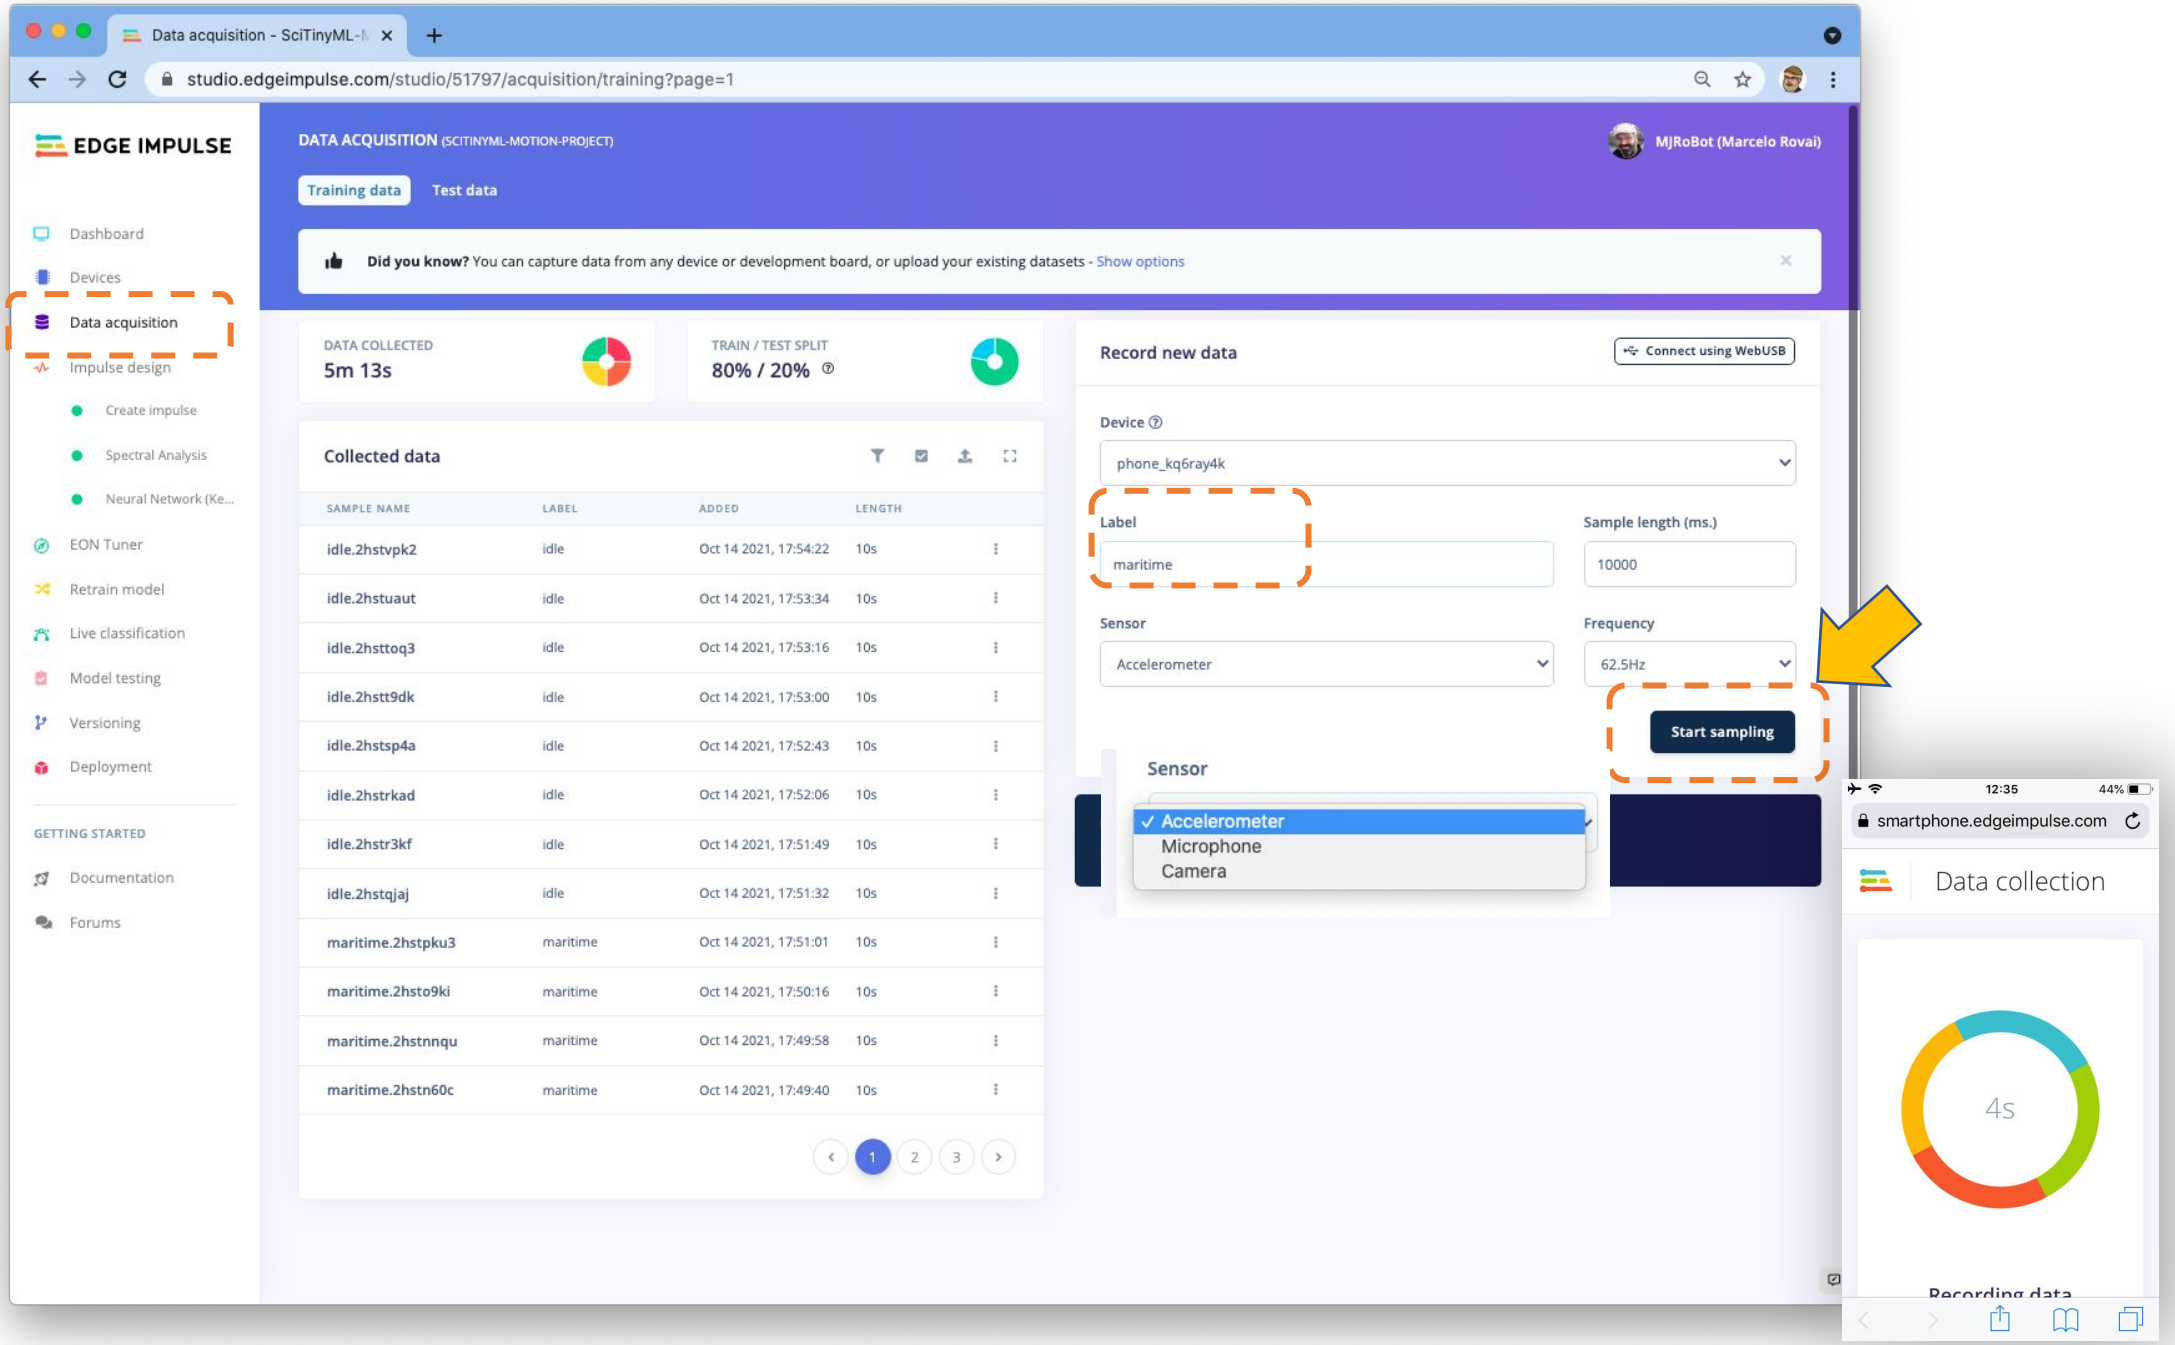

Collect Data

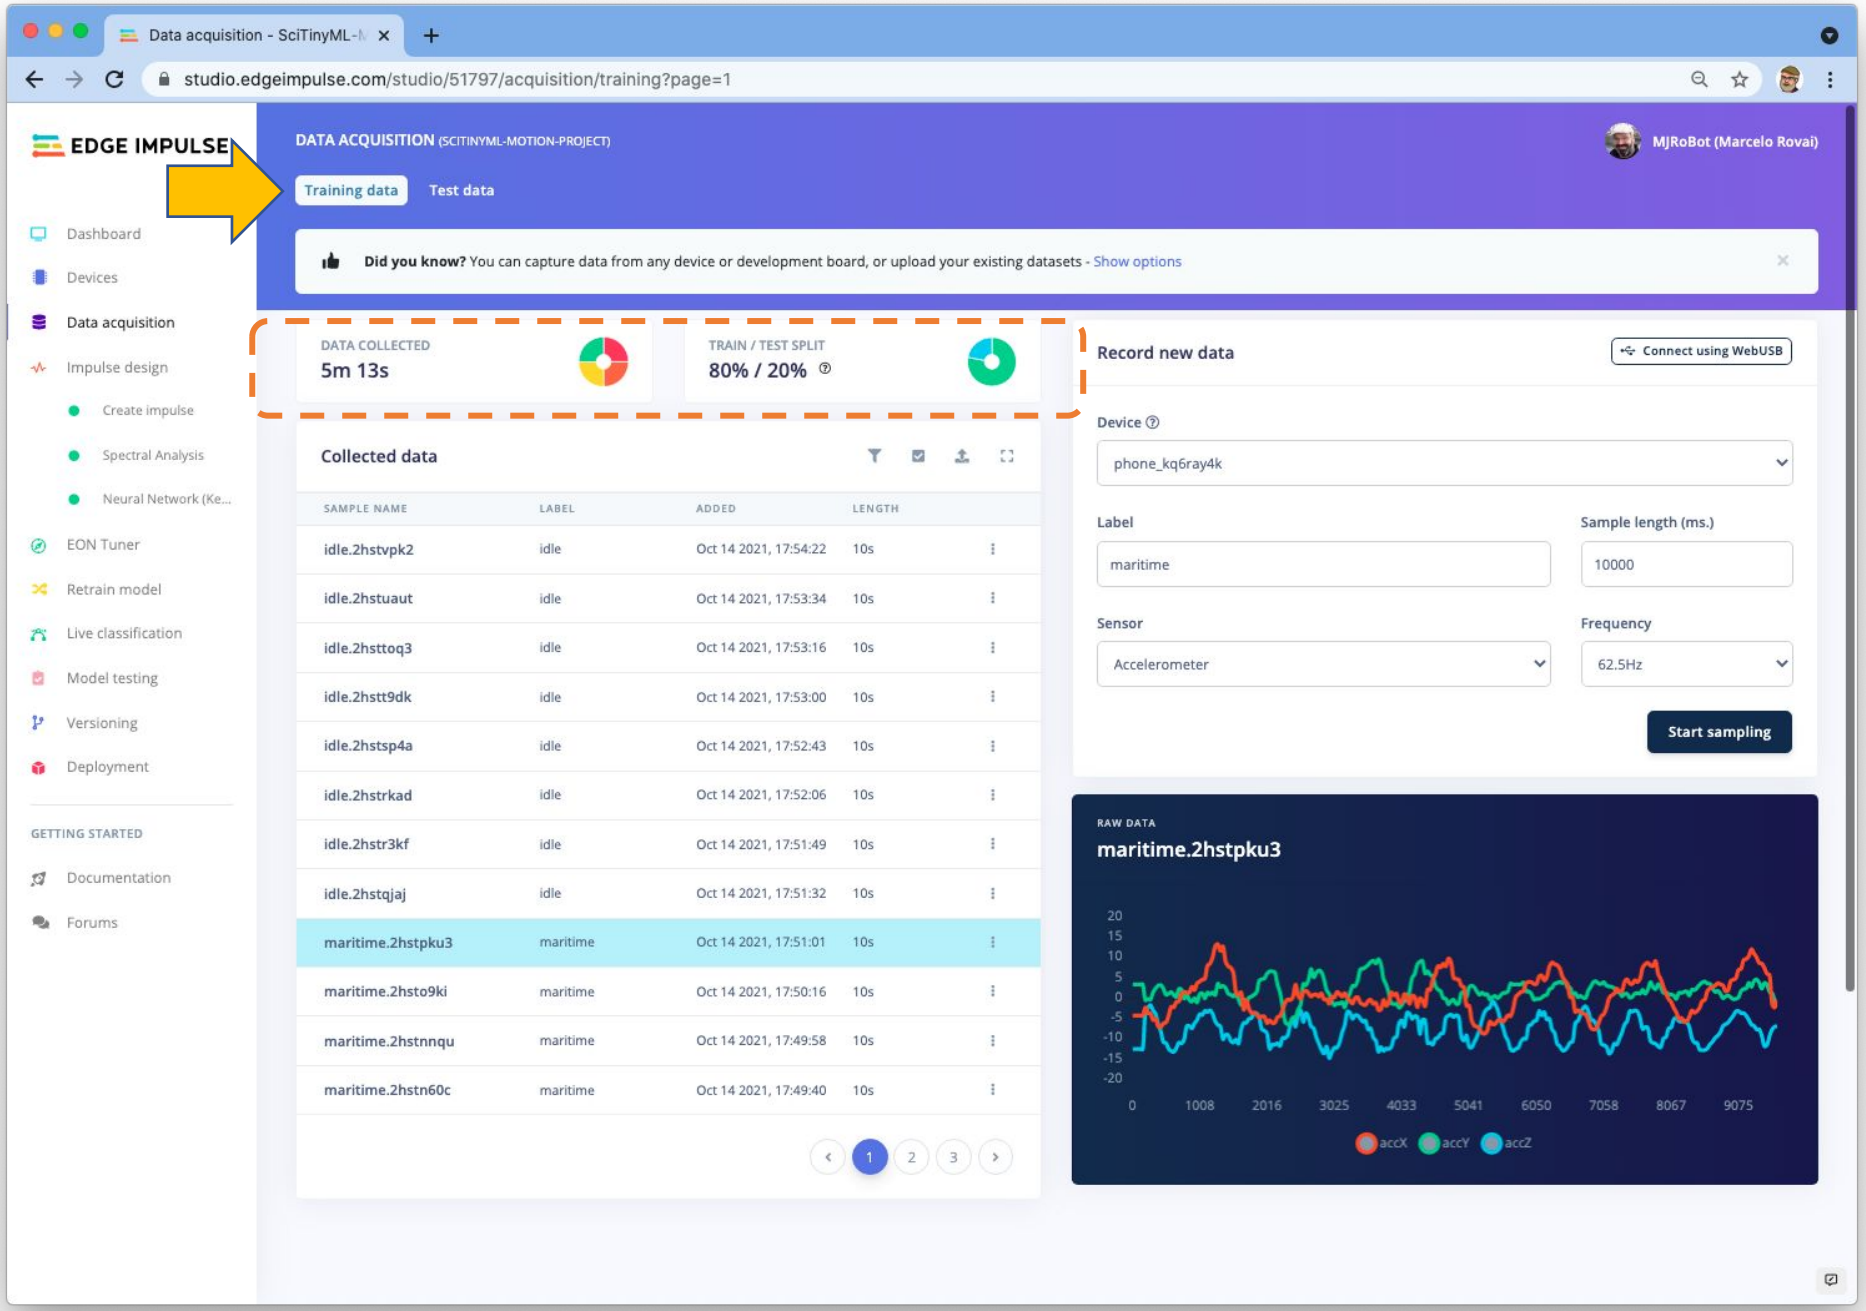

Collect Data

Original Dataset


Original Dataset
Collect Data
Training Set
Test Set

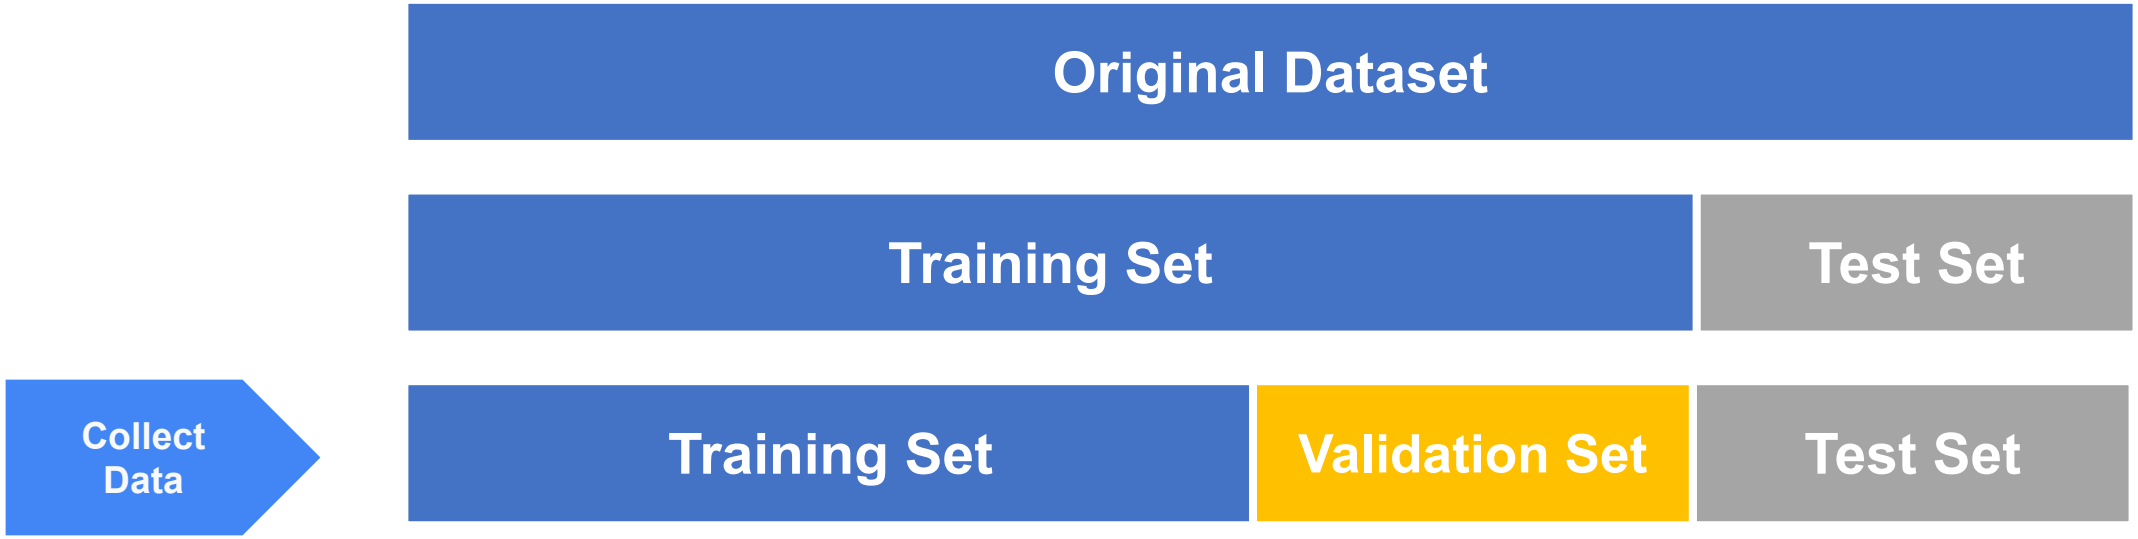
Original Dataset
Training Set
Test Set
Collect Data
Training Set
Validation Set
Test Set

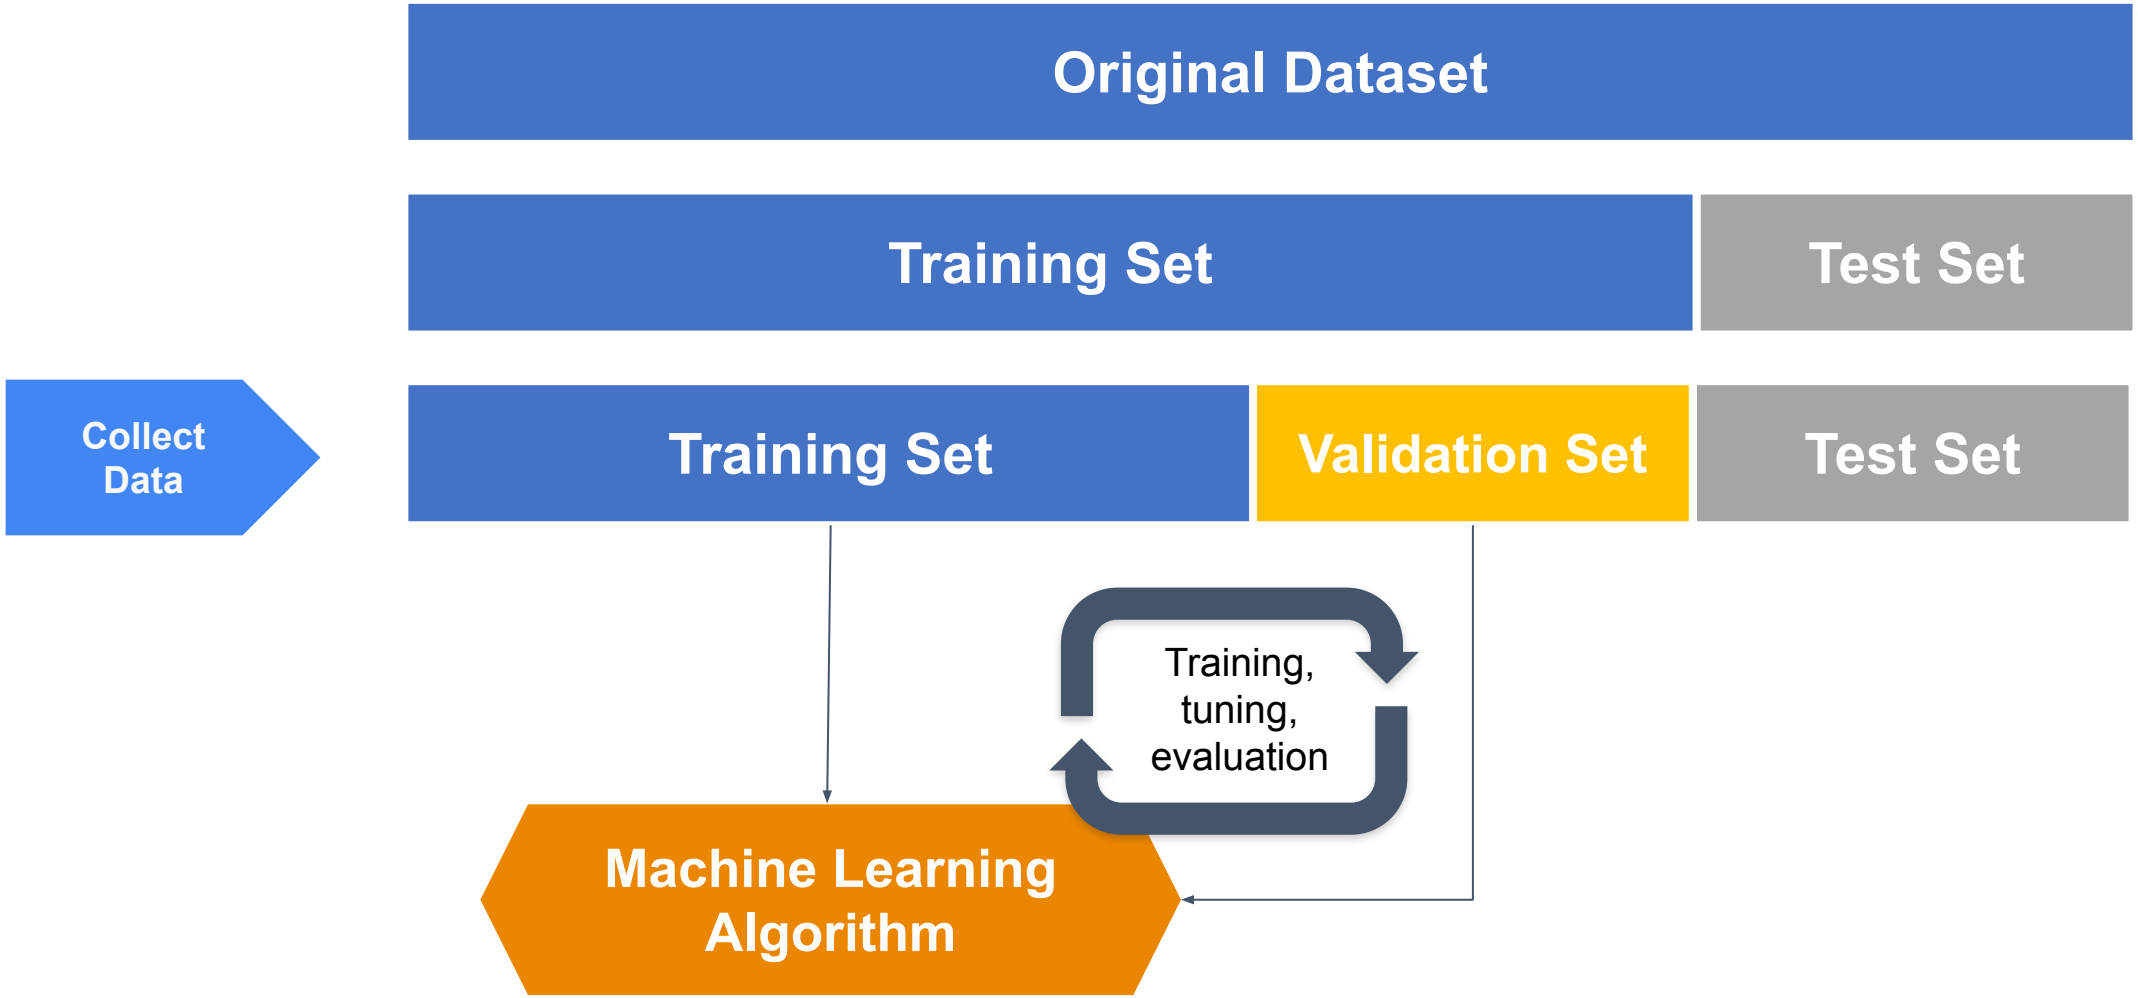
Original Dataset
Training Set
Test Set
Collect Data
Training Set
Validation Set
Test Set
Training, tuning, evaluation
Machine Learning Algorithm

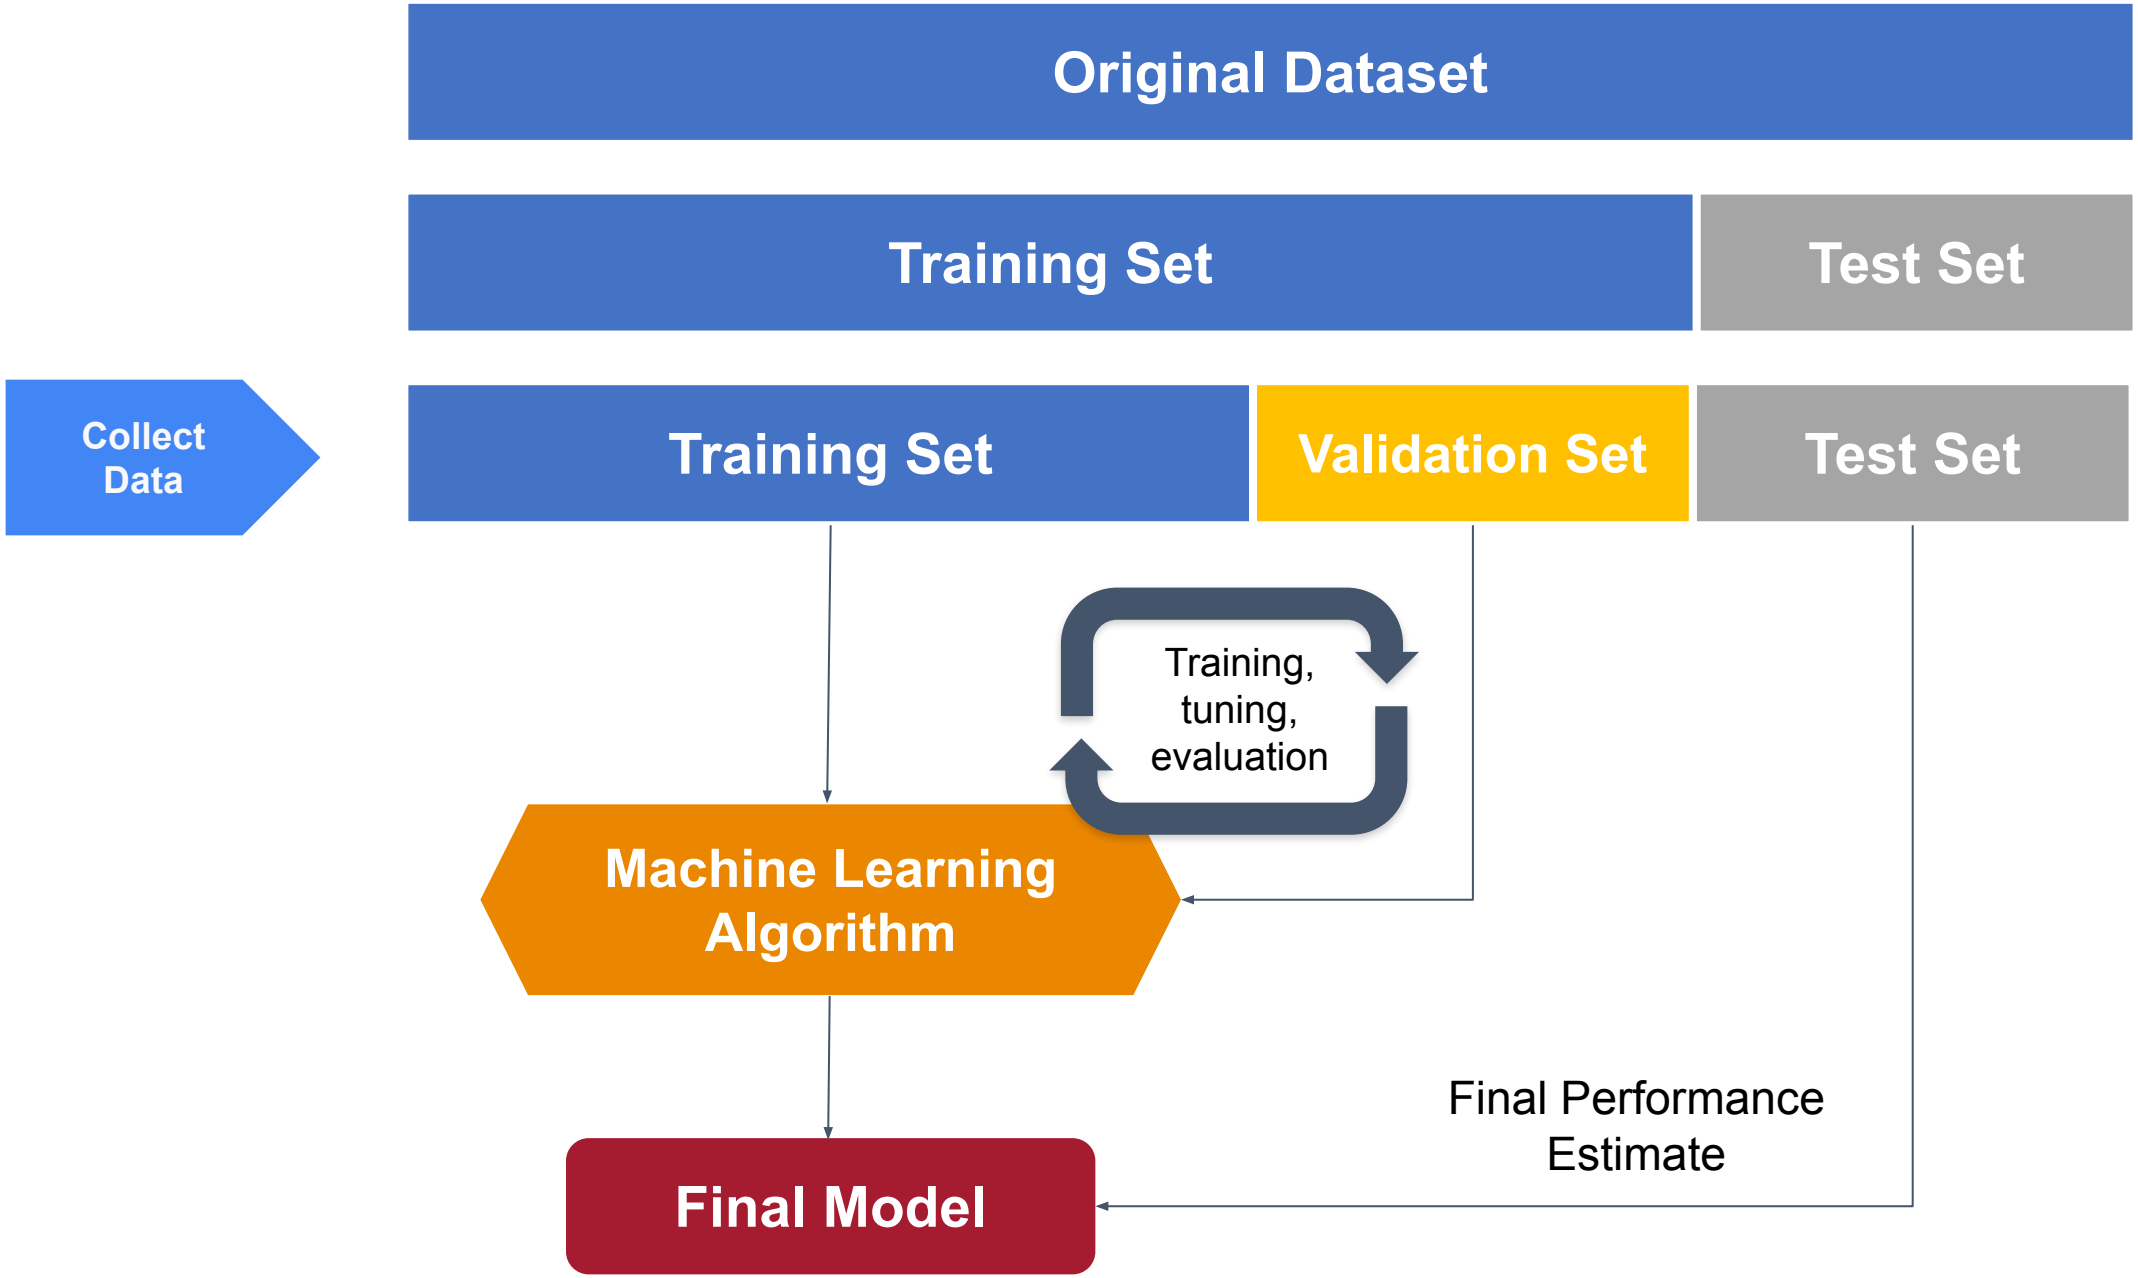
Original Dataset
Training Set
Test Set
Collect Data
Training Set
Validation Set
Test Set
Training, tuning, evaluation
Machine Learning Algorithm
Final Performance Estimate
Final Model

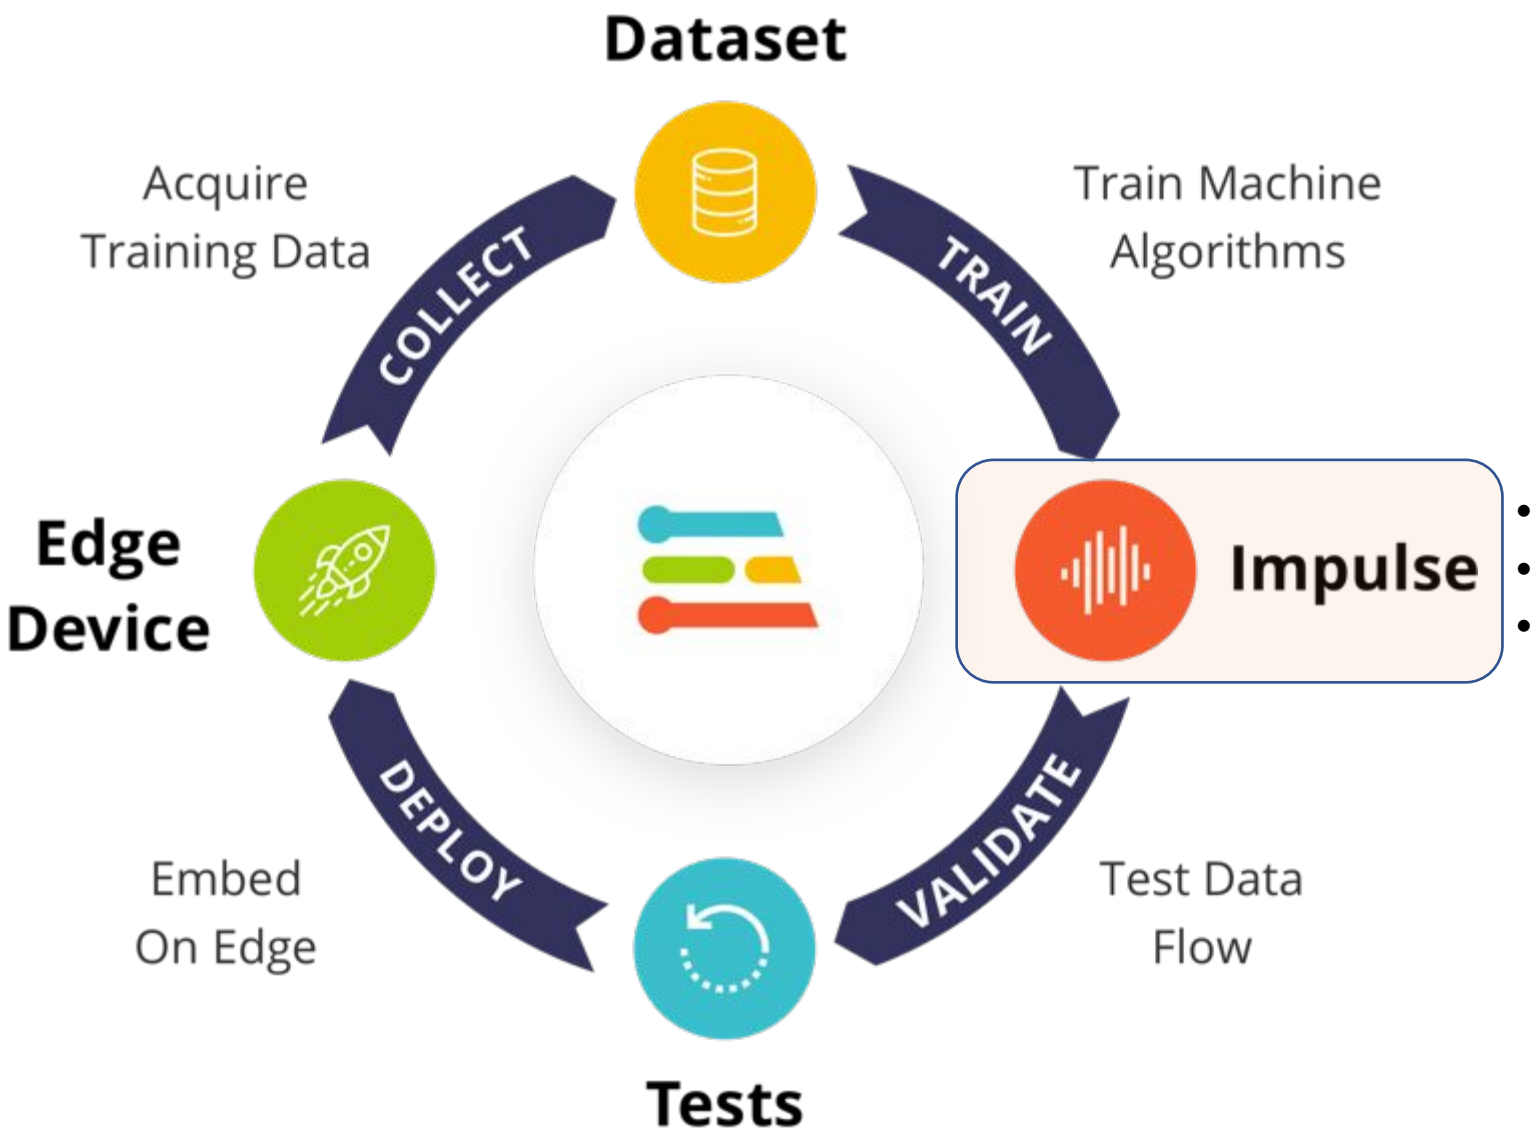

- Pre-Processing Data
- Design a Model
- Train a Model

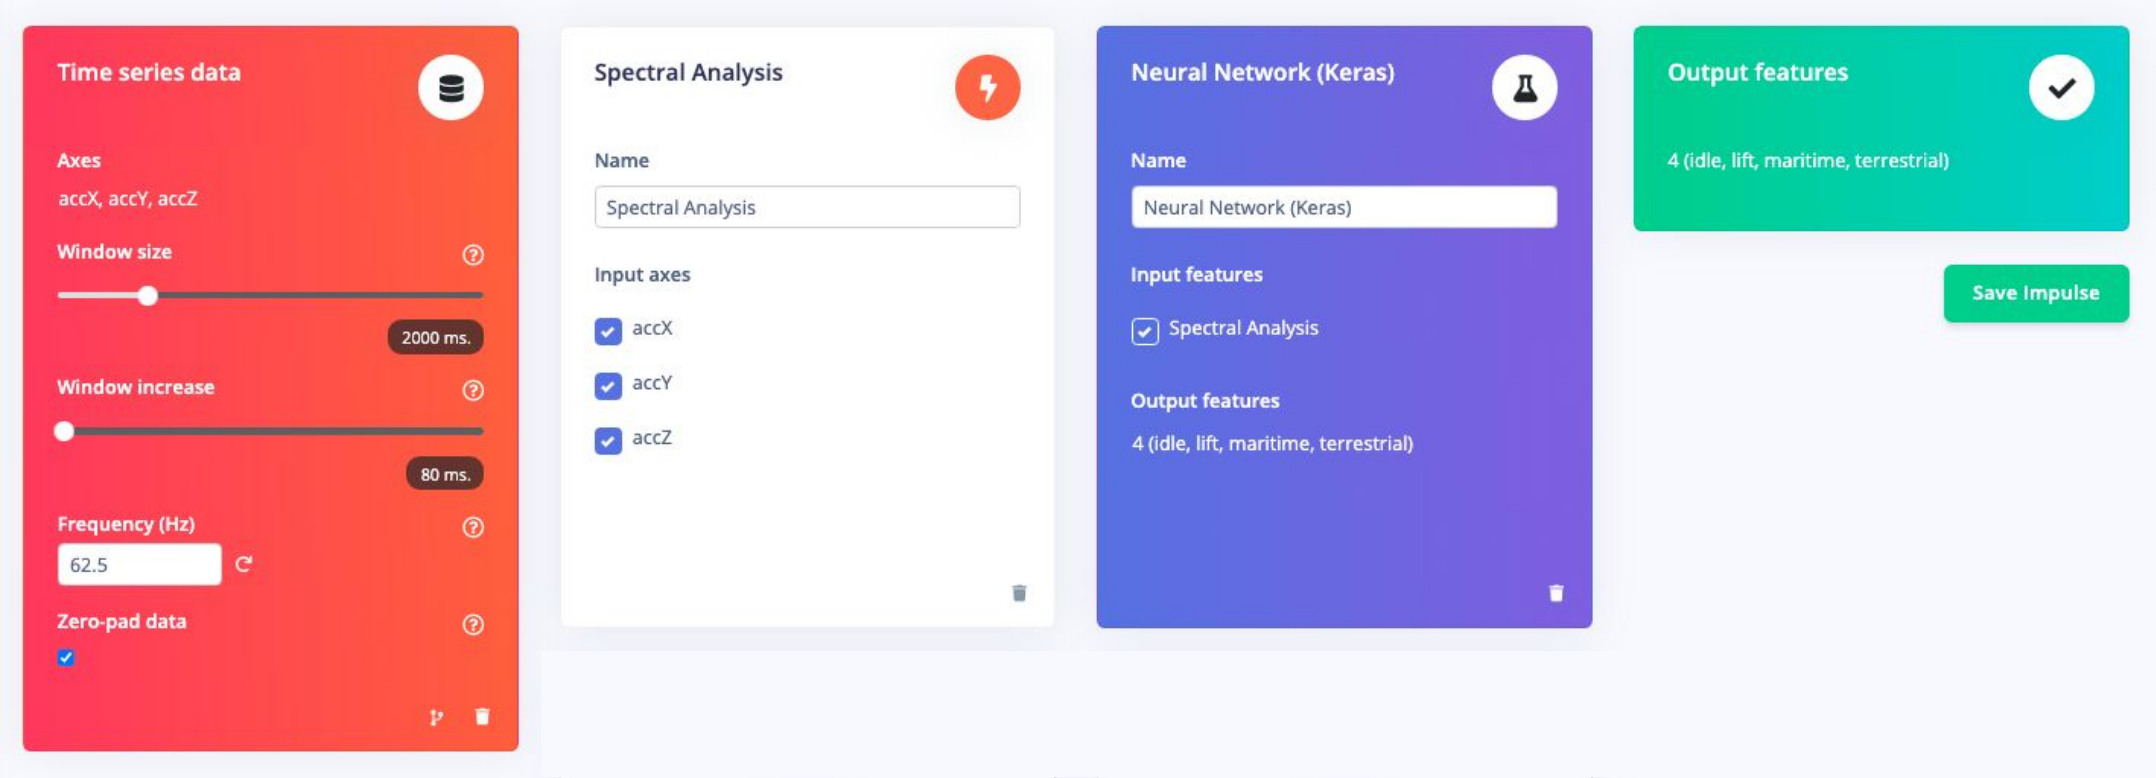

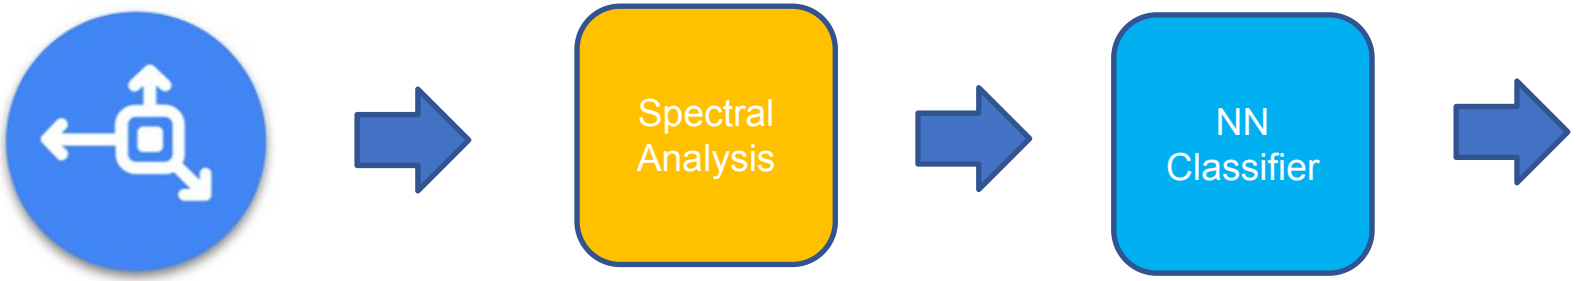

Classes

- Lift
- Terrestrial
- Maritime
- Idle

Preprocess Data

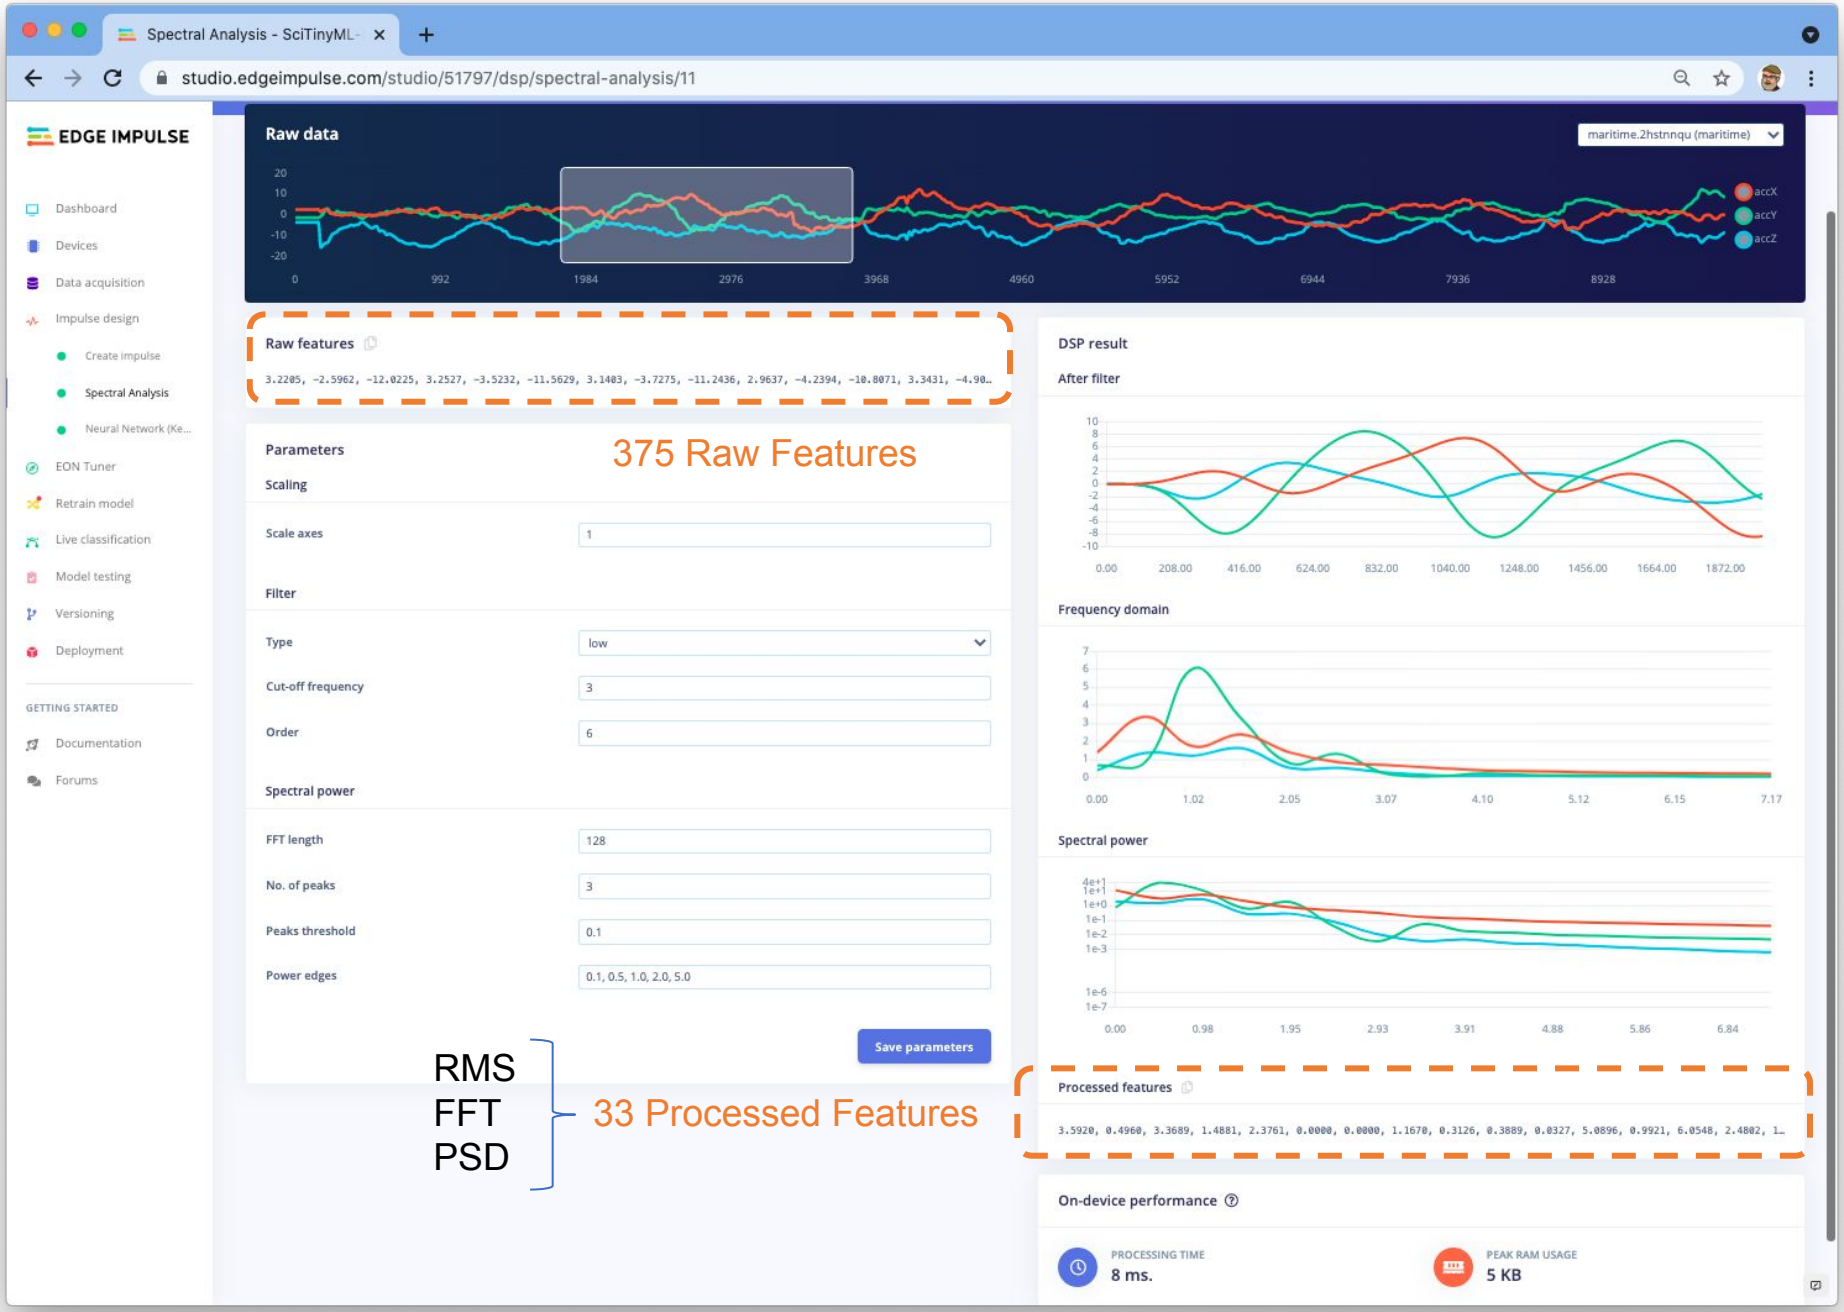

Preprocess Data

Spectral Analysis - SciTinyML

studio.edgeimpulse.com/studio/51797/dsp/spectral-analysis/11/generate-features

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Spectral Analysis
  - Neural Network (Ke...
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

SPECTRAL ANALYSIS (SCITINYML-MOTION-PROJECT)

#1 EON Tuner Primary

Parameters | Generate features

MJRoBot (Marcelo Rovai)

Training set

Data in training set	5m 22s
Classes	4 (idle, lift, maritime, terrestrial)
Window length	2000 ms.
Window increase	80 ms.
Training windows	3,230

Generate features

- ✓ accX RMS
- accX Peak 1 Freq
- accX Peak 1 Height
- accX Peak 2 Freq
- accX Peak 2 Height
- accX Peak 3 Freq
- accX Peak 3 Height
- accX Spectral Power 0.1 - 0.5
- accX Spectral Power 0.5 - 1.0
- accX Spectral Power 1.0 - 2.0
- accX Spectral Power 2.0 - 5.0
- accY RMS
- accY Peak 1 Freq
- accY Peak 1 Height
- accY Peak 2 Freq
- accY Peak 2 Height
- accY Peak 3 Freq
- accY Peak 3 Height
- accY Spectral Power 0.1 - 0.5
- accY Spectral Power 0.5 - 1.0
- accY Spectral Power 1.0 - 2.0
- accY Spectral Power 2.0 - 5.0
- accZ RMS
- accZ Peak 1 Freq
- accZ Peak 1 Height
- accZ Peak 2 Freq
- accZ Peak 2 Height
- accZ Peak 3 Freq
- accZ Peak 3 Height
- accZ Spectral Power 0.1 - 0.5
- accZ Spectral Power 0.5 - 1.0
- accZ Spectral Power 1.0 - 2.0
- accZ Spectral Power 2.0 - 5.0

Feature explorer (3,132 samples)

X Axis: accX RMS | Y Axis: accY RMS | Z Axis: accZ RMS

- idle
- lift
- maritime
- terrestrial

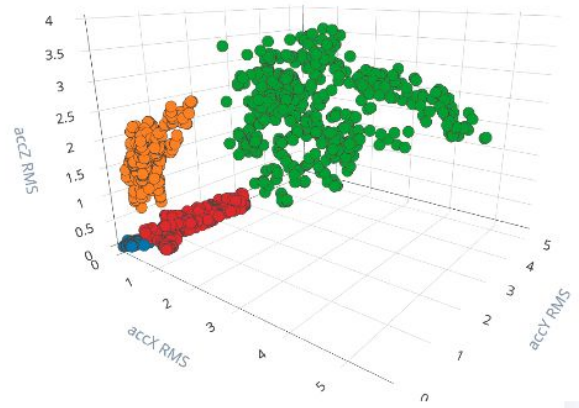

On-device performance

PROCESSING TIME
8 ms.

PEAK RAM USAGE
5 KB

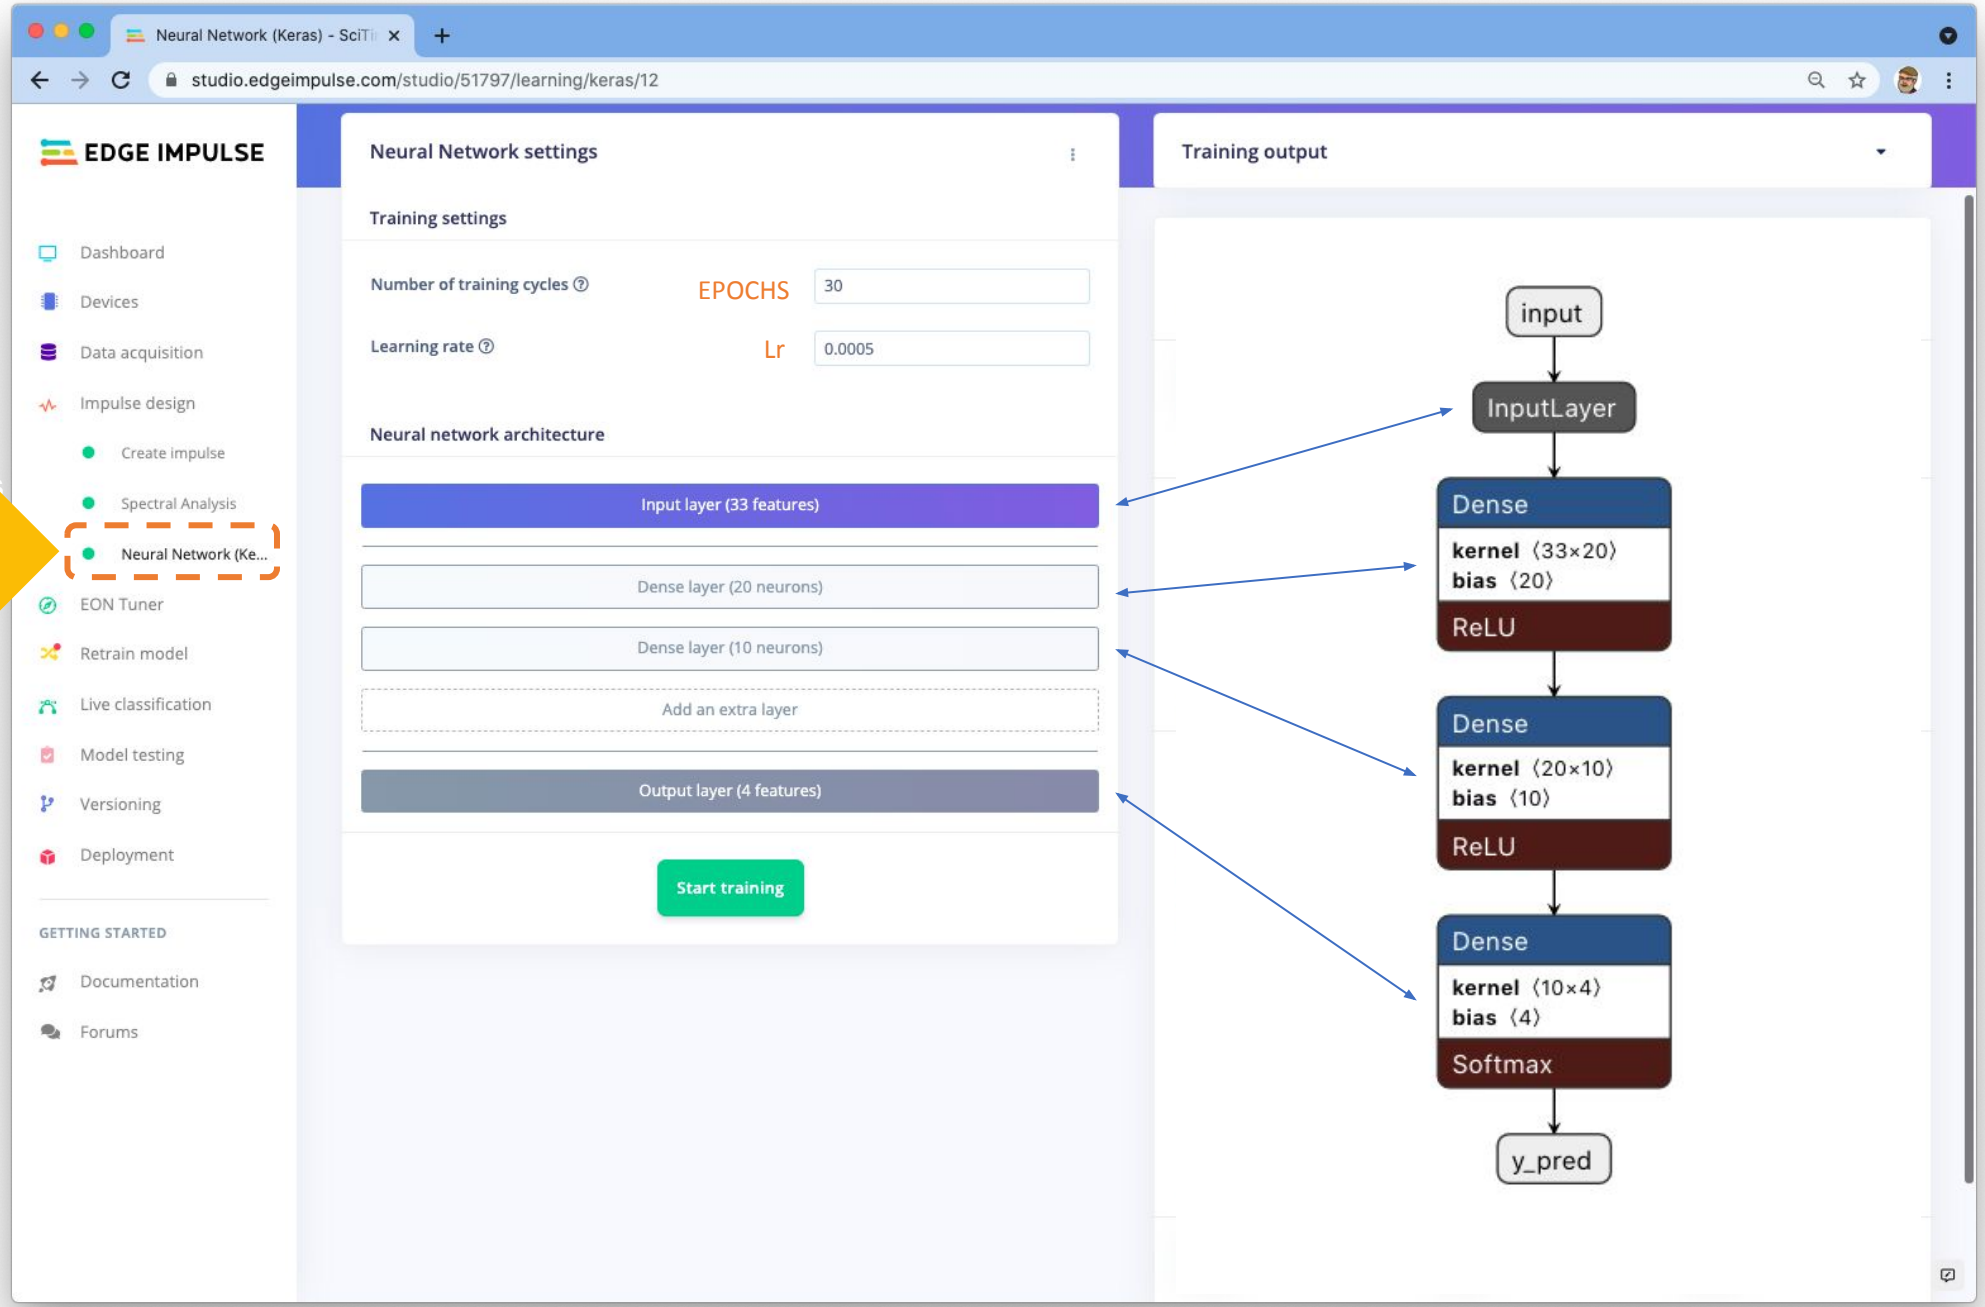

Design a Model

Train a Model

Neural Network (Keras) - SciTi

studio.edgeimpulse.com/studio/51797/learning/keras/12

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Spectral Analysis
  - Neural Network (Ke...
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

Neural Network settings

Training settings

Number of training cycles — EPOCHS: 30

Learning rate — Lr: 0.0005

Neural network architecture

Input layer (33 features)

Dense layer (20 neurons)

Dense layer (10 neurons)

Add an extra layer

Output layer (4 features)

Start training

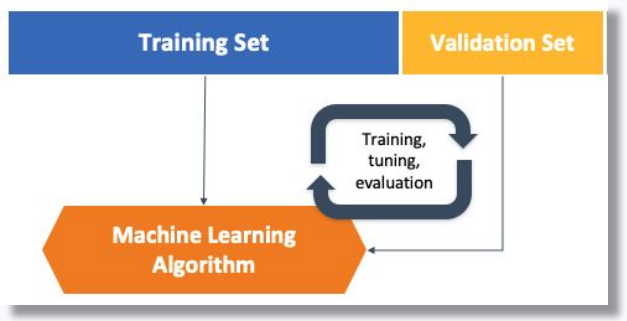

Training output

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY 99.7%

LOSS 0.01

Confusion matrix (validation set)

	IDLE	LIFT	MARITIME	TERRESTRIAL
IDLE	100%	0%	0%	0%
LIFT	0%	100%	0%	0%
MARITIME	0%	0.6%	99.4%	0%
TERRESTRIAL	0.6%	0%	0%	99.4%
F1 SCORE	1.00	1.00	1.00	1.00

Feature explorer (full training set)

accX RMS | accY RMS | accZ RMS

- idle - correct
- lift - correct
- maritime - correct
- terrestrial - correct
- maritime - incorrect
- terrestrial - incorrect

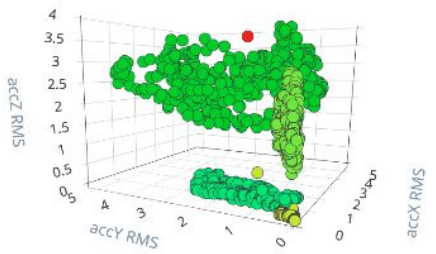

On-device performance

INFERENCING TIME 1 ms.

PEAK RAM USAGE 1.7K

FLASH USAGE 19.0K

Evaluate
Optimize

Neural Network (Keras) - SciTi

studio.edgeimpulse.com/studio/51797/learning/keras/12

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Spectral Analysis
  - Neural Network (Ke...
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

Neural Network settings

Training settings

Number of training cycles: 30

Learning rate: 0.0005

Neural network architecture

Input layer (33 features)

Dense layer (20 neurons)

Dense layer (10 neurons)

Add an extra layer

Output layer (4 features)

Start training

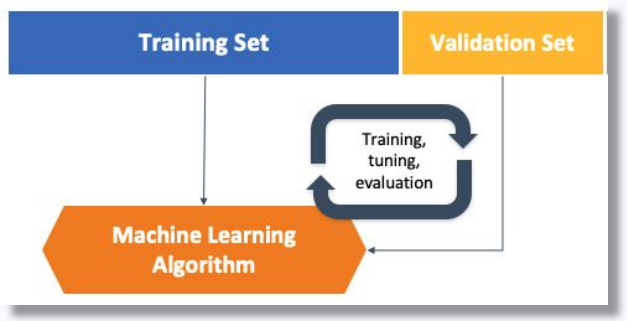

Training output

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY 99.7%

LOSS 0.01

Confusion matrix (validation set)

	IDLE	LIFT	MARITIME	TERRESTRIAL
IDLE	100%	0%	0%	0%
LIFT	0%	100%	0%	0%
MARITIME	0%	0.6%	99.4%	0%
TERRESTRIAL	0.6%	0%	0%	99.4%
F1 SCORE	1.00	1.00	1.00	1.00

Feature explorer (full training set)

accX RMS | accY RMS | accZ RMS

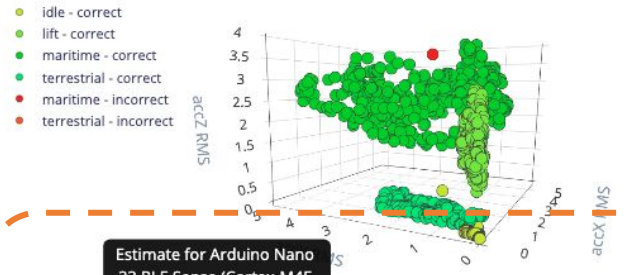

Estimate for Arduino Nano 33 BLE Sense (Cortex-M4F 64MHz), compiled with Edge Impulse EON™ compiler

On-device performance

INFERENCING TIME 1 ms.

PEAK RAM USAGE 1.7K

FLASH USAGE 19.0K

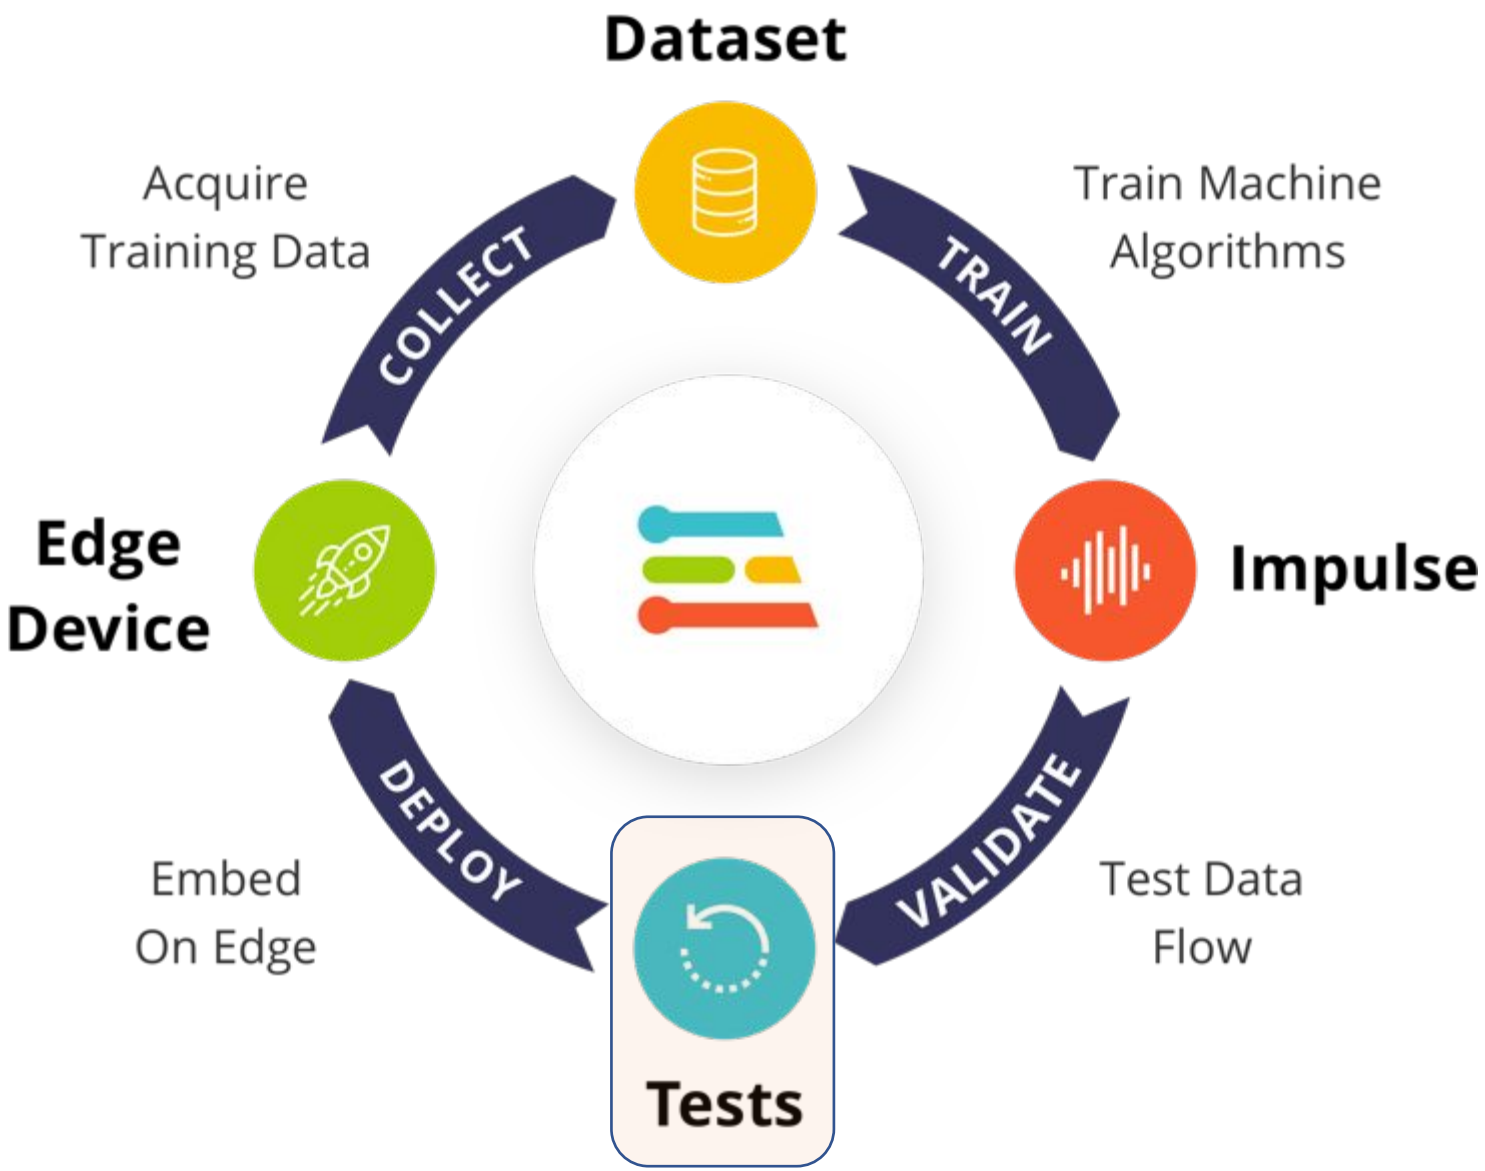

- Pre-Processing Data
- Design a Model
- Train a Model

Evaluate
Optimize

Model testing - SciTinyML-Mot

studio.edgeimpulse.com/studio/51797/validation

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Spectral Analysis
  - Neural Network (Ke...
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

MODEL TESTING (SCITINYML-MOTION-PROJECT)

MJRoBot (Marcelo Rovai)

This lists all test data. You can manage this data through Data acquisition.

Test data

Classify all

Set the 'expected outcome' for each sample to the desired outcome to automatically score the impulse.

SAMPLE NAME	EXPECTED OUTCO...	LENGTH	ACCURACY	RESULT
testing.2hvf...	testing	10s		
terrestrial.2...	terrestrial	10s	100%	98 terrestrial
terrestrial.2...	terrestrial	10s	100%	98 terrestrial
lift .2hssi1t6	lift	10s	100%	98 lift
lift .2hst8tvj	lift	10s	100%	98 lift

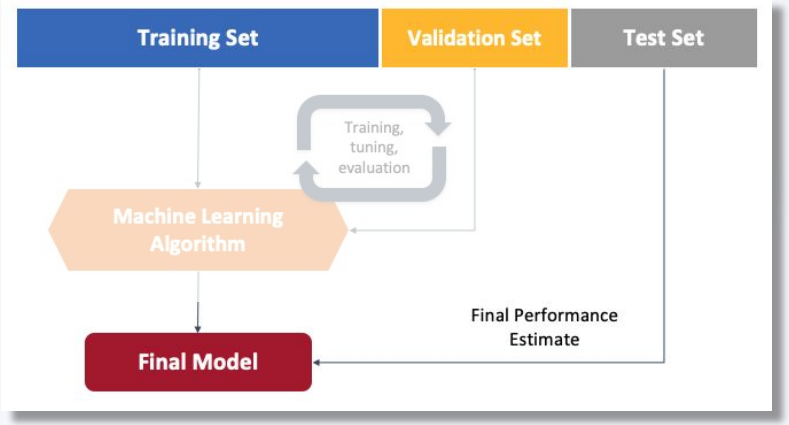

Model testing output

Model testing results

ACCURACY
99.74%

	IDLE	LIFT	MARITIME	TERRESTRIAL	UNCERTAIN
IDLE	99.5%	0.5%	0%	0%	0%
LIFT	0%	100%	0%	0%	0%
MARITIME	0%	0%	99.5%	0%	0.5%
TERRESTRIAL	0%	0%	0%	100%	0%
F1 SCORE	1.00	1.00	1.00	1.00	

Feature explorer

accX RMS | accY RMS | accZ RMS

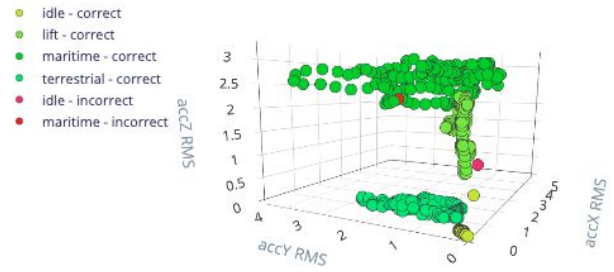

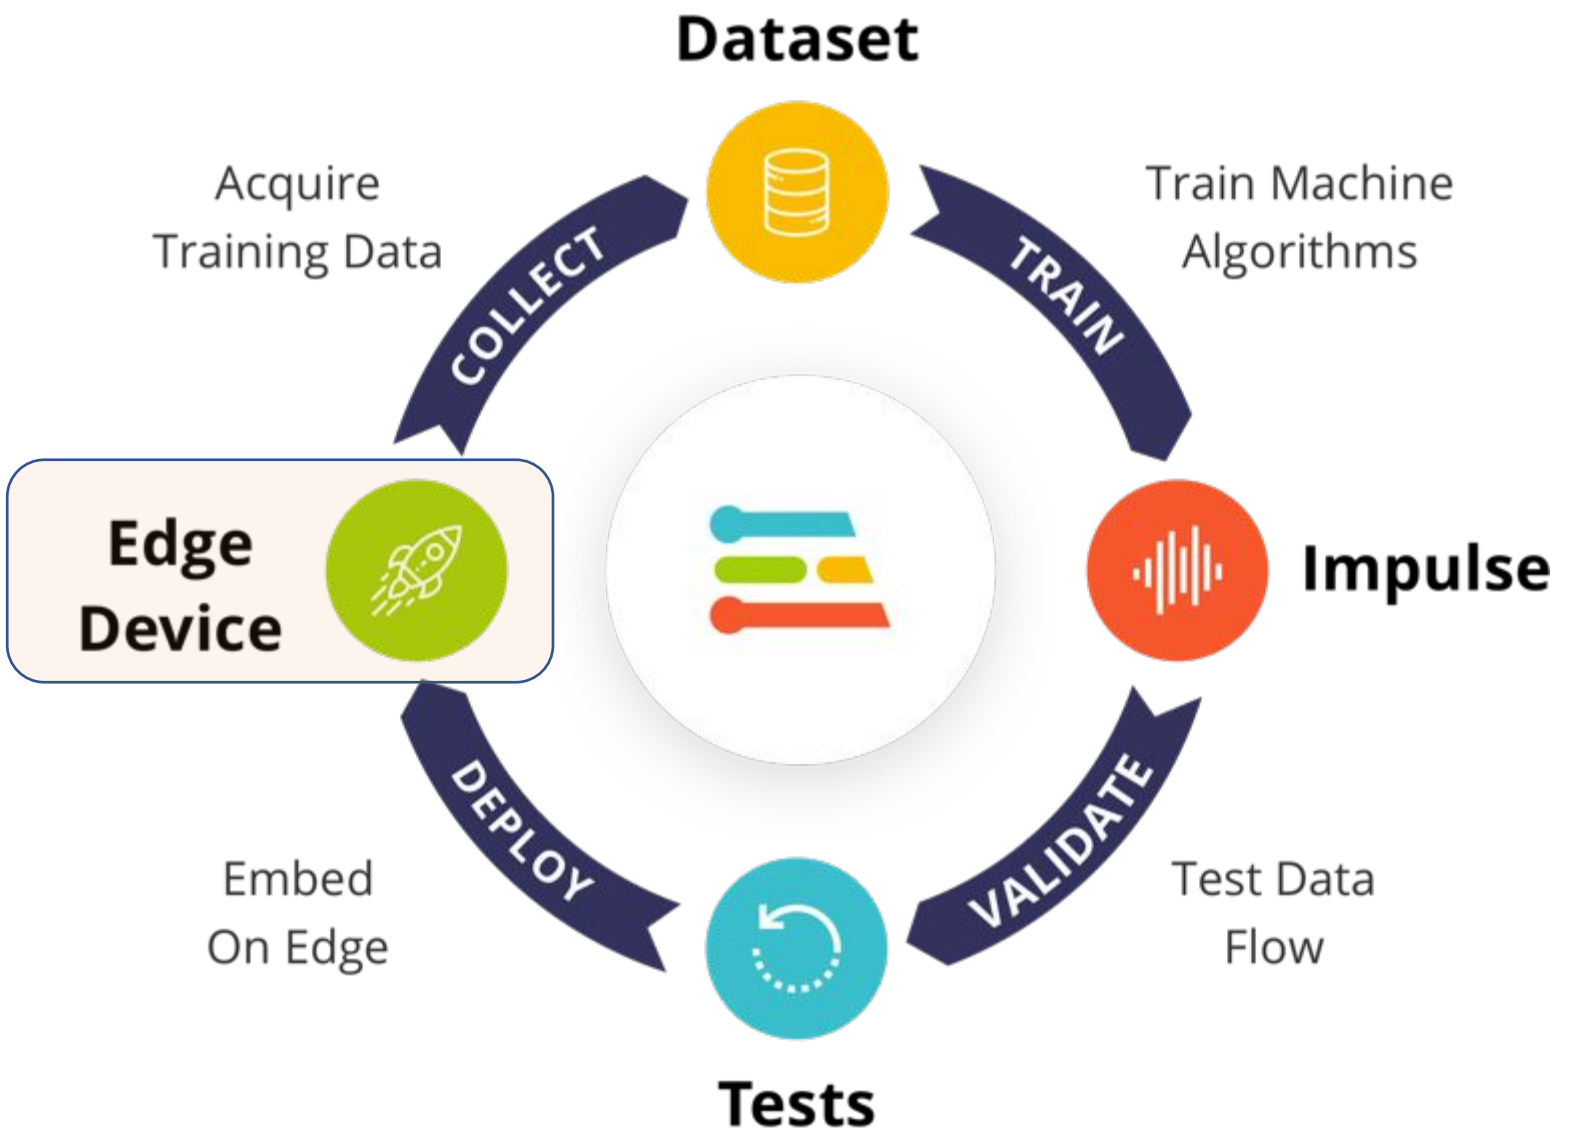

- Pre-Processing Data
- Design a Model
- Train a Model

Convert Model

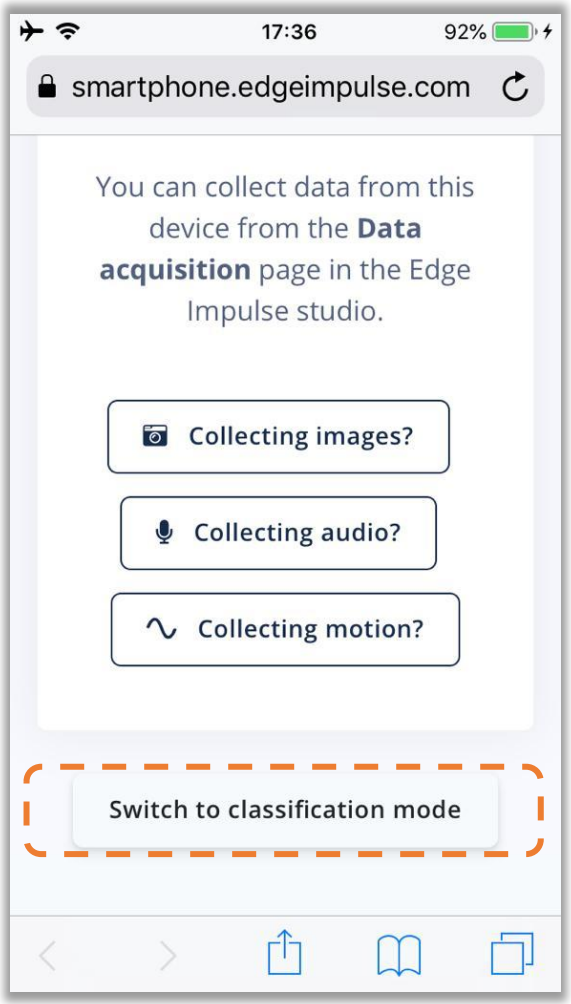

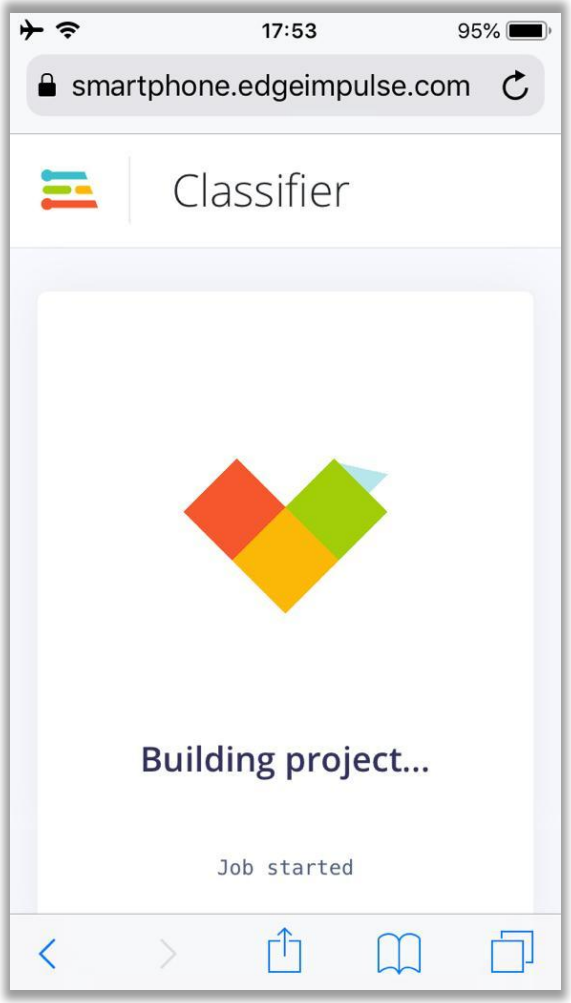

Deploy Model

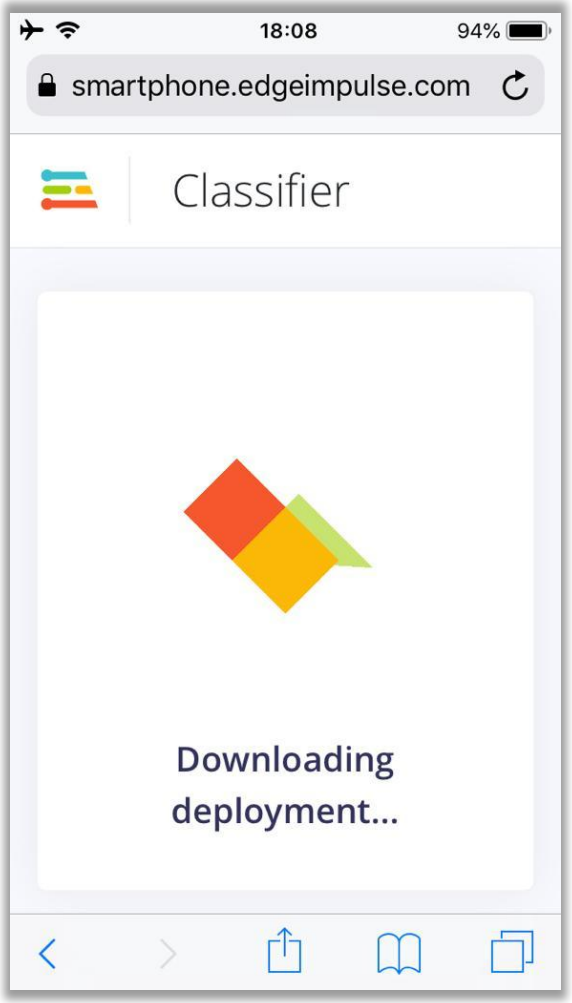

Make Inferences

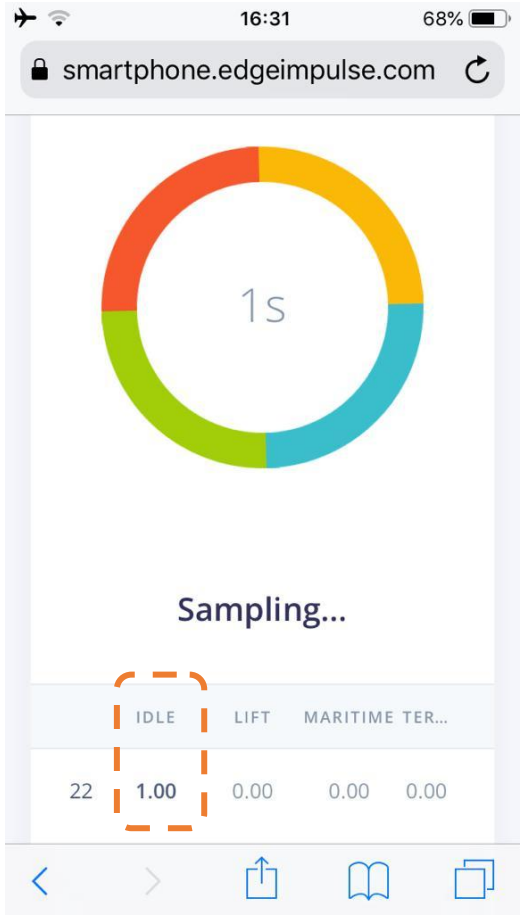

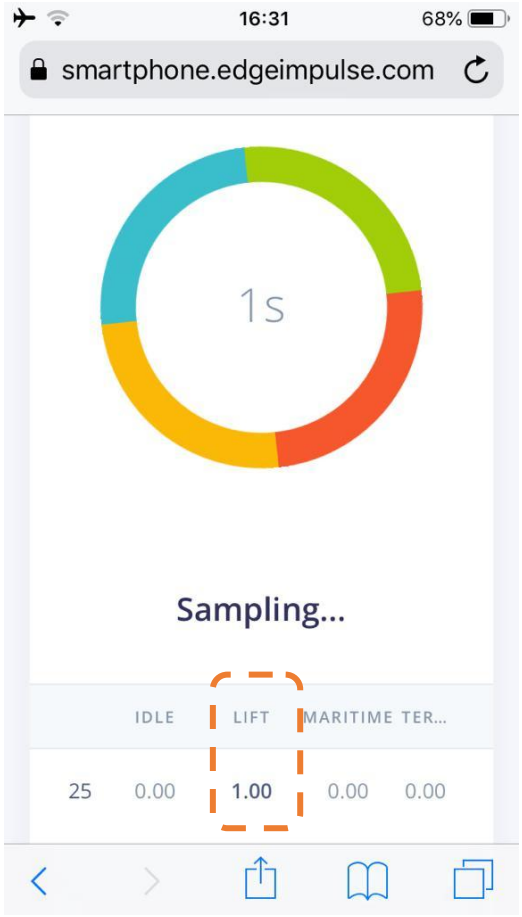

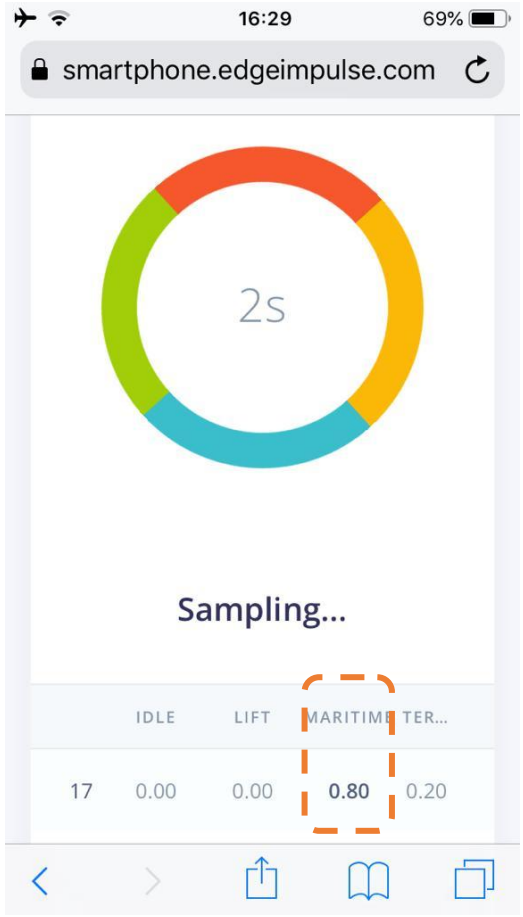

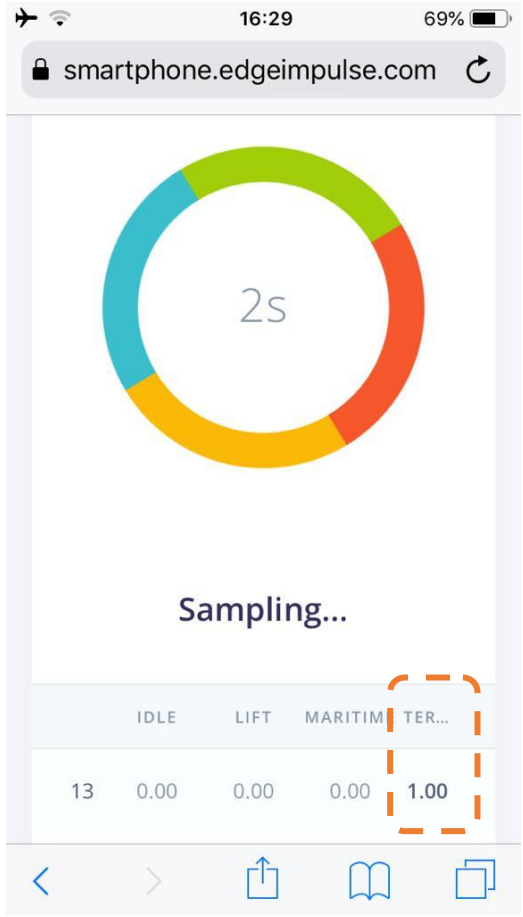

Motion Classification - Summary

Collect Data → Preprocess Data → Design a Model → Train a Model → Evaluate Optimize → Convert Model → Deploy Model → Make Inferences

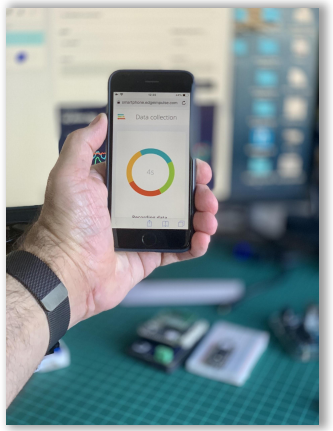

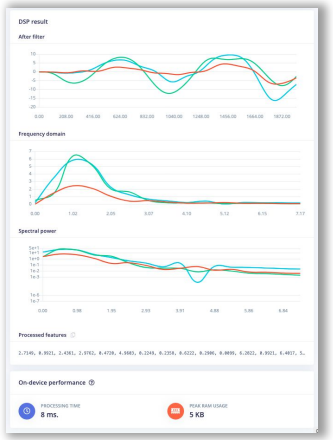

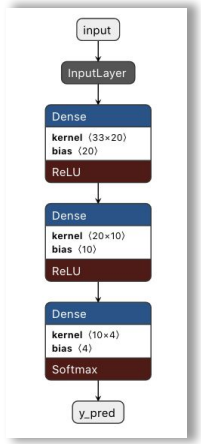

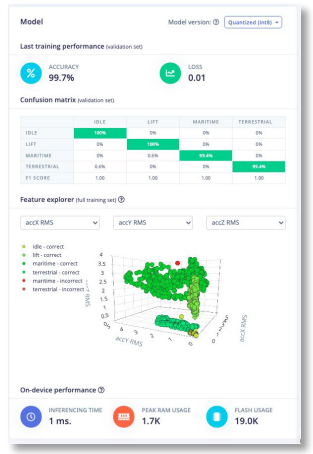

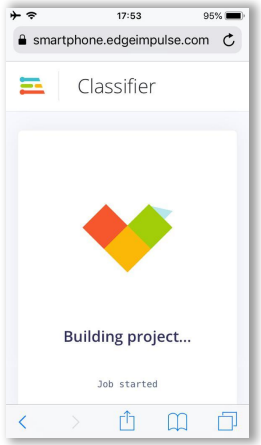

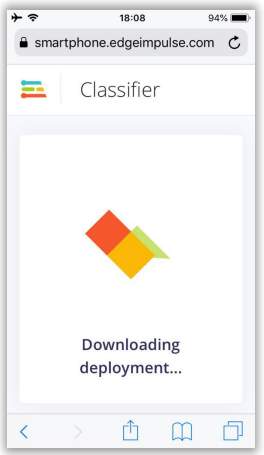

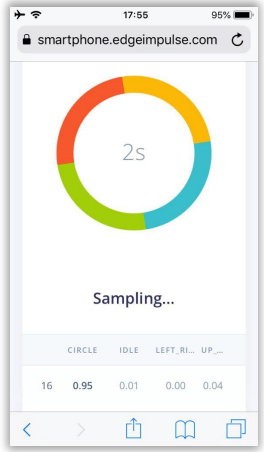

Spectral Analysis

NN Classifier

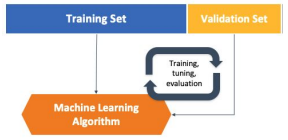

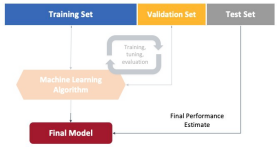

Anomaly Detection

What is **Anomaly Detection**?

In **data analysis, anomaly detection** is the **identification of rare** items, events or observations which **raise suspicions** because they **differing significantly** from the **majority of the data**.

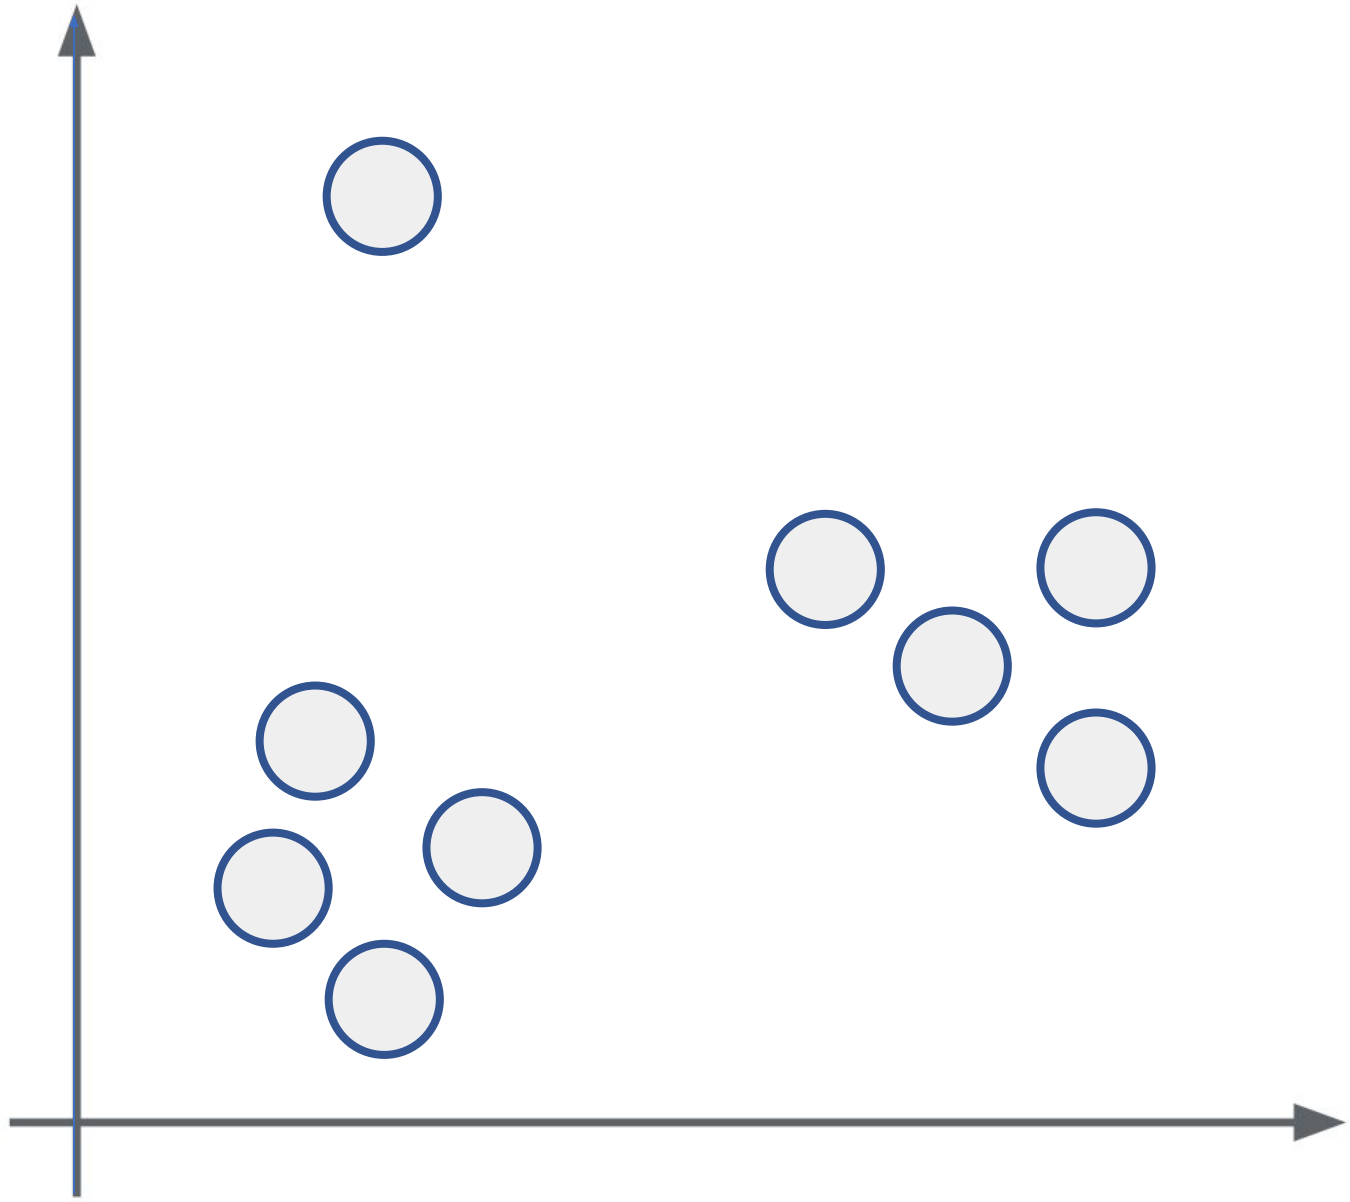

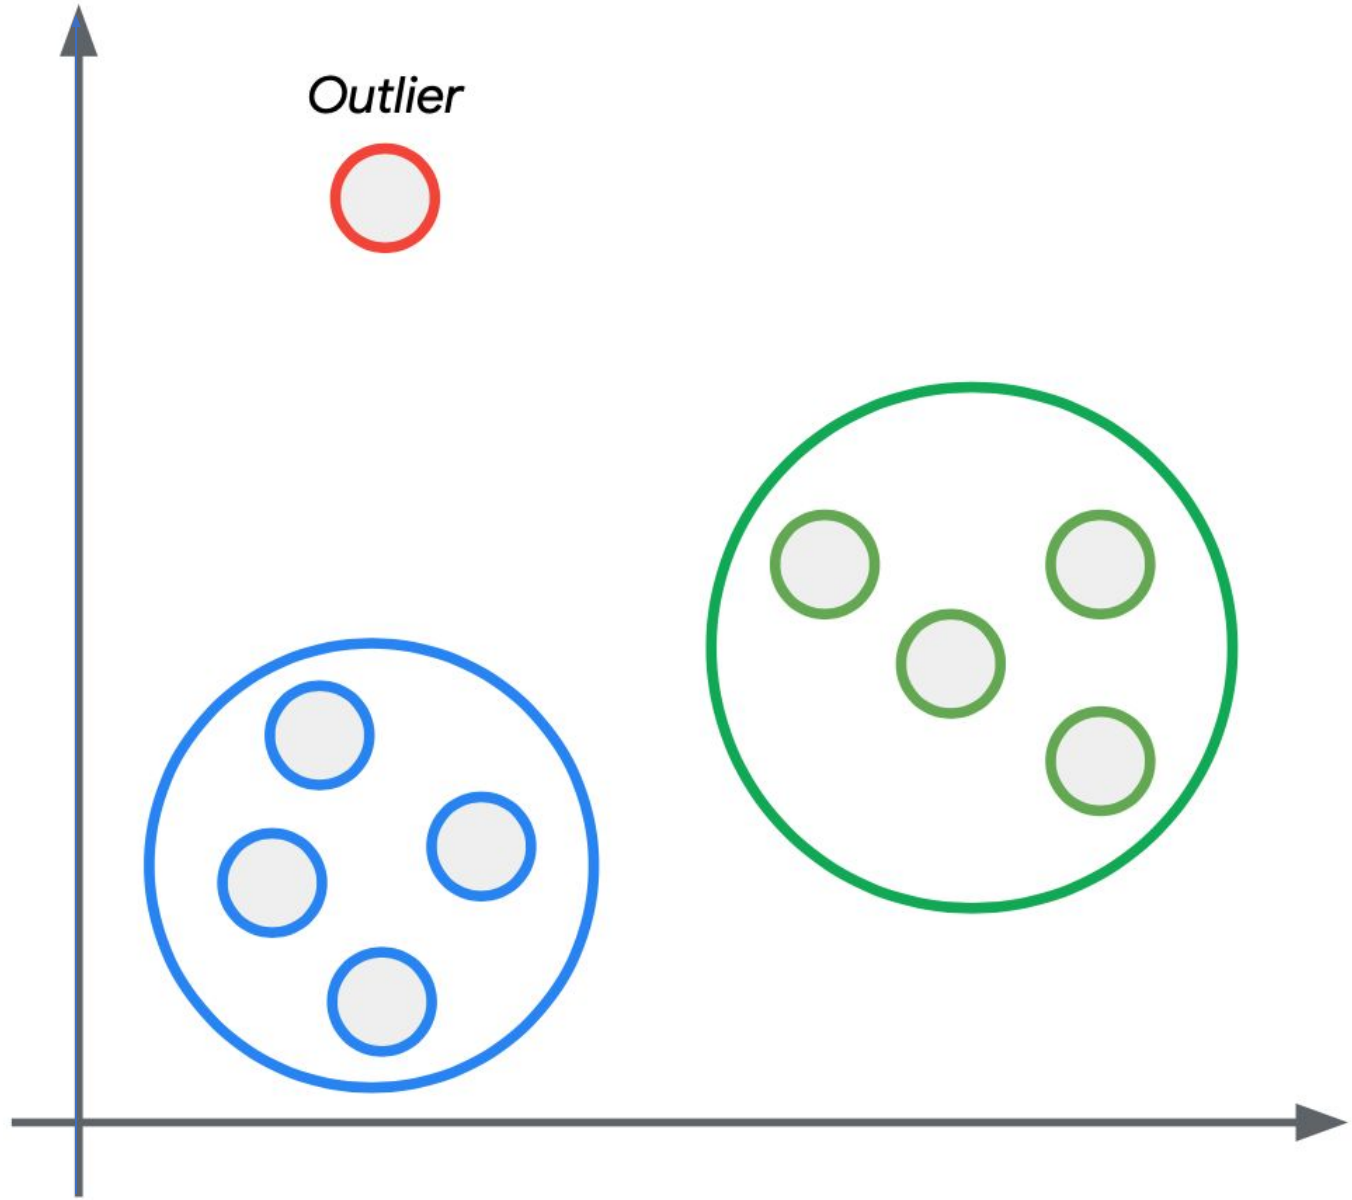
Outlier

Application: Factory machinery

Application: Factory machinery

Ball Bearings

Accelerometer

Sensor: **Accelerometer**

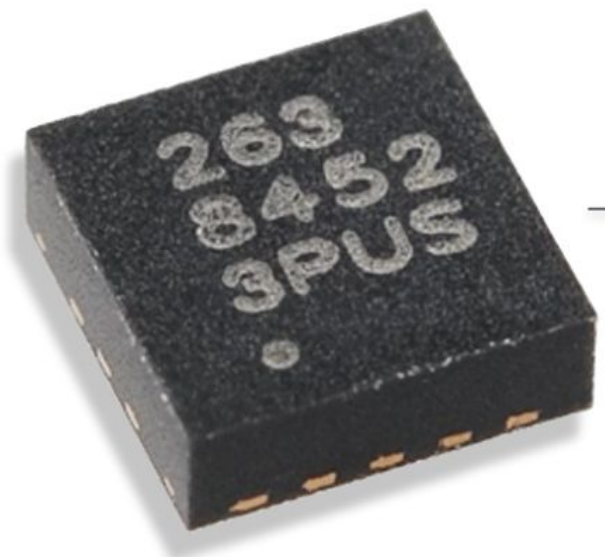

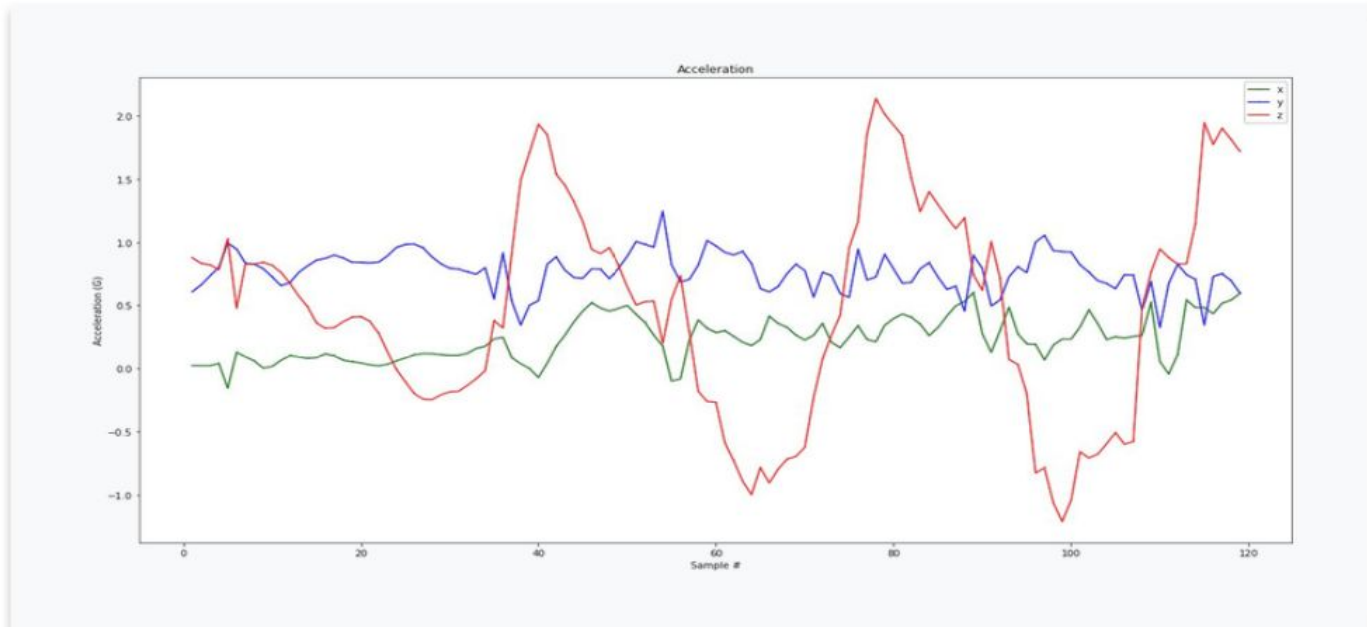

Sensor: **Accelerometer**

2 bytes × 8 × 20kHz = **320** KB / sec

Measurement

Sensors

Sample Rate

It’s too expensive to stream to the cloud

Need “**intelligence**” *close* **to sensors**

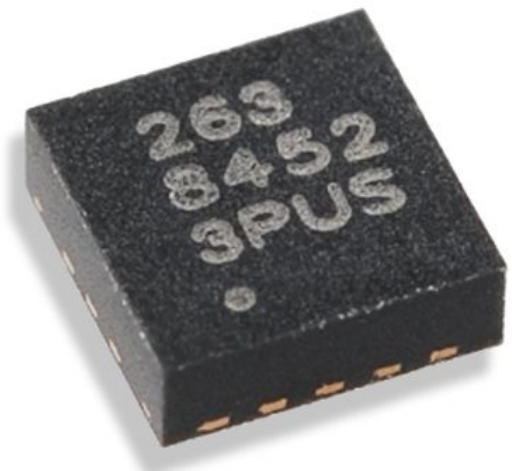

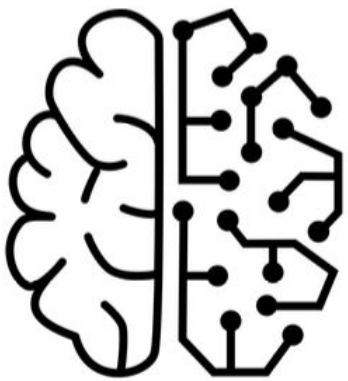

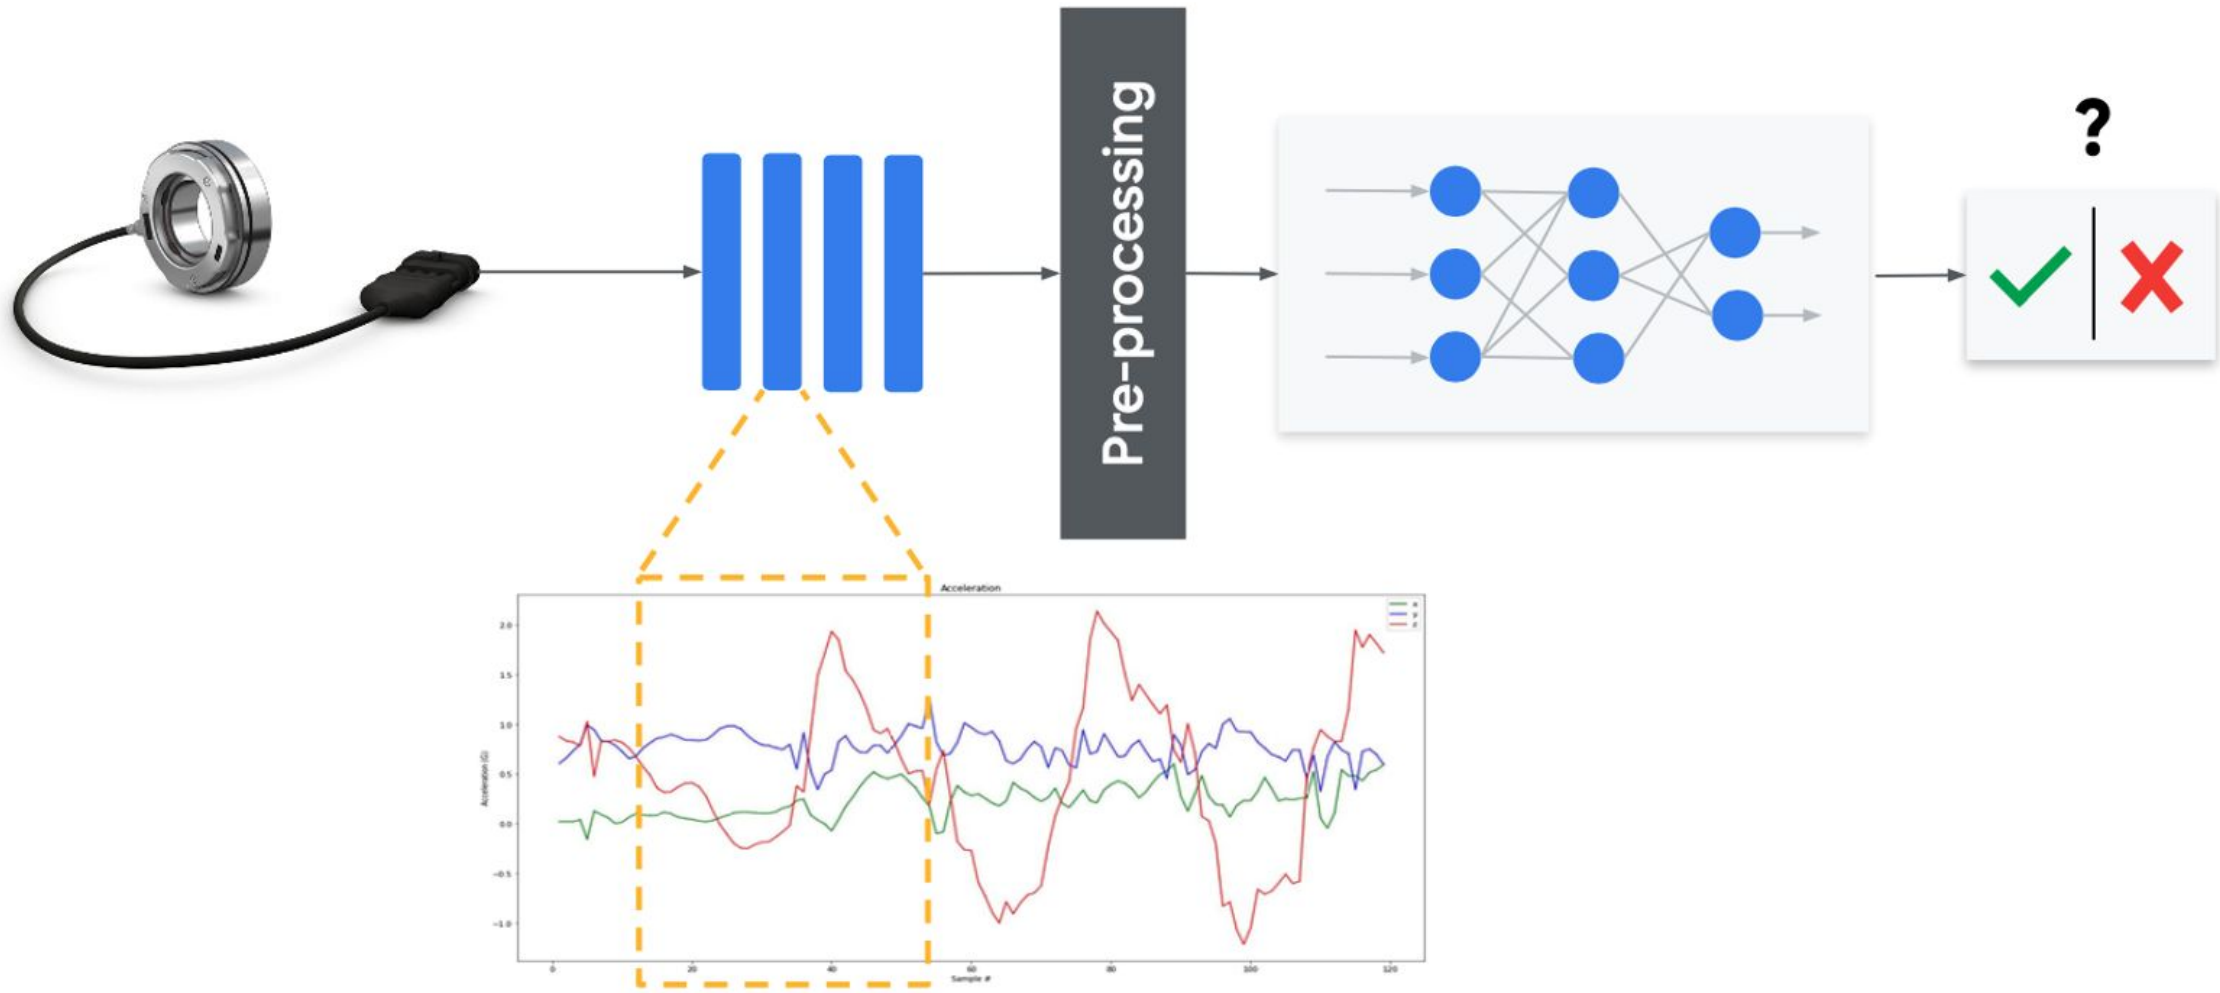
Pre-processing
?
Acceleration
Acceleration (G)
Sample #

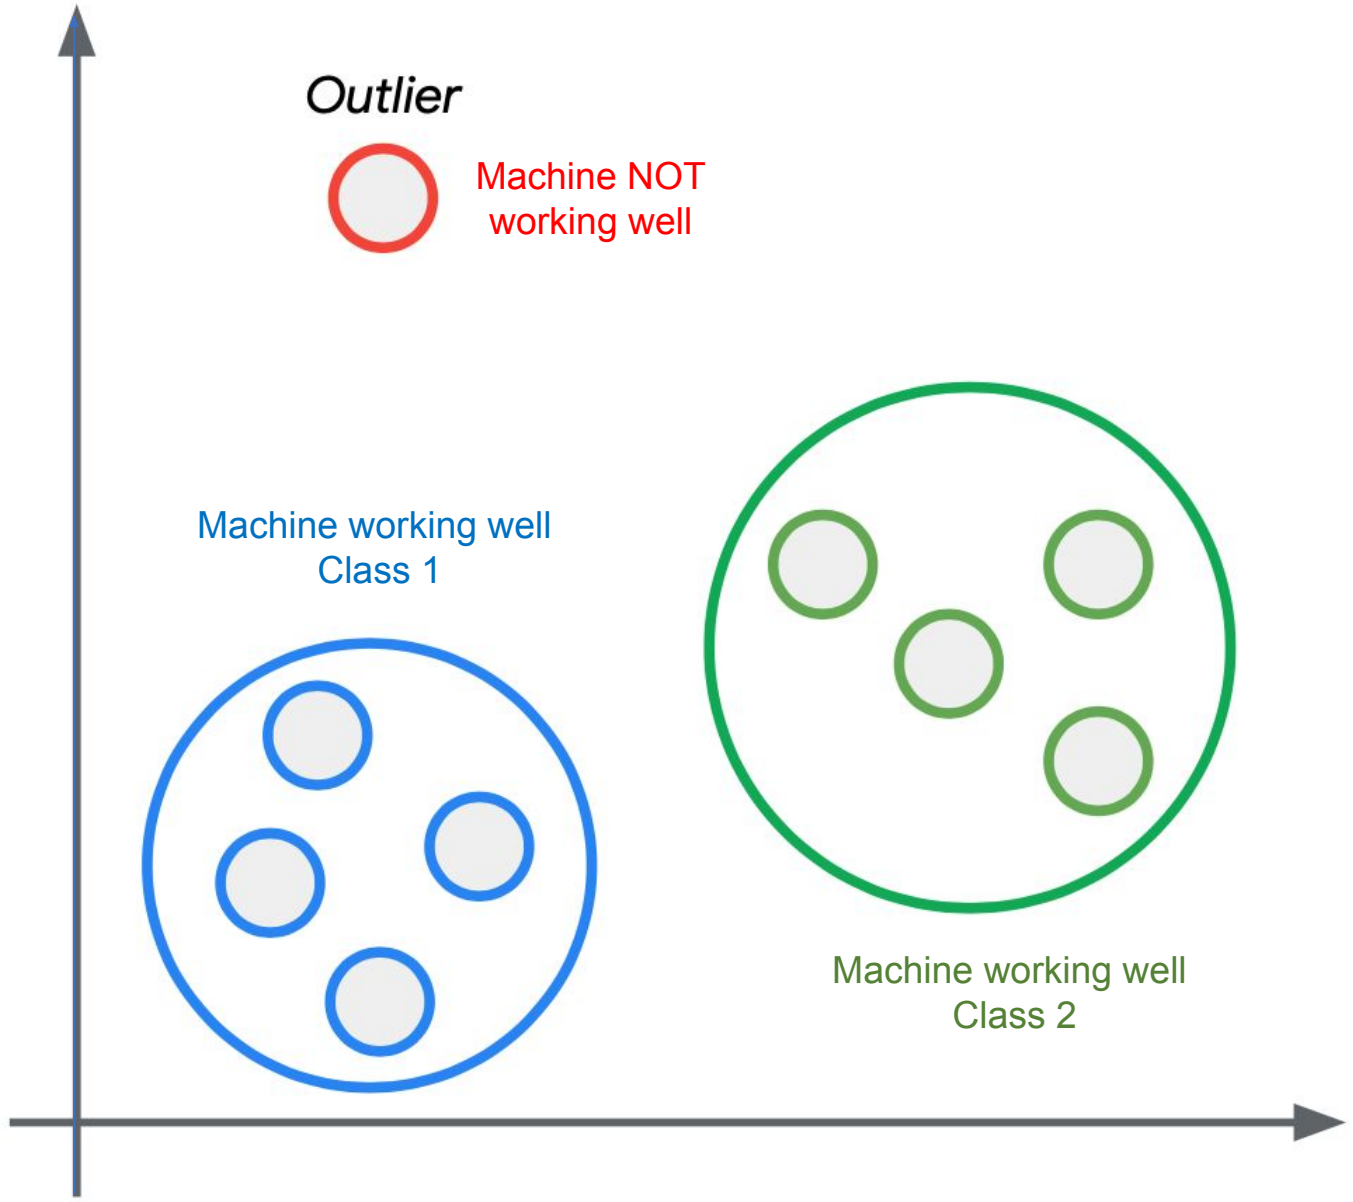
Outlier
Machine NOT working well
Machine working well Class 1
Machine working well Class 2

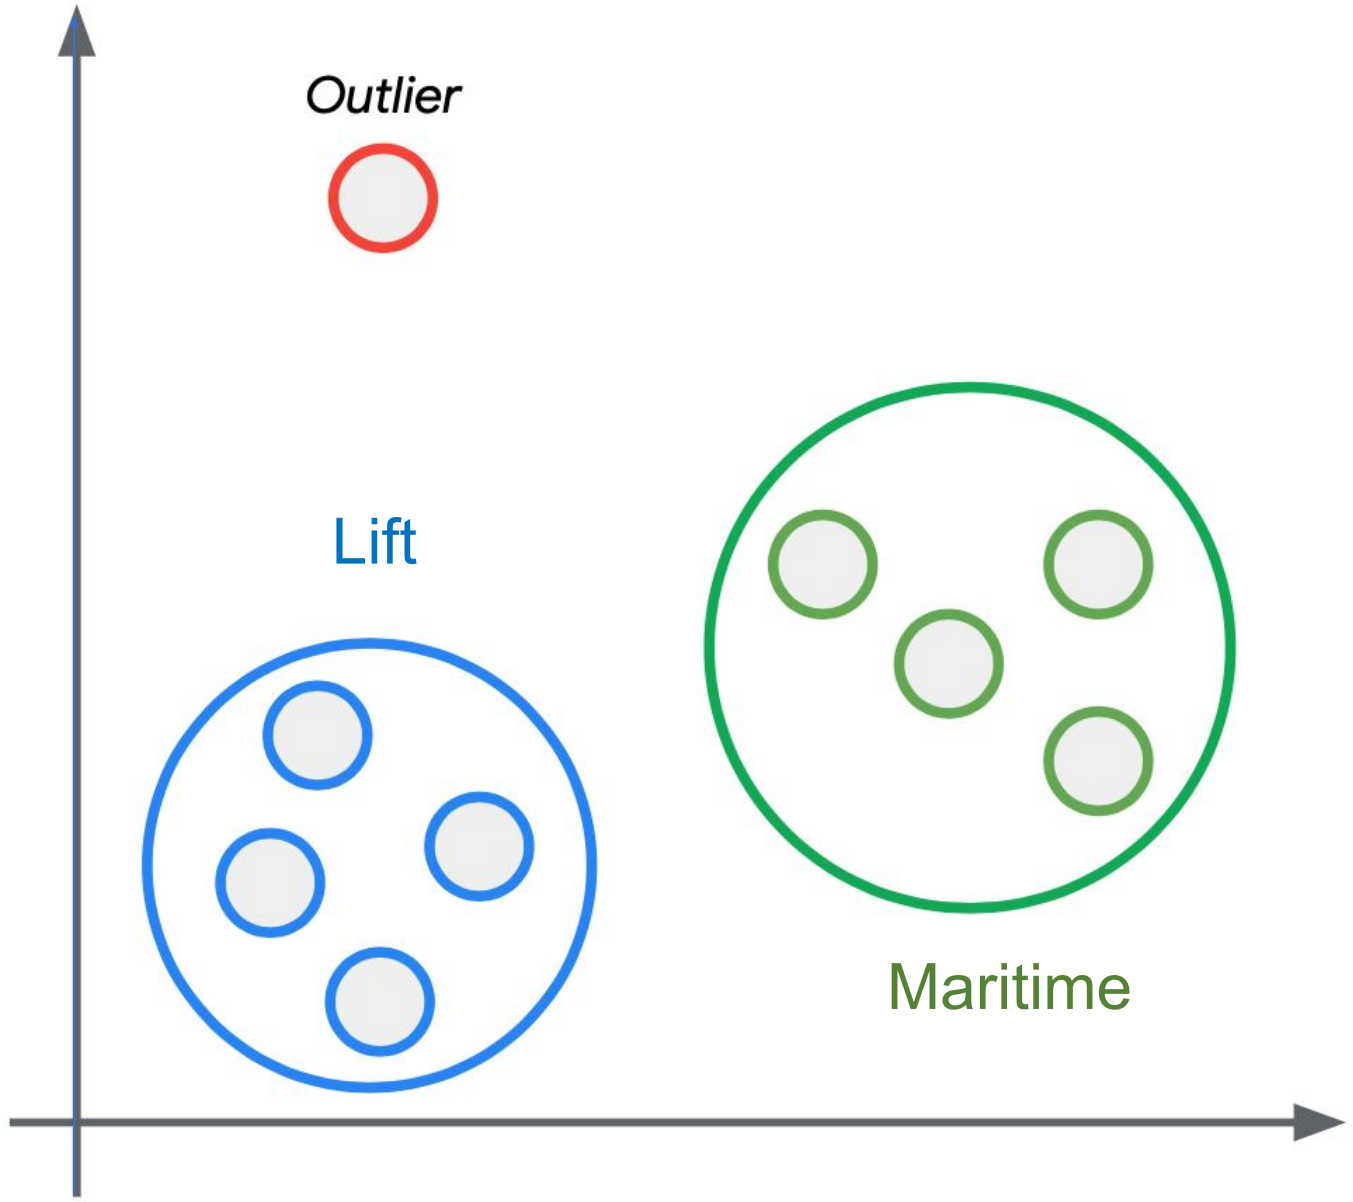
Outlier
Lift
Maritime

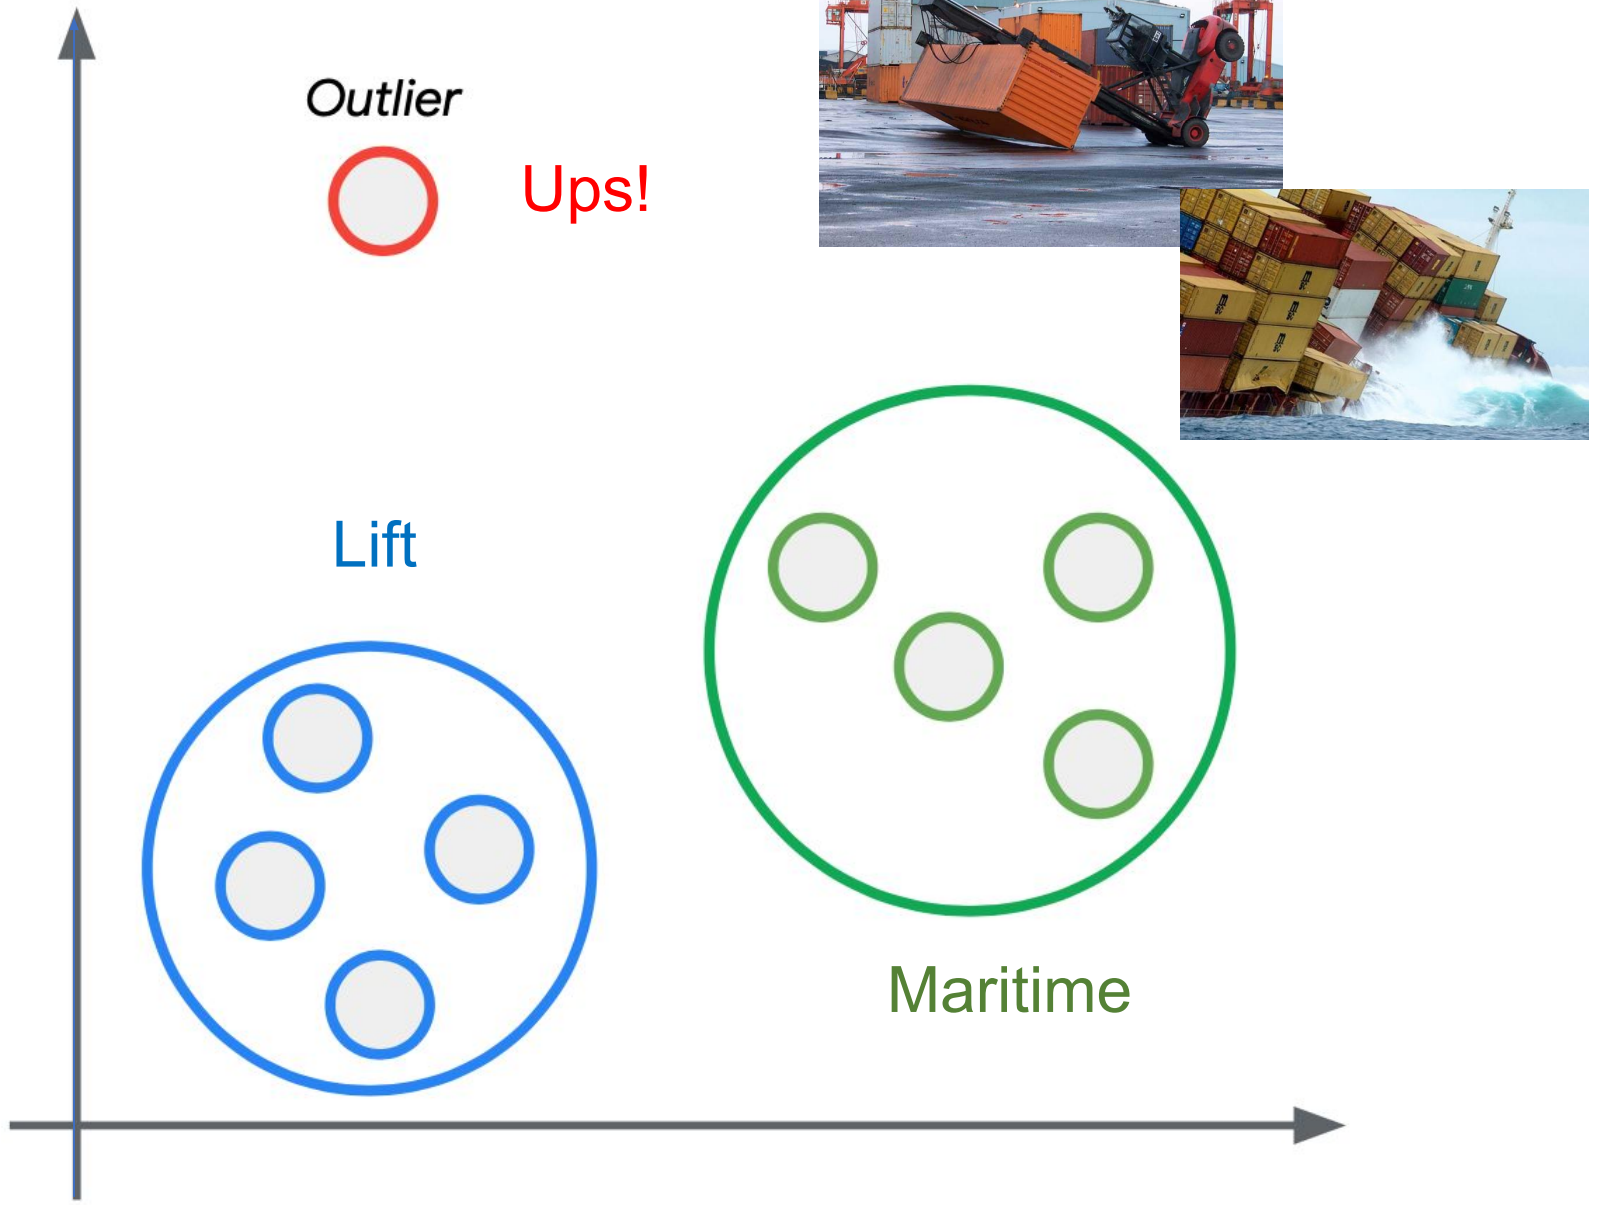
Outlier
Ups!
Lift
Maritime

It's **not all** deep learning

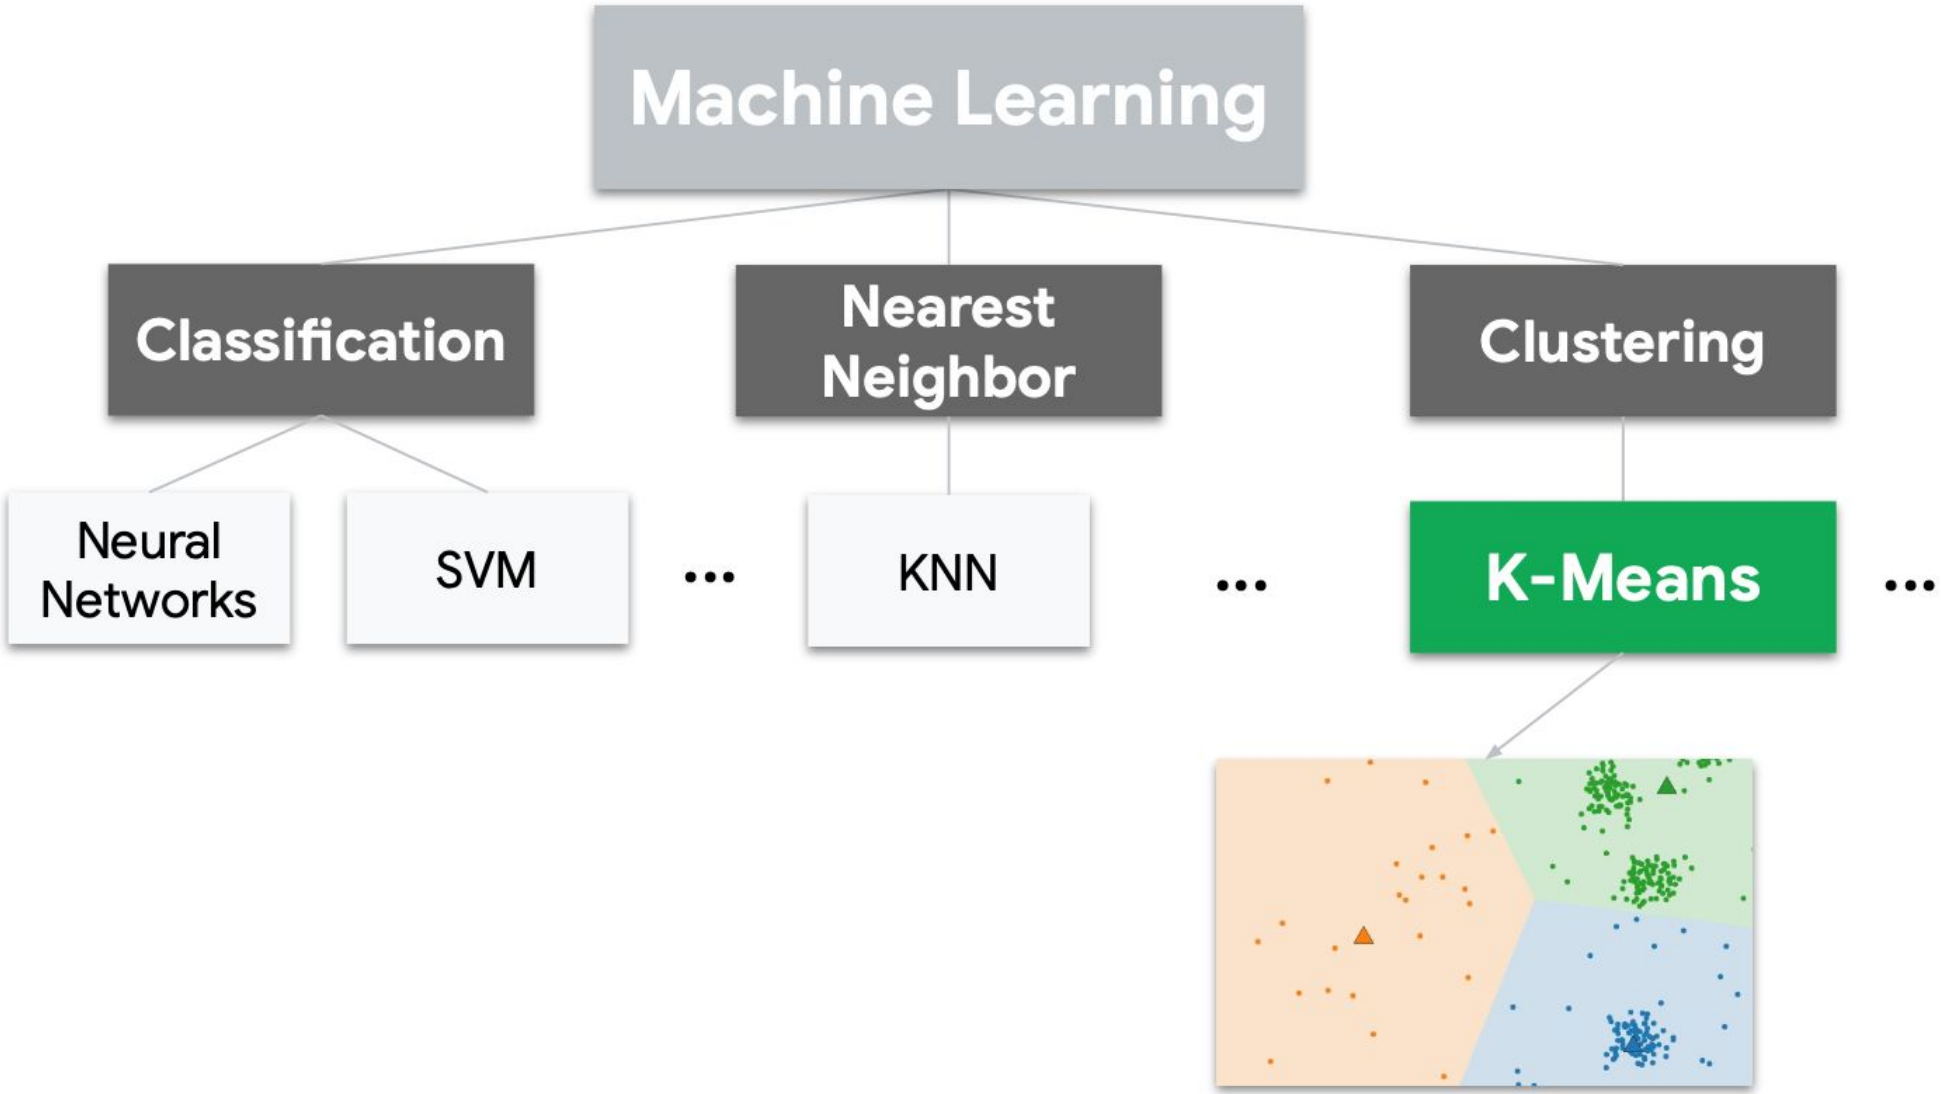

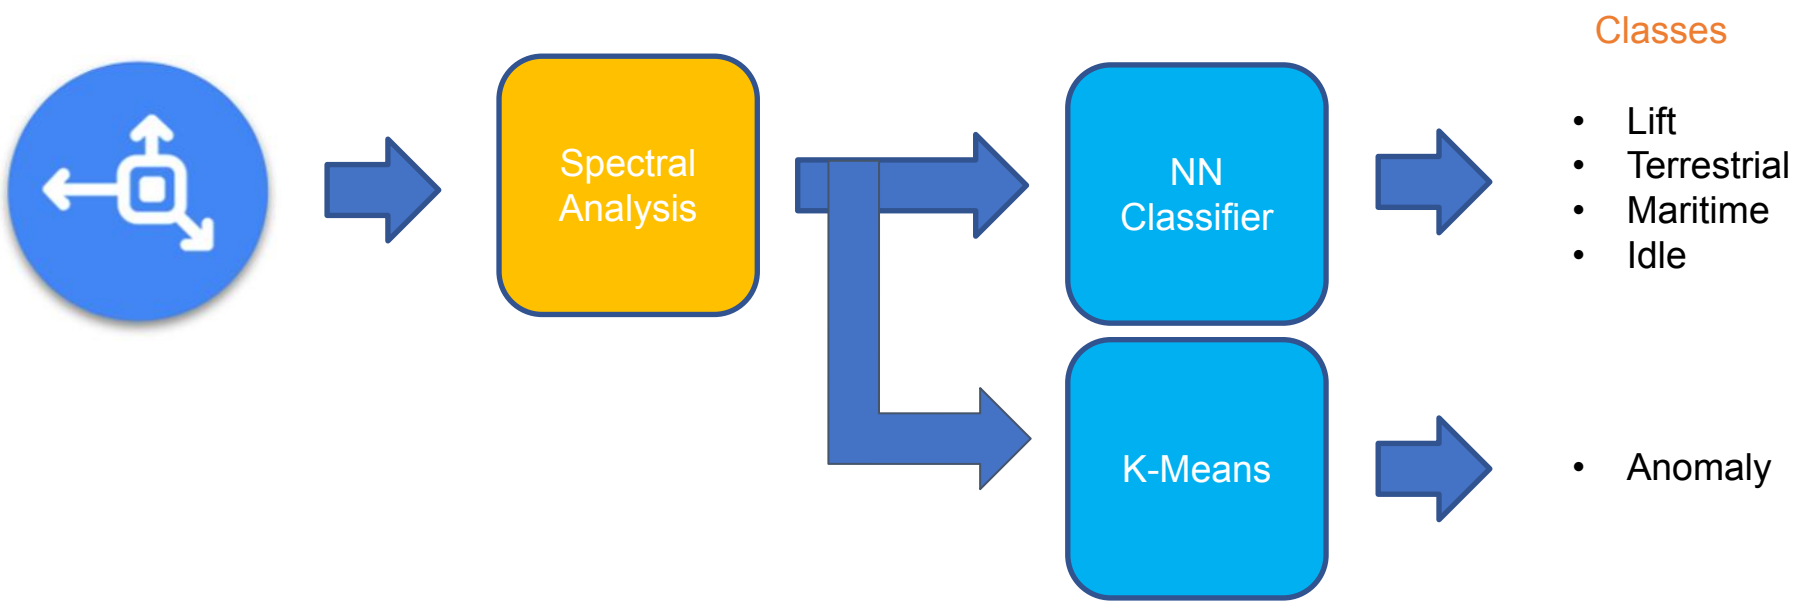
Spectral Analysis
NN Classifier
K-Means
Classes
Lift
Terrestrial
Maritime
Idle
Anomaly

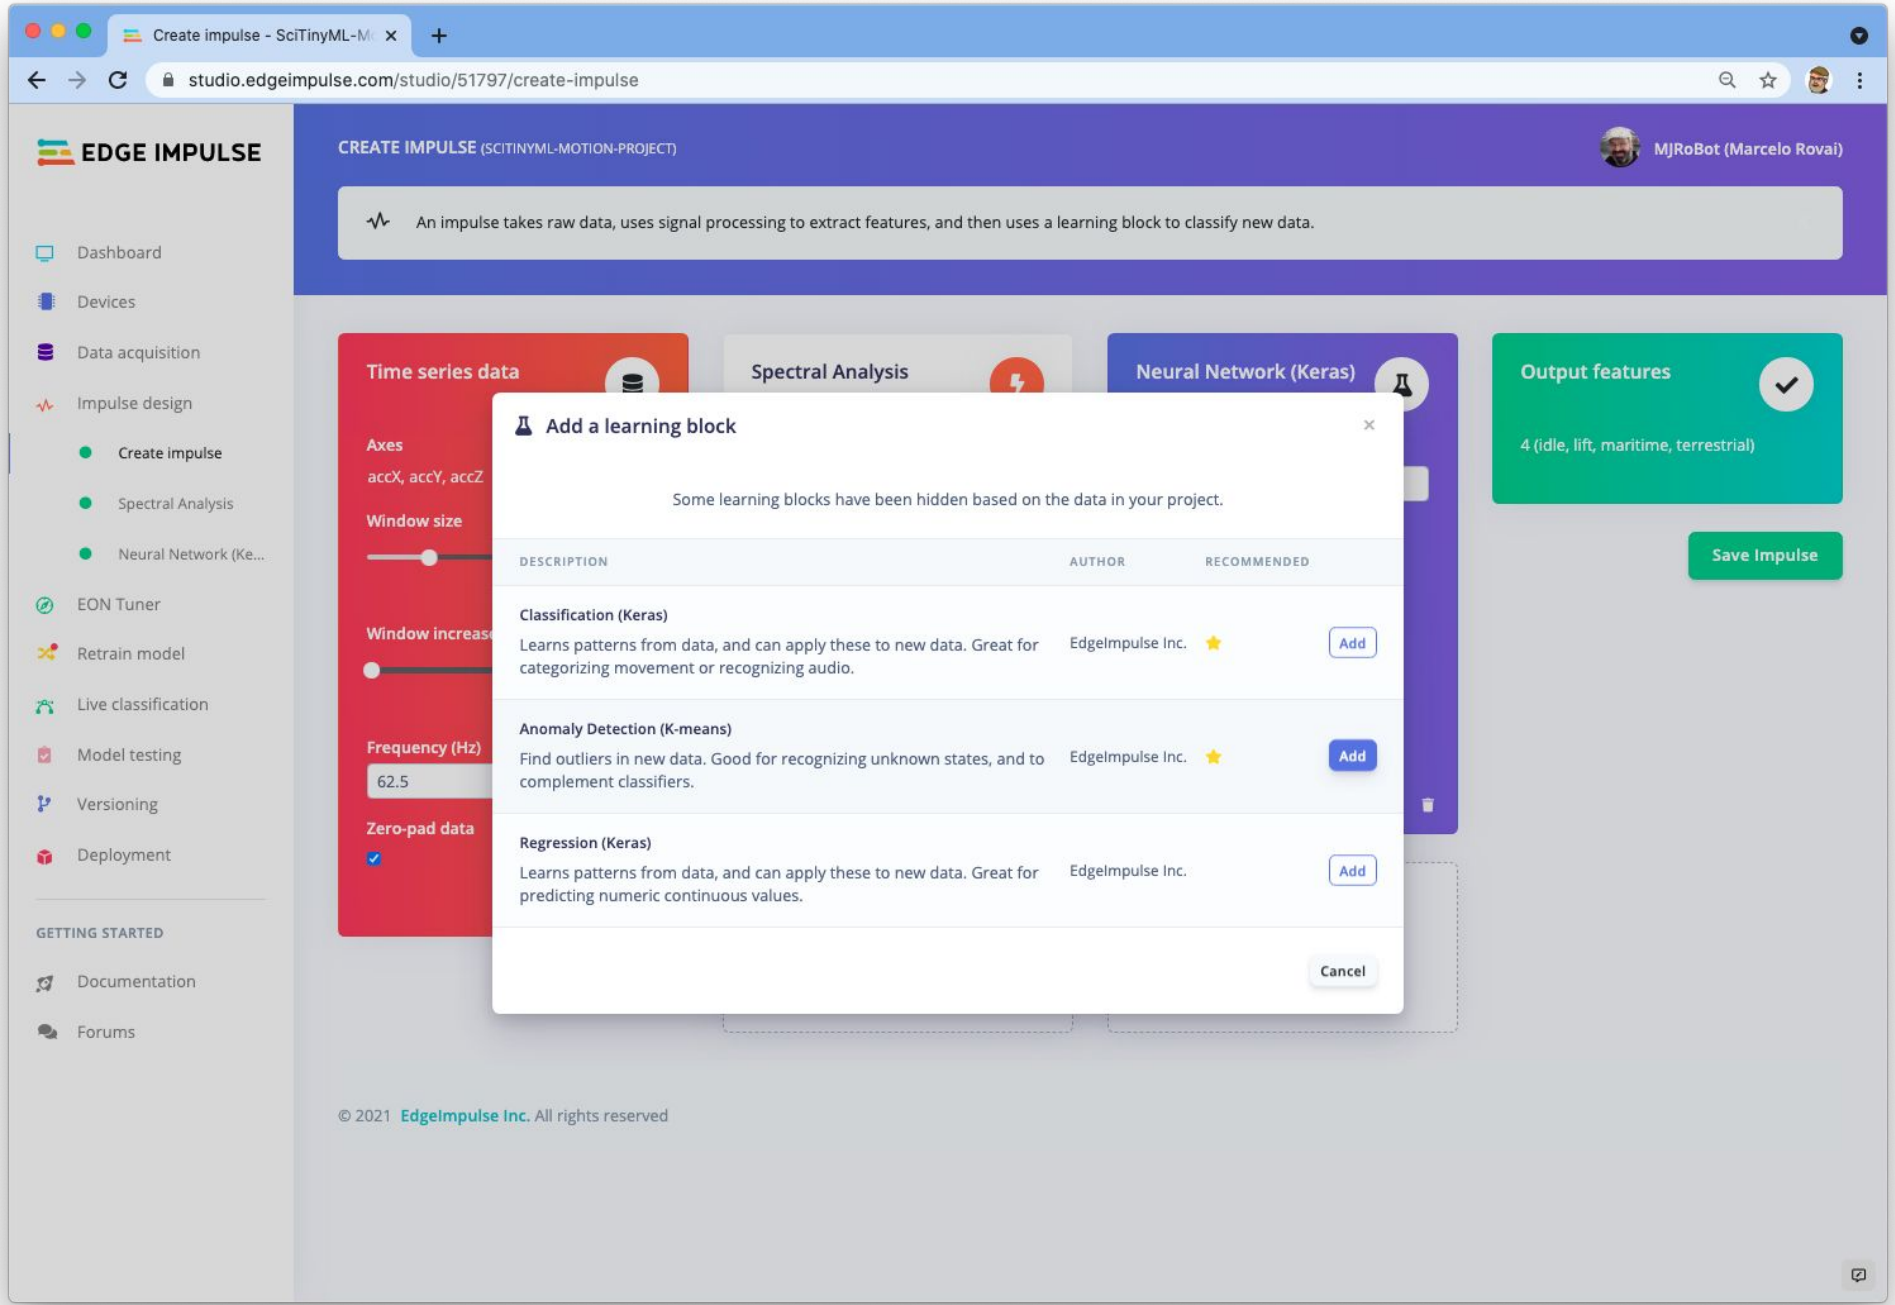
Create impulse - SciTinyML-M
studio.edgeimpulse.com/studio/51797/create-impulse
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Spectral Analysis
Neural Network (Ke...
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
CREATE IMPULSE (SCITINYML-MOTION-PROJECT)
MJRoBot (Marcelo Rovai)
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Time series data
Axes
accX, accY, accZ
Window size
Window increase
Frequency (Hz)
62.5
Zero-pad data
Spectral Analysis
Neural Network (Keras)
Output features
4 (idle, lift, maritime, terrestrial)
Save Impulse
Add a learning block
Some learning blocks have been hidden based on the data in your project.
DESCRIPTION
AUTHOR
RECOMMENDED
Classification (Keras)
Learns patterns from data, and can apply these to new data. Great for categorizing movement or recognizing audio.
EdgeImpulse Inc.
Add
Anomaly Detection (K-means)
Find outliers in new data. Good for recognizing unknown states, and to complement classifiers.
EdgeImpulse Inc.
Add
Regression (Keras)
Learns patterns from data, and can apply these to new data. Great for predicting numeric continuous values.
EdgeImpulse Inc.
Add
Cancel
© 2021 EdgeImpulse Inc. All rights reserved

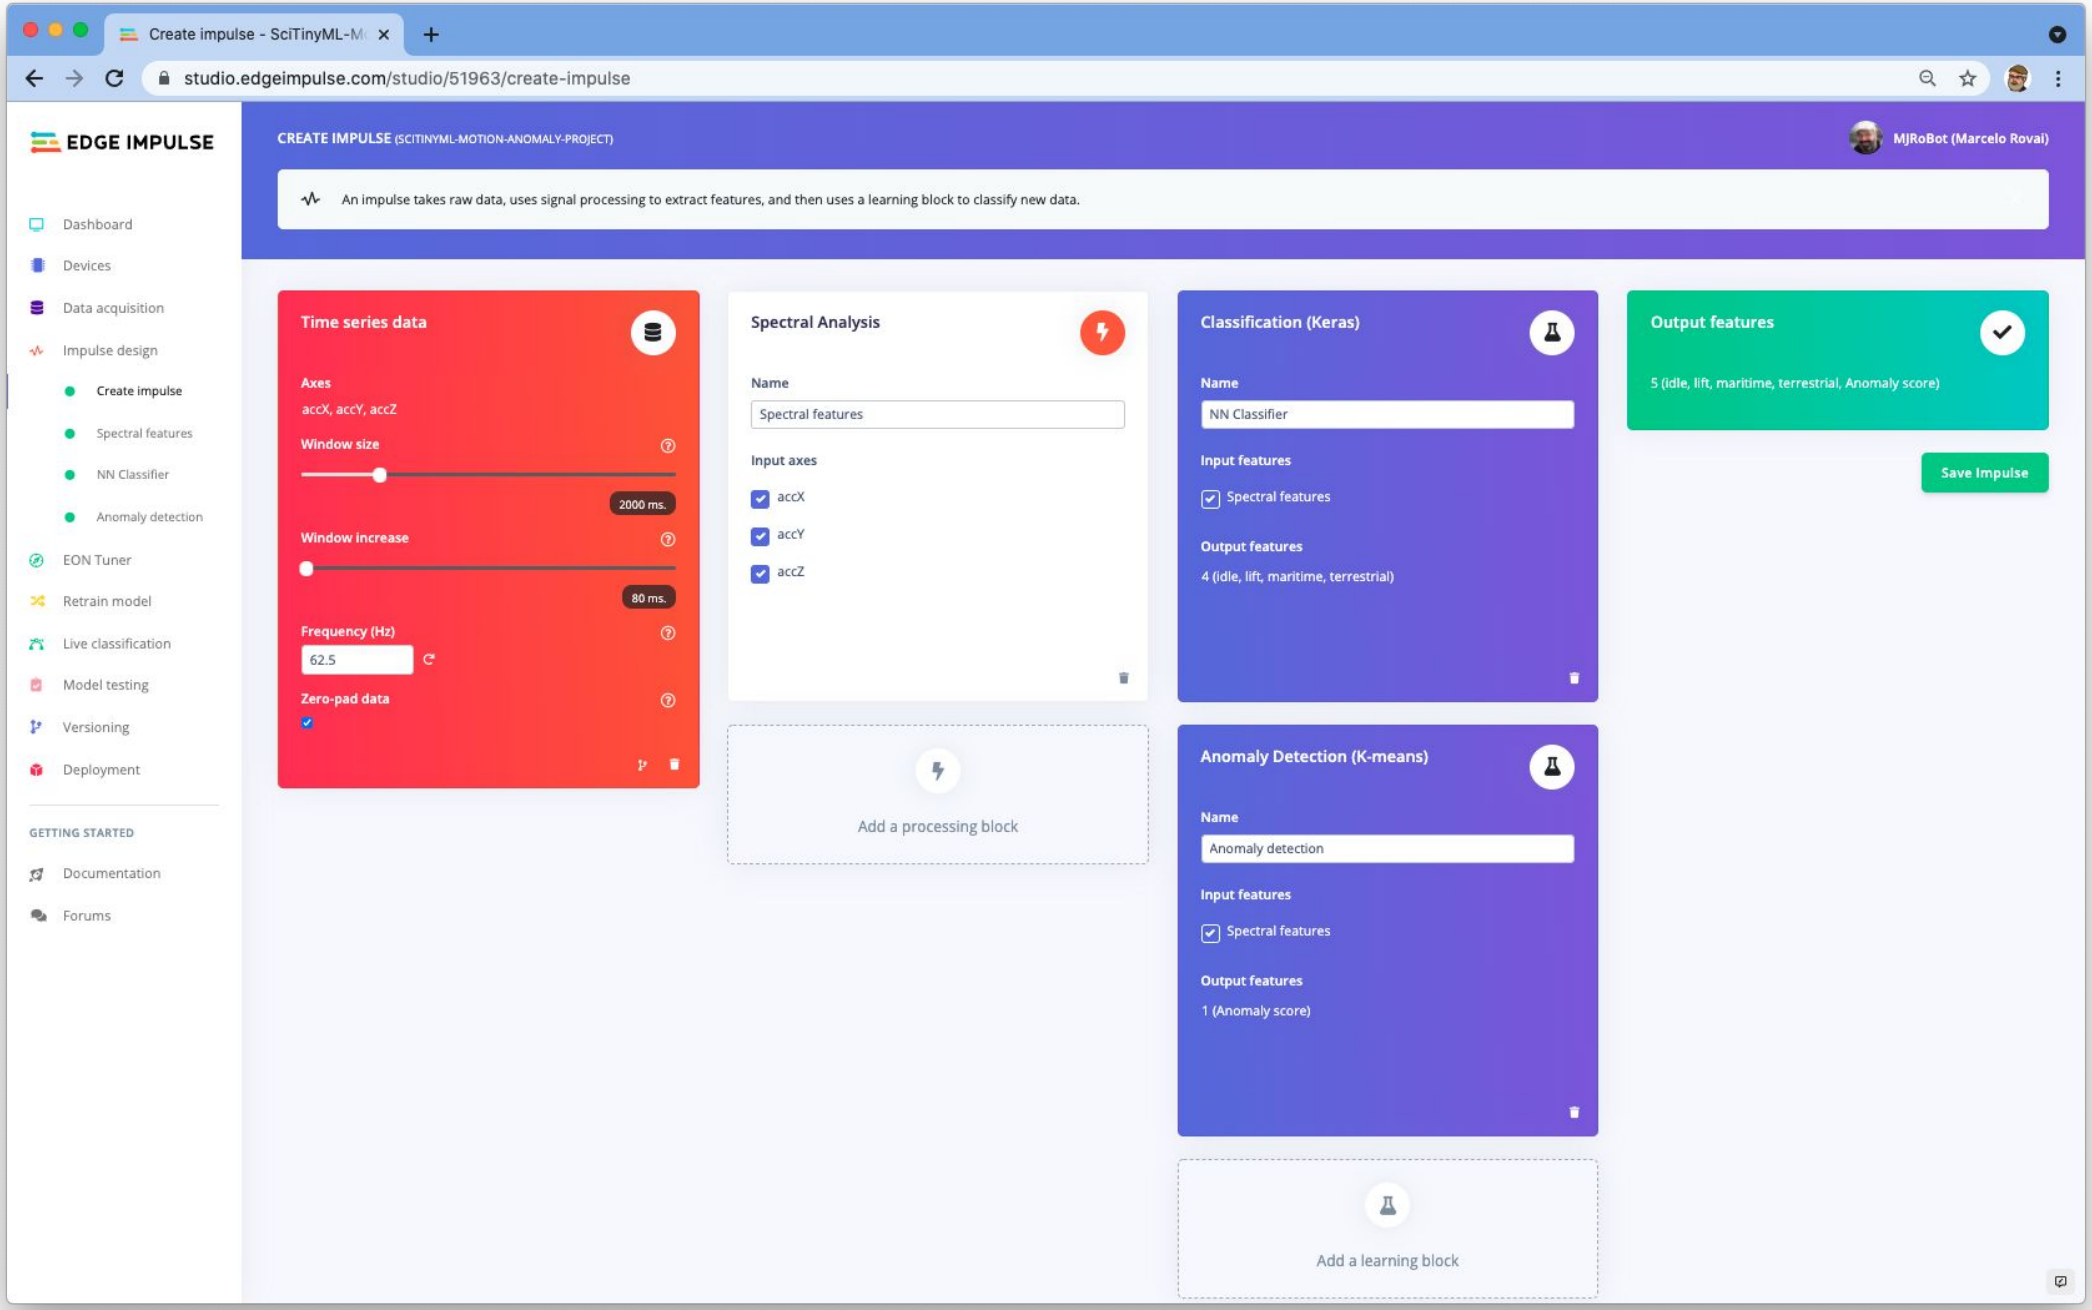
Create impulse - SciTinyML-M
studio.edgeimpulse.com/studio/51963/create-impulse
EDGE IMPULSE
CREATE IMPULSE (SCITINYML-MOTION-ANOMALY-PROJECT)
MJRoBot (Marcelo Rovai)
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Spectral features
NN Classifier
Anomaly detection
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Time series data
Axes
accX, accY, accZ
Window size
2000 ms.
Window increase
80 ms.
Frequency (Hz)
62.5
Zero-pad data
Spectral Analysis
Name
Spectral features
Input axes
accX
accY
accZ
Add a processing block
Classification (Keras)
Name
NN Classifier
Input features
Spectral features
Output features
4 (idle, lift, maritime, terrestrial)
Anomaly Detection (K-means)
Name
Anomaly detection
Input features
Spectral features
Output features
1 (Anomaly score)
Add a learning block
Output features
5 (idle, lift, maritime, terrestrial, Anomaly score)
Save Impulse

Anomaly detection - SciTinyML

studio.edgeimpulse.com/studio/51963/learning/anomaly/52

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Spectral features
  - NN Classifier
  - Anomaly detection
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

ANOMALY DETECTION (SCITINYML-MOTION-ANOMALY-PROJECT)

#1 Click to set a description for this version

MJRoBot (Marcelo Rovai)

Anomaly detection settings

Cluster count

4

Axes

Select all axes

- [x] accX RMS
- [ ] accX Peak 1 Freq
- [ ] accX Peak 1 Height
- [ ] accX Peak 2 Freq
- [ ] accX Peak 2 Height
- [ ] accX Peak 3 Freq
- [ ] accX Peak 3 Height
- [ ] accX Spectral Power 0.1 - 0.5
- [ ] accX Spectral Power 0.5 - 1.0
- [ ] accX Spectral Power 1.0 - 2.0
- [ ] accX Spectral Power 2.0 - 5.0
- [x] accY RMS
- [ ] accY Peak 1 Freq
- [ ] accY Peak 1 Height
- [ ] accY Peak 2 Freq
- [ ] accY Peak 2 Height
- [ ] accY Peak 3 Freq
- [ ] accY Peak 3 Height
- [ ] accY Spectral Power 0.1 - 0.5
- [ ] accY Spectral Power 0.5 - 1.0
- [ ] accY Spectral Power 1.0 - 2.0
- [ ] accY Spectral Power 2.0 - 5.0
- [x] accZ RMS
- [ ] accZ Peak 1 Freq
- [ ] accZ Peak 1 Height
- [ ] accZ Peak 2 Freq
- [ ] accZ Peak 2 Height
- [ ] accZ Peak 3 Freq
- [ ] accZ Peak 3 Height
- [ ] accZ Spectral Power 0.1 - 0.5
- [ ] accZ Spectral Power 0.5 - 1.0
- [ ] accZ Spectral Power 1.0 - 2.0
- [ ] accZ Spectral Power 2.0 - 5.0

Start training

Anomaly explorer (3,230 samples)

X Axis: accX RMS | Y Axis: accY RMS | Test data: -- No test data

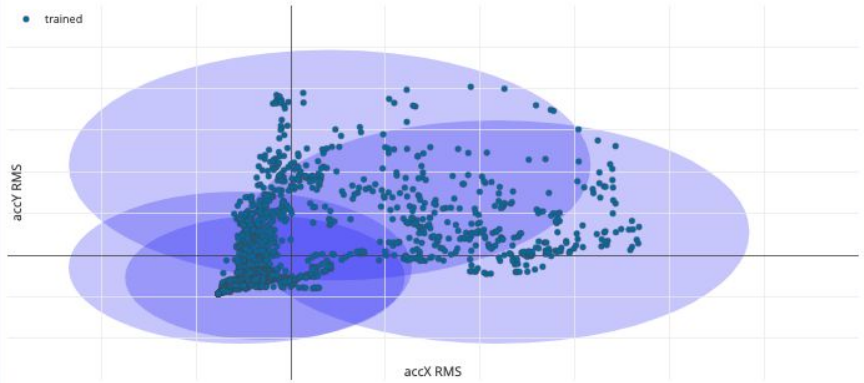

Training output

```
Copying features from processing blocks...
Copying features from DSP block...
Copying features from DSP block OK
Copying features from processing blocks OK

Training model
Job started
scaler scale [1.23777729 1.02773968 1.10088427] mean [0.95382248 0.94990646 1.12868147] var [1.53209261 1.05624885
1.21194617]
trained_clusters [{'center': [-0.5379795432090759, -0.30185389518737793, -0.8996922373771667], 'max_error':
1.8050675006641951}, {'center': [-0.2765962481498718, -0.5444689393043518, 0.5496397018432617], 'max_error':
1.4696349225868046}, {'center': [0.40855732560157776, 2.160626173019409, 1.2495908737182617], 'max_error':
2.7492433102802676}, {'center': [2.1753463745117188, 0.555717945098877, 1.391709804534912], 'max_error':
2.6628344654985634}]

Job completed
```

Anomaly detection - SciTinyML

studio.edgeimpulse.com/studio/51963/learning/anomaly/52

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Spectral features
  - NN Classifier
  - Anomaly detection
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

ANOMALY DETECTION (SCITINYML-MOTION-ANOMALY-PROJECT)

MJRoBot (Marcelo Rovai)

#1 Click to set a description for this version

Anomaly detection settings

Cluster count

32

Axes

Select all axes

- [x] accX RMS
- [ ] accX Peak 1 Freq
- [ ] accX Peak 1 Height
- [ ] accX Peak 2 Freq
- [ ] accX Peak 2 Height
- [ ] accX Peak 3 Freq
- [ ] accX Peak 3 Height
- [ ] accX Spectral Power 0.1 - 0.5
- [ ] accX Spectral Power 0.5 - 1.0
- [ ] accX Spectral Power 1.0 - 2.0
- [ ] accX Spectral Power 2.0 - 5.0
- [x] accY RMS
- [ ] accY Peak 1 Freq
- [ ] accY Peak 1 Height
- [ ] accY Peak 2 Freq
- [ ] accY Peak 2 Height
- [ ] accY Peak 3 Freq
- [ ] accY Peak 3 Height
- [ ] accY Spectral Power 0.1 - 0.5
- [ ] accY Spectral Power 0.5 - 1.0
- [ ] accY Spectral Power 1.0 - 2.0
- [ ] accY Spectral Power 2.0 - 5.0
- [x] accZ RMS
- [ ] accZ Peak 1 Freq
- [ ] accZ Peak 1 Height
- [ ] accZ Peak 2 Freq
- [ ] accZ Peak 2 Height
- [ ] accZ Peak 3 Freq
- [ ] accZ Peak 3 Height
- [ ] accZ Spectral Power 0.1 - 0.5
- [ ] accZ Spectral Power 0.5 - 1.0
- [ ] accZ Spectral Power 1.0 - 2.0
- [ ] accZ Spectral Power 2.0 - 5.0

Start training

Anomaly explorer (3,230 samples)

X Axis: accX RMS | Y Axis: accY RMS | Test data: -- No test data

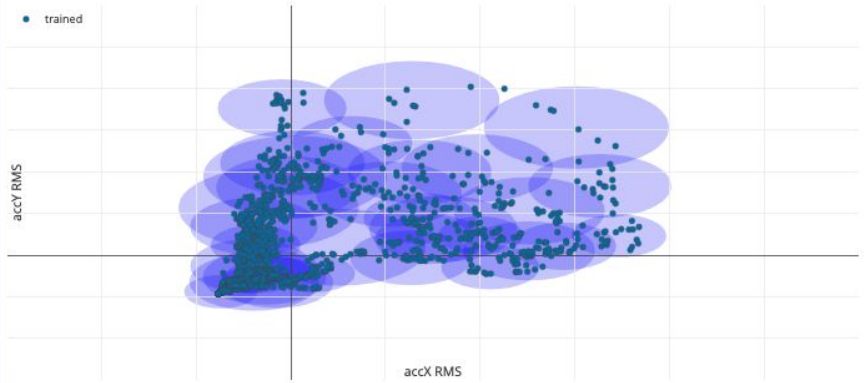

Training output

```
0.6358923488456604}, {'center': [0.2331821322441101, -0.44085508584976196, 1.193619966506958], 'max_error':
0.43379442394199896}, {'center': [-0.04211855307221413, 1.622160792350769, 1.5633275508880615], 'max_error':
0.7400662371471811}, {'center': [2.5153045654296875, 0.10167547315359116, 1.1958473920822144], 'max_error':
0.6889784598589724}, {'center': [0.6476534008979797, 2.6941537857055664, 1.7468148469924927], 'max_error':
0.6253040107657685}, {'center': [-0.09443876147270203, 3.532026529312134, 1.4222177267074585], 'max_error':
0.6722667084735653}, {'center': [2.1073172092437744, -0.2908220887184143, 0.9732017517089844], 'max_error':
0.5098888571637854}, {'center': [1.360011339187622, 2.0014262199401855, 0.6398685574531555], 'max_error':
0.750148069600636}, {'center': [1.3629376888275146, 0.8849844336509705, 1.6627943515777588], 'max_error':
0.5872878607683915}, {'center': [1.6086387634277344, 0.635323166847229, 0.2745974659919739], 'max_error':
0.7719379326342638}, {'center': [2.8824524879455566, 0.21148265898227692, 1.768933892250061], 'max_error':
0.5465850786758827}, {'center': [3.2645251750946045, 1.6626498699188232, 1.3556557893753052], 'max_error':
0.7446547869773675}, {'center': [1.2767690420150757, 3.724461078643799, 1.341691255569458], 'max_error':
0.9180358736727107}, {'center': [3.0170180797576904, 3.0672569274902344, 0.8234216570854187], 'max_error':
0.9686505165548877}]

Job completed
```

Spectral features (3,327 samples)

X Axis: accX RMS

Y Axis: accY RMS

Z Axis: accZ RMS

- classified
- idle
- lift
- maritime
- terrestrial
- classification 49

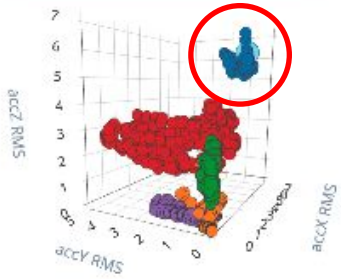

Processed features

6.8037, 0.4960, 8.2240, 1.4881, 3.2398, 2.9762, 0.4911, 7.0271, 0.9133, 0.6076, 0.0159, 1.0079, 0.4960, 1.0556, 2.4802, 0.1…

Anomaly explorer (3,327 samples)

X Axis: accX RMS

Y Axis: accY RMS

- trained
- classified
- classification 49

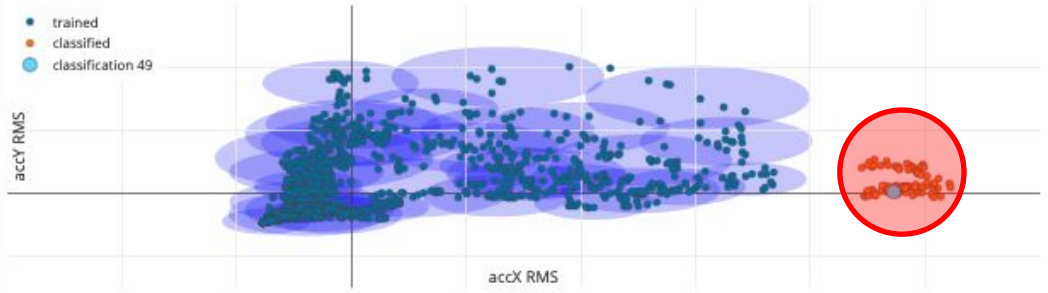

Distance from closest cluster

accX RMS: 1.2592, accY RMS: 0.4069, accZ RMS: 2.8992

EON Tuner

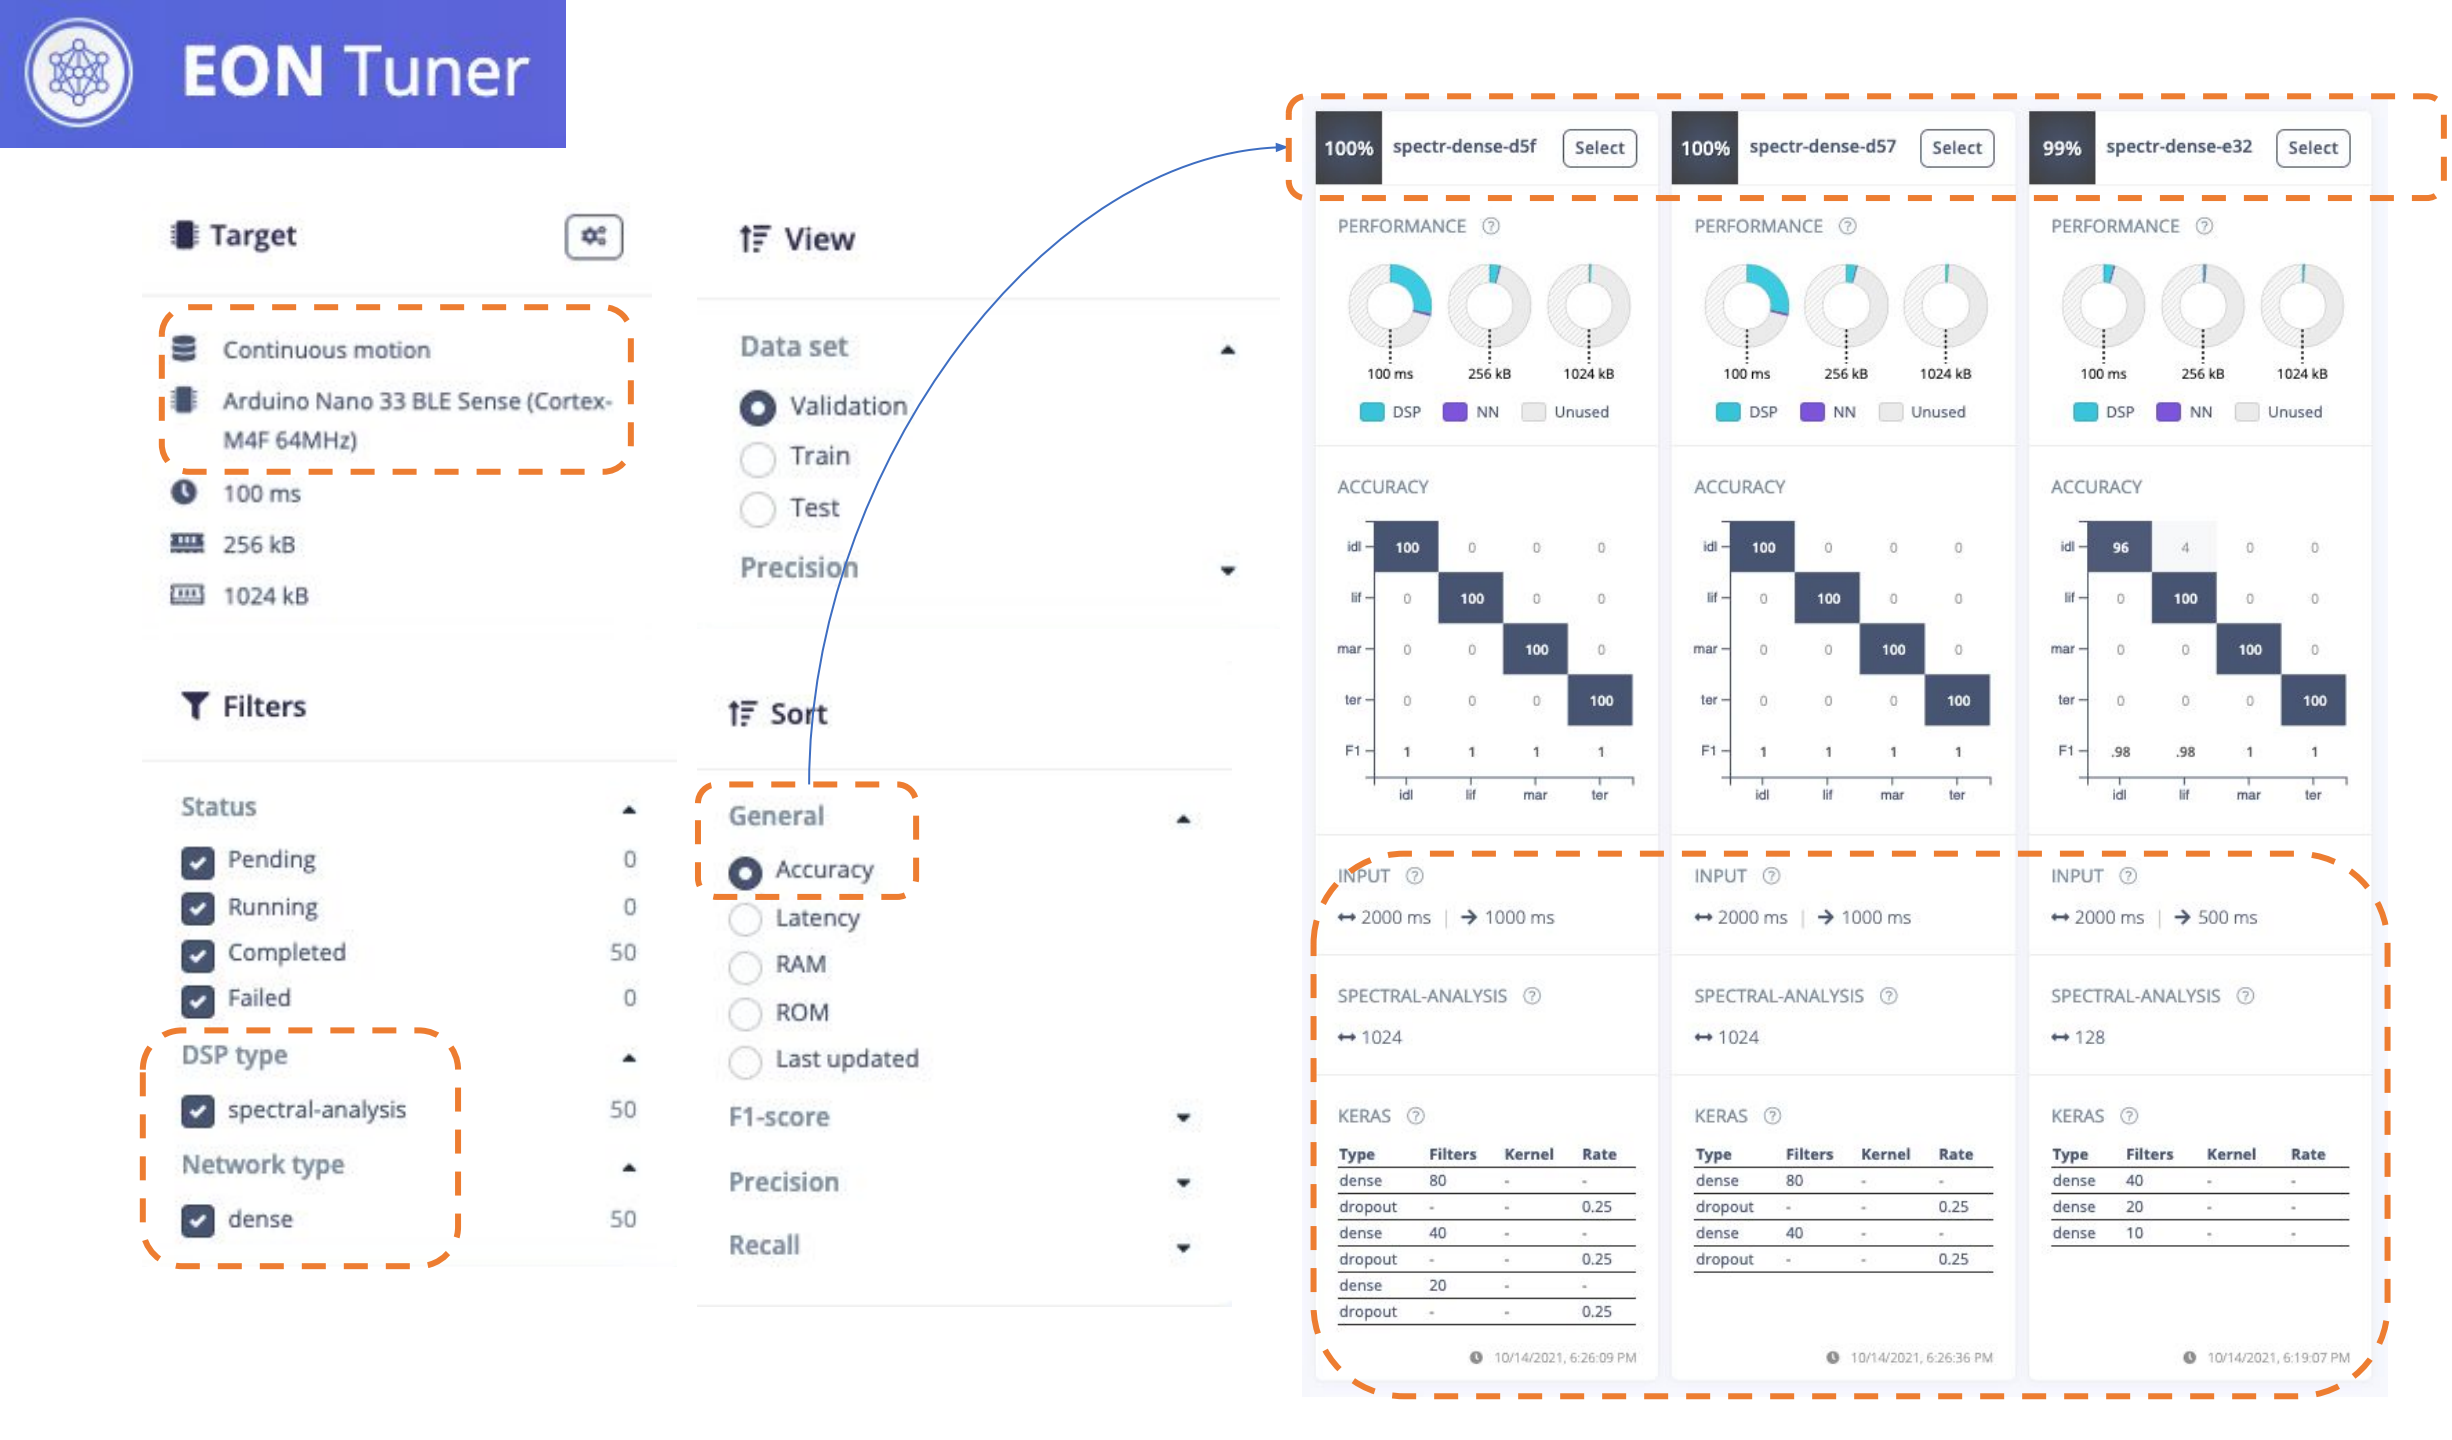
Target
Continuous motion
Arduino Nano 33 BLE Sense (Cortex-M4F 64MHz)
100 ms
256 kB
1024 kB
Filters
Status
Pending 0
Running 0
Completed 50
Failed 0
DSP type
spectral-analysis 50
Network type
dense 50
View
Data set
Validation
Train
Test
Precision
Sort
General
Accuracy
Latency
RAM
ROM
Last updated
F1-score
Precision
Recall
100% spectr-dense-d5f Select
100% spectr-dense-d57 Select
99% spectr-dense-e32 Select
PERFORMANCE
100 ms
256 kB
1024 kB
DSP
NN
Unused
ACCURACY
idl
lif
mar
ter
F1
INPUT
2000 ms | 1000 ms
2000 ms | 1000 ms
2000 ms | 500 ms
SPECTRAL-ANALYSIS
1024
1024
128
KERAS
Type Filters Kernel Rate
dense 80 - -
dropout - - 0.25
dense 40 - -
dropout - - 0.25
dense 20 - -
dropout - - 0.25
10/14/2021, 6:26:09 PM
Type Filters Kernel Rate
dense 80 - -
dropout - - 0.25
dense 40 - -
dropout - - 0.25
10/14/2021, 6:26:36 PM
Type Filters Kernel Rate
dense 40 - -
dense 20 - -
dense 10 - -
10/14/2021, 6:19:07 PM

Possible Scientific uses of TinyML on Real Life

Cow Monitoring

Using the Internet of Things for Agricultural Monitoring

"We aim to deploy a variety of sensors for agricultural monitoring. One of the projects involves using accelerometer sensors to monitor activity levels in dairy cows with a view to determining when the cows are on heat or when they are sick."

Ciira wa Maina, Ph.D.

Senior Lecturer
Department of Electrical and Electronic Engineering
Dedan Kimathi University of Technology
Nyeri Kenya
Email: ciira.maina@dkut.ac.ke

Kenia

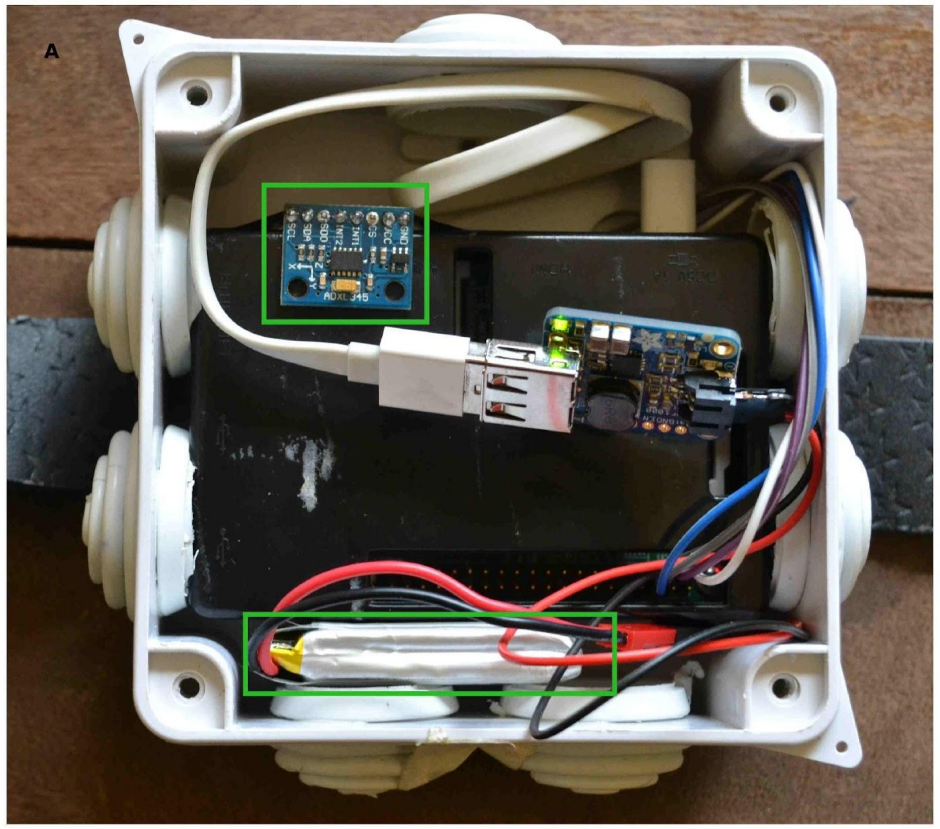

https://sites.google.com/site/cwamainadekut/research

iBean

Detecting Diseases in the Bean plants

AIR Lab Makerere University

UGANDA

Sound

Vibration

Vision

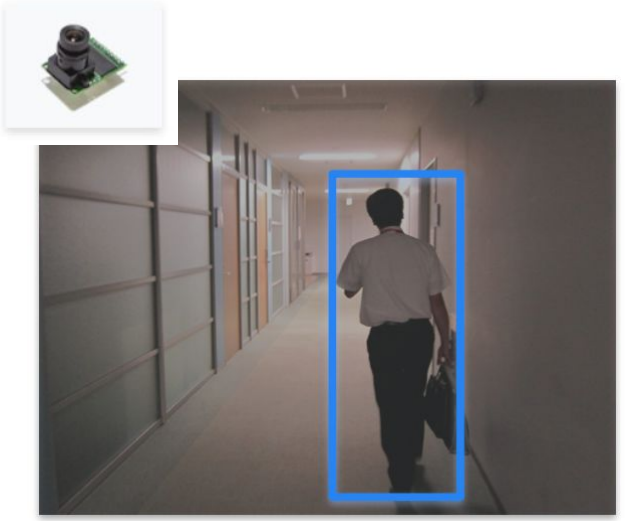

iBean Dataset

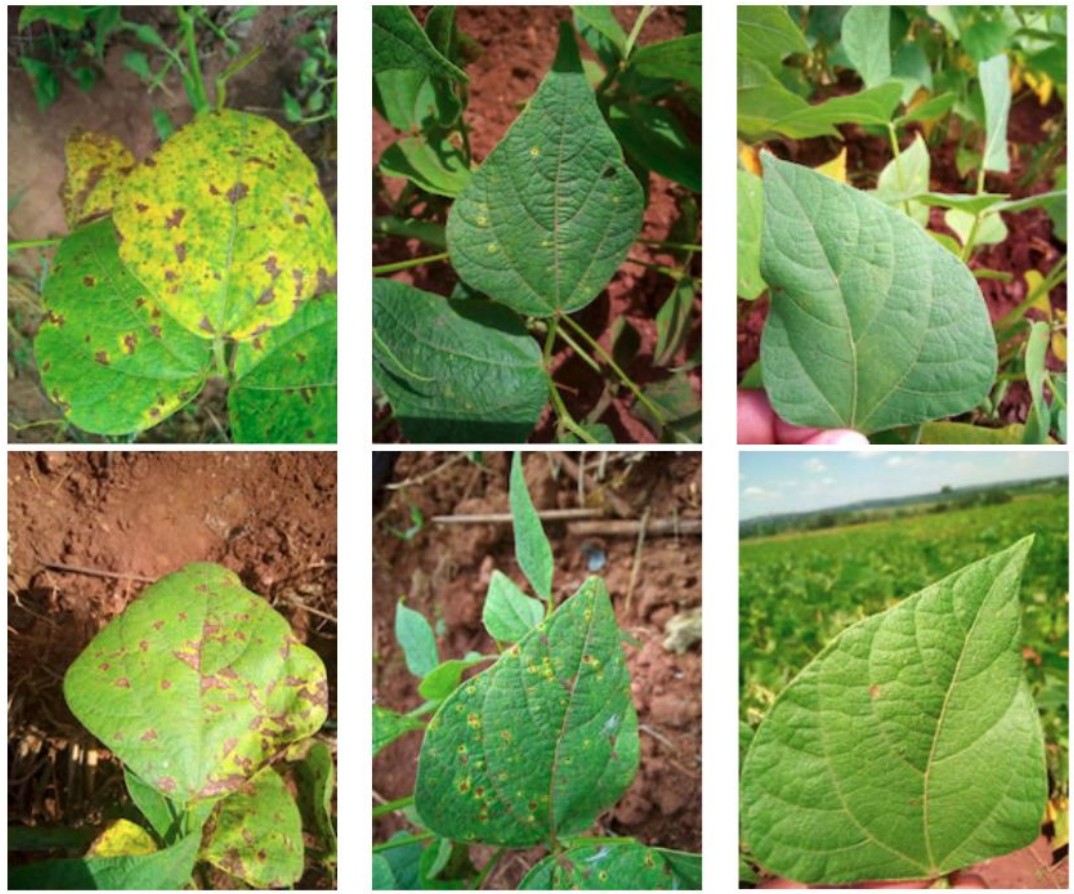

Angular Leaf Spot Bean Rust Healthy

This dataset is of leaf images taken in the field in different districts in Uganda by the Makerere AI lab in collaboration with the National Crops Resources Research Institute (NaCRRI), the national body in charge of research in agriculture in Uganda.

Goal:

To build a neural network that can tell the difference between the healthy and diseased leaves.

Dataset:

Training, Test and Validation data based on 224x224 pixel color images taken of bean plants in Uganda.

Class	Examples
Healthy class	428
Angular Leaf Spot	432
Bean Rust	436
Total:	1,296

Dataset: https://github.com/AI-Lab-Makerere/ibean/

iBean Dataset

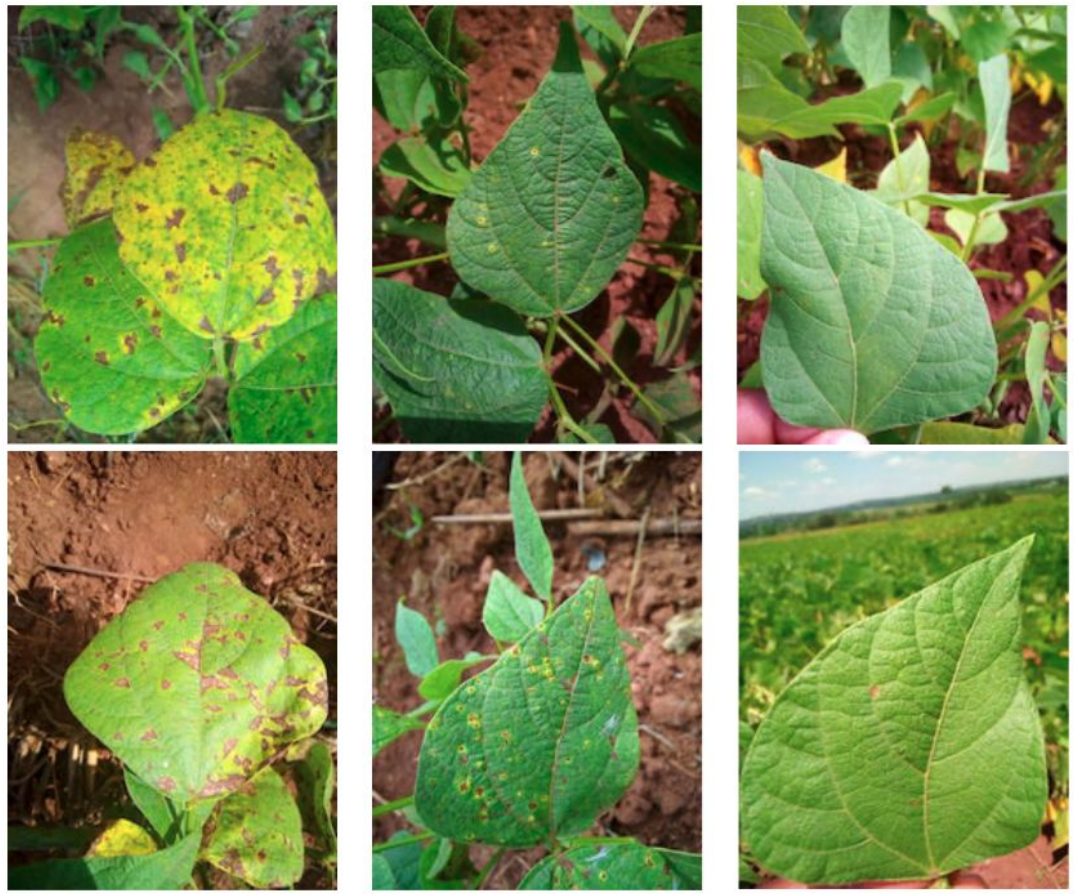

Angular Leaf Spot | Bean Rust | Healthy

This dataset is of leaf images taken in the field in different districts in Uganda by the Makerere AI lab in collaboration with the National Crops Resources Research Institute (NaCRRI), the national body in charge of research in agriculture in Uganda.

Goal:

To build a neural network that can tell the difference between the healthy and diseased leaves.

Dataset:

Training, Test and Validation data based on 224x224 pixel color images taken of bean plants in Uganda.

Class	Examples
Healthy class	428
Angular Leaf Spot	432
Bean Rust	436
Total:	1,296

Problem!

Not a lot of data

Dataset: https://github.com/AI-Lab-Makerere/ibean/

But what to do if we do not have more data?

- Data Augmentation (artificial)
- Transfer Learning

Data Augmentation (artificial)

Data augmentation takes the approach of generating additional training data from your existing examples by augmenting them using random transformations that yield believable-looking images. This helps expose the model to more aspects of the data and generalize better (avoiding overfitting).

Using tf.image

```
1 flipped = tf.image.flip_left_right(image)
2 visualize(image, flipped)
```

Original image

Augmented image

```
1 rotated = tf.image.rot90(image)
2 visualize(image, rotated)
```

Original image

Augmented image

Using tf.image

```
saturated = tf.image.adjust_saturation(image, 3)
visualize(image, saturated)
```

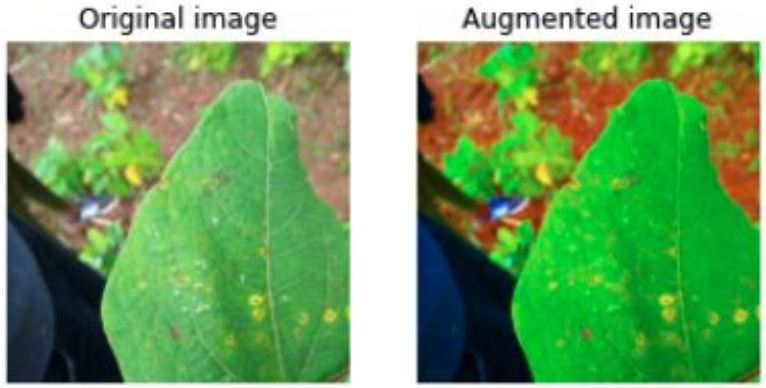


```
bright = tf.image.adjust_brightness(image, 0.4)
visualize(image, bright)
```

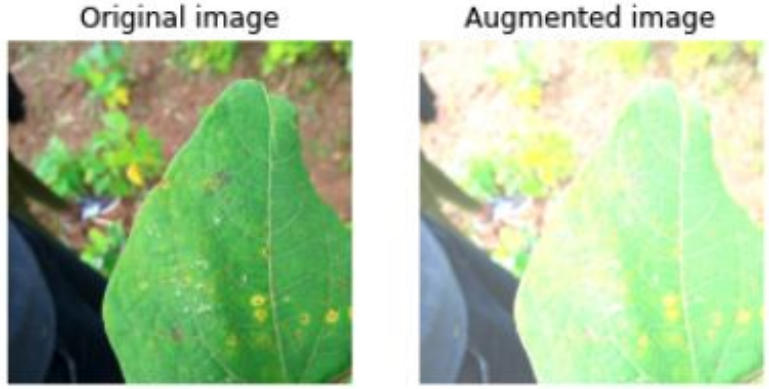


```
for i in range(3):
  seed = (i, 0)  # tuple of size (2,)
  stateless_random_crop = tf.image.stateless_random_crop(
      image, size=[210, 300, 3], seed=seed)
  visualize(image, stateless_random_crop)
```

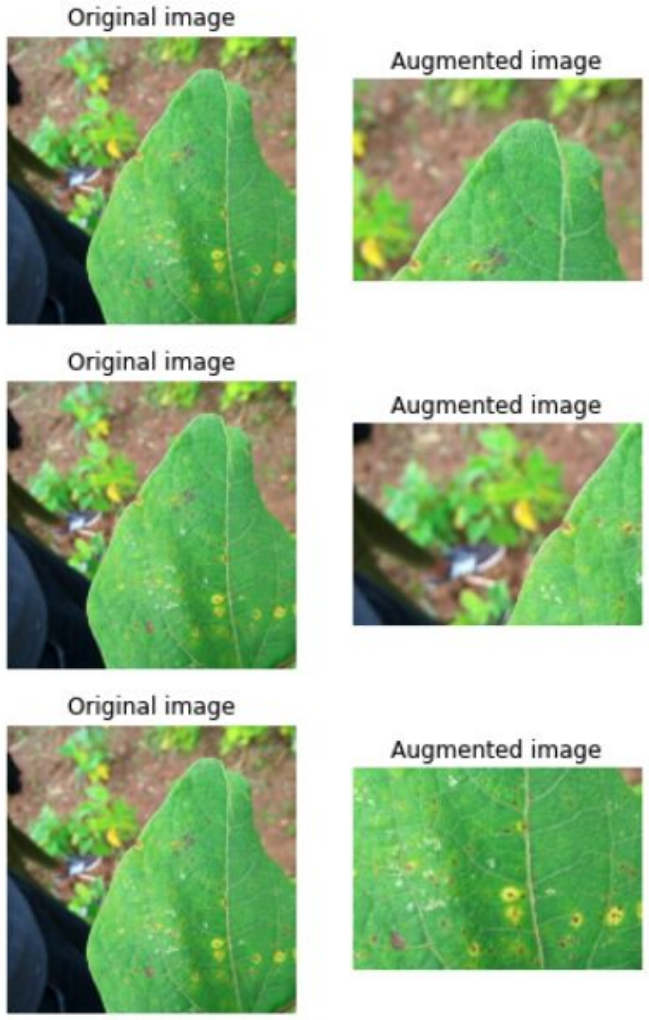

Transfer Learning

End Result of Training

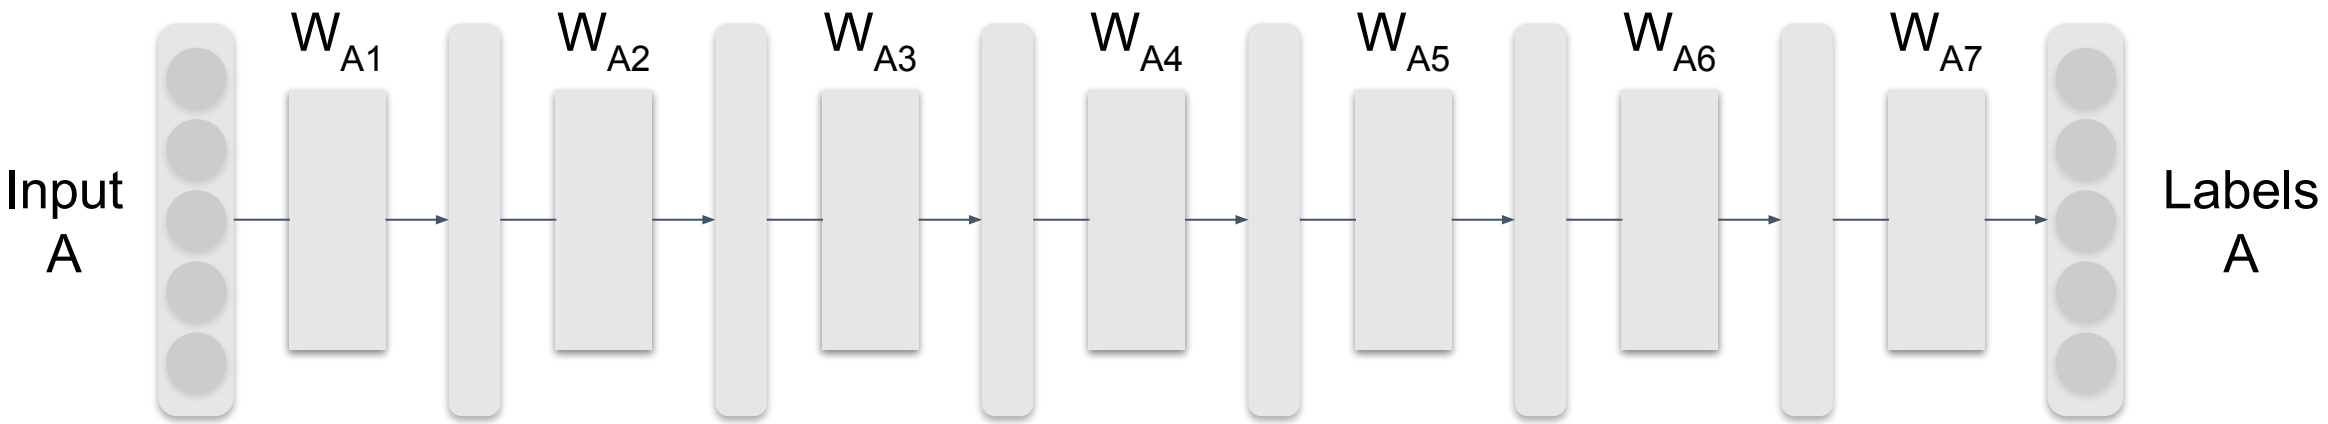

The end result of the training is to learn the weights of the neural network model.

End Result of Training

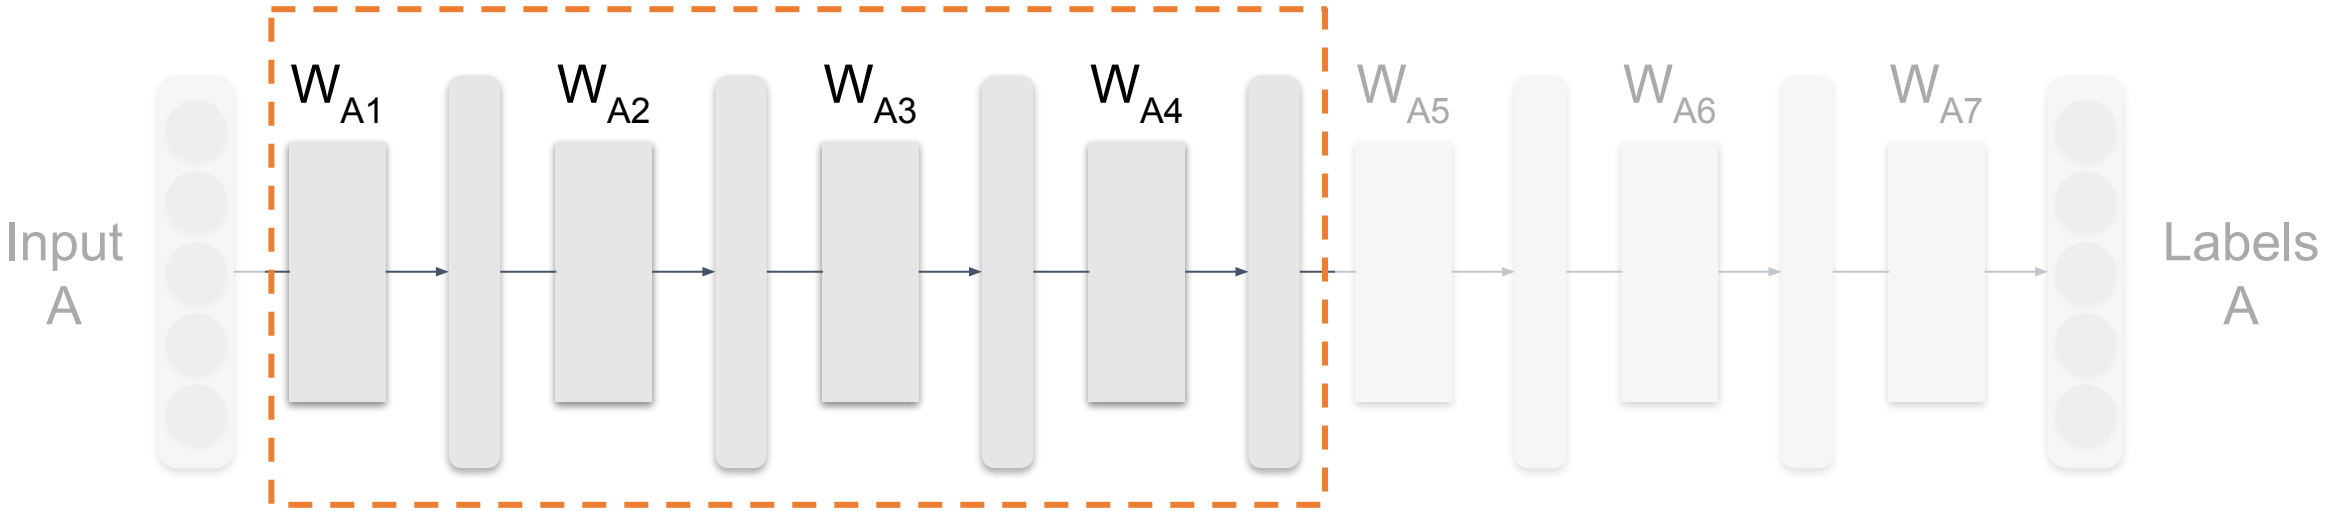

Learns ***general features*** irrespective of task

End Result of Training

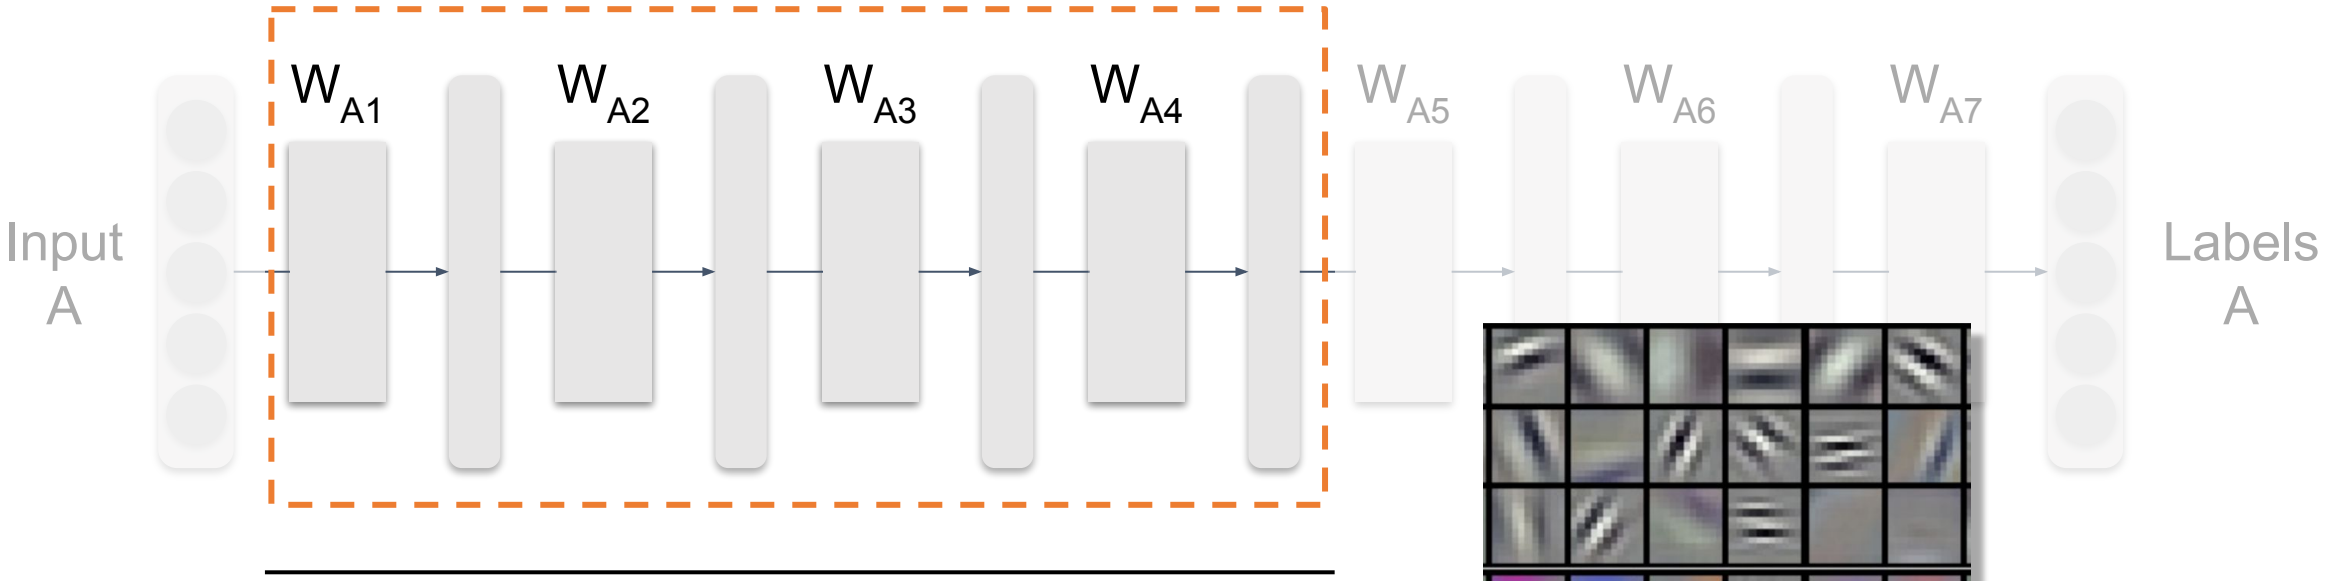

Learns ***general features*** irrespective of task

End Result of Training

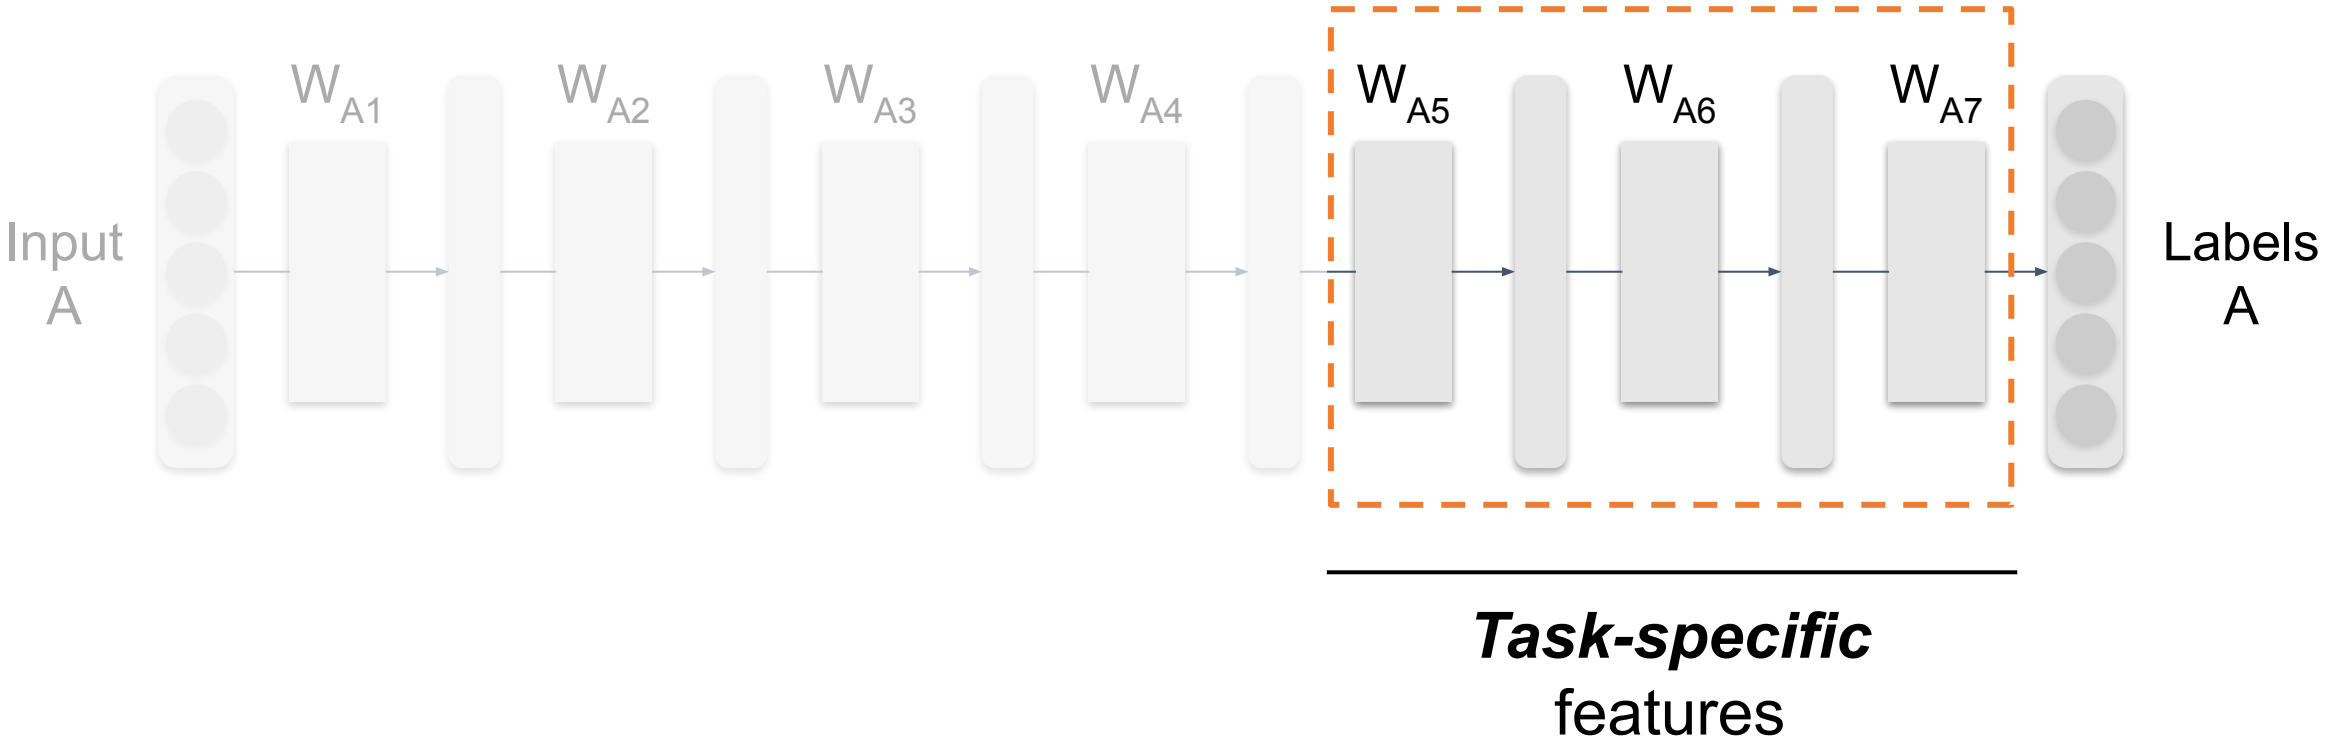


Source: Google

Transfer Learning

Reuse (freeze general feature extraction)

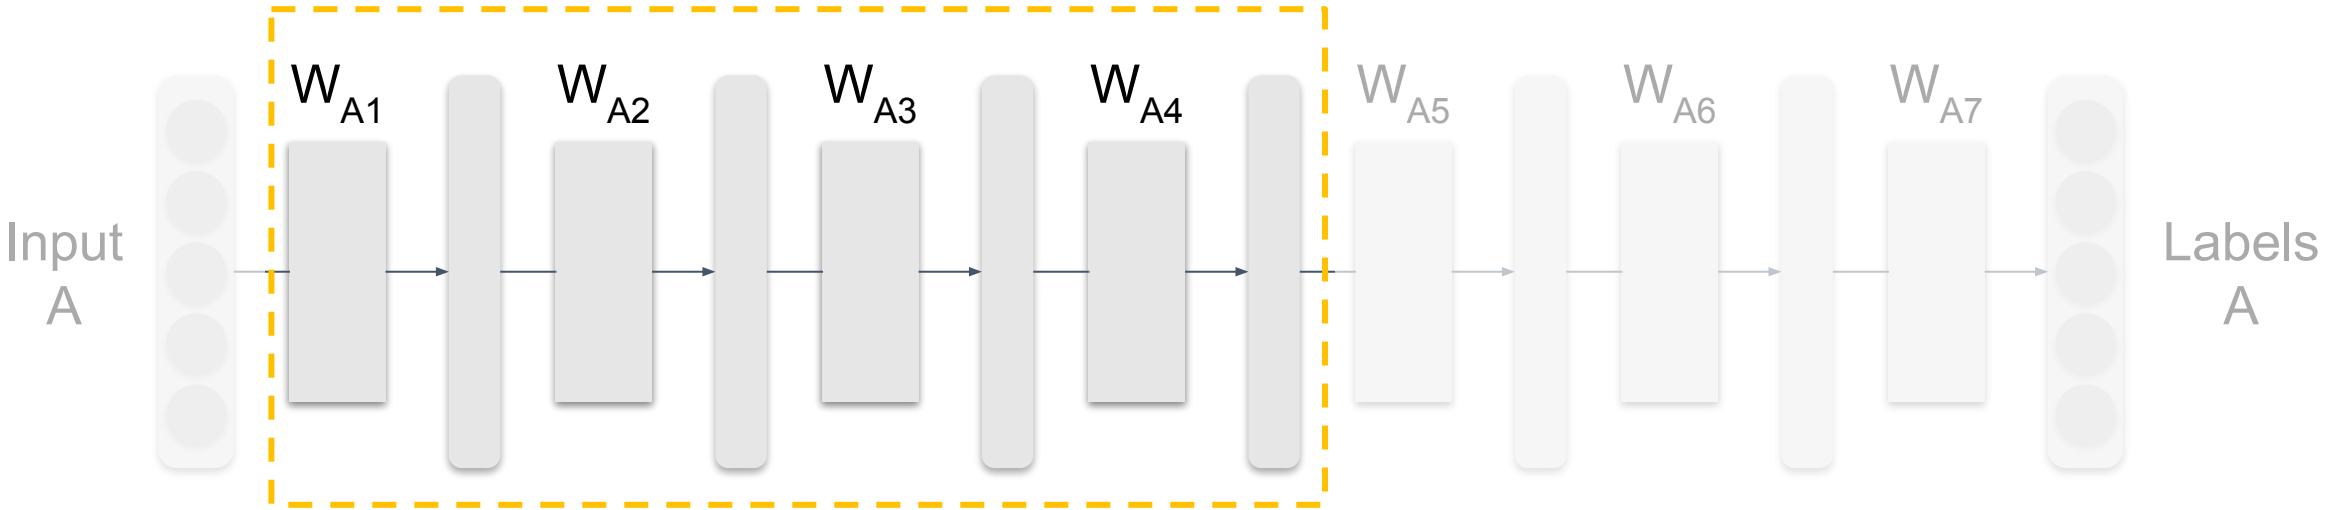

Learns ***general features*** irrespective of task

Transfer Learning

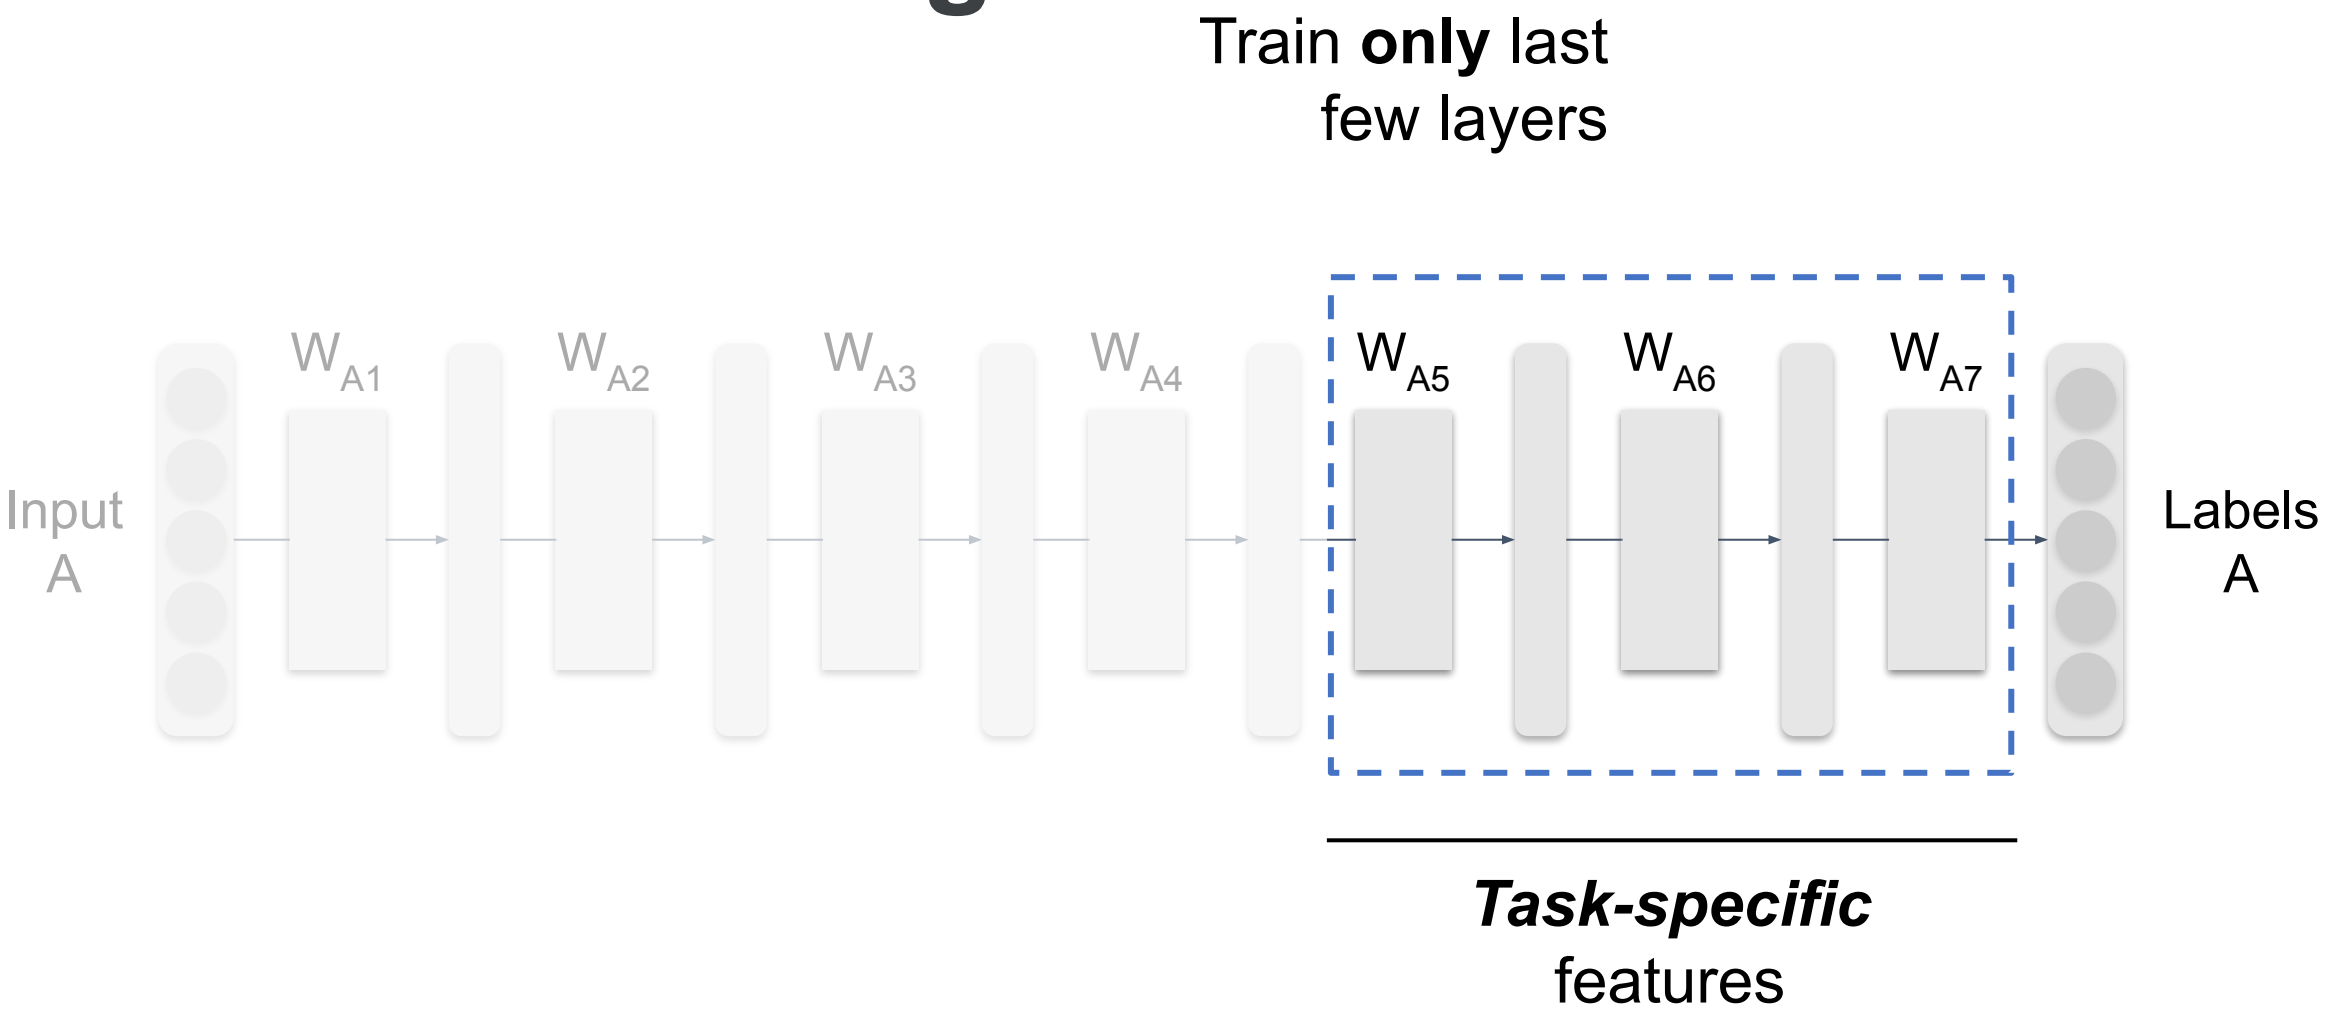

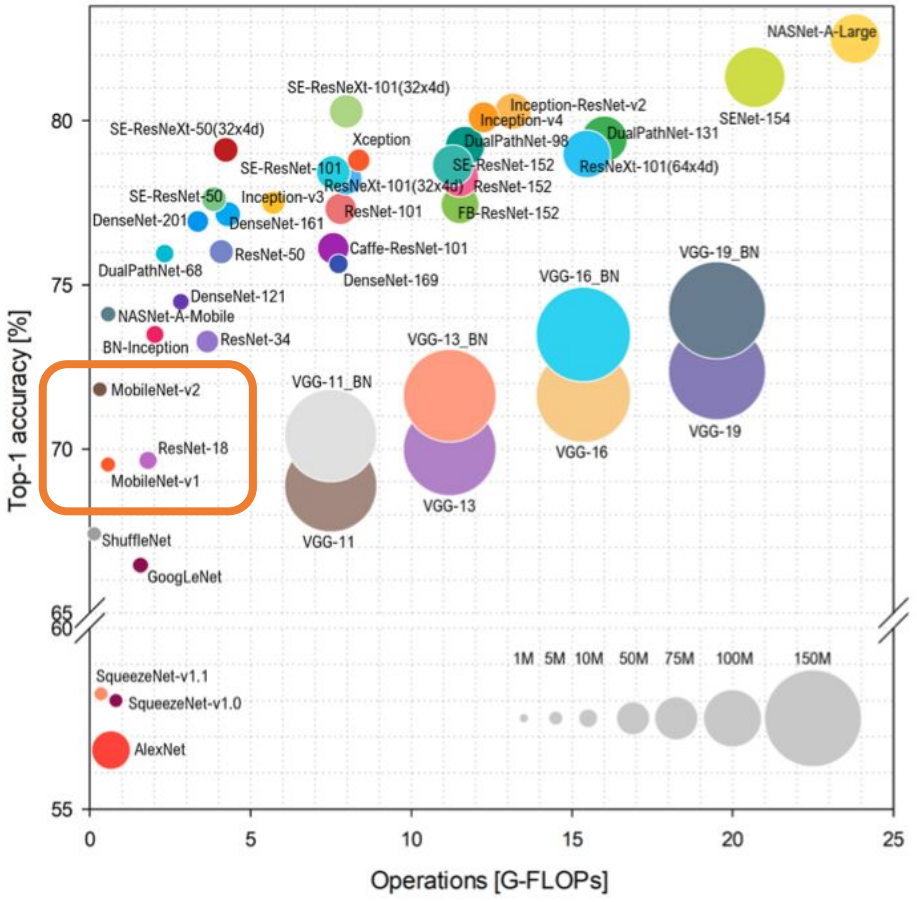

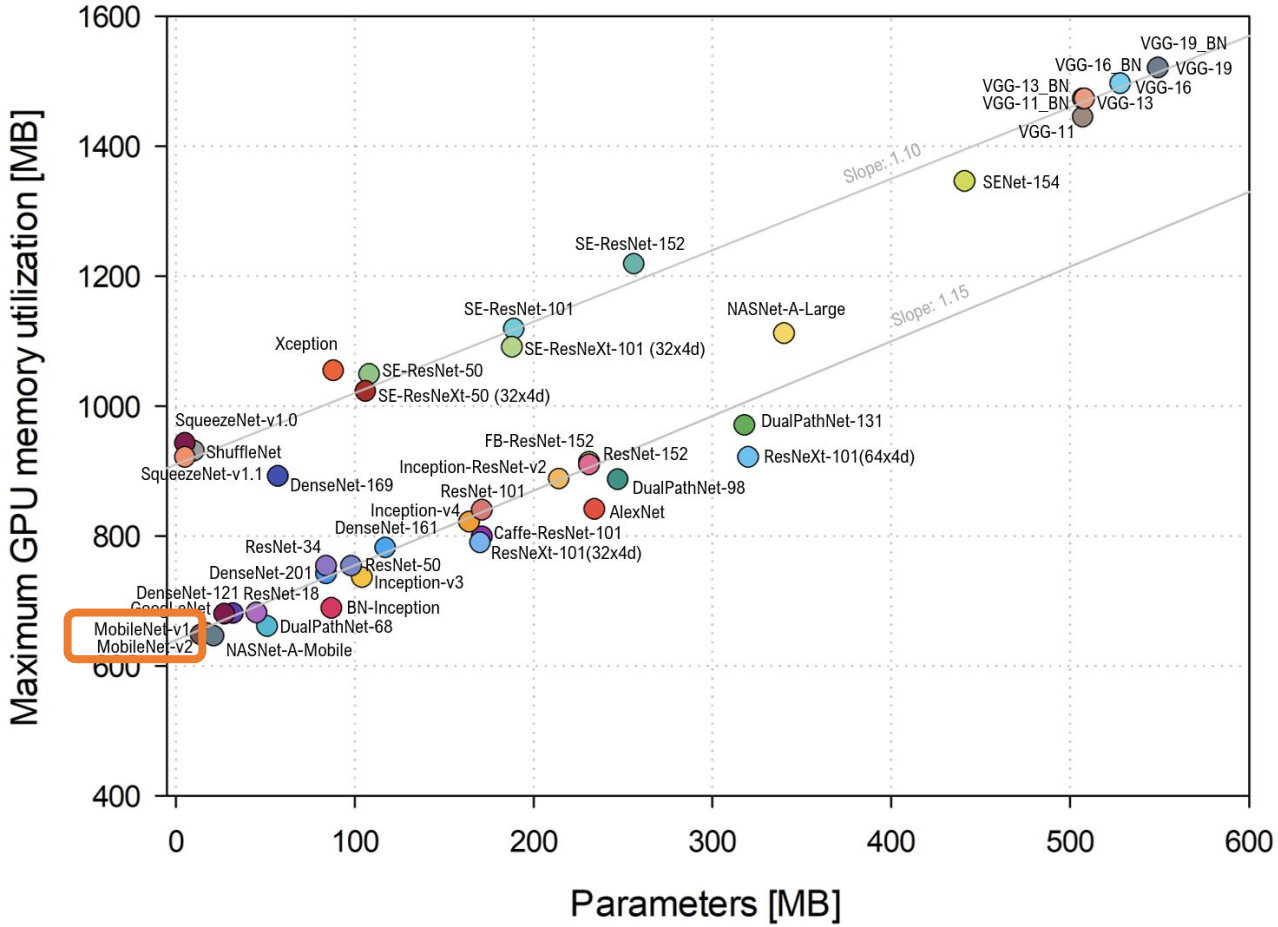

https://arxiv.org/pdf/1810.00736.pdf

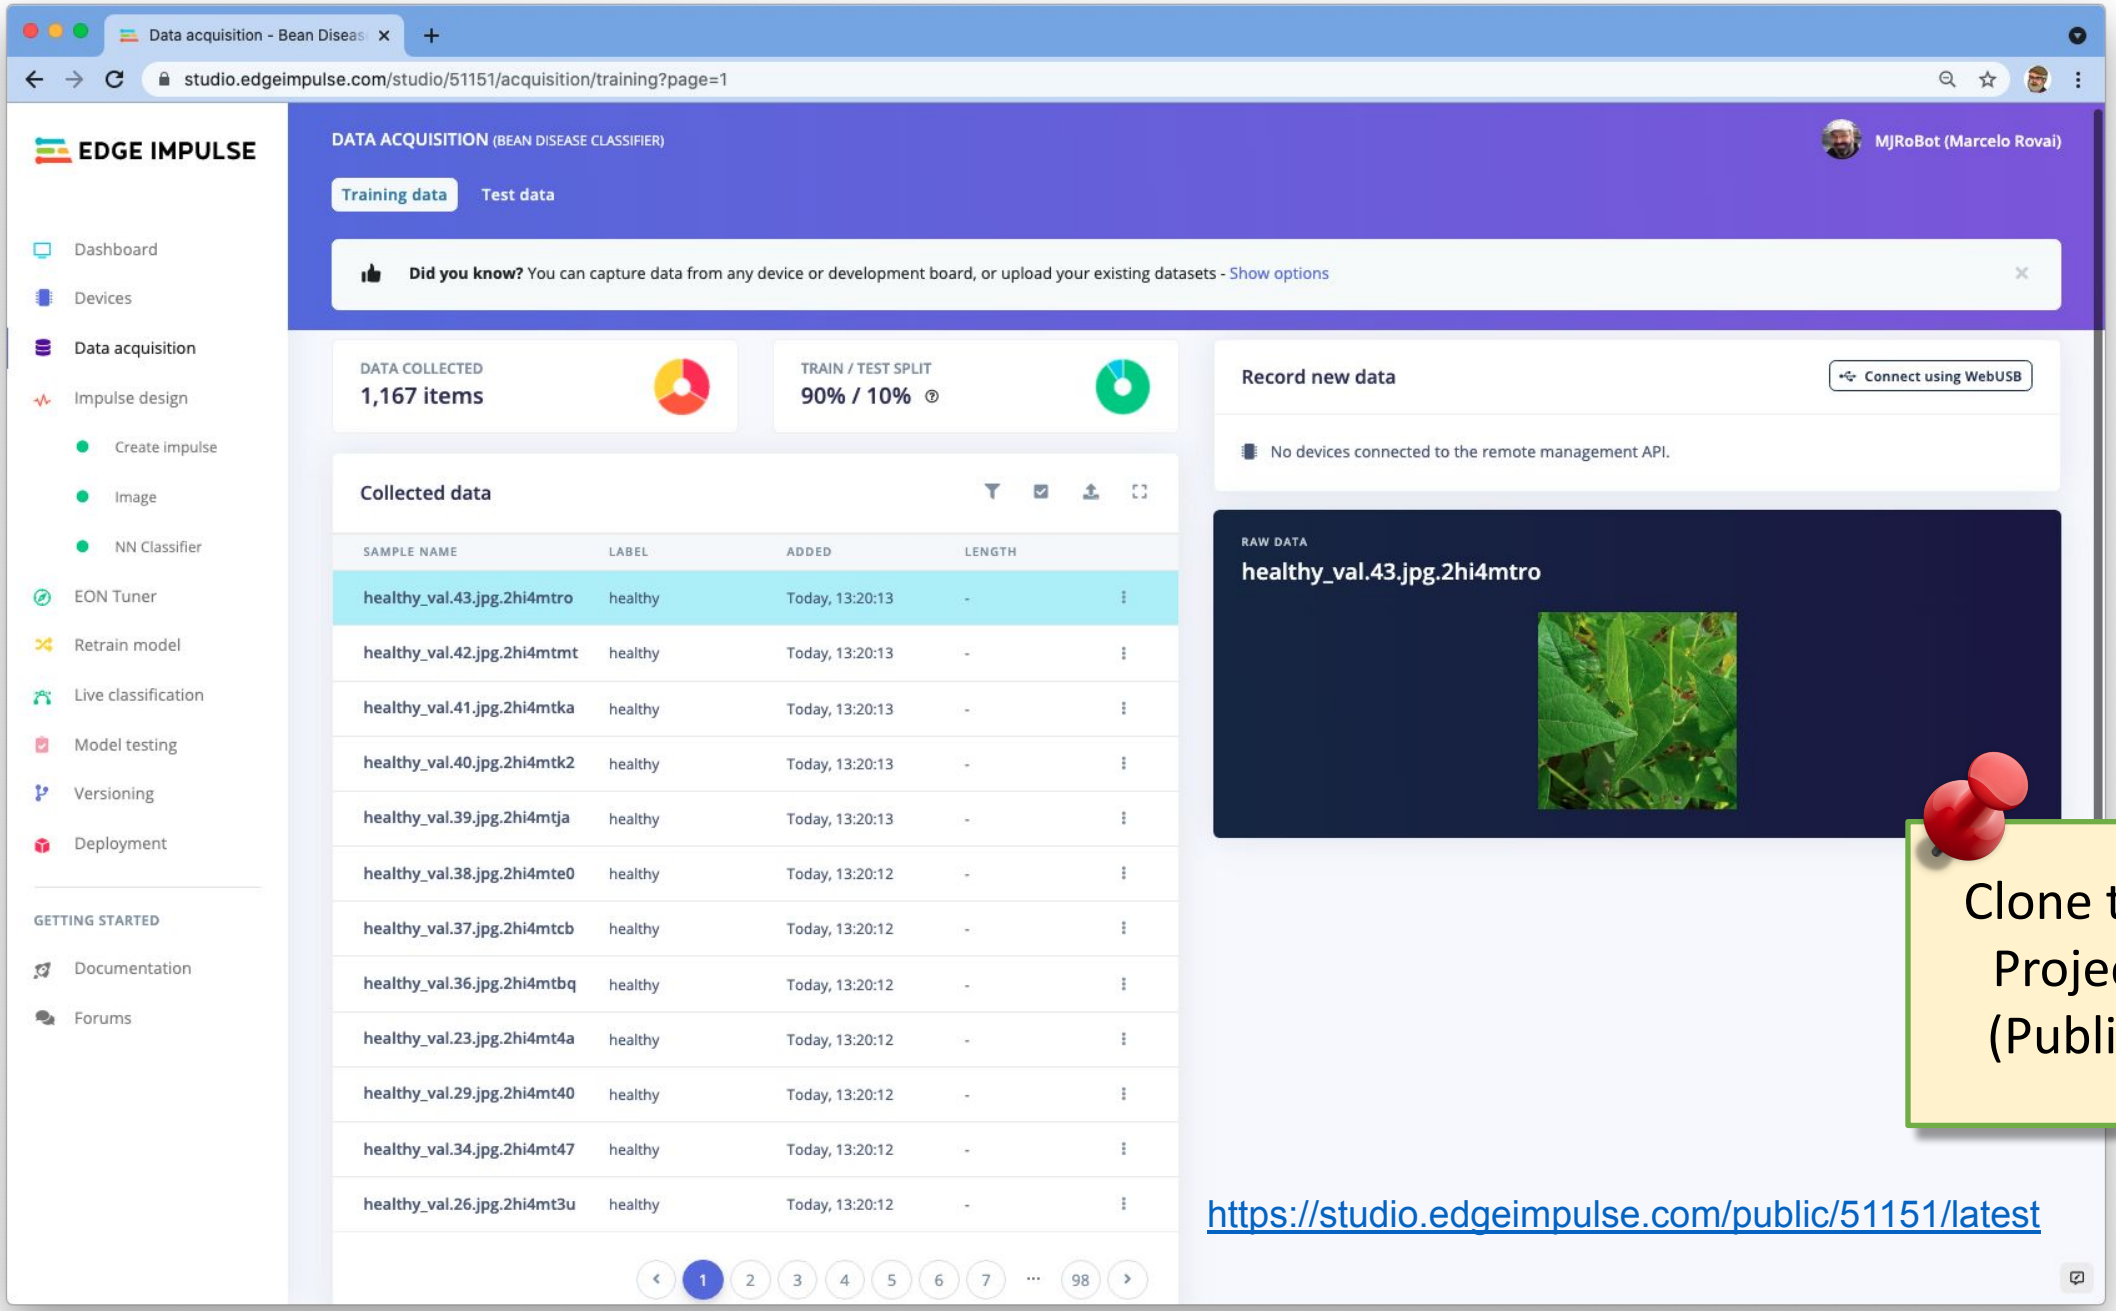
Data acquisition - Bean Diseas
studio.edgeimpulse.com/studio/51151/acquisition/training?page=1
EDGE IMPULSE
DATA ACQUISITION (BEAN DISEASE CLASSIFIER)
MJRoBot (Marcelo Rovai)
Training data
Test data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
NN Classifier
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
DATA COLLECTED
1,167 items
TRAIN / TEST SPLIT
90% / 10%
Record new data
Connect using WebUSB
No devices connected to the remote management API.
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
healthy_val.43.jpg.2hi4mtro healthy Today, 13:20:13 -
healthy_val.42.jpg.2hi4mtmt healthy Today, 13:20:13 -
healthy_val.41.jpg.2hi4mtka healthy Today, 13:20:13 -
healthy_val.40.jpg.2hi4mtk2 healthy Today, 13:20:13 -
healthy_val.39.jpg.2hi4mtja healthy Today, 13:20:13 -
healthy_val.38.jpg.2hi4mte0 healthy Today, 13:20:12 -
healthy_val.37.jpg.2hi4mtcb healthy Today, 13:20:12 -
healthy_val.36.jpg.2hi4mtbq healthy Today, 13:20:12 -
healthy_val.23.jpg.2hi4mt4a healthy Today, 13:20:12 -
healthy_val.29.jpg.2hi4mt40 healthy Today, 13:20:12 -
healthy_val.34.jpg.2hi4mt47 healthy Today, 13:20:12 -
healthy_val.26.jpg.2hi4mt3u healthy Today, 13:20:12 -
1 2 3 4 5 6 7 ... 98
RAW DATA
healthy_val.43.jpg.2hi4mtro
Clone the Project (Public)
https://studio.edgeimpulse.com/public/51151/latest

EON Tuner

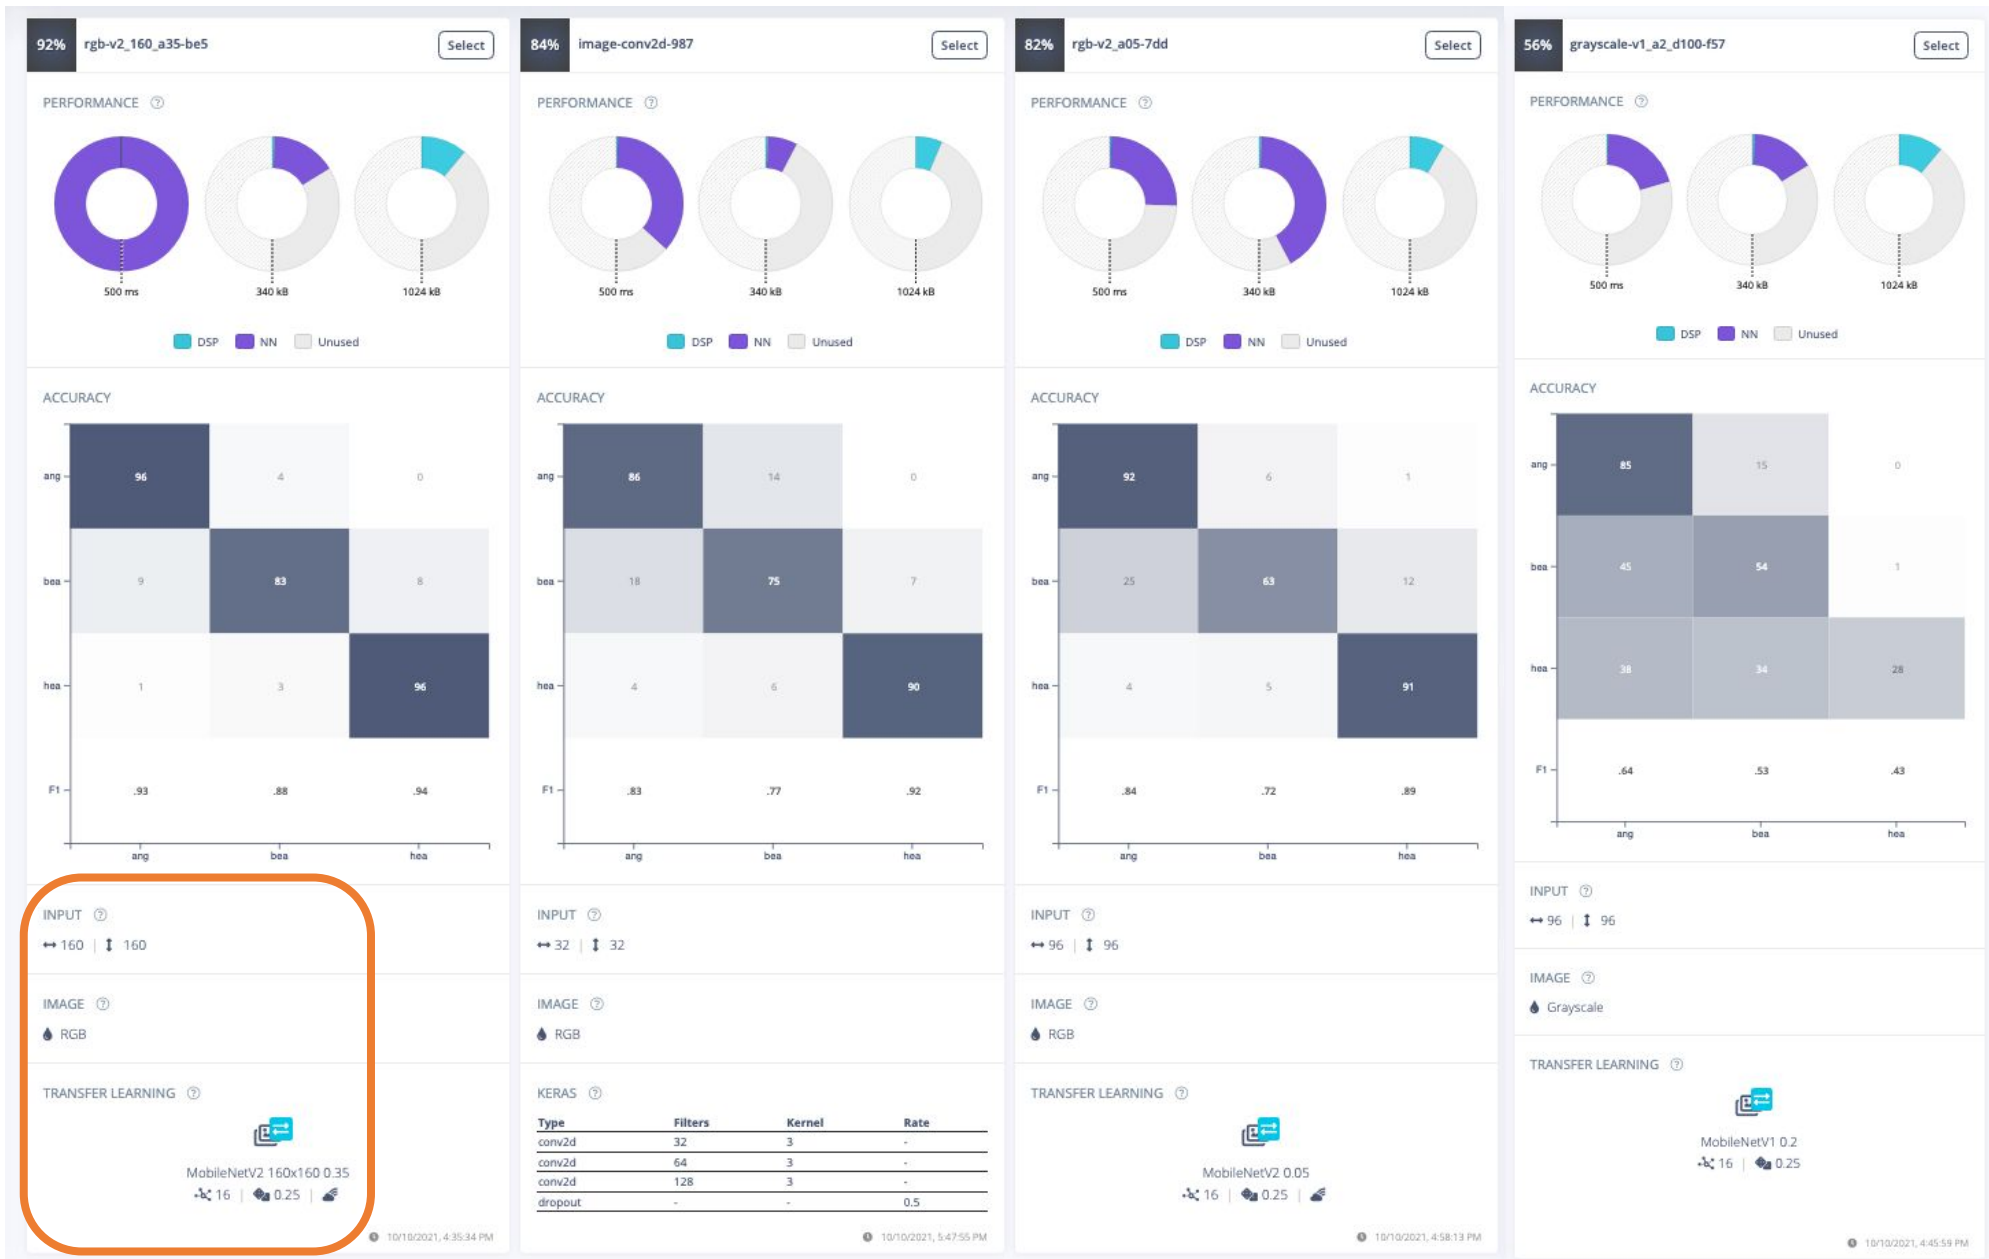
92% rgb-v2_160_a35-be5 Select
PERFORMANCE
500 ms 340 kB 1024 kB
DSP NN Unused
ACCURACY
ang bea hea F1
96 4 0
9 83 8
1 3 96
.93 .88 .94
INPUT
160 | 160
IMAGE
RGB
TRANSFER LEARNING
MobileNetV2 160x160 0.35
16 | 0.25
10/10/2021, 4:35:34 PM
84% image-conv2d-987 Select
PERFORMANCE
500 ms 340 kB 1024 kB
DSP NN Unused
ACCURACY
ang bea hea F1
86 14 0
18 75 7
4 6 90
.83 .77 .92
INPUT
32 | 32
IMAGE
RGB
KERAS
Type Filters Kernel Rate
conv2d 32 3 -
conv2d 64 3 -
conv2d 128 3 -
dropout - - 0.5
10/10/2021, 5:47:55 PM
82% rgb-v2_a05-7dd Select
PERFORMANCE
500 ms 340 kB 1024 kB
DSP NN Unused
ACCURACY
ang bea hea F1
92 6 1
25 63 12
4 5 91
.84 .72 .89
INPUT
96 | 96
IMAGE
RGB
TRANSFER LEARNING
MobileNetV2 0.05
16 | 0.25
10/10/2021, 4:58:13 PM
56% grayscale-v1_a2_d100-f57 Select
PERFORMANCE
500 ms 340 kB 1024 kB
DSP NN Unused
ACCURACY
ang bea hea F1
85 15 0
45 54 1
38 34 28
.64 .53 .43
INPUT
96 | 96
IMAGE
Grayscale
TRANSFER LEARNING
MobileNetV1 0.2
16 | 0.25
10/10/2021, 4:45:59 PM

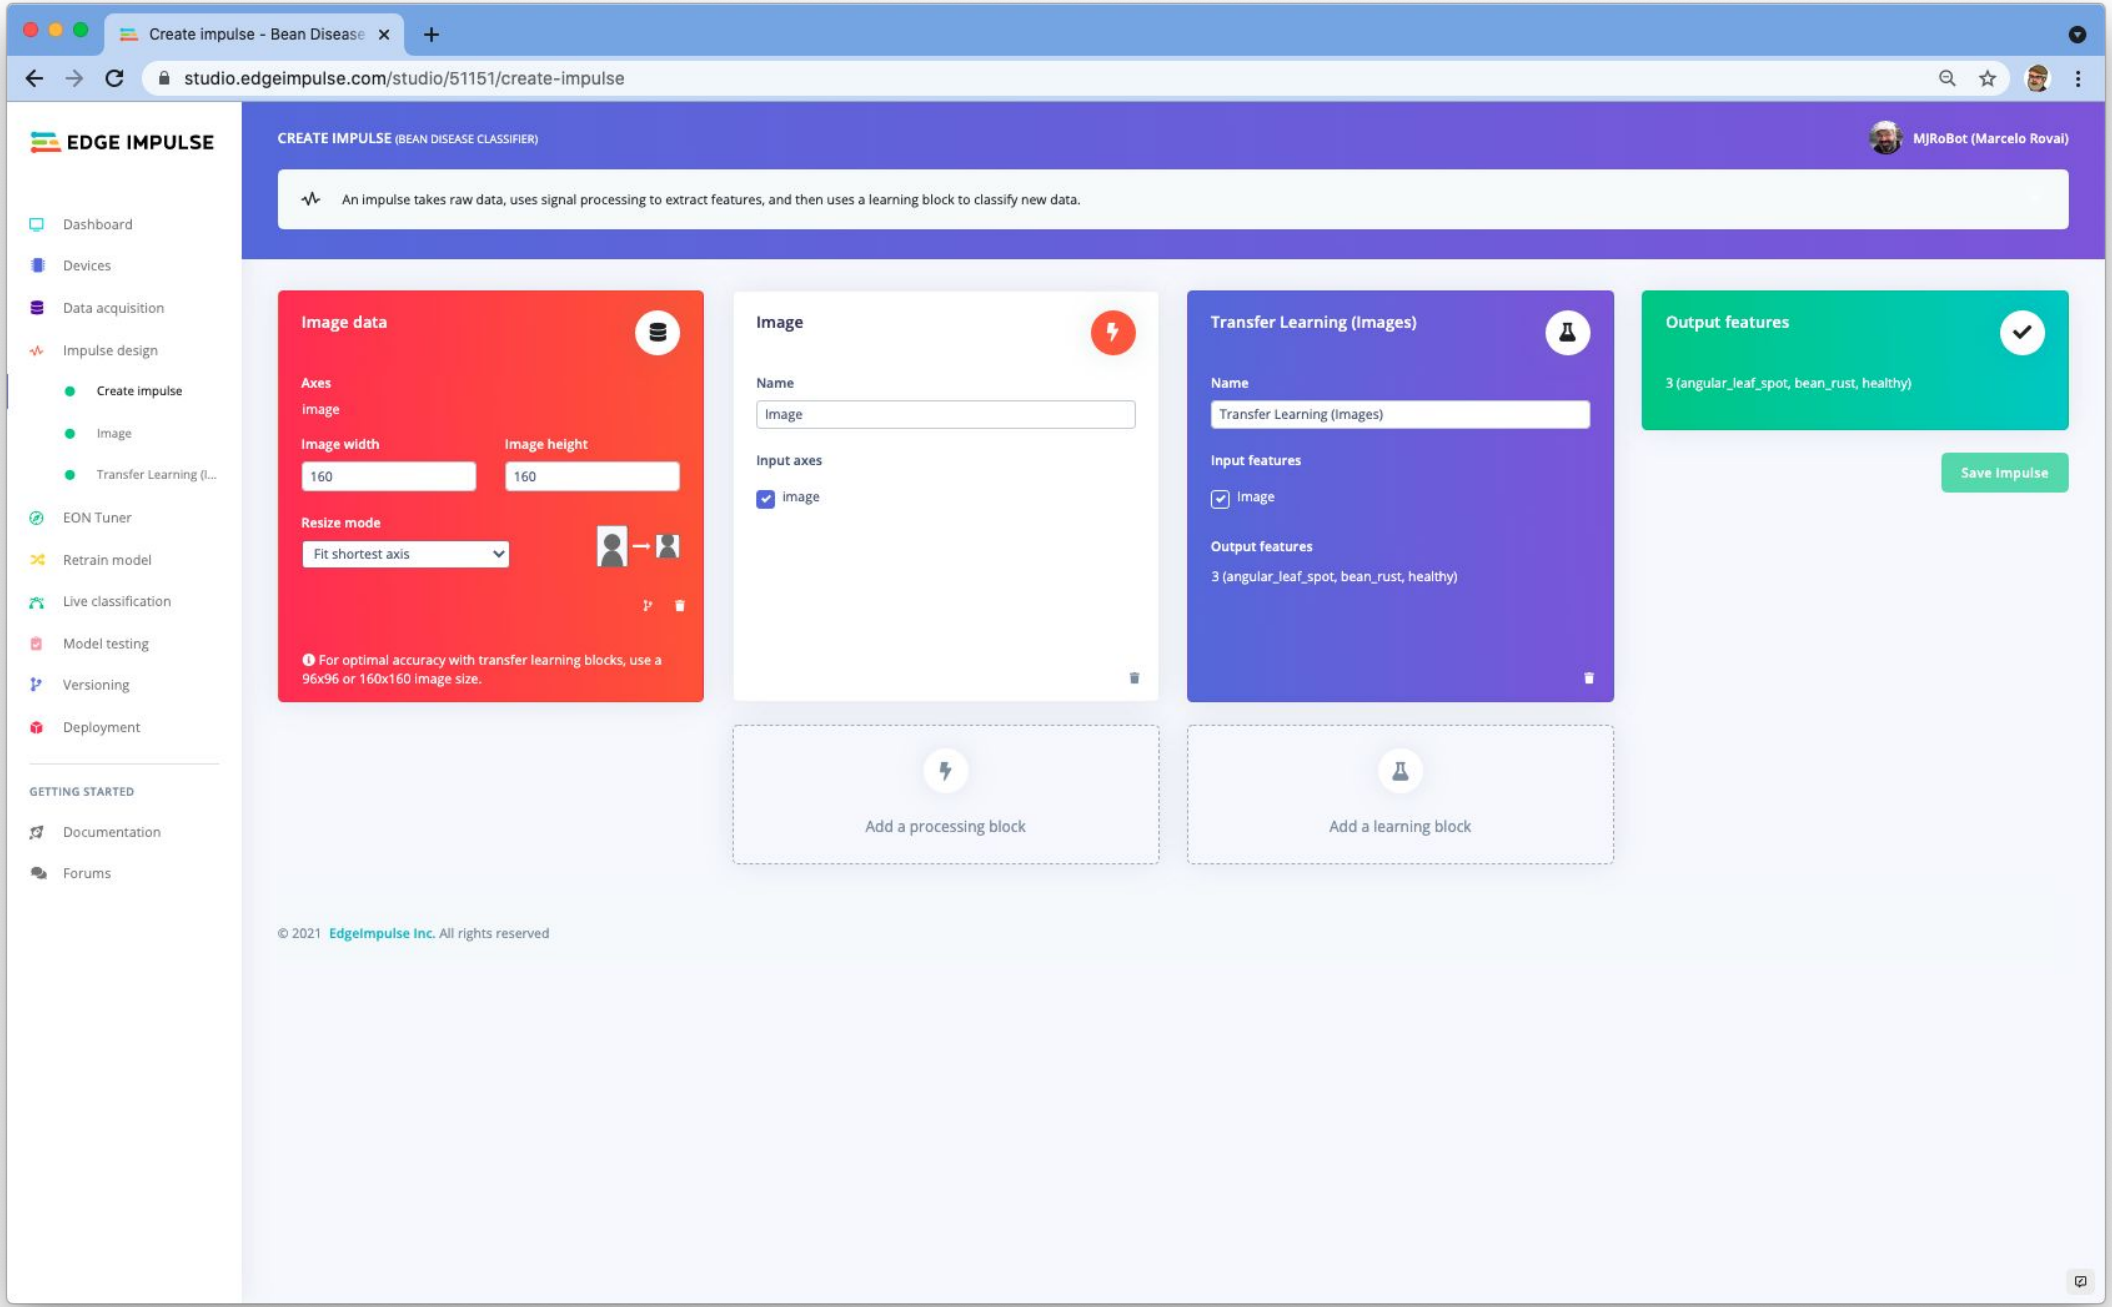
Create impulse - Bean Disease
studio.edgeimpulse.com/studio/51151/create-impulse
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
Transfer Learning (I...
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
CREATE IMPULSE (BEAN DISEASE CLASSIFIER)
MJRoBot (Marcelo Rovai)
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Image data
Axes
image
Image width
160
Image height
160
Resize mode
Fit shortest axis
For optimal accuracy with transfer learning blocks, use a 96x96 or 160x160 image size.
Image
Name
Image
Input axes
image
Transfer Learning (Images)
Name
Transfer Learning (Images)
Input features
Image
Output features
3 (angular_leaf_spot, bean_rust, healthy)
Output features
3 (angular_leaf_spot, bean_rust, healthy)
Save Impulse
Add a processing block
Add a learning block
© 2021 EdgeImpulse Inc. All rights reserved

Transfer Learning (Images) - B

studio.edgeimpulse.com/studio/51151/learning/keras-transfer-image/9

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Image
  - Transfer Learning (I...
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

TRANSFER LEARNING (IMAGES) (BEAN DISEASE CLASSIFIER)

#1 EON Tuner Primary

MJRoBot (Marcelo Rovai)

Neural Network settings

Training settings

Number of training cycles: 20

Learning rate: 0.0005

Data augmentation: ☑

Neural network architecture

Input layer (76,800 features)

MobileNetV2 160x160 0.35 (final layer: 16 neurons, 0.25 dropout)

Choose a different model

Output layer (3 features)

Start training

Training output

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY 91.9%

LOSS 0.23

Confusion matrix (validation set)

	ANGULAR_LEAF_SPOT	BEAN_RUST	HEALTHY
ANGULAR_LEAF_SPOT	96.2%	3.8%	0%
BEAN_RUST	9.2%	82.9%	7.9%
HEALTHY	1.3%	2.5%	96.2%
F1 SCORE	0.93	0.88	0.94

Feature explorer (full training set)

- angular_leaf_spot - correct
- bean_rust - correct
- healthy - correct
- angular_leaf_spot - incorrect
- bean_rust - incorrect
- healthy - incorrect

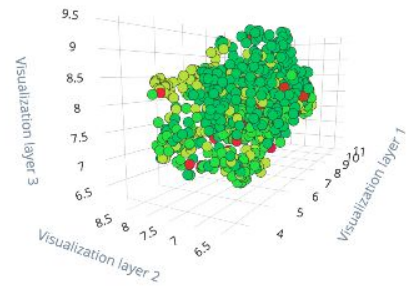

On-device performance

INFERENCING TIME 2,386 ms.

PEAK RAM USAGE 106.6K

FLASH USAGE 226.8K

Model testing - Bean Disease

studio.edgeimpulse.com/studio/51151/validation

EDGE IMPULSE

- Dashboard
- Devices
- Data acquisition
- Impulse design
  - Create impulse
  - Image
  - Transfer Learning (I...
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

MODEL TESTING (BEAN DISEASE CLASSIFIER)

MJRoBot (Marcelo Rovai)

This lists all test data. You can manage this data through Data acquisition.

Test data

Classify all

Set the 'expected outcome' for each sample to the desired outcome to automatically score the impulse.

SAMPLE NAME	EXPECTED OUTCOME	LENGTH	ACCURACY	RESULT
testing.2hiolkem	testing	-		
testing.2hiojo5d	testing	-		
testing.2hiohgsc	testing	-		
testing.2hiofu77	testing	-		
testing.2hiodaqp	testing	-		
testing.2hio9ui3	testing	-		
healthy_test.41.jpg...	healthy	-	100%	1 healthy
healthy_test.40.jpg...	healthy	-	100%	1 healthy
healthy_test.39.jpg...	healthy	-	100%	1 healthy
healthy_test.38.jpg...	healthy	-	100%	1 healthy
healthy_test.37.jpg...	healthy	-	100%	1 healthy
healthy_test.36.jpg...	healthy	-	100%	1 healthy
healthy_test.35.jpg...	healthy	-	100%	1 healthy
healthy_test.34.jpg...	healthy	-	100%	1 healthy
healthy_test.33.jpg...	healthy	-	100%	1 healthy
healthy_test.32.jpg...	healthy	-	100%	1 healthy

Model testing output

Model testing results

ACCURACY
84.38%

	ANGULAR_LEAF_SPOT	BEAN_RUST	HEALTHY	UNCERTAIN
ANGULAR_LEAF_SPOT	79.1%	14.0%	0%	7.0%
BEAN_RUST	7.0%	76.7%	2.3%	14.0%
HEALTHY	0%	2.4%	97.6%	0%
F1 SCORE	0.85	0.80	0.98	

Feature explorer

- angular_leaf_spot - correct
- bean_rust - correct
- healthy - correct
- angular_leaf_spot - incorrect
- bean_rust - incorrect
- healthy - incorrect

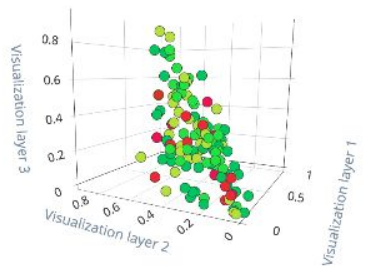

19:22
◂ Camera

AA ne.edgeimpulse.com

Classifier

Next photo

	ANGULAR_LEAF...	BEAN_RUST	HEALTHY
2	0.01	0.02	**0.97**
1	0.00	0.00	**1.00**

19:27

AA ne.edgeimpulse.com

Classifier

Next photo

	ANGULAR_LEAF...	BEAN_RUST	HEALTHY
26	0.00	**1.00**	0.00
25	0.01	**0.99**	0.00

19:27

AA ne.edgeimpulse.com

Classifier

Next photo

	ANGULAR_LEAF...	BEAN_RUST	HEALTHY
25	0.01	**0.99**	0.00
24	0.00	**1.00**	0.00

19:29

AA ne.edgeimpulse.com

Classifier

Next photo

	ANGULAR_LEAF...	BEAN_RUST	HEALTHY
33	**0.99**	0.01	0.00
32	**0.88**	0.12	0.00

SciTinyML - ICTP workshop

Scientific Use of Machine Learning on Low Power Devices

Motion Classification – Anomaly Detection

Prof. Marcelo José Rovai
UNIFEI - Universidade Federal de Itajubá, Brazil
Web: https://github.com/Mjrovai
Email: rovai@unifei.edu.br