Computer Vision (CV) at the Edge

Prof. Marcelo J. Rovai
rovai@unifei.edu.br

UNIFEI - Federal University of Itajuba, Brazil
TinyML4D Academic Network Co-Chair

Edge Computer Vision has a high impact potential, and it is for now!

*"**Visual recognition** has undergone the largest changes and fastest development in the last decade, due in part to the availability of much larger labeled datasets as well as breakthroughs in deep learning."*

Computer Vision: Algorithms and Applications, 2nd ed.

Computer Vision Recognition Tasks

Image Classification

(**Multi-Class** Classification)

Cat: 70%

Dog: 80%

Object Detection

Multi-Label Classification + Object **Localization**

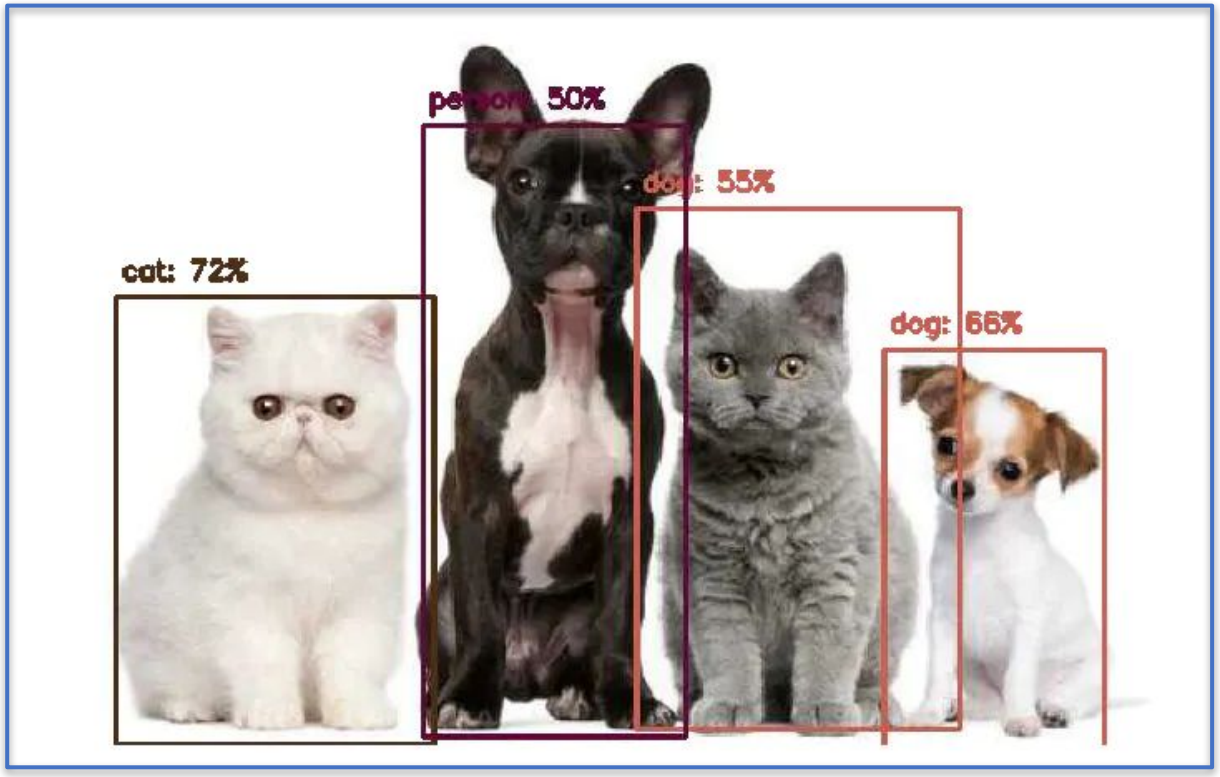

Computer Vision Recognition Tasks

Instance Segmentation

Each pixel in an image IS CLASSIFIED into a predefined category.

Pose Estimation

Key points (or landmarks) on the object, such as joints on a human body are detect

Computer Vision Recognition Tasks

Image Classification

(**Multi-Class** Classification)

Cat: 70%

Dog: 80%

Object Detection

Multi-Label Classification + Object **Localization**

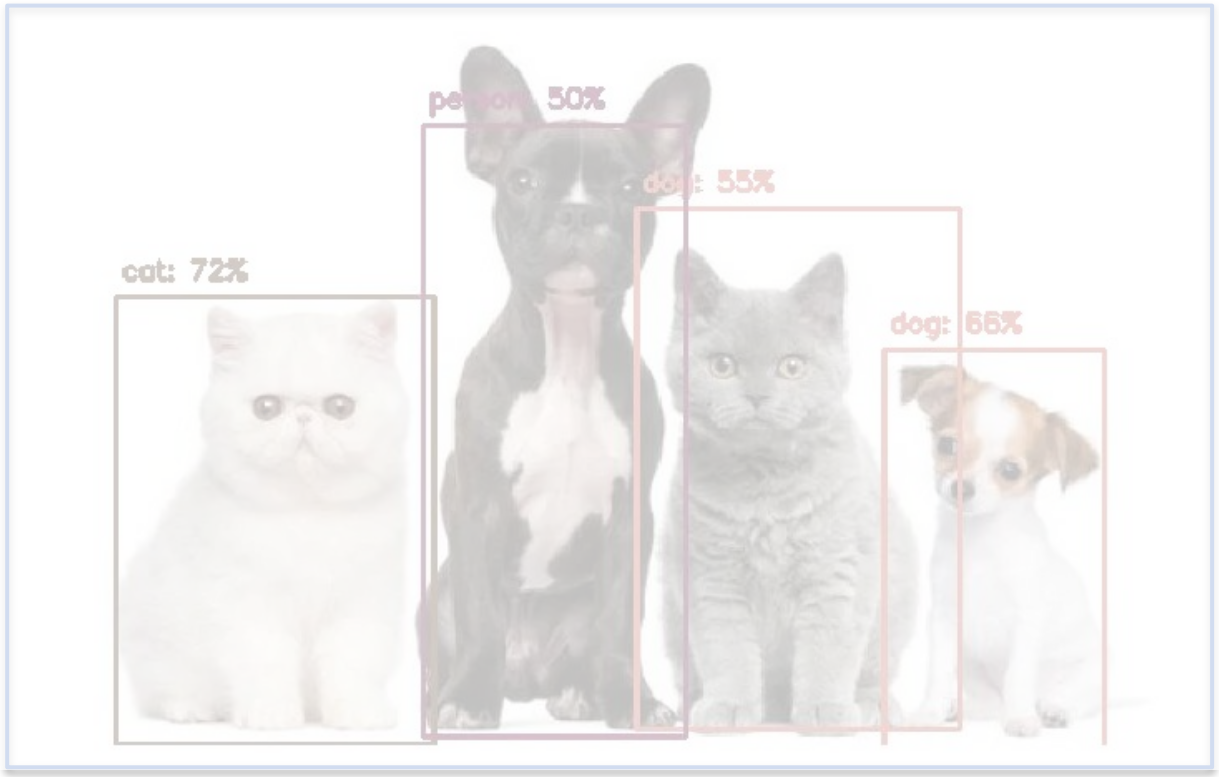

Embedded Machine Learning (TinyML) Workflow Review

Machine Learning Workflow

7

Tiny Machine Learning Workflow ("What")

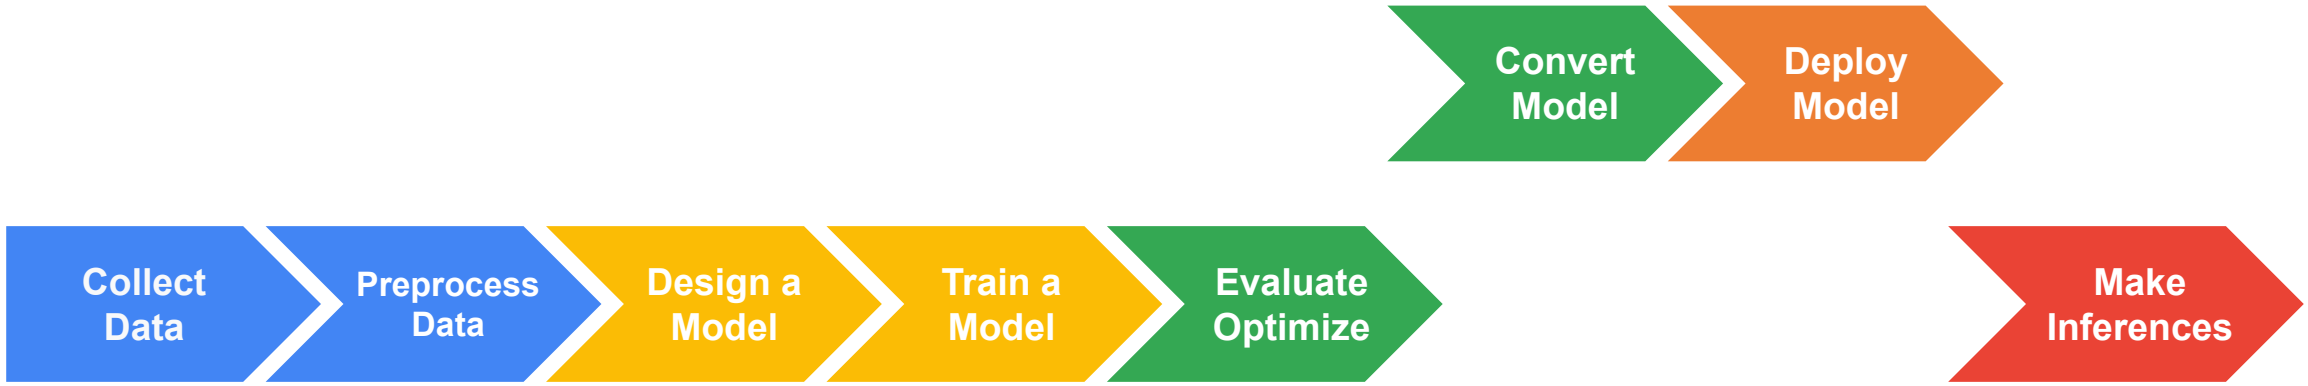

Tiny Machine Learning Workflow ("Where")

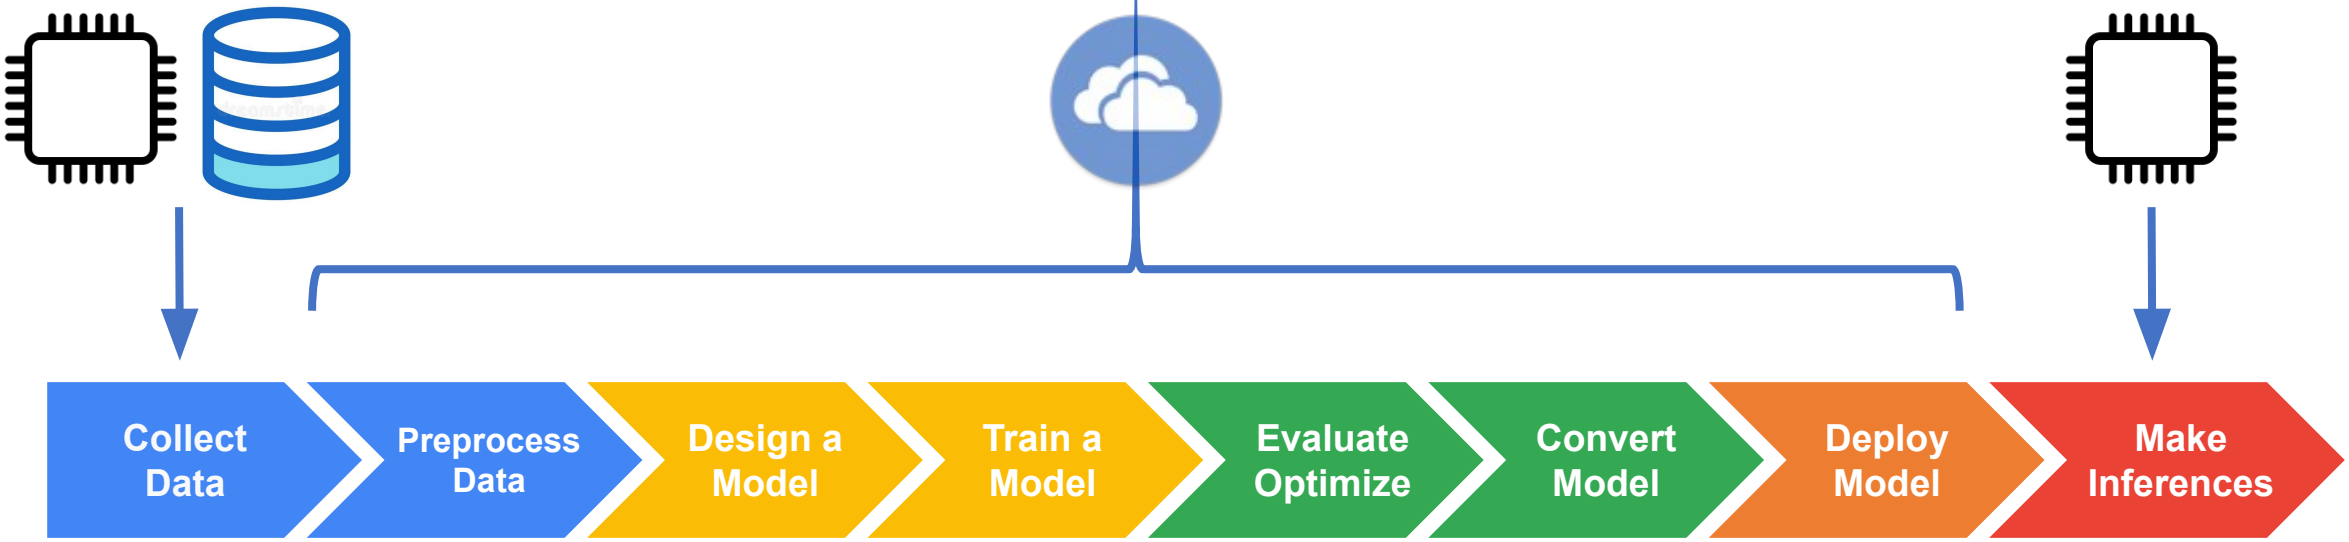

9

Machine Learning Workflow ("How")

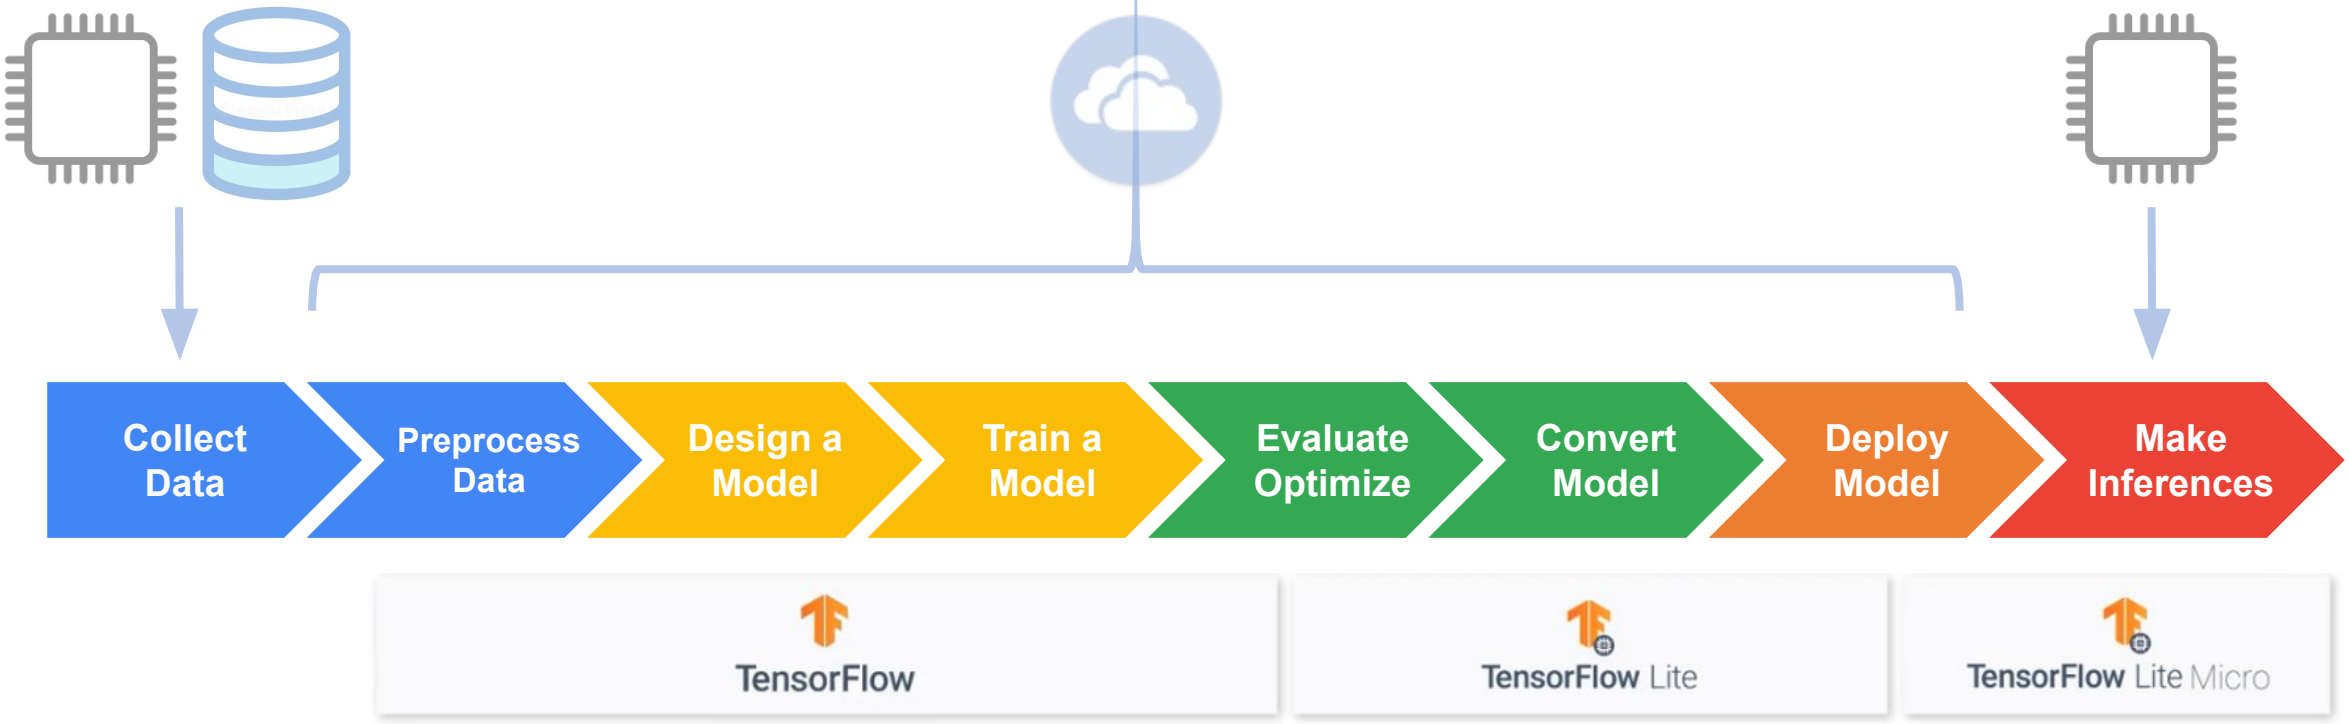

10

Machine Learning Workflow ("How")

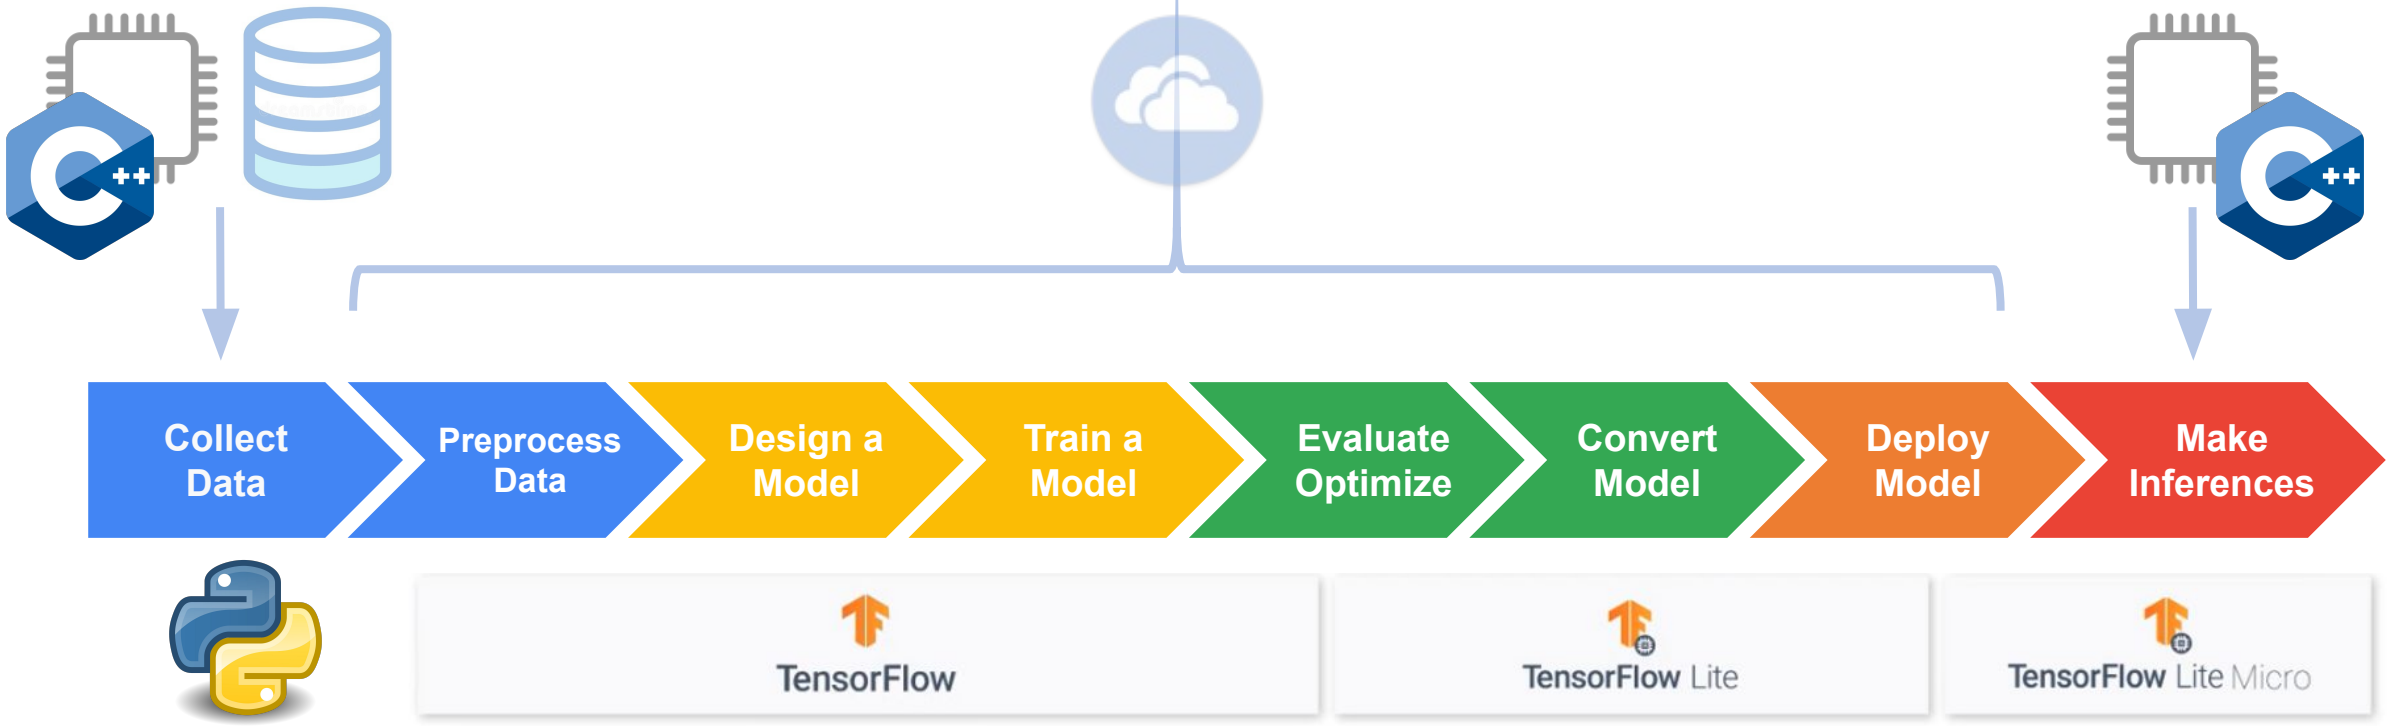

11

Machine Learning Workflow ("How")

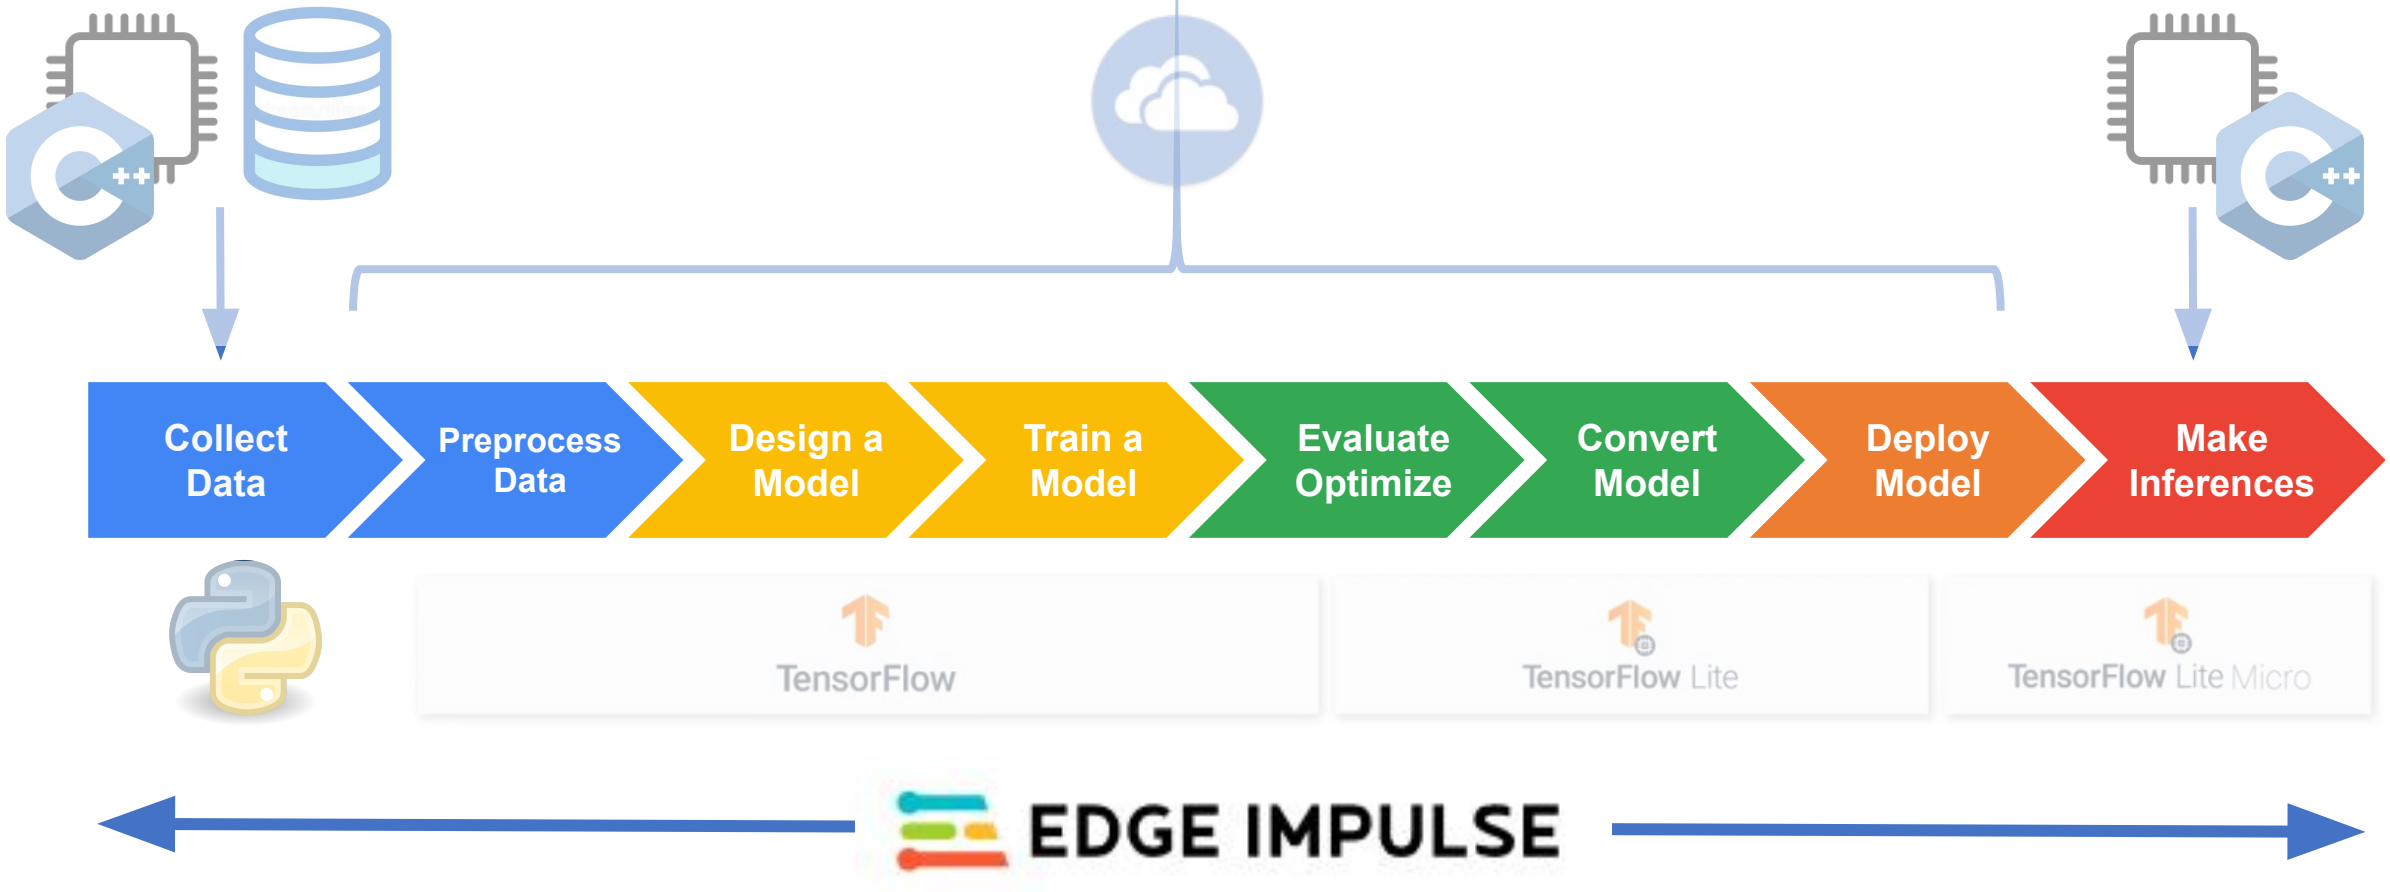

12

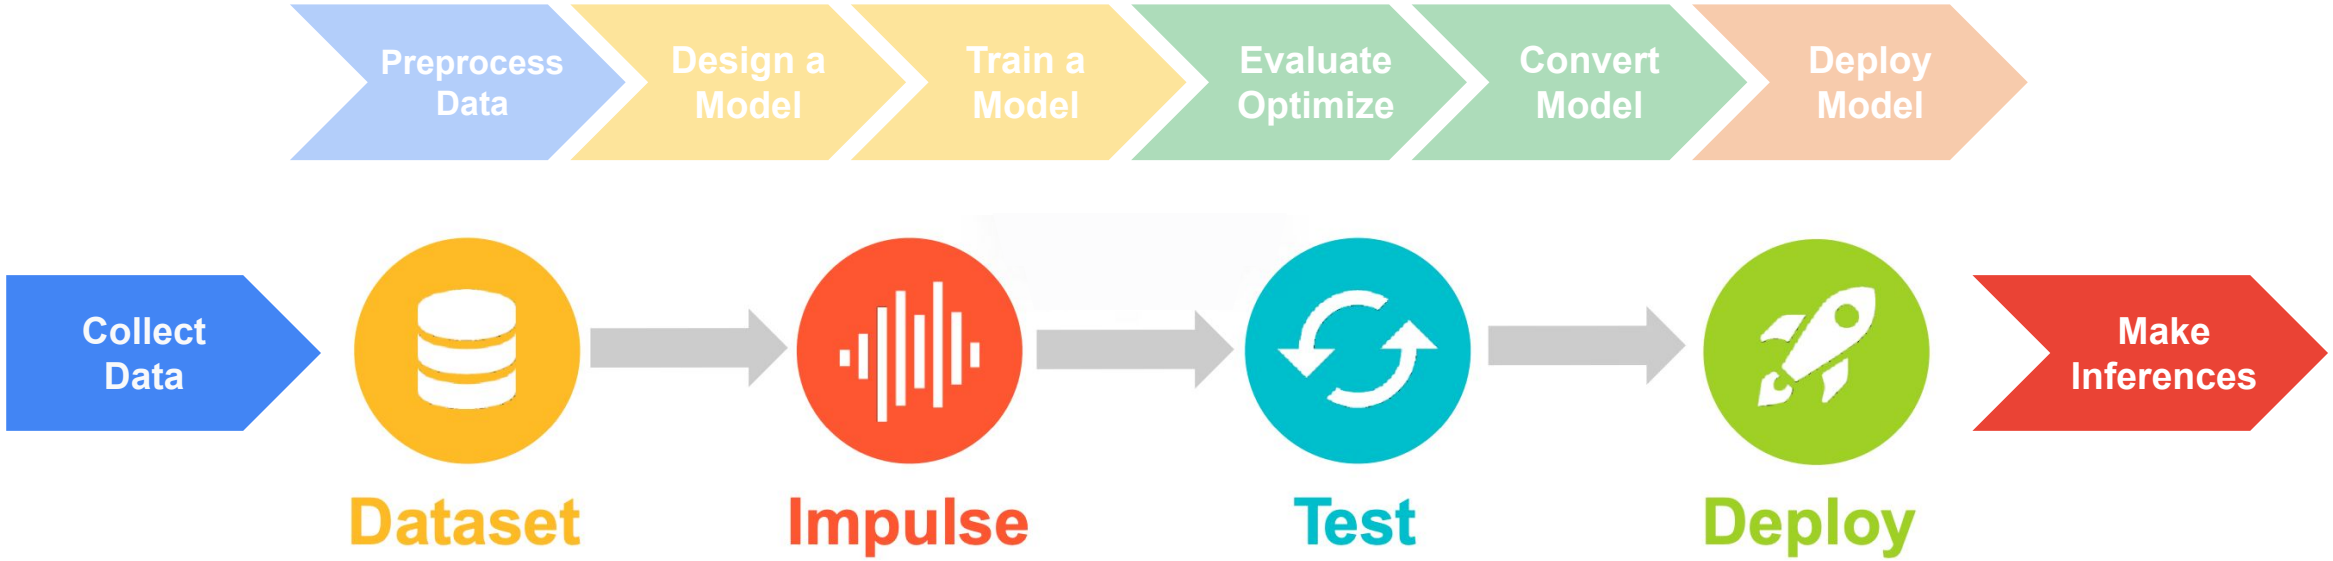
Preprocess Data
Design a Model
Train a Model
Evaluate Optimize
Convert Model
Deploy Model
Collect Data
Dataset
Impulse
Test
Deploy
Make Inferences

13

EI Studio - Embedded ML platform ("AutoML")

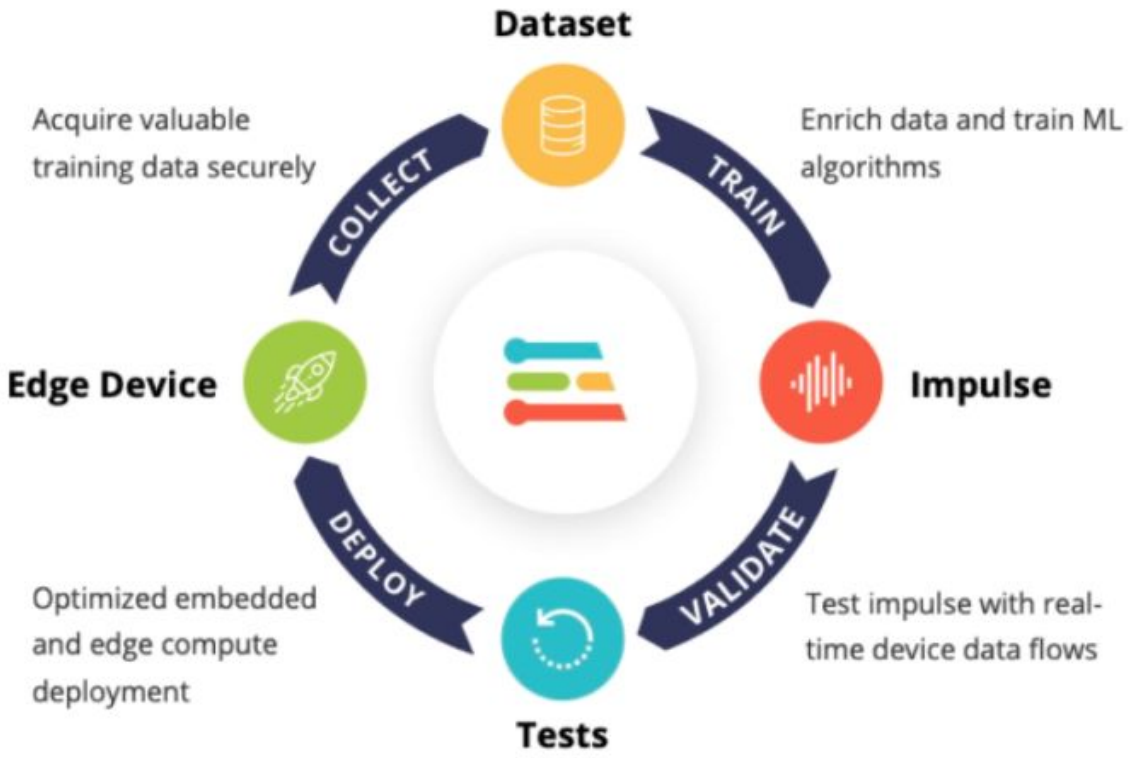

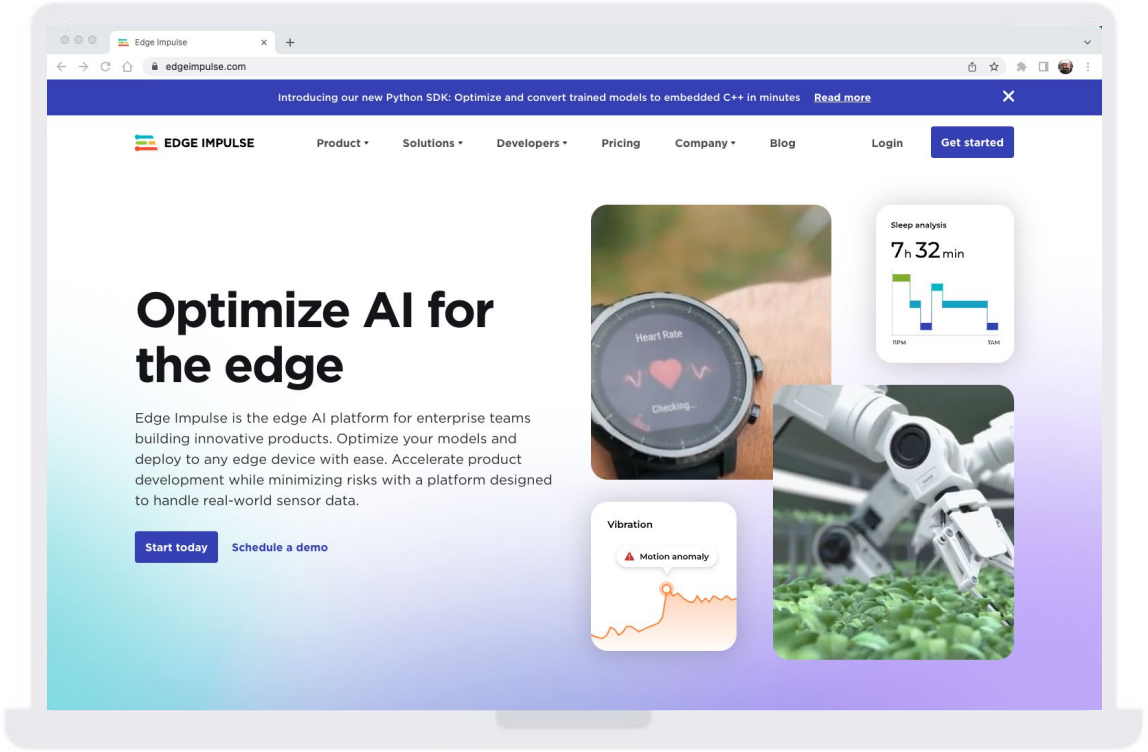

Learn more at http://edgeimpulse.com

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

Image Classification Project

Edge Impulse Studio

https://studio.edgeimpulse.com/public/353482/live

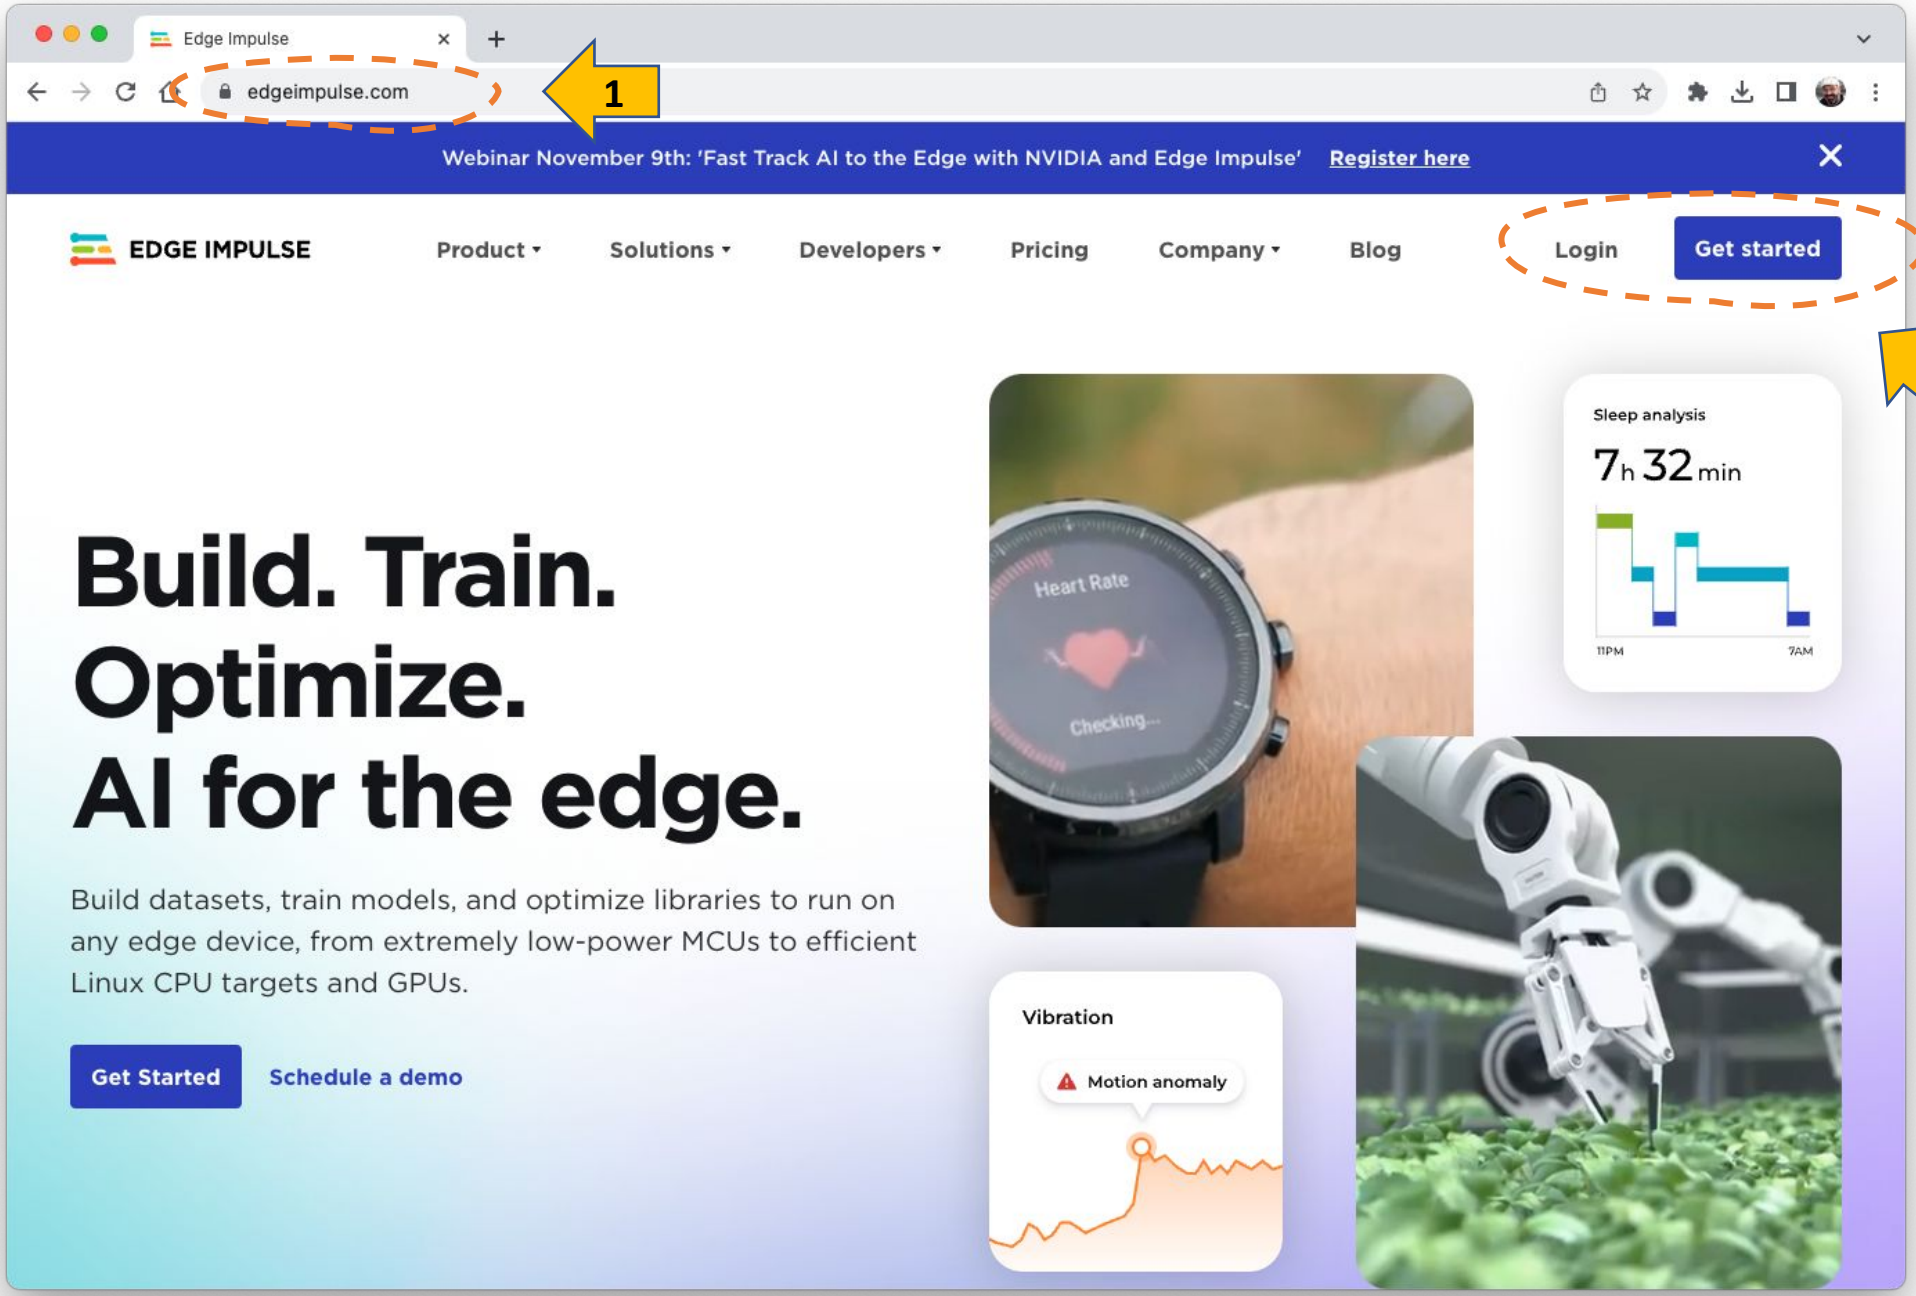
Edge Impulse
edgeimpulse.com
1
Webinar November 9th: 'Fast Track AI to the Edge with NVIDIA and Edge Impulse' Register here
EDGE IMPULSE
Product
Solutions
Developers
Pricing
Company
Blog
Login
Get started
2
Build. Train.
Optimize.
AI for the edge.
Build datasets, train models, and optimize libraries to run on any edge device, from extremely low-power MCUs to efficient Linux CPU targets and GPUs.
Get Started
Schedule a demo
Heart Rate
Checking...
Sleep analysis
7h 32 min
11PM
7AM
Vibration
Motion anomaly

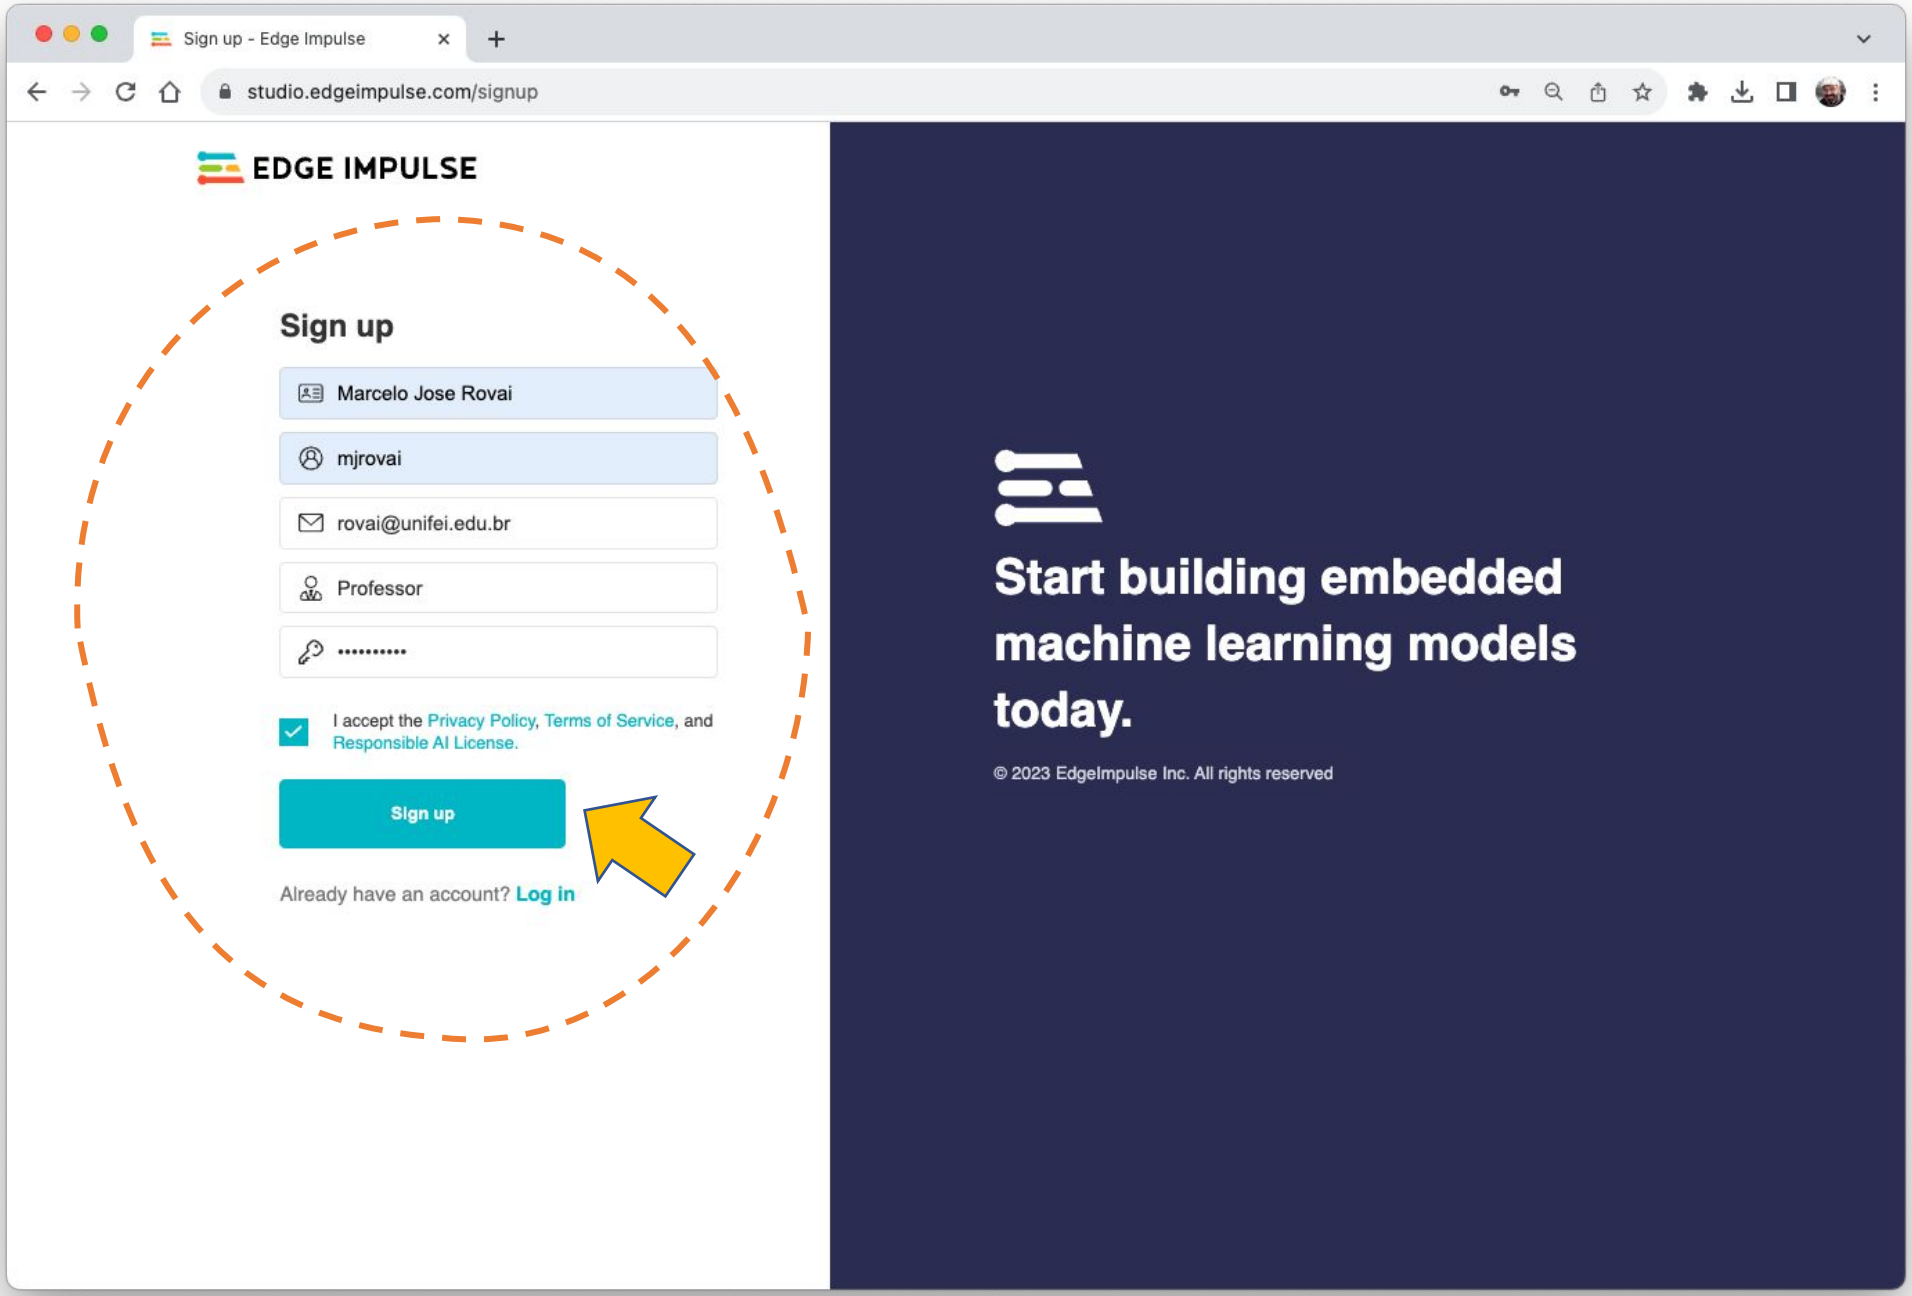
Sign up - Edge Impulse
studio.edgeimpulse.com/signup
EDGE IMPULSE
Sign up
Marcelo Jose Rovai
mjrovai
rovai@unifei.edu.br
Professor
I accept the Privacy Policy, Terms of Service, and Responsible AI License.
Sign up
Already have an account? Log in
Start building embedded machine learning models today.
© 2023 EdgeImpulse Inc. All rights reserved

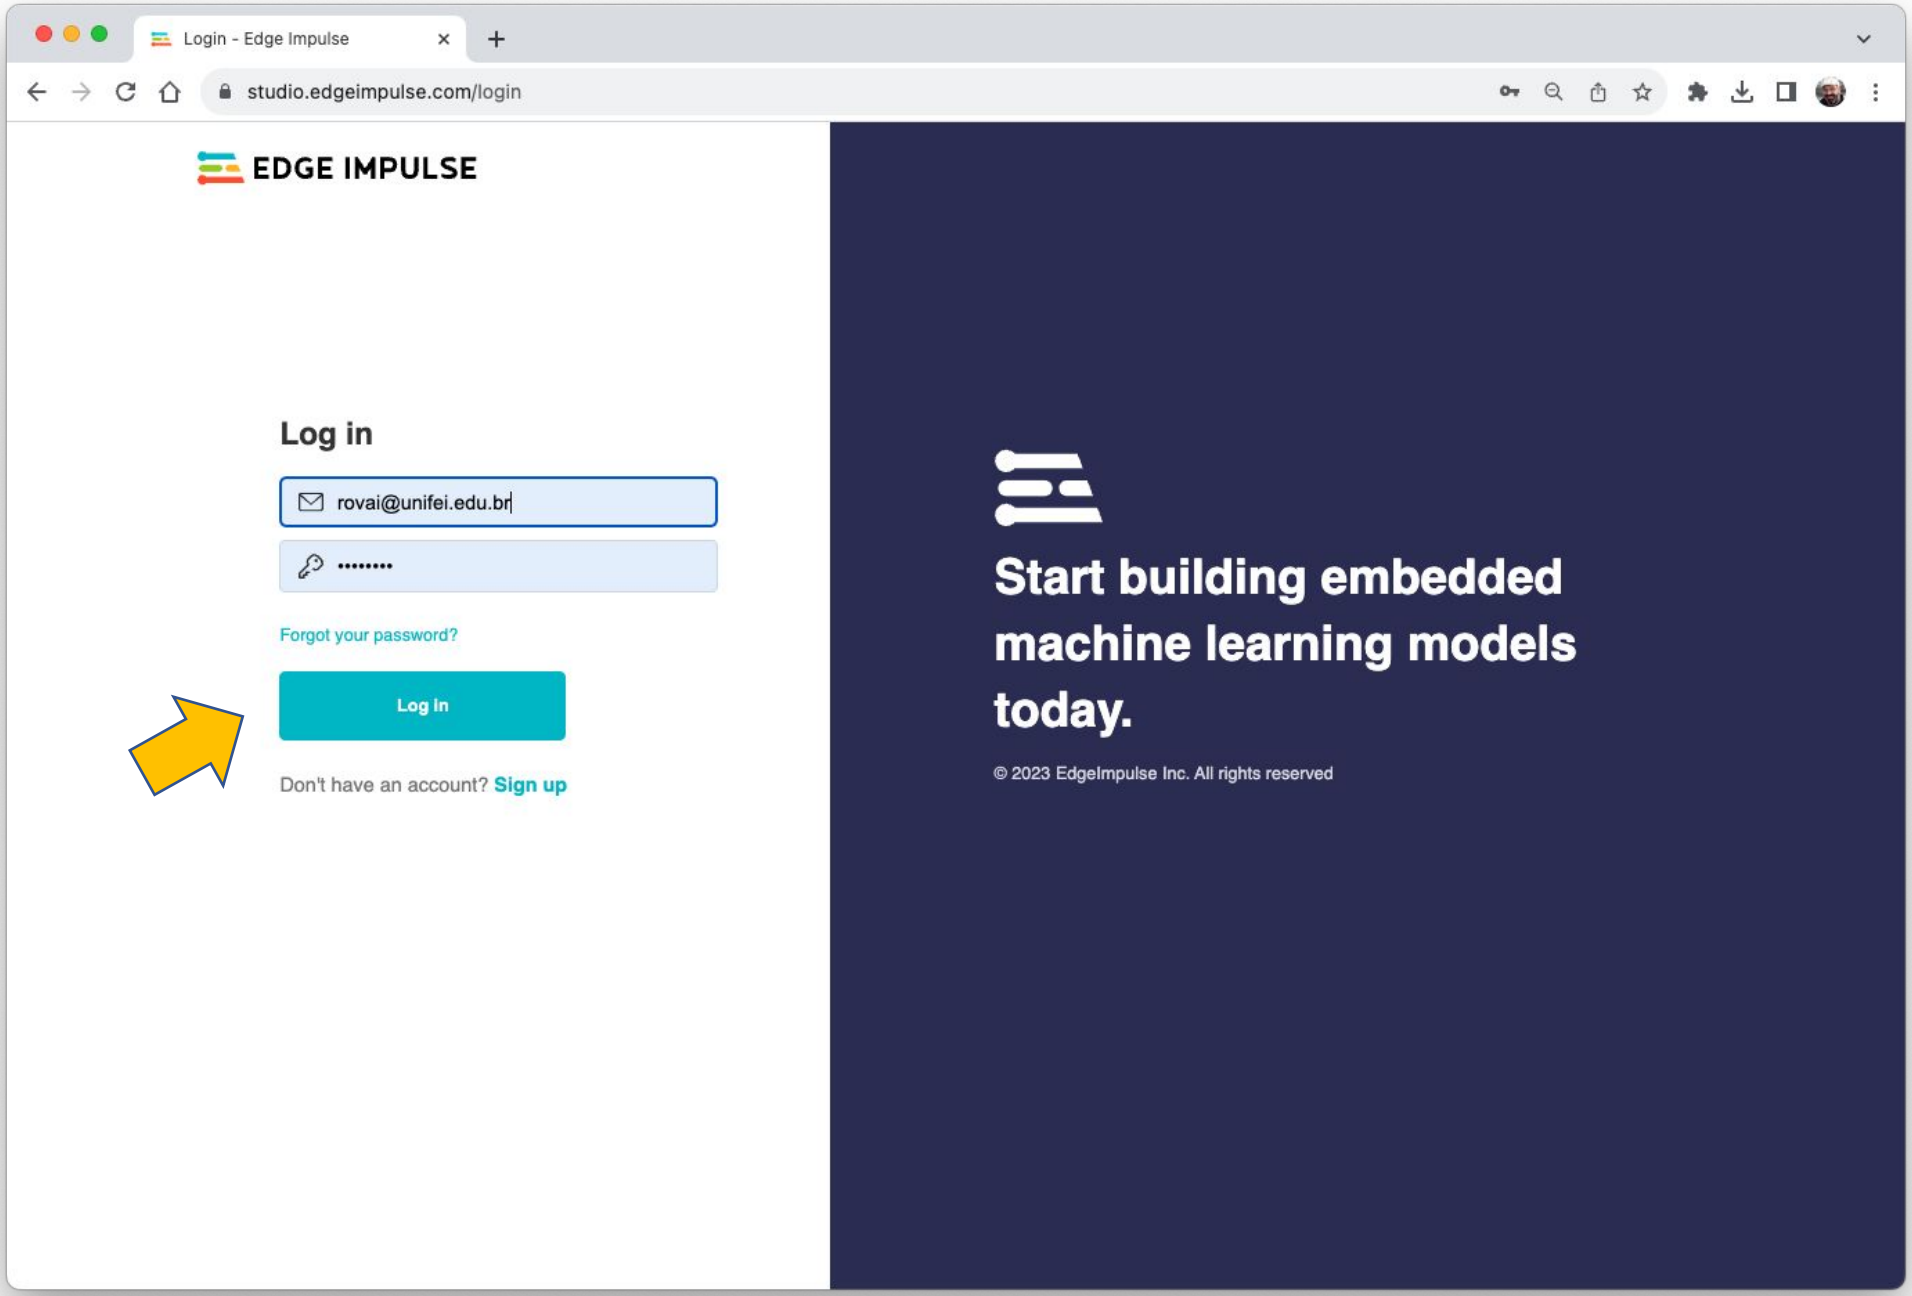
Login - Edge Impulse
studio.edgeimpulse.com/login
EDGE IMPULSE
Log in
rovai@unifei.edu.br
Forgot your password?
Log In
Don't have an account? Sign up
Start building embedded machine learning models today.
© 2023 EdgeImpulse Inc. All rights reserved

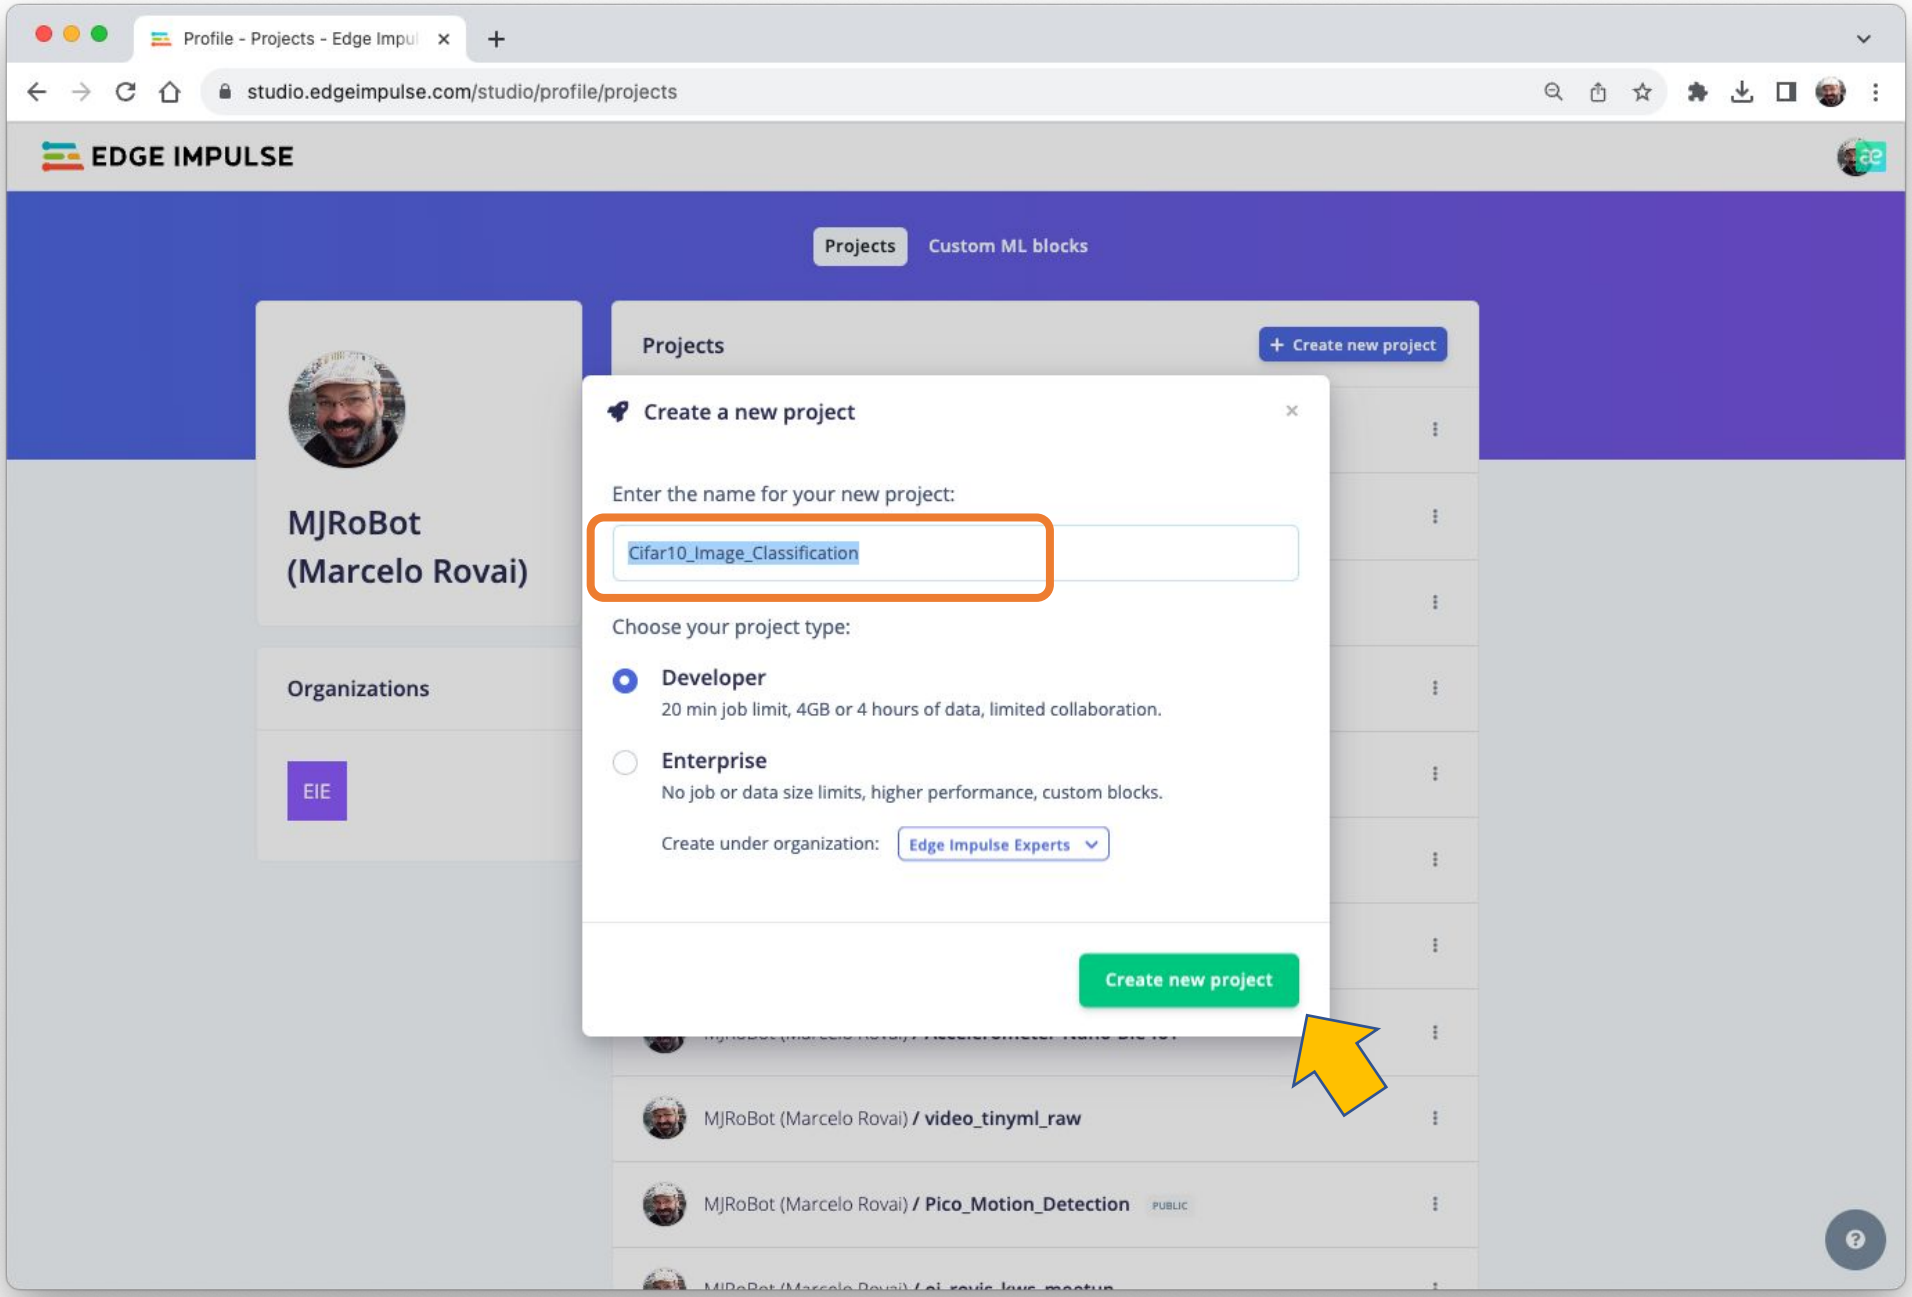
Profile - Projects - Edge Impul
studio.edgeimpulse.com/studio/profile/projects
EDGE IMPULSE
Projects
Custom ML blocks
MJRoBot
(Marcelo Rovai)
Organizations
EIE
Projects
+ Create new project
Create a new project
Enter the name for your new project:
Cifar10_Image_Classification
Choose your project type:
Developer
20 min job limit, 4GB or 4 hours of data, limited collaboration.
Enterprise
No job or data size limits, higher performance, custom blocks.
Create under organization:
Edge Impulse Experts
Create new project
MJRoBot (Marcelo Rovai) / video_tinyml_raw
MJRoBot (Marcelo Rovai) / Pico_Motion_Detection
PUBLIC

Decide a Goal

- Classes:
  - background
  - periquito
  - robot

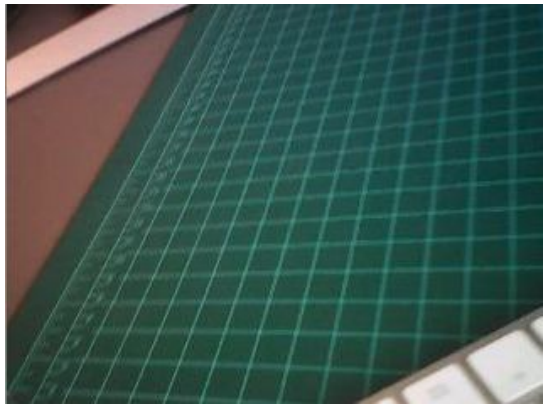

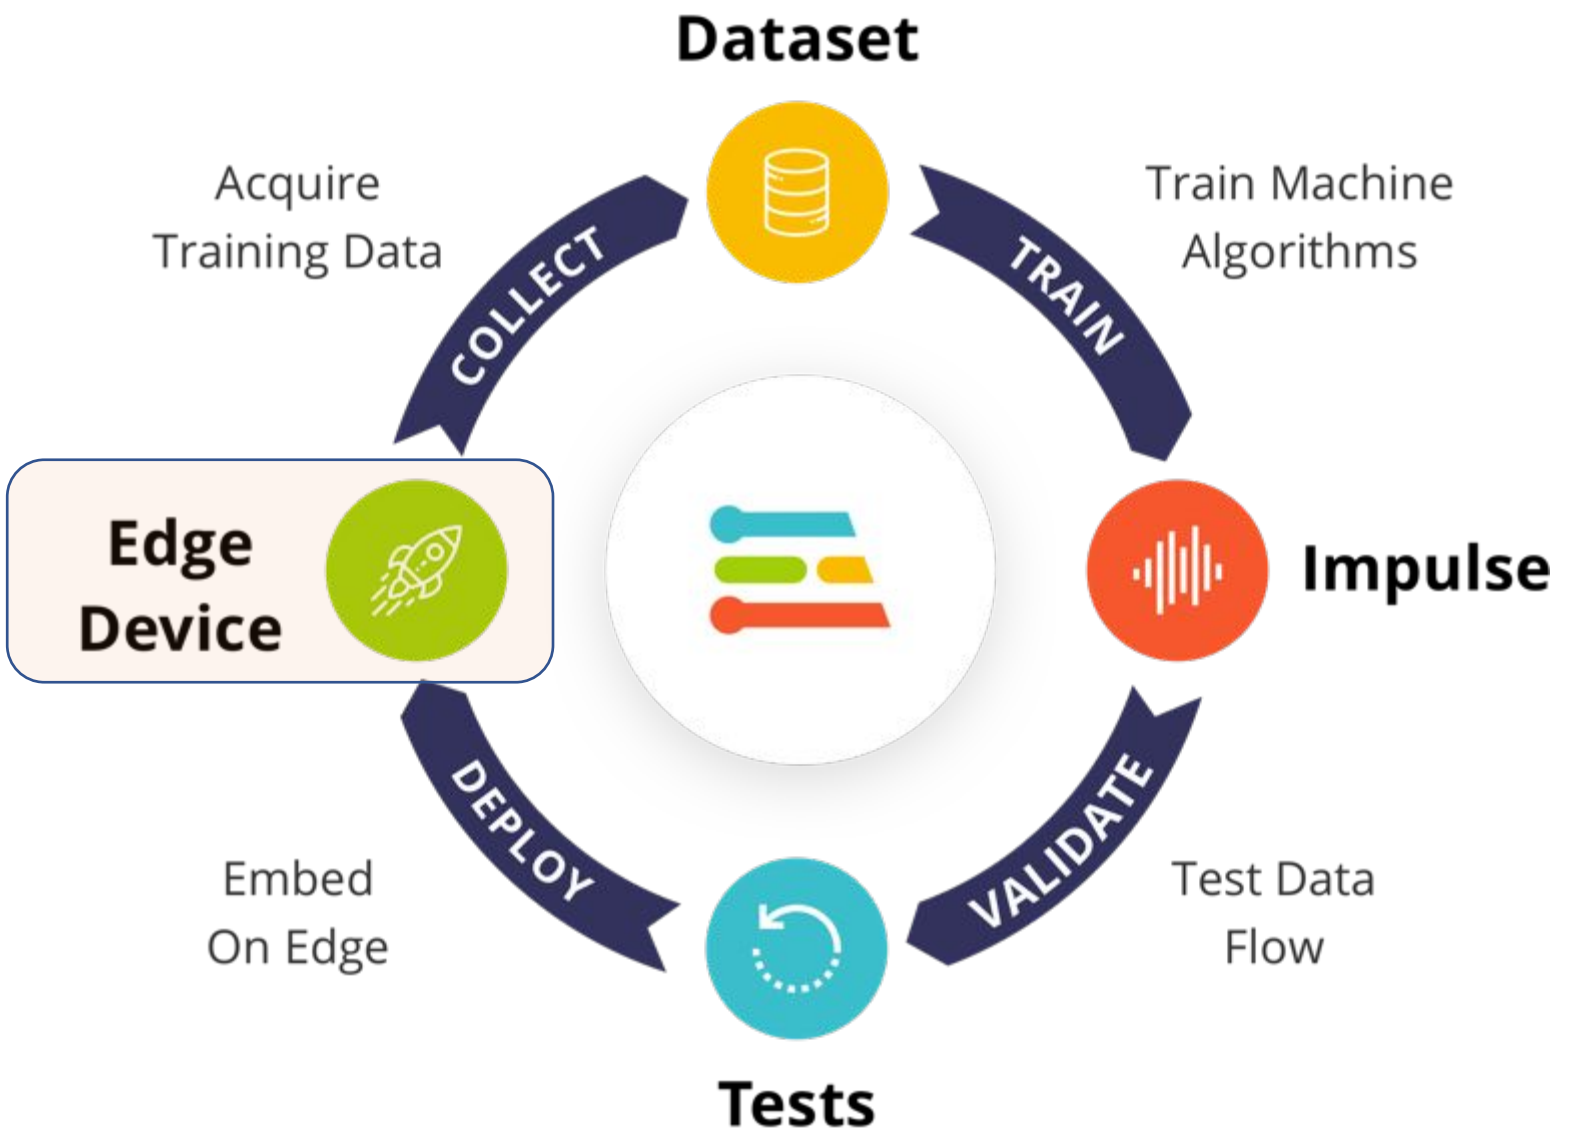

- Pre-Processing Data
- Design a Model
- Train a Model

Edge Device & Sensors

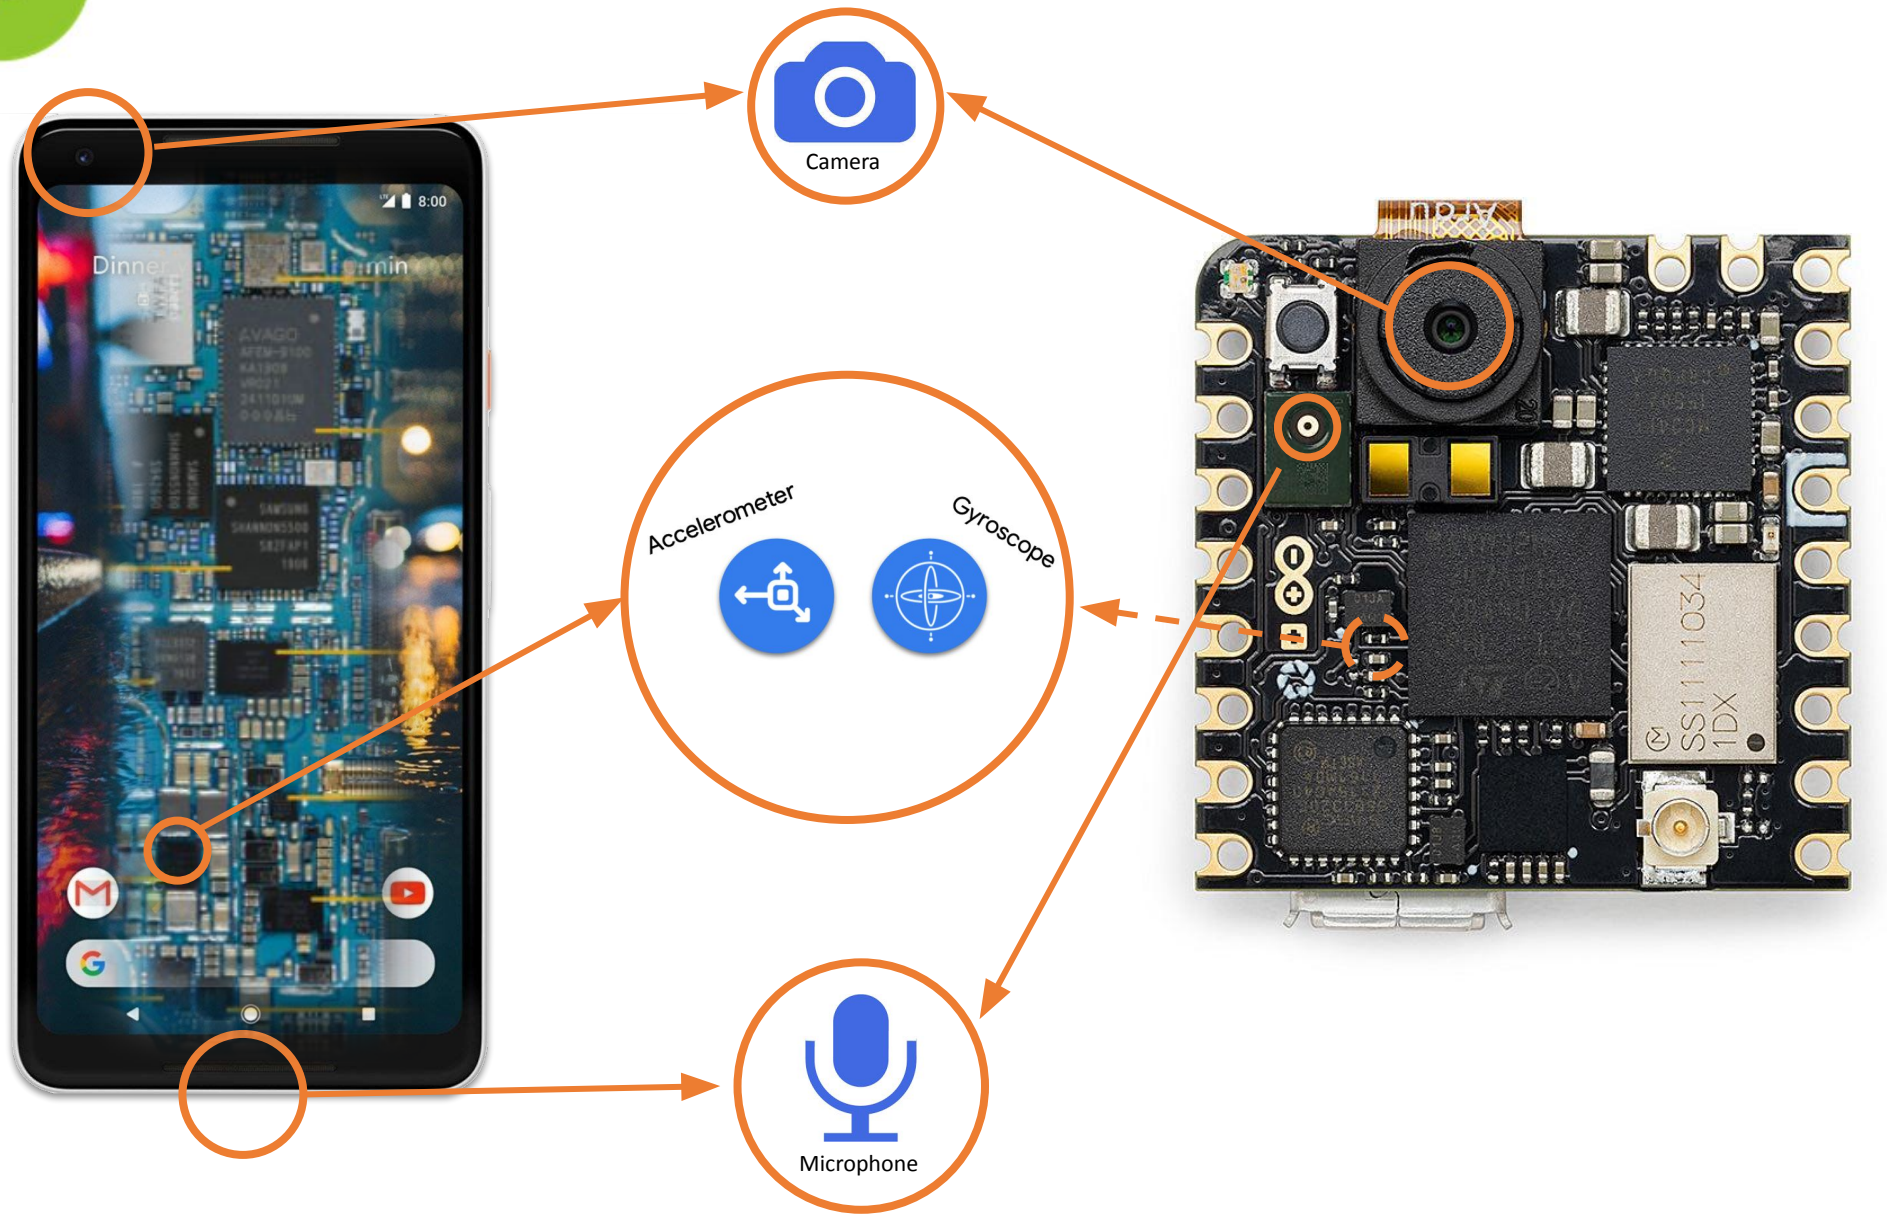

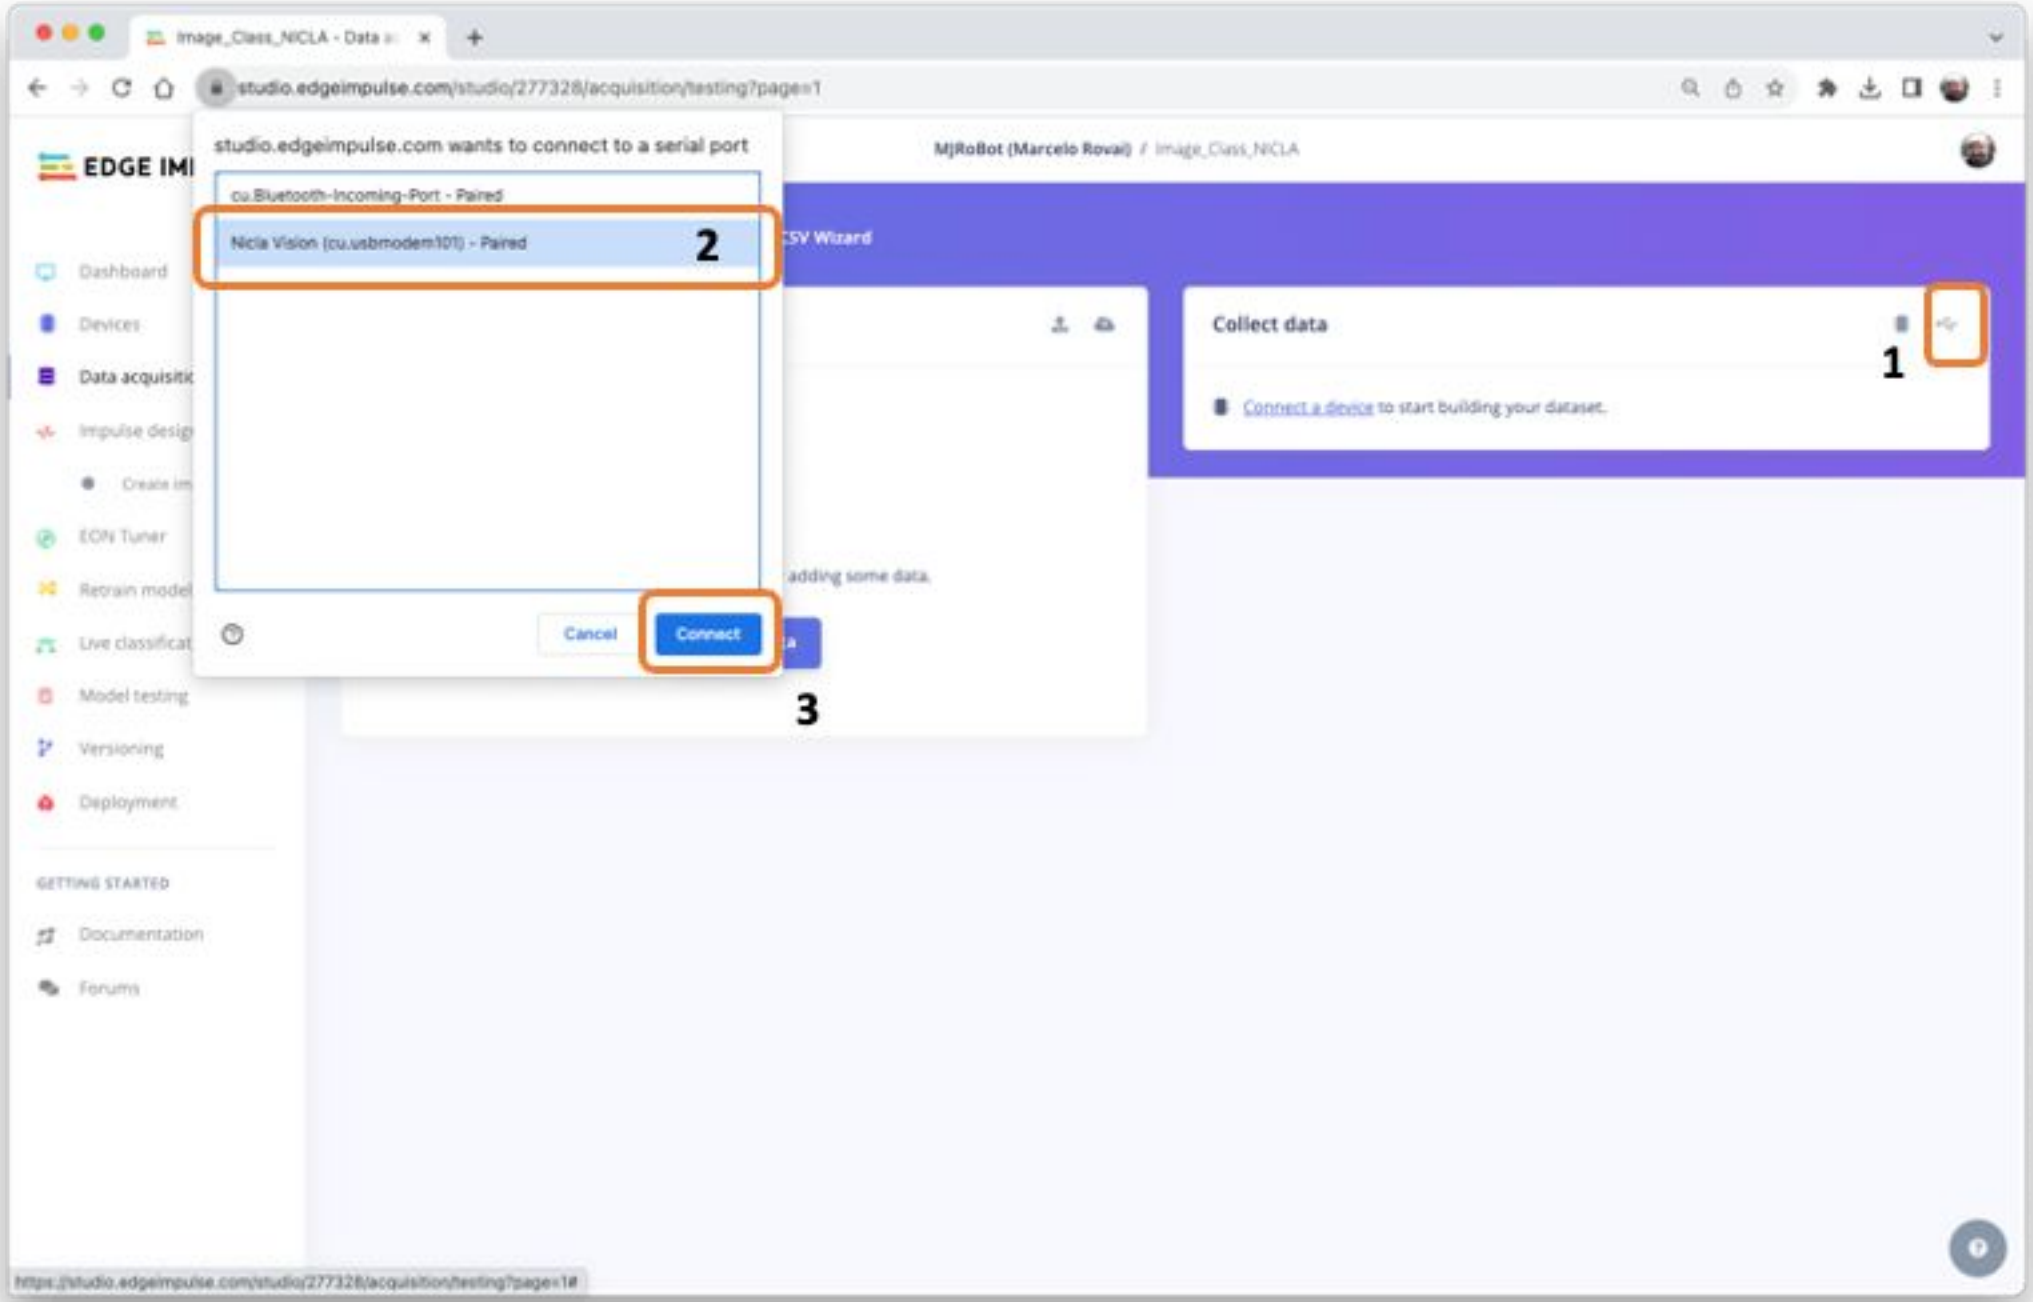
studio.edgeimpulse.com wants to connect to a serial port
cu.Bluetooth-Incoming-Port - Paired
Nicla Vision (cu.usbmodem101) - Paired
2
Cancel
Connect
3
Collect data
1
Connect a device to start building your dataset.
MjRoBot (Marcelo Rovai) / Image_Class_NICLA
Dashboard
Devices
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums

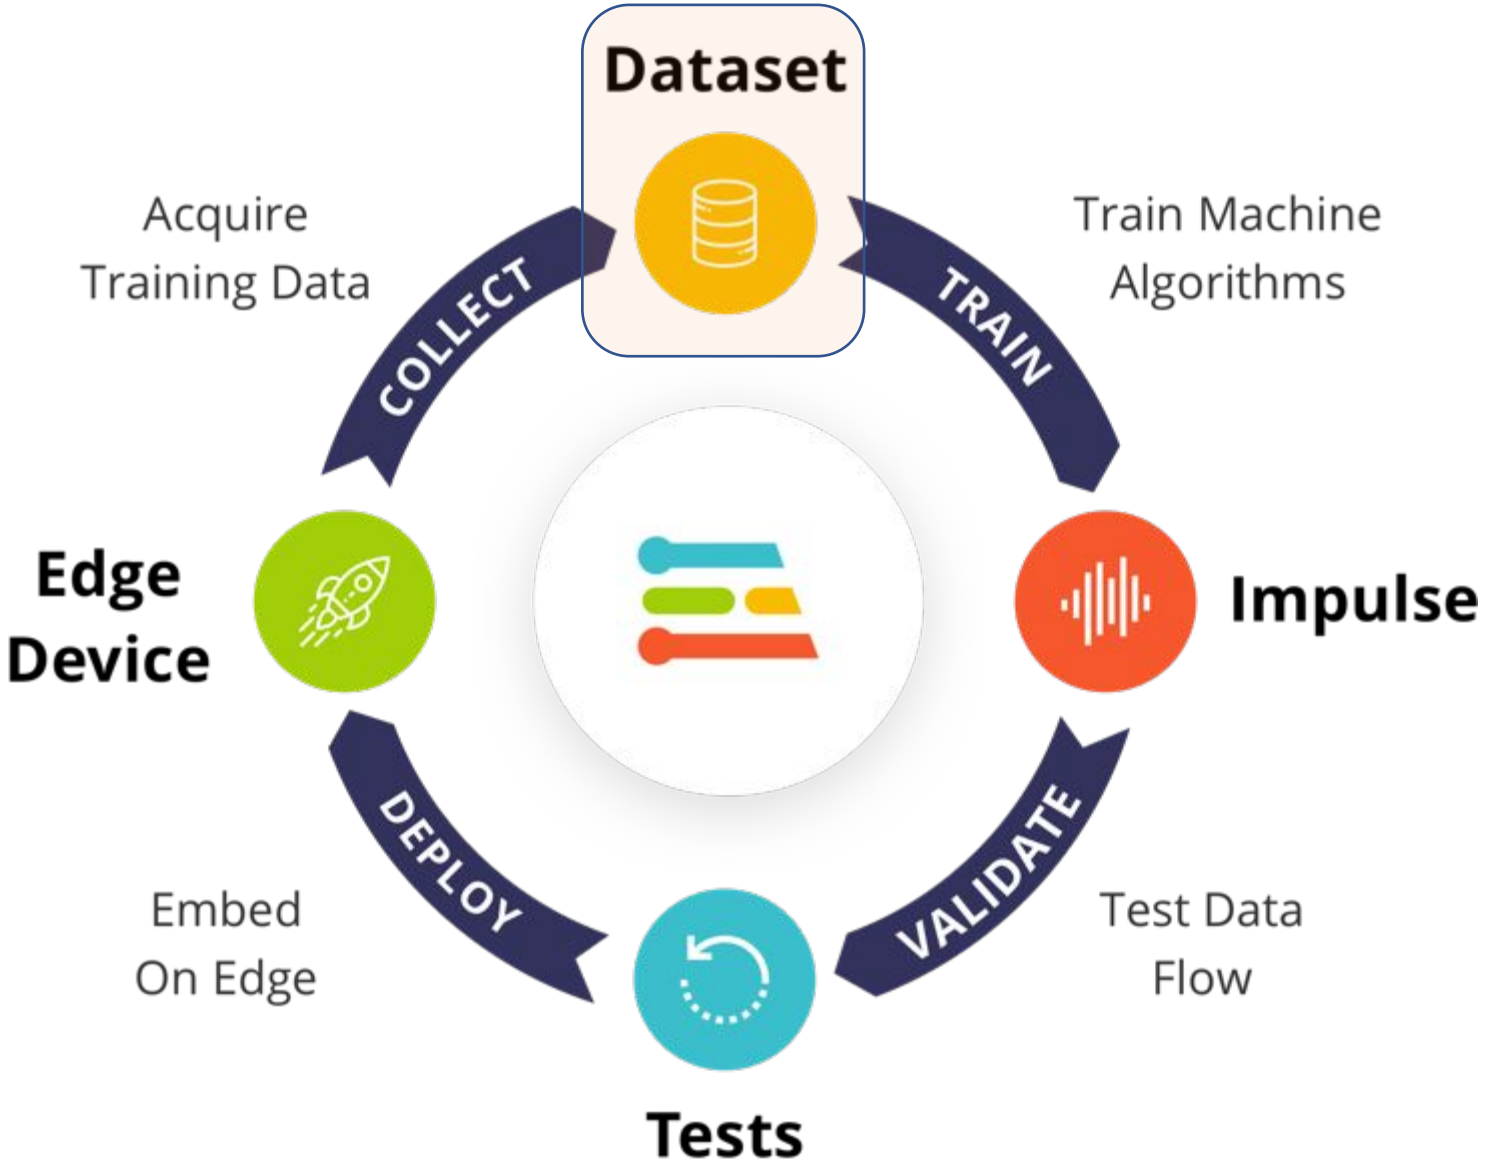

- Pre-Processing Data
- Design a Model
- Train a Model

Collect Data

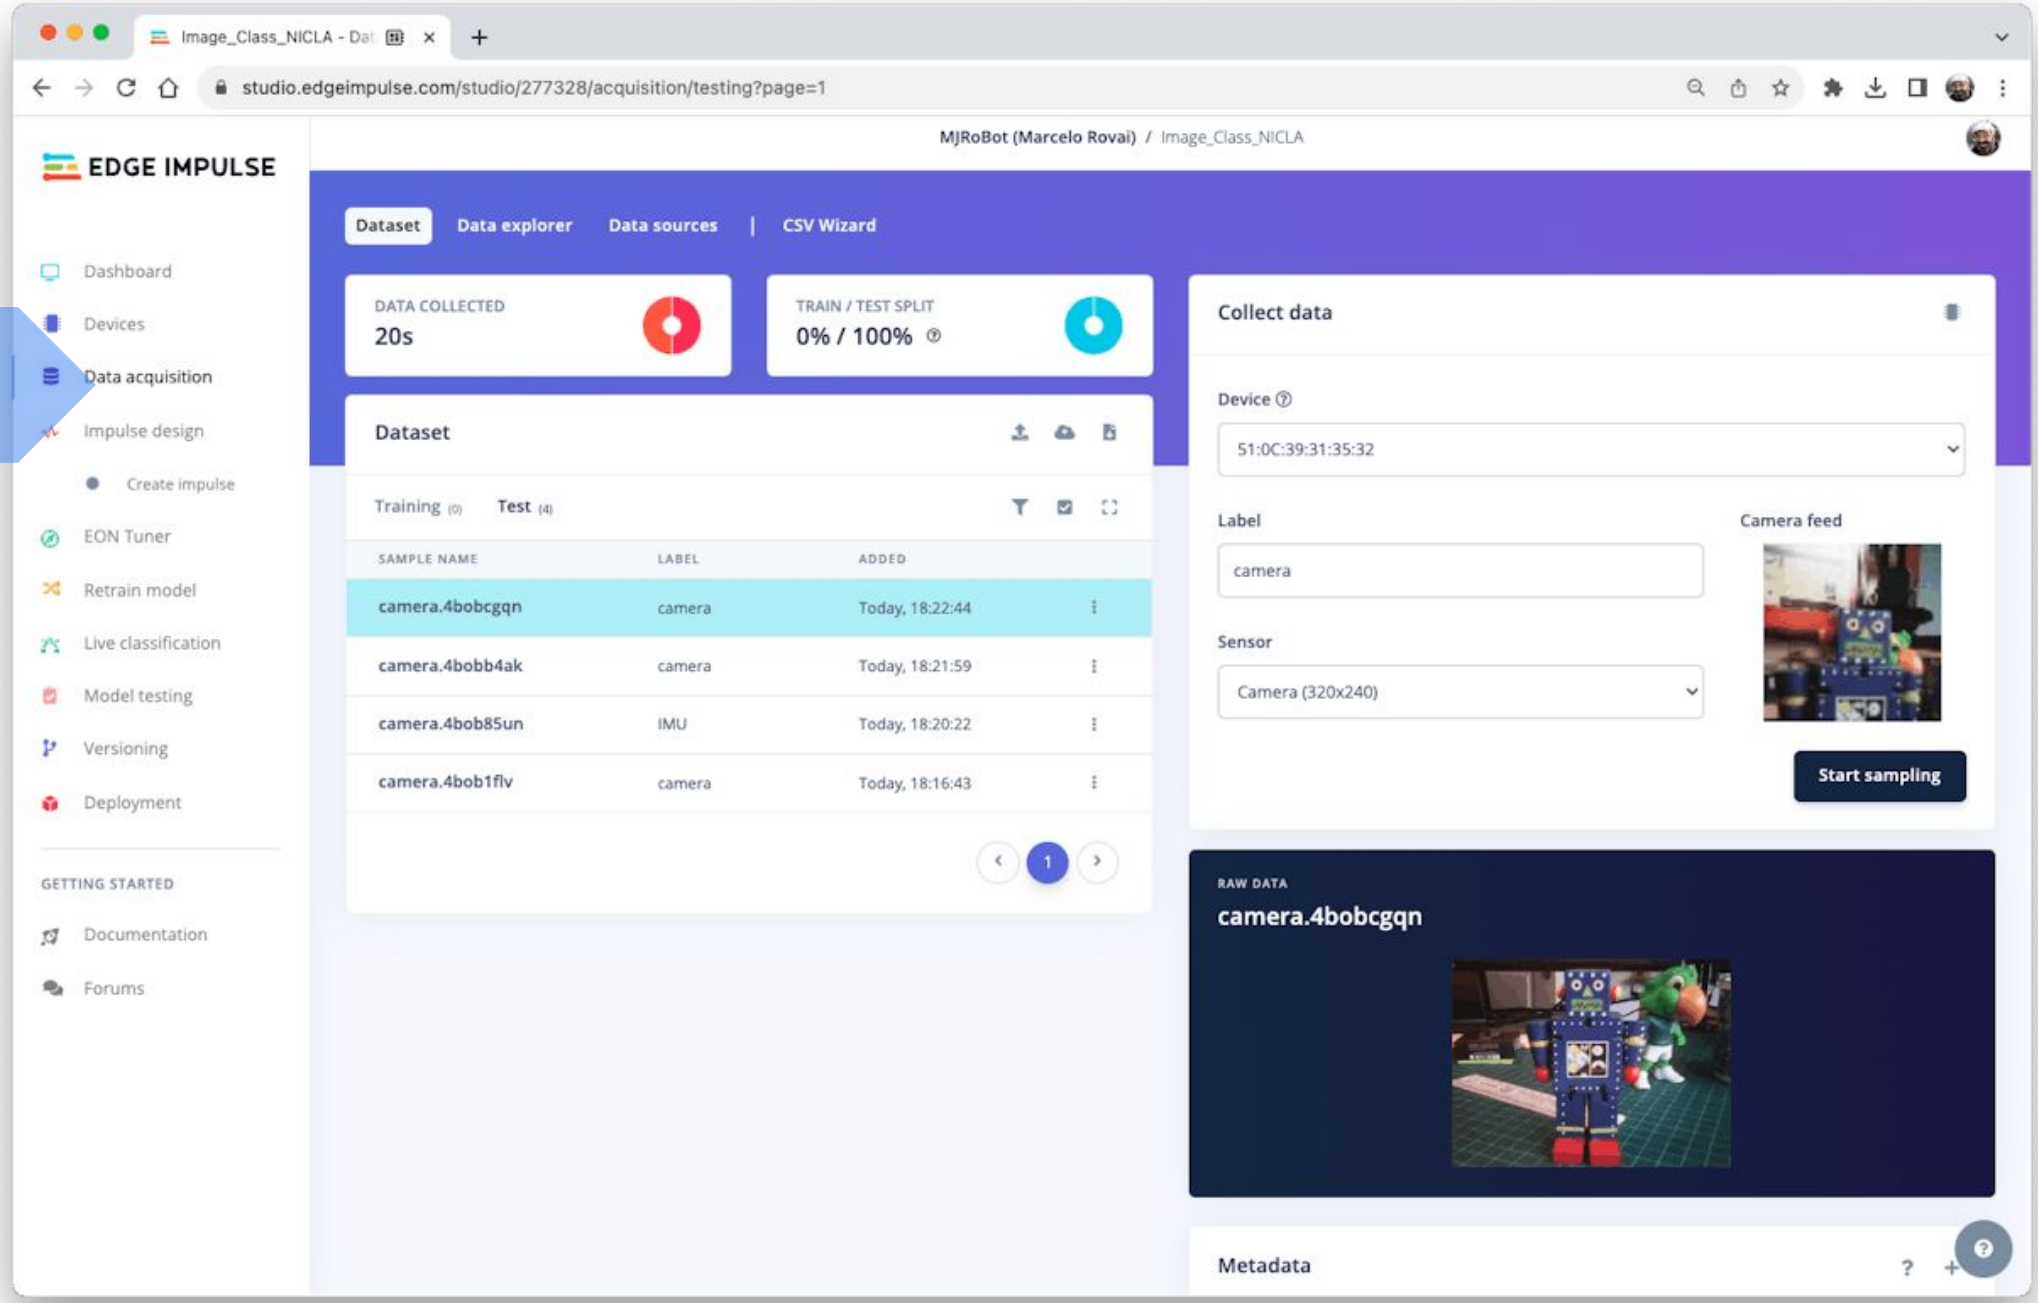

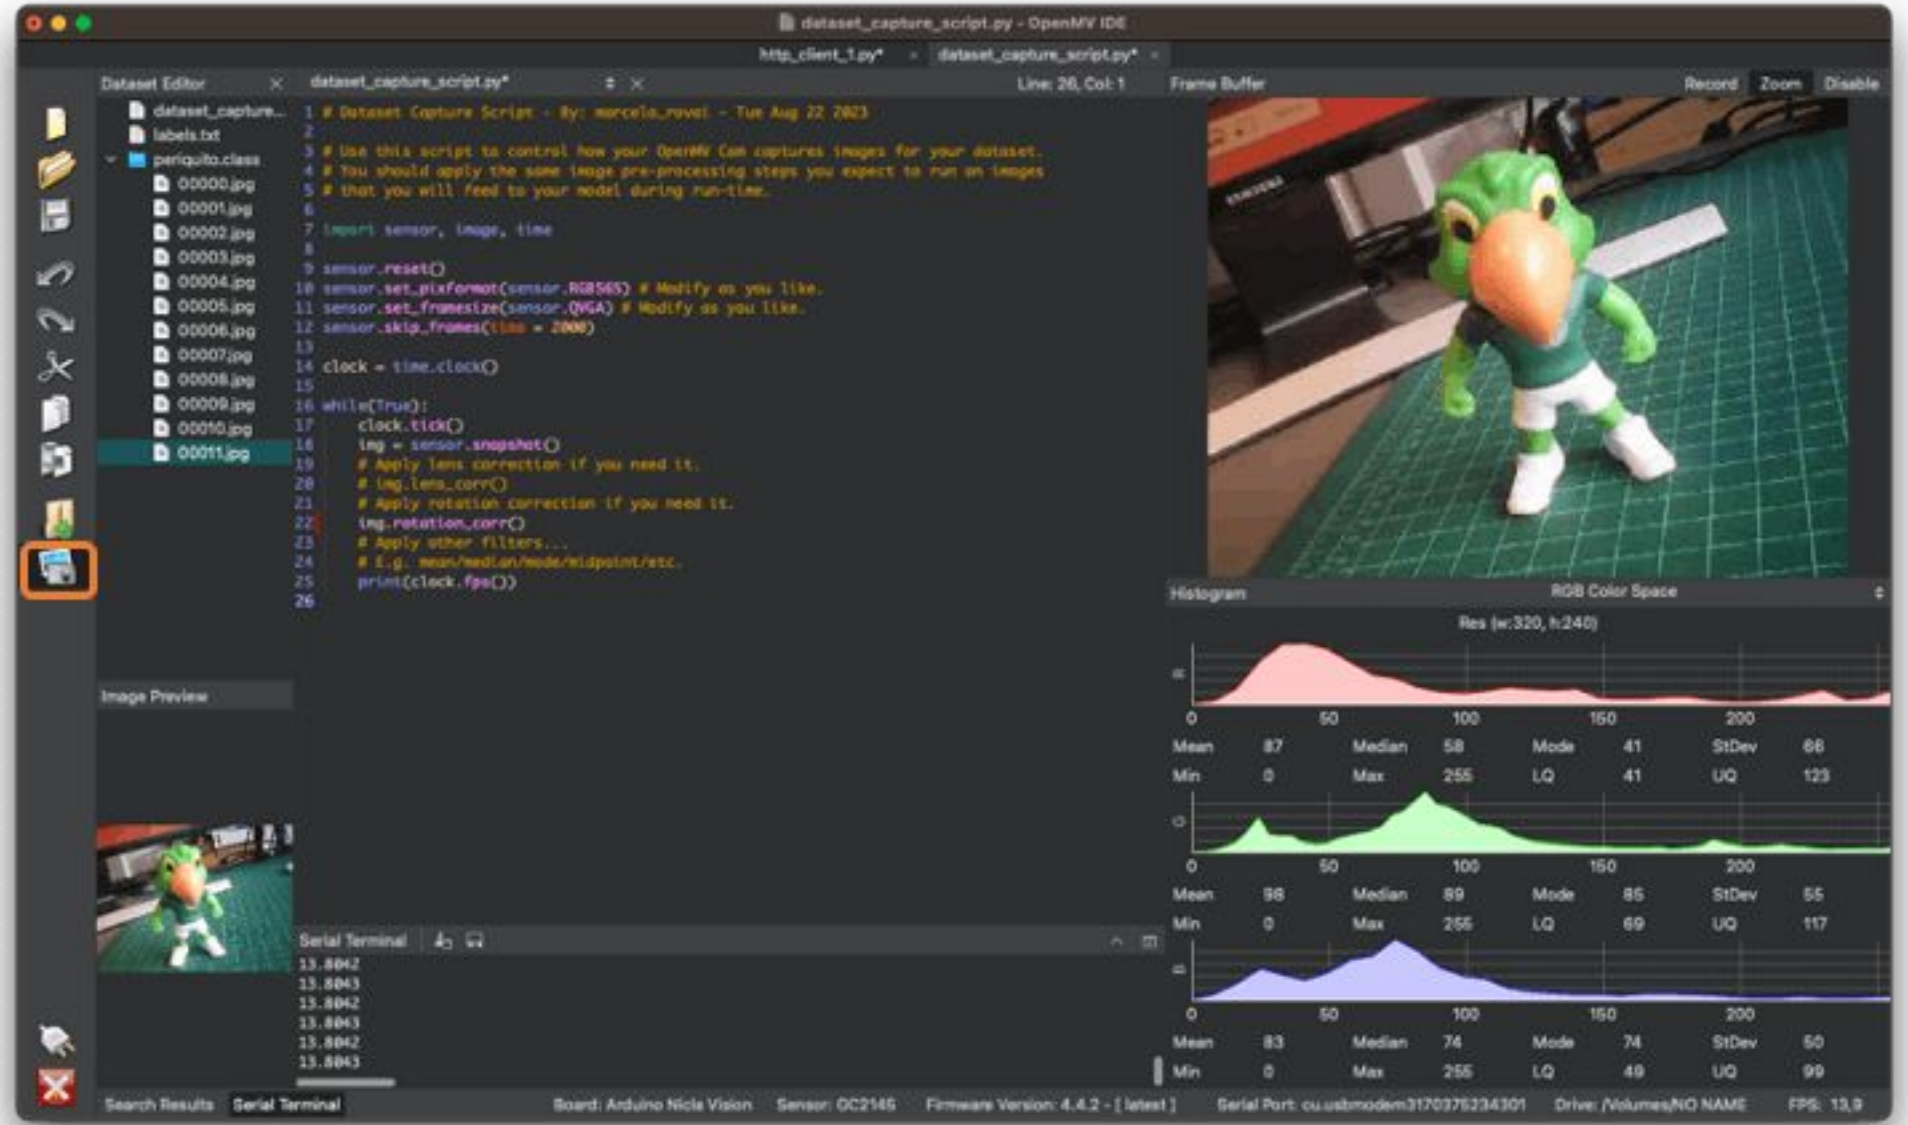
dataset_capture_script.py - OpenMV IDE
Dataset Editor
dataset_capture_script.py*
Line: 26, Col: 1
Frame Buffer
Record
Zoom
Disable
dataset_capture...
labels.txt
periquito.class
00000.jpg
00001.jpg
00002.jpg
00003.jpg
00004.jpg
00005.jpg
00006.jpg
00007.jpg
00008.jpg
00009.jpg
00010.jpg
00011.jpg
# Dataset Capture Script - By: marcela_rovai - Tue Aug 22 2023
# Use this script to control how your OpenMV Cam captures images for your dataset.
# You should apply the same image pre-processing steps you expect to run on images
# that you will feed to your model during run-time.
import sensor, image, time
sensor.reset()
sensor.set_pixformat(sensor.RGB565) # Modify as you like.
sensor.set_framesize(sensor.QVGA) # Modify as you like.
sensor.skip_frames(time = 2000)
clock = time.clock()
while(True):
clock.tick()
img = sensor.snapshot()
# Apply lens correction if you need it.
# img.lens_corr()
# Apply rotation correction if you need it.
img.rotation_corr()
# Apply other filters...
# E.g. mean/median/mode/midpoint/etc.
print(clock.fps())
Image Preview
Histogram
RGB Color Space
Res (w:320, h:240)
Mean 87 Median 58 Mode 41 StDev 66
Min 0 Max 255 LQ 41 UQ 123
Mean 98 Median 89 Mode 85 StDev 55
Min 0 Max 255 LQ 69 UQ 117
Mean 83 Median 74 Mode 74 StDev 50
Min 0 Max 255 LQ 49 UQ 99
Serial Terminal
13.8042
13.8043
13.8042
13.8043
13.8042
13.8043
Search Results
Serial Terminal
Board: Arduino Nicla Vision
Sensor: OC2145
Firmware Version: 4.4.2 - [latest]
Drive: /Volumes/NO NAME
FPS: 13.9

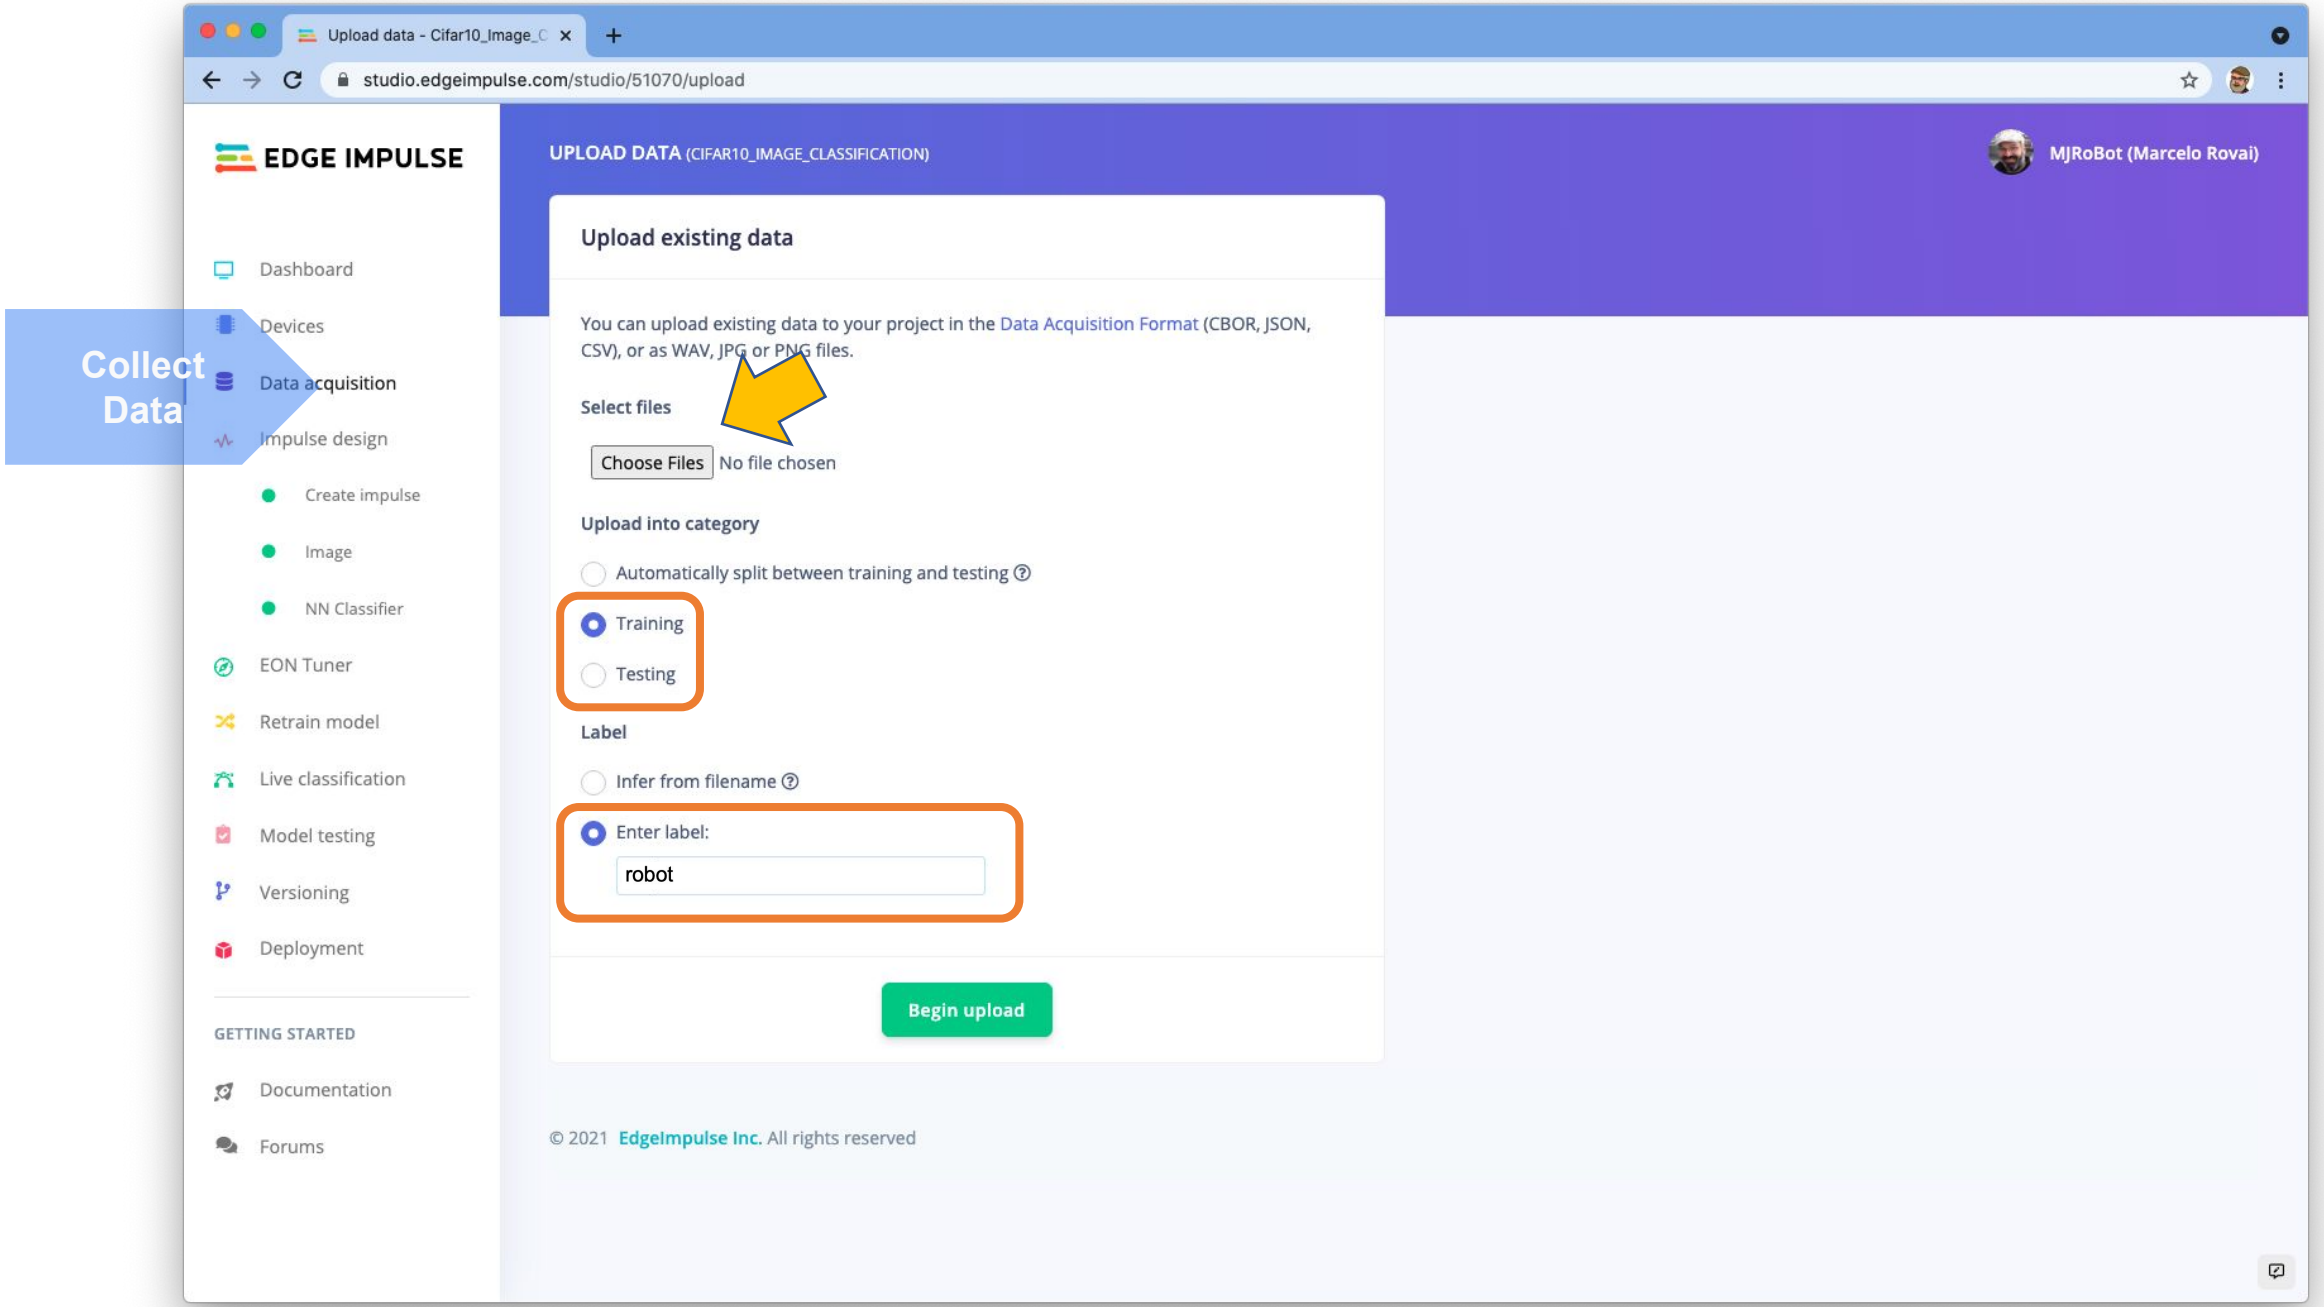
Upload data - Cifar10_Image_C
studio.edgeimpulse.com/studio/51070/upload
EDGE IMPULSE
UPLOAD DATA (CIFAR10_IMAGE_CLASSIFICATION)
MJRoBot (Marcelo Rovai)
Collect Data
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
NN Classifier
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Upload existing data
You can upload existing data to your project in the Data Acquisition Format (CBOR, JSON, CSV), or as WAV, JPG or PNG files.
Select files
Choose Files
No file chosen
Upload into category
Automatically split between training and testing
Training
Testing
Label
Infer from filename
Enter label:
robot
Begin upload
© 2021 EdgeImpulse Inc. All rights reserved

Collect Data

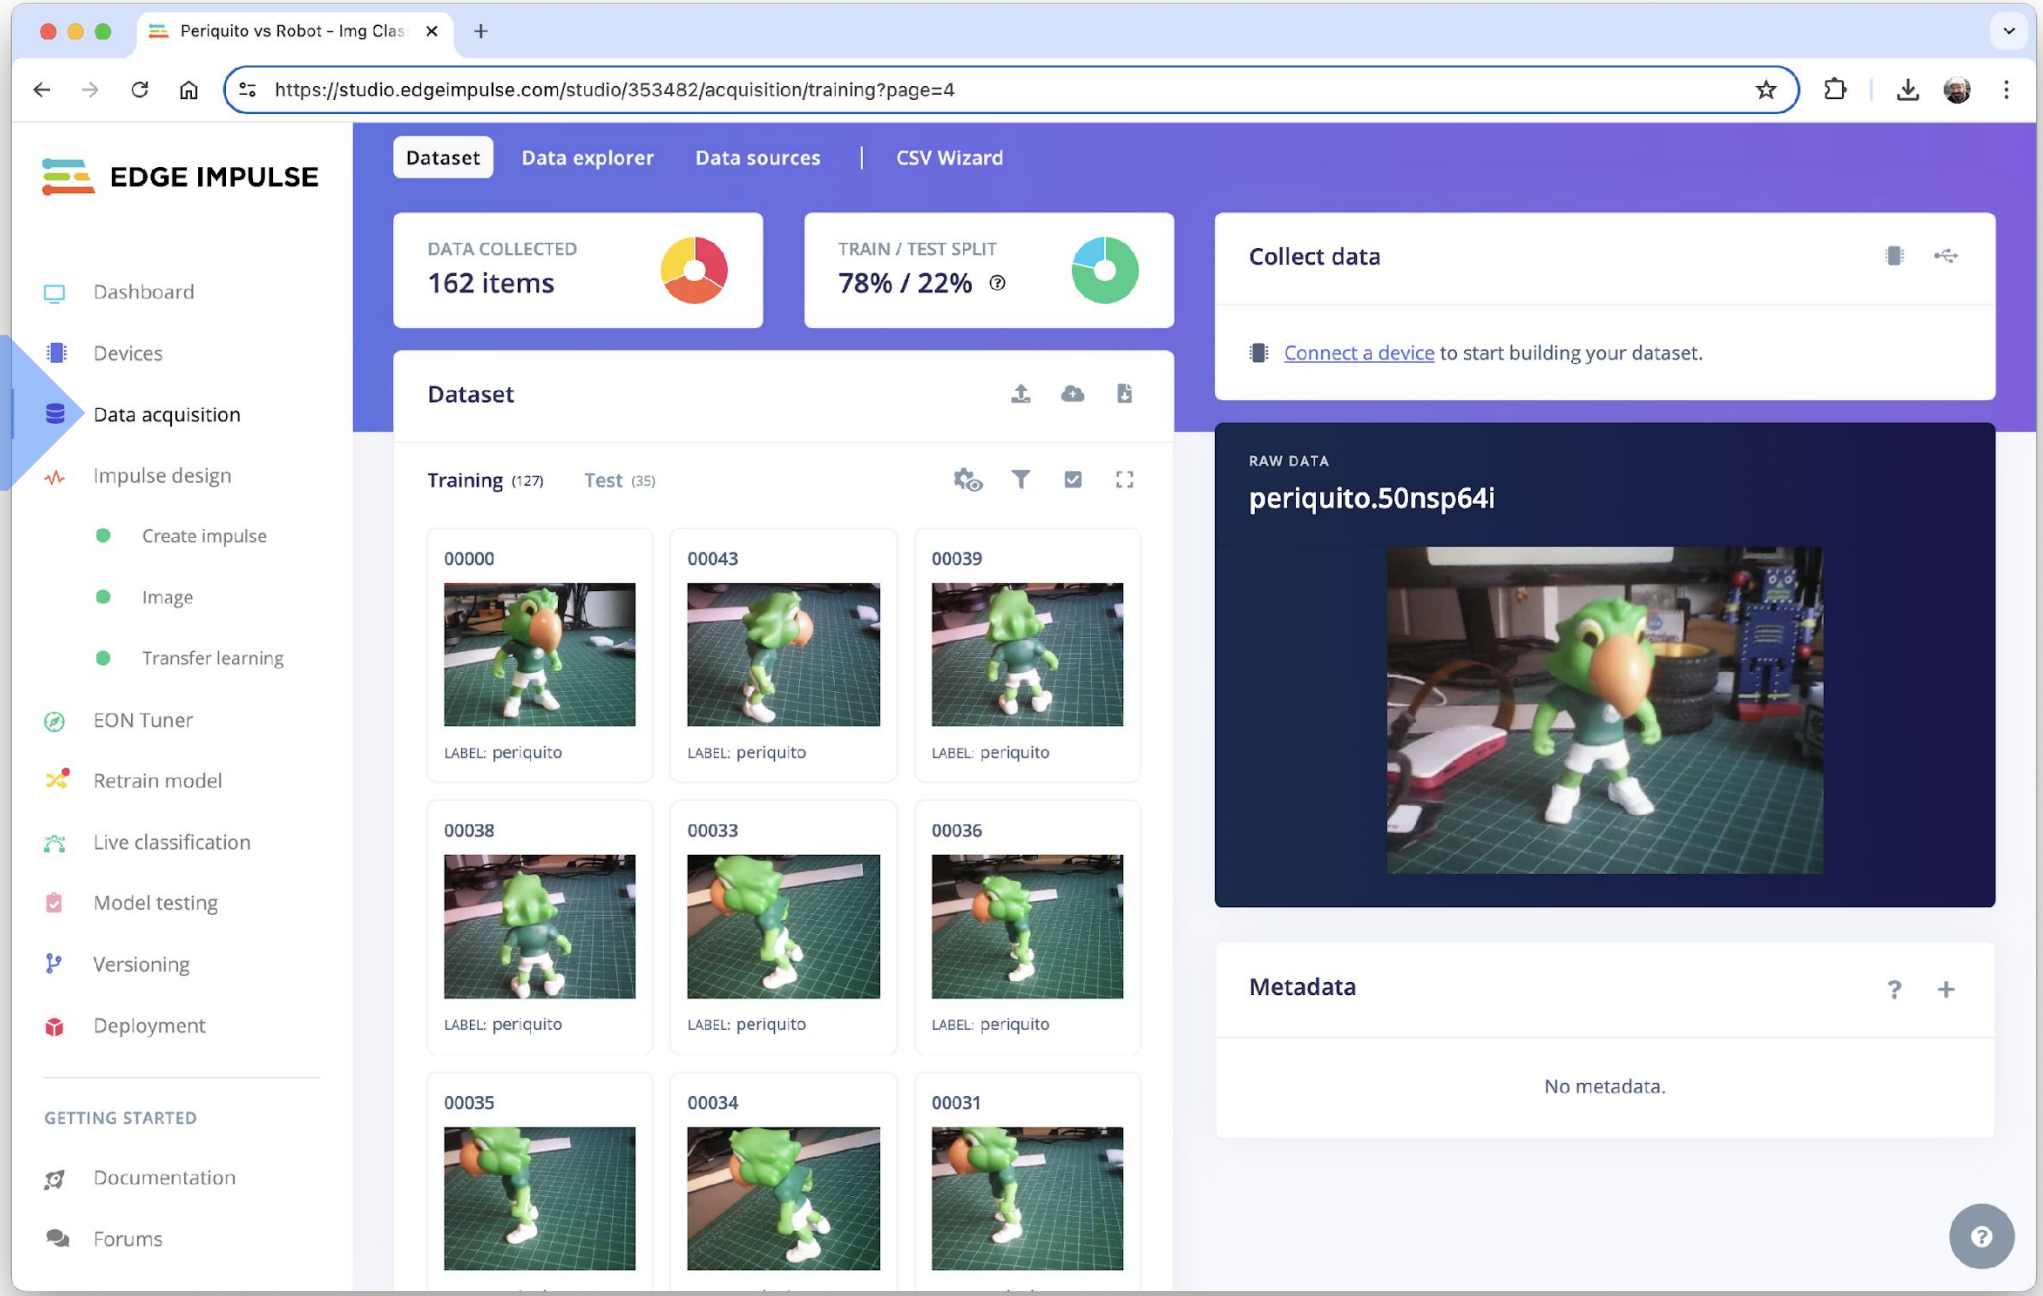

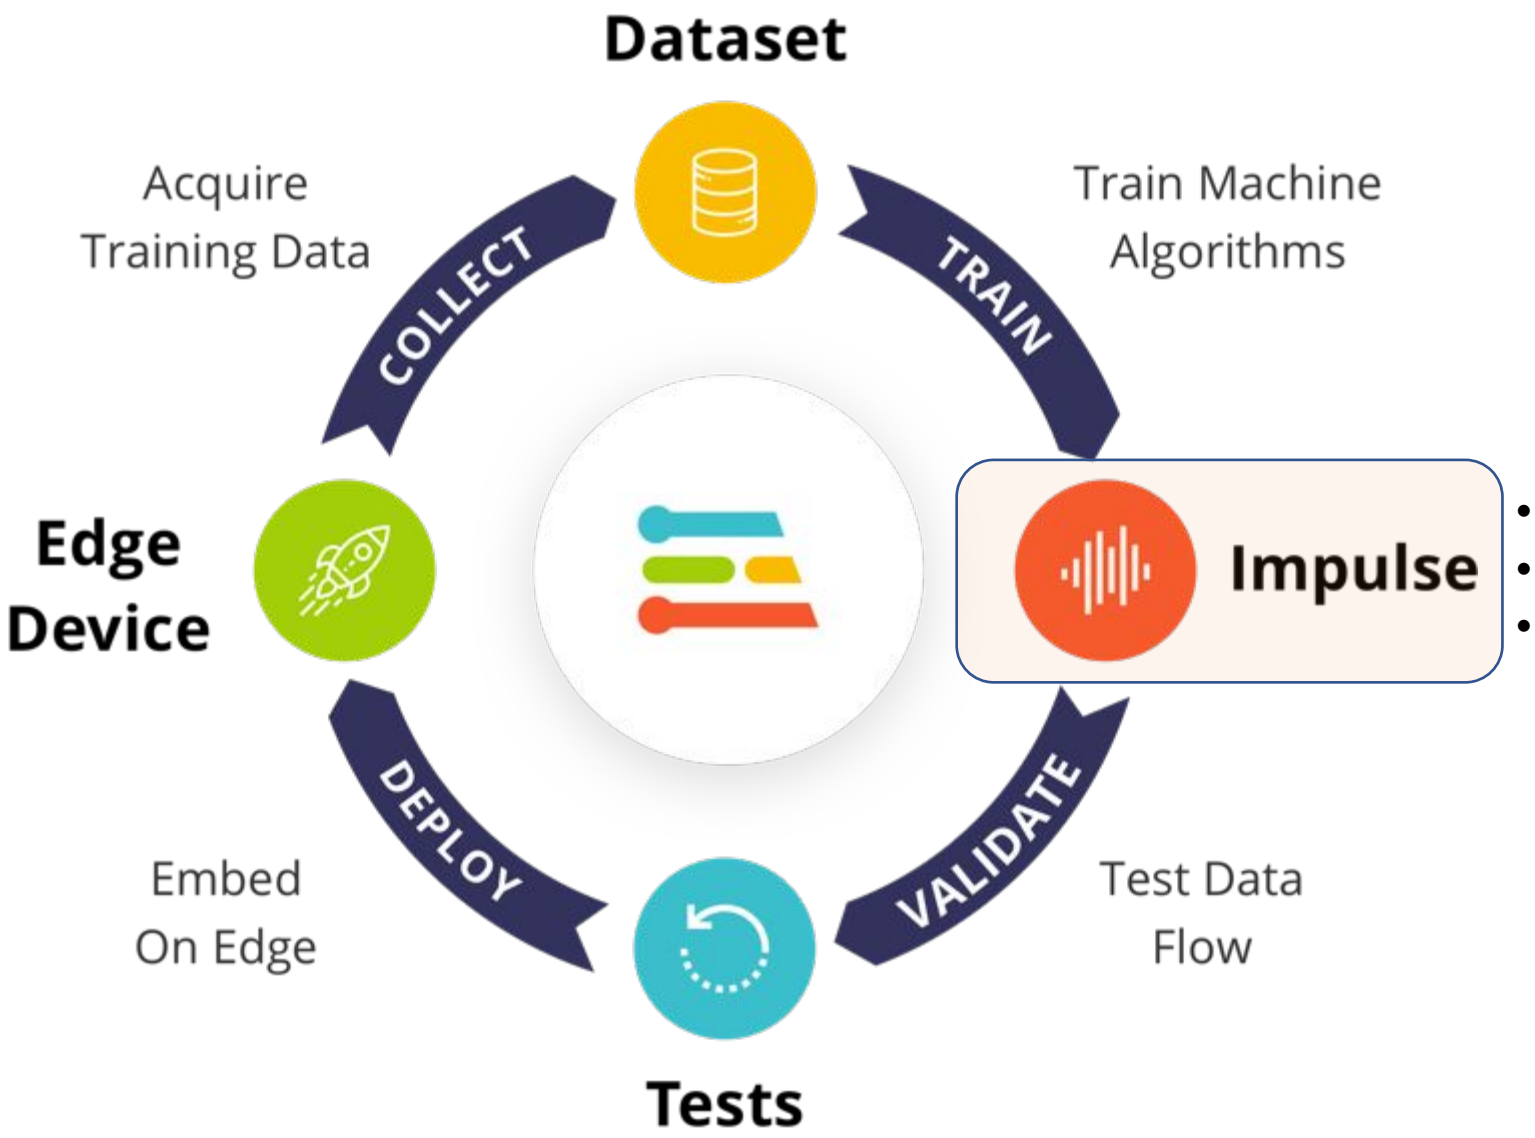

- Pre-Processing Data
- Design a Model
- Train a Model

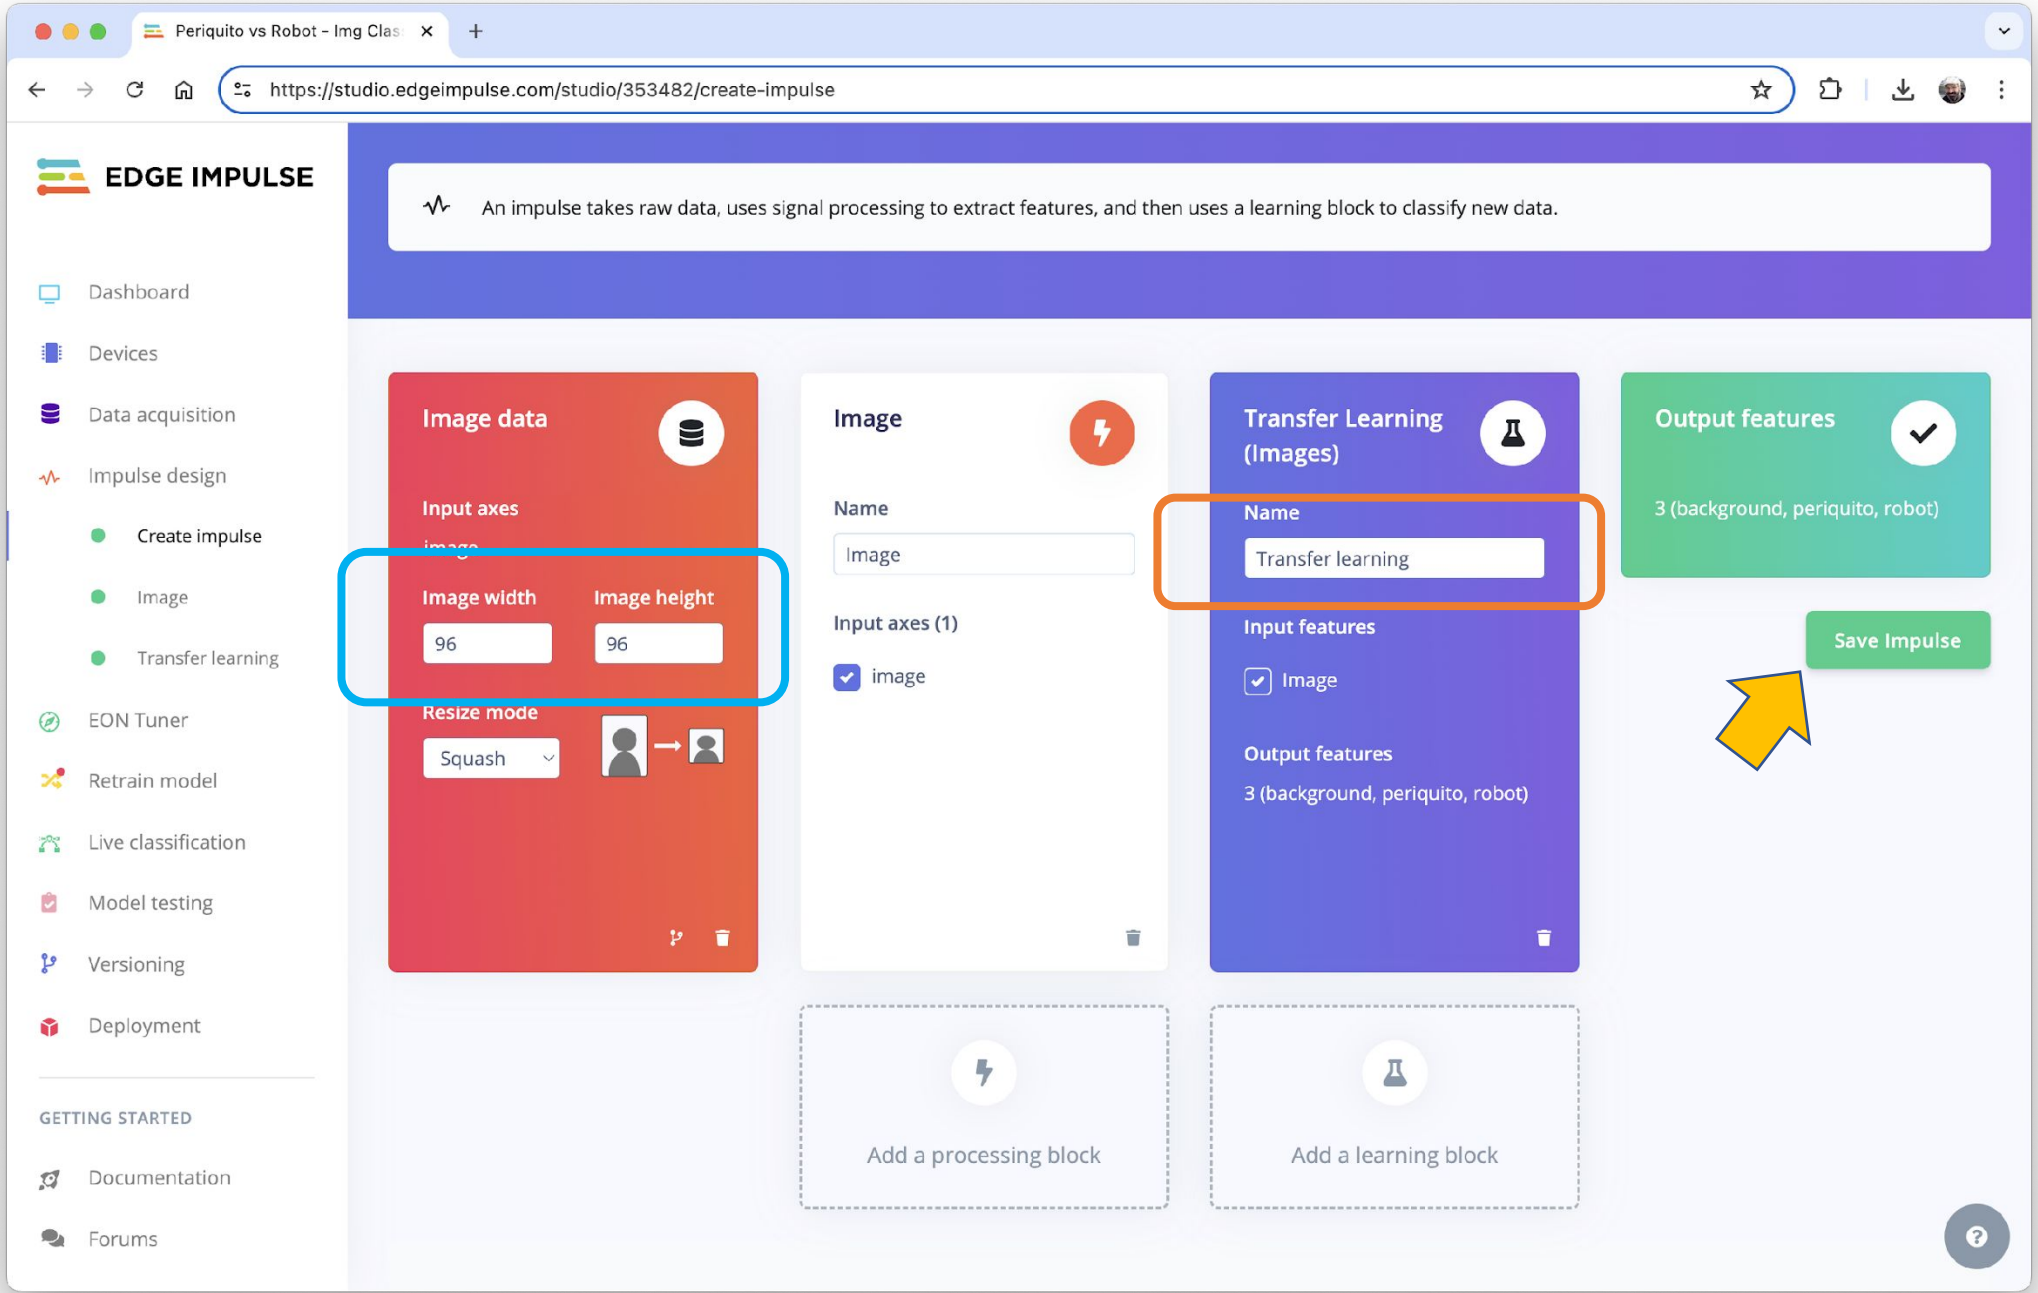
Periquito vs Robot - Img Clas
https://studio.edgeimpulse.com/studio/353482/create-impulse
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
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 width
96
Image height
96
Resize mode
Squash
Image
Name
Image
Input axes (1)
image
Transfer Learning (Images)
Name
Transfer learning
Input features
Image
Output features
3 (background, periquito, robot)
Output features
3 (background, periquito, robot)
Save Impulse
Add a processing block
Add a learning block

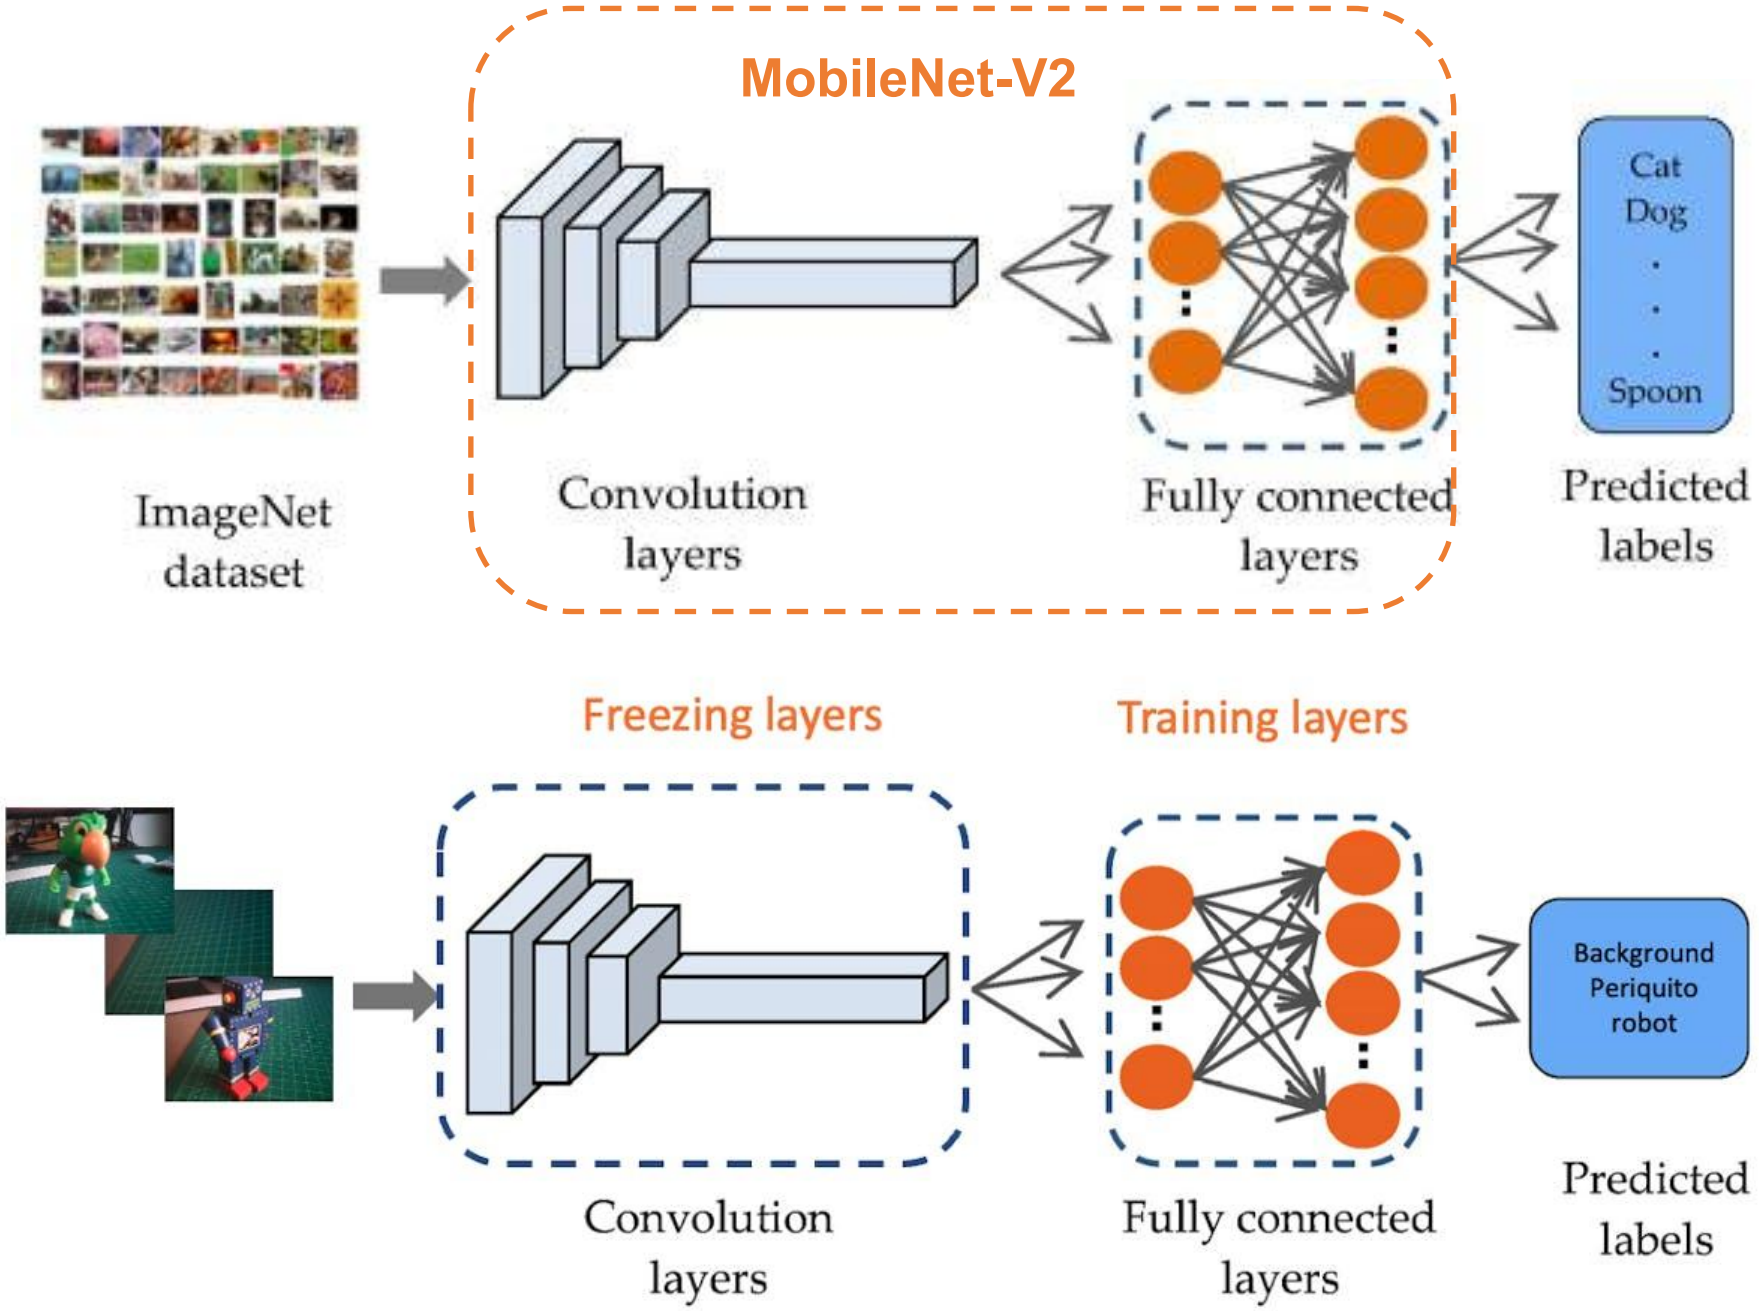
MobileNet-V2
ImageNet dataset
Convolution layers
Fully connected layers
Cat
Dog
Spoon
Predicted labels
Freezing layers
Training layers
Convolution layers
Fully connected layers
Background
Periquito
robot
Predicted labels

Preprocess Data

Periquito vs Robot - Img Clas

https://studio.edgeimpulse.com/studio/353482/dsp/image/3

EDGE IMPULSE

- 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

Raw data

Show: All labels | periquito.50nsp64i (periqu

Raw features

0x2a241f, 0x2e2422, 0x302423, 0x312727, 0x30282a, 0x322a29, 0x3b2e28, 0…

Parameters

Image

Color depth: RGB

Save parameters

DSP result

Image

96x96x3
= 27,648

Copy 27648 features to clipboard

Processed features

0.1647, 0.1412, 0.1216, 0.1804, 0.1412, 0.1333, 0.1882, 0.1412, 0.1373,…

On-device performance

PROCESSING TIME
1 ms.

PEAK RAM USAGE
4 KB

https://studio.edgeimpulse.com/studio/353482/dsp/image/3#

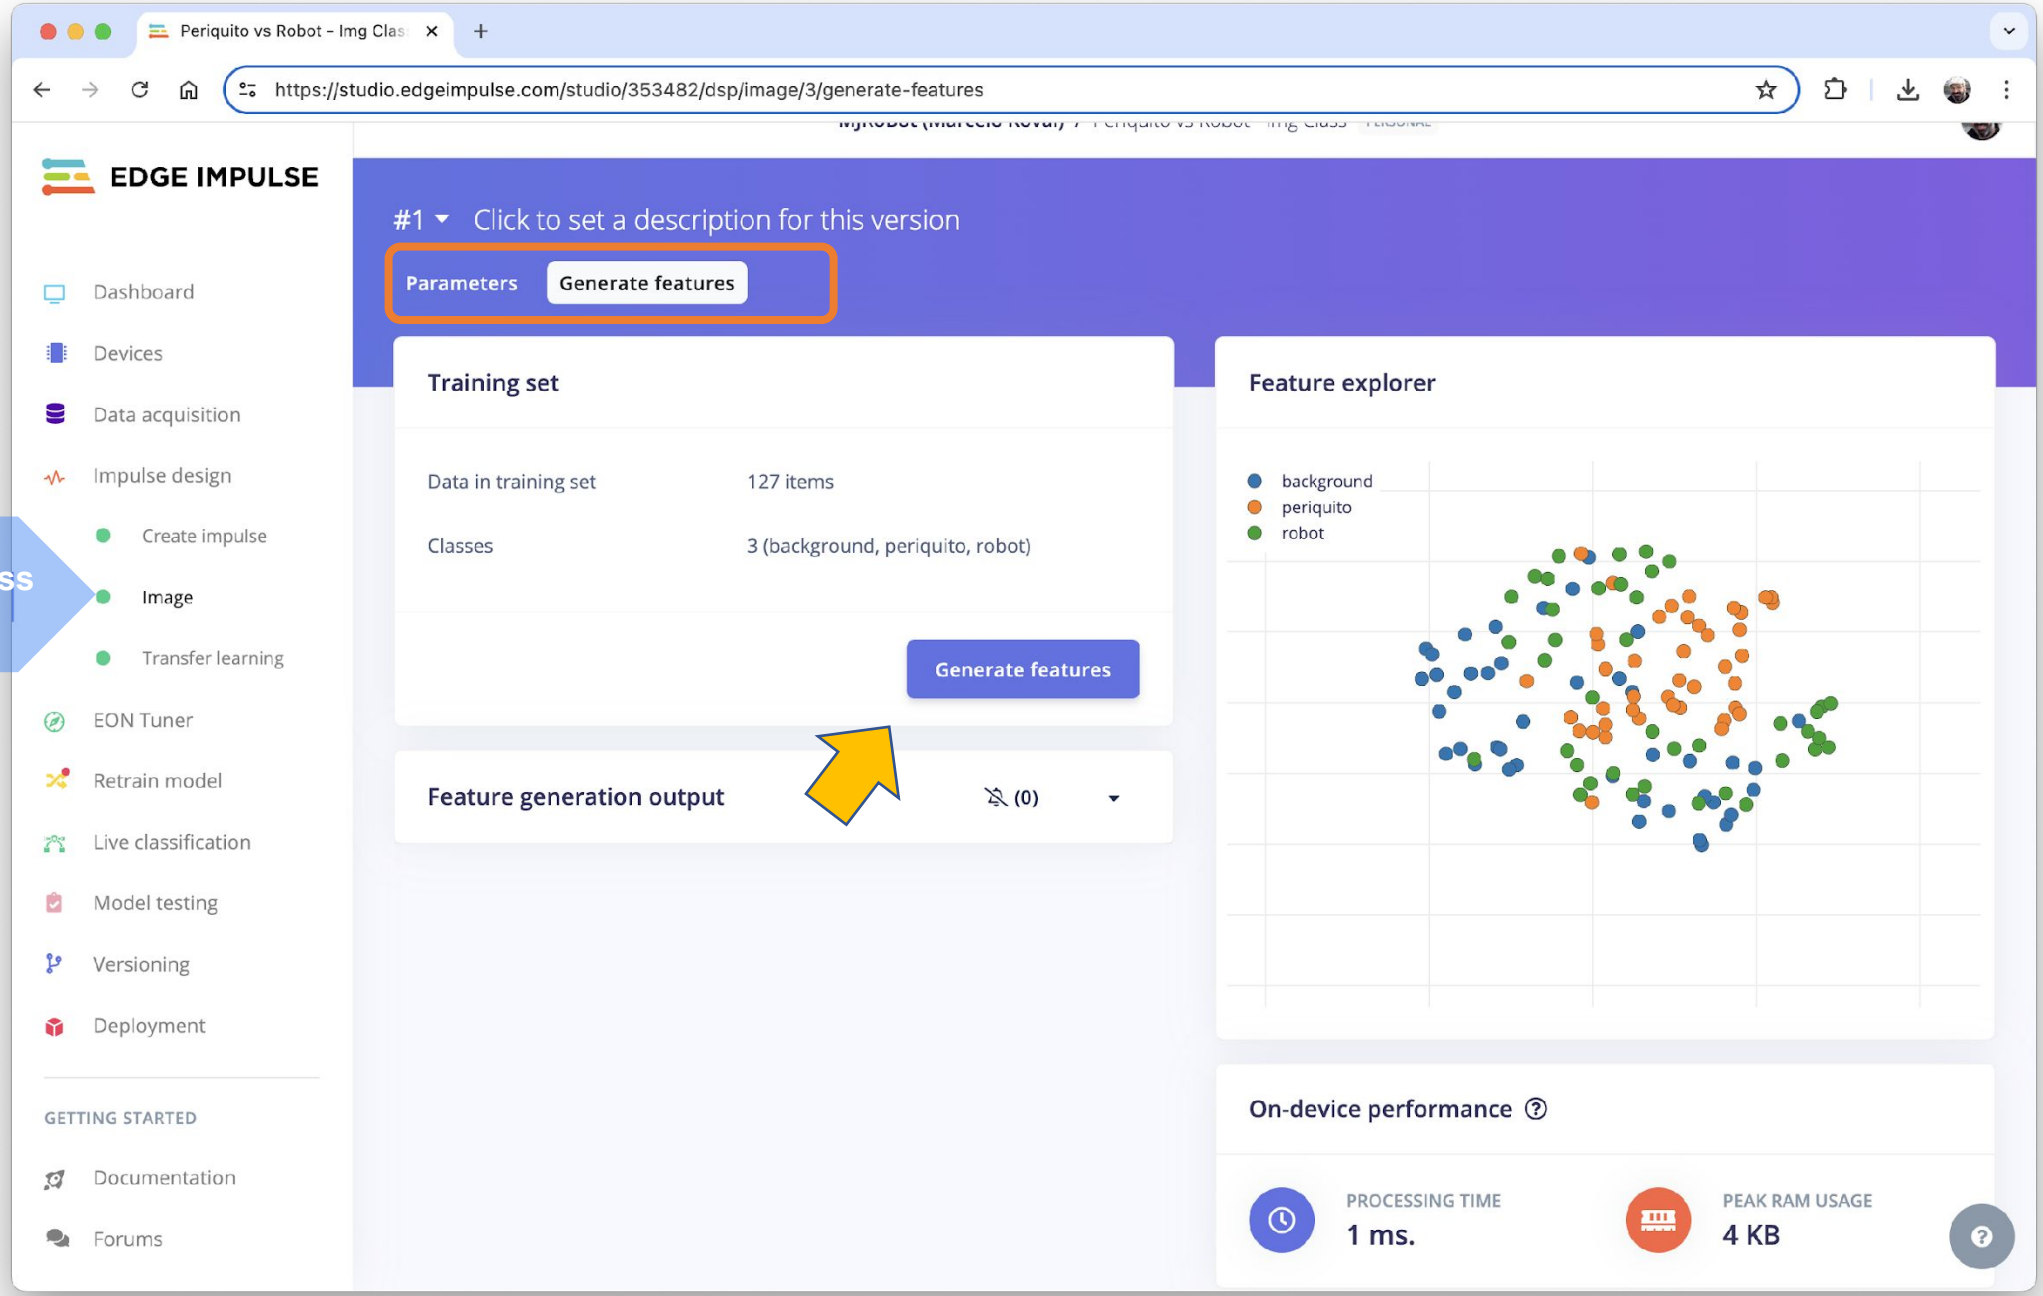

Preprocess Data

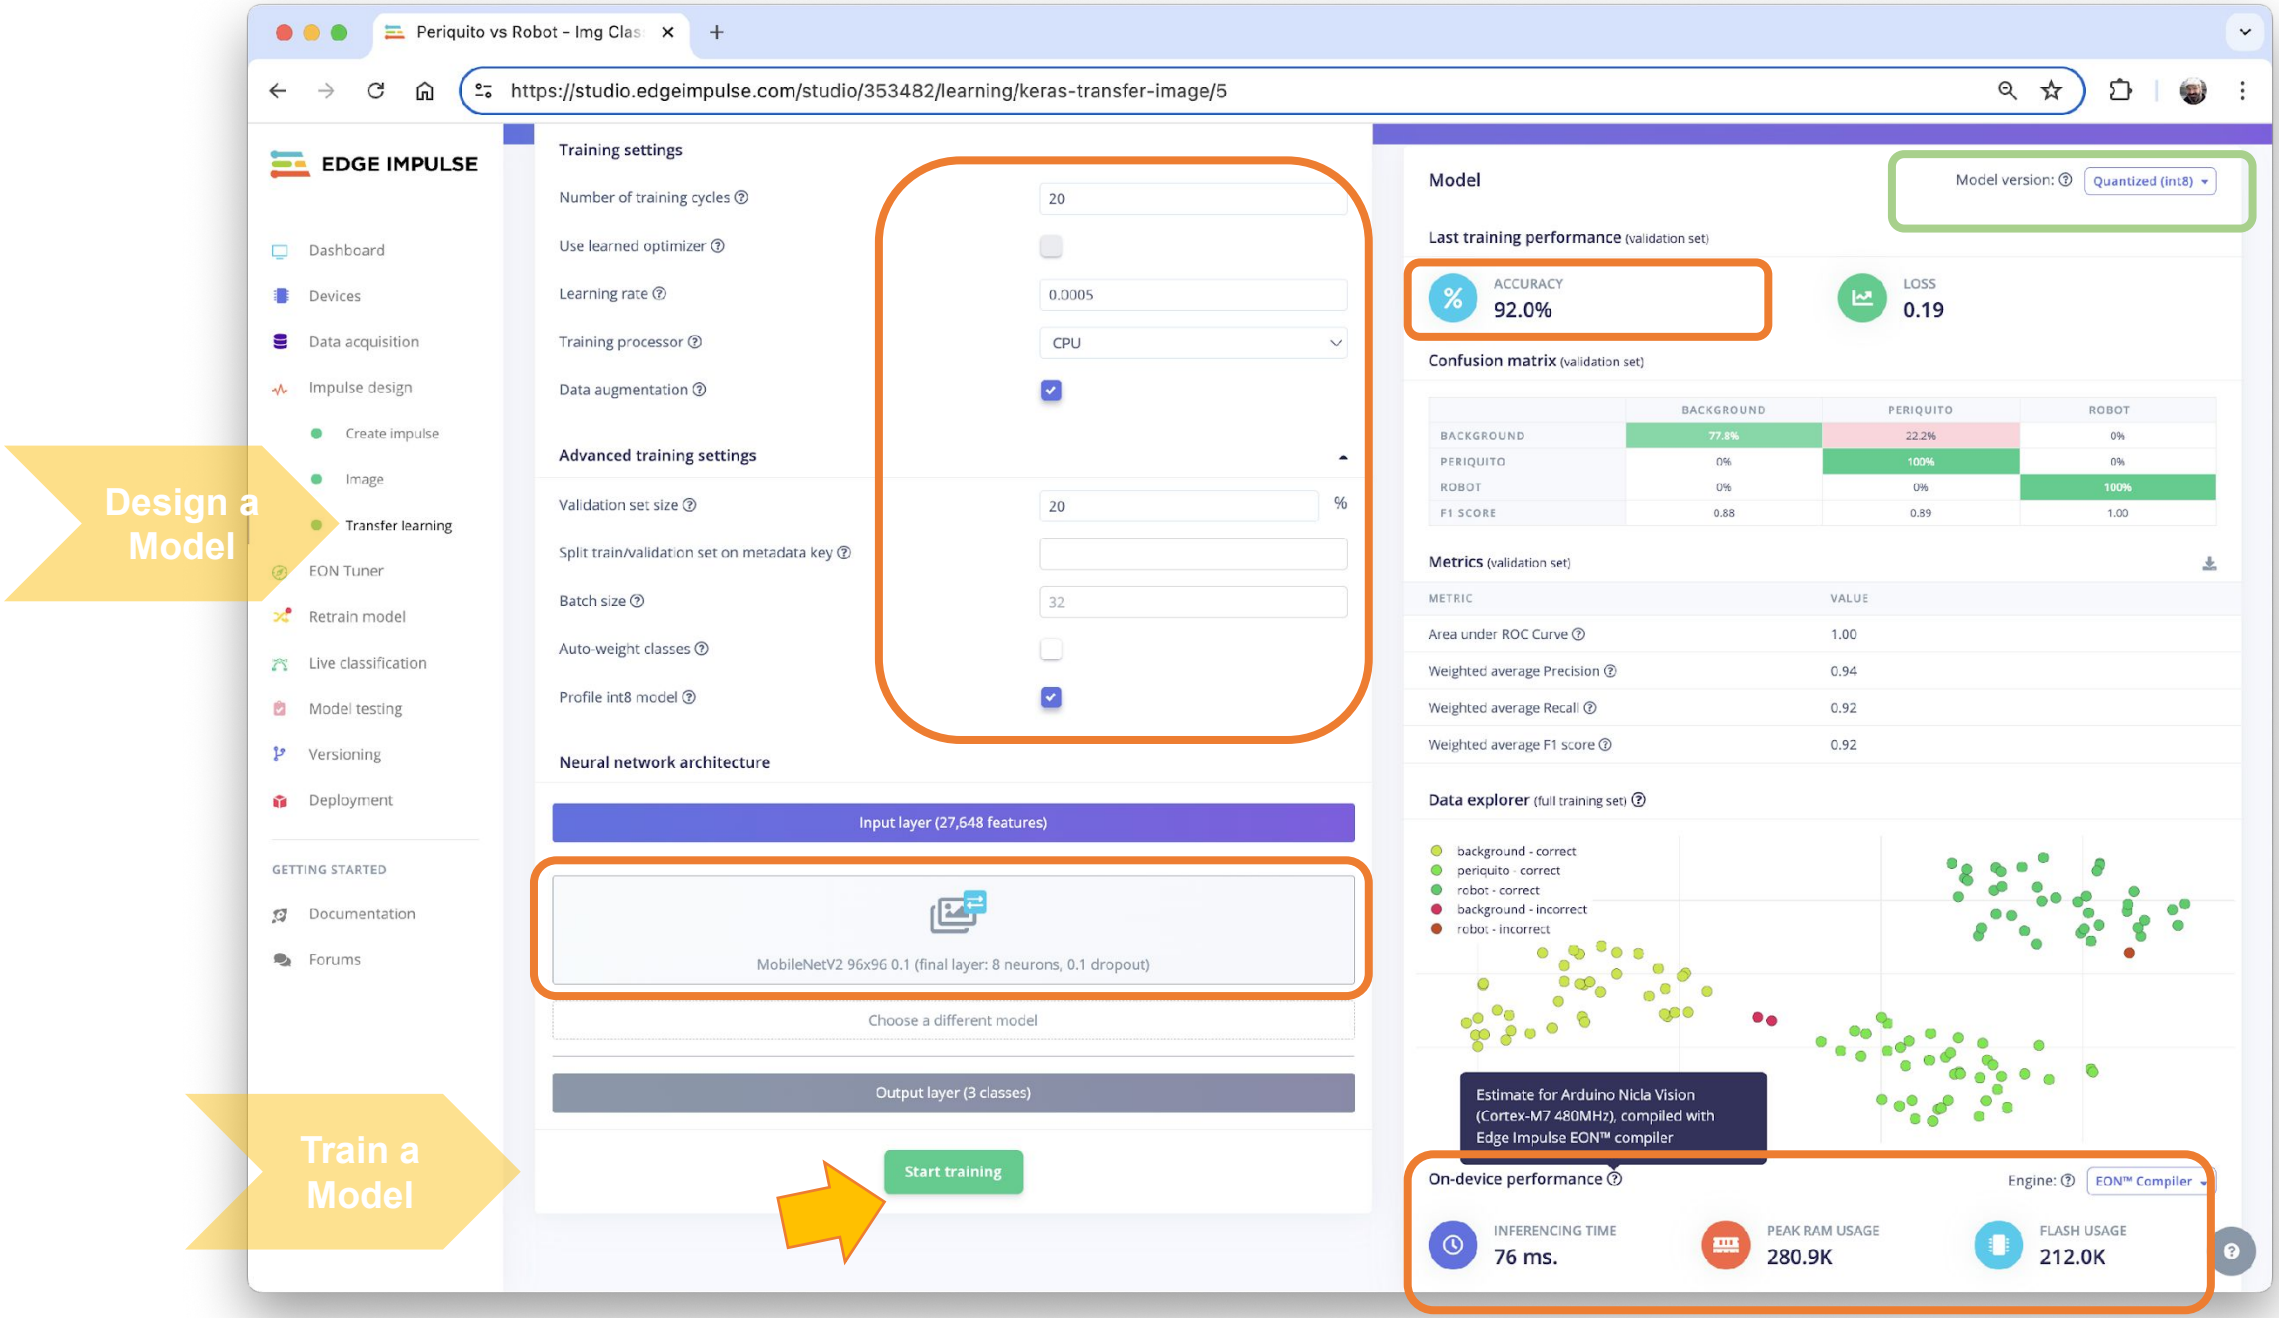

Design a Model

Train a Model

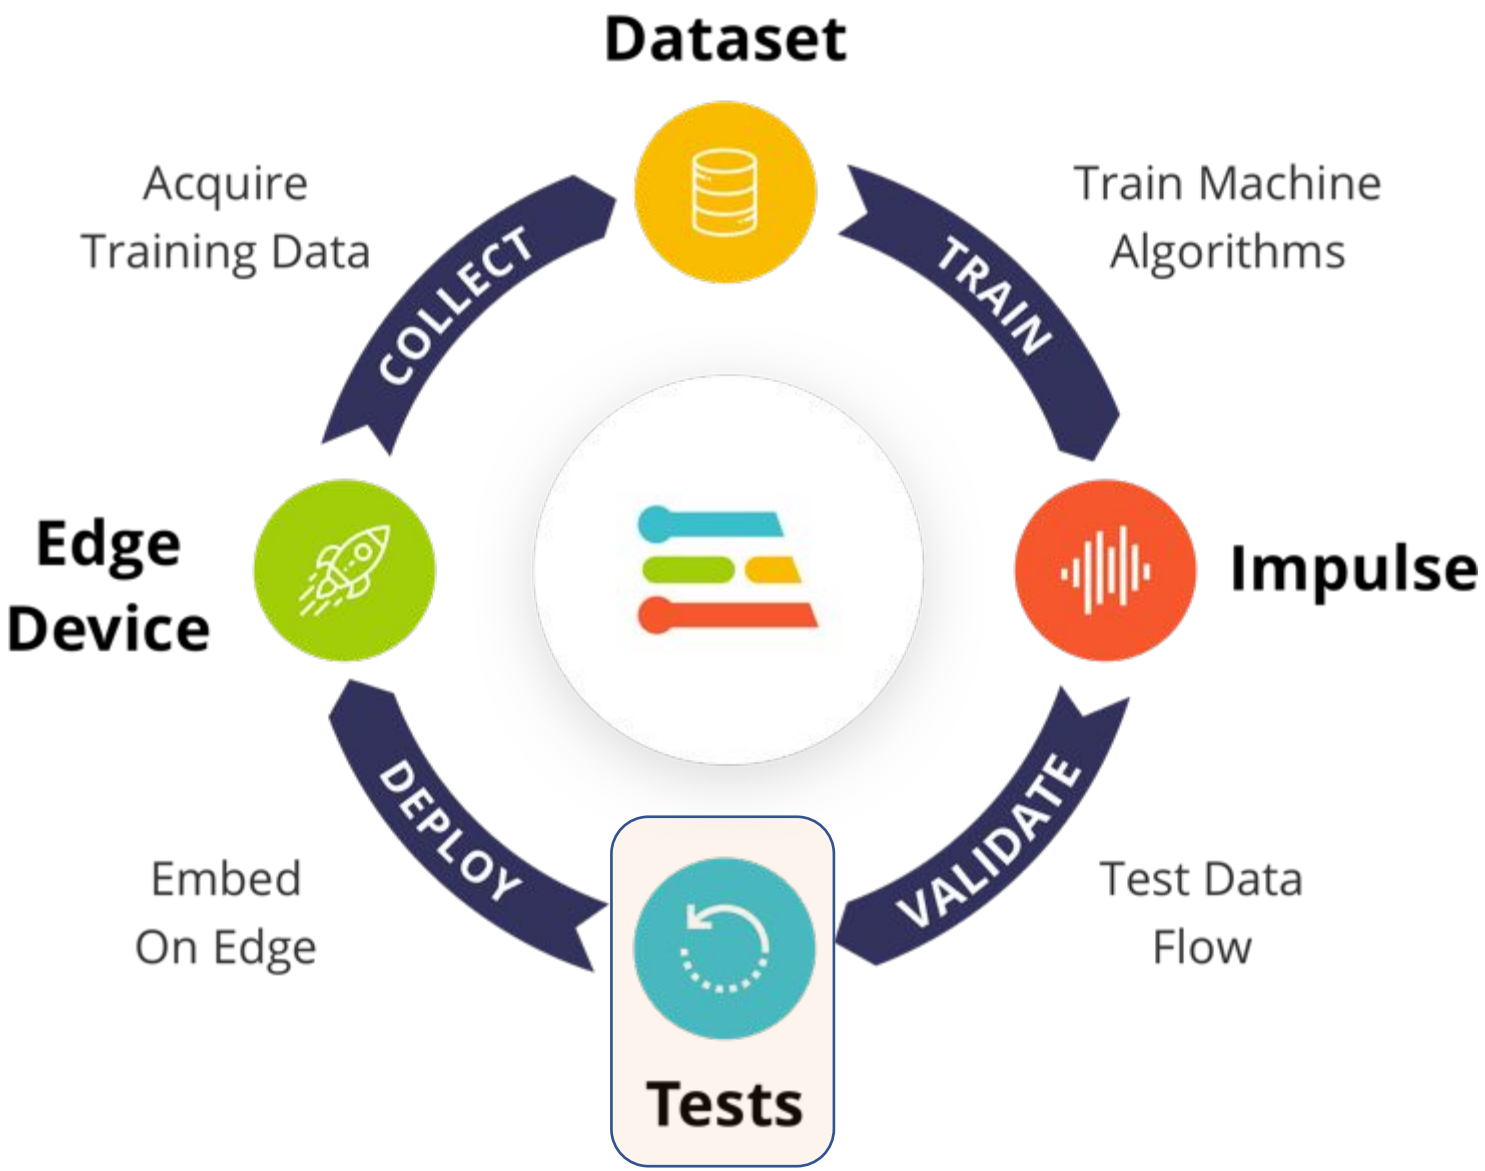

- Pre-Processing Data
- Design a Model
- Train a Model

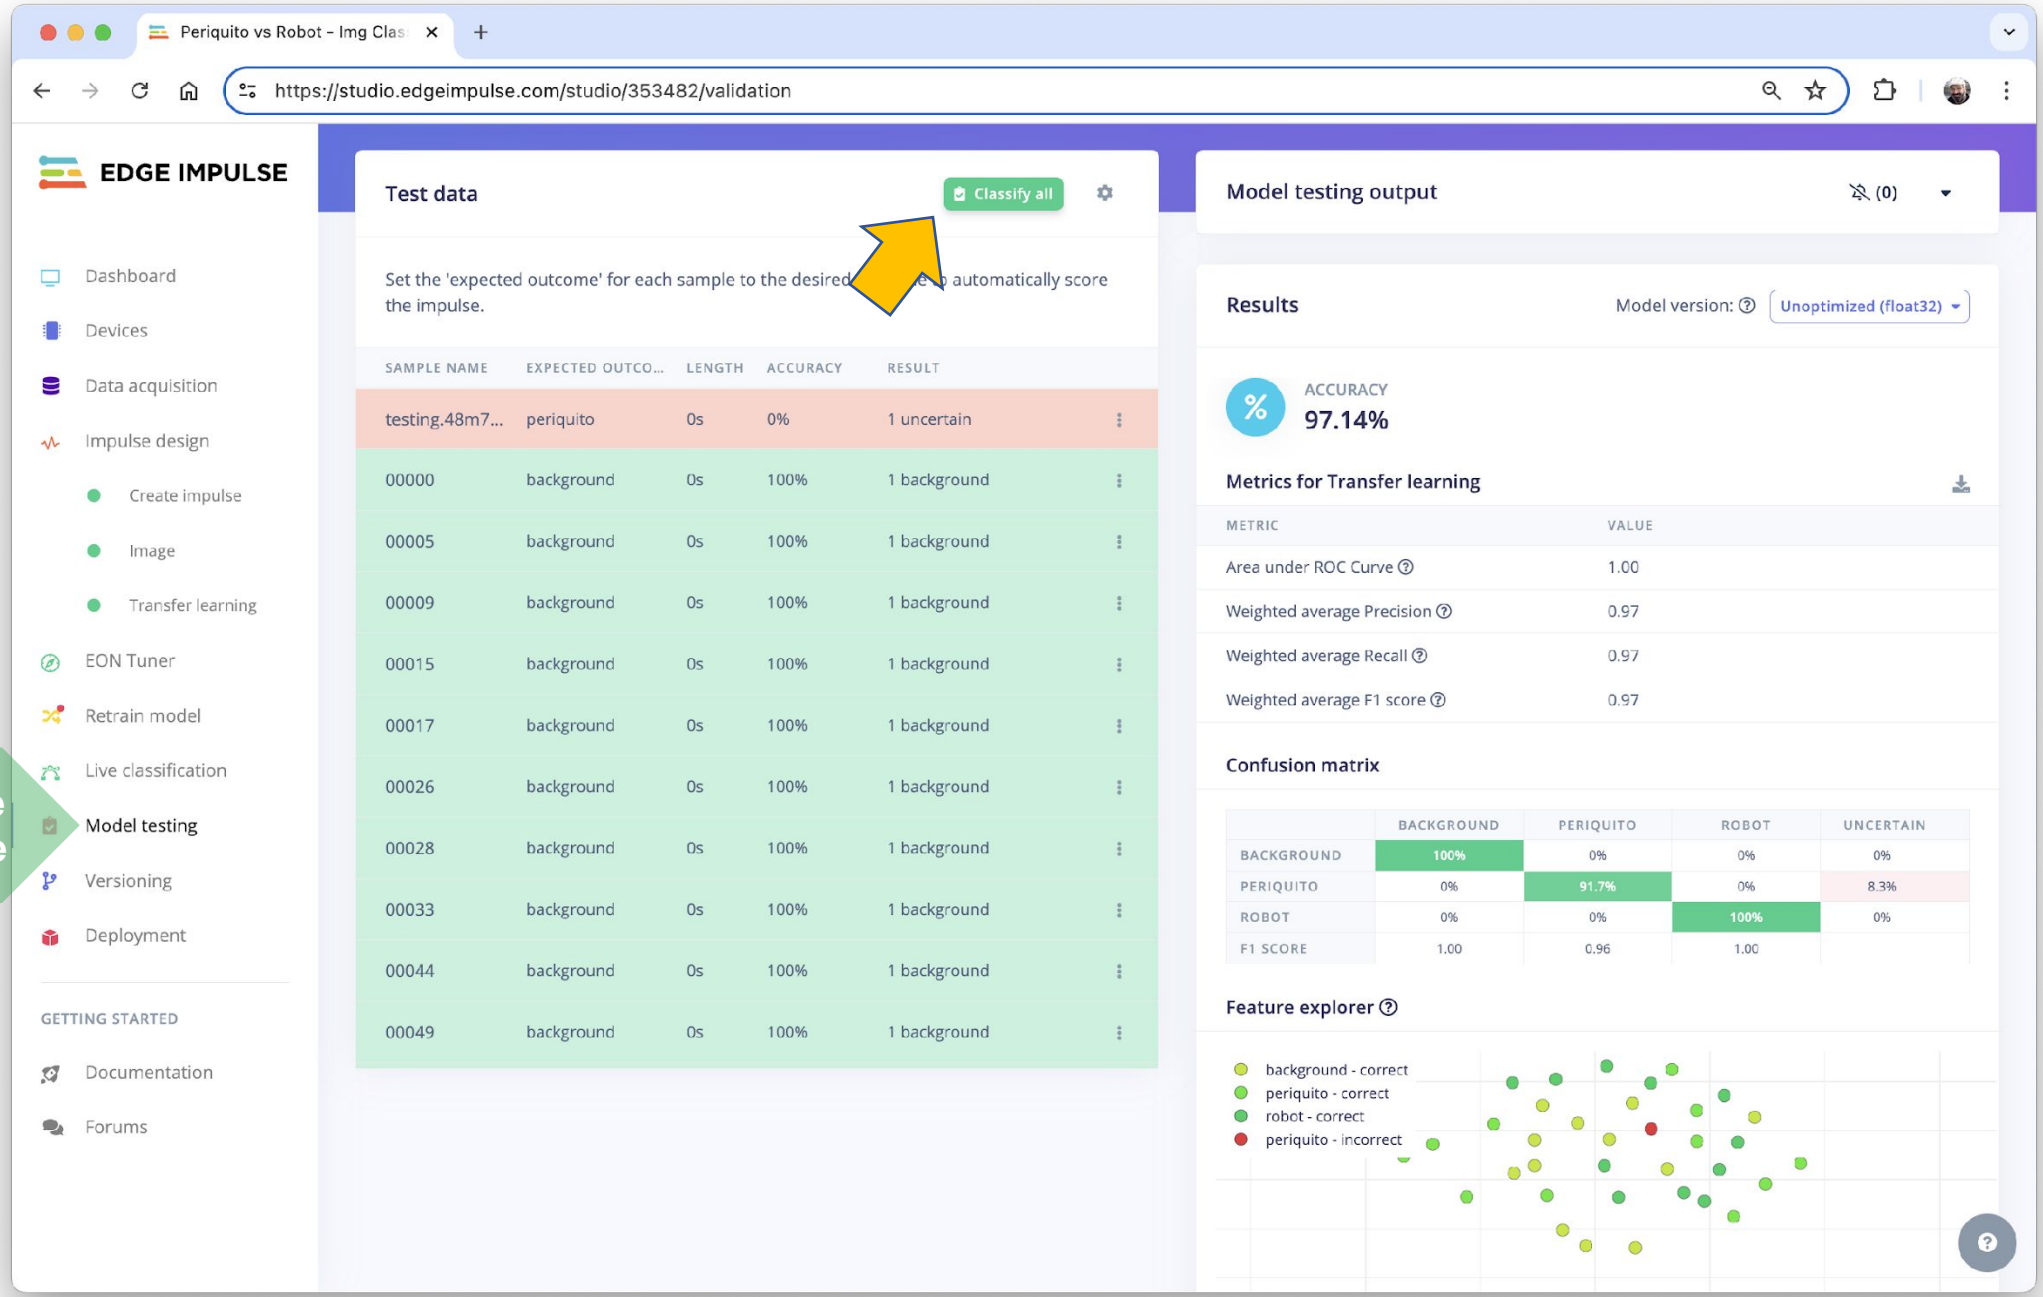

Evaluate
Optimize

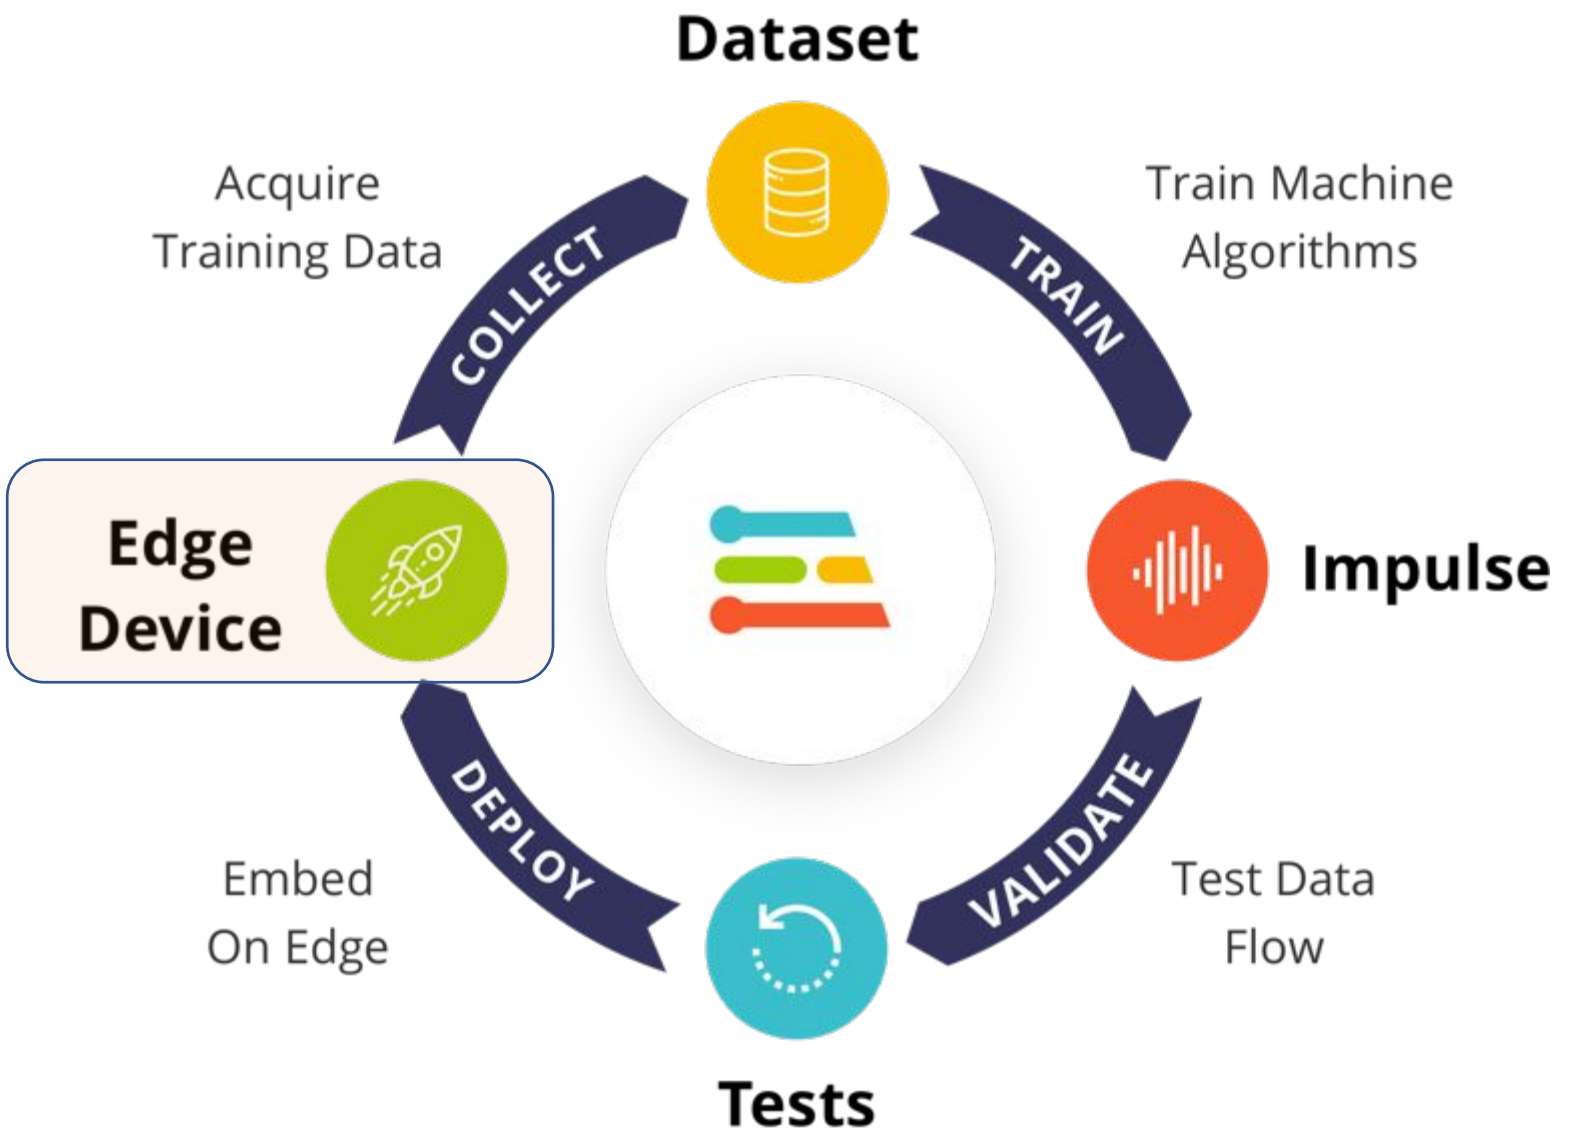

- Pre-Processing Data
- Design a Model
- Train a Model

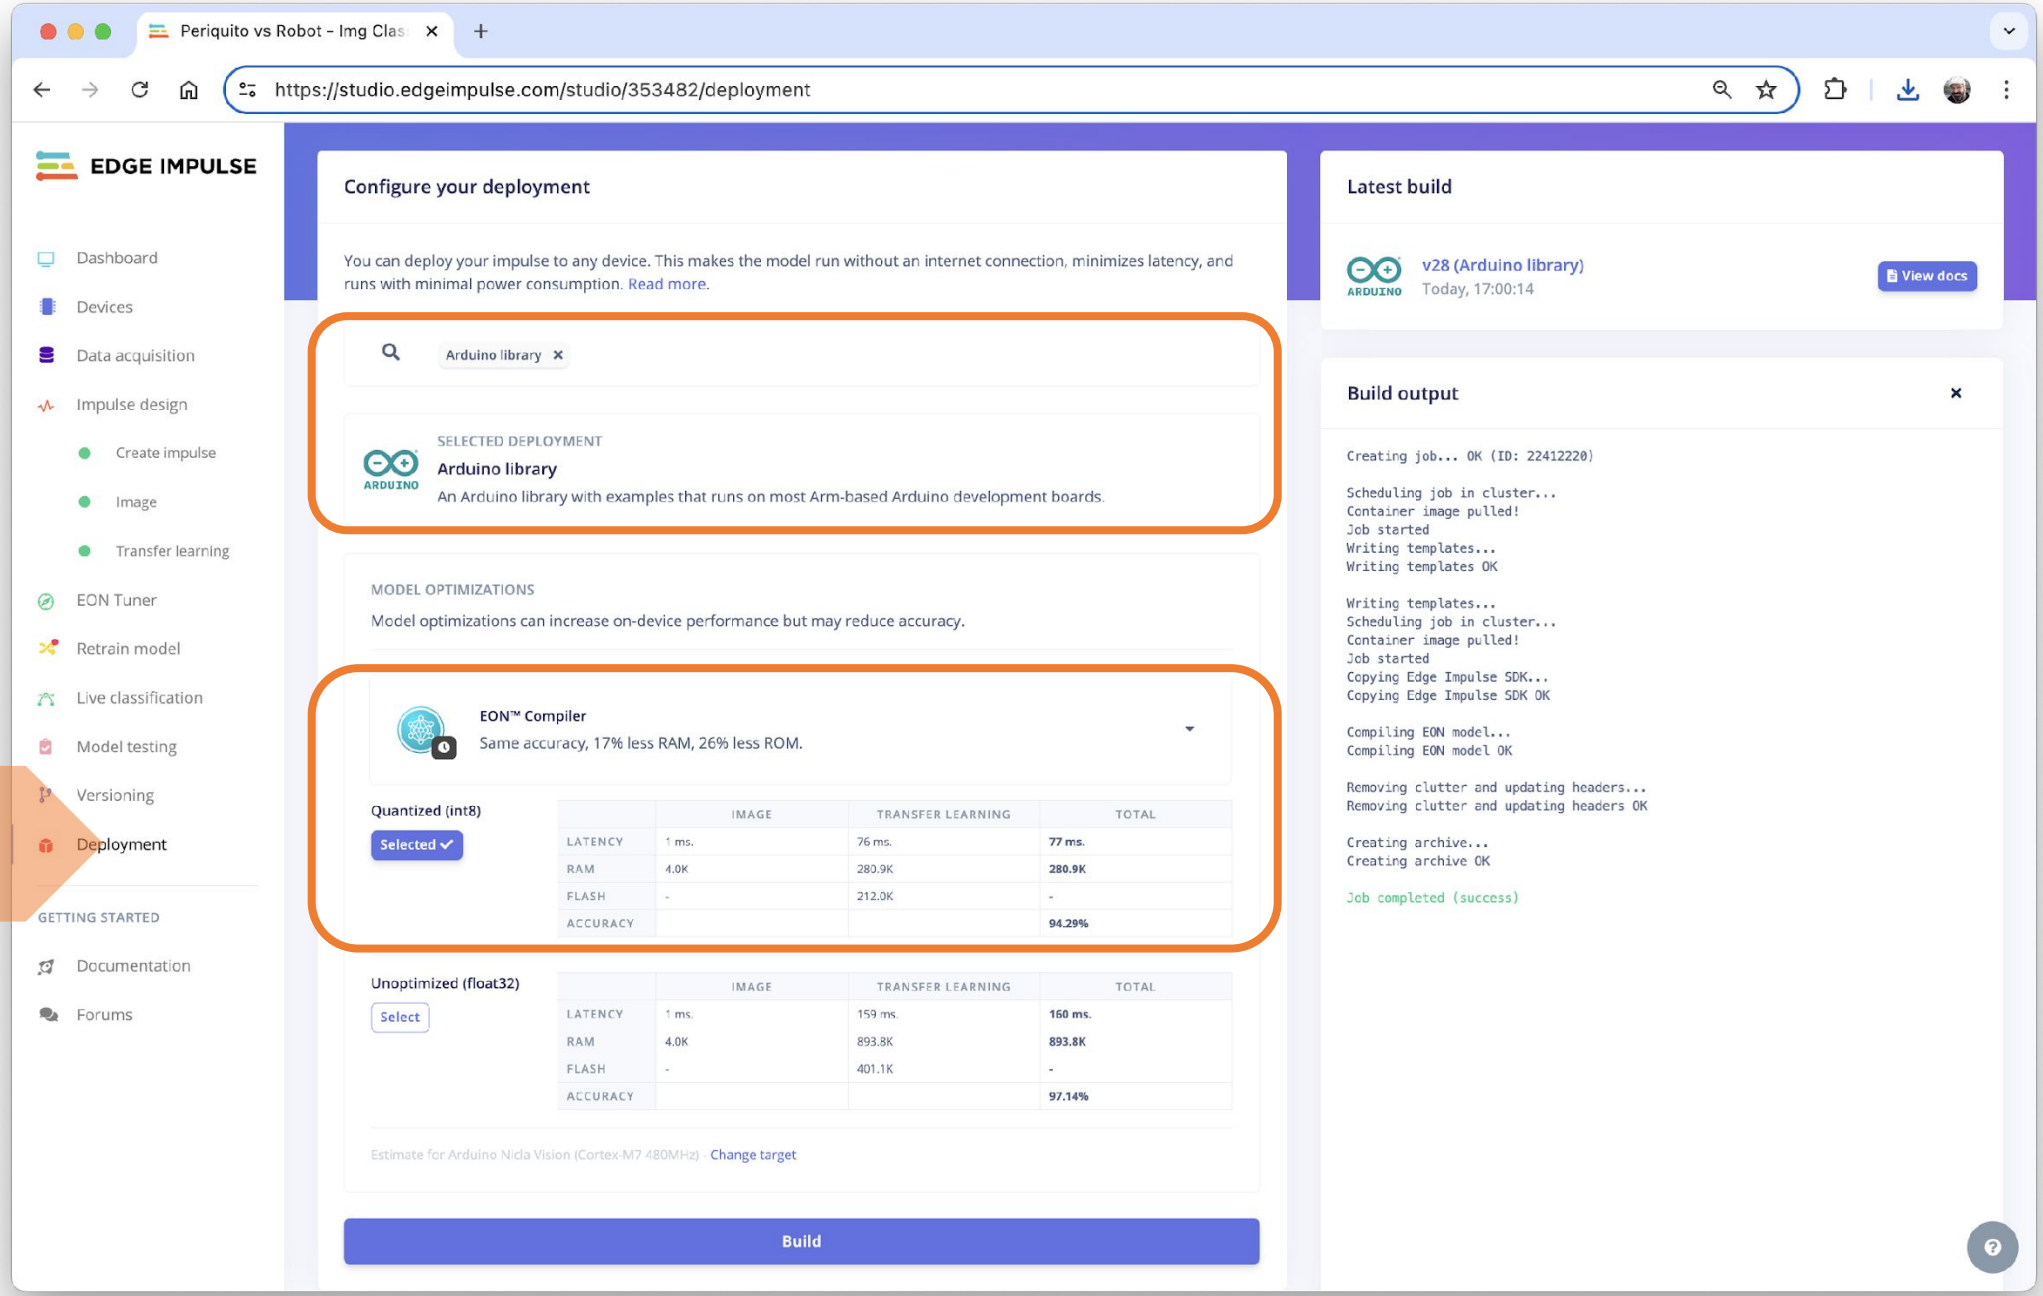

Deploy Model

Make Inferences

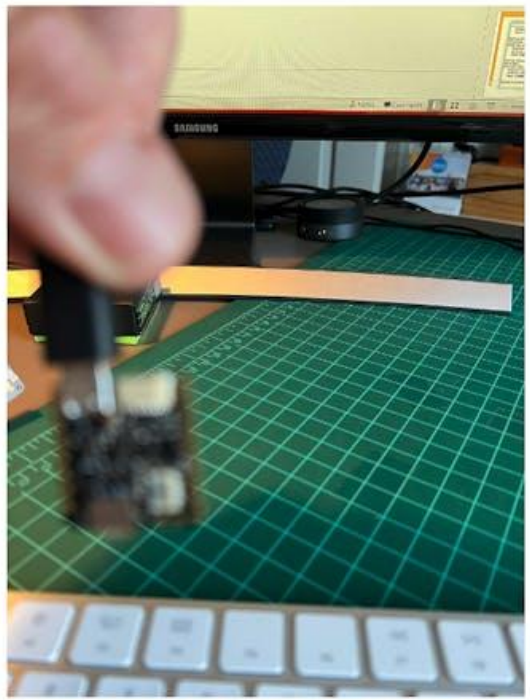

```
Predictions (DSP: 1 ms., Classification: 86 ms.)
    background: 0.992187
    periquito: 0.007813
    robot: 0.000000
```

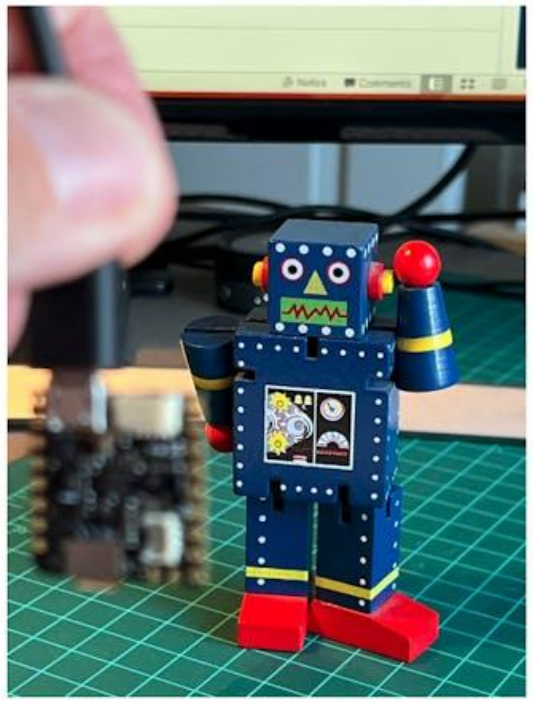

```
Predictions (DSP: 1 ms., Classification: 86 ms.)
    background: 0.015625
    periquito: 0.210937
    robot: 0.773437
```

```
Predictions (DSP: 1 ms., Classification: 86 ms.)
    background: 0.000000
    periquito: 0.992187
    robot: 0.007813
```

https://youtu.be/bZPZZJblU-o?si=grsv71cER-clvsSd

Post-Processing with LEDs:

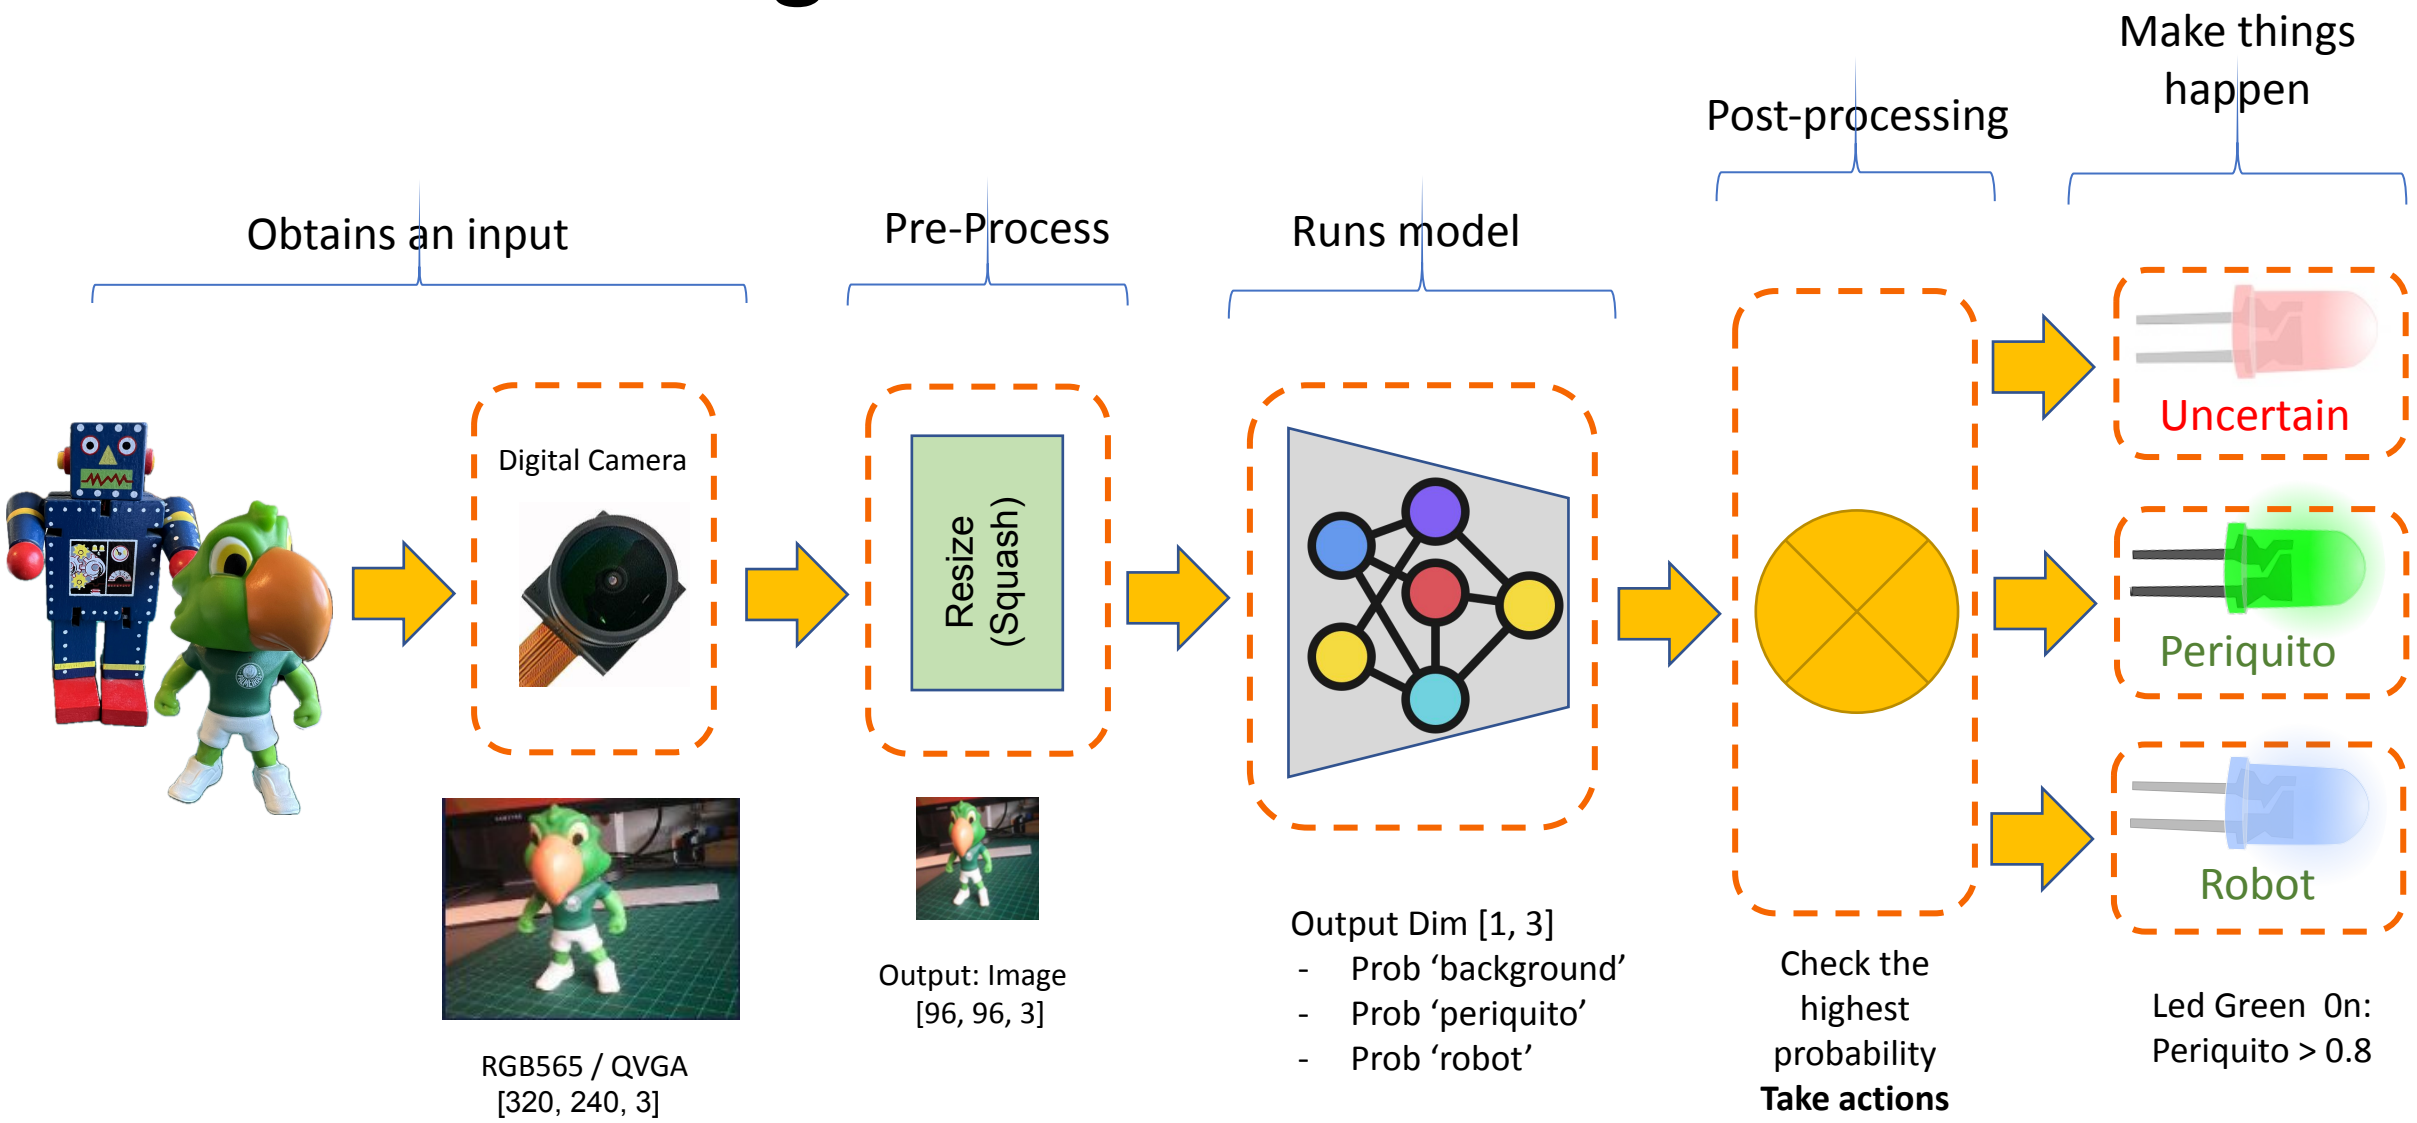


```
Serial Terminal
**********
Prediction:
periquito with a prob of 0.99
FPS: 7.37 fps ==> latency: 136 ms
```


```
Serial Terminal
**********
Prediction:
robot with a prob of 0.96
FPS: 7.34 fps ==> latency: 136 ms
```

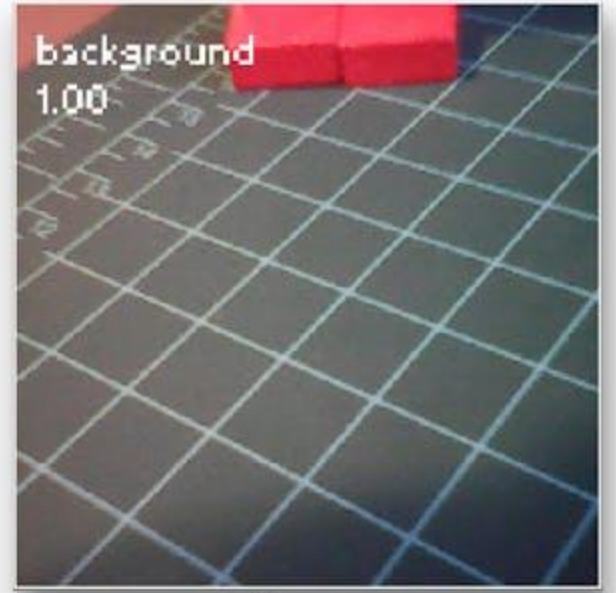


```
Serial Terminal
**********
Prediction:
background with a prob of 1.00
FPS: 7.34 fps ==> latency: 136 ms
```


```
Serial Terminal
**********
Prediction:
uncertain with a prob of 0.76
FPS: 7.42 fps ==> latency: 135 ms
```

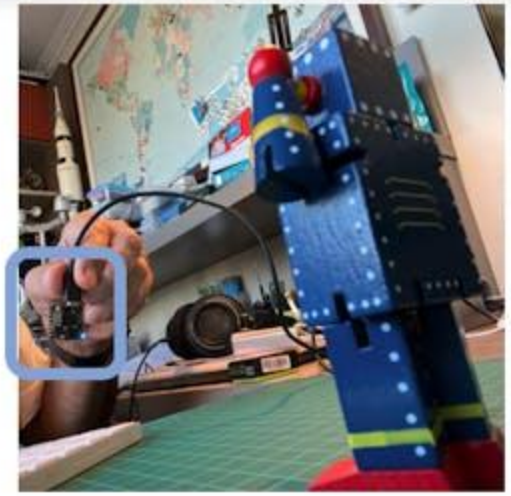

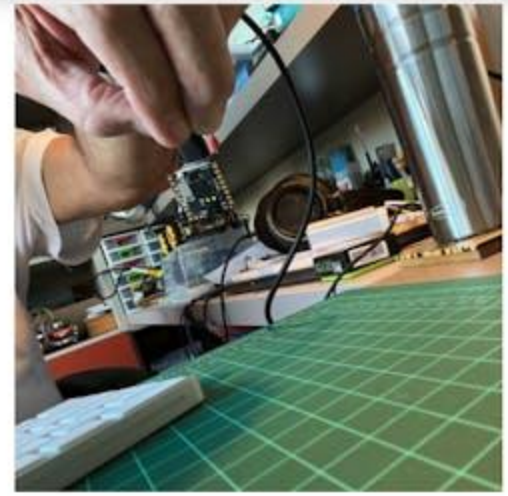

Computer Vision Recognition Tasks

Image Classification

(**Multi-Class** Classification)

Cat: 70%

Dog: 80%

Object Detection

Multi-Label Classification + Object **Localization**

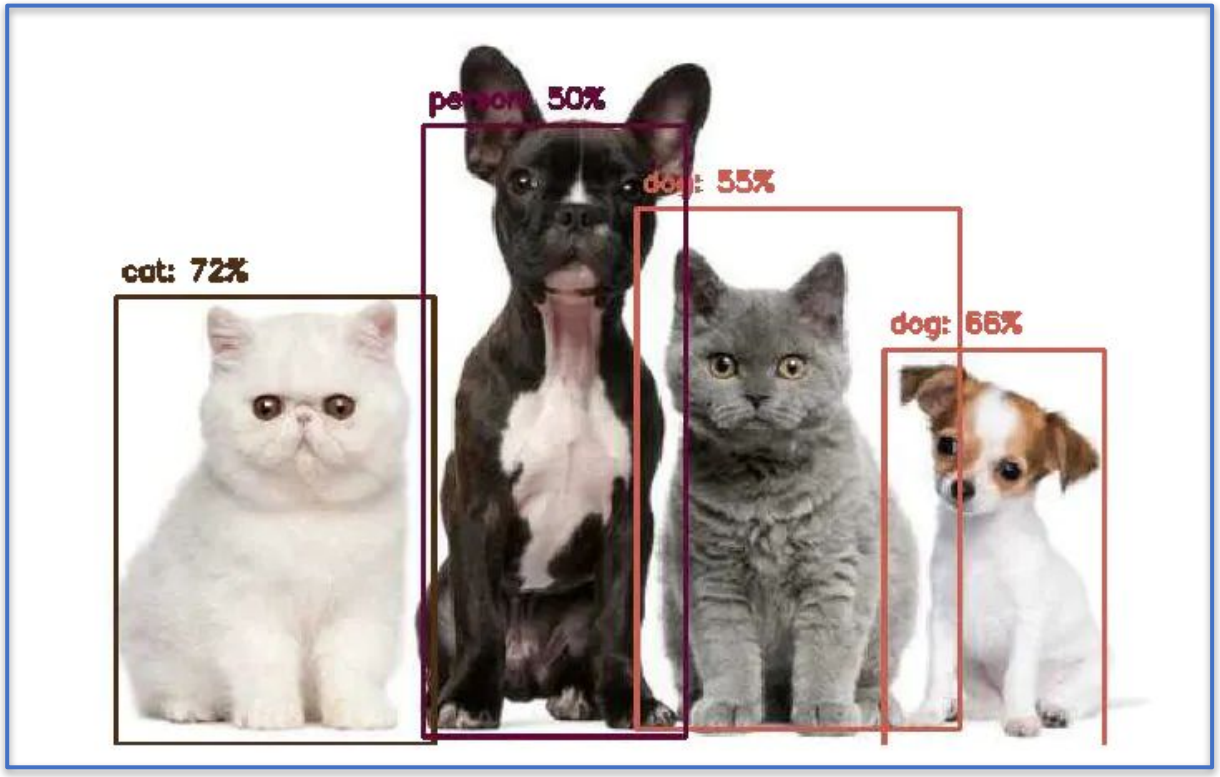

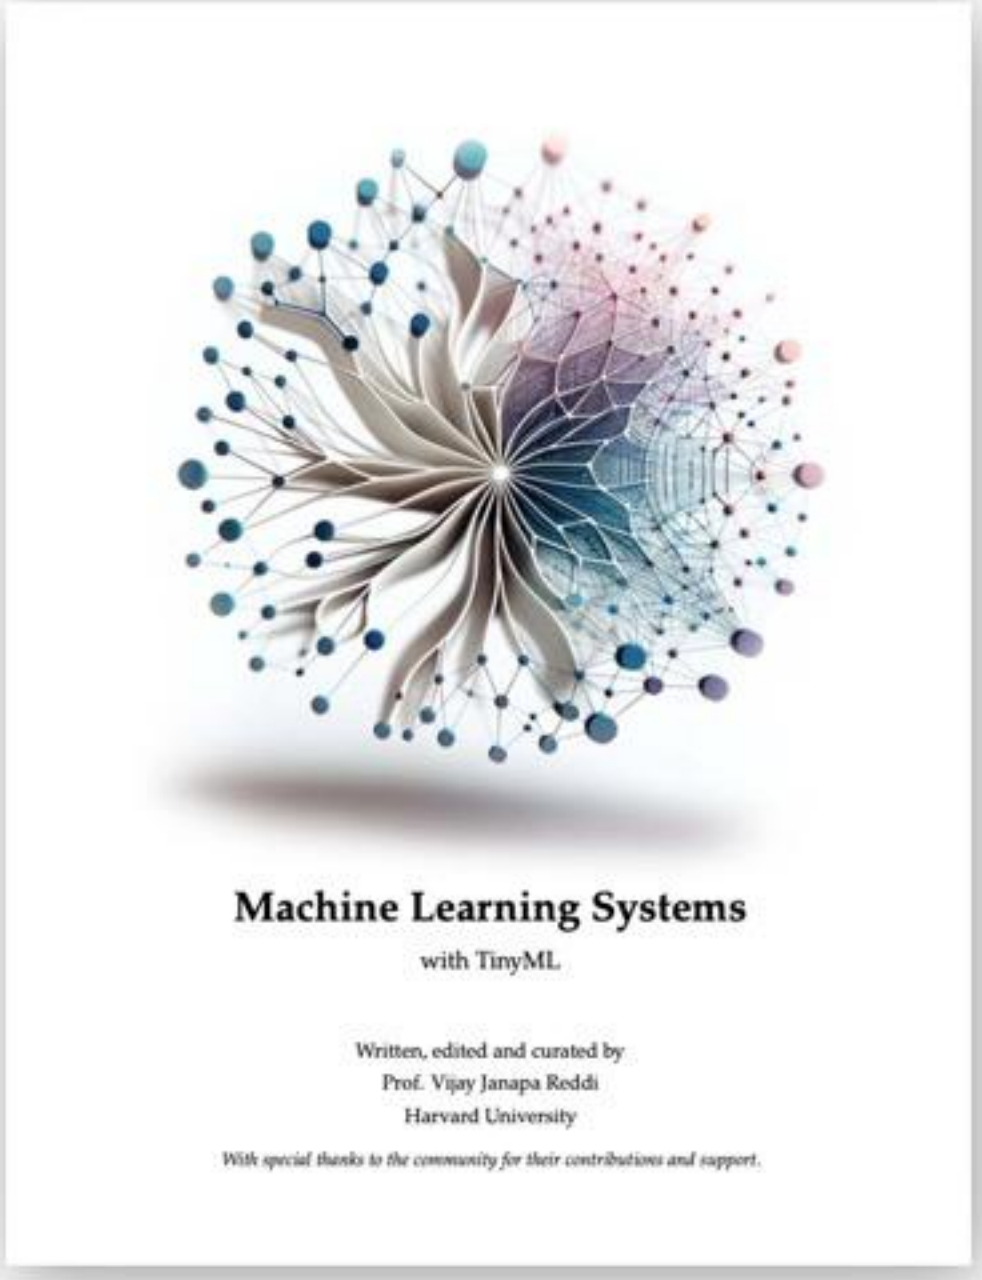

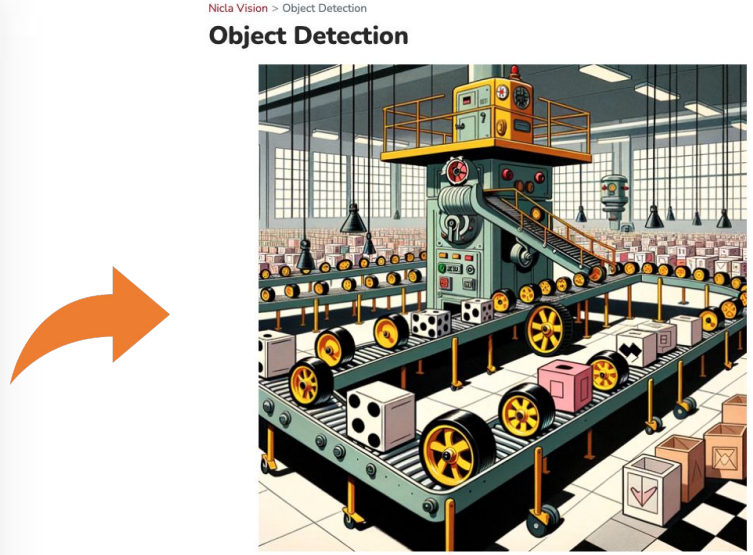

Nicla Vision

XIAO ESP32S3

FOMO

Object Detection model

Labeling

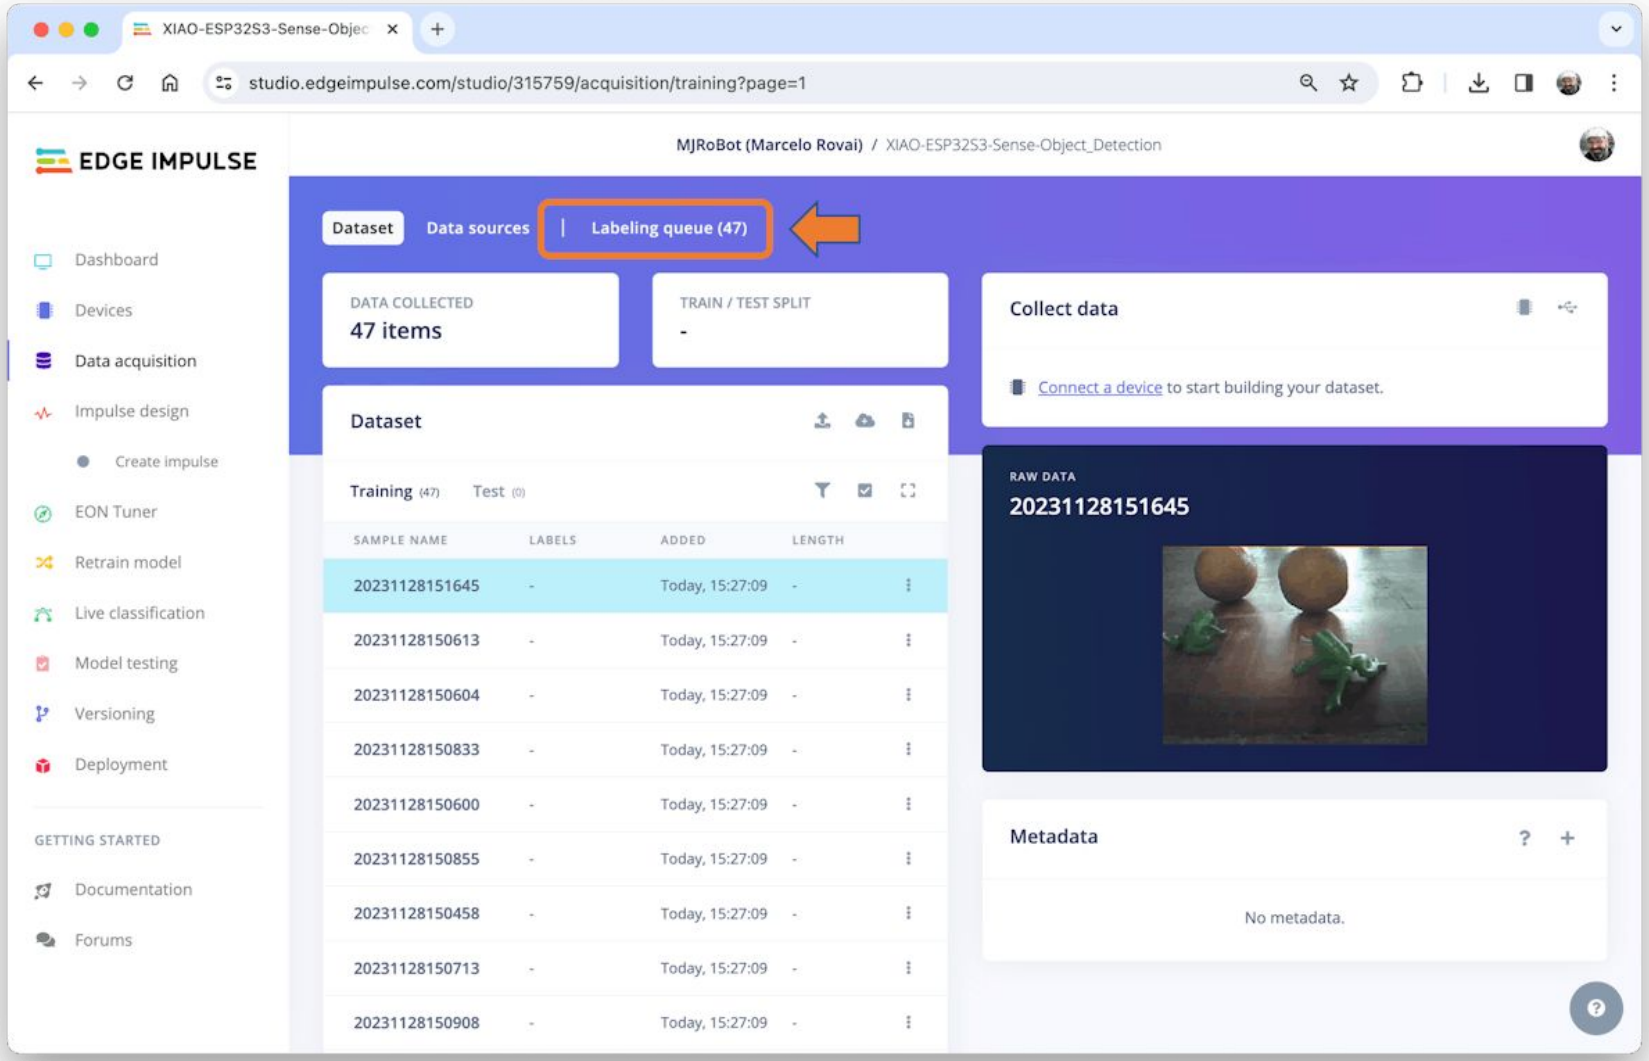
XIAO-ESP32S3-Sense-Objec
studio.edgeimpulse.com/studio/315759/acquisition/training?page=1
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / XIAO-ESP32S3-Sense-Object_Detection
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Dataset
Data sources
Labeling queue (47)
DATA COLLECTED
47 items
TRAIN / TEST SPLIT
-
Collect data
Connect a device to start building your dataset.
Dataset
Training (47)
Test (0)
SAMPLE NAME
LABELS
ADDED
LENGTH
20231128151645 - Today, 15:27:09 -
20231128150613 - Today, 15:27:09 -
20231128150604 - Today, 15:27:09 -
20231128150833 - Today, 15:27:09 -
20231128150600 - Today, 15:27:09 -
20231128150855 - Today, 15:27:09 -
20231128150458 - Today, 15:27:09 -
20231128150713 - Today, 15:27:09 -
20231128150908 - Today, 15:27:09 -
RAW DATA
20231128151645
Metadata
No metadata.

Labeling

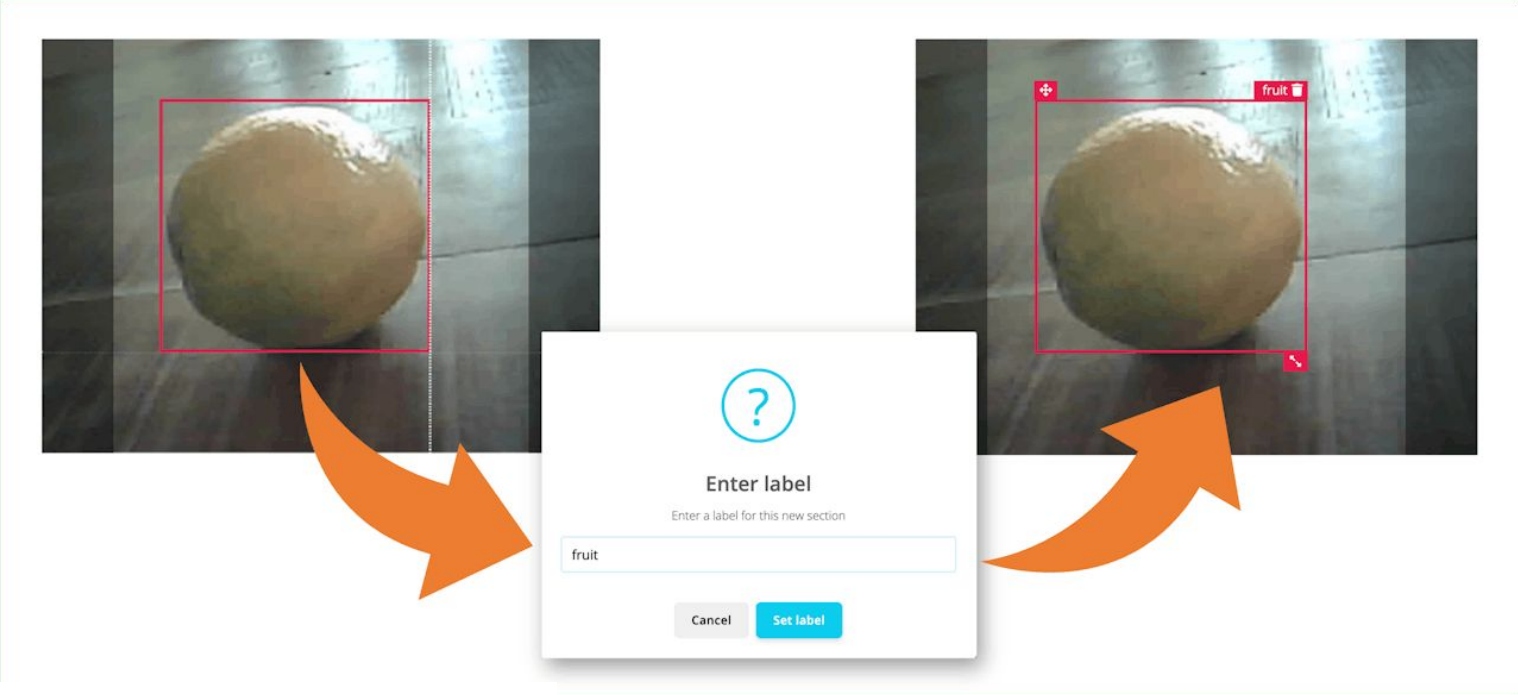
fruit
Enter label
Enter a label for this new section
fruit
Cancel
Set label

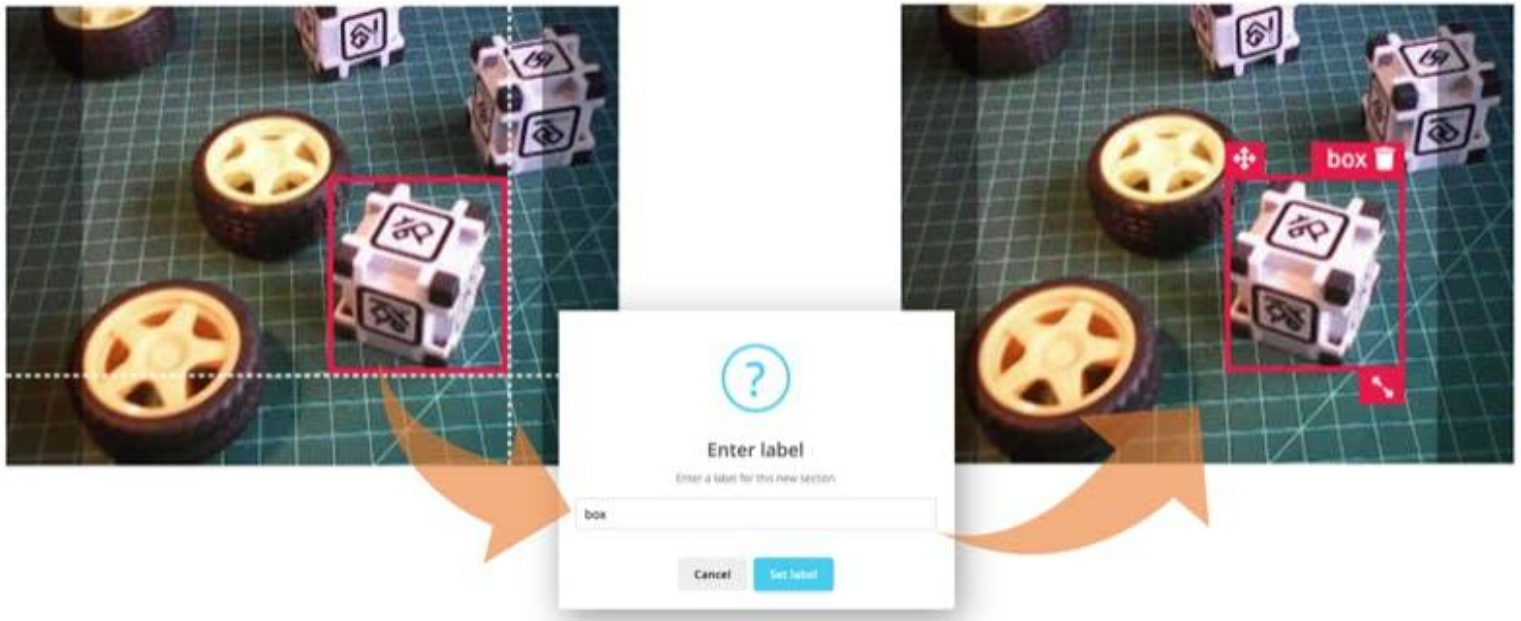
box
Enter label
box
Cancel
Set label

Labeling

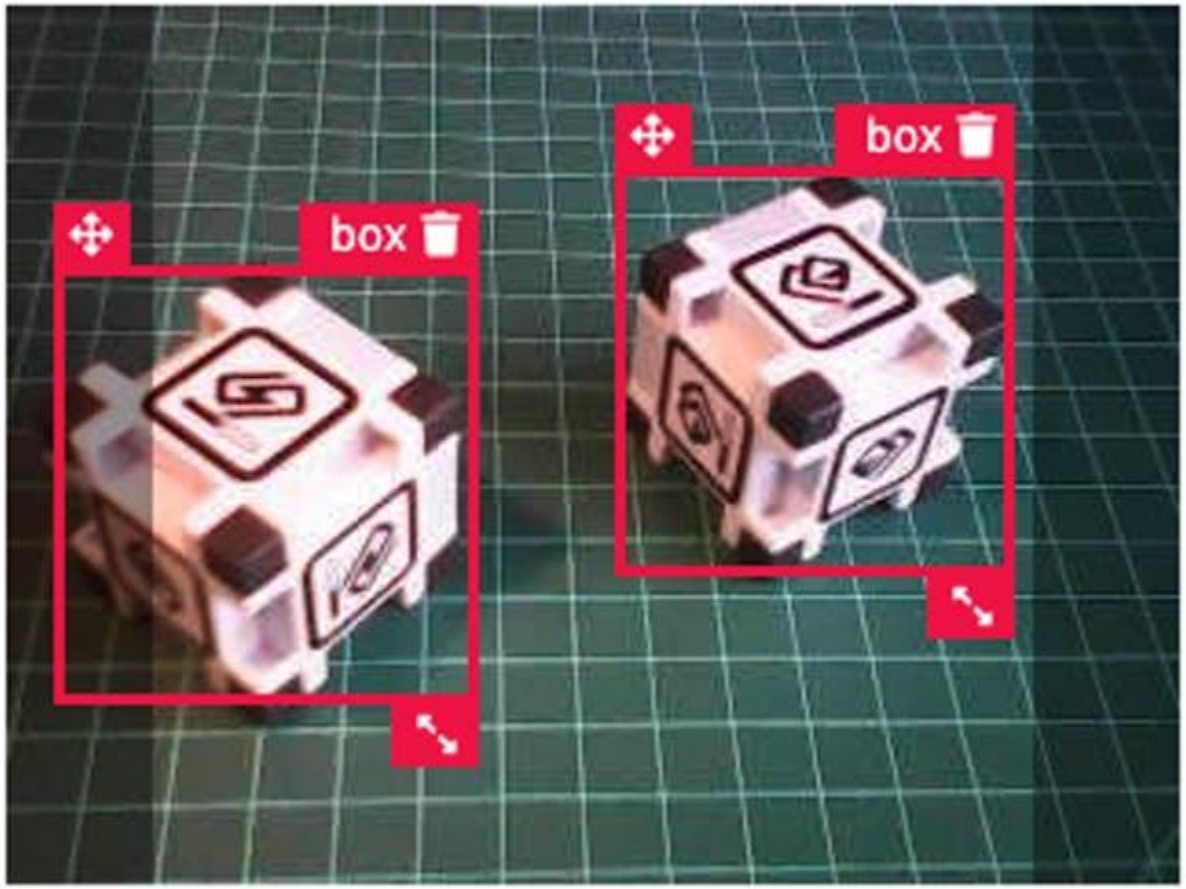

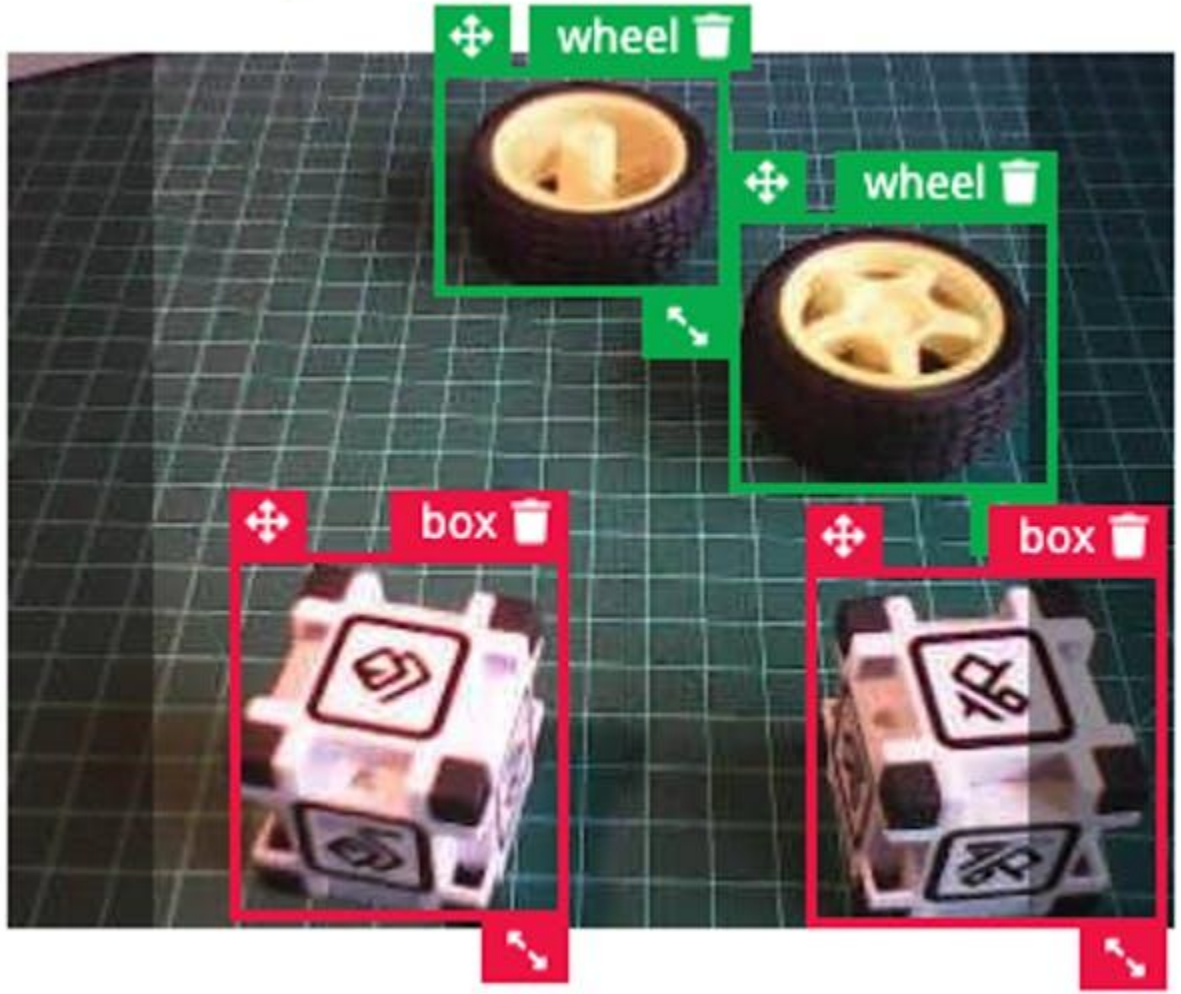

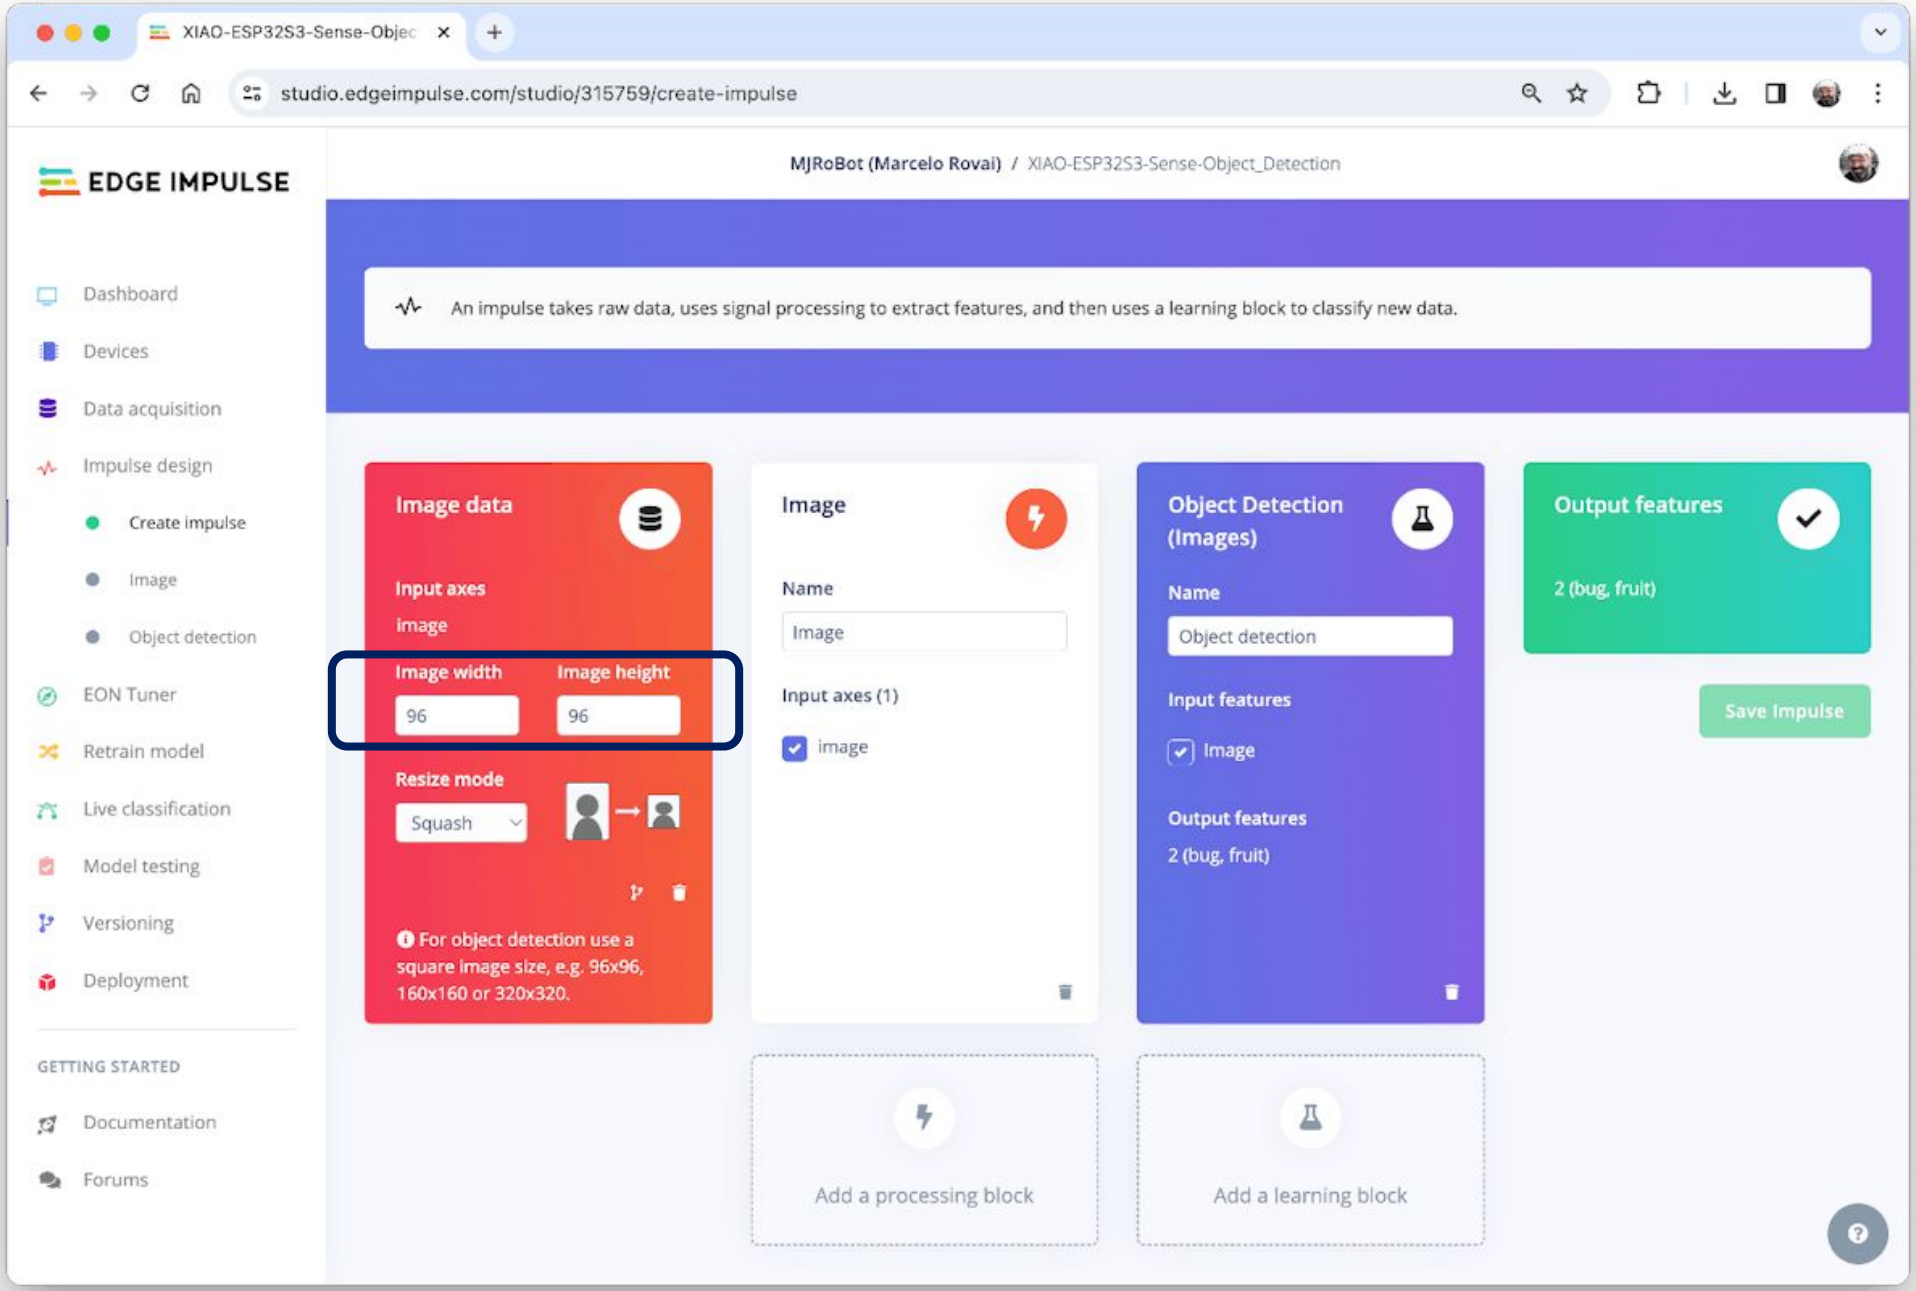
XIAO-ESP32S3-Sense-Objec
studio.edgeimpulse.com/studio/315759/create-impulse
EDGE IMPULSE
MJRoBot (Marcelo Rovai) / XIAO-ESP32S3-Sense-Object_Detection
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
Object detection
EON Tuner
Retrain model
Live classification
Model testing
Versioning
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
Squash
For object detection use a square image size, e.g. 96x96, 160x160 or 320x320.
Image
Name
Image
Input axes (1)
image
Object Detection (Images)
Name
Object detection
Input features
Image
Output features
2 (bug, fruit)
Output features
2 (bug, fruit)
Save Impulse
Add a processing block
Add a learning block

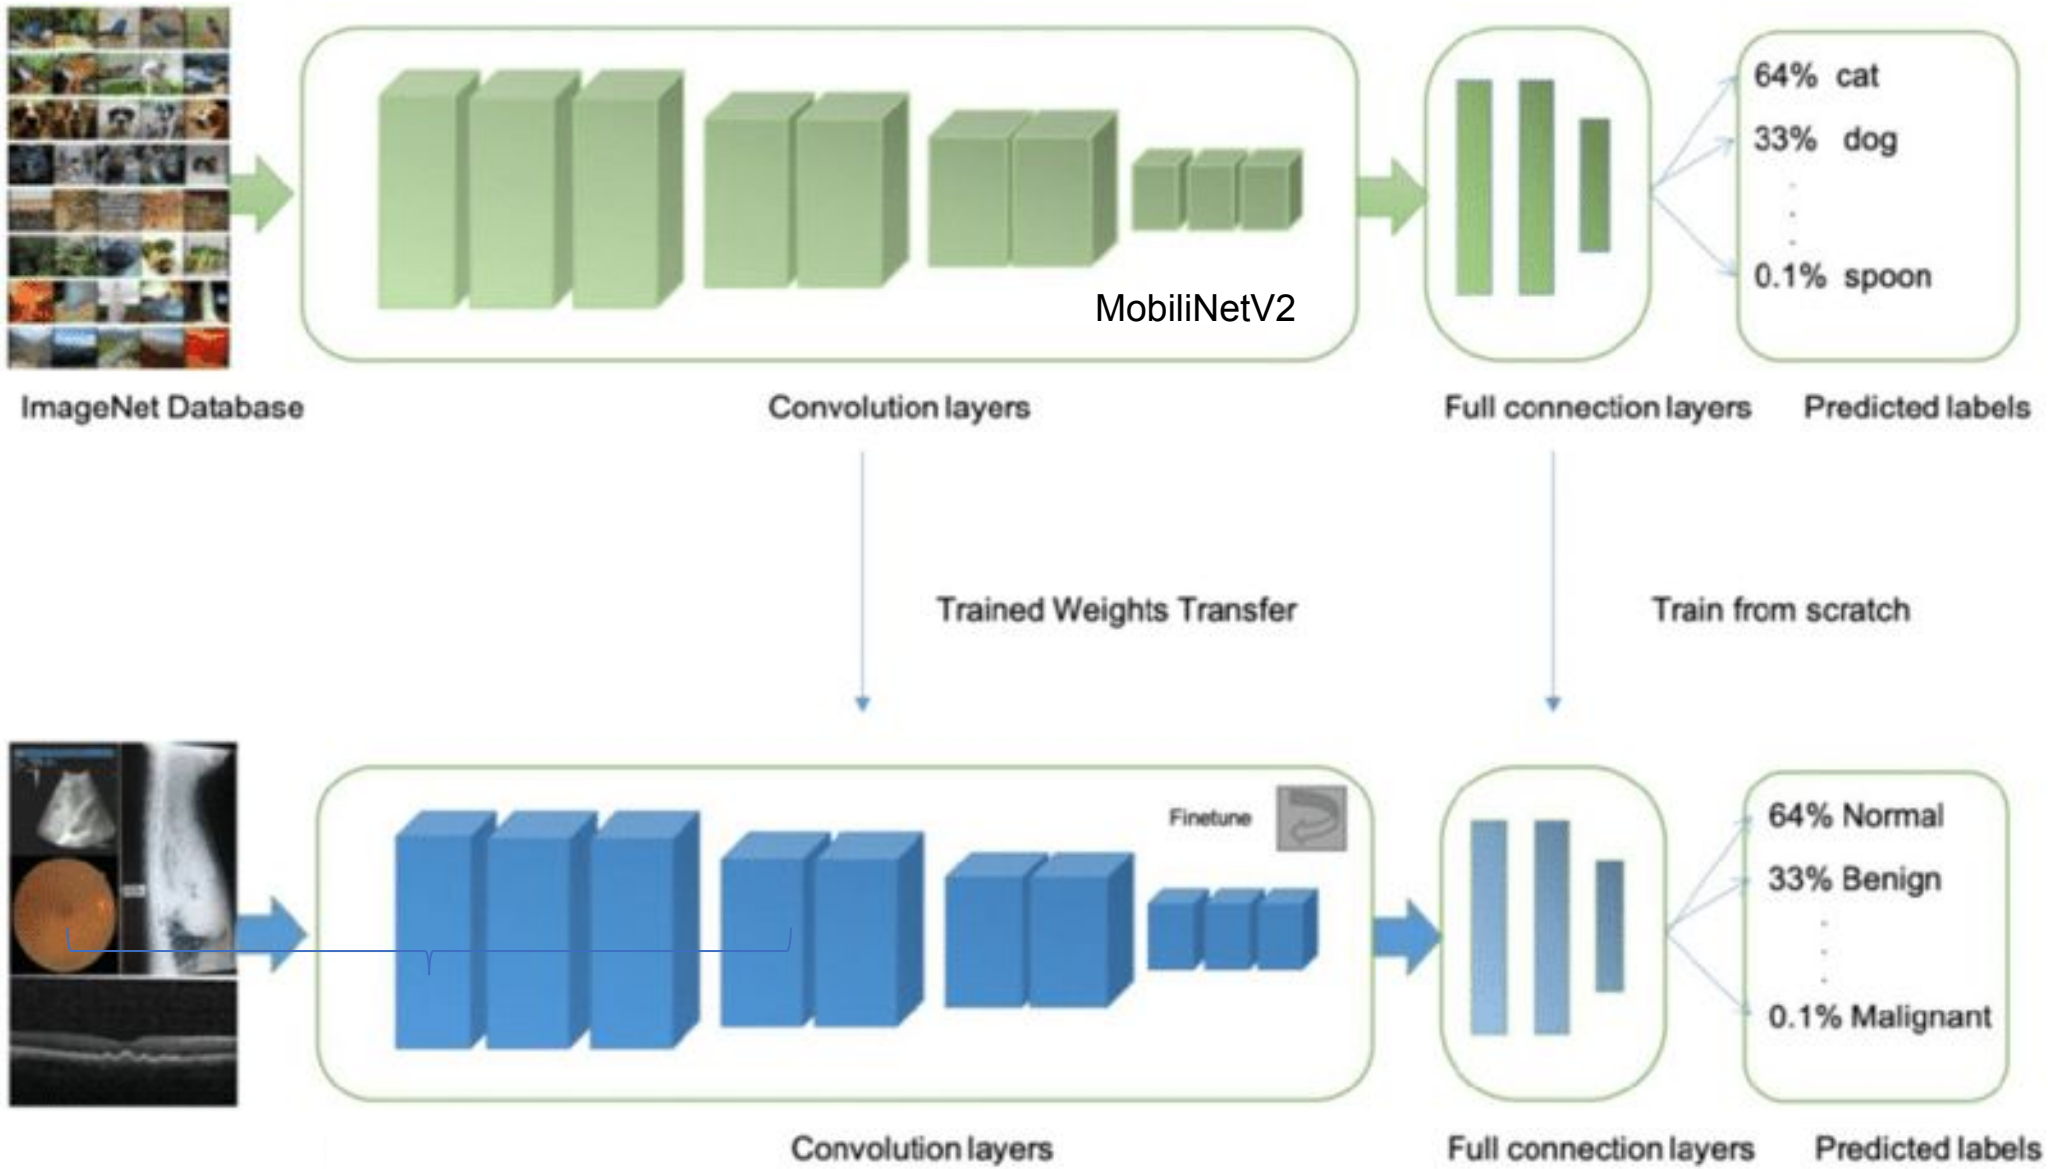
64% cat
33% dog
0.1% spoon
MobiliNetV2
ImageNet Database
Convolution layers
Full connection layers
Predicted labels
Trained Weights Transfer
Train from scratch
Finetune
64% Normal
33% Benign
0.1% Malignant
Convolution layers
Full connection layers
Predicted labels

Model: FOMO

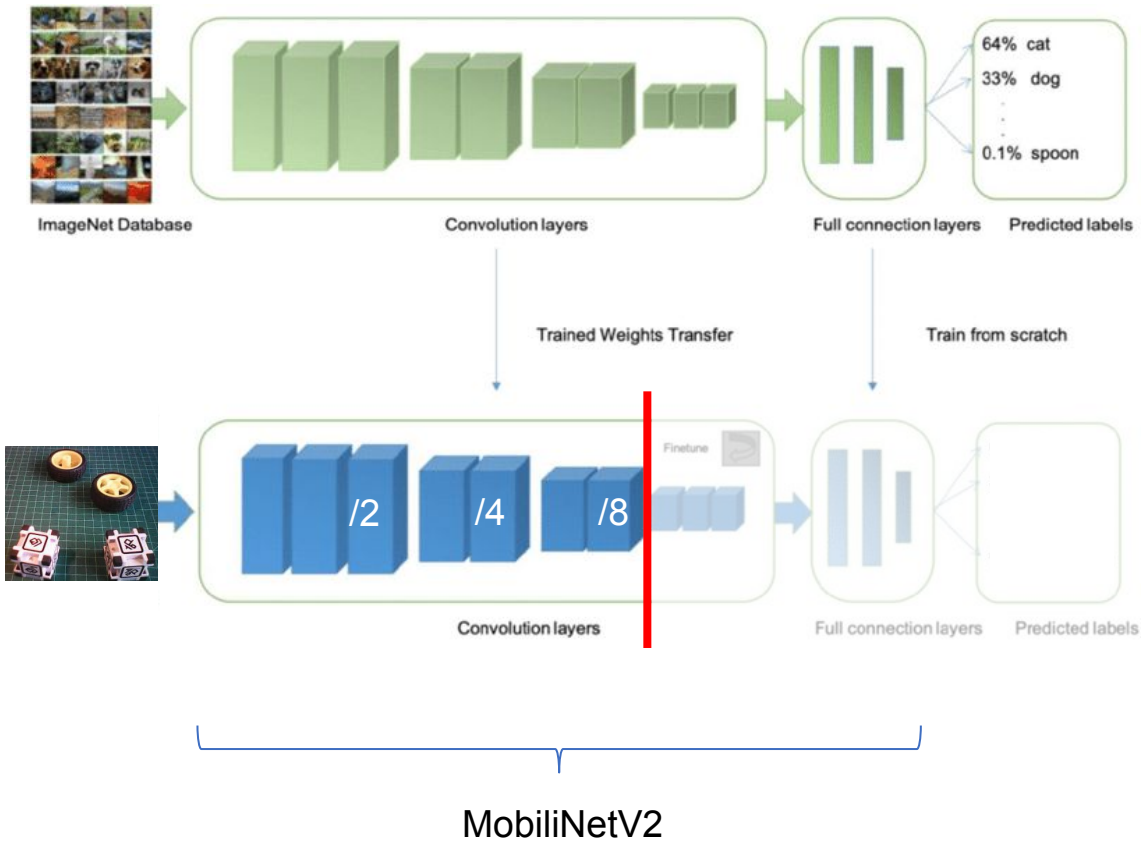
64% cat
33% dog
0.1% spoon
ImageNet Database
Convolution layers
Full connection layers
Predicted labels
Trained Weights Transfer
Train from scratch
/2
/4
/8
Finetune
Convolution layers
Full connection layers
Predicted labels
MobiliNetV2

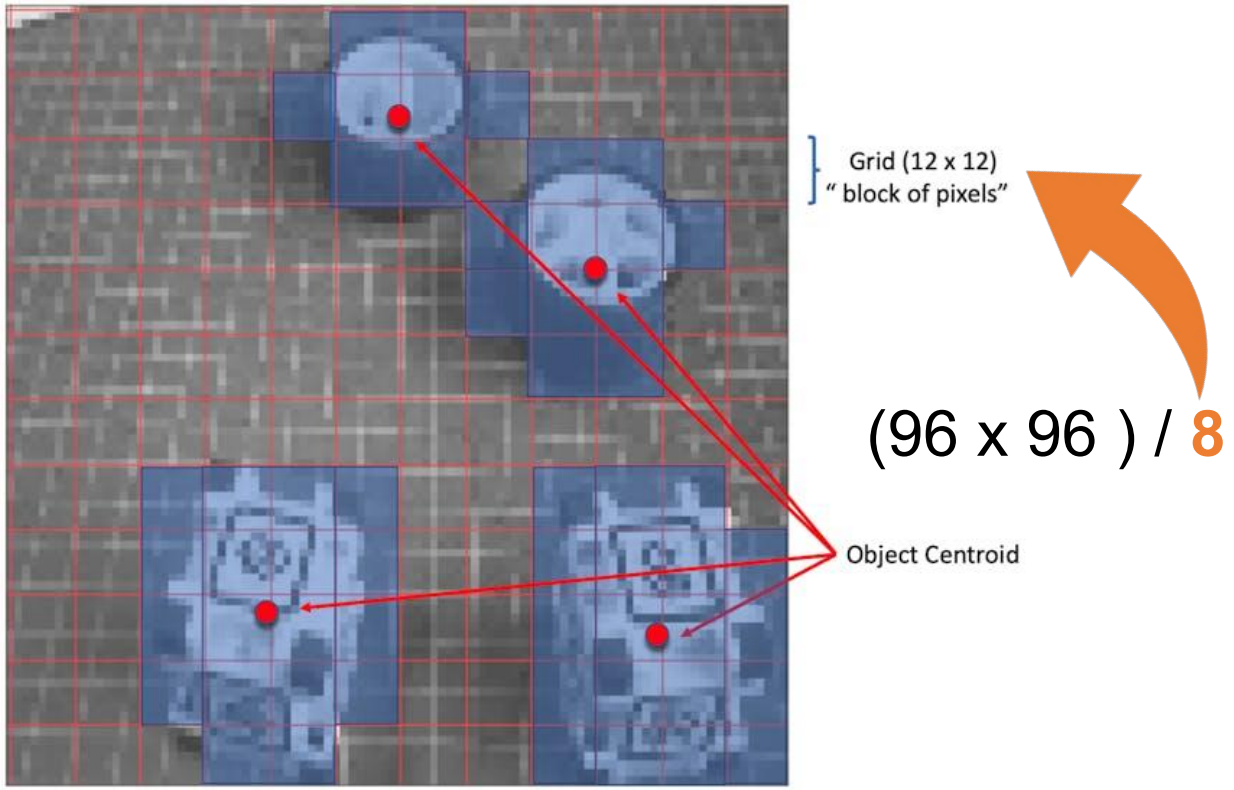
Grid (12 x 12)
" block of pixels"
(96 x 96) / 8
Object Centroid

Inference Test

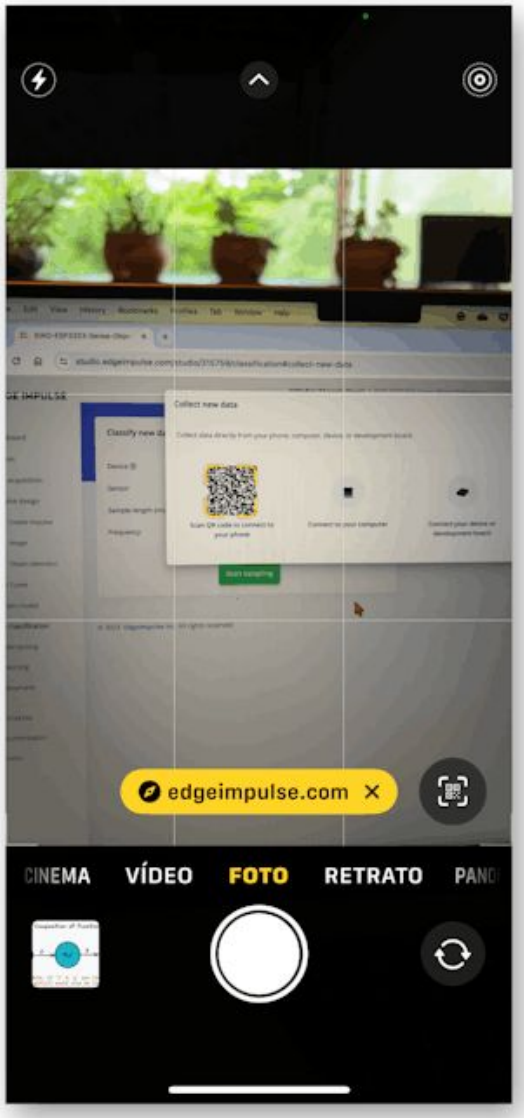

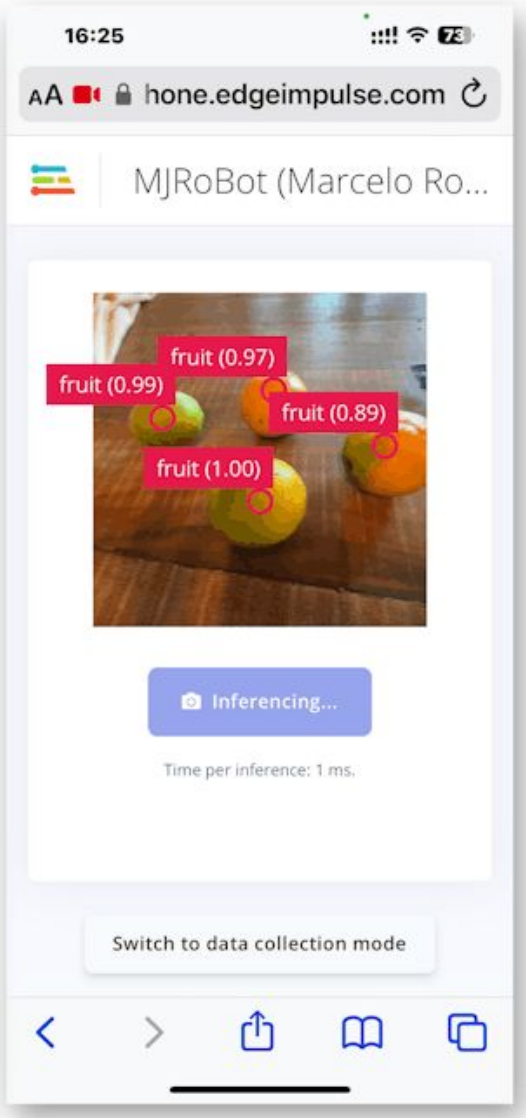

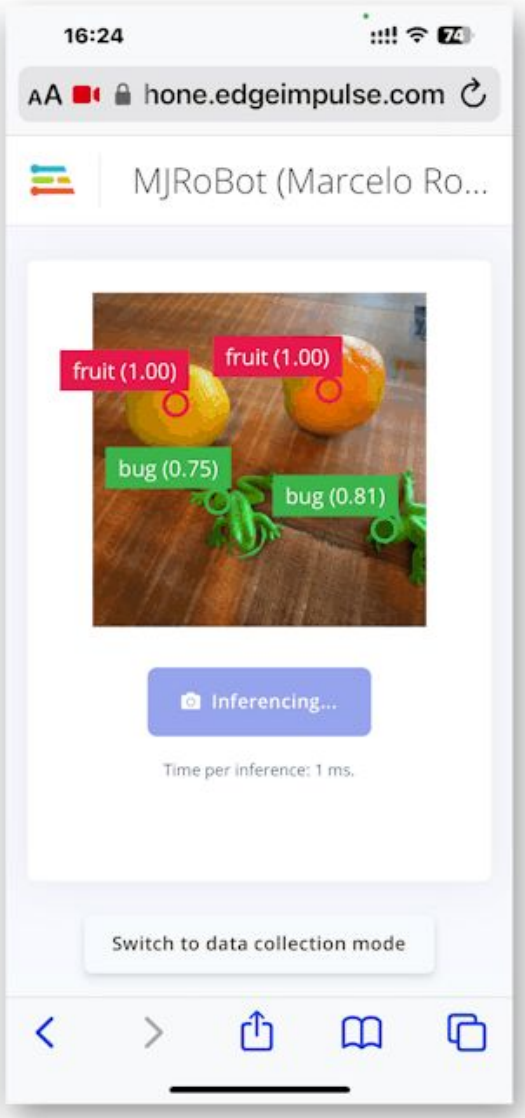

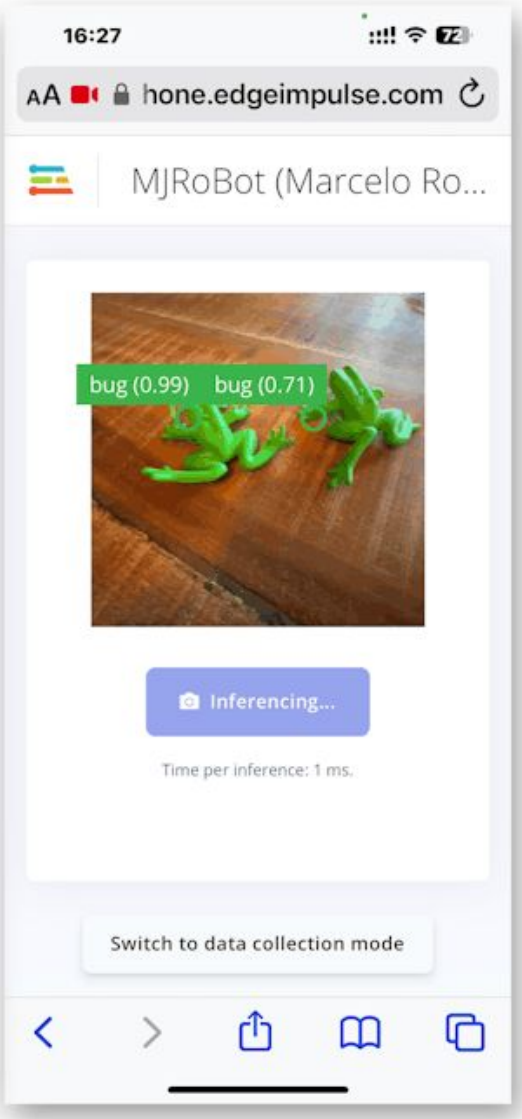

Deploy

OpenMV

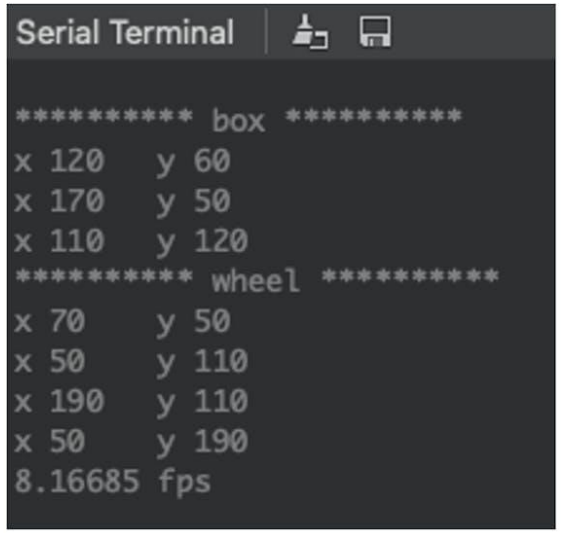

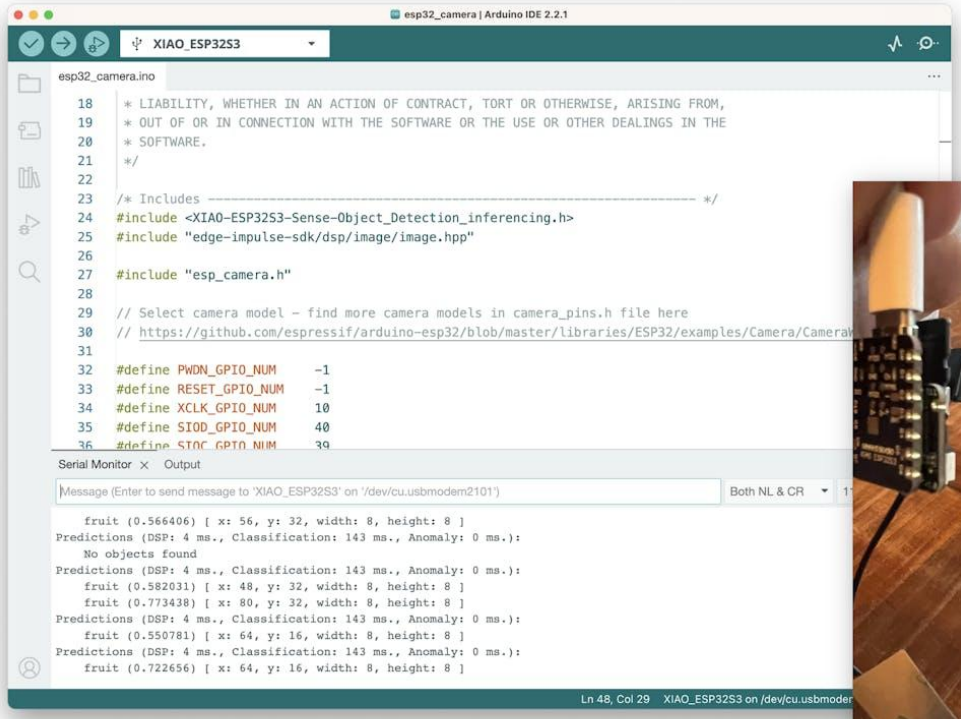

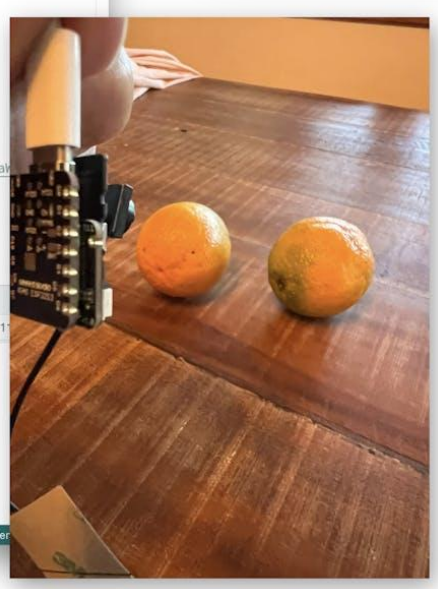

Detecting Objects using TinyML (FOMO)

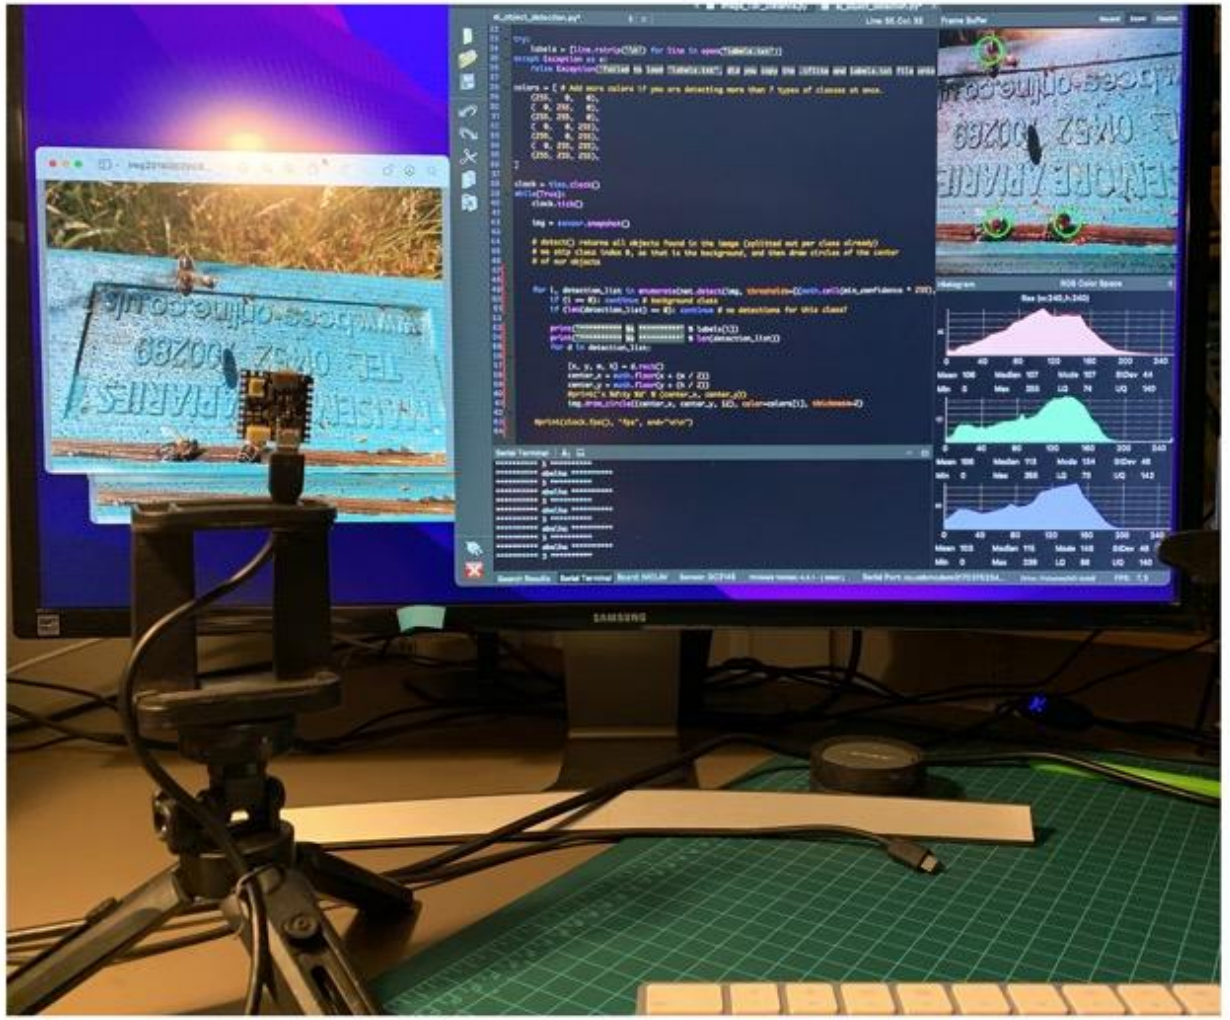

OpenMV

MicroPython

EDGE IMPULSE

microNPU

a neural network unit for TinyML

ML- optimized Solutions

	Ethos-U55	Ethos-U65	Ethos-U85
Performance (At 1 GHz)	64 to 512 GOP/s	512 GOP/s to 1 TOP/s	256 GOPS/s to 4 TOP/s
MACs (8x8)	32, 64, 128, 256	256, 512	128, 256, 512, 1024, 2048

TOPS (Tera Operations per Second)

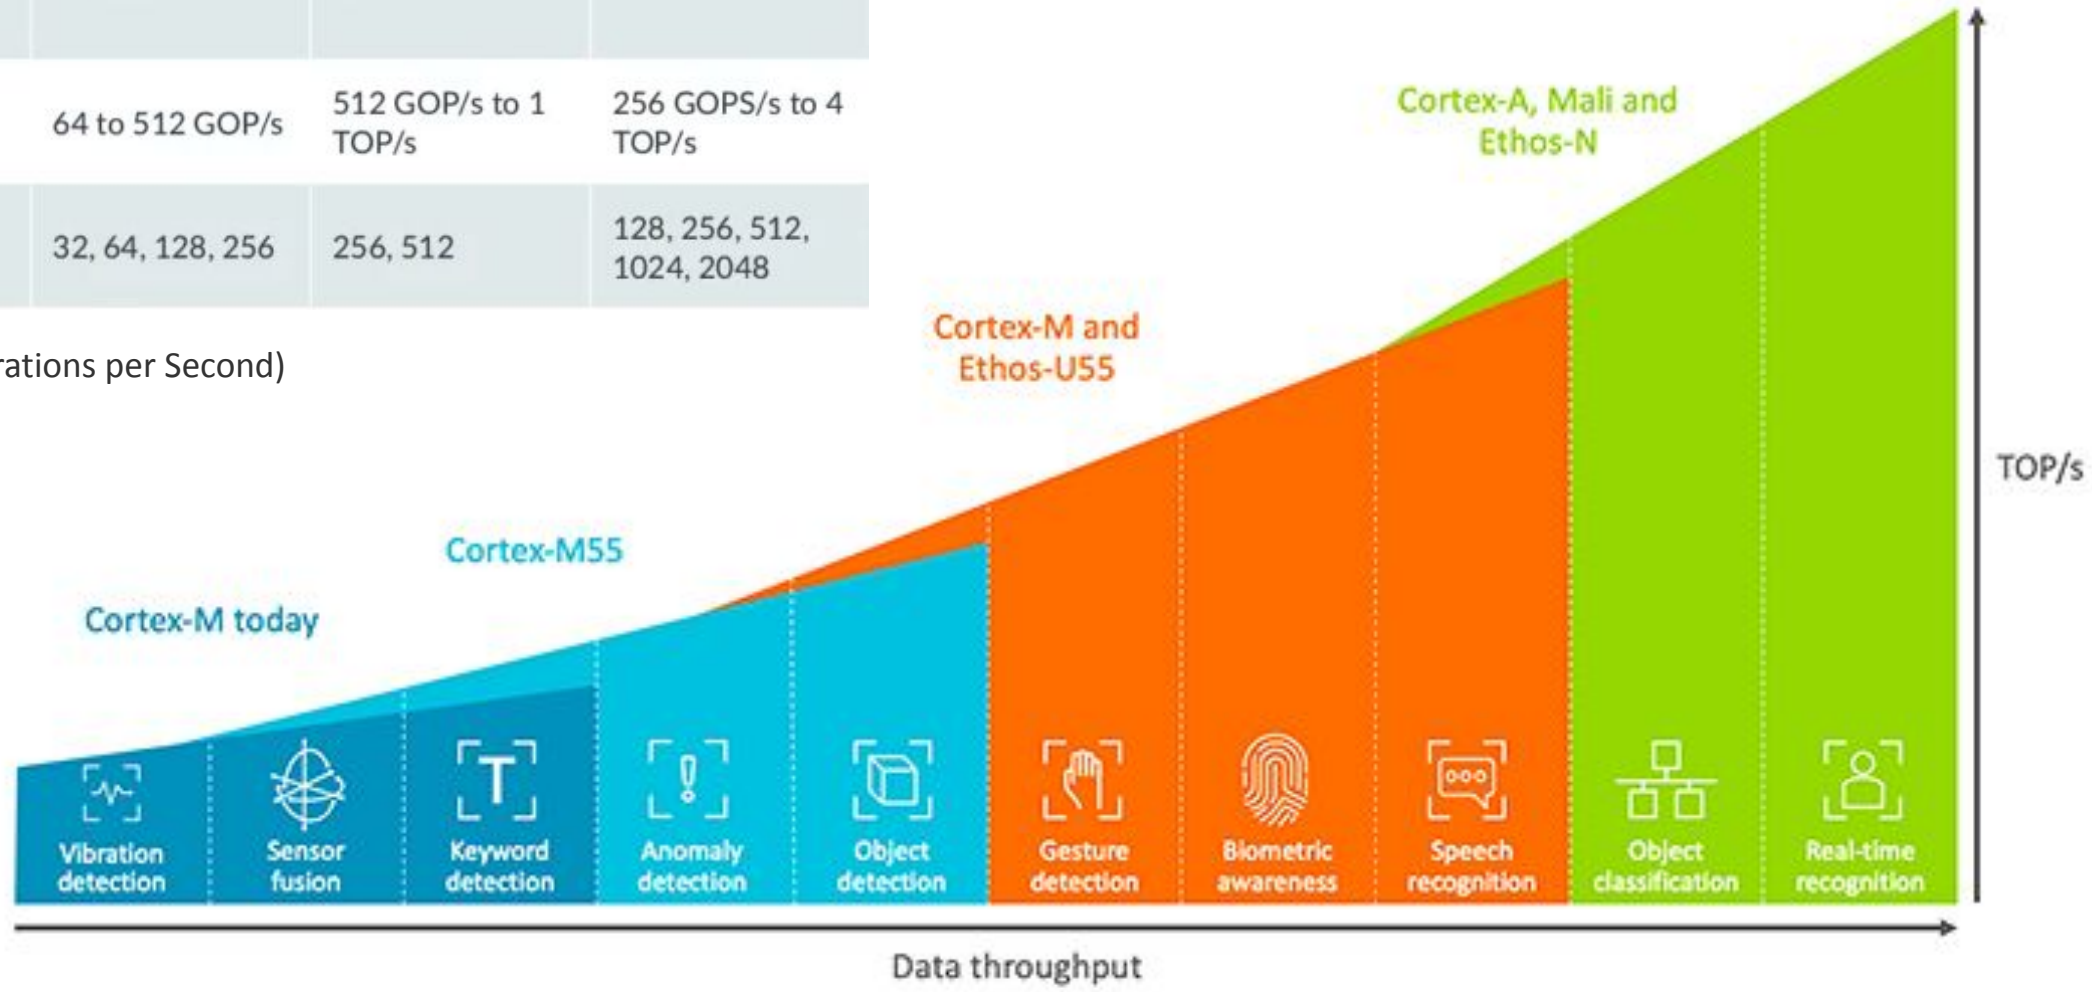

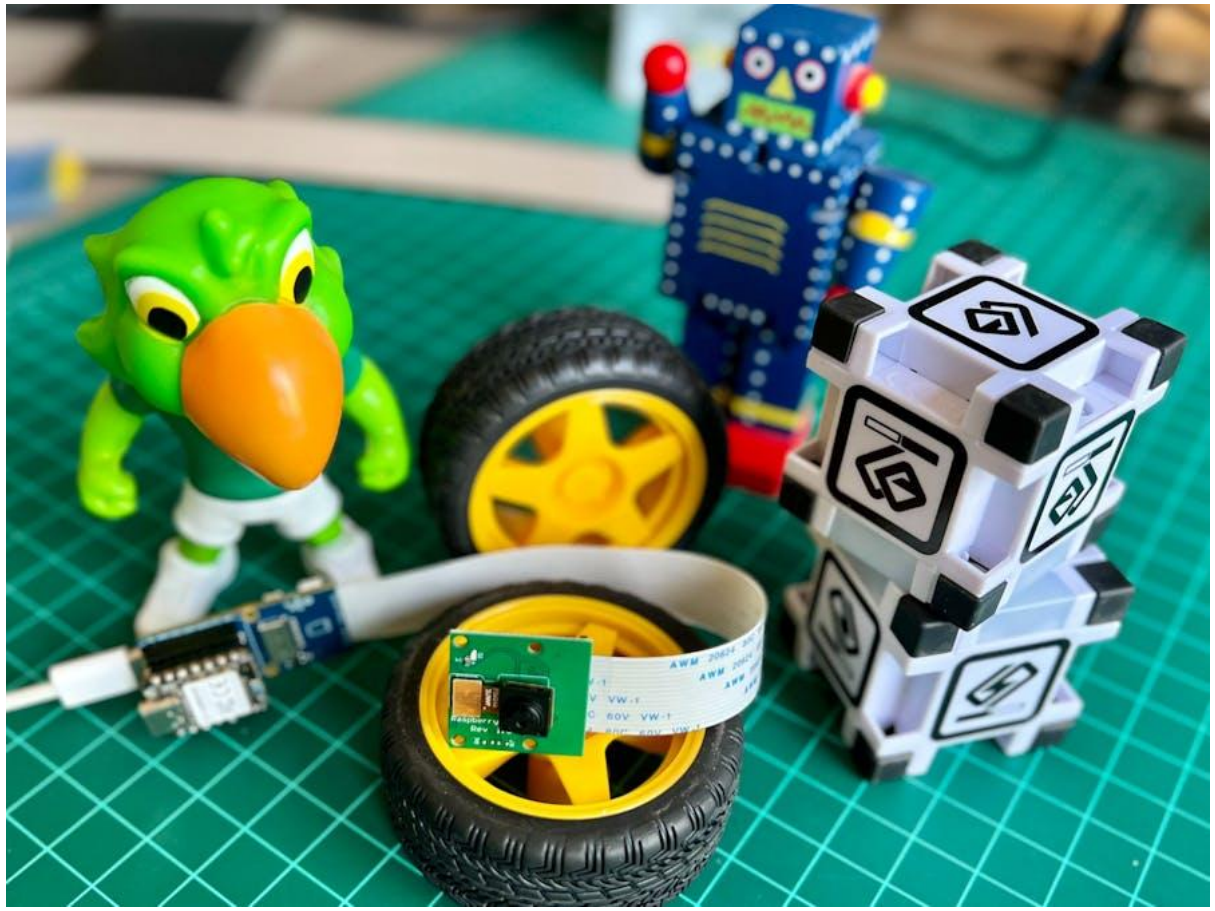

Computer Vision at the Edge with Grove Vision AI Module V2

Exploring Computer Vision applications such as Image Classification, Object Detection, and Pose estimation.

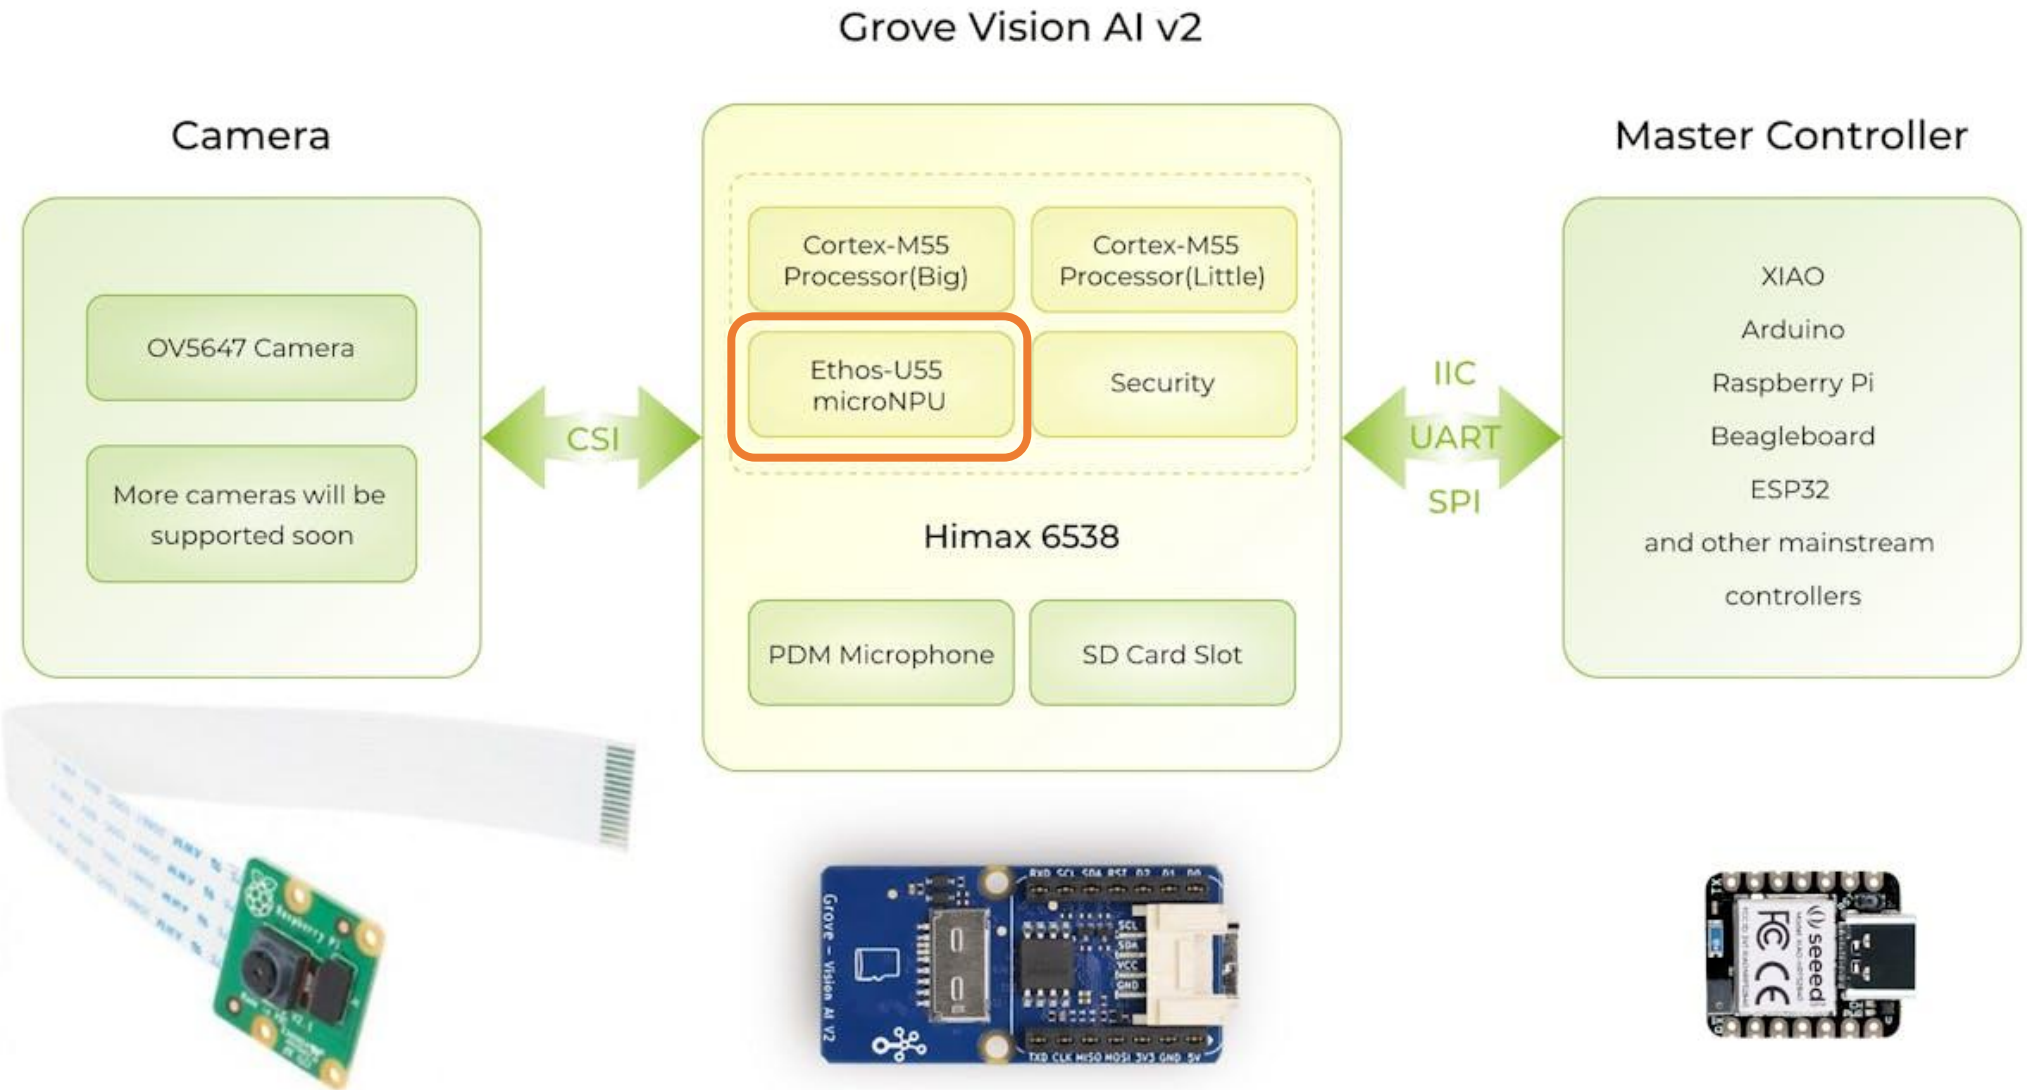
Grove Vision AI v2
Camera
OV5647 Camera
More cameras will be supported soon
CSI
Cortex-M55 Processor(Big)
Cortex-M55 Processor(Little)
Ethos-U55 microNPU
Security
Himax 6538
PDM Microphone
SD Card Slot
IIC
UART
SPI
Master Controller
XIAO
Arduino
Raspberry Pi
Beagleboard
ESP32
and other mainstream controllers

Person Classification

Face Detection

Person Detection

Human Pose Detection

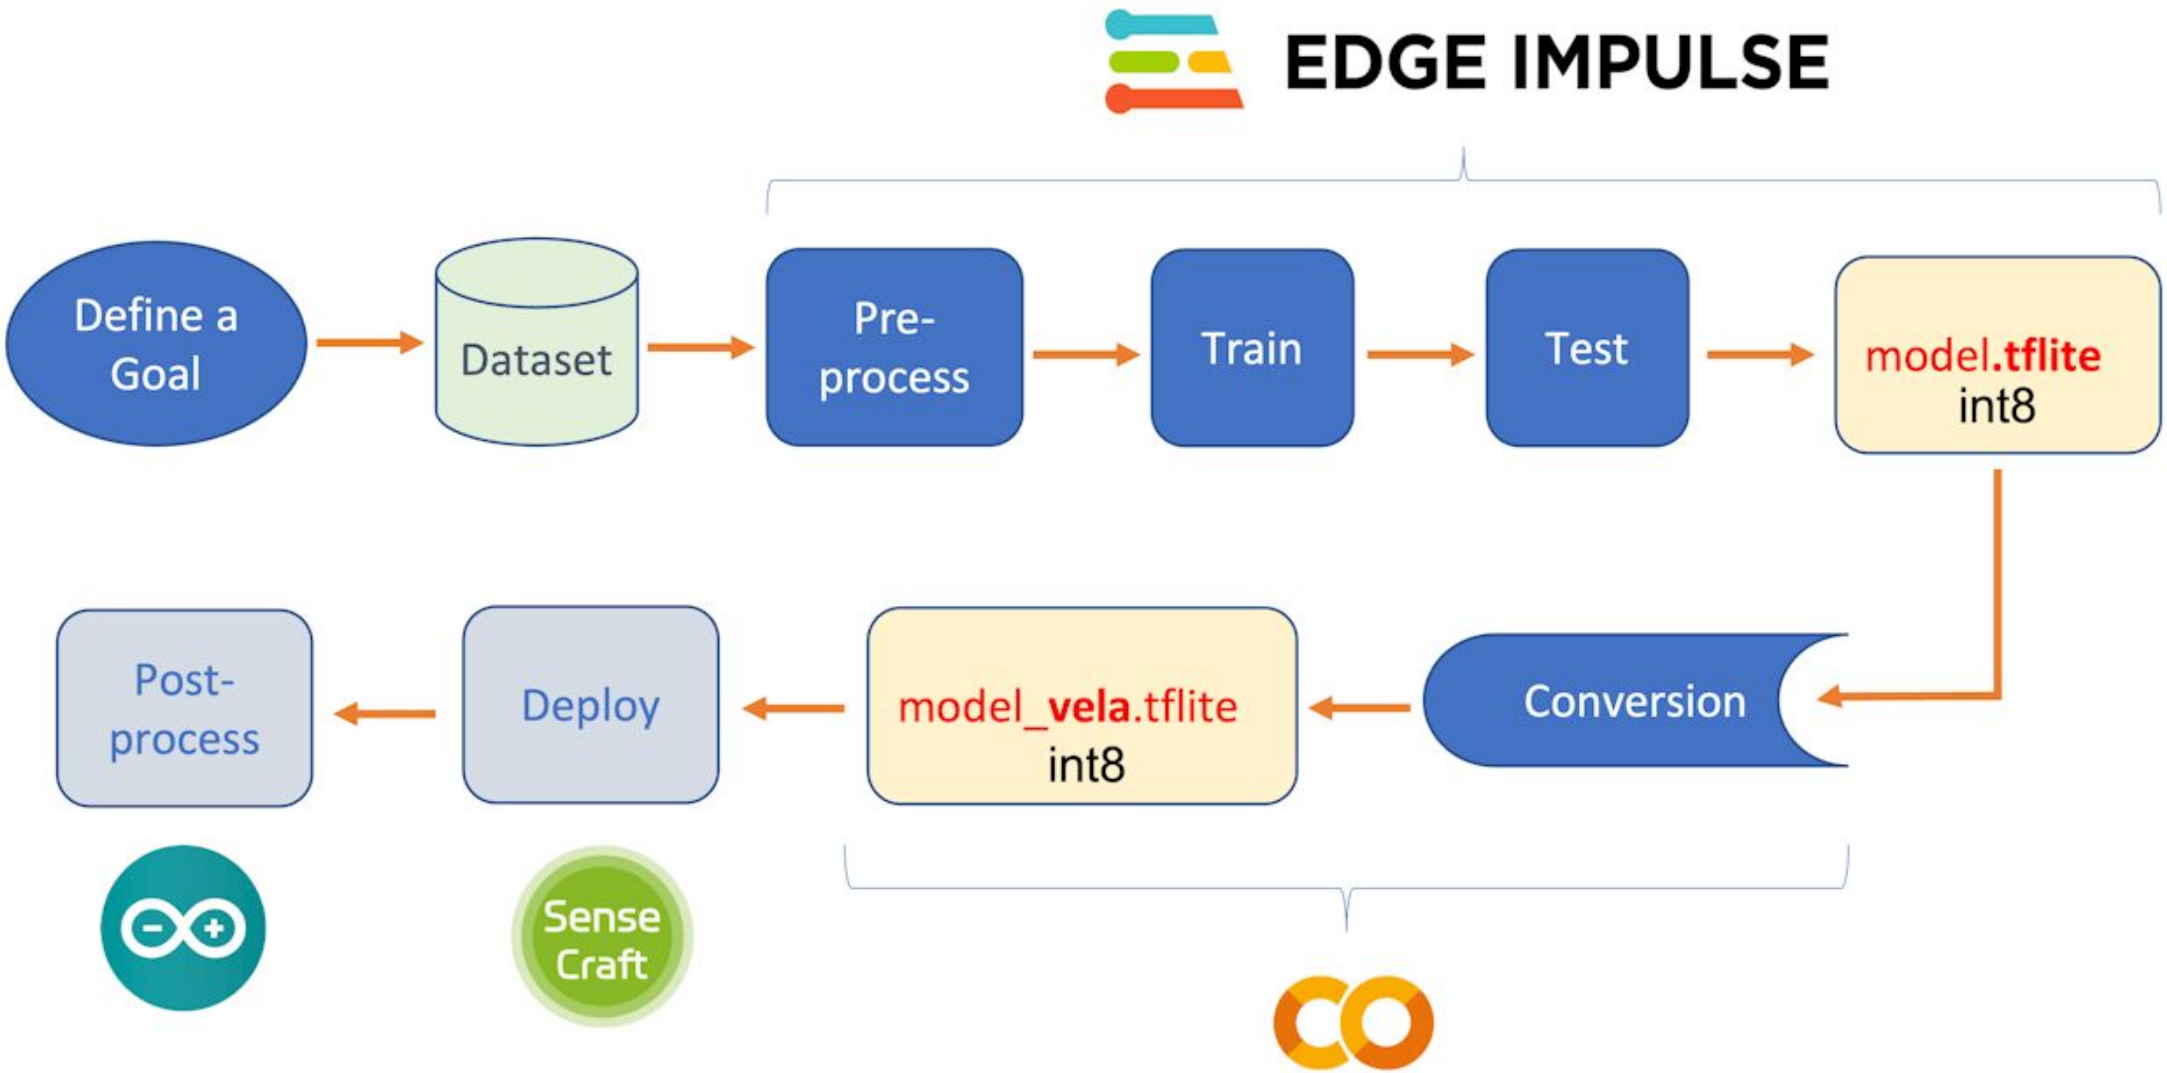
EDGE IMPULSE
Define a Goal
Dataset
Pre-process
Train
Test
model.tflite int8
Conversion
model_vela.tflite int8
Deploy
Post-process
Sense Craft

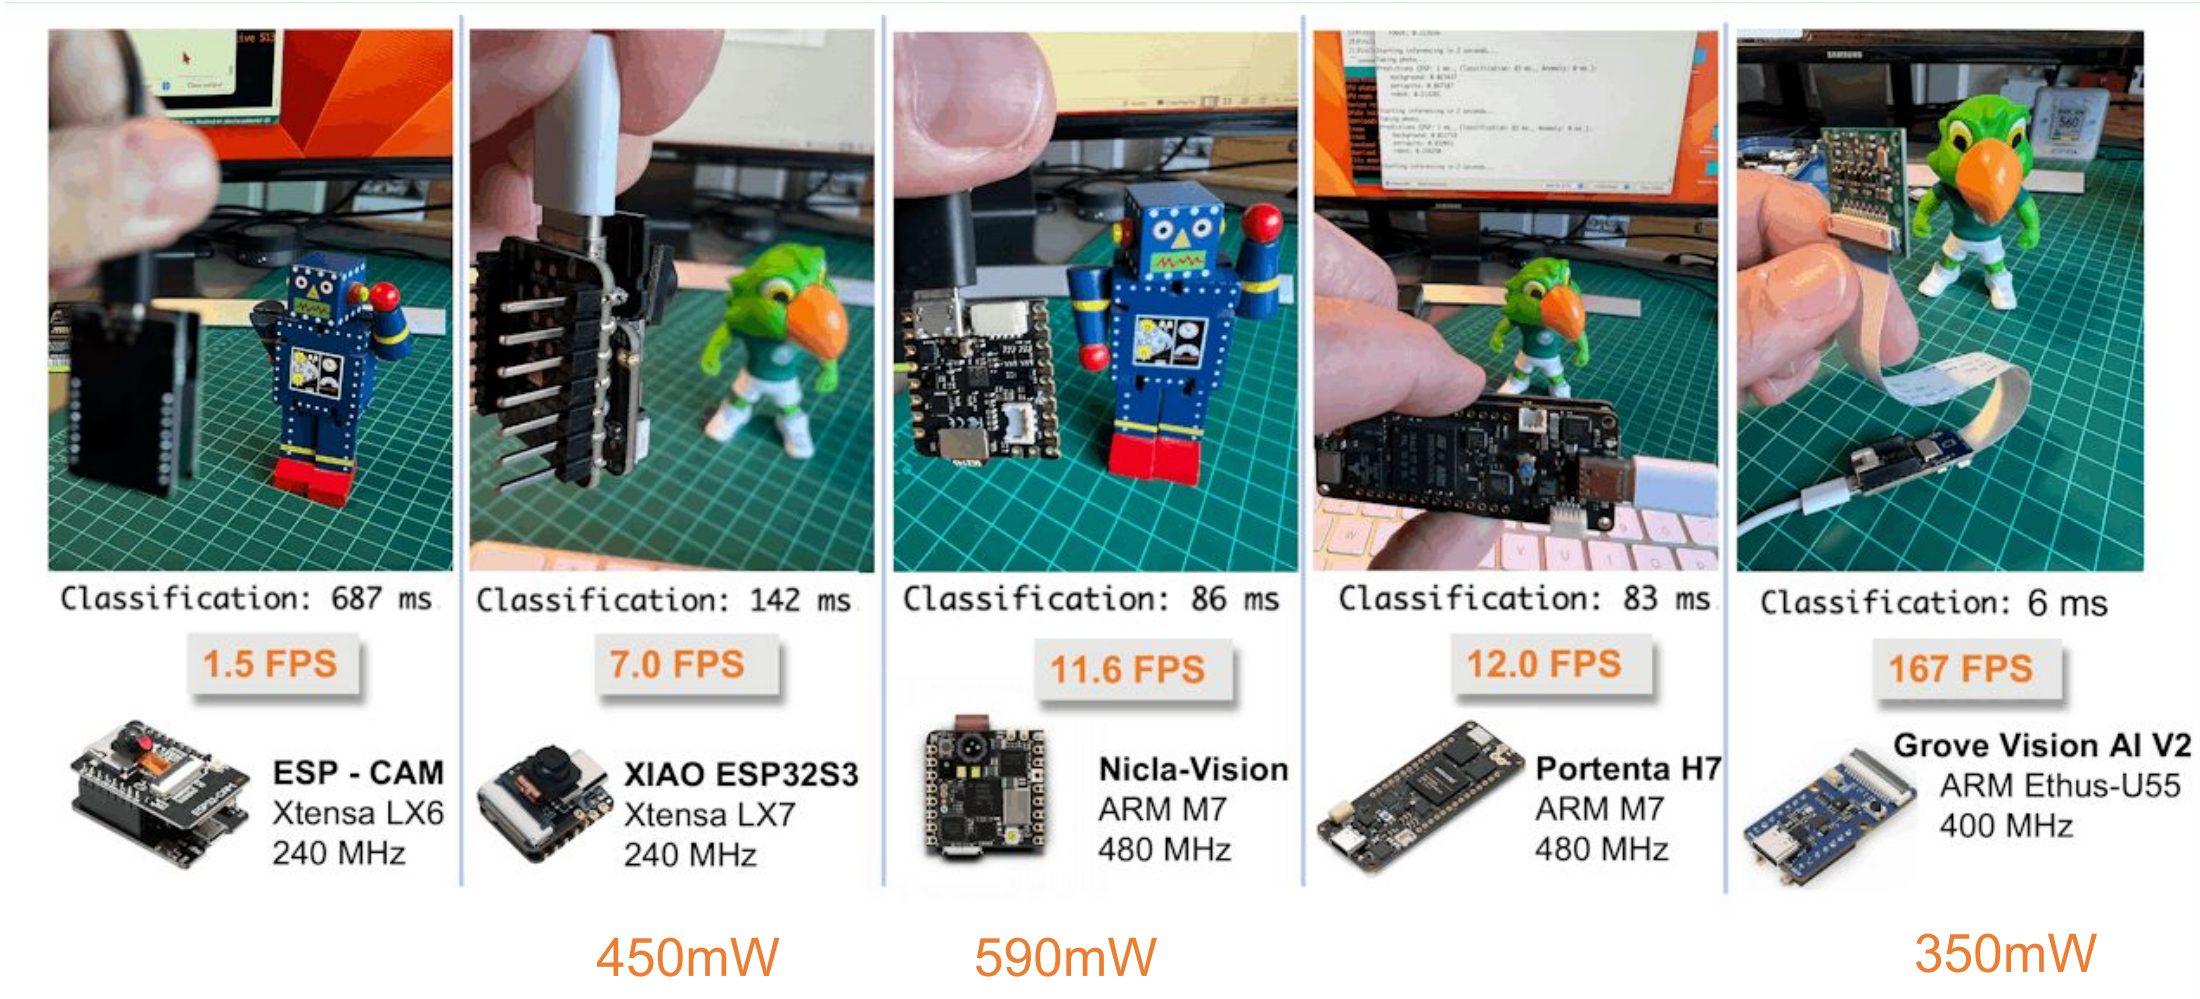
Classification: 687 ms
1.5 FPS
ESP - CAM
Xtensa LX6
240 MHz
Classification: 142 ms
7.0 FPS
XIAO ESP32S3
Xtensa LX7
240 MHz
450mW
Classification: 86 ms
11.6 FPS
Nicla-Vision
ARM M7
480 MHz
590mW
Classification: 83 ms
12.0 FPS
Portenta H7
ARM M7
480 MHz
Classification: 6 ms
167 FPS
Grove Vision AI V2
ARM Ethus-U55
400 MHz
350mW

YOLO

Object Detection Model for Bee Counting

Ultralytics YOLO (You Only Look Once)

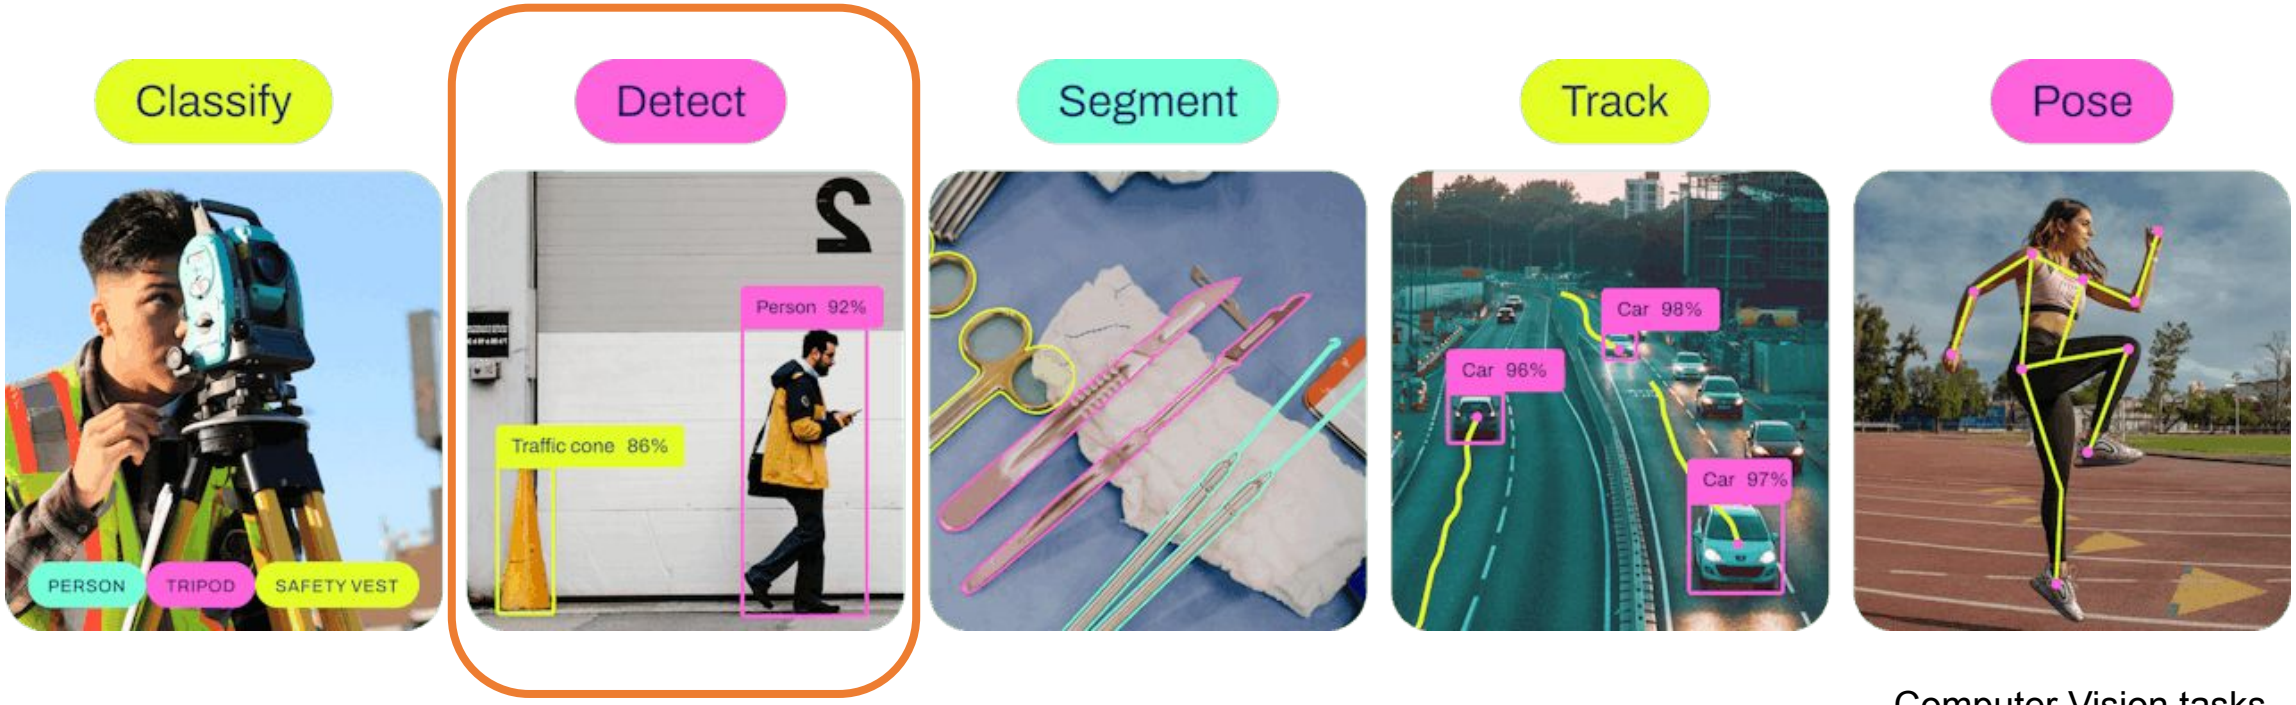

Computer Vision tasks

Real-time **object detection** systems that identify and classify many objects **very fast** in a single image pass.

BuzzTech: Machine Learning at the Edge

Deploying YOLOv8 on Raspberry Pi Zero 2W for Real-Time Bee Counting at the Hive Entrance.

Goal: Estimate the number of bees

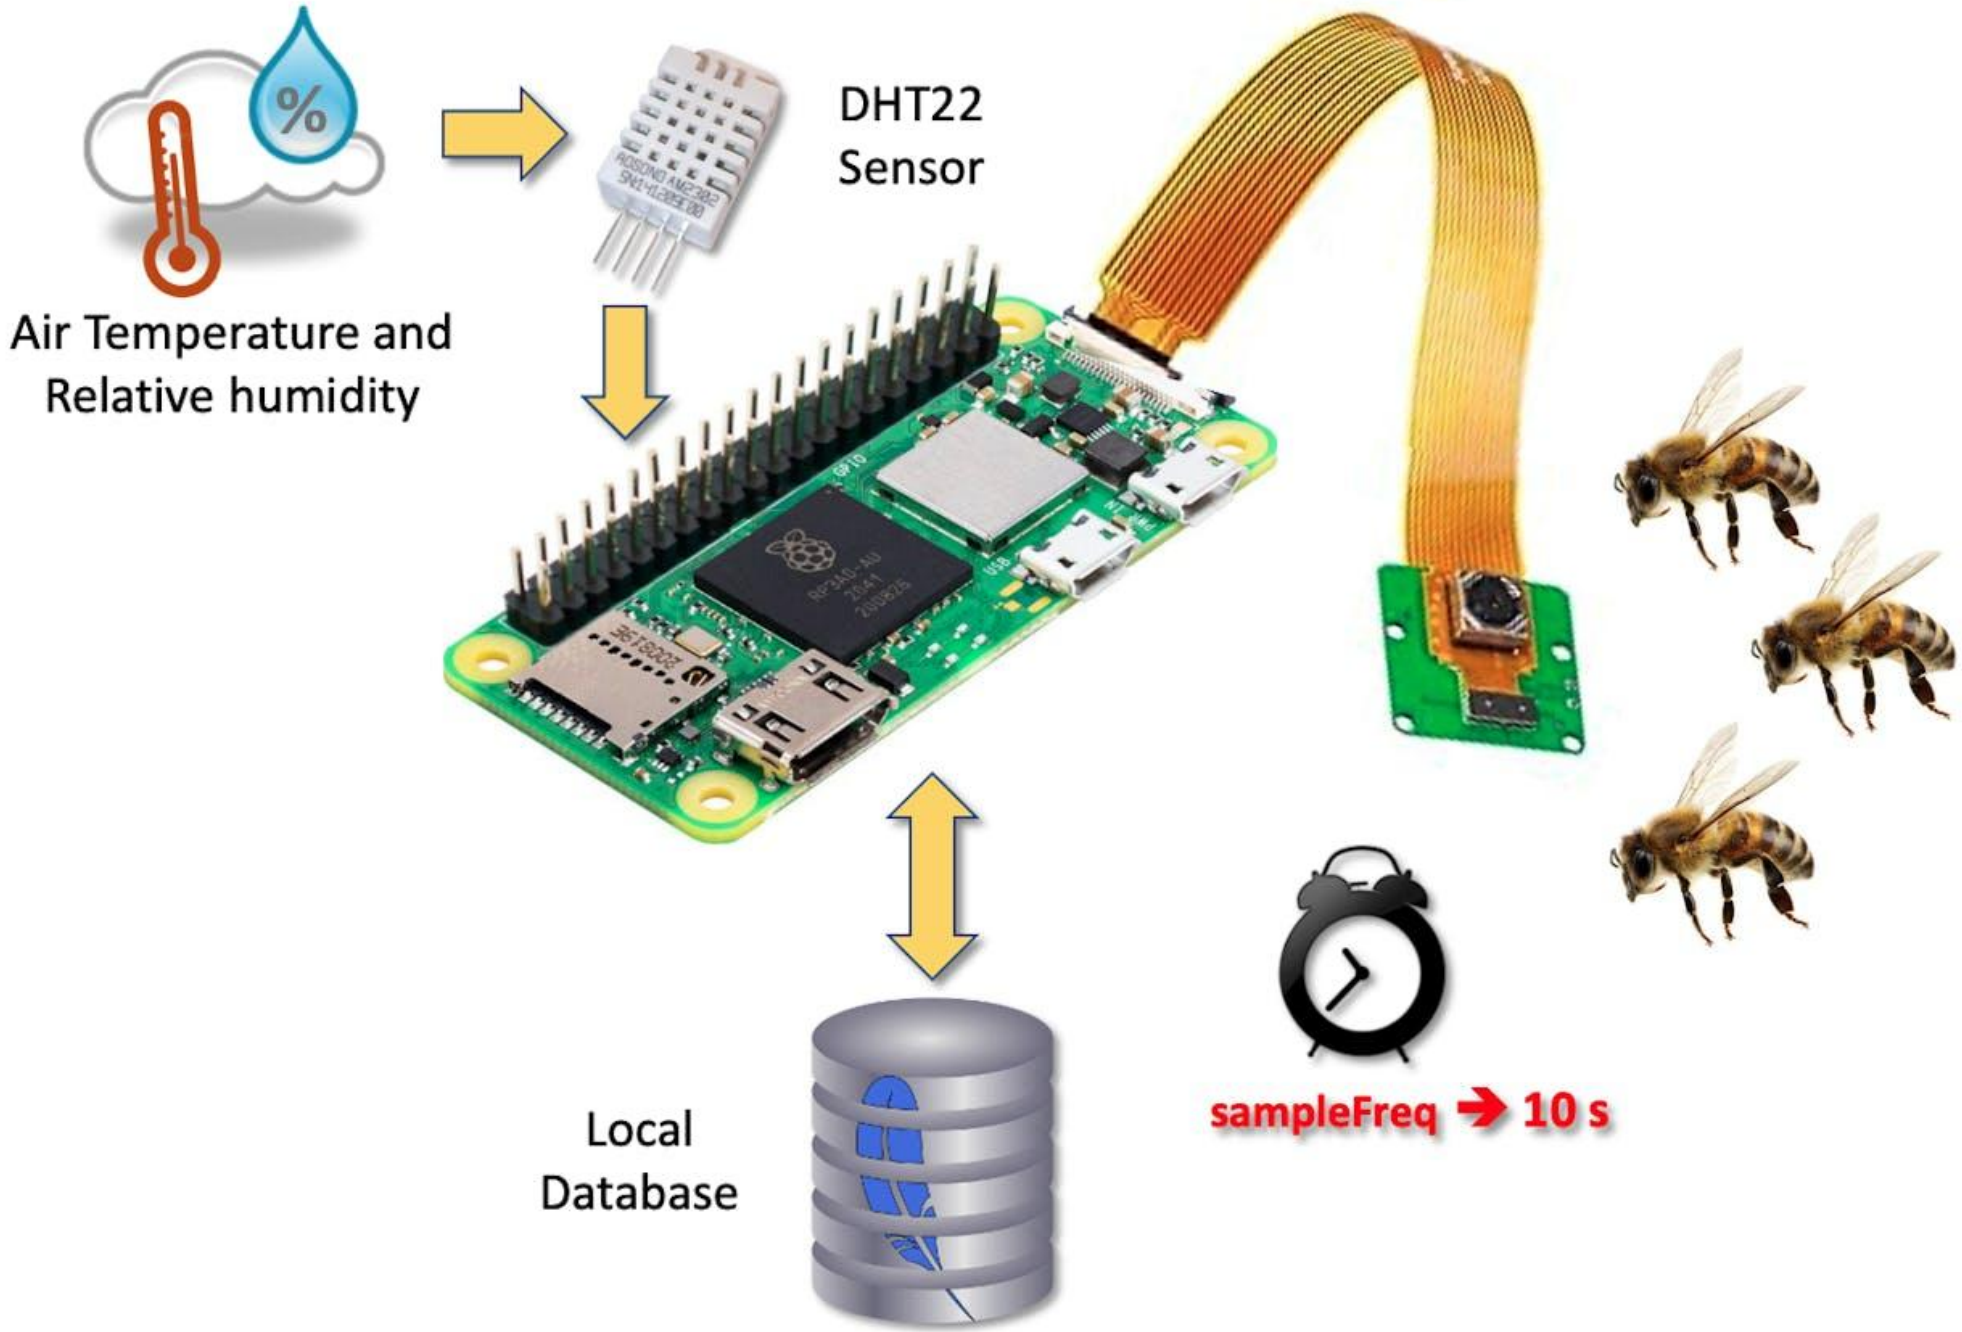
Air Temperature and
Relative humidity
DHT22
Sensor
Local
Database
sampleFreq → 10 s

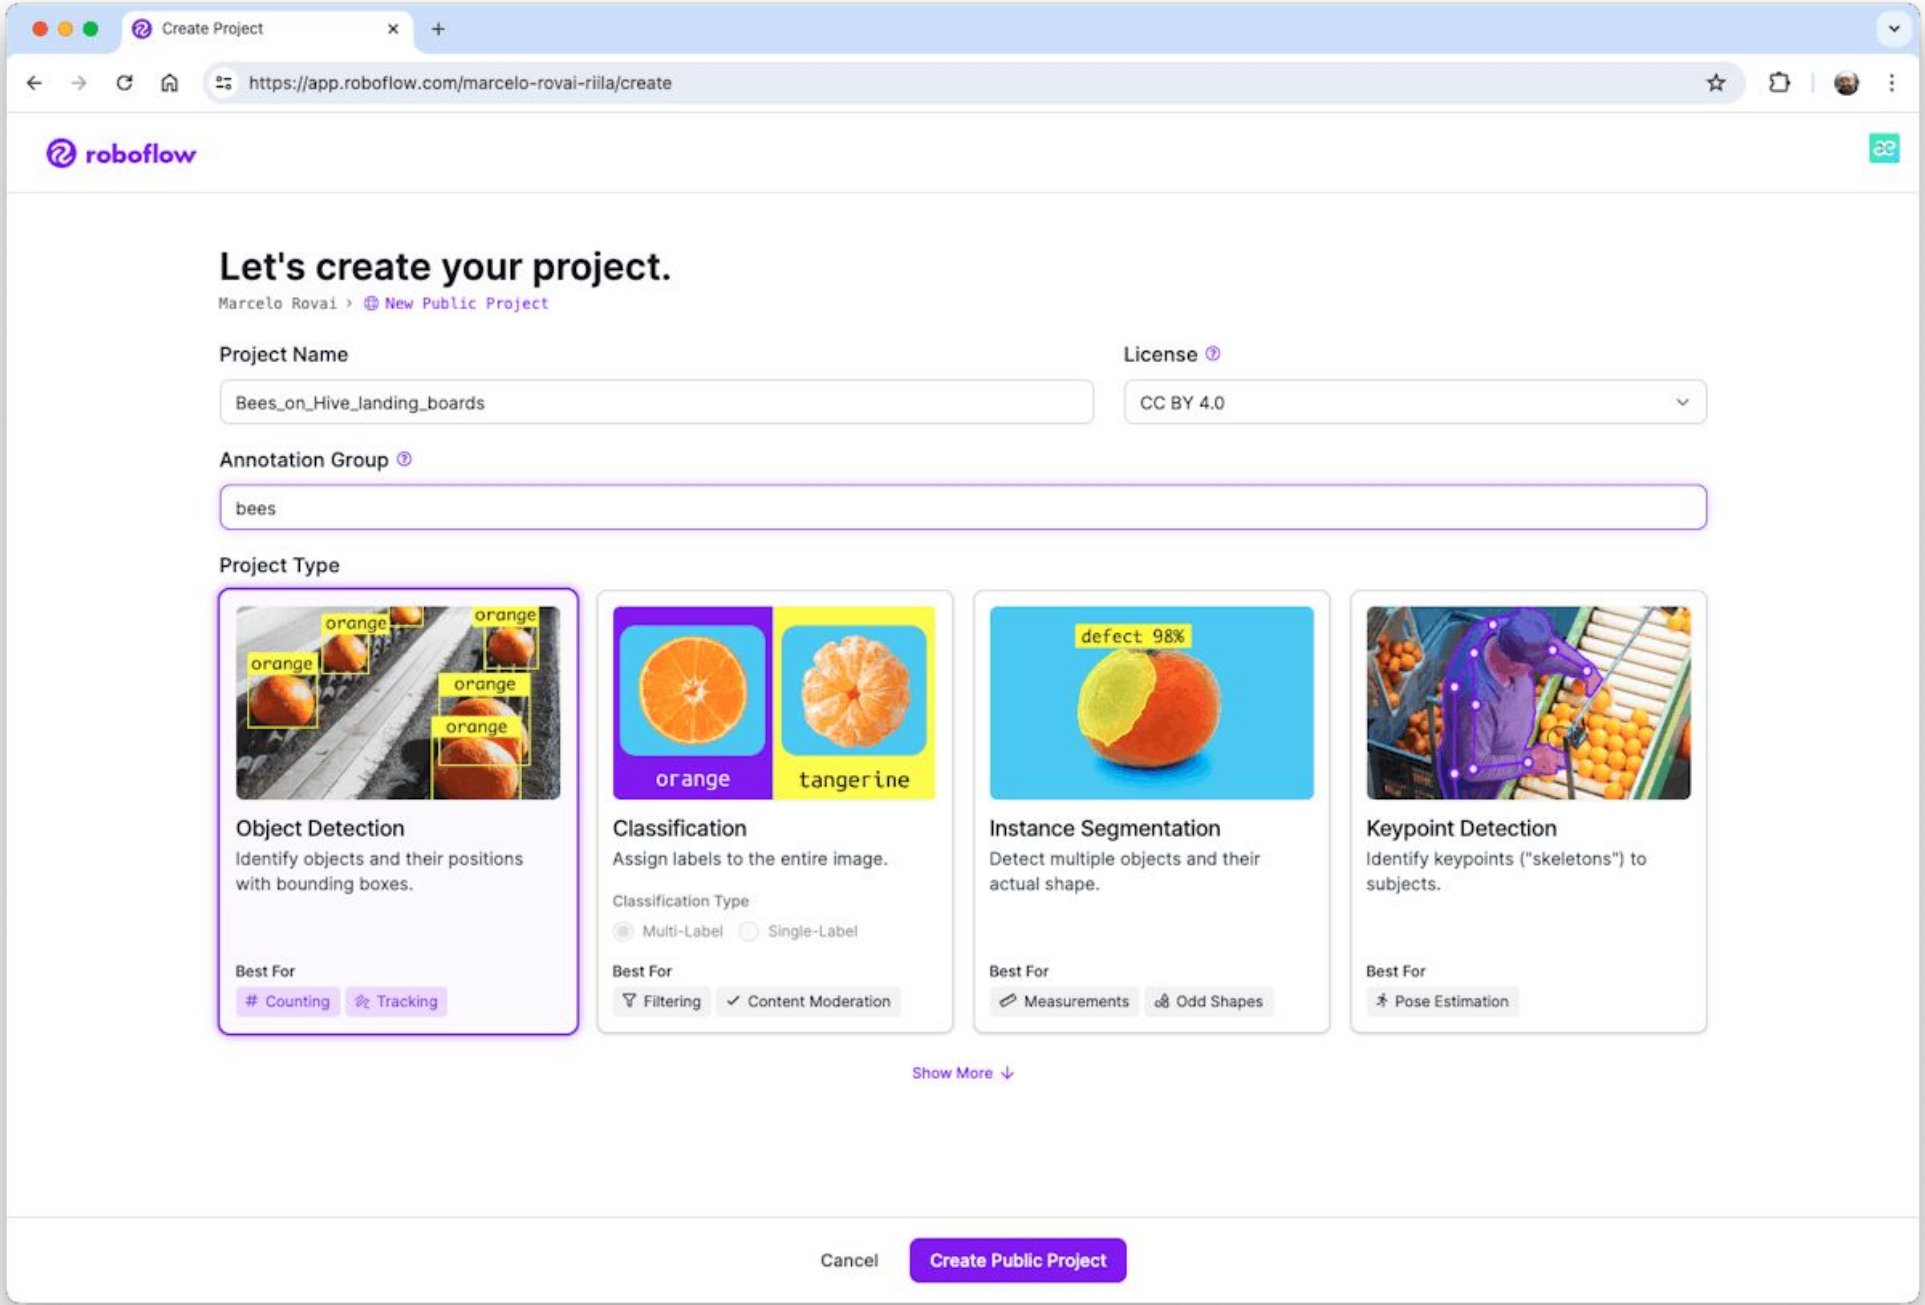
Create Project
https://app.roboflow.com/marcelo-rovai-riila/create
roboflow
Let's create your project.
Marcelo Rovai › New Public Project
Project Name
Bees_on_Hive_landing_boards
License
CC BY 4.0
Annotation Group
bees
Project Type
orange
orange
orange
orange
orange
Object Detection
Identify objects and their positions with bounding boxes.
Best For
Counting
Tracking
orange
tangerine
Classification
Assign labels to the entire image.
Classification Type
Multi-Label
Single-Label
Best For
Filtering
Content Moderation
defect 98%
Instance Segmentation
Detect multiple objects and their actual shape.
Best For
Measurements
Odd Shapes
Keypoint Detection
Identify keypoints ("skeletons") to subjects.
Best For
Pose Estimation
Show More
Cancel
Create Public Project

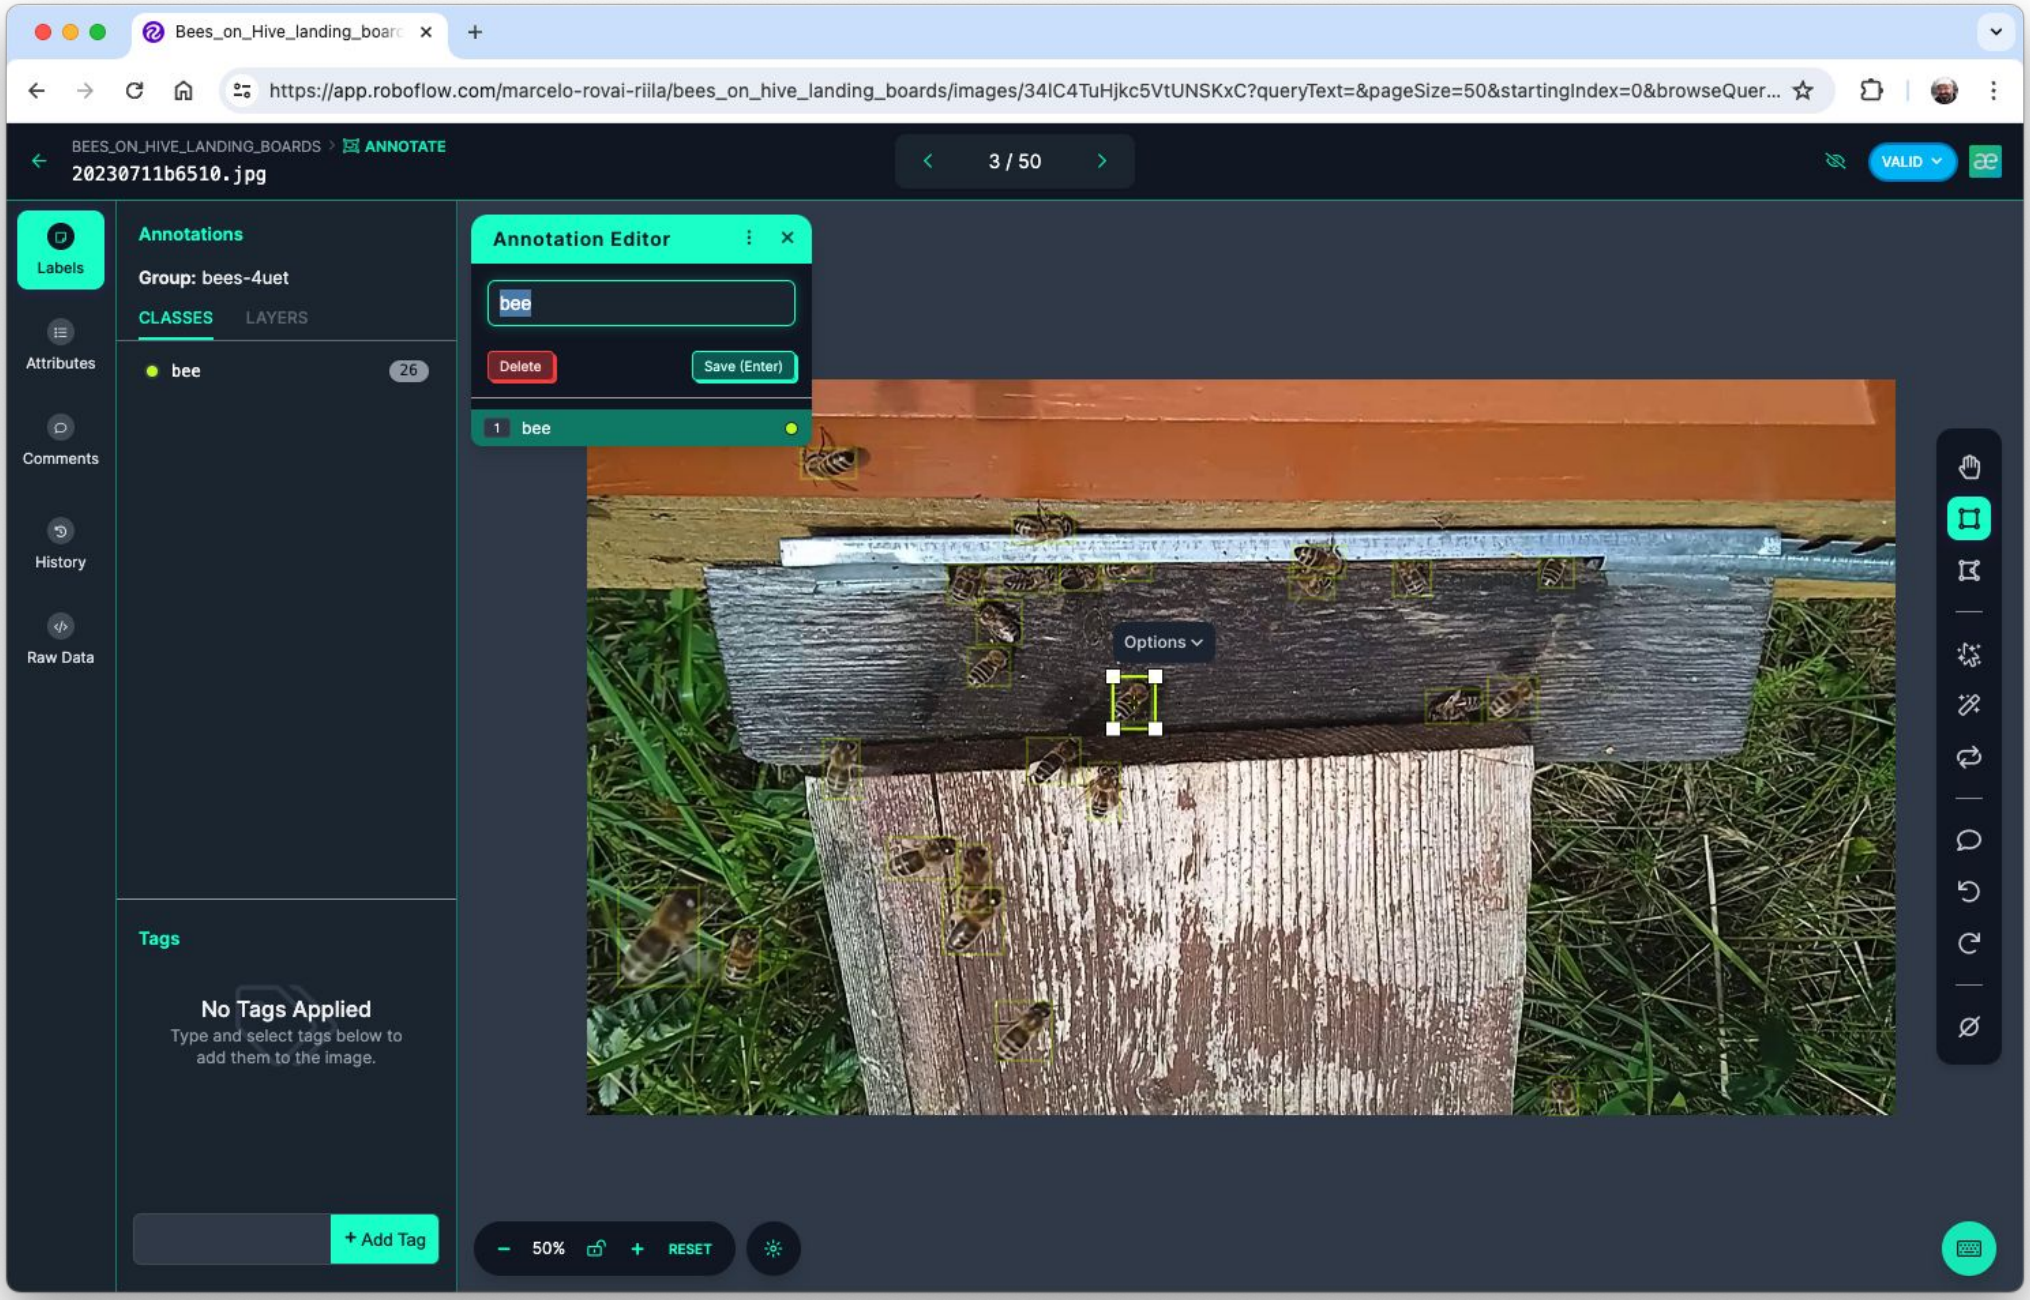
Bees_on_Hive_landing_boarc
https://app.roboflow.com/marcelo-rovai-riila/bees_on_hive_landing_boards/images/34lC4TuHjkc5VtUNSKxC?queryText=&pageSize=50&startingIndex=0&browseQuer...
BEES_ON_HIVE_LANDING_BOARDS
ANNOTATE
20230711b6510.jpg
3 / 50
VALID
Labels
Attributes
Comments
History
Raw Data
Annotations
Group: bees-4uet
CLASSES
LAYERS
bee
26
Tags
No Tags Applied
Type and select tags below to add them to the image.
+ Add Tag
Annotation Editor
bee
Delete
Save (Enter)
1 bee
Options
50%
RESET

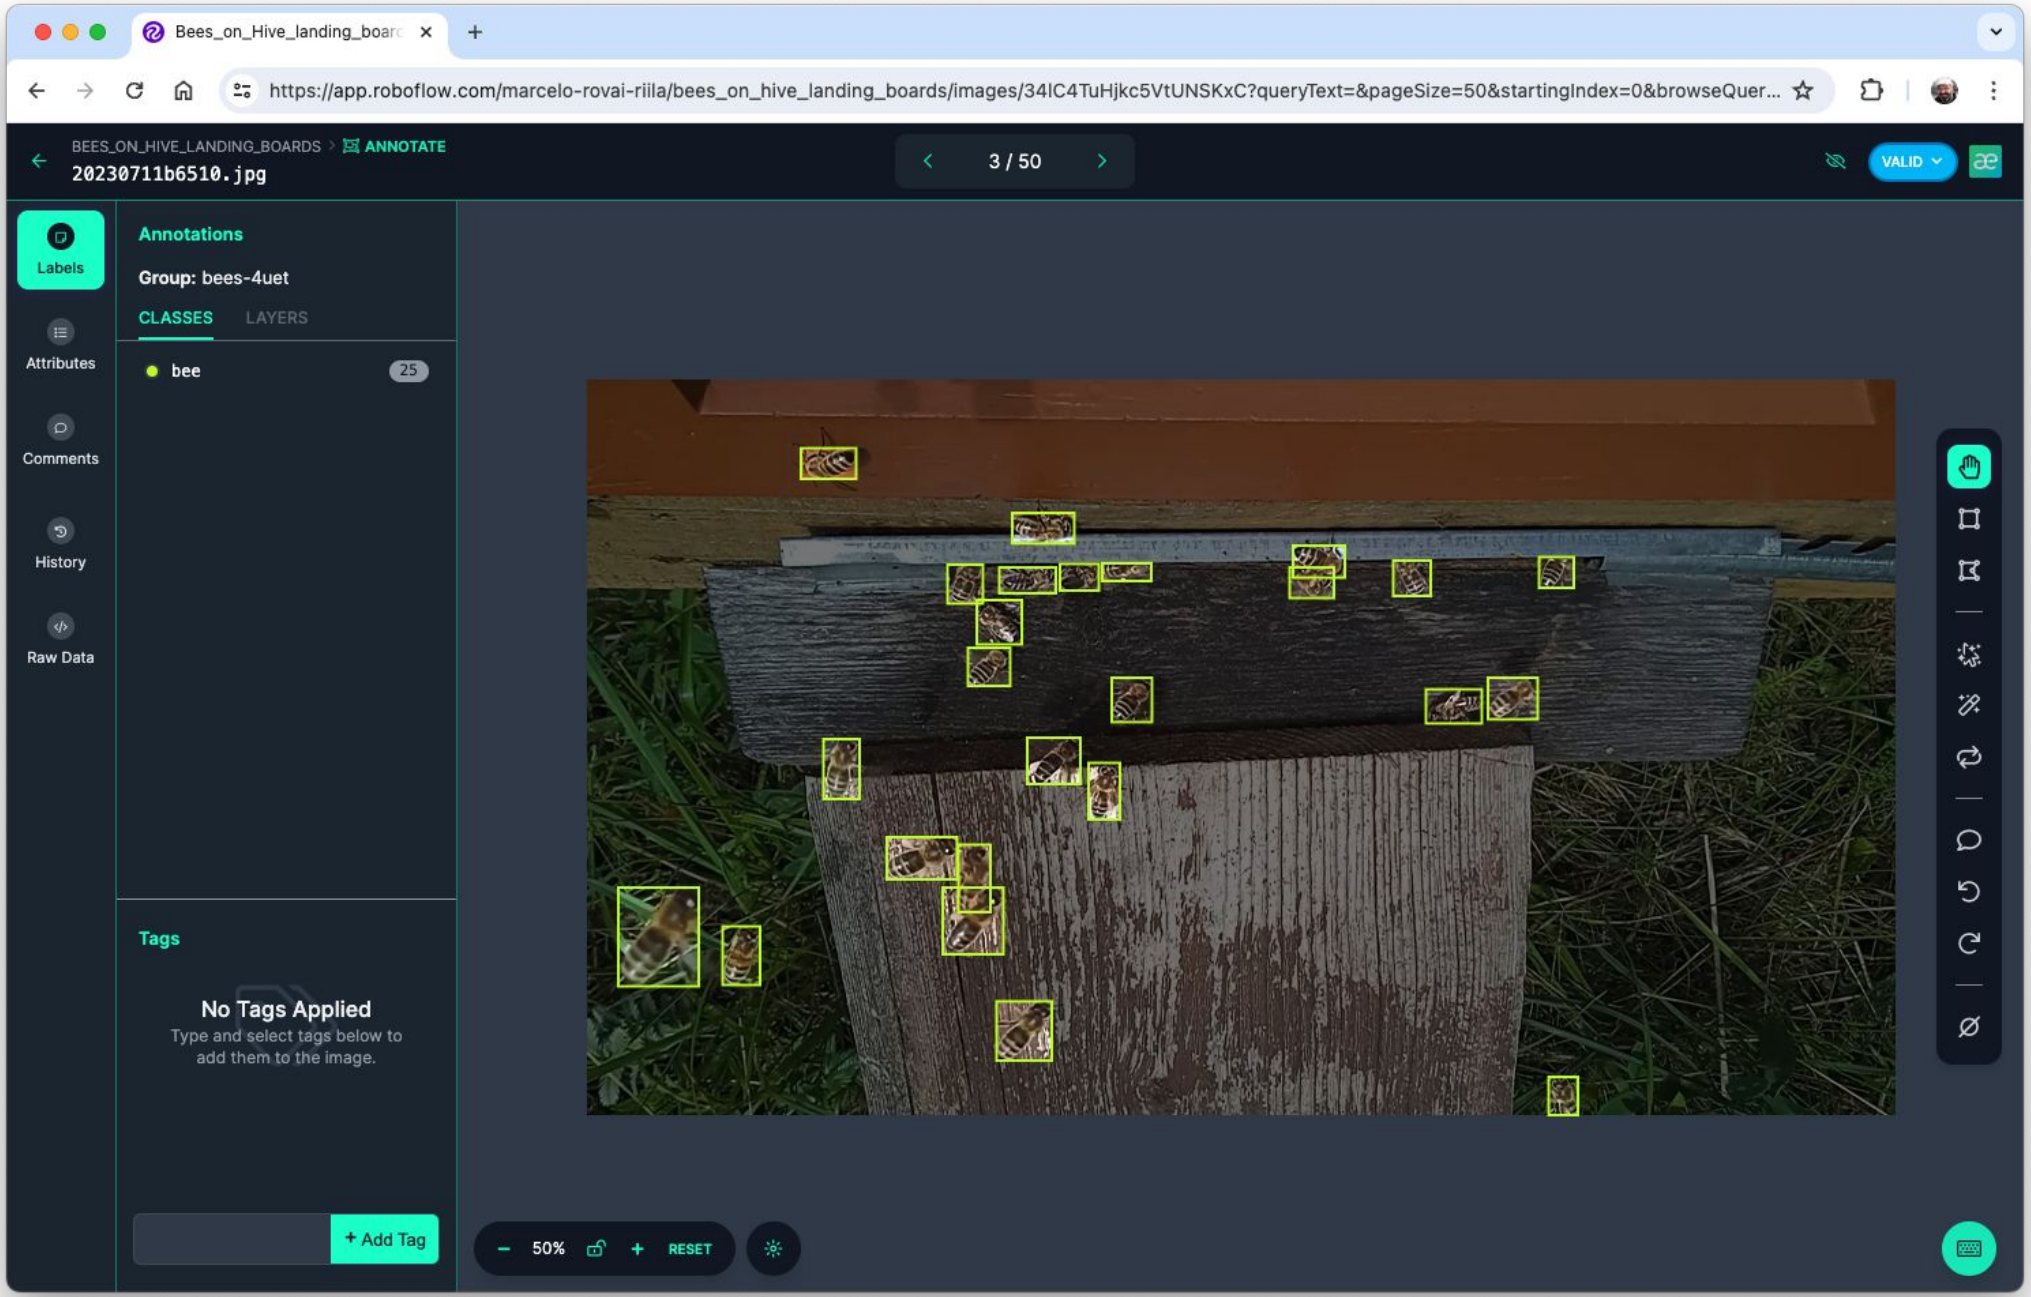
Bees_on_Hive_landing_boarc
https://app.roboflow.com/marcelo-rovai-riila/bees_on_hive_landing_boards/images/34lC4TuHjkc5VtUNSKxC?queryText=&pageSize=50&startingIndex=0&browseQuer...
BEES_ON_HIVE_LANDING_BOARDS
ANNOTATE
20230711b6510.jpg
3 / 50
VALID
Labels
Attributes
Comments
History
Raw Data
Annotations
Group: bees-4uet
CLASSES
LAYERS
bee
25
Tags
No Tags Applied
Type and select tags below to add them to the image.
+ Add Tag
50%
RESET

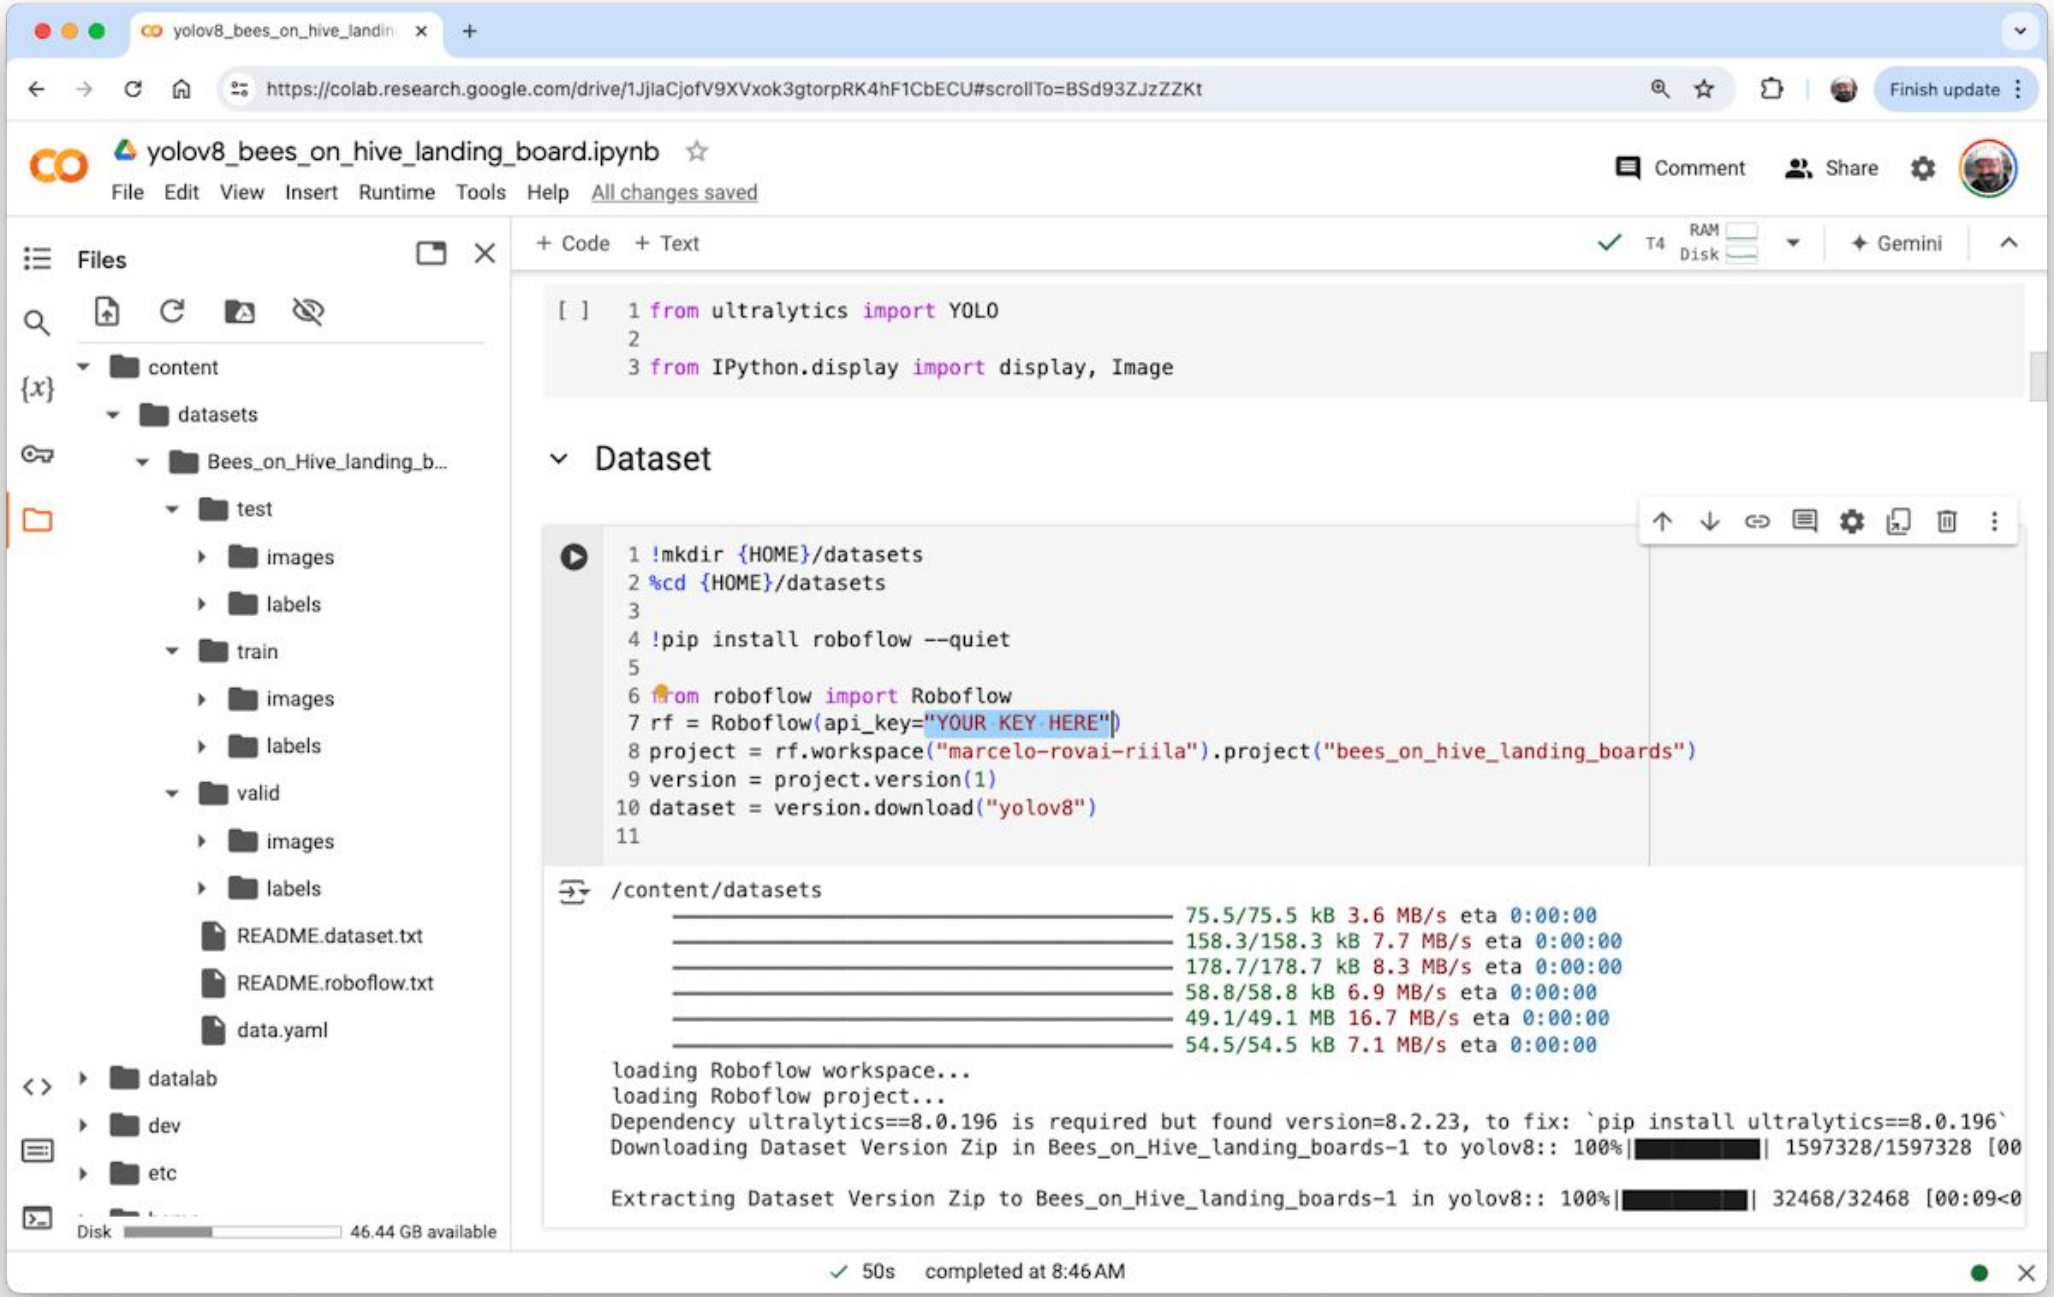
yolov8_bees_on_hive_landin
https://colab.research.google.com/drive/1JjlaCjofV9XVxok3gtorpRK4hF1CbECU#scrollTo=BSd93ZJzZZKt
Finish update
yolov8_bees_on_hive_landing_board.ipynb
File Edit View Insert Runtime Tools Help All changes saved
Comment
Share
Files
+ Code + Text
T4
RAM
Disk
Gemini
content
datasets
Bees_on_Hive_landing_b...
test
images
labels
train
images
labels
valid
images
labels
README.dataset.txt
README.roboflow.txt
data.yaml
datalab
dev
etc
Disk 46.44 GB available
[ ] 1 from ultralytics import YOLO
2
3 from IPython.display import display, Image
Dataset
1 !mkdir {HOME}/datasets
2 %cd {HOME}/datasets
3
4 !pip install roboflow --quiet
5
6 from roboflow import Roboflow
7 rf = Roboflow(api_key="YOUR KEY HERE")
8 project = rf.workspace("marcelo-rovai-riila").project("bees_on_hive_landing_boards")
9 version = project.version(1)
10 dataset = version.download("yolov8")
11
/content/datasets
75.5/75.5 kB 3.6 MB/s eta 0:00:00
158.3/158.3 kB 7.7 MB/s eta 0:00:00
178.7/178.7 kB 8.3 MB/s eta 0:00:00
58.8/58.8 kB 6.9 MB/s eta 0:00:00
49.1/49.1 MB 16.7 MB/s eta 0:00:00
54.5/54.5 kB 7.1 MB/s eta 0:00:00
loading Roboflow workspace...
loading Roboflow project...
Dependency ultralytics==8.0.196 is required but found version=8.2.23, to fix: `pip install ultralytics==8.0.196`
Downloading Dataset Version Zip in Bees_on_Hive_landing_boards-1 to yolov8:: 100%| 1597328/1597328 [00
Extracting Dataset Version Zip to Bees_on_Hive_landing_boards-1 in yolov8:: 100%| 32468/32468 [00:09<0
50s completed at 8:46 AM

Raw Dataset

Dataset Labeling (Roboflow)

Model Training **Ultralytics** (PyTorch)

Model Inference **Ultralytics - NCNN** (Rasp)

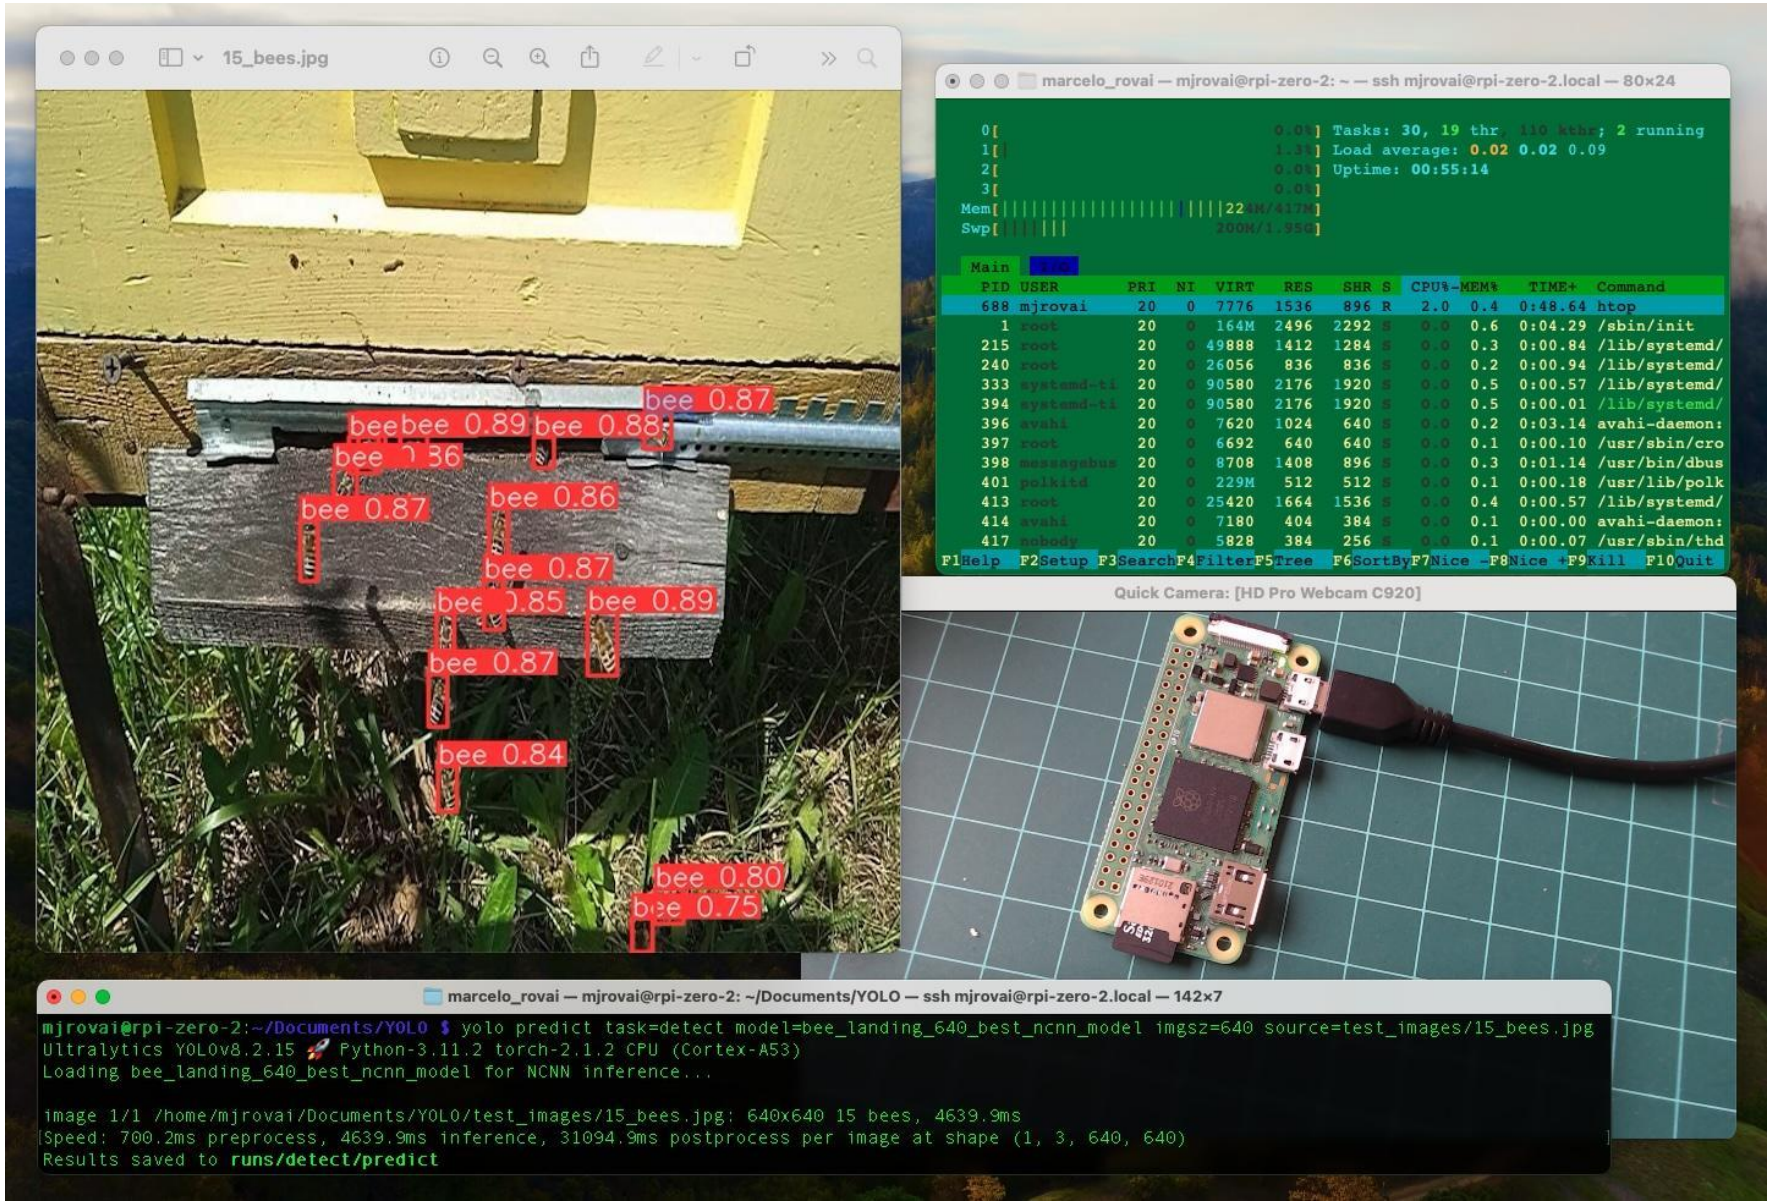

Thanks

TINYML4D