tinyML® Talks

Enabling Ultra-low Power Machine Learning at the Edge

"Unleashing the Power of the New XIAO ESP32S3 Sense: Tackling Anomaly Detection, Image Classification, and Keyword Spotting with TinyML"

Marcelo Rovai – Co-Chair, TinyML4D group

June 13, 2023

www.tinyML.org

TINY ML TALKS webcast

Thank you, **tinyML Strategic Partners**, for committing to take tinyML to the next Level, together

TINY ML TALKS webcast

Executive Strategic Partners

TINY ML TALKS webcast

EDGE IMPULSE

The Leading Development Platform for Edge ML

edgeimpulse.com

Qualcomm
AI research

Advancing AI research to make efficient AI ubiquitous

Power efficiency

Model design, compression, quantization, algorithms, efficient hardware, software tool

Personalization

Continuous learning, contextual, always-on, privacy-preserved, distributed learning

Efficient learning

Robust learning through minimal data, unsupervised learning, on-device learning

A platform to scale AI across the industry

Perception

Object detection, speech recognition, contextual fusion

Reasoning

Scene understanding, language understanding, behavior prediction

Action

Reinforcement learning for decision making

Edge cloud

Cloud

IoT/IIoT

Automotive

Mobile

Accelerate Your Edge Compute

SYNTIANT

Making Edge AI A Reality

www.syntiant.com

TINY ML TALKS webcast

Platinum Strategic Partners

Renesas is enabling the next generation of AI-powered solutions that will revolutionize every industry sector.

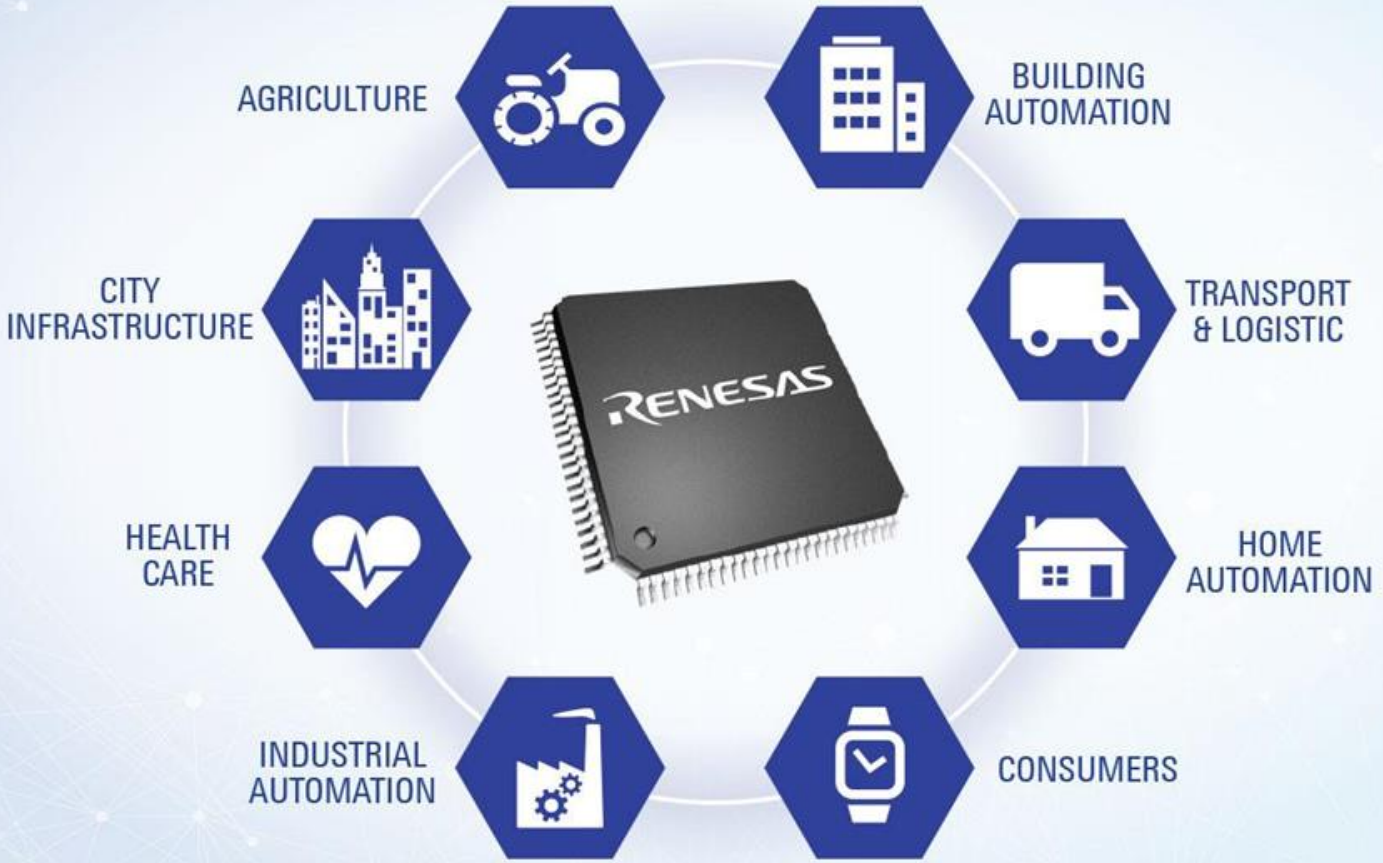

renesas.com

AITRIOS

DEPLOY VISION AI AT THE EDGE AT SCALE

SONY

TINY ML TALKS webcast

Gold Strategic Partners

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

Where what if becomes what is.

Witness potential made possible at analog.com.

ARDUINO PRO™

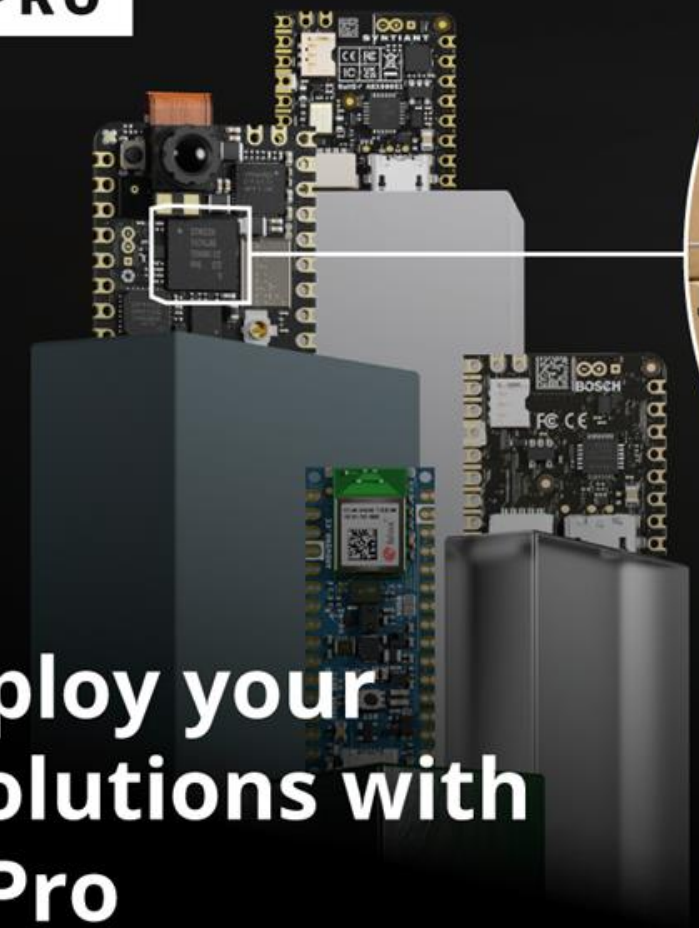

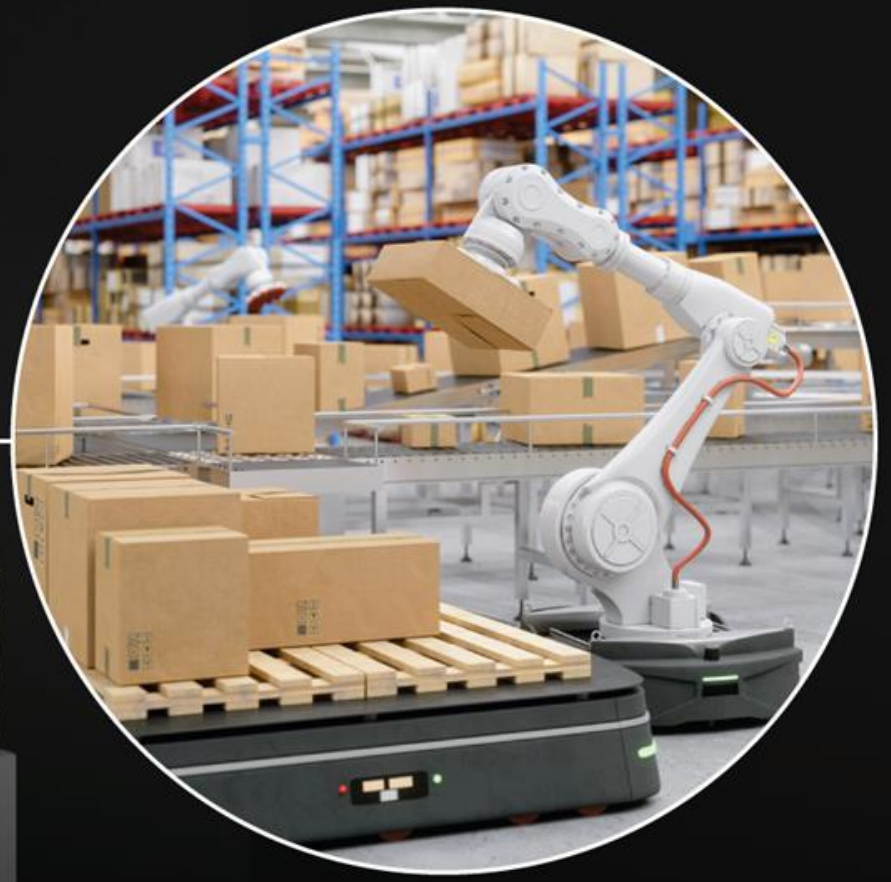

Easily deploy your tinyML solutions with Arduino Pro

arduino.cc/pro

Made In Italy

arm AI

Powering tinyML Innovation

Arm AI Virtual Tech Talks

The latest in AI trends, technologies & best practices from Arm and our Ecosystem Partners.

Demos, code examples, workshops, panel sessions and much more!

Fortnightly Tuesday @ 4pm GMT/8am PT

Find out more:
www.arm.com/techtalks

Driving decarbonization and digitalization. Together.

Infineon serving all target markets as
Leader in Power Systems and IoT

www.infineon.com

innatera

NEUROMORPHIC INTELLIGENCE FOR THE SENSOR-EDGE

www.innatera.com


Microsoft

The Right Edge AI Tools Can Make or Break Your Next Smart IoT Product

SensiML

Analytics Toolkit Suite

sensiml.com/tinyML

STMicroelectronics provides extensive solutions to make tiny Machine Learning easy

life.augmented

www.st.com/ai

synaptics®

ENGINEERING EXCEPTIONAL EXPERIENCES

We engineer exceptional experiences for consumers in the home, at work, in the car, or on the go.

www.synaptics.com

TINY ML TALKS webcast

Silver Strategic Partners

brainchip

GREENWAVES TECHNOLOGIES

Groovety Inc.

IBM

imagimob

Nota AI

NXP

POLYN TECHNOLOGY

Qeexo

Schneider Electric

SILICON LABS

Join Growing tinyML Communities:

145k members in
48 Groups in 40 Countries

tinyML - Enabling ultra-low Power ML at the Edge

https://www.meetup.com/tinyML-Enabling-ultra-low-Power-ML-at-the-Edge/

4k members
&
11.6k followers

The tinyML Community

https://www.linkedin.com/groups/13694488/

TINY ML TALKS webcast

YouTube

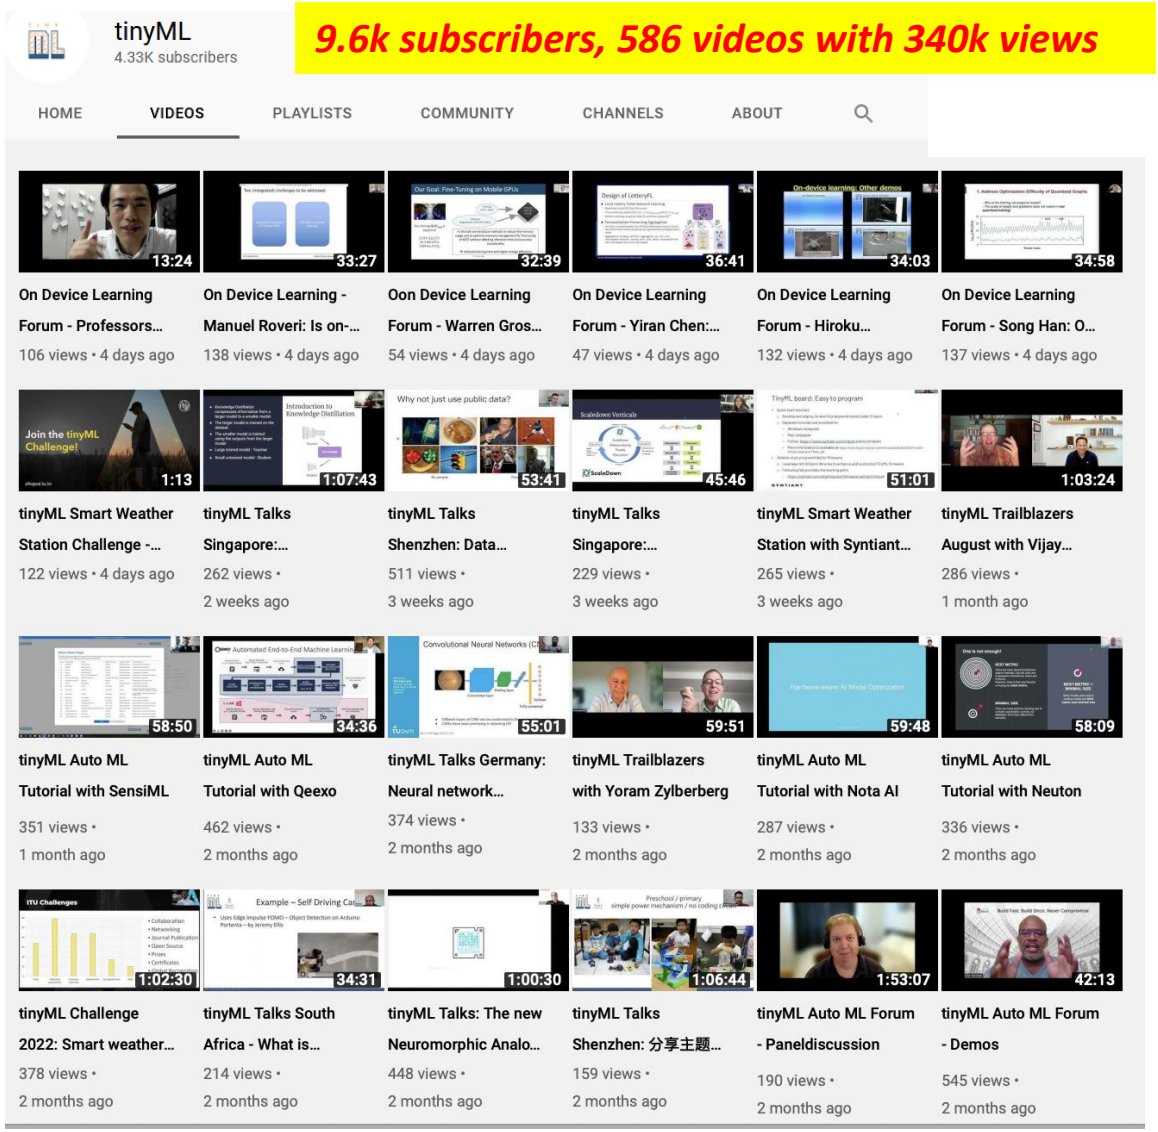

Subscribe to tinyML YouTube Channel for updates and notifications

(including this video)

www.youtube.com/tinyML

TINY ML TALKS webcast

TINY ML FOUNDATION

EMEA 2023

https://www.tinyml.org/event/emea-2023

More sponsorships are available: sponsorships@tinyML.org

TINY ML TALKS webcast

Reminders

Slides & Videos will be posted tomorrow

PDF

tinyml.org/forums

youtube.com/tinyml

Please use the Q&A window for your questions

TINY ML TALKS webcast

Marcelo Rovai

Marcelo Rovai is born in São Paulo and held a Master's in Data Science from the Universidad del Desarrollo (UDD) in Chile and an MBA from IBMEC (INSPER) in Brazil. He graduated in 1982 as an Engineer from UNIFEI, Federal University of Itajubá, with a specialization from Escola Politécnica de Engenharia of São Paulo University (USP), both institutions located in Brazil. Rovai has experience 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 VP and a Senior Advisor for Latin America. Marcelo Rovai publishes articles about electronics on websites such as MJRoBot.org, Hackster.io, Instructables.com, and Medium.com. Furthermore, he is a volunteer Professor at the UNIFEI in Brazil and a lecturer at several Congresses and Universities on IoT and TinyML. He is an active member and a Co-Chair of the TinyML4D group, an initiative to bring TinyML education to developing countries.

TINY ML TALKS webcast

Unleashing the Power of the New XIAO ESP32S3 Sense

Tackling Anomaly Detection, Image Classification, and Keyword Spotting with TinyML

Prof. Marcelo José Rovai
UNIFEI - Federal University of Itajubá, Brazil
TinyML4D Academic Network Co-Chair

UNIFEI

TinyML4D Academic Network

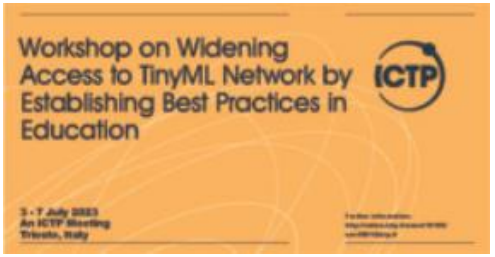

TINY ML TALKS webcast

Sound

TinyML Made Easy: KeyWord Spotting (KWS)

MJRoBot (Marcelo Rovai)

Vibration

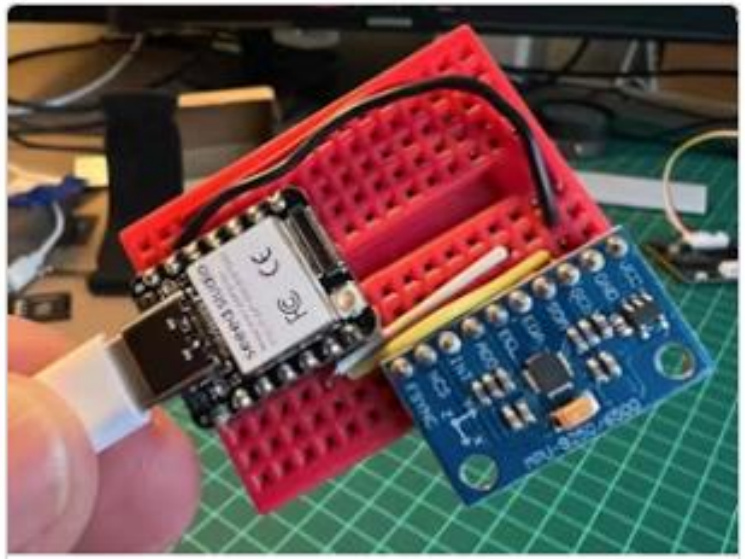

Exploring Machine Learning with the new XIAO ESP32S3

MJRoBot (Marcelo Rovai)

Vision

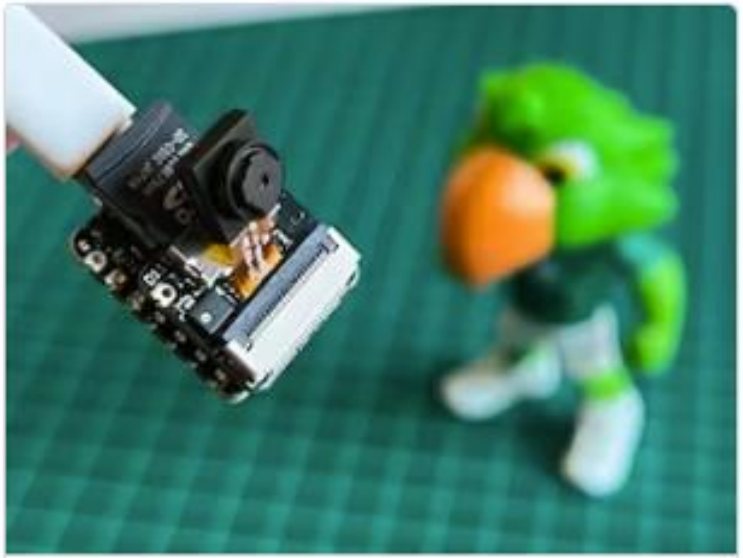

TinyML Made Easy: Image Classification

MJRoBot (Marcelo Rovai)

Embedded ML (TinyML)

Introduction

TINY ML TALKS webcast

What Makes TinyML?

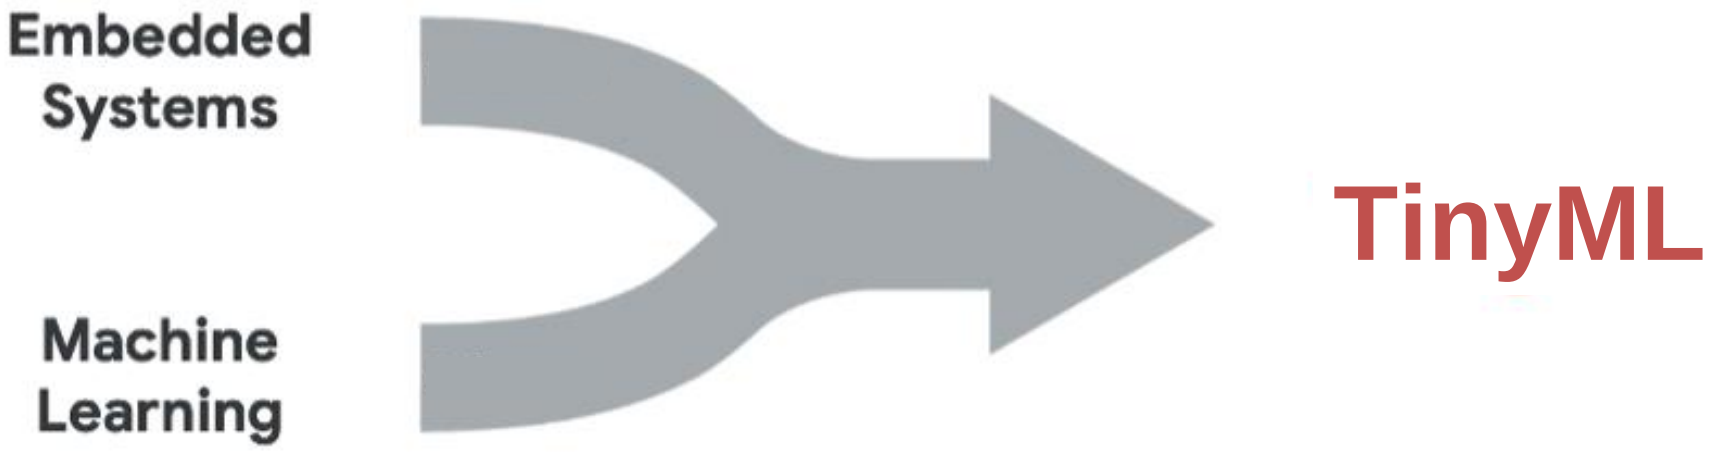

TINY ML TALKS webcast

What Makes TinyML?

TensorFlow Lite

TinyML

Hardware (Dev. Boards)





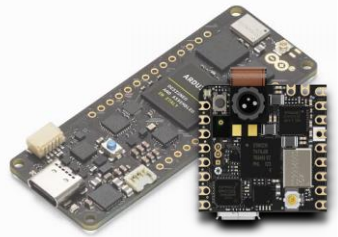

	Raspberry Pico (W)	Arduino Nano Sense	Espressif ESP 32	Seeed XIAO ESP32S3 Sense	Arduino Pro
32Bits CPU	Dual-core Arm Cortex-M0+	Arm Cortex-M4F	Xtensa LX6 Dual Core	Xtensa LX7 Dual Core	Dual Core Arm Cortex M7/M4
CLOCK	133MHz	64MHz	240MHz	240MHz	480/240MHz
RAM	264KB	256KB	520KB (part available)	8MB (PSRAM)	1MB
ROM	2MB	1MB	2MB	8MB	2MB
Radio	(Yes for W)	BLE	BLE / WiFi	BLE / WiFi	BLE / WiFi
Sensors	No	Yes	No	Yes	Yes (Nicla)
Bat. Power Manag.	No	No	No	Yes	Yes
Price	$	$$$	$	$$	$$$$$

https://media.digikey.com/Resources/Maker/the-original-guide-to-boards-2022.pdf

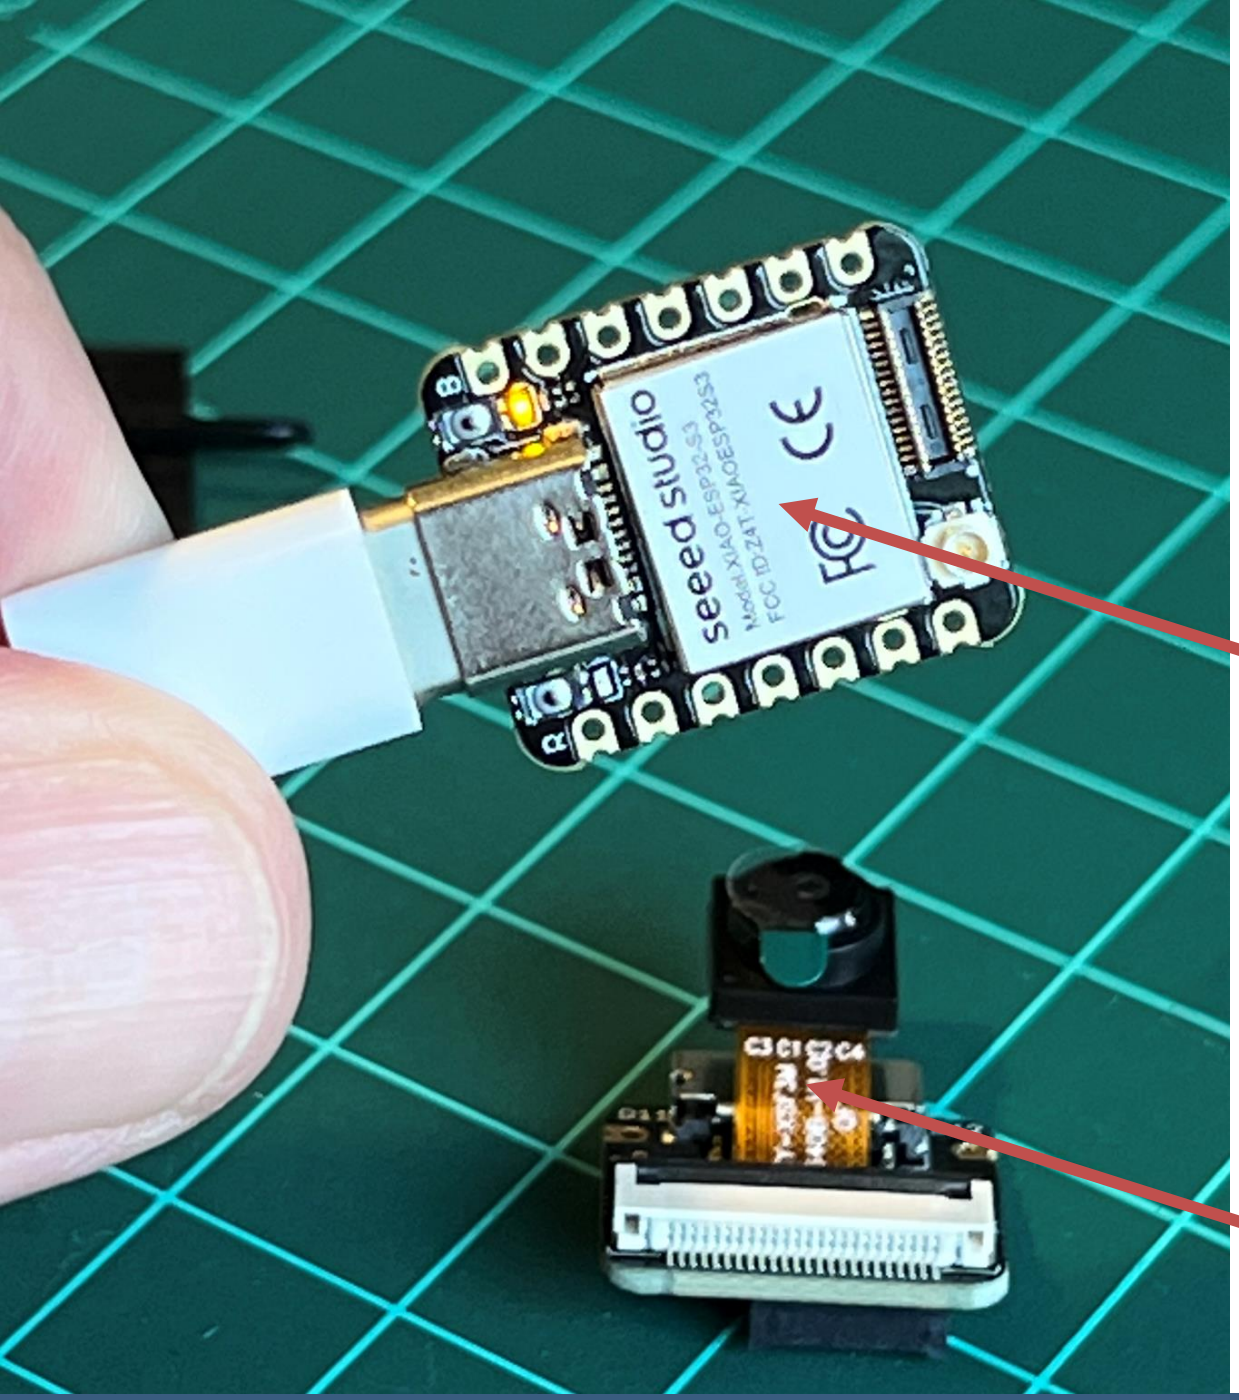

Powerful MCU Board: ESP32S3 32-bit, dual-core, Xtensa processor chip operating up to 240 MHz.

Elaborate Power Design: Lithium battery charge management capability (deep sleep mode with power consumption as low as 14μA)

Great Memory for more Possibilities: Offer 8MB PSRAM and 8MB FLASH

Outstanding RF performance: Support 2.4GHz Wi-Fi and BLE dual wireless communication, support 100m+ remote communication when connected with U.FL antenna

Thumb-sized Compact Design: 21 x 17.5mm, adopting the classic form factor of XIAO, suitable for space-limited projects like wearable devices

Advanced Functionality: Detachable OV2640 camera sensor for 1600*1200 resolution, compatible with OV5640 camera sensor, integrating an additional digital microphone and an SD card slot for external 32GB FAT memory.

TINY ML TALKS webcast

Software

TensorFlow

TensorFlow Lite

Train a model → Convert model → Optimize model → Deploy model at Edge → Make inferences at Edge

Raspberry Pi

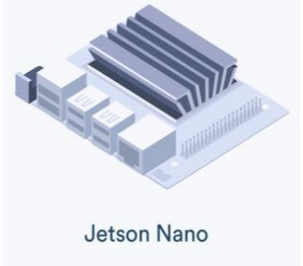

Jetson Nano

Linux

iOS

TensorFlow Lite Micro

TINY ML TALKS webcast

Datasets Preprocessing

Sound

Vision

Motion

Quantization Pruning

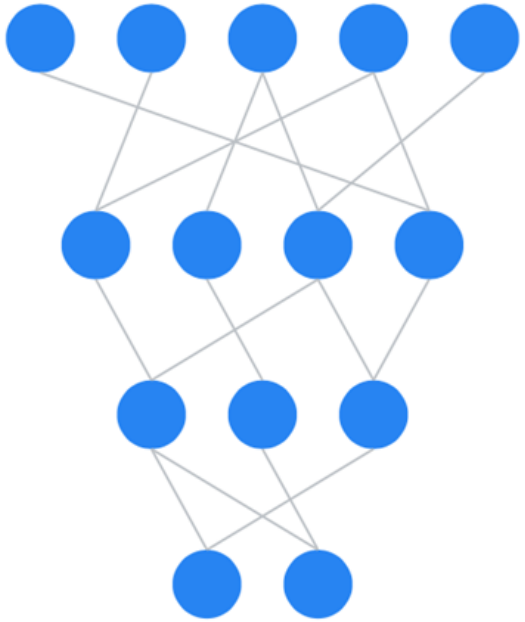

Resource constraints

End-to-end **TinyML** application design

TINY ML TALKS webcast

How to Train a ML Model?

TINY ML TALKS webcast

Machine Learning Workflow

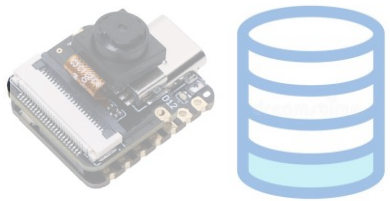

Collect Data → Preprocess Data → Design a Model → Train a Model → Evaluate Optimize → Convert Model → Deploy Model → Make Inferences

TensorFlow | TensorFlow Lite | TensorFlow Lite Micro

EDGE IMPULSE

TINY ML TALKS webcast

Machine Learning Workflow

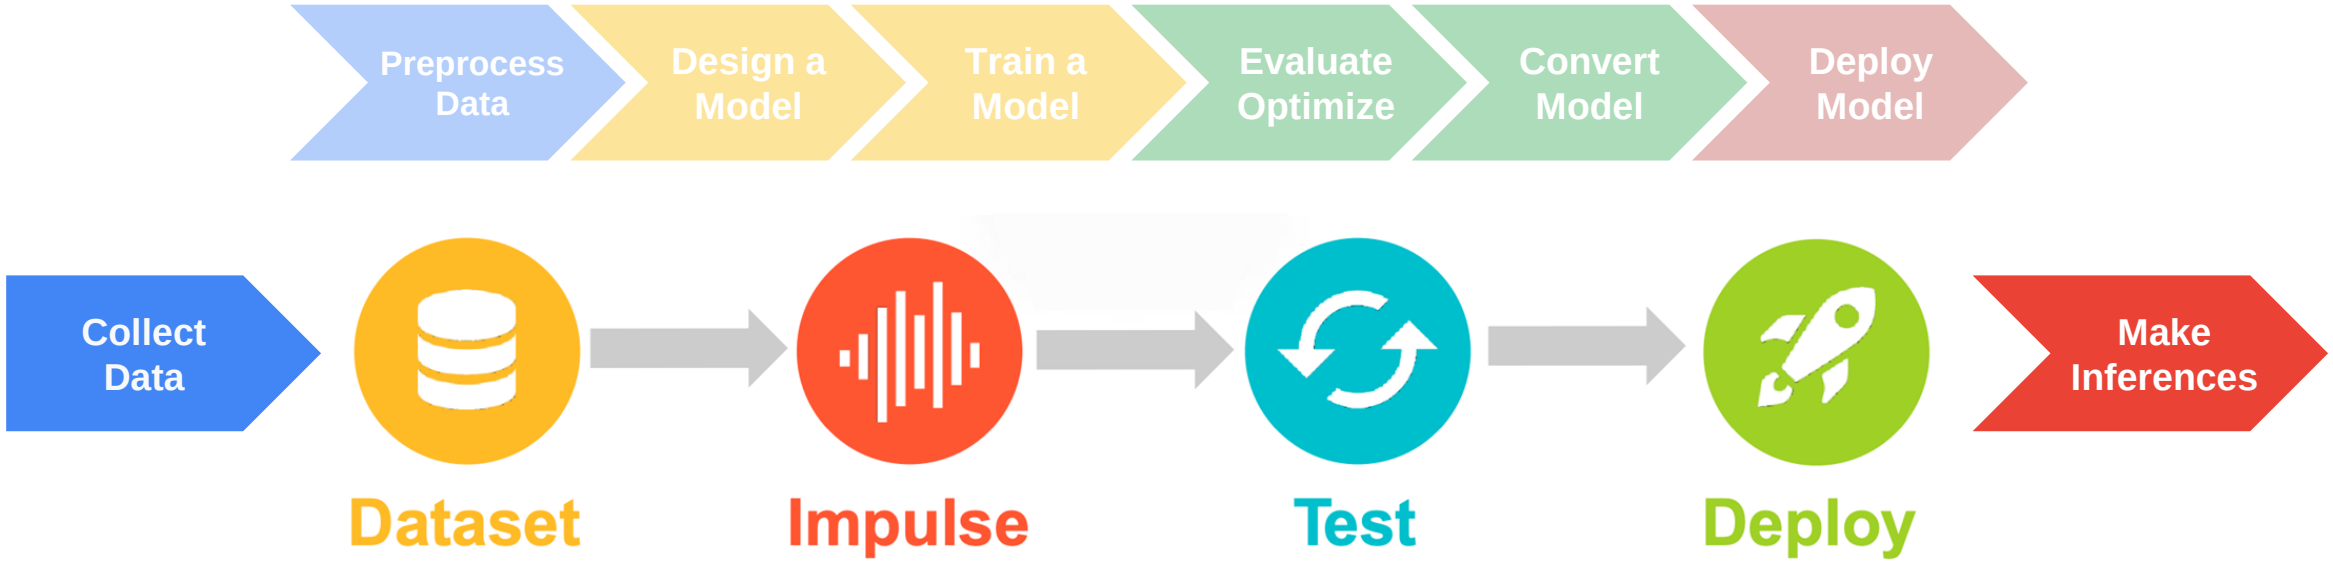

TINY ML TALKS webcast

EI Studio - Embedded ML platform

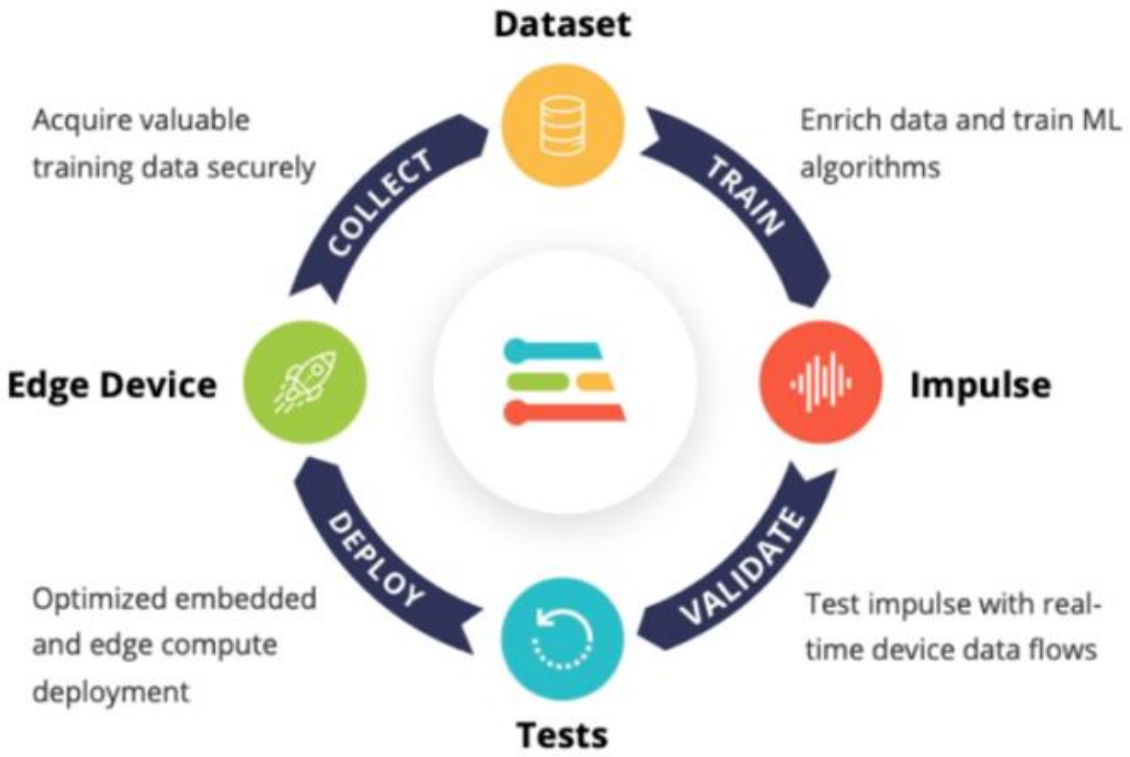

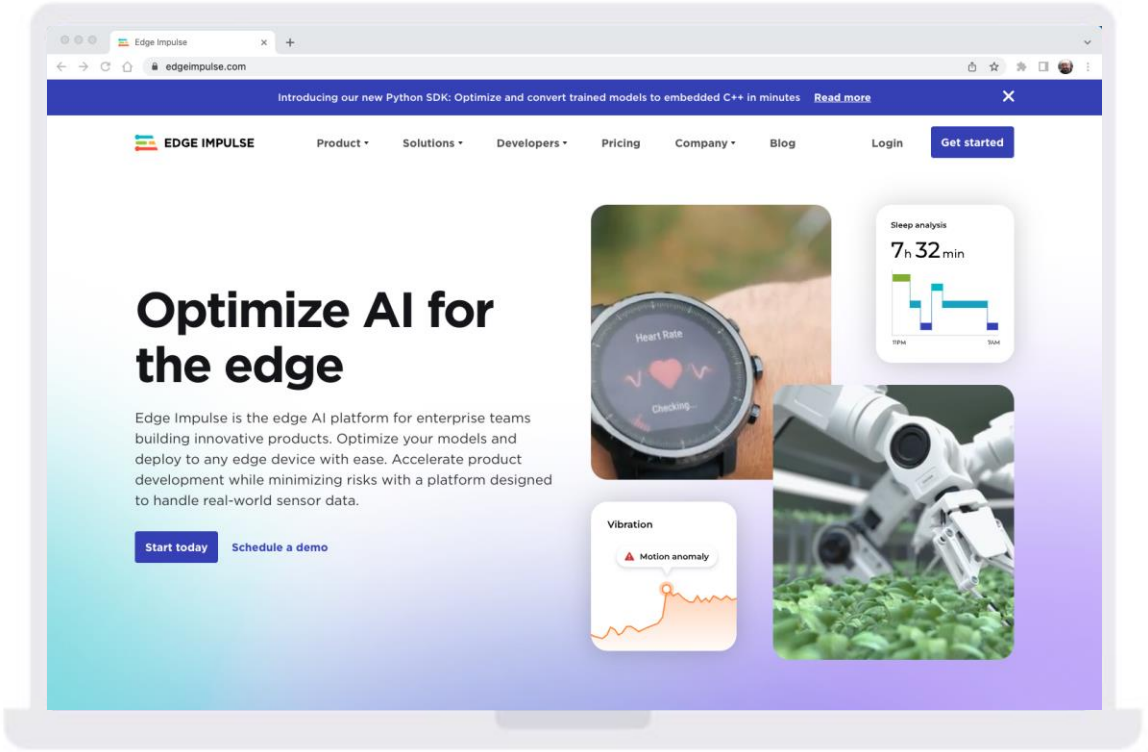

Learn more at http://edgeimpulse.com

[2023 SciTinyML Workshop - Edge Impulse Overview by Shawn Hymel]

TINY ML TALKS webcast

TinyML Applications

TINY ML TALKS webcast

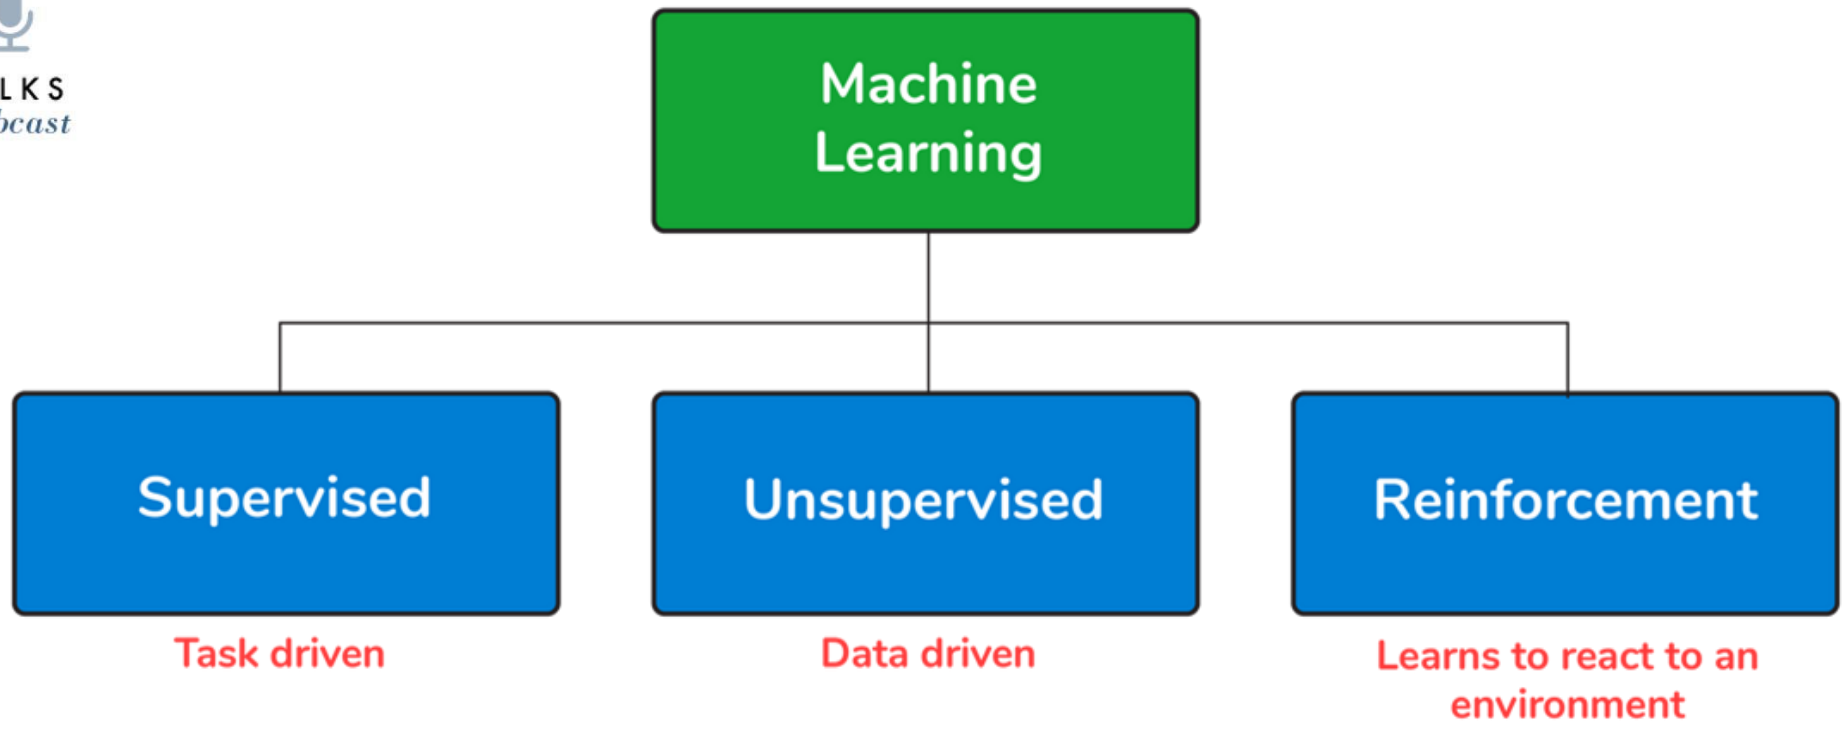

- Regression
- **Classification**
- Object Detection

- **Anomaly Detection**

- Autonomous Navigation

TINY ML TALKS webcast

Sound

Vibration

Vision

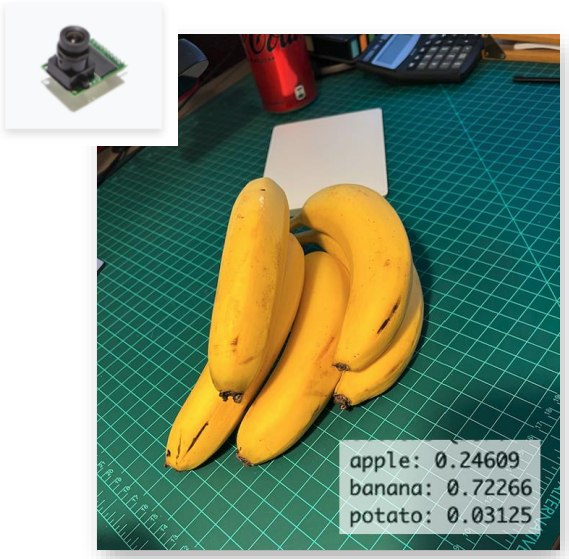

TINY ML TALKS webcast

Sound

Vibration

Vision

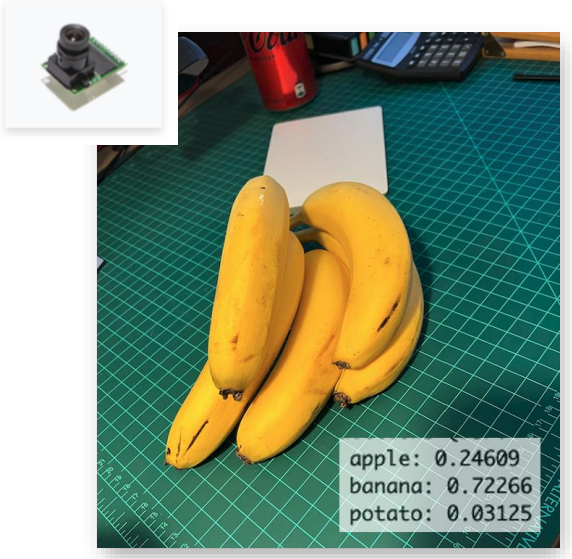

TINY ML TALKS webcast

Sound

Vibration

Vision

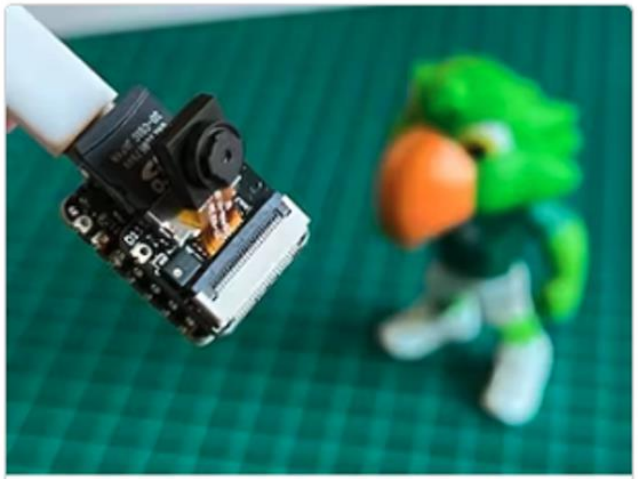

TinyML Made Easy: Image Classification

MJRoBot (Marcelo Rovai)

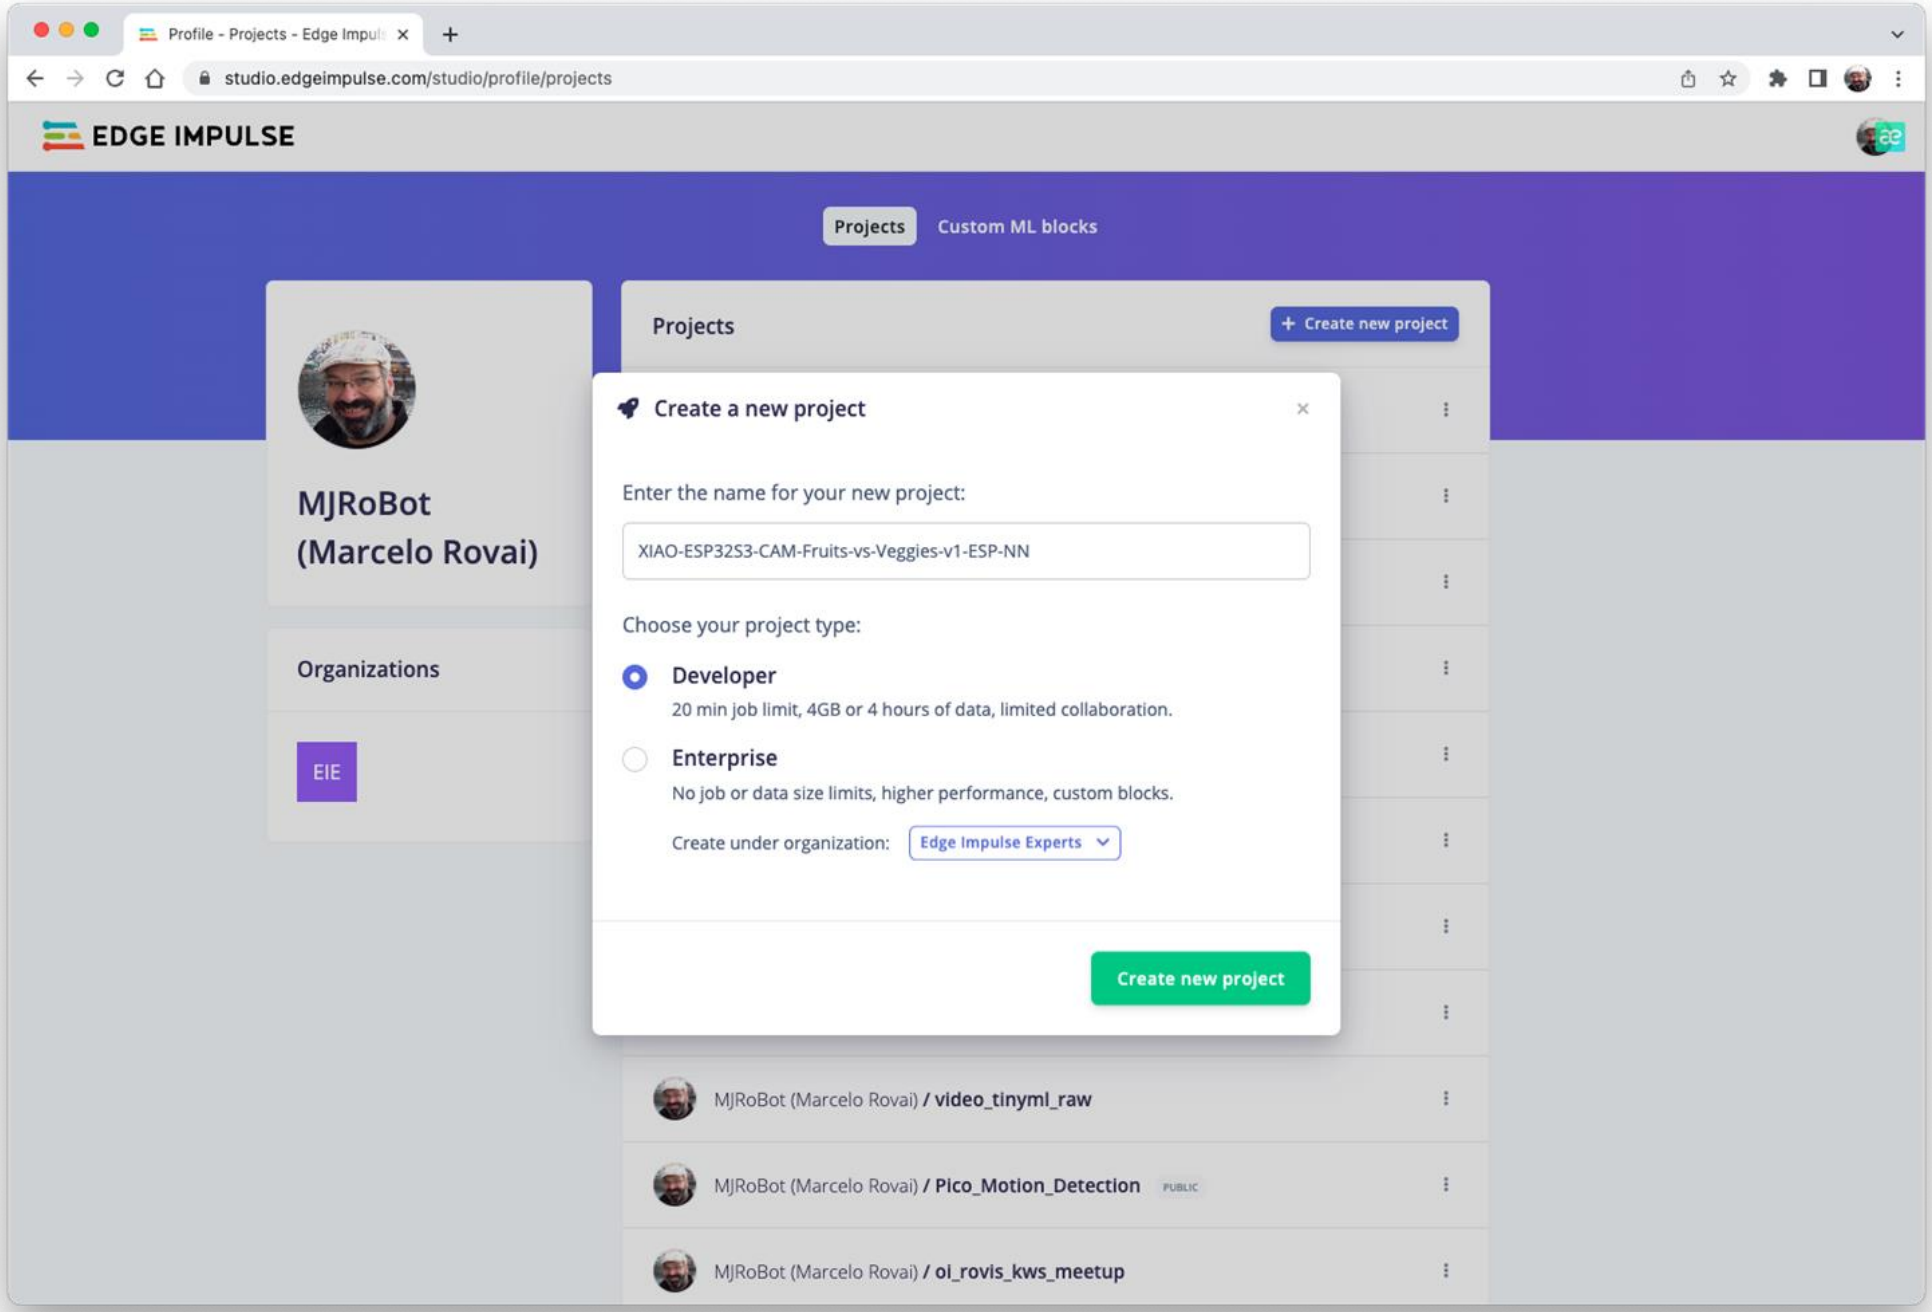

XIAO-ESP32S3-CAM-Fruits-vs-Veggies-v1-ESP-NN (Edge Impulse)

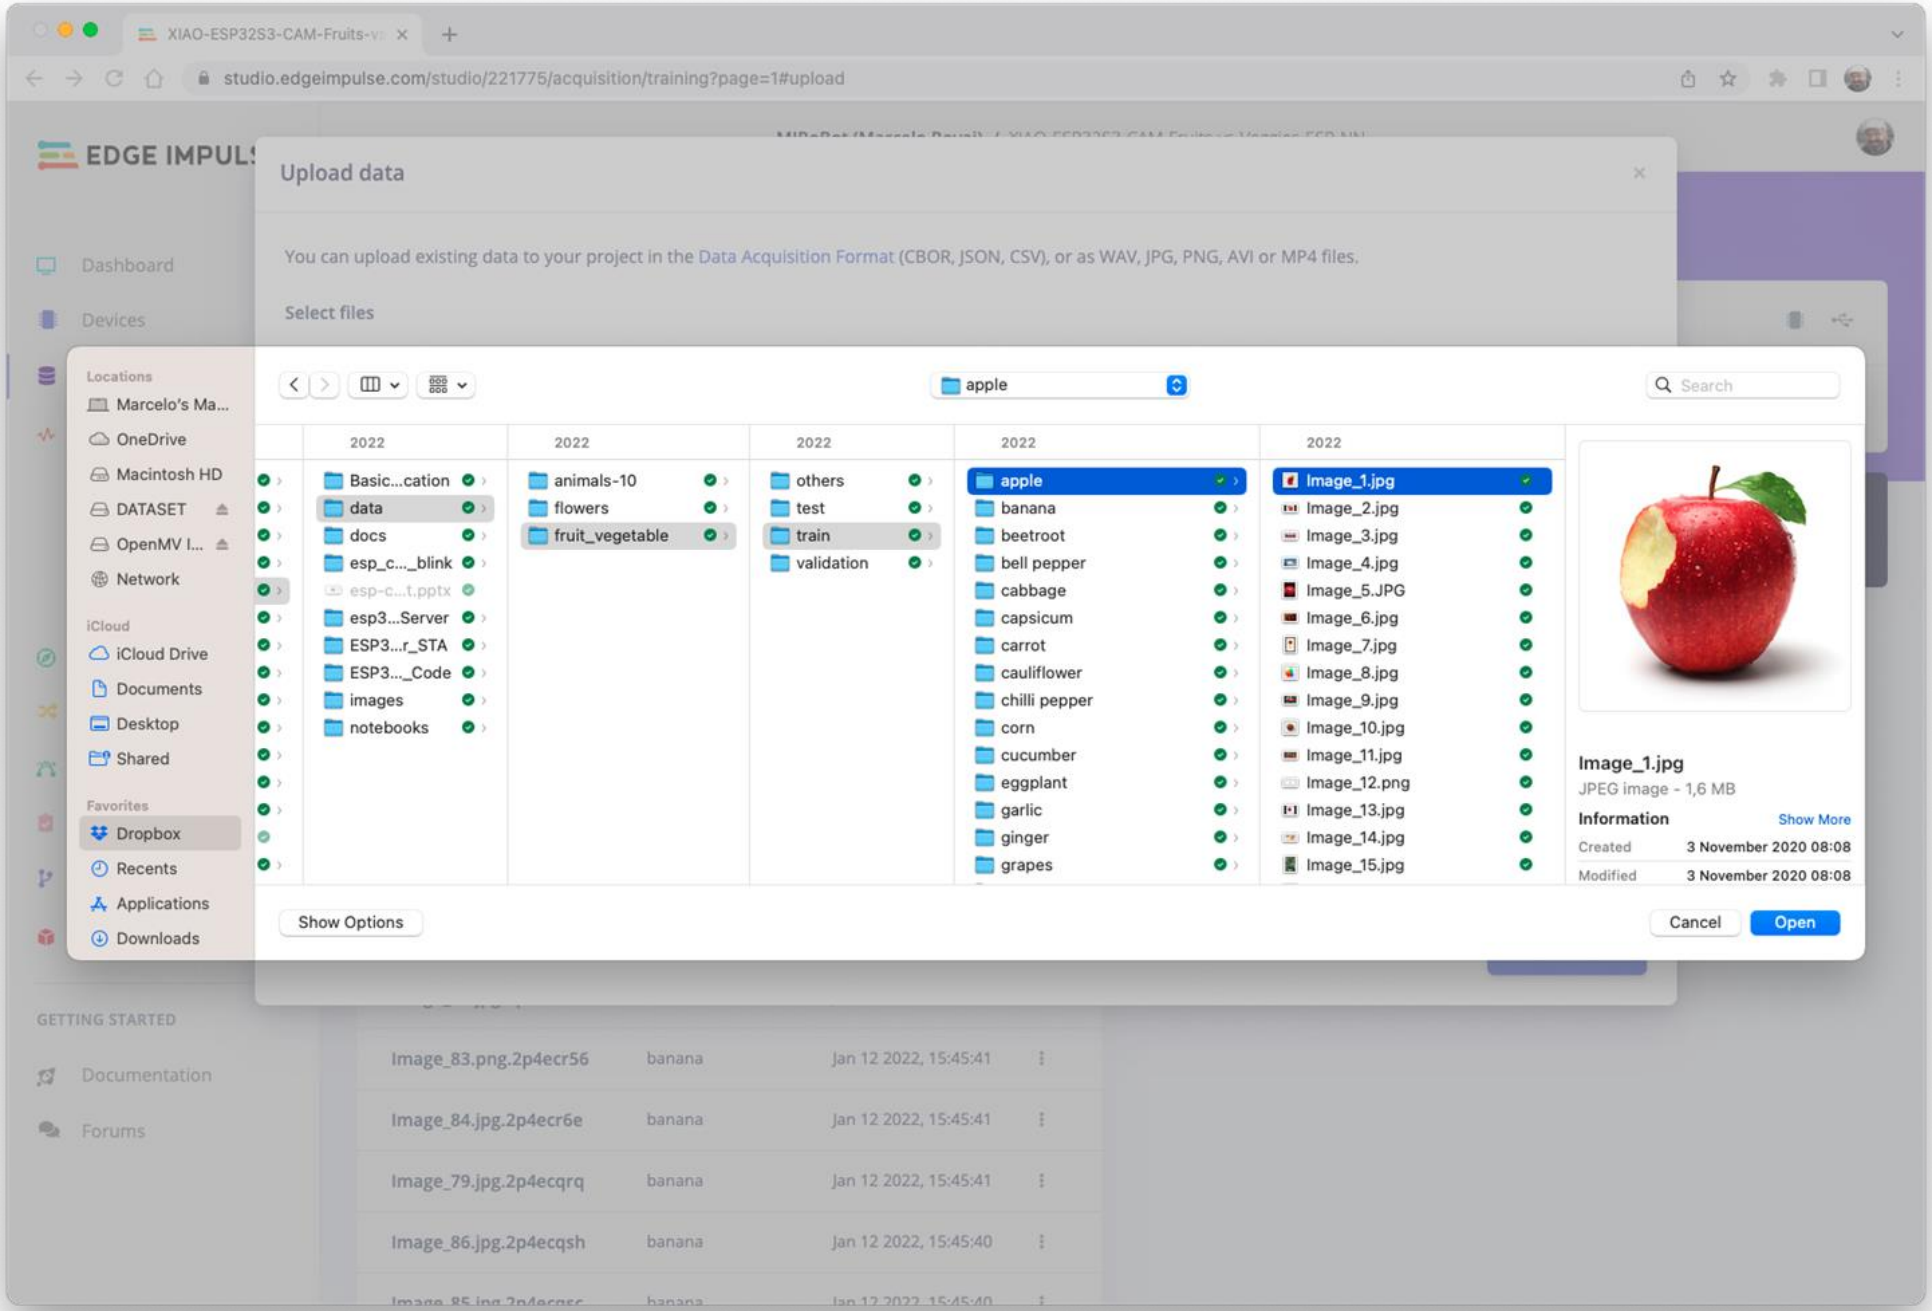

Fruits and Vegetables Image Recognition Dataset (Kaggle)

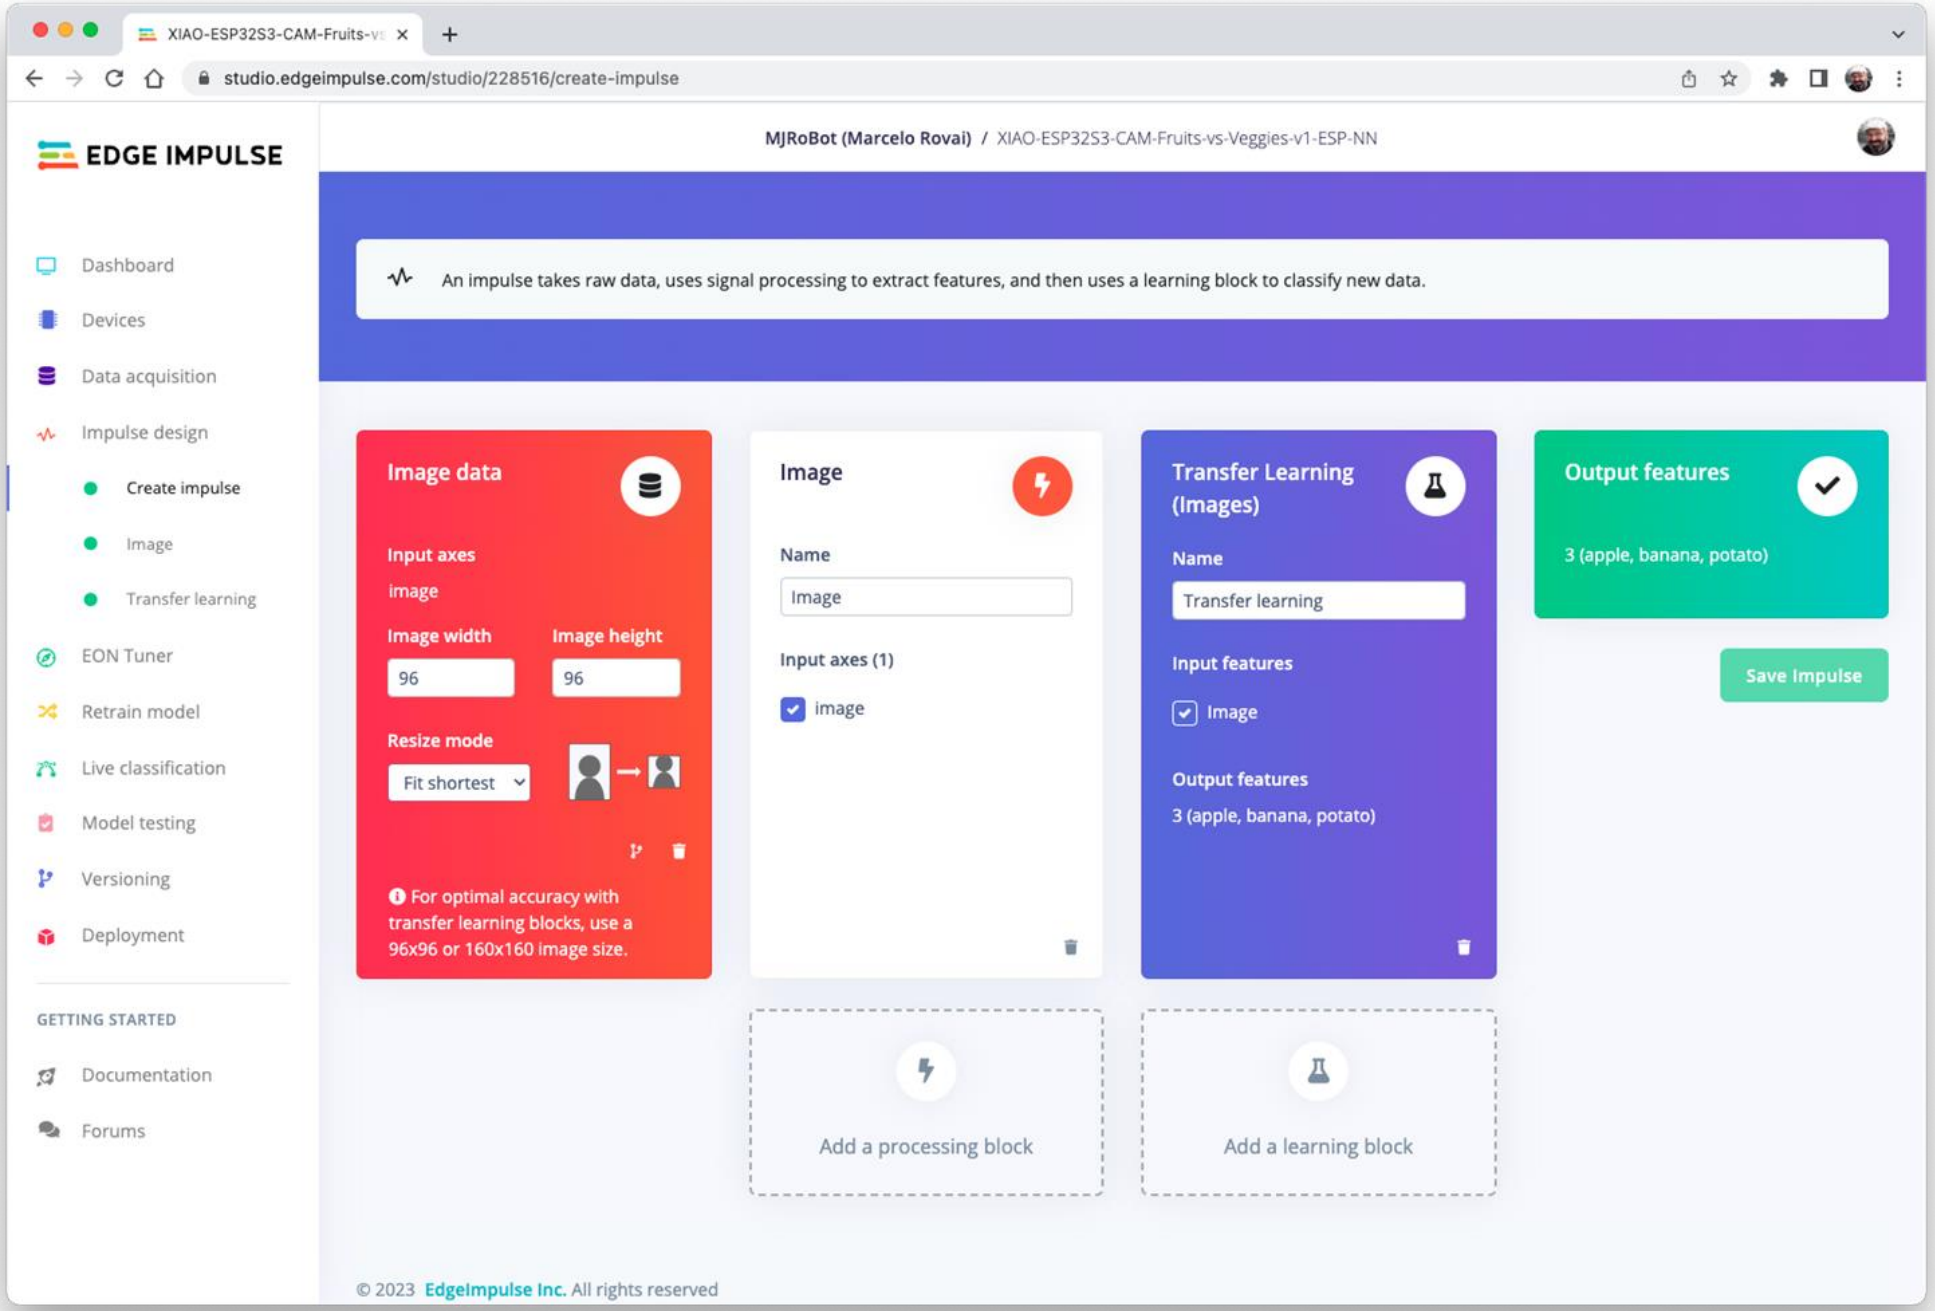
XIAO-ESP32S3-CAM-Fruits-v
studio.edgeimpulse.com/studio/228516/create-impulse
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / XIAO-ESP32S3-CAM-Fruits-vs-Veggies-v1-ESP-NN
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Image data
Input axes
image
Image width
96
Image height
96
Resize mode
Fit shortest
For optimal accuracy with transfer learning blocks, use a 96x96 or 160x160 image size.
Image
Name
Image
Input axes (1)
image
Transfer Learning (Images)
Name
Transfer learning
Input features
Image
Output features
3 (apple, banana, potato)
Output features
3 (apple, banana, potato)
Save Impulse
Add a processing block
Add a learning block
© 2023 EdgeImpulse Inc. All rights reserved

Neural Network settings

Training settings

Number of training cycles: 20

Learning rate: 0.0005

Data augmentation: ☑

Advanced training settings

Validation set size: 20 %

Split train/validation set on metadata key:

Auto-balance dataset: ☐

Profile int8 model: ☑

Neural network architecture

Input layer (27,648 features)

MobileNetV2 96x96 0.05 (final layer: 8 neurons, 0.1 dropout)

Choose a different model

Output layer (3 classes)

Start training

Target: Arduino Portenta H7 (Cortex-M7 480MHz)

Training output

(0)

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY
80.4%

LOSS
0.46

Confusion matrix (validation set)

	APPLE	BANANA	POTATO
APPLE	94.7%	5.3%	0%
BANANA	22.7%	54.5%	22.7%
POTATO	0%	0%	100%
F1 SCORE	0.86	0.69	0.86

Data explorer (full training set)

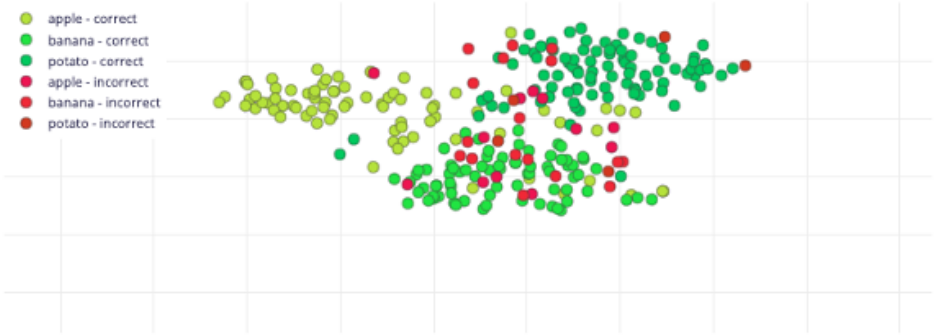

On-device performance

INFERENCING TIME
45 ms.

PEAK RAM USAGE
270.1K

FLASH USAGE
156.9K

TINY ML TALKS webcast

Serial Monitor

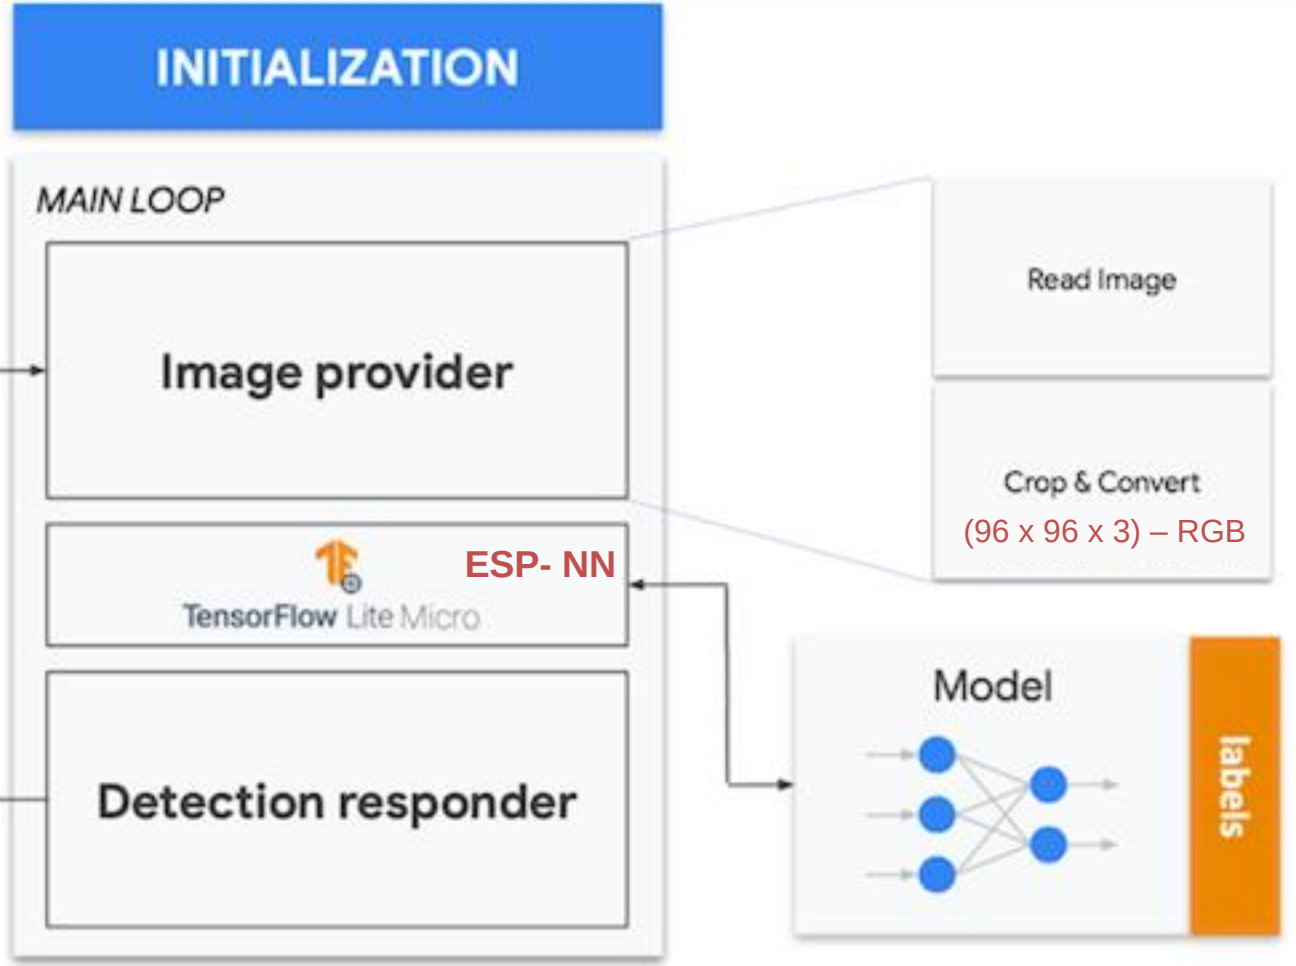

TINY ML TALKS webcast

/dev/cu.usbmodem1101

Send

```
10:44:47.849 ->     banana: 0.01953
10:44:47.849 ->     potato: 0.12891
10:44:48.103 -> Predictions (DSP: 3 ms., Classification: 135 ms., Anomaly: 0 ms.):
10:44:48.103 ->     apple: 0.86328
10:44:48.103 ->     banana: 0.03906
10:44:48.103 ->     potato: 0.10156
10:44:48.356 -> Predictions (DSP: 3 ms., Classification: 135 ms., Anomaly: 0 ms.):
10:44:48.356 ->     apple: 0.90234
10:44:48.356 ->     banana: 0.02344
10:44:48.356 ->     potato: 0.07422
10:44:48.612 -> Predictions (DSP: 3 ms., Classification: 135 ms., Anomaly: 0 ms.):
10:44:48.612 ->     apple: 0.91797
10:44:48.612 ->     banana: 0.02344
10:44:48.612 ->     potato: 0.05859
10:44:48.861 -> Predictions (DSP: 3 ms., Classification: 135 ms., Anomaly: 0 ms.):
10:44:48.861 ->     apple: 0.88281
10:44:48.861 ->     banana: 0.03516
10:44:48.861 ->     potato: 0.08203
10:44:49.114 -> Predictions (DSP: 3 ms., Classification: 135 ms., Anomaly: 0 ms.):
```

Autoscroll Show timestamp Both NL & CR 115200 baud Clear output

135 ms

XIAO ESP32S3
Xtensa LX7
240 MHz

171 ms

ESP - CAM
Xtensa LX6
240 MHz

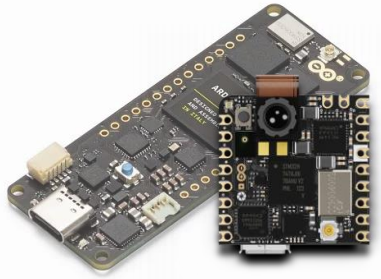

45 ms

ARDUINO Pro
ARM H7
480 MHz

TINY ML TALKS webcast

Sound

Vibration

Vision

TINY ML TALKS webcast

Sound

Vibration

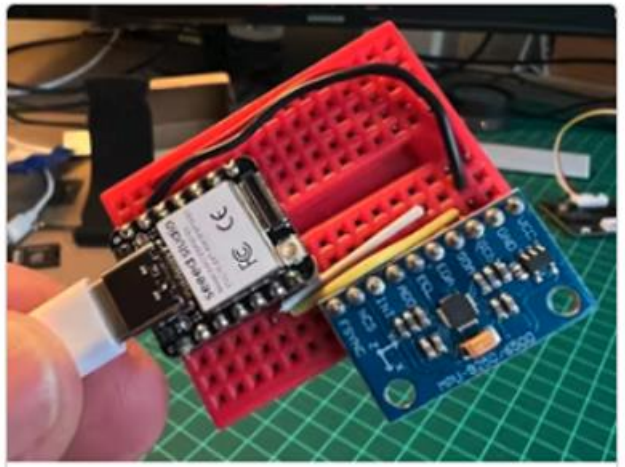

Exploring Machine Learning with the new XIAO ESP32S3

MJRoBot (Marcelo Rovai)

Vision

TINY ML TALKS webcast

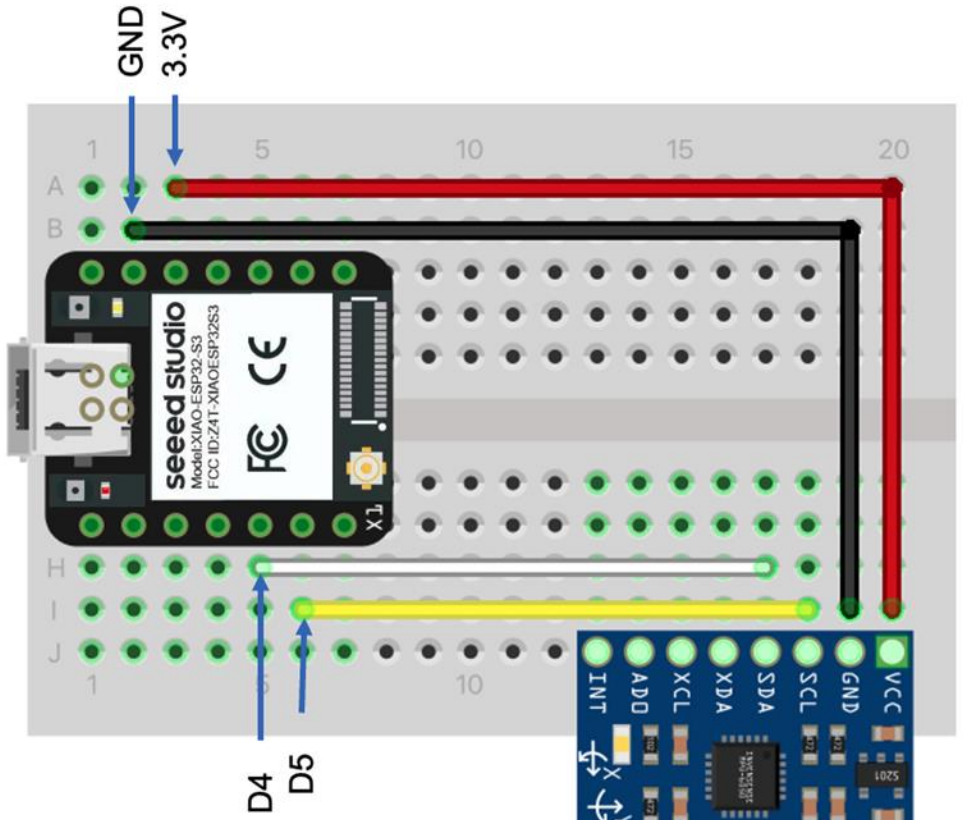

MPU6050 **SCL** --> XIAO **D5**
MPU6050 **SDA** --> XIAO **D4**
MPU6050 **VCC** --> XIAO **3.3V**
MPU6050 **GND** --> XIAO **GND**

fritzing

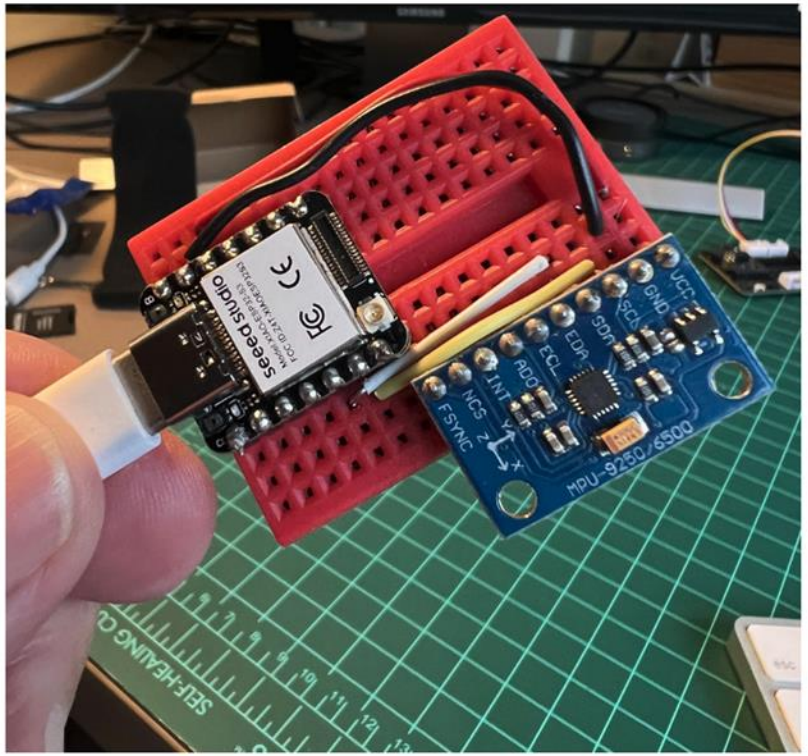

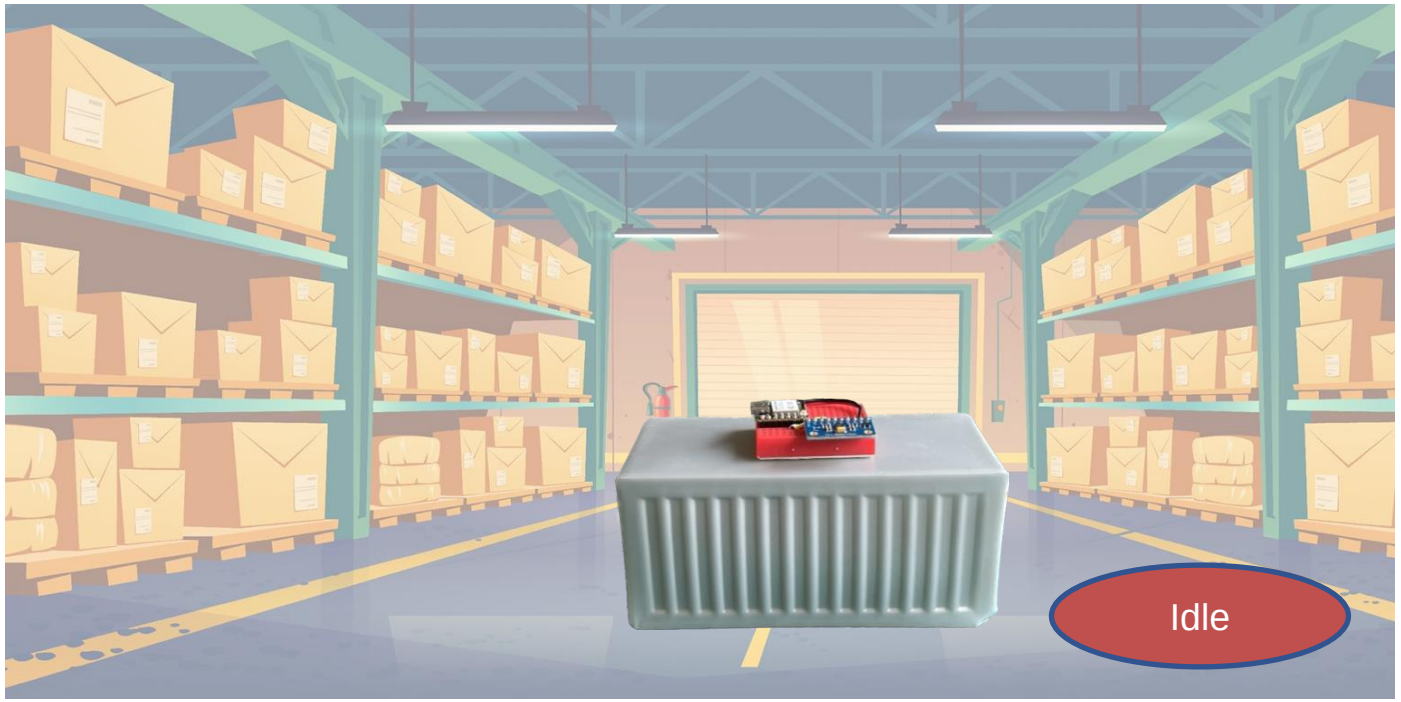
Idle


Lift

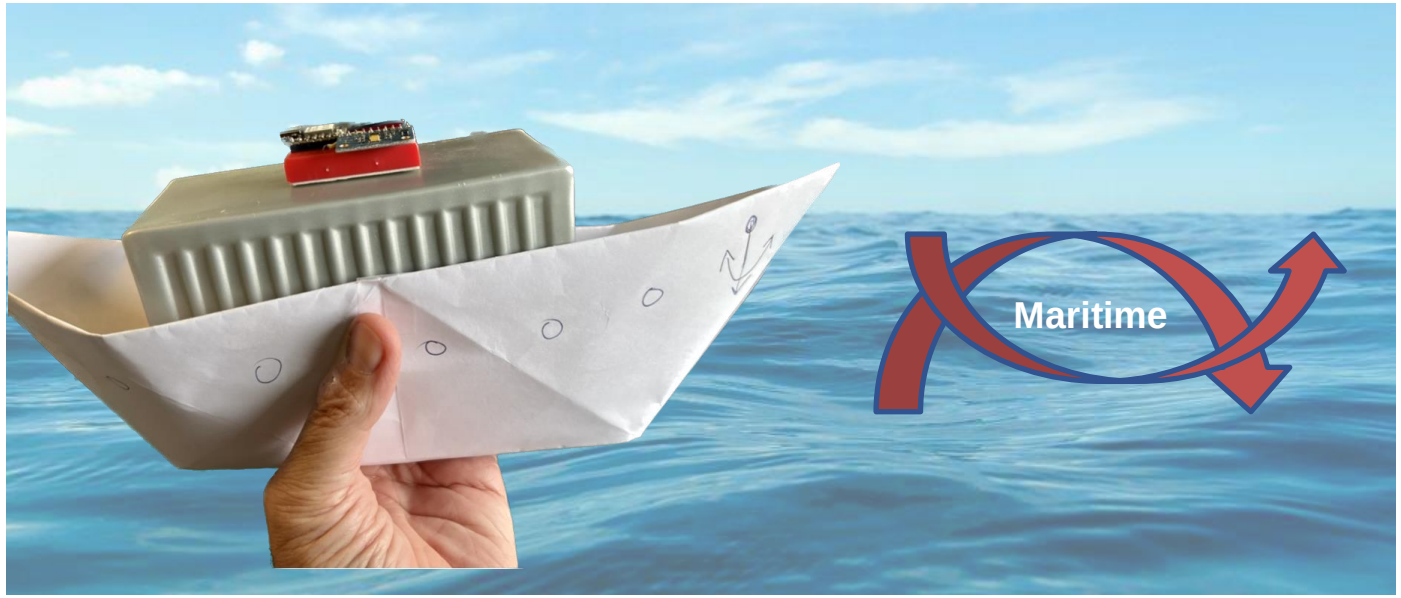
Maritime

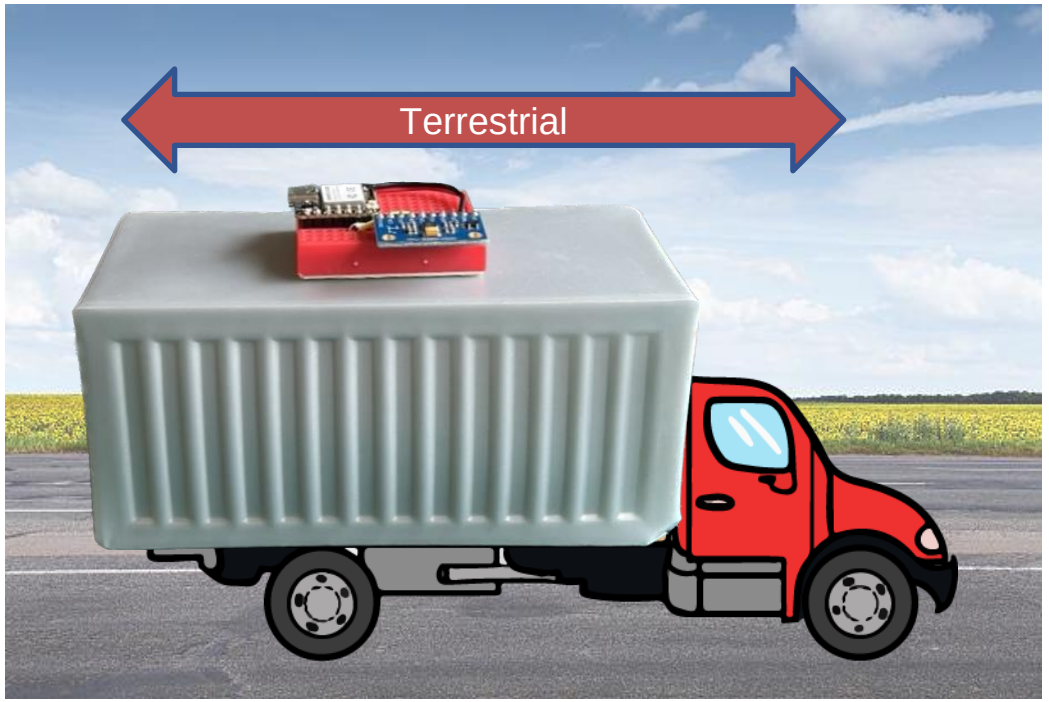
Terrestrial

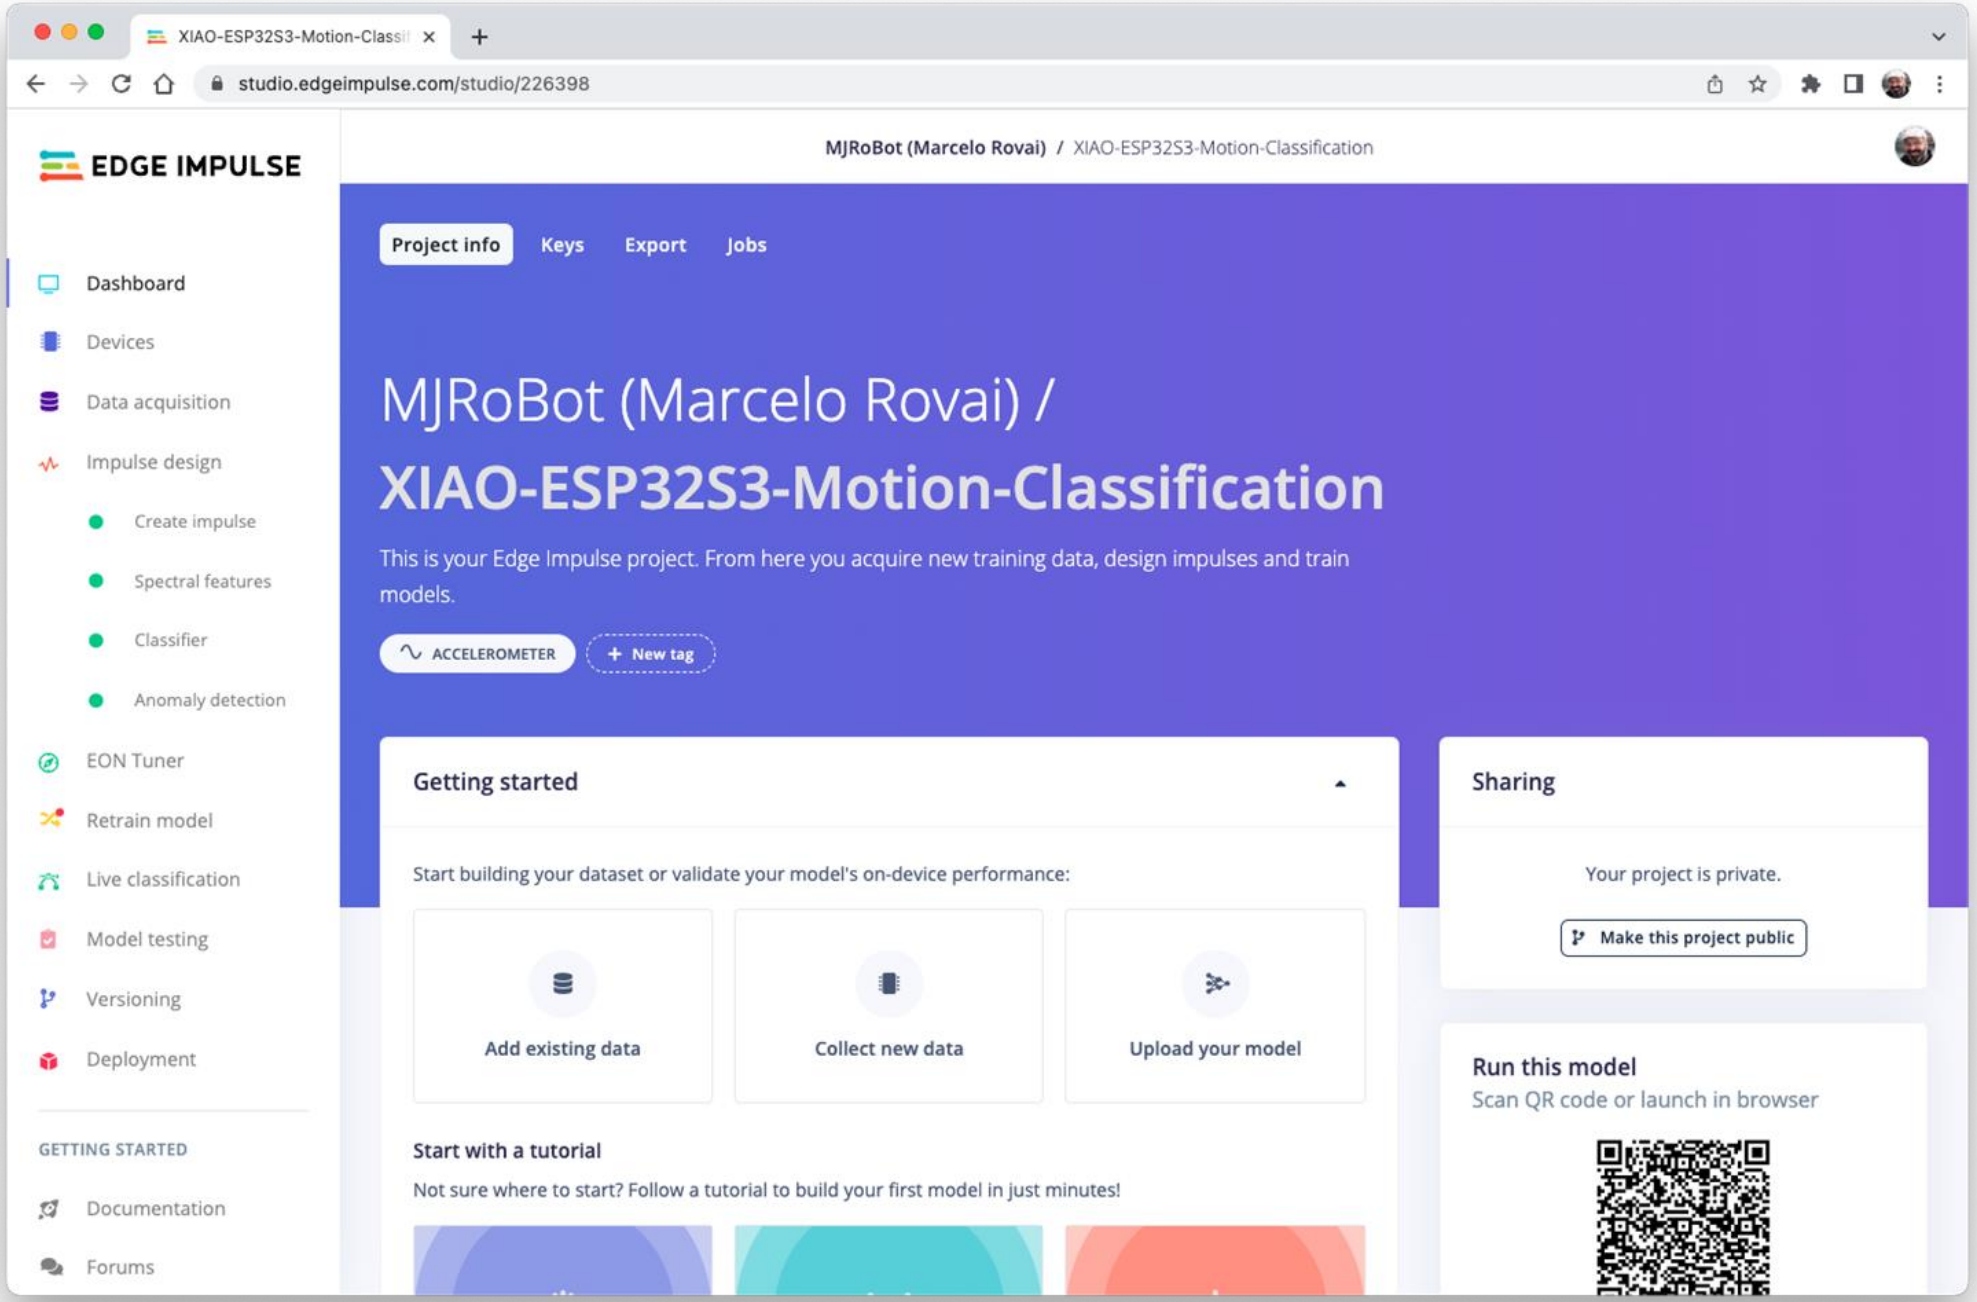

XIAO-ESP32S3-Motion-Classification (Edge Impulse)

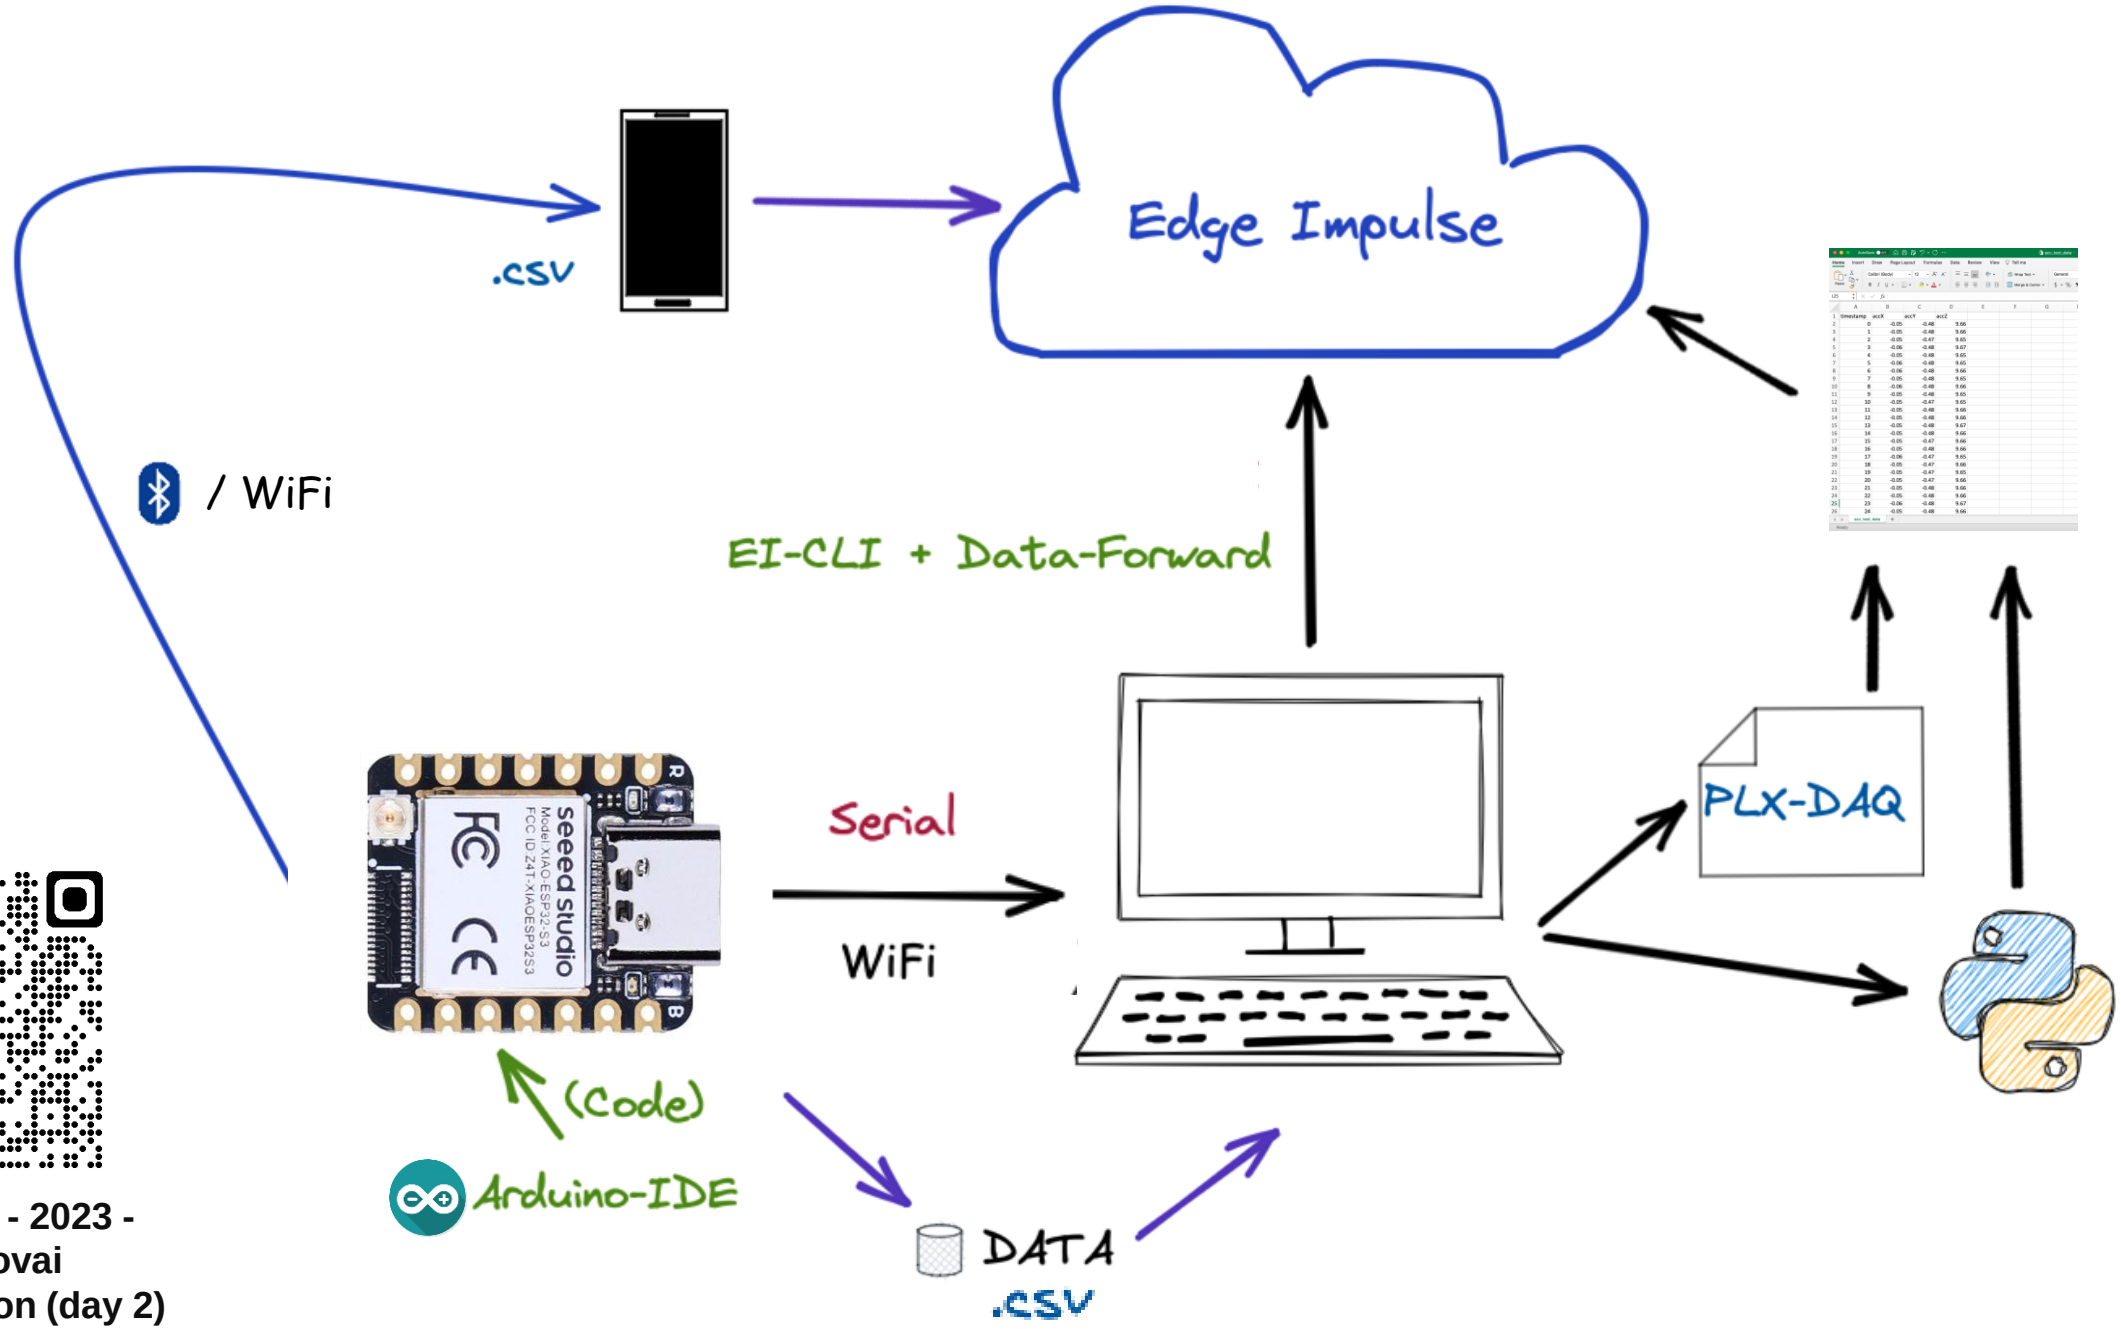

SciTinyML - 2023 -
Marcelo Rovai
Presentation (day 2)

TINY ML TALKS webcast

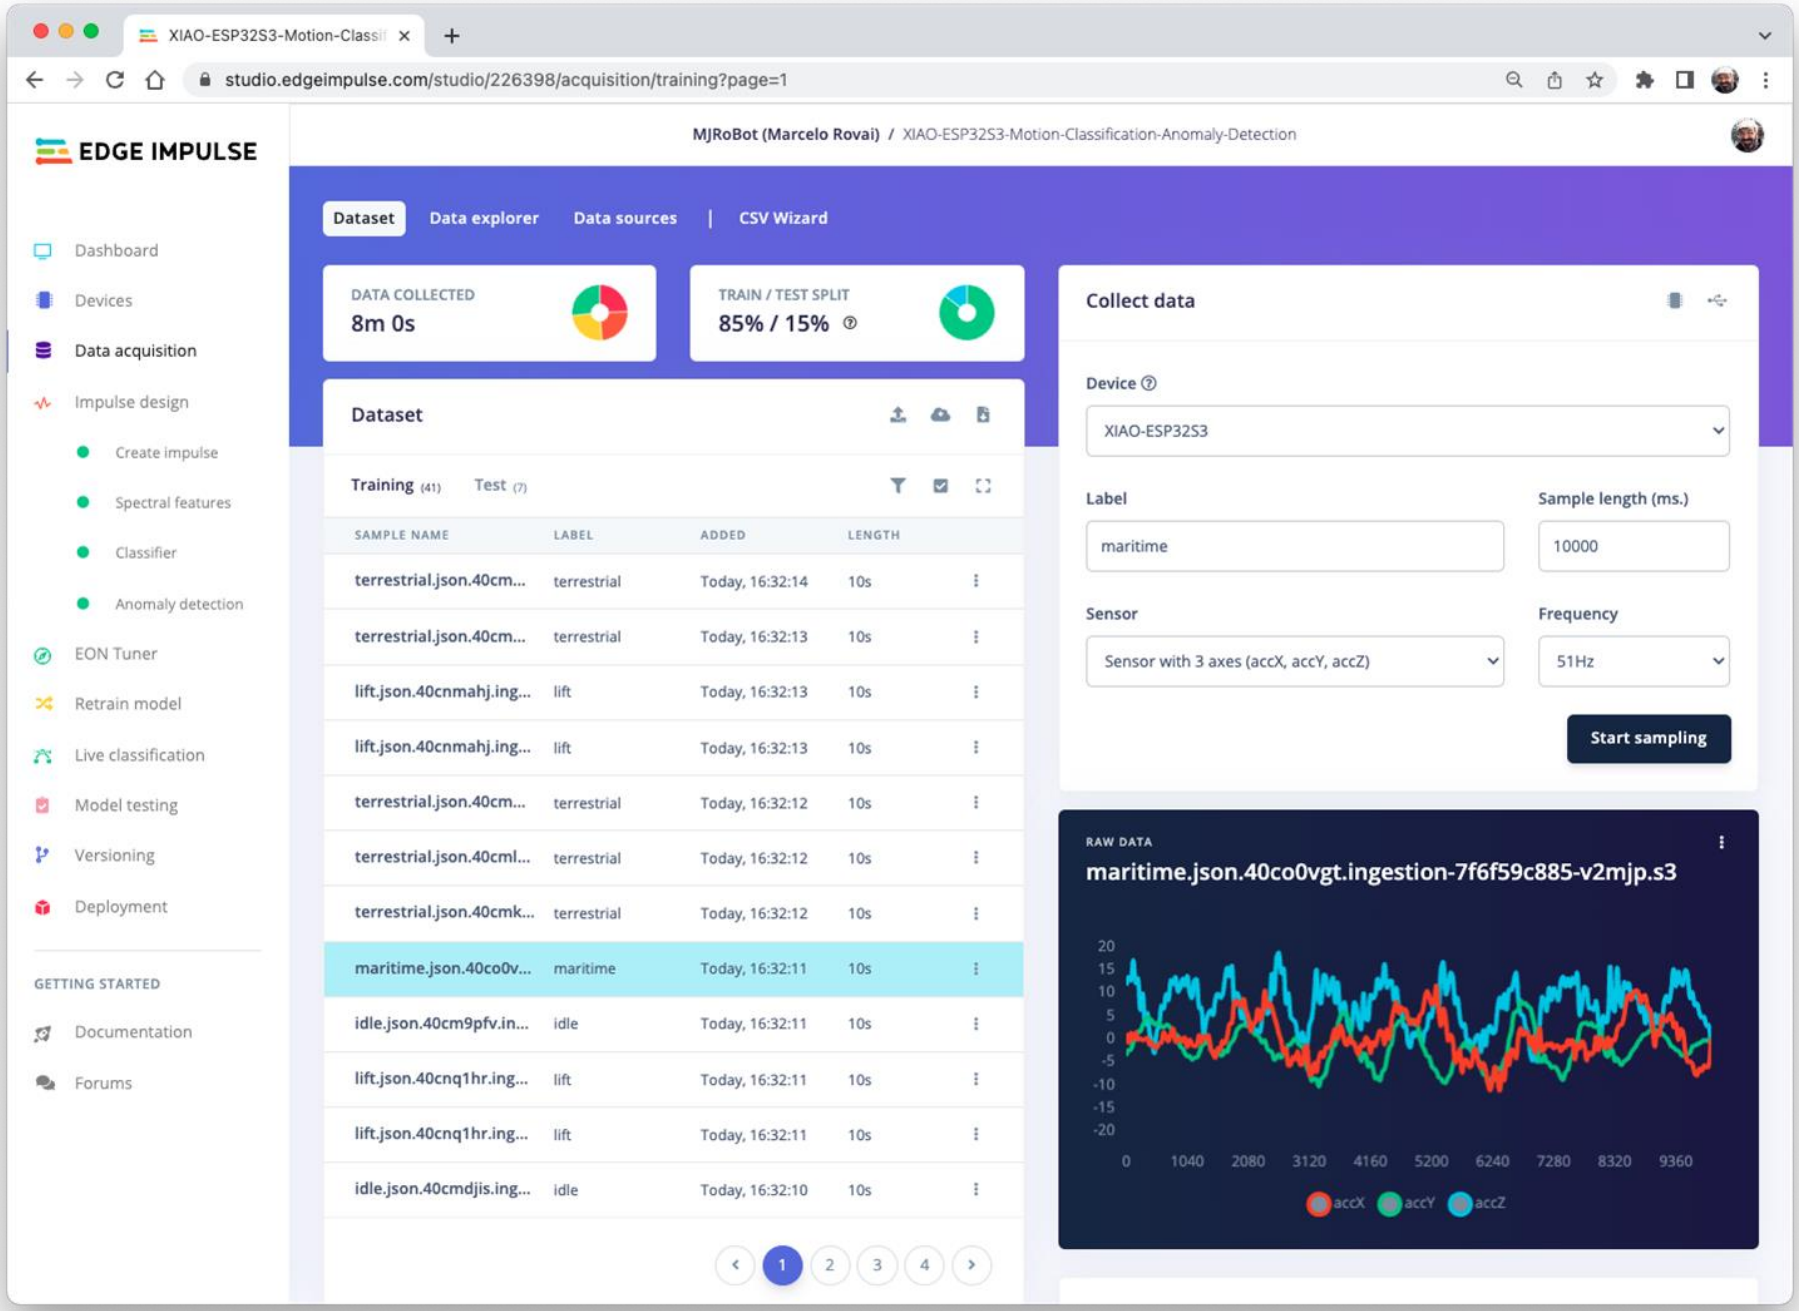

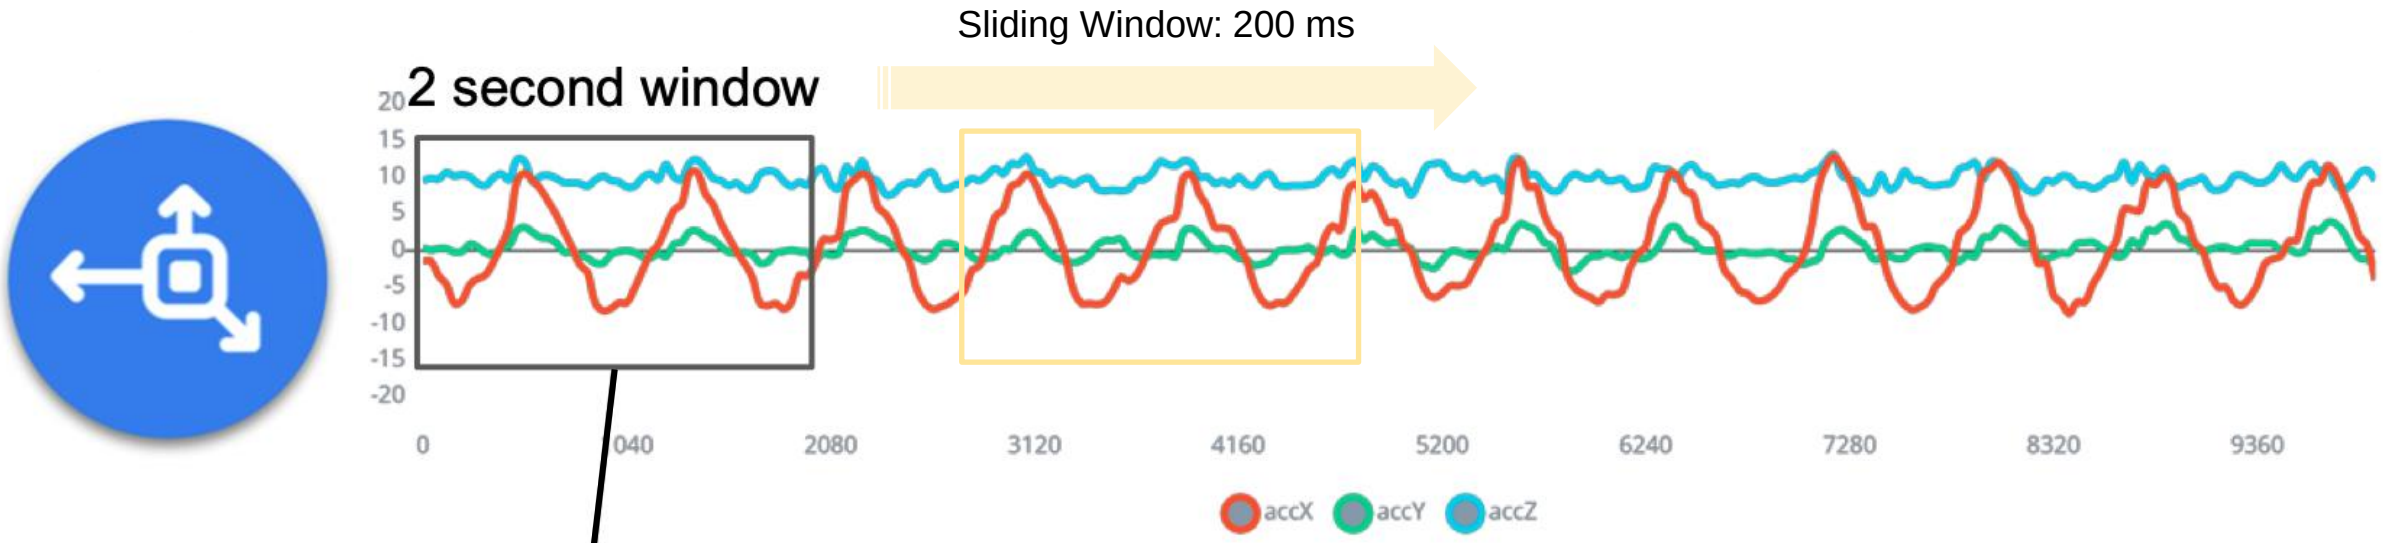

300 Raw Features

Manual Feature Extraction

Raw Data from sensor

Spectral Analysis

Features

- RMS
- SKEW
- KURT
- FFT
- PSD

NN Classifier

Classes

- Lift
- Terrestrial
- Maritime
- Idle

TinyML under the hood: Spectral Analysis

TINY ML TALKS webcast

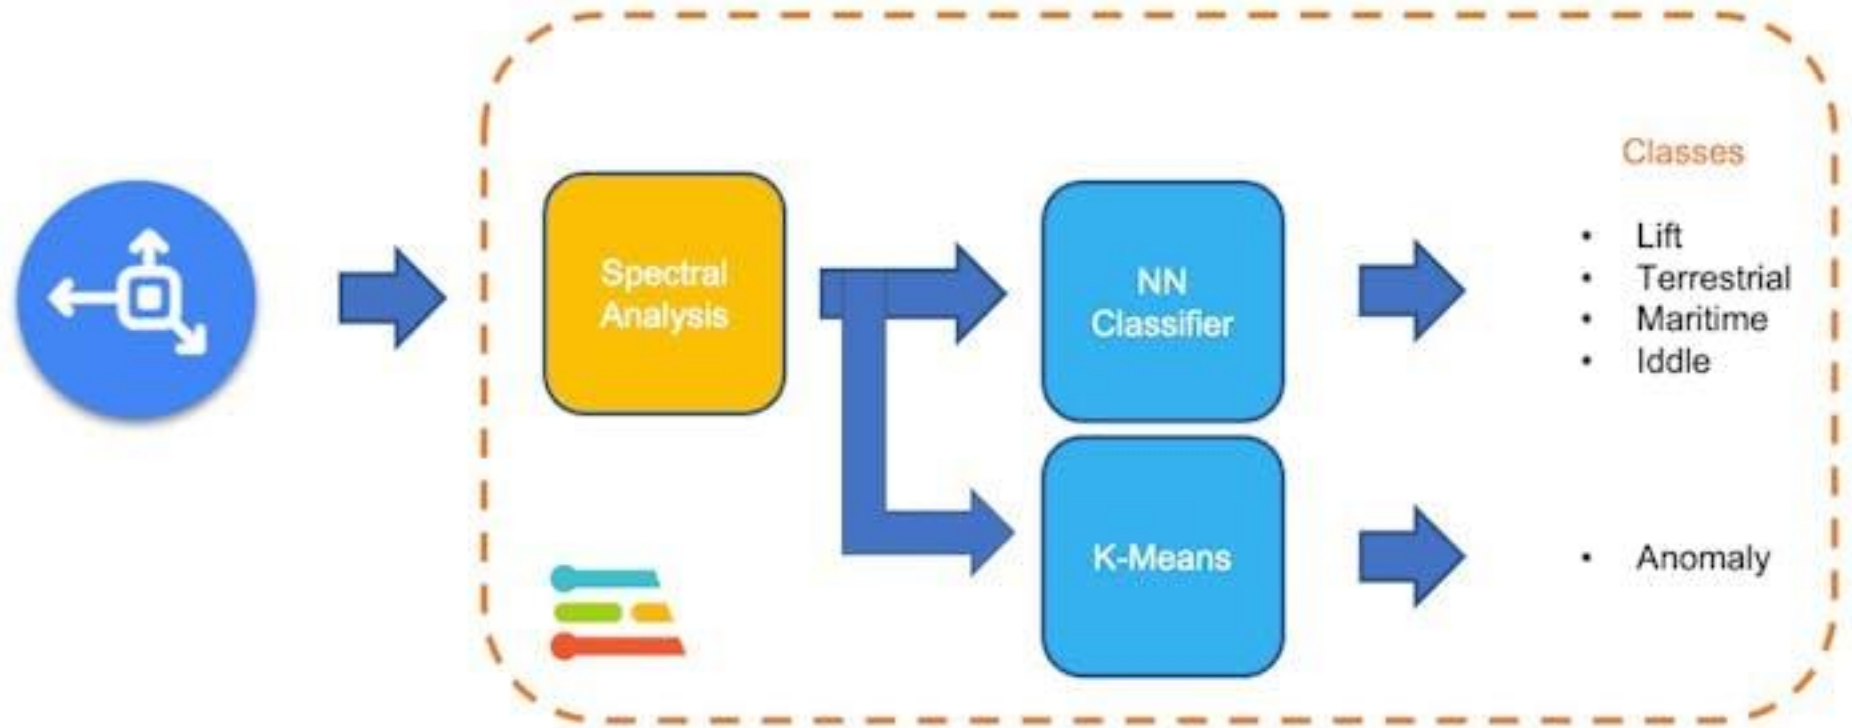

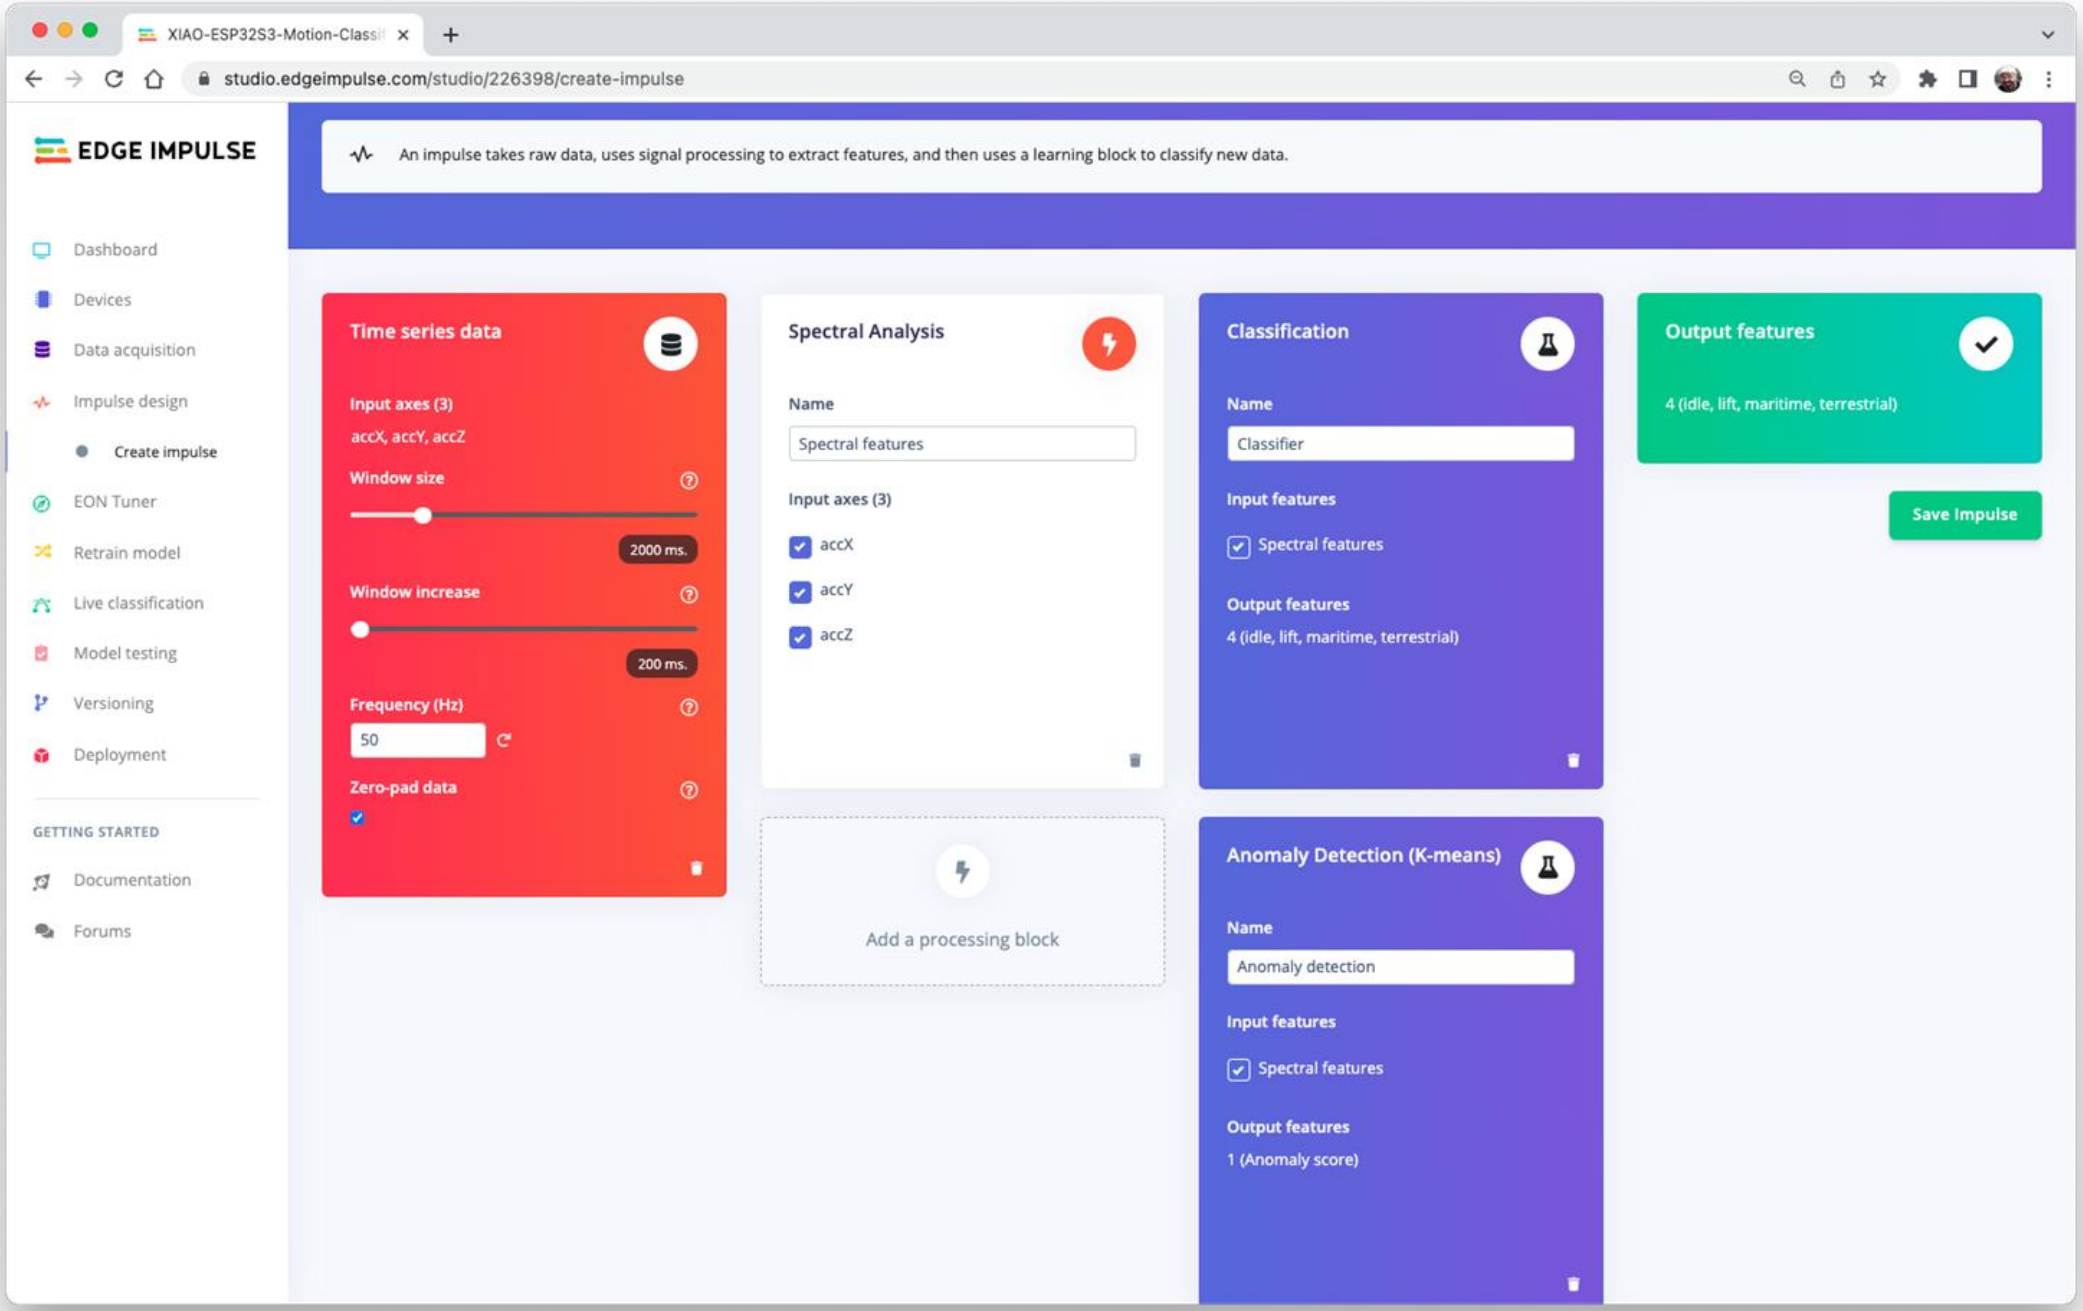
XIAO-ESP32S3-Motion-Classif
studio.edgeimpulse.com/studio/226398/create-impulse
EDGE IMPULSE
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Time series data
Input axes (3)
accX, accY, accZ
Window size
2000 ms.
Window increase
200 ms.
Frequency (Hz)
50
Zero-pad data
Spectral Analysis
Name
Spectral features
Input axes (3)
accX
accY
accZ
Add a processing block
Classification
Name
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)
Output features
4 (idle, lift, maritime, terrestrial)
Save Impulse

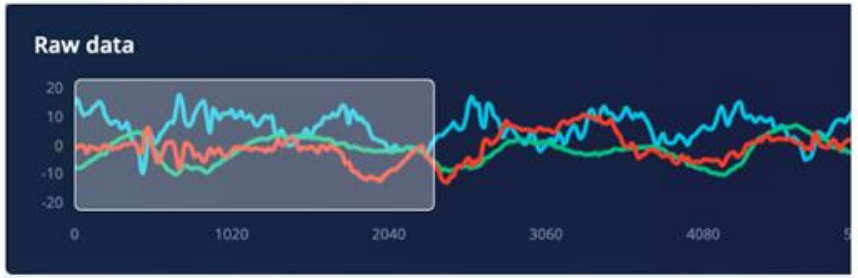

Raw features **300 Features**

```
0.2000, -7.6700, 14.3700, -0.3000, -7.3100, 16.4100, 0.3300, -7.3300, 13.1200, 0.0600, -6.8300,…
```

Parameters

Autotune parameters

Filter

Scale axes	0.03559985740718736
Input decimation ratio	1
Type	none

Analysis

Type	FFT
FFT length	32
Take log of spectrum?	☑
Overlap FFT frames?	☐
Improve low frequency resolution?	☐

Save parameters

DSP result

After filter

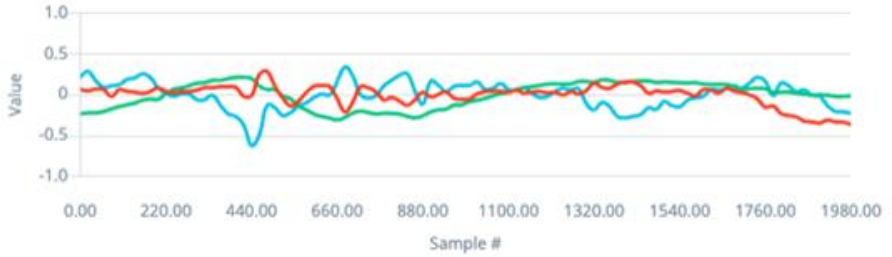

Spectral power (log)

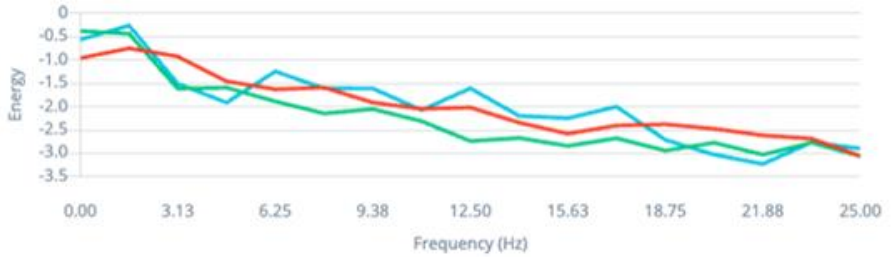

Processed features **63 Features**

```
0.1263, -1.2548, 1.5810, 1.8394, 2.0510, -0.7463, -0.9212, -1.4551, -1.6268, -1.5890, -1.9100, …
```

On-device performance

PROCESSING TIME
2 ms.

PEAK RAM USAGE
2 KB

Feature importance

All data

- accZ RMS
- accZ Spectral Power 0.78 - 2.34 Hz
- accX RMS
- accY Spectral Power 10.16 - 11.72 Hz
- accY Spectral Power 17.97 - 19.53 Hz
- accX Spectral Power 7.03 - 8.59 Hz
- accY RMS
- accX Spectral Power 0.78 - 2.34 Hz
- accZ Spectral Power 2.34 - 3.91 Hz
- accY Spectral Power 3.91 - 5.47 Hz
- accY Spectral Power 11.72 - 13.28 Hz
- accY Spectral Power 8.59 - 10.16 Hz
- accY Spectral Power 0.78 - 2.34 Hz
- accY Spectral Power 14.84 - 16.41 Hz
- accY Spectral Power 16.41 - 17.97 Hz

Feature explorer

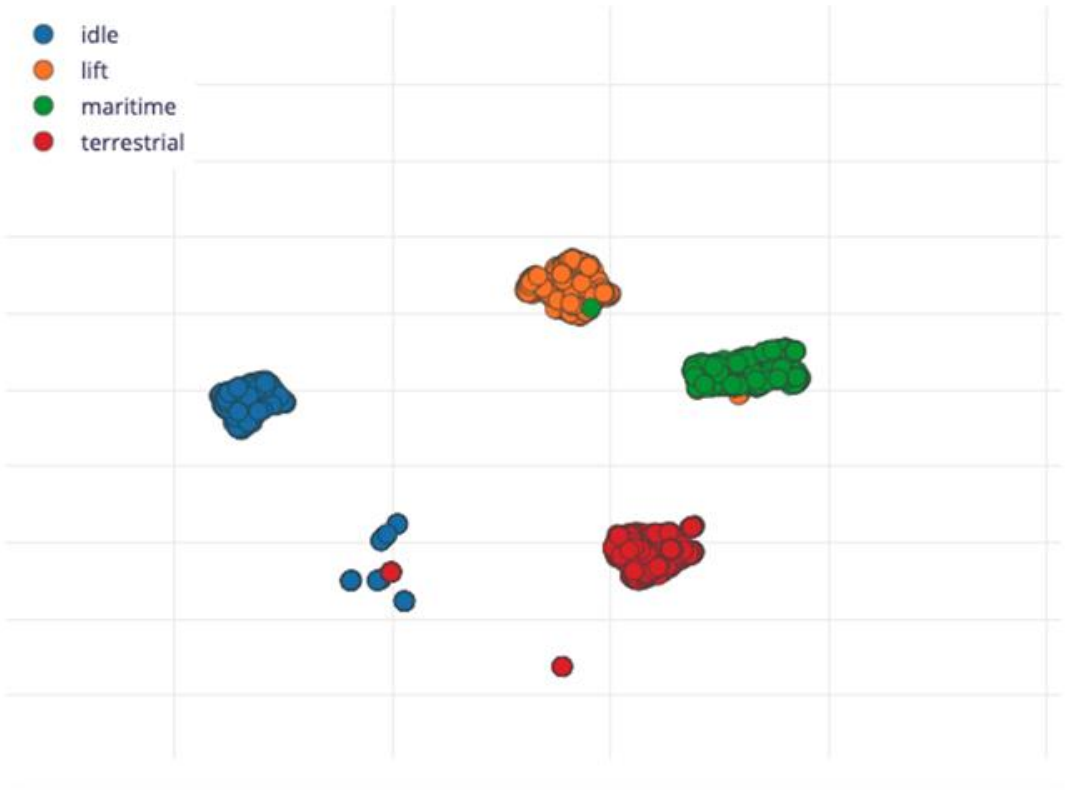

TINY ML TALKS webcast

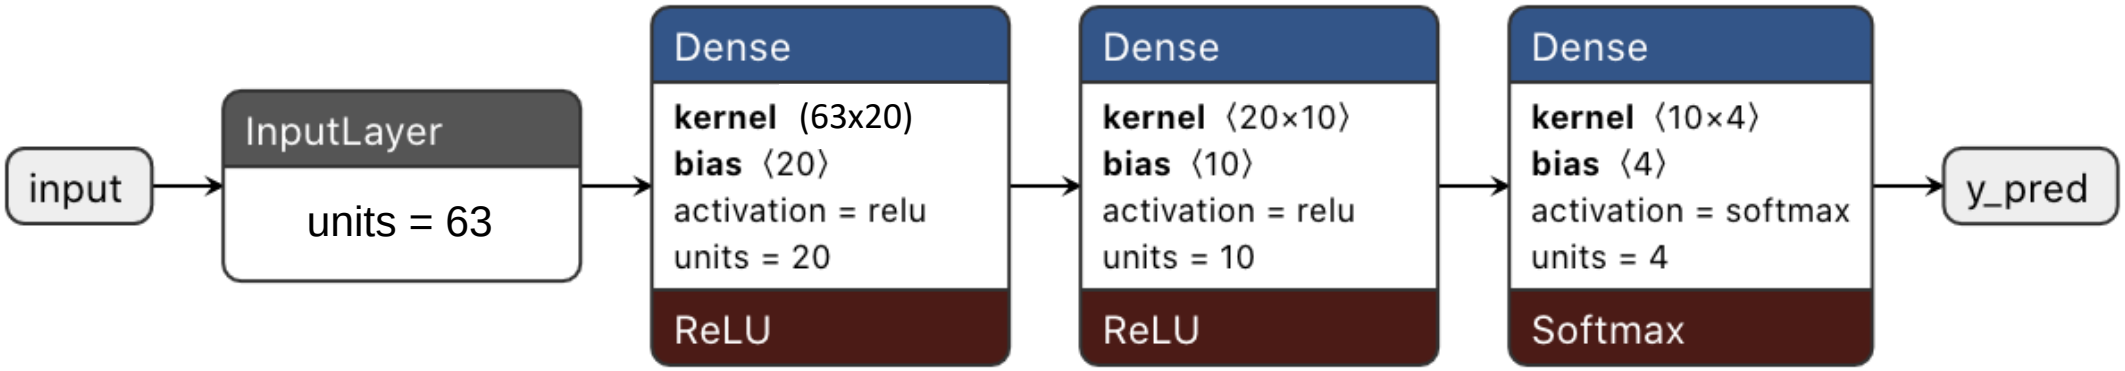

63 Features

- RMS
- SKEW
- KURT
- FFT
- PSD

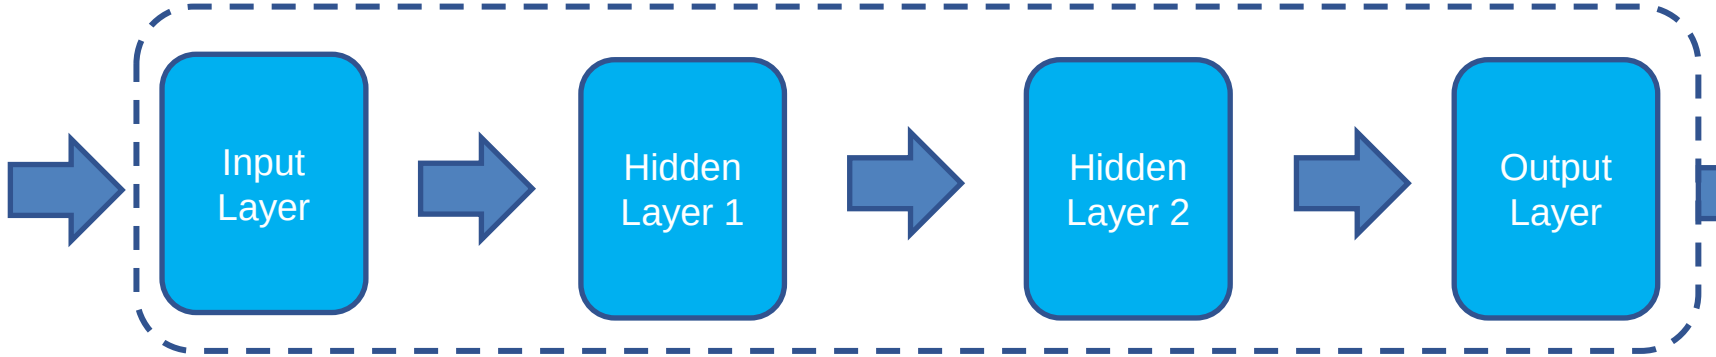

Classes

- Lift
- Terrestrial
- Maritime
- Idle

Neural Network settings

Training settings

Number of training cycles: 30

Learning rate: 0.0005

Advanced training settings

Validation set size: 20 %

Split train/validation set on metadata key:

Auto-balance dataset: ☐

Profile int8 model: ☑

Neural network architecture

Input layer (63 features)

Dense layer (20 neurons)

Dense layer (10 neurons)

Add an extra layer

Output layer (4 classes)

Start training

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY
97.0%

LOSS
0.28

Confusion matrix (validation set)

	IDLE	LIFT	MARITIME	TERRESTRIAL
IDLE	100%	0%	0%	0%
LIFT	0%	98.7%	1.3%	0%
MARITIME	0%	3.5%	96.5%	0%
TERRESTRIAL	4.4%	1.1%	1.1%	93.4%
F1 SCORE	0.98	0.97	0.97	0.97

Data explorer (full training set)

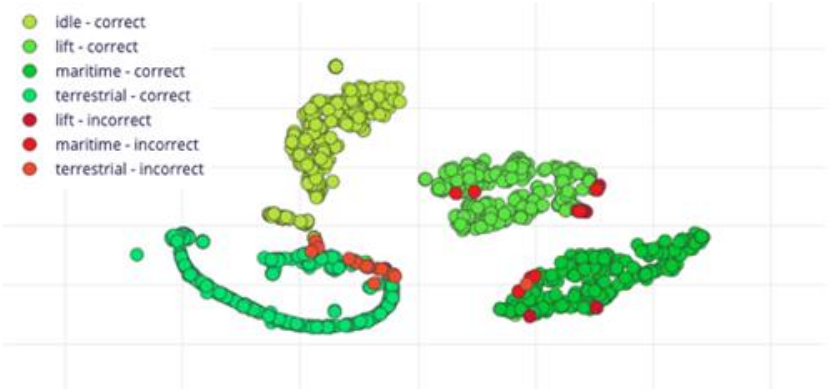

On-device performance

INFERENCING TIME
1 ms.

PEAK RAM USAGE
1.3K

FLASH USAGE
15.3K

Anomaly detection settings

Cluster count

32

Axes

★ Select suggested axes

- [x] accX RMS
- [ ] accX Skewness
- [ ] accX Kurtosis
- [ ] accX Spectral Skewness
- [ ] accX Spectral Kurtosis
- [ ] accX Spectral Power 0.78 - 2.34 Hz
- [ ] accX Spectral Power 2.34 - 3.91 Hz
- [ ] accX Spectral Power 3.91 - 5.47 Hz
- [ ] accX Spectral Power 5.47 - 7.03 Hz
- [ ] accX Spectral Power 7.03 - 8.59 Hz
- [ ] accX Spectral Power 8.59 - 10.16 Hz
- [ ] accX Spectral Power 10.16 - 11.72 Hz
- [ ] accX Spectral Power 11.72 - 13.28 Hz
- [ ] accX Spectral Power 13.28 - 14.84 Hz
- [ ] accX Spectral Power 14.84 - 16.41 Hz
- [ ] accX Spectral Power 16.41 - 17.97 Hz
- [ ] accX Spectral Power 17.97 - 19.53 Hz
- [ ] accX Spectral Power 19.53 - 21.09 Hz
- [ ] accX Spectral Power 21.09 - 22.66 Hz
- [ ] accX Spectral Power 22.66 - 24.22 Hz
- [ ] accX Spectral Power 24.22 - 25.78 Hz
- [ ] accY RMS
- [ ] accY Skewness
- [ ] accY Kurtosis
- [ ] accY Spectral Skewness
- [ ] accY Spectral Kurtosis
- [ ] accY Spectral Power 0.78 - 2.34 Hz
- [ ] accY Spectral Power 2.34 - 3.91 Hz
- [ ] accY Spectral Power 3.91 - 5.47 Hz
- [ ] accY Spectral Power 5.47 - 7.03 Hz
- [ ] accY Spectral Power 7.03 - 8.59 Hz
- [ ] accY Spectral Power 8.59 - 10.16 Hz
- [x] accY Spectral Power 10.16 - 11.72 Hz
- [ ] accY Spectral Power 11.72 - 13.28 Hz
- [ ] accY Spectral Power 13.28 - 14.84 Hz
- [ ] accY Spectral Power 14.84 - 16.41 Hz
- [ ] accY Spectral Power 16.41 - 17.97 Hz
- [ ] accY Spectral Power 17.97 - 19.53 Hz
- [ ] accY Spectral Power 19.53 - 21.09 Hz
- [ ] accY Spectral Power 21.09 - 22.66 Hz
- [ ] accY Spectral Power 22.66 - 24.22 Hz
- [ ] accY Spectral Power 24.22 - 25.78 Hz
- [x] accZ RMS
- [ ] accZ Skewness
- [ ] accZ Kurtosis
- [ ] accZ Spectral Skewness
- [ ] accZ Spectral Kurtosis
- [x] accZ Spectral Power 0.78 - 2.34 Hz
- [ ] accZ Spectral Power 2.34 - 3.91 Hz
- [ ] accZ Spectral Power 3.91 - 5.47 Hz
- [ ] accZ Spectral Power 5.47 - 7.03 Hz
- [ ] accZ Spectral Power 7.03 - 8.59 Hz
- [ ] accZ Spectral Power 8.59 - 10.16 Hz
- [ ] accZ Spectral Power 10.16 - 11.72 Hz
- [ ] accZ Spectral Power 11.72 - 13.28 Hz
- [ ] accZ Spectral Power 13.28 - 14.84 Hz
- [ ] accZ Spectral Power 14.84 - 16.41 Hz
- [ ] accZ Spectral Power 16.41 - 17.97 Hz
- [ ] accZ Spectral Power 17.97 - 19.53 Hz
- [ ] accZ Spectral Power 19.53 - 21.09 Hz
- [ ] accZ Spectral Power 21.09 - 22.66 Hz
- [ ] accZ Spectral Power 22.66 - 24.22 Hz
- [ ] accZ Spectral Power 24.22 - 25.78 Hz

Start training

Anomaly explorer (1,681 samples)

X Axis: accX RMS

Y Axis: accY Spectral Power 1

Test data: -- No test data

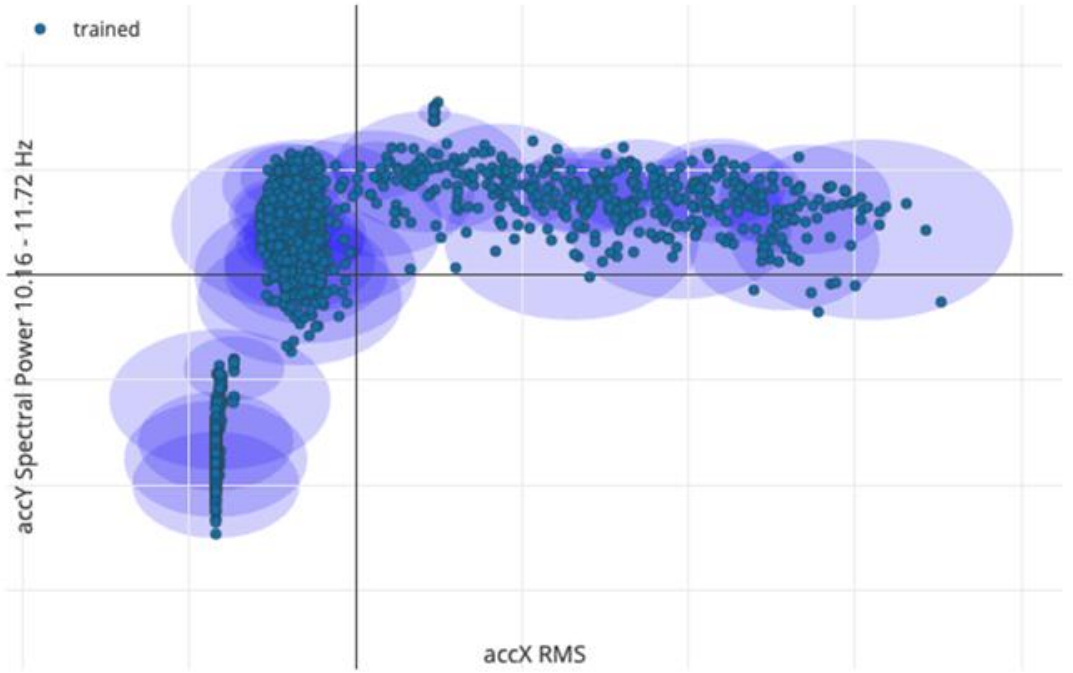

Configure your deployment

You can deploy your impulse to any device. This makes the model run without an internet connection, minimizes latency, and runs with minimal power consumption. Read more.

Arduino library ×

SELECTED DEPLOYMENT

Arduino library

An Arduino library with examples that runs on most Arm-based Arduino development boards.

MODEL OPTIMIZATIONS

Model optimizations can increase on-device performance but may reduce accuracy.

Enable EON™ Compiler *Same accuracy, up to 50% less memory. Open source. Learn more*

Quantized (int8) ⭐

Selected ✓

	SPECTRAL FEATU...	CLASSIFIER	TOTAL
LATENCY	2 ms.	1 ms.	**3 ms.**
RAM	1.7K	1.3K	**1.7K**
FLASH	-	15.3K	-
ACCURACY			**96.86%**

Unoptimized (float32)

Select

	SPECTRAL FEATU...	CLASSIFIER	TOTAL
LATENCY	2 ms.	9 ms.	**11 ms.**
RAM	1.7K	1.4K	**1.7K**
FLASH	-	15.5K	-
ACCURACY			**97.21%**

Estimate for Cortex-M4F 80MHz - Change target

Build

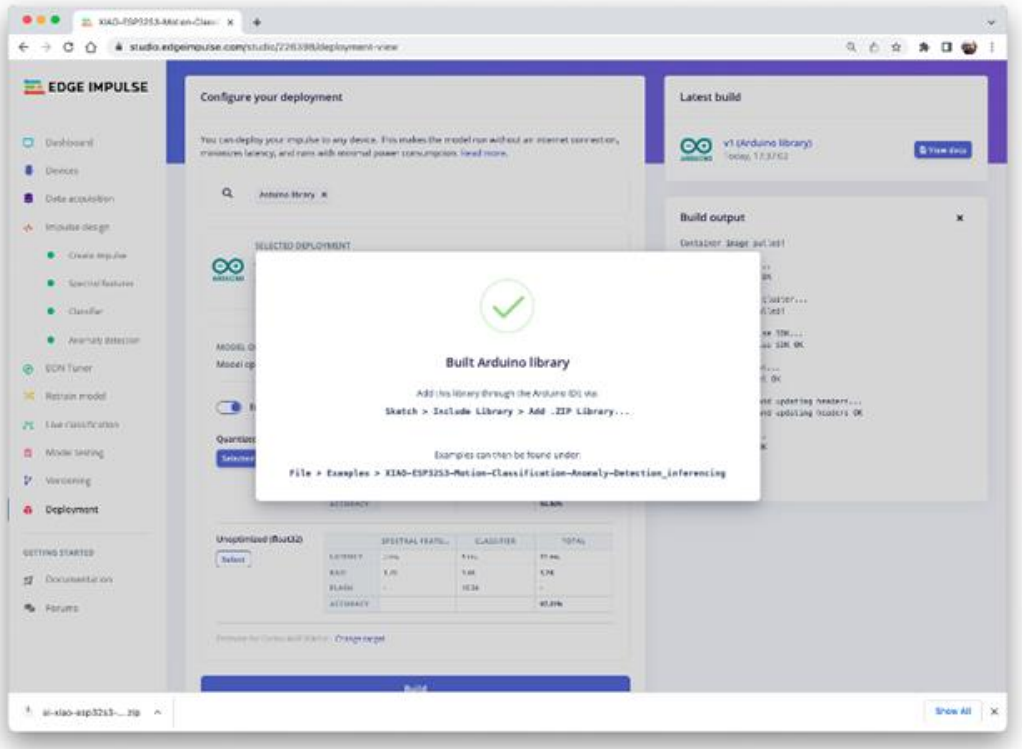

/dev/cu.usbmodem1101

Send

```
09:28:30.557 -> Sampling...
09:28:32.559 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:28:32.559 ->     idle: 0.14844
09:28:32.559 ->     lift: 0.18359
09:28:32.559 ->     maritime: 0.20312
09:28:32.559 ->     terrestrial: 0.46484
09:28:32.559 ->     anomaly score: -0.123
09:28:32.559 ->
09:28:32.559 -> Starting inferencing in 2 seconds...
09:28:34.562 -> Sampling...
09:28:36.567 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:28:36.567 ->     idle: 0.16016
09:28:36.567 ->     lift: 0.17969
09:28:36.567 ->     maritime: 0.19922
09:28:36.567 ->     terrestrial: 0.45703
09:28:36.567 ->     anomaly score: -0.107
09:28:36.567 ->
09:28:36.567 -> Starting inferencing in 2 seconds...
```

Autoscroll Show timestamp Both NL & CR 115200 baud Clear output

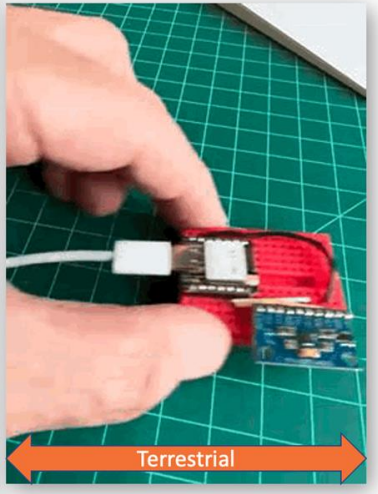

/dev/cu.usbmodem1101

Send

```
09:29:04.641 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:29:04.641 ->     idle: 0.00000
09:29:04.641 ->     lift: 0.02734
09:29:04.641 ->     maritime: 0.96875
09:29:04.641 ->     terrestrial: 0.00391
09:29:04.641 ->     anomaly score: 0.989
09:29:04.641 ->
09:29:04.641 -> Starting inferencing in 2 seconds...
09:29:06.628 -> Sampling...
09:29:08.690 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:29:08.690 ->     idle: 0.00000
09:29:08.690 ->     lift: 0.03906
09:29:08.690 ->     maritime: 0.92578
09:29:08.690 ->     terrestrial: 0.03516
09:29:08.690 ->     anomaly score: 0.697
09:29:08.690 ->
09:29:08.690 -> Starting inferencing in 2 seconds...
09:29:10.706 -> Sampling...
```

Autoscroll Show timestamp Both NL & CR 115200 baud Clear output

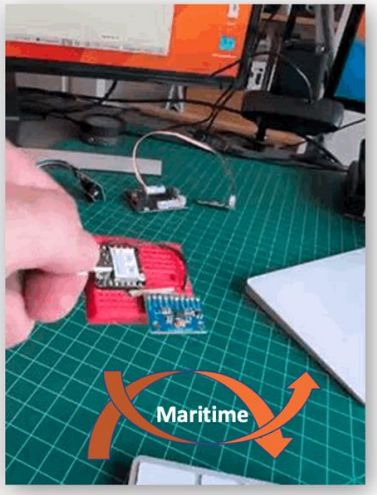

/dev/cu.usbmodem1101

Send

```
09:26:08.258 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:26:08.258 ->     idle: 0.98828
09:26:08.258 ->     lift: 0.00781
09:26:08.258 ->     maritime: 0.00000
09:26:08.258 ->     terrestrial: 0.00000
09:26:08.258 ->     anomaly score: -0.273
09:26:08.258 ->
09:26:08.258 -> Starting inferencing in 2 seconds...
09:26:10.230 -> Sampling...
09:26:12.270 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:26:12.270 ->     idle: 0.99219
09:26:12.270 ->     lift: 0.00391
09:26:12.270 ->     maritime: 0.00000
09:26:12.270 ->     terrestrial: 0.00391
09:26:12.270 ->     anomaly score: -0.345
09:26:12.270 ->
09:26:12.270 -> Starting inferencing in 2 seconds...
09:26:14.262 -> Sampling...
```

Autoscroll Show timestamp Both NL & CR 115200 baud Clear output

/dev/cu.usbmodem1101

Send

```
09:27:36.424 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:27:36.424 ->     idle: 0.00000
09:27:36.424 ->     lift: 0.98828
09:27:36.424 ->     maritime: 0.01172
09:27:36.424 ->     terrestrial: 0.00000
09:27:36.424 ->     anomaly score: -0.093
09:27:36.424 ->
09:27:36.424 -> Starting inferencing in 2 seconds...
09:27:38.432 -> Sampling...
09:27:40.446 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:27:40.446 ->     idle: 0.00000
09:27:40.446 ->     lift: 0.98828
09:27:40.446 ->     maritime: 0.01172
09:27:40.446 ->     terrestrial: 0.00000
09:27:40.446 ->     anomaly score: -0.203
09:27:40.446 ->
09:27:40.446 -> Starting inferencing in 2 seconds...
09:27:42.442 -> Sampling...
```

Autoscroll Show timestamp Both NL & CR 115200 baud Clear output

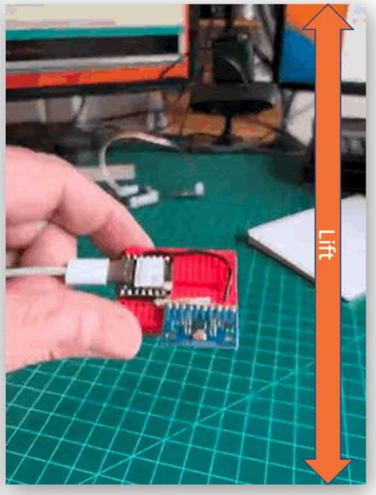

TINY ML TALKS webcast

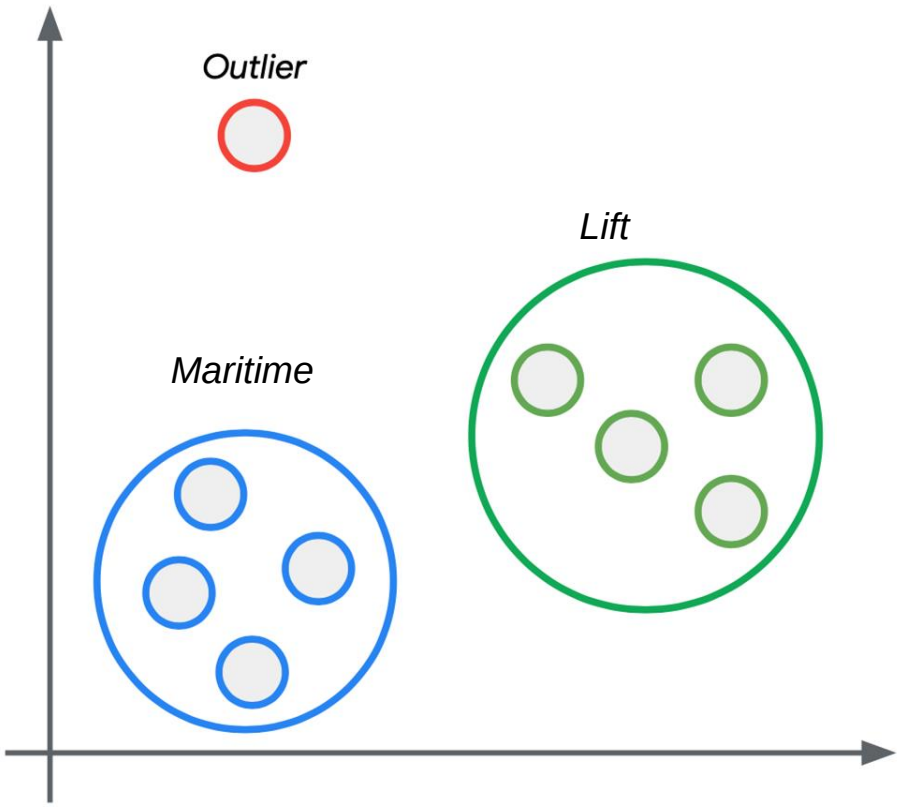


```
/dev/cu.usbmodem1101
Send
09:30:30.876 -> Sampling...
09:30:32.872 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:30:32.872 ->     idle: 0.00000
09:30:32.872 ->     lift: 0.05078
09:30:32.872 ->     maritime: 0.94922
09:30:32.872 ->     terrestrial: 0.00000
09:30:32.872 ->     anomaly score: 1.736
09:30:32.872 ->
09:30:32.872 -> Starting inferencing in 2 seconds...
09:30:34.895 -> Sampling...
09:30:36.881 -> Predictions (DSP: 7 ms., Classification: 0 ms., Anomaly: 0
09:30:36.881 ->     idle: 0.00000
09:30:36.881 ->     lift: 0.07031
09:30:36.881 ->     maritime: 0.92578
09:30:36.881 ->     terrestrial: 0.00391
09:30:36.881 ->     anomaly score: 3.605
09:30:36.881 ->
09:30:36.881 -> Starting inferencing in 2 seconds...
Autoscroll  Show timestamp  Both NL & CR  115200 baud  Clear output
```

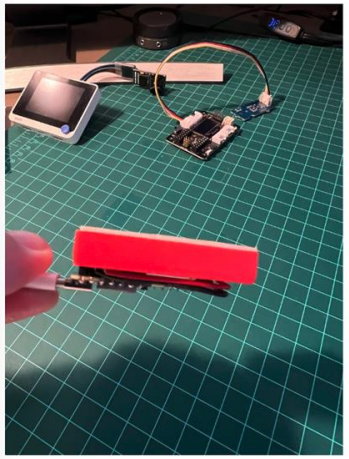

TINY ML TALKS webcast

Sound

Vibration

Vision

TINY ML TALKS webcast

Sound

TinyML Made Easy: KeyWord Spotting (KWS)

MJRoBot (Marcelo Rovai)

Vibration

Vision

TINY ML TALKS webcast

Personal Assistant

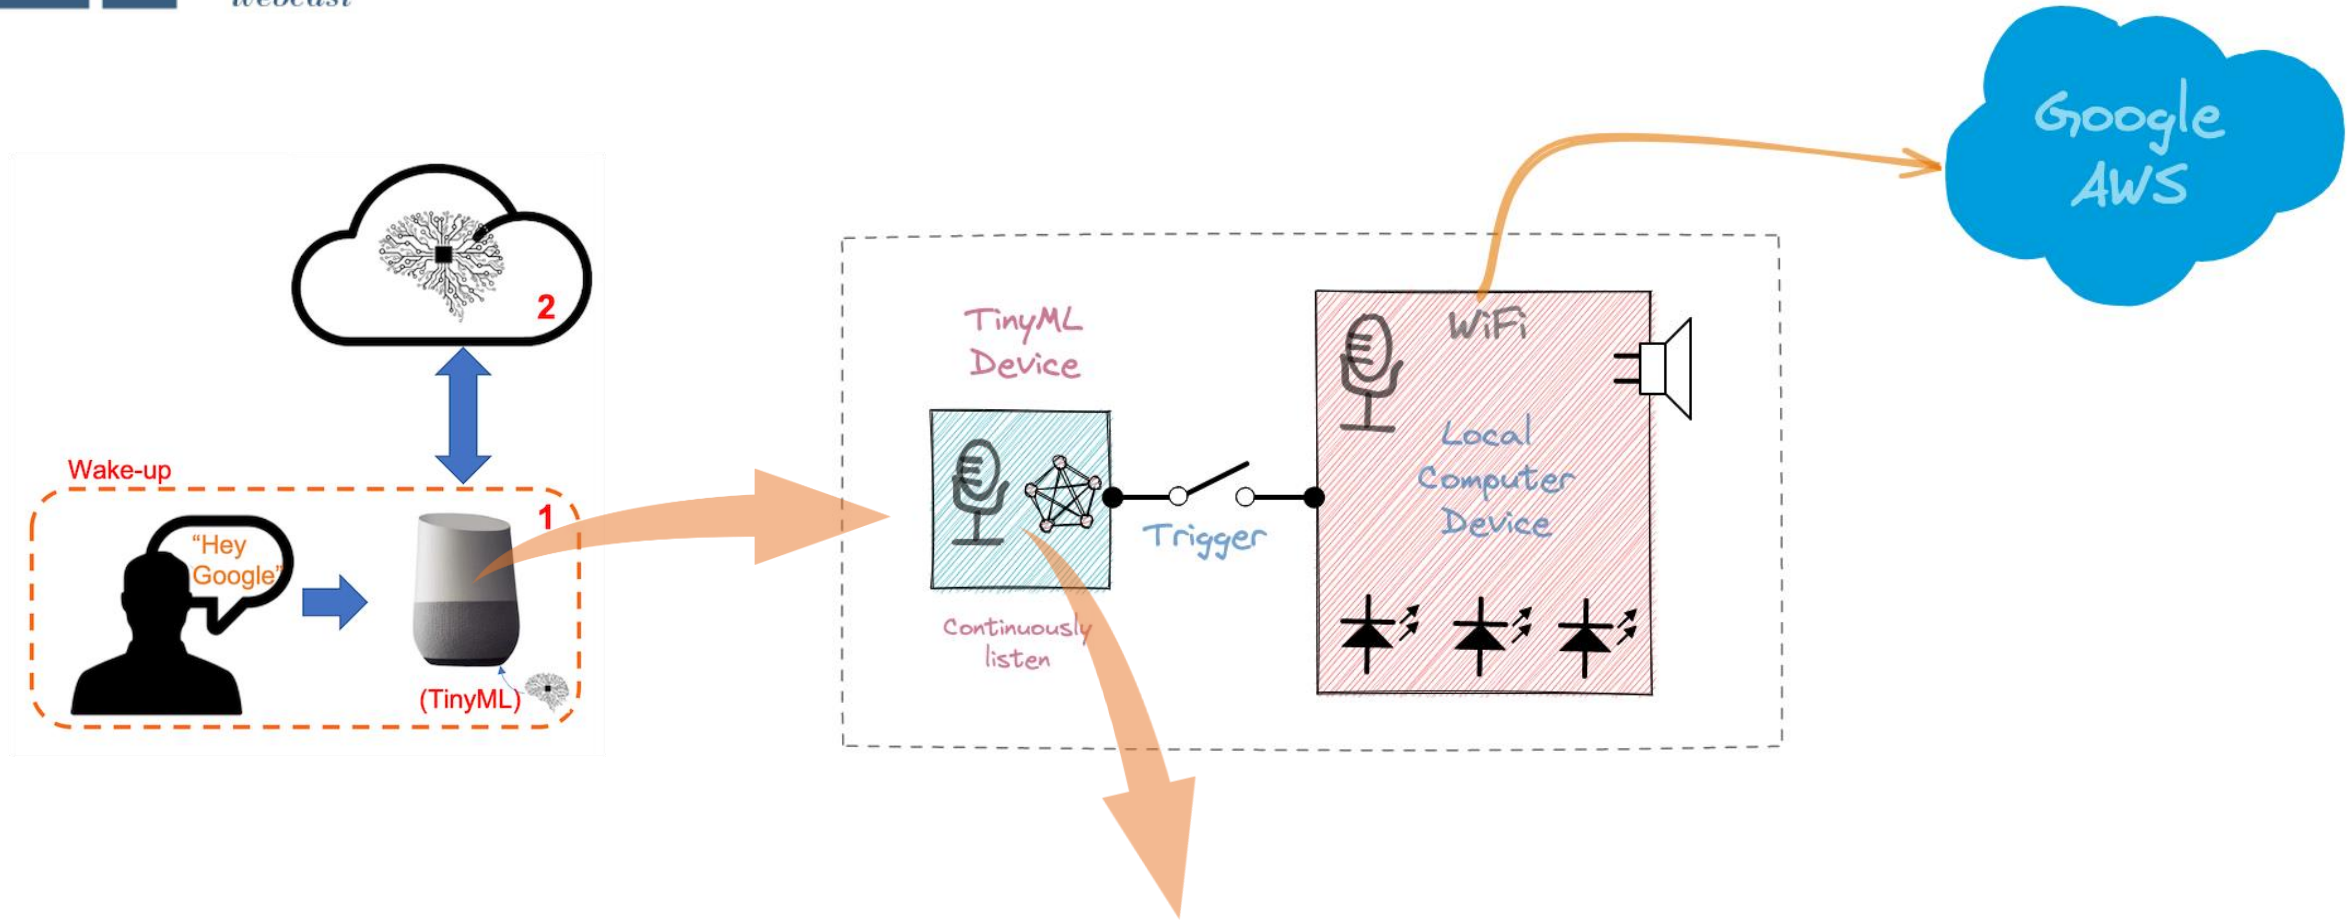

Keyword Spotting (KWS) - Inference

Obtains an input

Pre-Process

Runs model

Post-Processes

Make things happen

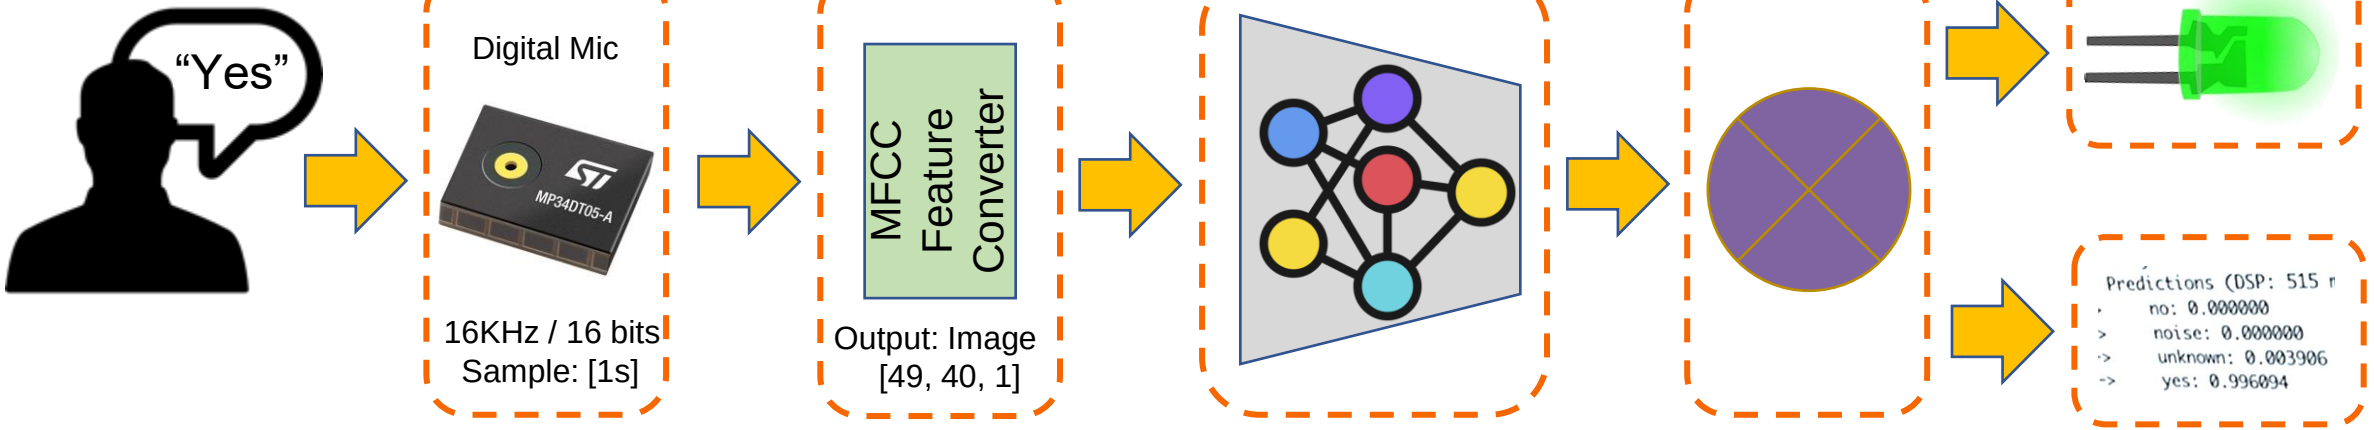

Output Dim [1, 4]
- Prob 'Noise'
- Prob 'Unknown'
- Prob 'Yes'
- Prob 'No'

If Probability of YES is greater than 80%
Take actions

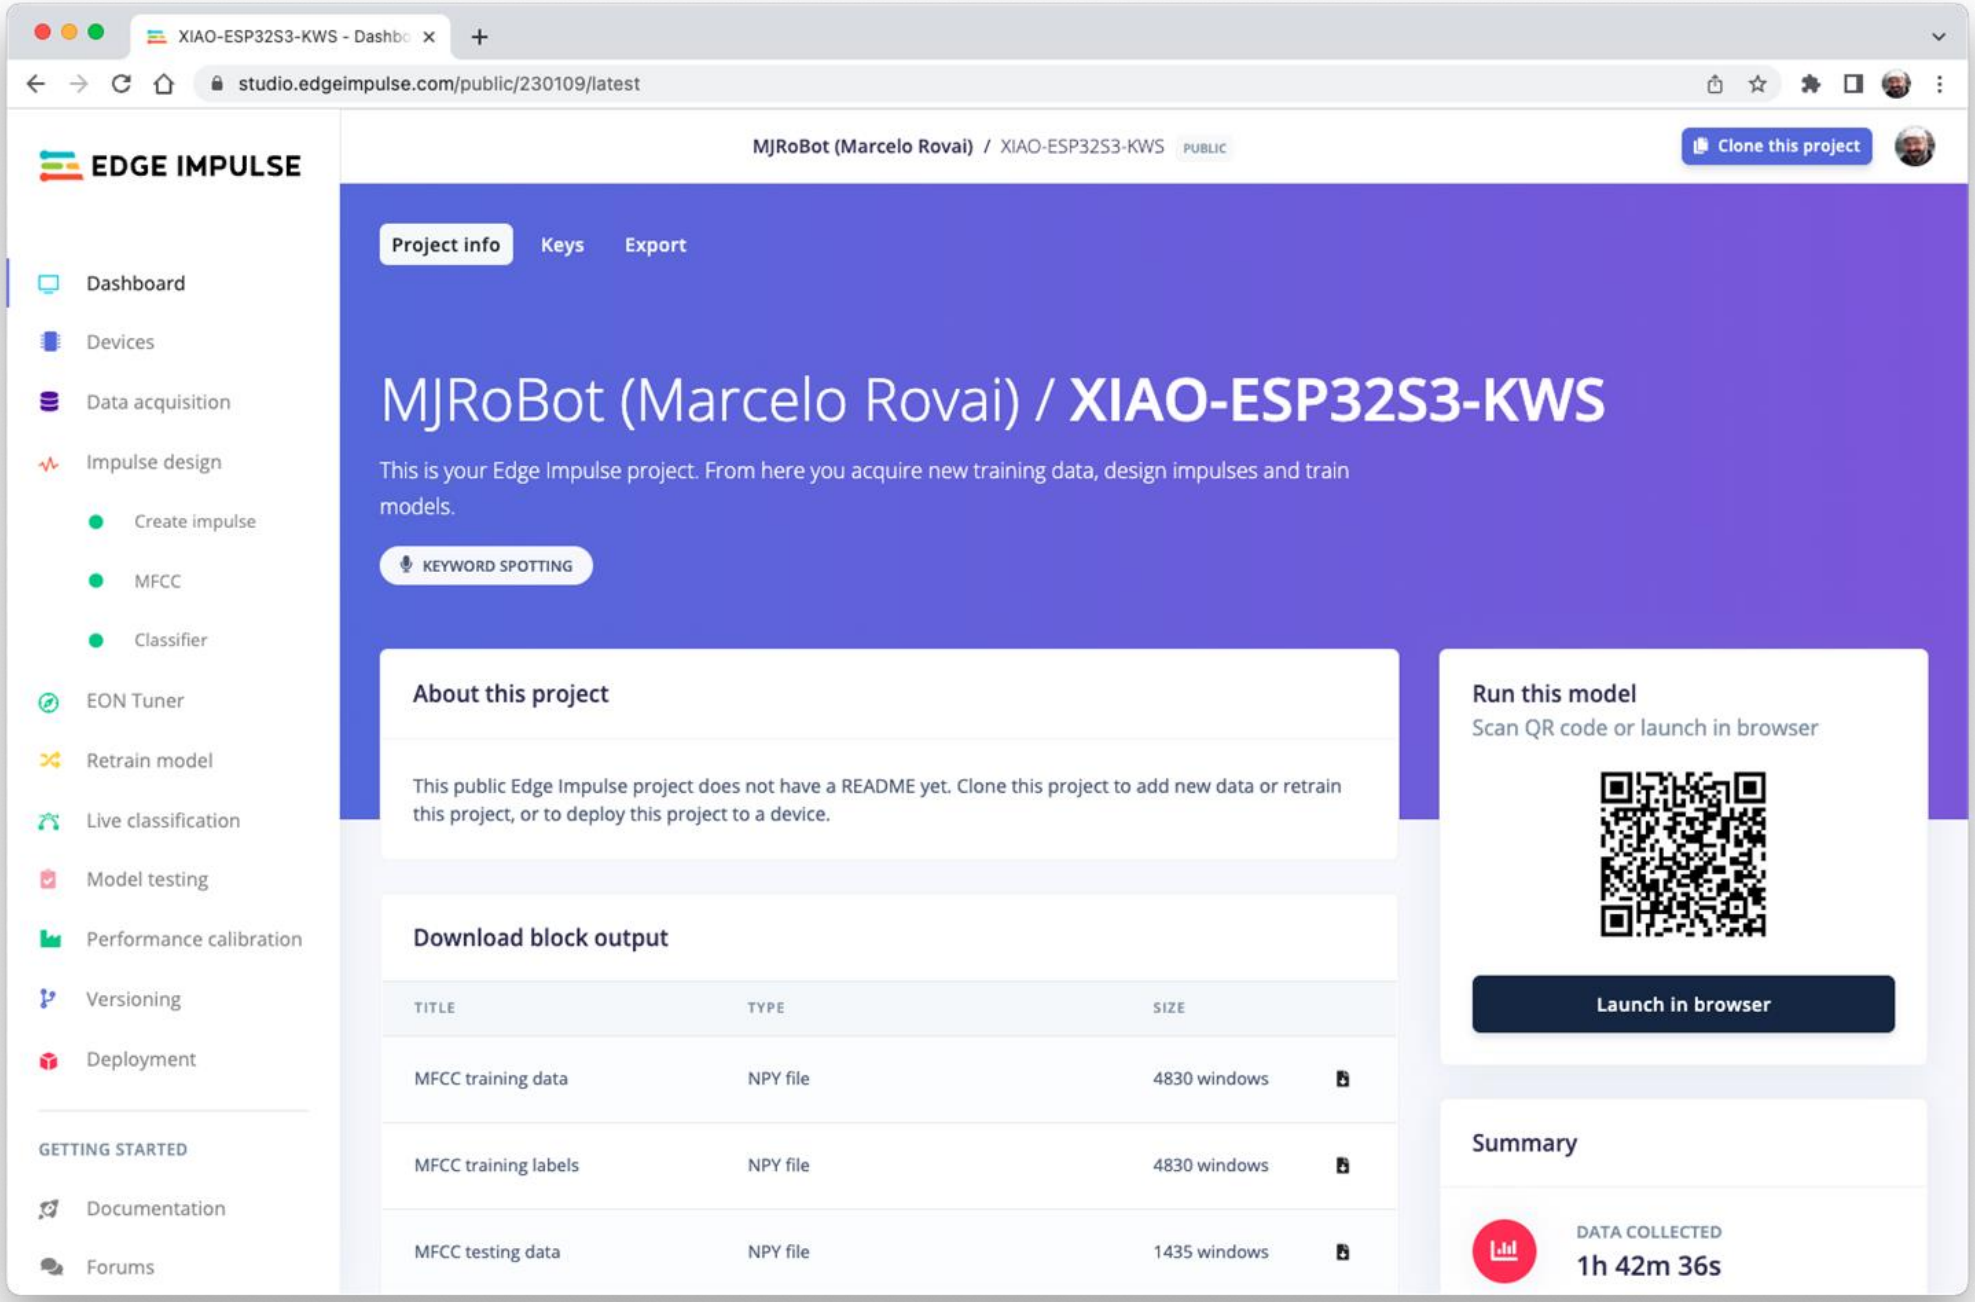

XIAO-ESP32S3-KWS (Edge Impulse)

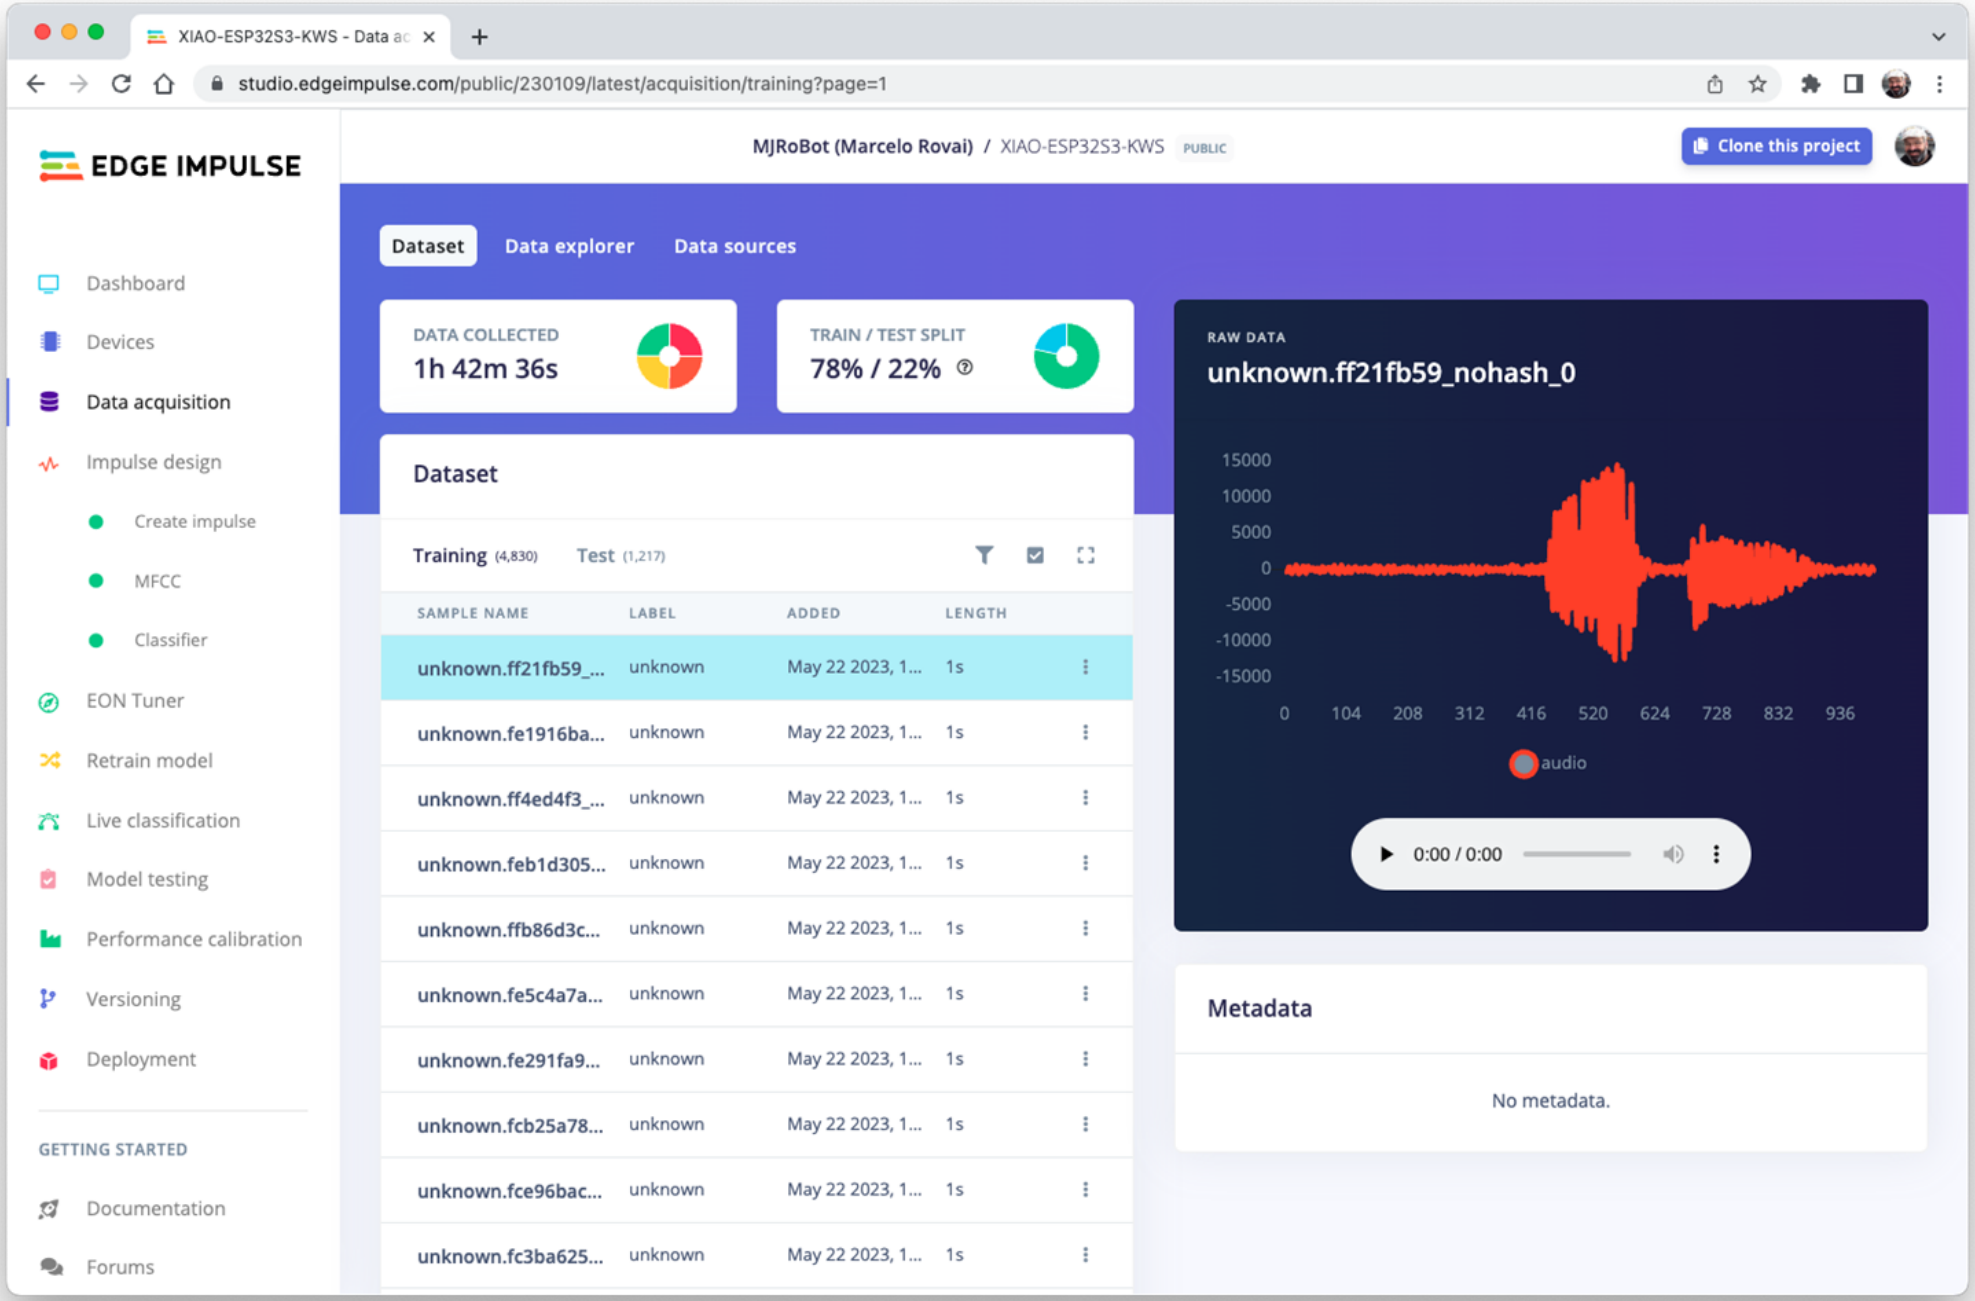

Speech Commands Dataset (reduced set)

Pre-Processing (MFCC)

1 second sample@16KHz raw data -> 16,000 features

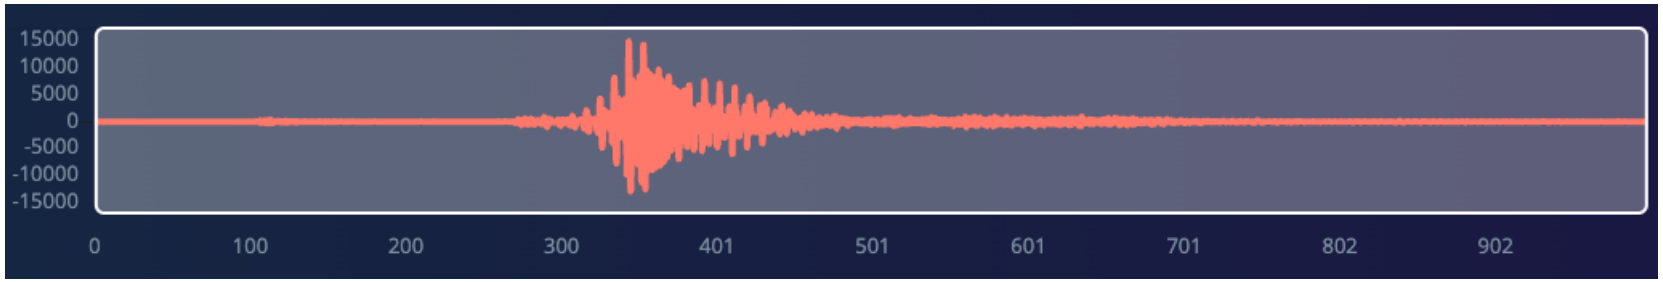

Processed features -> 637 features (13 x 49)

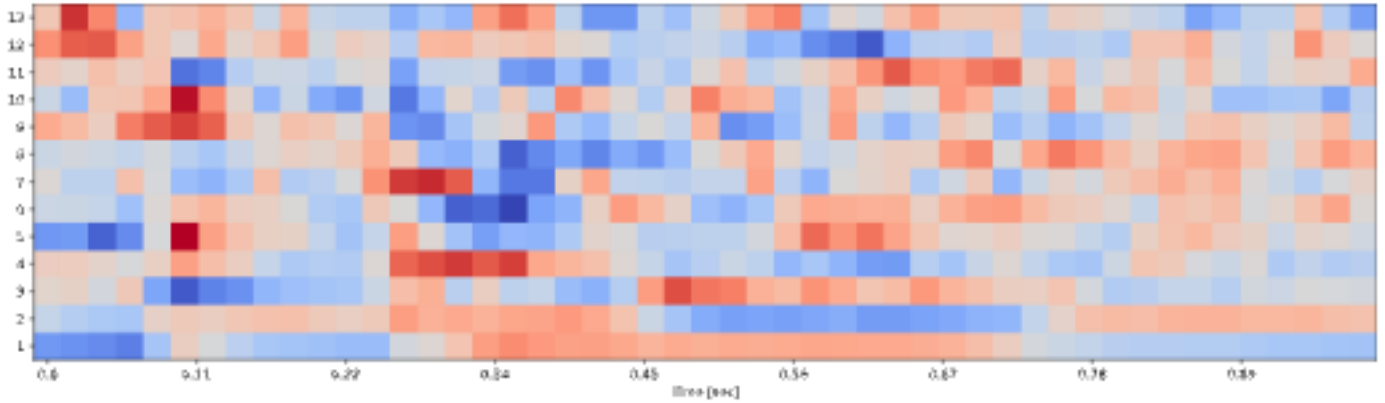

13

49

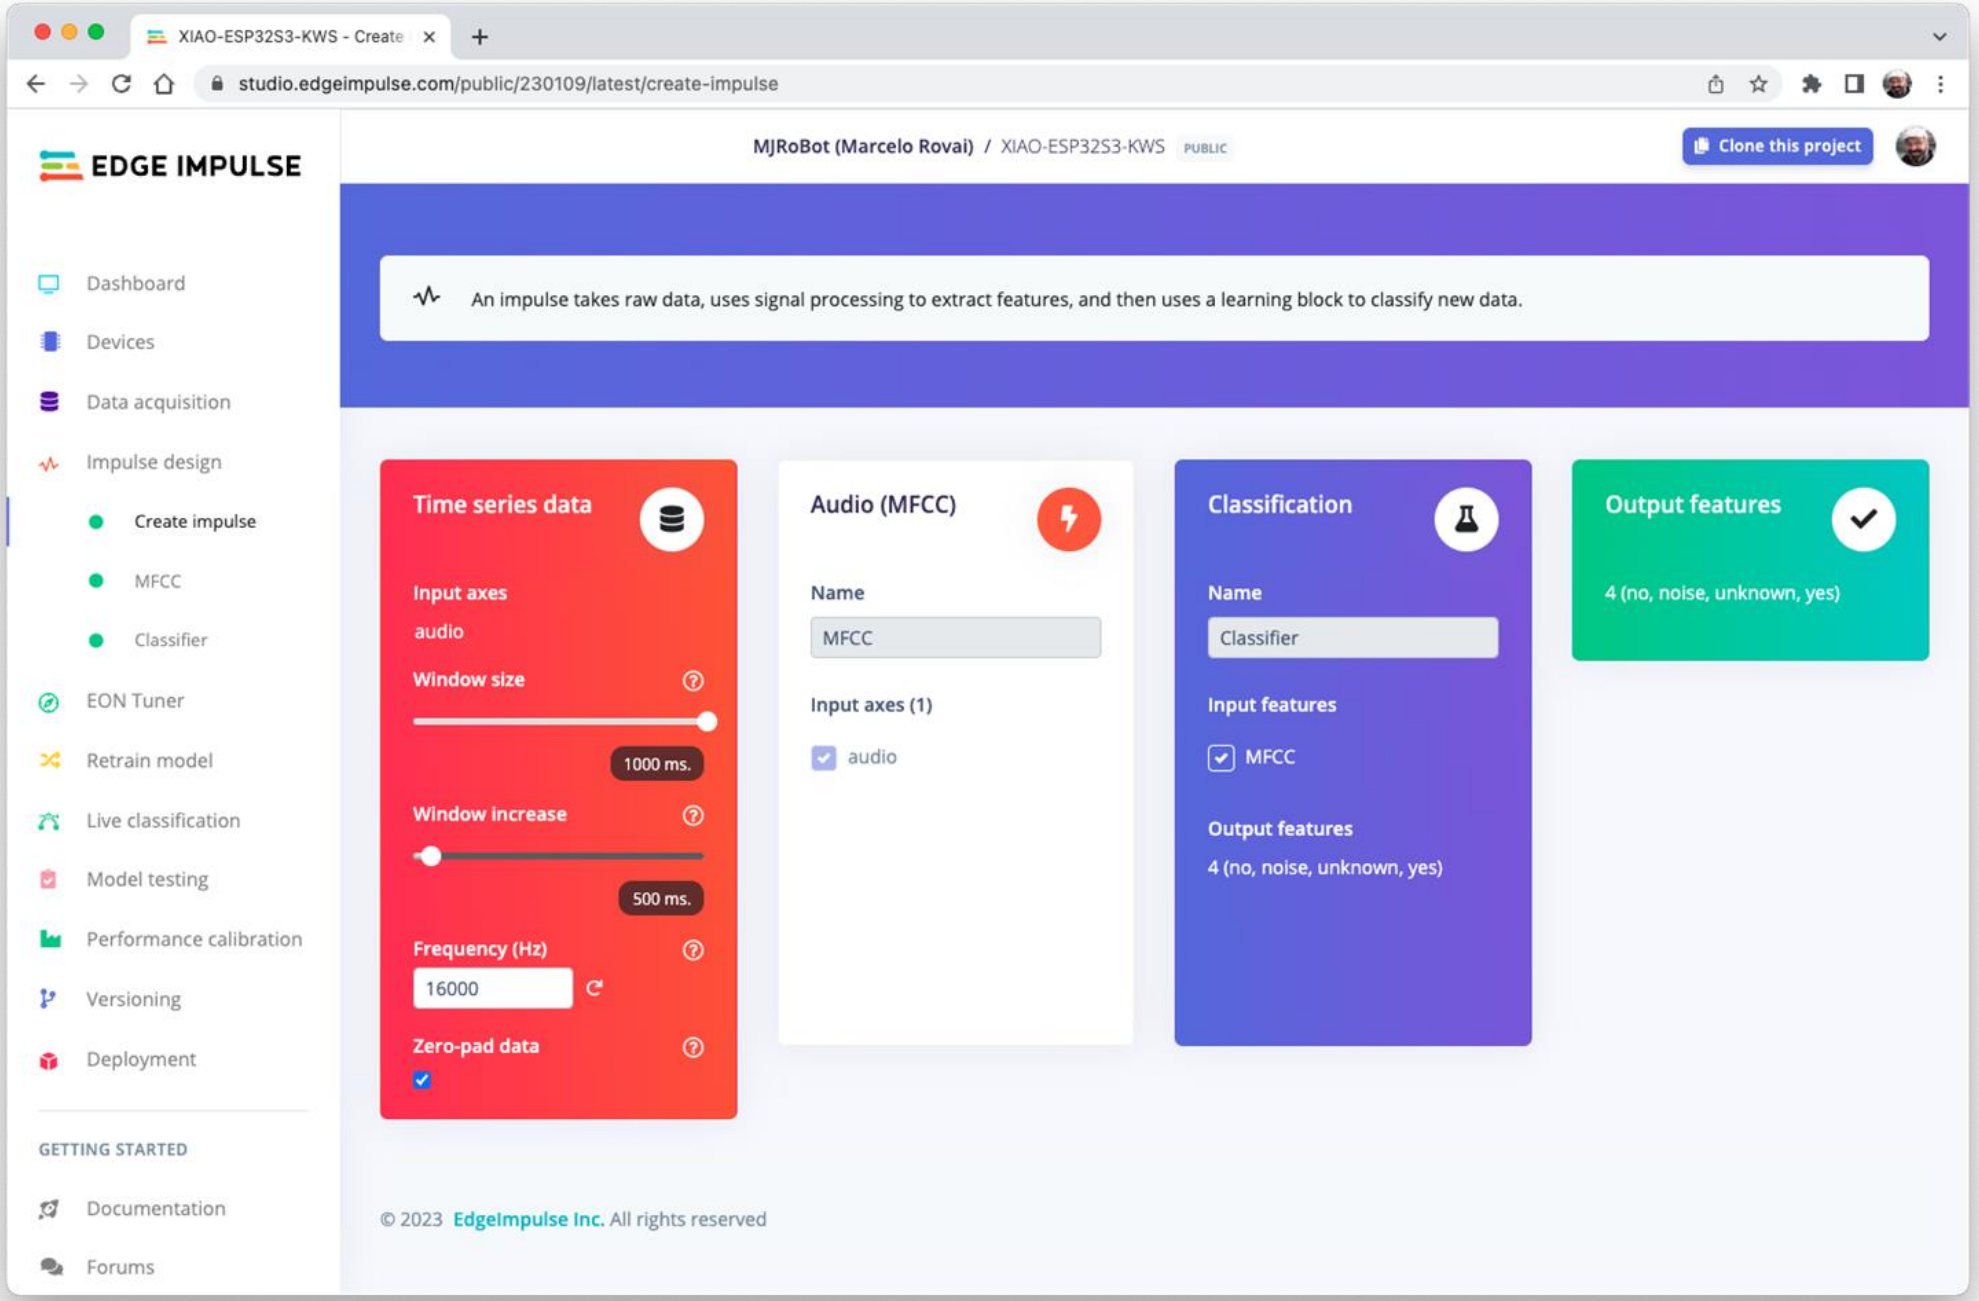
XIAO-ESP32S3-KWS - Create
studio.edgeimpulse.com/public/230109/latest/create-impulse
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / XIAO-ESP32S3-KWS PUBLIC
Clone this project
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
MFCC
Classifier
EON Tuner
Retrain model
Live classification
Model testing
Performance calibration
Versioning
Deployment
GETTING STARTED
Documentation
Forums
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Time series data
Input axes
audio
Window size
1000 ms.
Window increase
500 ms.
Frequency (Hz)
16000
Zero-pad data
Audio (MFCC)
Name
MFCC
Input axes (1)
audio
Classification
Name
Classifier
Input features
MFCC
Output features
4 (no, noise, unknown, yes)
Output features
4 (no, noise, unknown, yes)
© 2023 EdgeImpulse Inc. All rights reserved

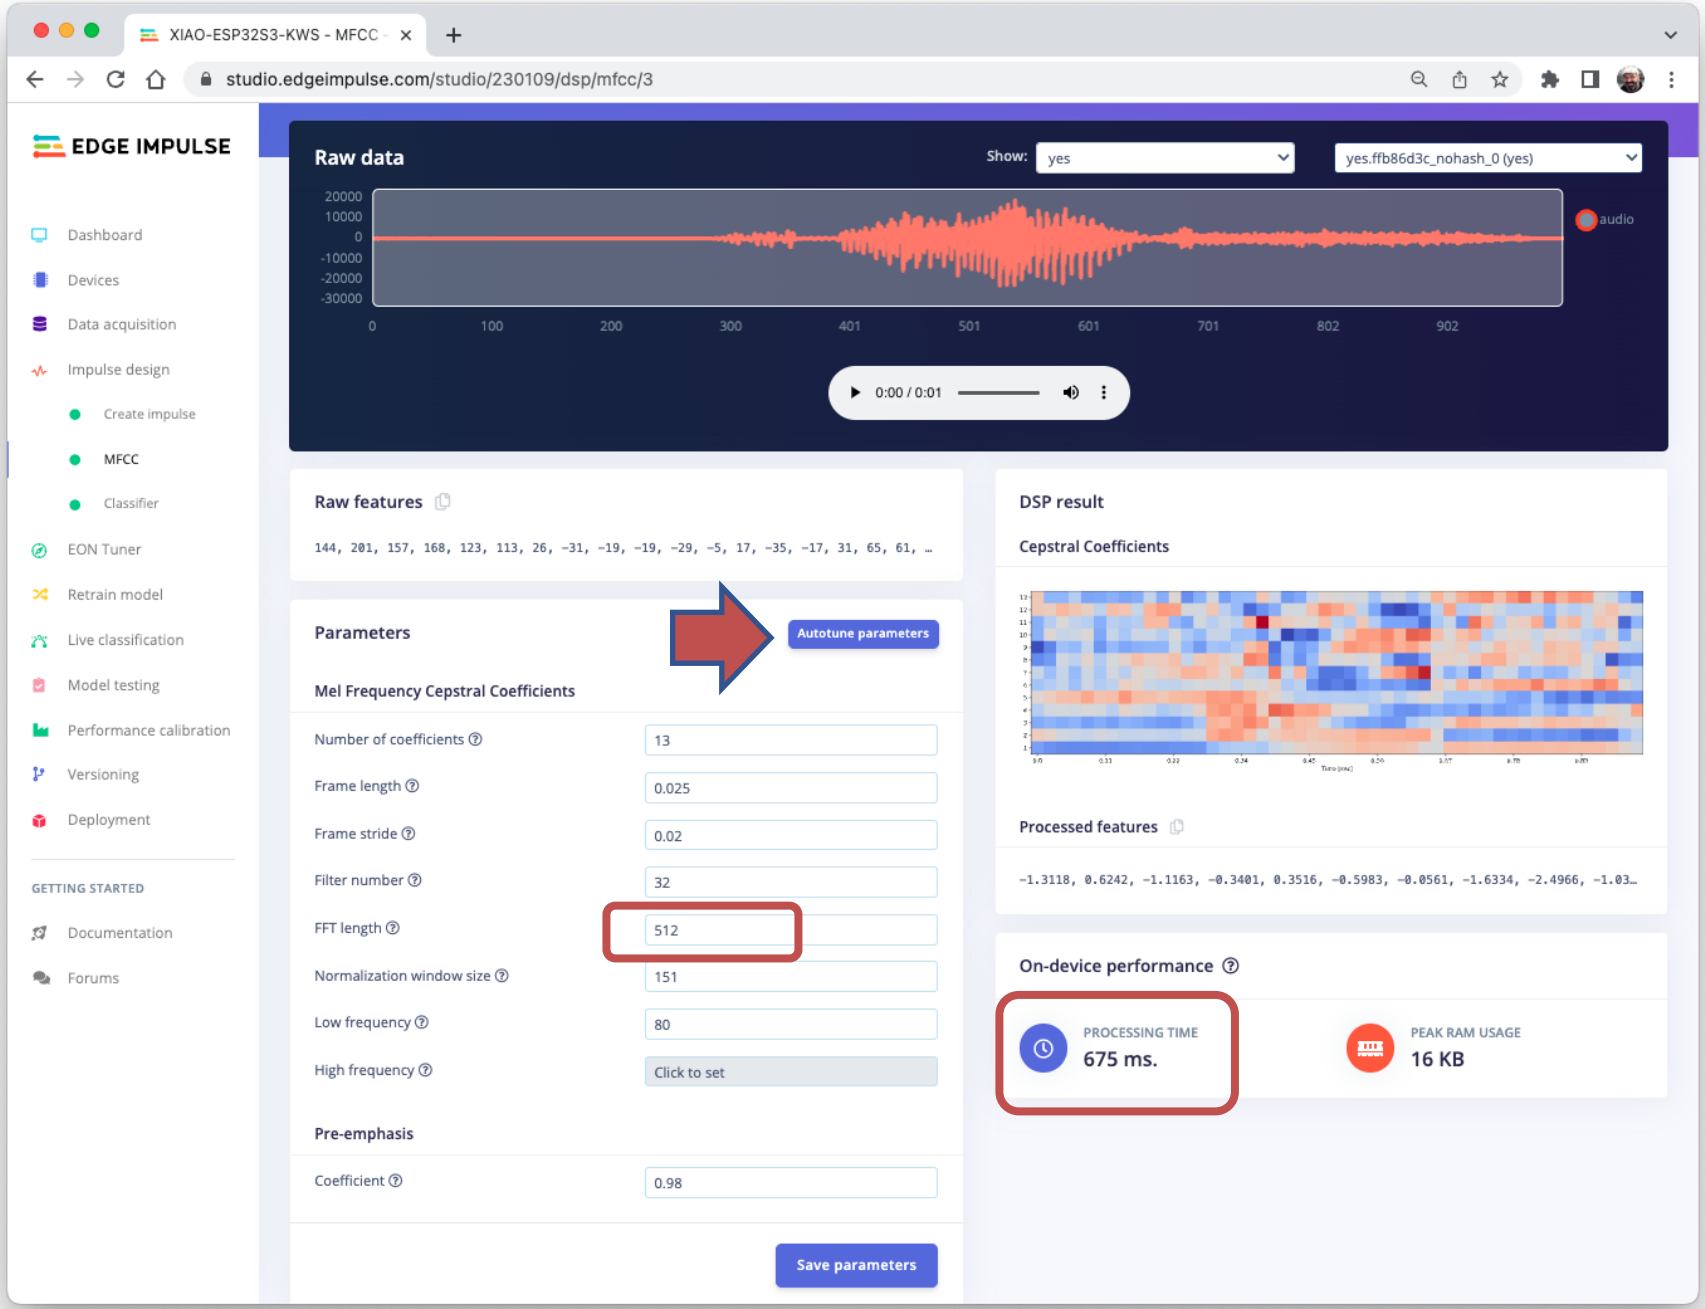
XIAO-ESP32S3-KWS - MFCC
studio.edgeimpulse.com/studio/230109/dsp/mfcc/3
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
MFCC
Classifier
EON Tuner
Retrain model
Live classification
Model testing
Performance calibration
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Raw data
Show:
yes
yes.ffb86d3c_nohash_0 (yes)
audio
0:00 / 0:01
Raw features
144, 201, 157, 168, 123, 113, 26, -31, -19, -19, -29, -5, 17, -35, -17, 31, 65, 61, …
Parameters
Autotune parameters
Mel Frequency Cepstral Coefficients
Number of coefficients 13
Frame length 0.025
Frame stride 0.02
Filter number 32
FFT length 512
Normalization window size 151
Low frequency 80
High frequency Click to set
Pre-emphasis
Coefficient 0.98
Save parameters
DSP result
Cepstral Coefficients
Processed features
-1.3118, 0.6242, -1.1163, -0.3401, 0.3516, -0.5983, -0.0561, -1.6334, -2.4966, -1.03…
On-device performance
PROCESSING TIME
675 ms.
PEAK RAM USAGE
16 KB

Neural network architecture

Architecture presets (1D Convolutional (Default)) (2D Convolutional)

Input layer (637 features)

Reshape layer (13 columns)

1D conv / pool layer (8 neurons, 3 kernel size, 1 layer)

Dropout (rate 0.25)

1D conv / pool layer (16 neurons, 3 kernel size, 1 layer)

Dropout (rate 0.25)

Flatten layer

Add an extra layer

Output layer (4 classes)

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY
90.7%

LOSS
0.25

Confusion matrix (validation set)

	NO	NOISE	UNKNOWN	YES
NO	92.2%	0.8%	5.3%	1.6%
NOISE	0.4%	95.2%	4.0%	0.4%
UNKNOWN	10.2%	5.1%	82.0%	2.7%
YES	2.1%	0.4%	3.3%	94.1%
F1 SCORE	0.90	0.94	0.85	0.95

Data explorer (full training set)

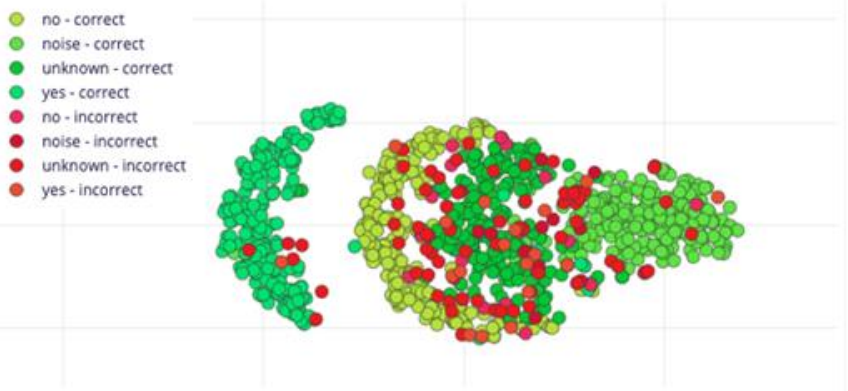

On-device performance

INFERENCING TIME
6 ms.

PEAK RAM USAGE
3.7K

FLASH USAGE
27.1K

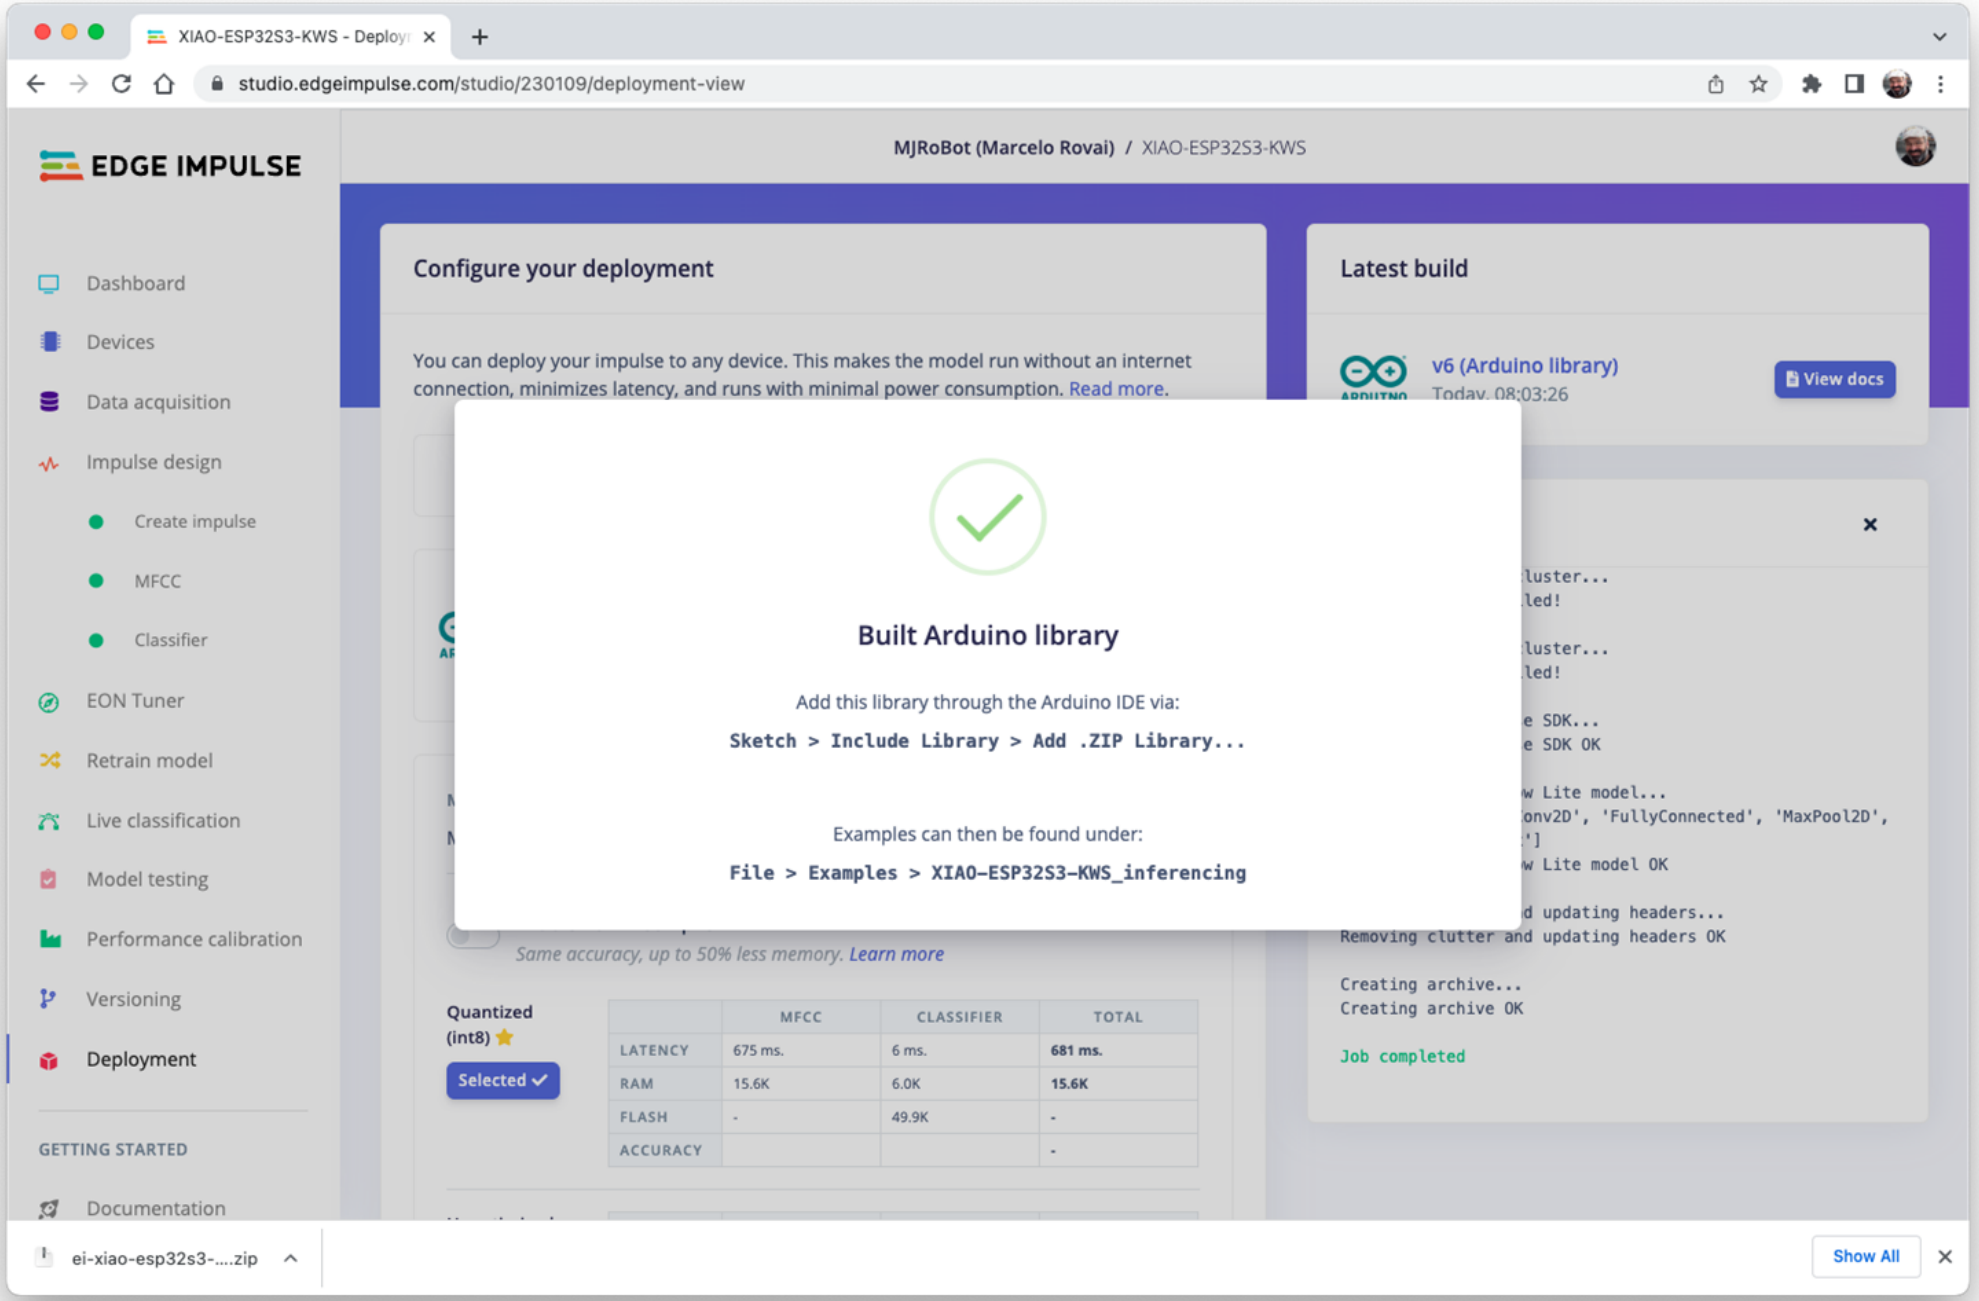
XIAO-ESP32S3-KWS - Deploy
studio.edgeimpulse.com/studio/230109/deployment-view
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / XIAO-ESP32S3-KWS
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
MFCC
Classifier
EON Tuner
Retrain model
Live classification
Model testing
Performance calibration
Versioning
Deployment
GETTING STARTED
Documentation
Configure your deployment
You can deploy your impulse to any device. This makes the model run without an internet connection, minimizes latency, and runs with minimal power consumption. Read more.
Latest build
v6 (Arduino library)
Today, 08:03:26
View docs
Built Arduino library
Add this library through the Arduino IDE via:
Sketch > Include Library > Add .ZIP Library...
Examples can then be found under:
File > Examples > XIAO-ESP32S3-KWS_inferencing
Same accuracy, up to 50% less memory. Learn more
Quantized (int8)
Selected
MFCC
CLASSIFIER
TOTAL
LATENCY
675 ms.
6 ms.
681 ms.
RAM
15.6K
6.0K
15.6K
FLASH
-
49.9K
-
ACCURACY
-
Removing clutter and updating headers OK
Creating archive...
Creating archive OK
Job completed
ei-xiao-esp32s3-....zip
Show All

TINY ML TALKS webcast

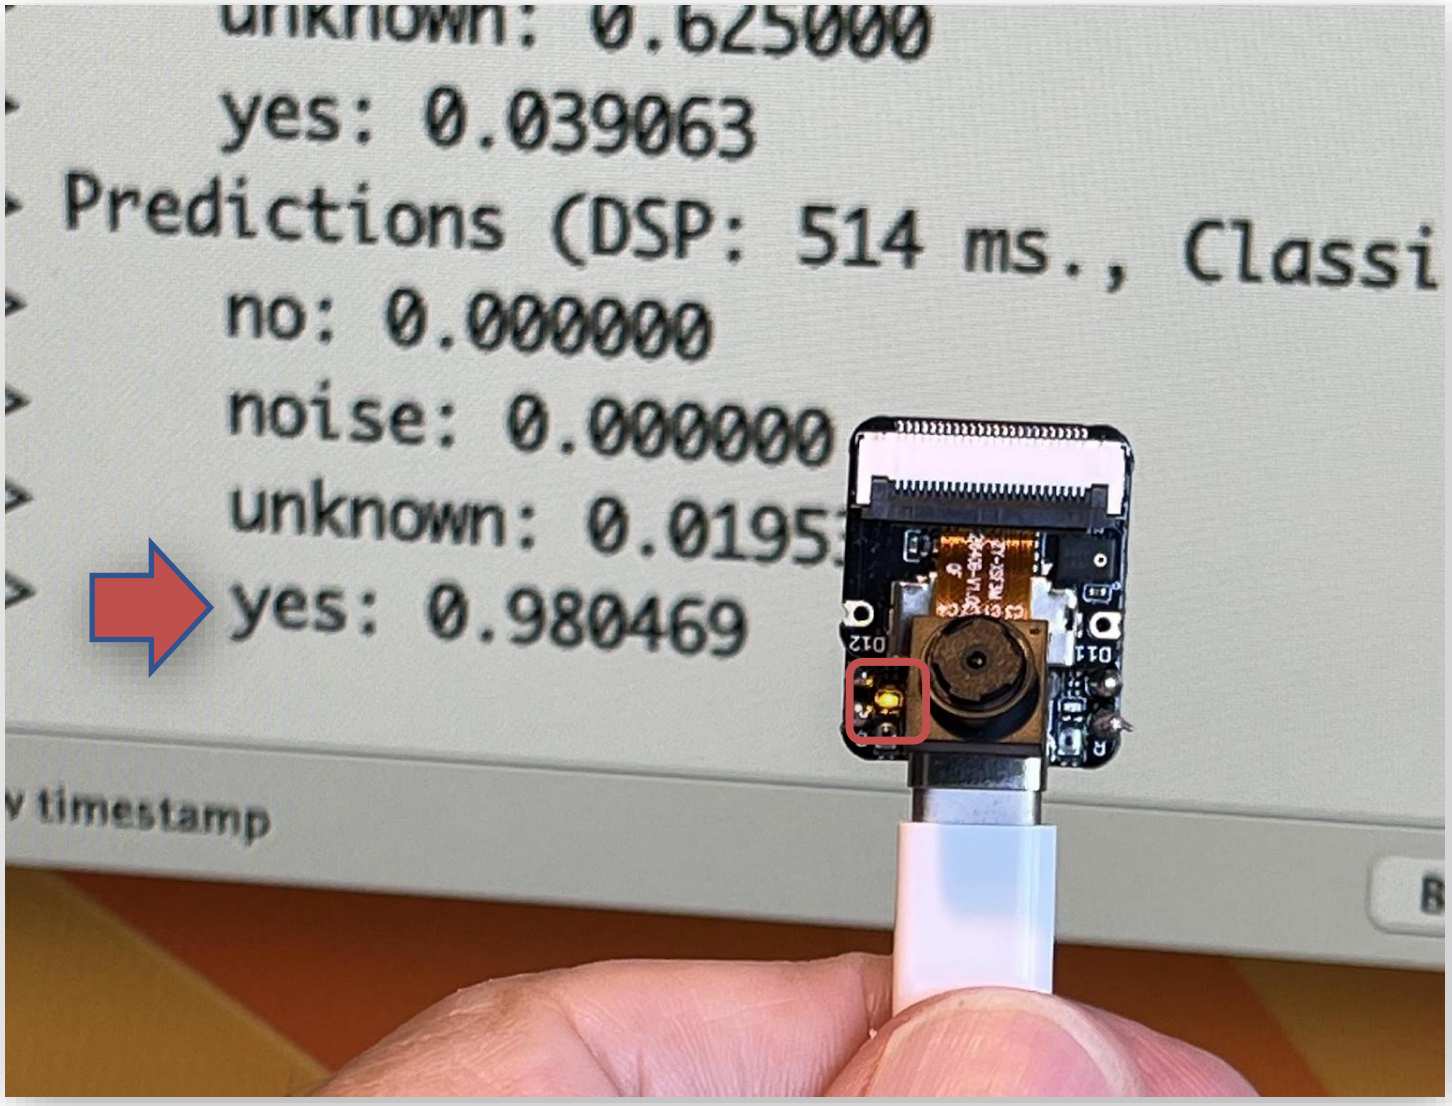

To learn more …

- IESTI01 TinyML - Machine Learning for Embedding Devices (Videos: Pt)
- WALC 22 – Applied AI - TinyML (Videos in Spanish)
- Professional Certificate in Tiny Machine Learning (TinyML) – edX/Harvard
- Introduction to Embedded Machine Learning - Coursera/Edge Impulse
- Computer Vision with Embedded Machine Learning - Coursera/Edge Impulse
- "Deep Learning with Python" book by François Chollet
- "TinyML" book by Pete Warden, Daniel Situnayake
- "TinyML Cookbook" by Gian Marco Iodice
- "AI at the Edge" book by Daniel Situnayake, Jenny Plunkett

On the TinyML4D website, You can find lots of educational materials on TinyML. They are all free and open-source for educational uses – we ask that if you use the material, please cite them! TinyML4D is an initiative to make TinyML education available to everyone globally.

TinyML4D **Show&Tell**

Date	Thread	Video
August 31st, 2023	TBD	Video here when ready
May 25th, 2023	Thread here	Video here when ready
April 20th, 2023	Thread here	https://youtu.be/uoM_IjXjDFY
March 30th, 2023	thread here	https://youtu.be/UQ0l-SwBwUY
February 23rd, 2023	thread here	https://youtu.be/BAEdil7X68Y
January 26th, 2023	thread here 17	https://youtu.be/-0xRZ-5UYUc 9
December 1st, 2022	thread here 2	https://youtu.be/e49pkjnIMlQ 8
October 27th, 2022	thread here 2	https://youtu.be/s8_hKpOWUwY 1

TinymML4D Academic Network Show and Tell Main Index.

The TinyML4D Academic Network Students should use this form to sign up for the latest presentations.
https://forms.gle/ic52HZMqVv4pBrkP7 2

The Show and Tell are typically held at 2 pm UTC on the last Thursday of each month and will take place in this Zoom room.
https://zoom.us/j/95229860797 1
Meeting ID: 952 2986 0797
Passcode: 141278

Copyright Notice

www.tinyml.org

Copyright Notice

www.tinyml.org