



Image Classification

Hands-On

Prof. Marcelo J. Rovai
UNIFEI – Federal University of Itajuba, Brazil
rovai@unifei.edu.br



Computer Vision Main Types

Image Classification (Multi-Class Classification)




Cat: 70%



Dog: 80%

Object Detection Multi-Label Classification + Object Localization



Computer Vision Main Types

Image Classification (Multi-Class Classification)



Cat: 70%



Dog: 80%

Object Detection Multi-Label Classification + Object Localization




Image Classification Project: Package Inspection



[Deep Learning at the Edge Simplifies Package Inspection](#)



Image Classification Project

Decide a Goal

- Possible Images:
 - medicine
 - background



Image Classification Application: Design, Train, Test and Deploy

<https://studio.edgeimpulse.com/public/114253/latest>



Select project - Edge Impulse x +

studio.edgeimpulse.com/studio/select-project

EDGE IMPULSE

Marcelo Rovai

Select project

Create a new project

Enter the name for your new project:

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. [Learn more](#)

Create new project

Select y

NAME

Marcel

Marcel

Marcel

Marcel

Marcel

Marcel

Marcel

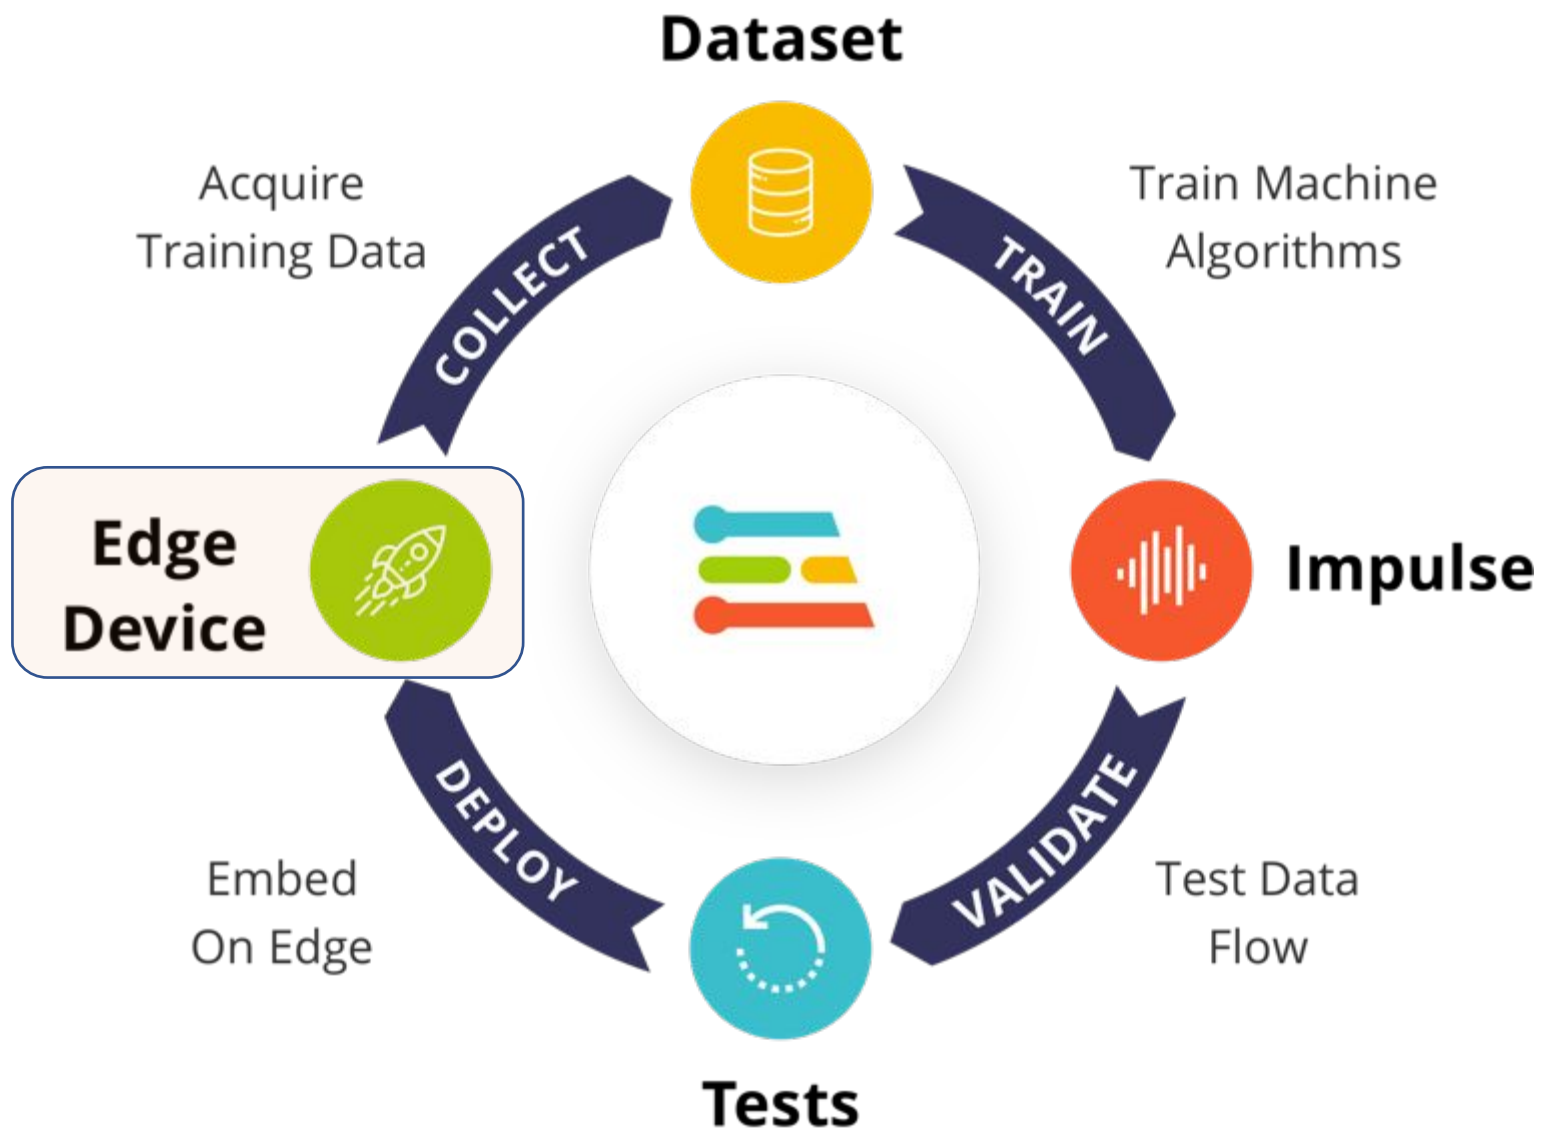
Marcel

Marcelo Rovai / SciTinyML22-KWS


Marcelo Rovai / Cifar10_Image_Classification

Marcelo Rovai / IESTI01-Cifar10_Classification

Marcelo Rovai / Bean Disease Classifier




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



← Label: **medicine**

Label: **background**



EDGE IMPULSE Home API Reference Log In

Guides > Arduino Nano 33 BLE Sense Search 🔍

DOCUMENTATION

- Getting Started
- API and SDK references
- What is embedded ML, anyway?
- Frequently asked questions

DEVELOPMENT BOARDS

- Overview
- ST B-L475E-IOT01A
- Arduino Nano 33 BLE Sense**
- Eta Compute ECM3532 AI Sensor
- Eta Compute ECM3532 AI Vision
- OpenMV Cam H7 Plus
- Himax WE-I Plus
- Nordic Semi nRF52840 DK
- Nordic Semi nRF5340 DK
- SiLabs Thunderboard Sense 2
- Sony's Spresense
- Arduino Portenta H7 + Vision shield (preview)
- Raspberry Pi 4
- NVIDIA Jetson Nano
- Mobile phone
- Porting guide

COMMUNITY BOARDS

- Seeed Wio Terminal
- Agora Product Development Kit


EDGE IMPULSE FOR LINUX

- Edge Impulse for Linux
- Linux Node.js SDK
- Linux Go SDK
- Linux C++ SDK
- Linux Python SDK

Arduino Nano 33 BLE Sense

The Arduino Nano 33 BLE Sense is a tiny development board with a Cortex-M4 microcontroller, motion sensors, a microphone and BLE - and it's fully supported by Edge Impulse. You'll be able to sample raw data, build models, and deploy trained machine learning models directly from the studio. It's available for around 30 USD from [Arduino](#) and a wide range of distributors.

The Edge Impulse firmware for this development board is open source and hosted on GitHub: [edgeimpulse/firmware-arduino-nano-33-ble-sense](#).



Arduino Nano 33 BLE Sense

Installing dependencies

To set this device up in Edge Impulse, you will need to install the following software:

1. [Edge Impulse CLI](#).
2. [Arduino CLI](#).
 - Here's an [instruction video for Windows](#).
 - The [Arduino website](#) has instructions for macOS and Linux.
3. On Linux:
 - GNU Screen: install for example via `sudo apt install screen`.

Note that the **1. Edge Impulse CLI** is not necessary for Arduino Nano-33 if you use WebUSB.

DOCUMENTATION

- Getting Started
- API and SDK references
- What is embedded ML, anyway?
- Frequently asked questions



DEVELOPMENT BOARDS

- Overview
 - ST B-L475E-IOT01A
 - Arduino Nano 33 BLE Sense**
 - Eta Compute ECM3532 AI Sensor
 - Eta Compute ECM3532 AI Vision
 - OpenMV Cam H7 Plus
 - Himax WE-I Plus
 - Nordic Semi nRF52840 DK
 - Nordic Semi nRF5340 DK
 - SiLabs Thunderboard Sense 2
 - Sony's Spresense
 - Arduino Portenta H7 + Vision shield (preview)
 - Raspberry Pi 4
 - NVIDIA Jetson Nano
 - Mobile phone
 - Porting guide
- ## COMMUNITY BOARDS
- Seeed Wio Terminal
 - Agora Product Development Kit



1. Connect the development board to your computer

Use a micro-USB cable to connect the development board to your computer. Then press RESET twice to launch into the bootloader. The on-board LED should start pulsating to indicate this.



Press RESET twice quickly to launch the bootloader on the Arduino Nano 33 BLE Sense.

2. Update the firmware

The development board does not come with the right firmware yet. To update the firmware:

1. [Download the latest Edge Impulse firmware](#), and unzip the file.
2. Open the flash script for your operating system (`flash_windows.bat` , `flash_mac.command` or `flash_linux.sh`) to flash the firmware.
3. Wait until flashing is complete, and press the RESET button once to launch the new firmware.

3. Setting keys

DOCUMENTATION


- Getting Started
- API and SDK references
- What is embedded ML, anyway?
- Frequently asked questions

DEVELOPMENT BOARDS

- Overview
- ST B-L475E-IOT01A
- Arduino Nano 33 BLE Sense**
- Eta Compute ECM3532 AI Sensor
- Eta Compute ECM3532 AI Vision
- OpenMV Cam H7 Plus
- Himax WE-I Plus
- Nordic Semi nRF52840 DK
- Nordic Semi nRF5340 DK
- SiLabs Thunderboard Sense 2
- Sony's Spresense
- Arduino Portenta H7 + Vision shield (preview)
- Raspberry Pi 4
- NVIDIA Jetson Nano
- Mobile phone
- Porting guide
- COMMUNITY BOARDS**
- Seeed Wio Terminal
- Agora Product Development Kit

**1. Connect the development board to your computer**

Use a micro-USB cable to connect the development board to your computer. Then press RESET twice to launch into the bootloader. The on-board LED should start pulsating to indicate this.

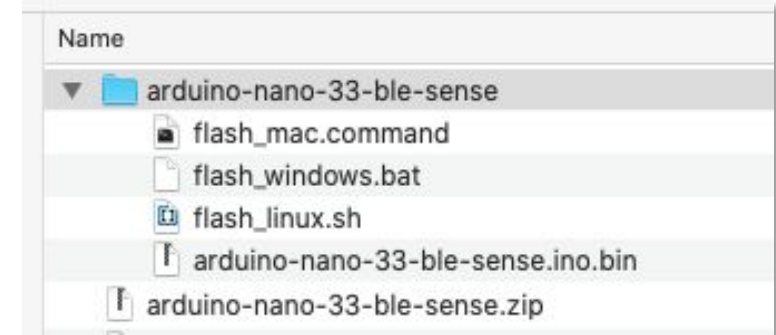
**1**

Press **RESET** twice quickly to launch the bootloader on the Arduino Nano 33 BLE Sense.

2. Update the firmware

The development board does not come with the right firmware yet. To update the firmware:

1. [Download the latest Edge Impulse firmware](#), and unzip the file.
2. Open the flash script for your operating system (`flash_windows.bat` , `flash_mac.command` or `flash_linux.sh`) to flash the firmware.
3. Wait until flashing is complete, and press the RESET button once to launch the new firmware.

3. Setting keys**2**

```

mjrovai — flash_mac.command — 124x43
Last login: Mon Jun 28 08:58:22 on ttys002
You have new mail.
/Users/mjrovai/Downloads/arduino-nano-33-ble-sense/flash_mac.command ; exit;

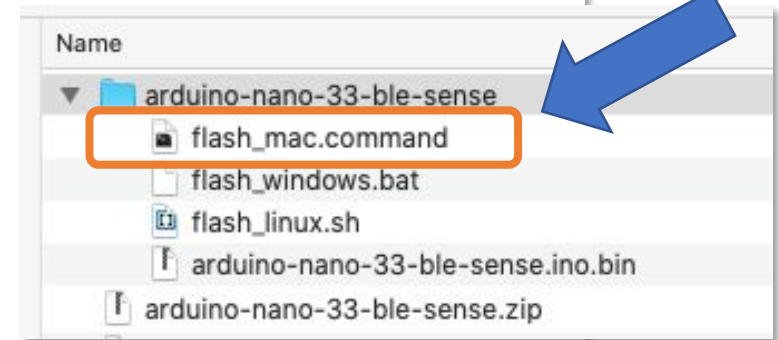
The default interactive shell is now zsh.
To update your account to use zsh, please run `chsh -s /bin/zsh`.
For more details, please visit https://support.apple.com/kb/HT208050.
(base) MacBook-Pro-de-Marcelo:~ mjrovai$ /Users/mjrovai/Downloads/arduino-nano-33-ble-sense/flash_mac.command ; exit;
Finding Arduino Mbed core...
Finding Arduino Mbed OK
Finding Arduino Nano 33 BLE...
Finding Arduino Nano 33 BLE OK
Flashing board...
Device       : nRF52840-QIAA
Version      : Arduino Bootloader (SAM-BA extended) 2.0 [Arduino:IKXYZ]
Address      : 0x0
Pages        : 256
Page Size    : 4096 bytes
Total Size   : 1024KB
Planes       : 1
Lock Regions : 0
Locked       : none
Security     : false
Erase flash

Done in 0.001 seconds
Write 525440 bytes to flash (129 pages)
[=====] 100% (129/129 pages)
Done in 20.533 seconds

Flashed your Arduino Nano 33 BLE development board.
To set up your development with Edge Impulse, run 'edge-impulse-daemon'
To run your impulse on your development board, run 'edge-impulse-run-impulse'
logout
Saving session...
...copying shared history...
...saving history...truncating history files...
...completed.

[Process completed]
```

1. Press Nano-33 Reset button Twice
2. With Nano-33 LED Flashing:



3. Nano-33 LED Stop Flashing

Windows 10

```
Microsoft Windows [versão 10.0.19041.1052]
(c) Microsoft Corporation. Todos os direitos reservados.

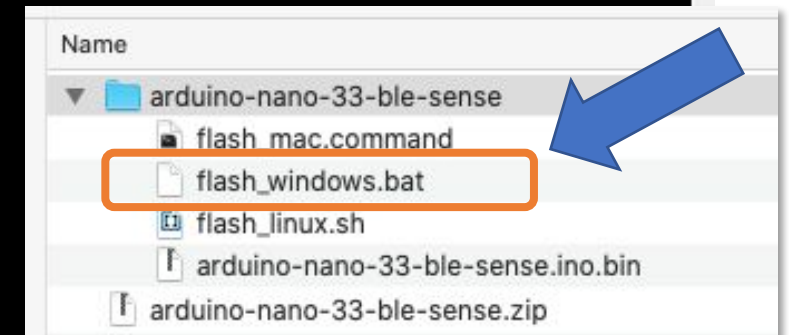
C:\Users\GUILH>arduino-cli
Arduino Command Line Interface (arduino-cli).

Usage:
  arduino-cli [command]

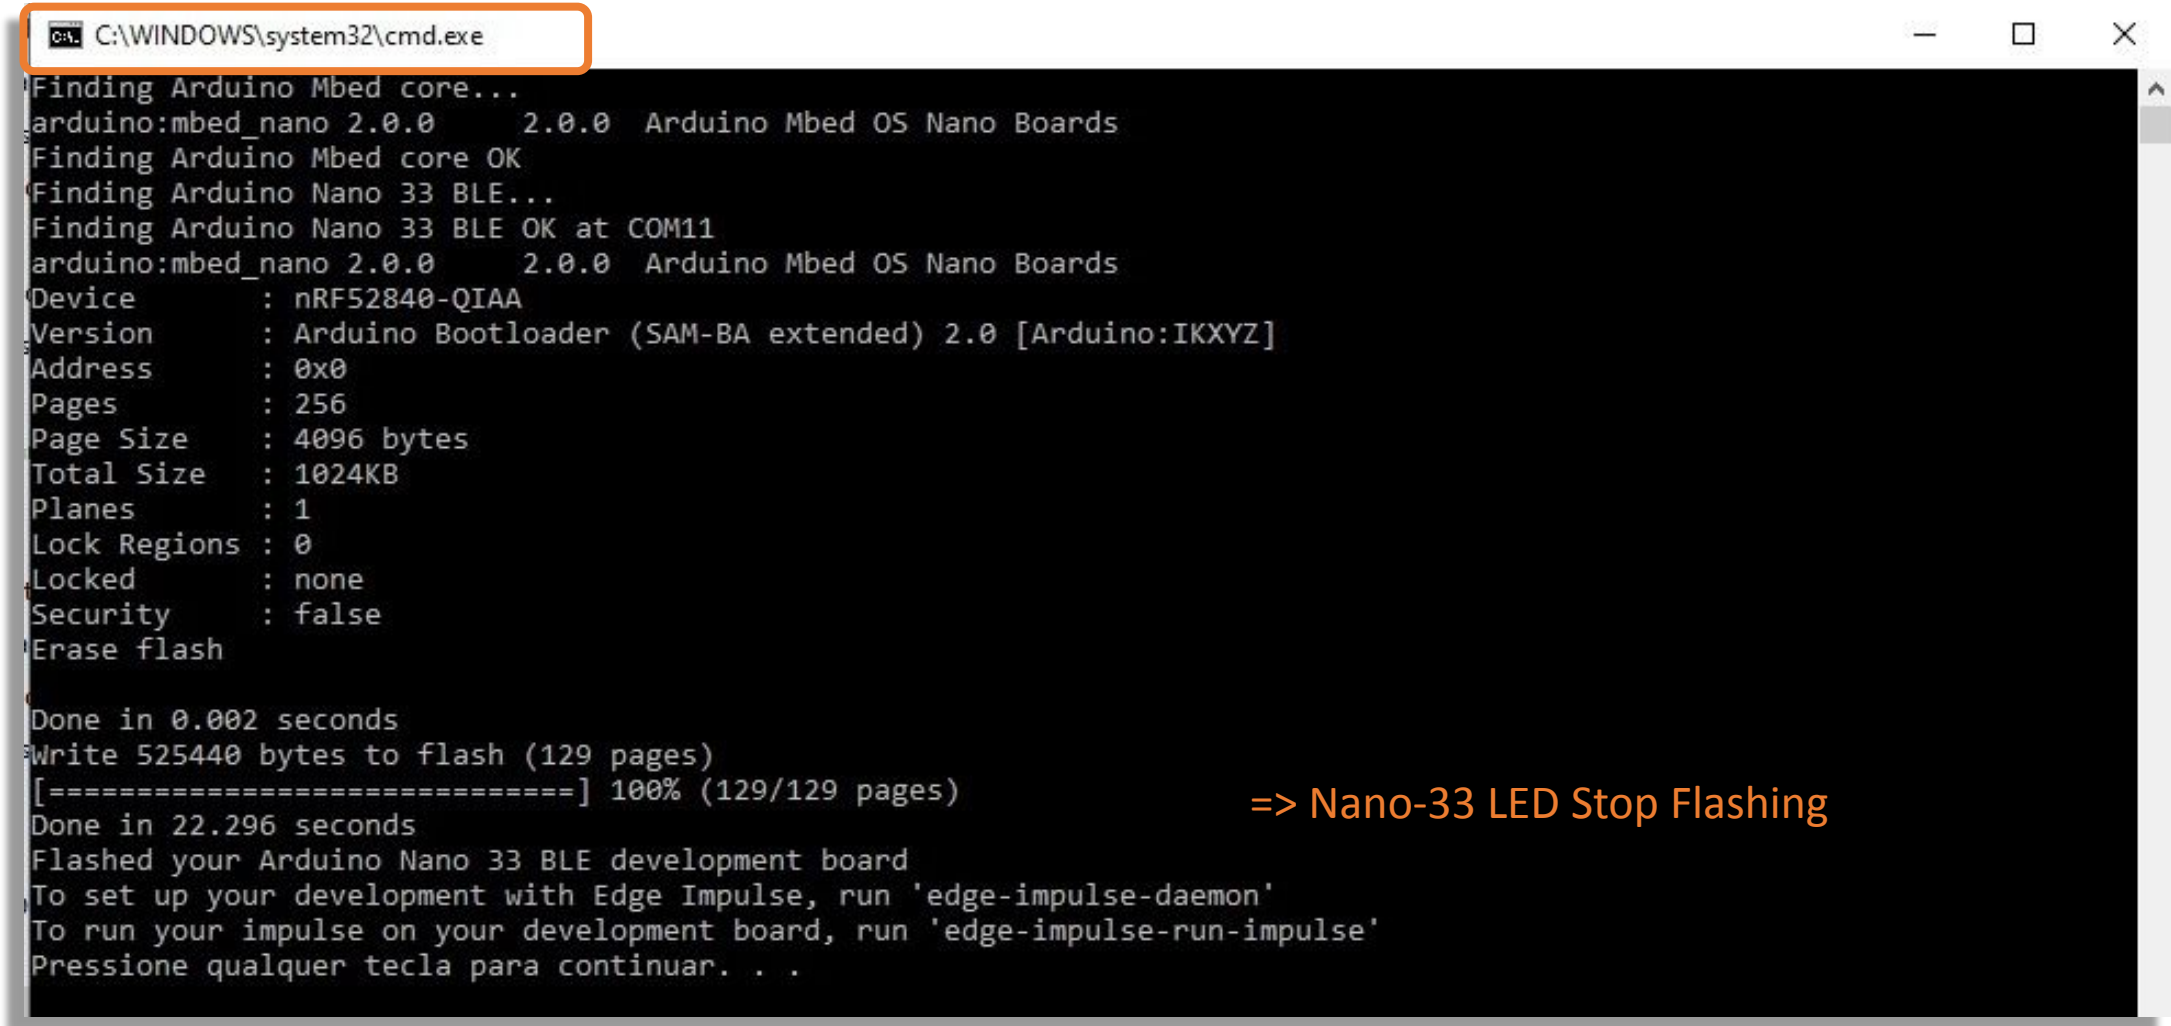
Examples:
  arduino-cli <command> [flags...]

Available Commands:
  board           Arduino board commands.
  burn-bootloader Upload the bootloader.
  cache          Arduino cache commands.
  compile         Compiles Arduino sketches.
  completion     Generates completion scripts
  config         Arduino configuration commands.
  core           Arduino core operations.
  daemon         Run as a daemon on port 50051
  debug          Debug Arduino sketches.
  help           Help about any command
  lib            Arduino commands about libraries.
  outdated       Lists cores and libraries that can be upgraded
  sketch         Arduino CLI sketch commands.
  update         Updates the index of cores and libraries
  upgrade        Upgrades installed cores and libraries.
  upload         Upload Arduino sketches.
  version        Shows version number of Arduino CLI.
```

1. Press Nano-33 Reset button Twice
2. With Nano-33 LED Flashing:



Windows 10



```
C:\WINDOWS\system32\cmd.exe
Finding Arduino Mbed core...
arduino:mbed_nano 2.0.0      2.0.0  Arduino Mbed OS Nano Boards
Finding Arduino Mbed core OK
Finding Arduino Nano 33 BLE...
Finding Arduino Nano 33 BLE OK at COM11
arduino:mbed_nano 2.0.0      2.0.0  Arduino Mbed OS Nano Boards
Device       : nRF52840-QIAA
Version      : Arduino Bootloader (SAM-BA extended) 2.0 [Arduino:IKXYZ]
Address      : 0x0
Pages       : 256
Page Size   : 4096 bytes
Total Size  : 1024KB
Planes     : 1
Lock Regions : 0
Locked      : none
Security    : false
Erase flash

Done in 0.002 seconds
Write 525440 bytes to flash (129 pages)
[=====] 100% (129/129 pages)
Done in 22.296 seconds
Flashed your Arduino Nano 33 BLE development board
To set up your development with Edge Impulse, run 'edge-impulse-daemon'
To run your impulse on your development board, run 'edge-impulse-run-impulse'
Pressione qualquer tecla para continuar. . .
```

=> Nano-33 LED Stop Flashing


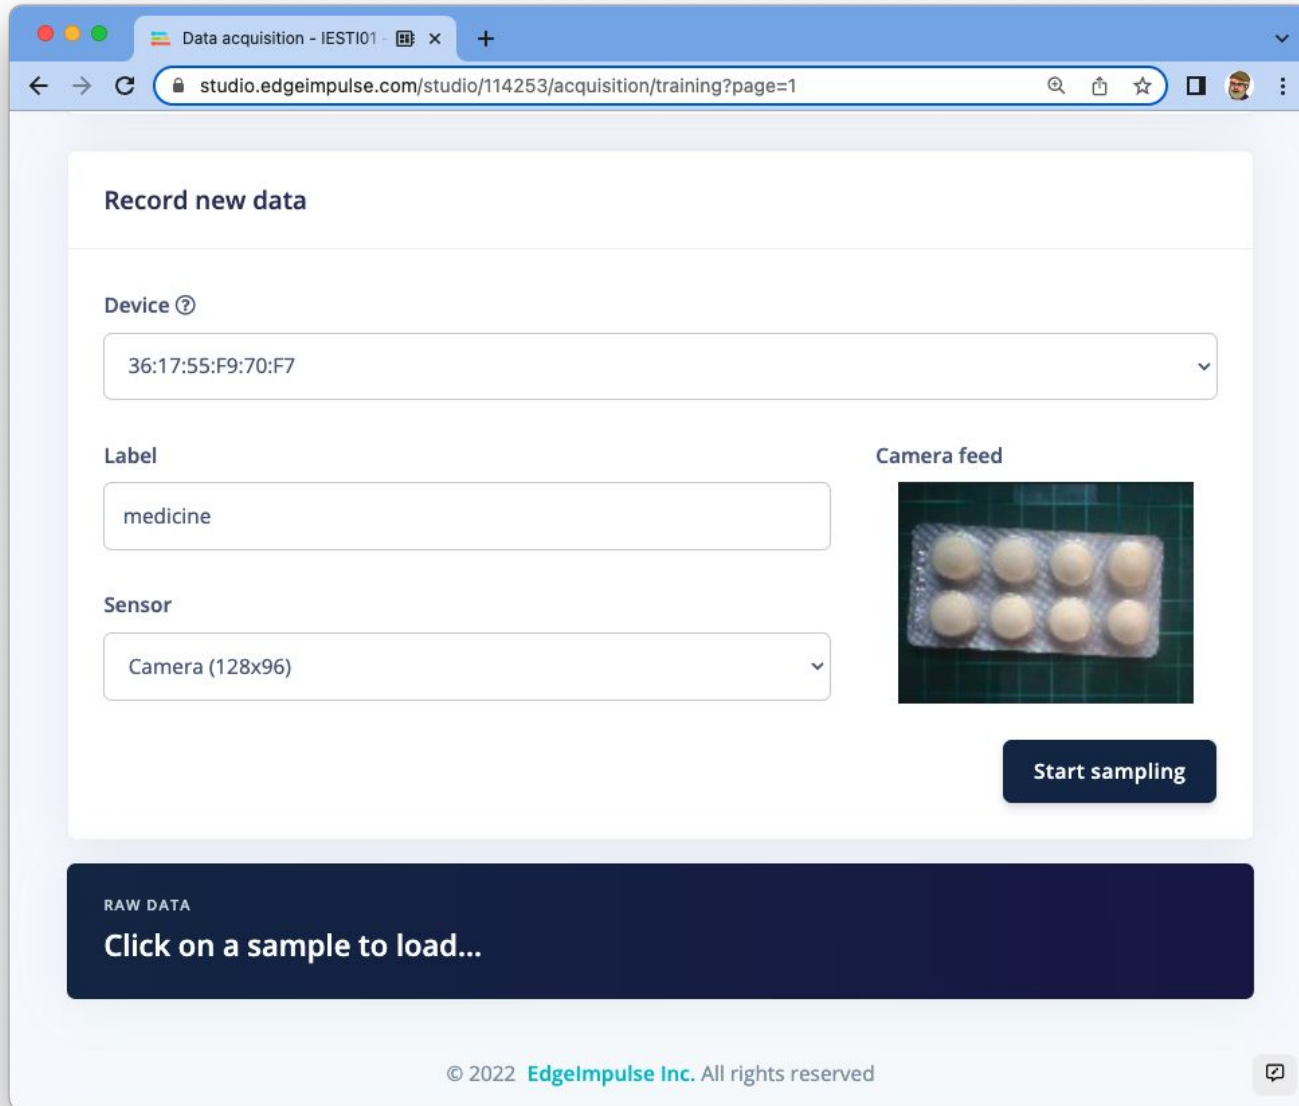
Follow this [video](#) to install the Arduino CLI on Windows 10 and add the 'arduino-cli' binary to your PATH. **This makes the CLI available from any location through the command prompt.**

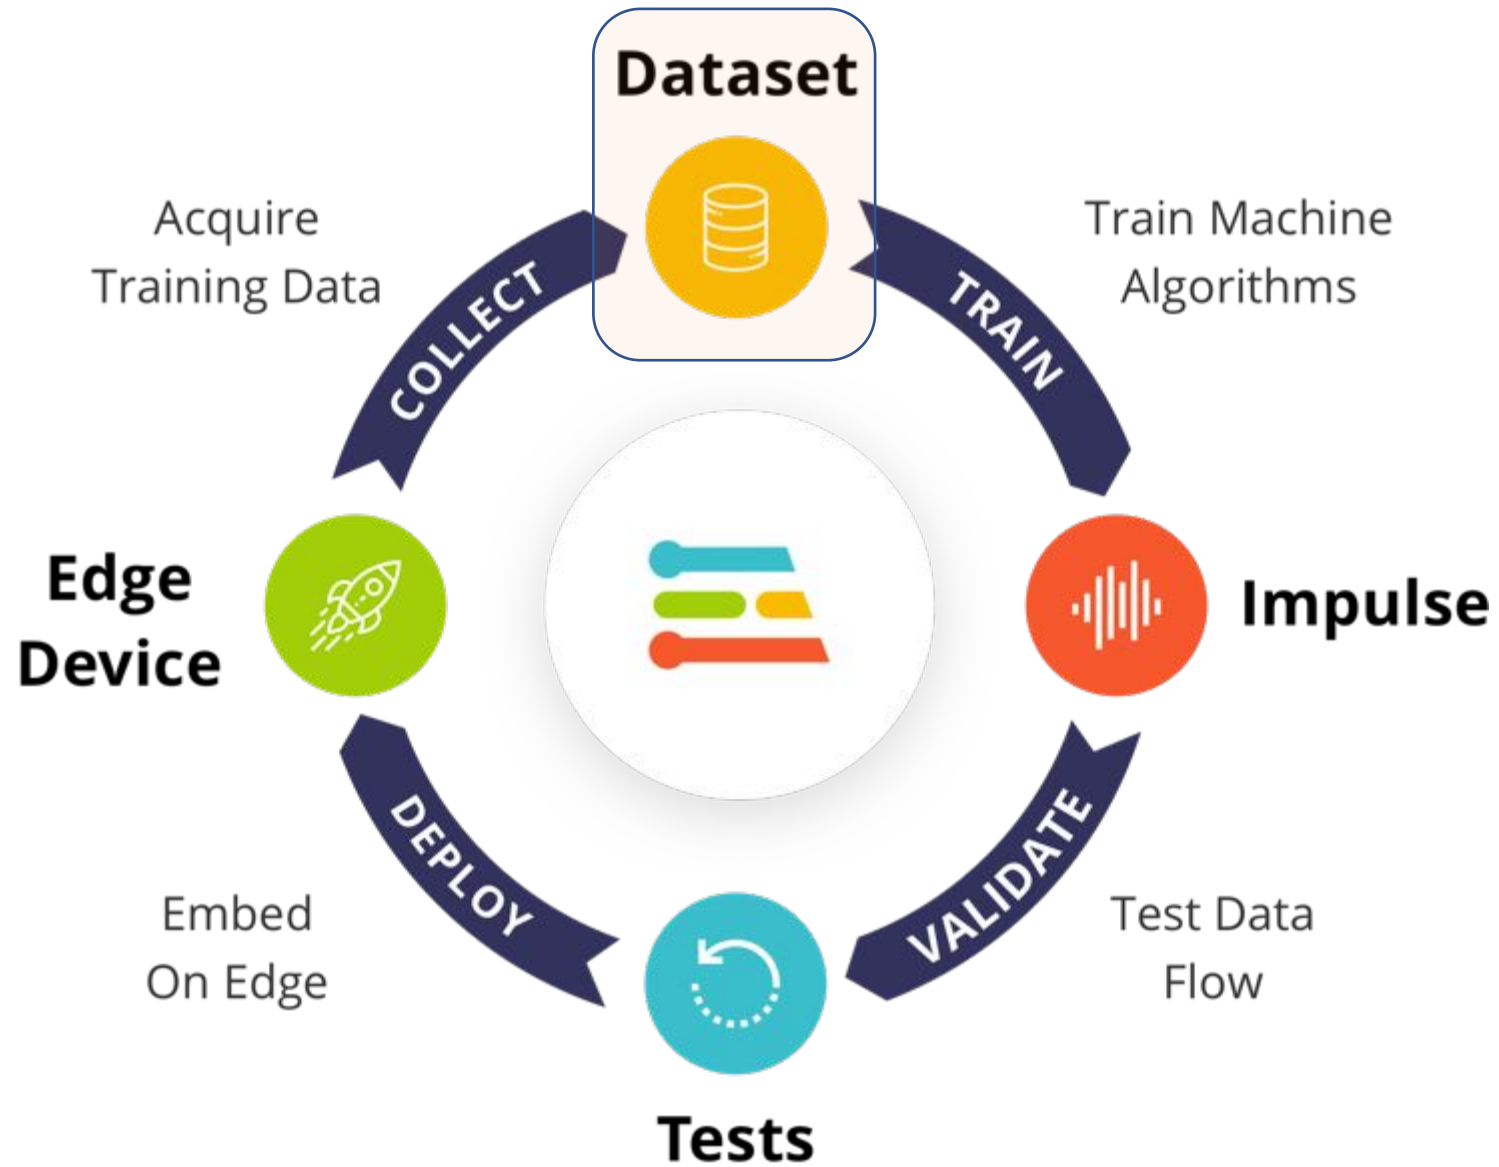
The screenshot shows the EdgeImpulse Studio web interface. A modal dialog box is open, titled "studio.edgeimpulse.com wants to connect to a serial port". It lists several serial ports: "cu.Bluetooth-Incoming-Port", "cu.MALS", "cu.RovaisAirPods-Wireless", "cu.SOC", and "Nano 33 BLE (cu.usbmodem144301) - Paired". The "Nano 33 BLE" port is selected and highlighted. Below the list are "Cancel" and "Connect" buttons, with the "Connect" button circled in orange. In the background, the main interface shows a "Record new data" section with a "Connect using WebUSB" button also circled in orange. A "RAW DATA" section below it says "Click on a sample to load...".

1


Do not forget to connect the Nano with your PC serial via Arduino-CLI, before starting collecting data.

Name
arduino-nano-33-ble-sense
flash_mac.command
flash_windows.bat
flash_linux.sh
arduino-nano-33-ble-sense.ino.bin
arduino-nano-33-ble-sense.zip





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



EDGE IMPULSE

DATA ACQUISITION (IESTI01 - IMAGE CLASSIFICATION)

Training data | Test data | Data explorer | Upload data | Export data

Did you know? You can capture data from any device or development board, or upload your existing datasets - [Show options](#)


DATA COLLECTED
68 items

TRAIN / TEST SPLIT
85% / 15%

Record new data [Connect using WebUSB](#)

No devices connected to the remote management API.

RAW DATA
Medicine.35vvi6da



SAMPLE NAME	LABEL	ADDED	LENGTH
Medicine.35vvi6da	medicine	Today, 10:33:34	-
Medicine.35vvh070	medicine	Today, 10:33:19	-
Medicine.35vvhep2	medicine	Today, 10:33:09	-
Medicine.35vvh3qf	medicine	Today, 10:32:58	-
Medicine.35vvtfs	medicine	Today, 10:32:52	-
Medicine.35vvg6j	medicine	Today, 10:32:45	-
Medicine.35vvgcbr	medicine	Today, 10:32:34	-
Medicine.35vvg2s0	medicine	Today, 10:32:24	-
Medicine.35vfvfith	medicine	Today, 10:32:08	-
Medicine.35vfv6ij	medicine	Today, 10:31:55	-
Medicine.35vver8u	medicine	Today, 10:31:44	-

Marcelo Rovai

Dashboard
Devices
Data sources
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums

Navigation: < 1 2 3 4 5 6 >

studio.edgeimpulse.com/studio/114253/acquisition/training?page=1

EDGE IMPULSE DATA ACQUISITION (IESTI01 - IMAGE CLASSIFICATION) Marcelo Rovai

Training data | Test data | Data explorer | Upload data | Export data

Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options

DATA COLLECTED: 68 items

TRAIN / TEST SPLIT: 85% / 15%

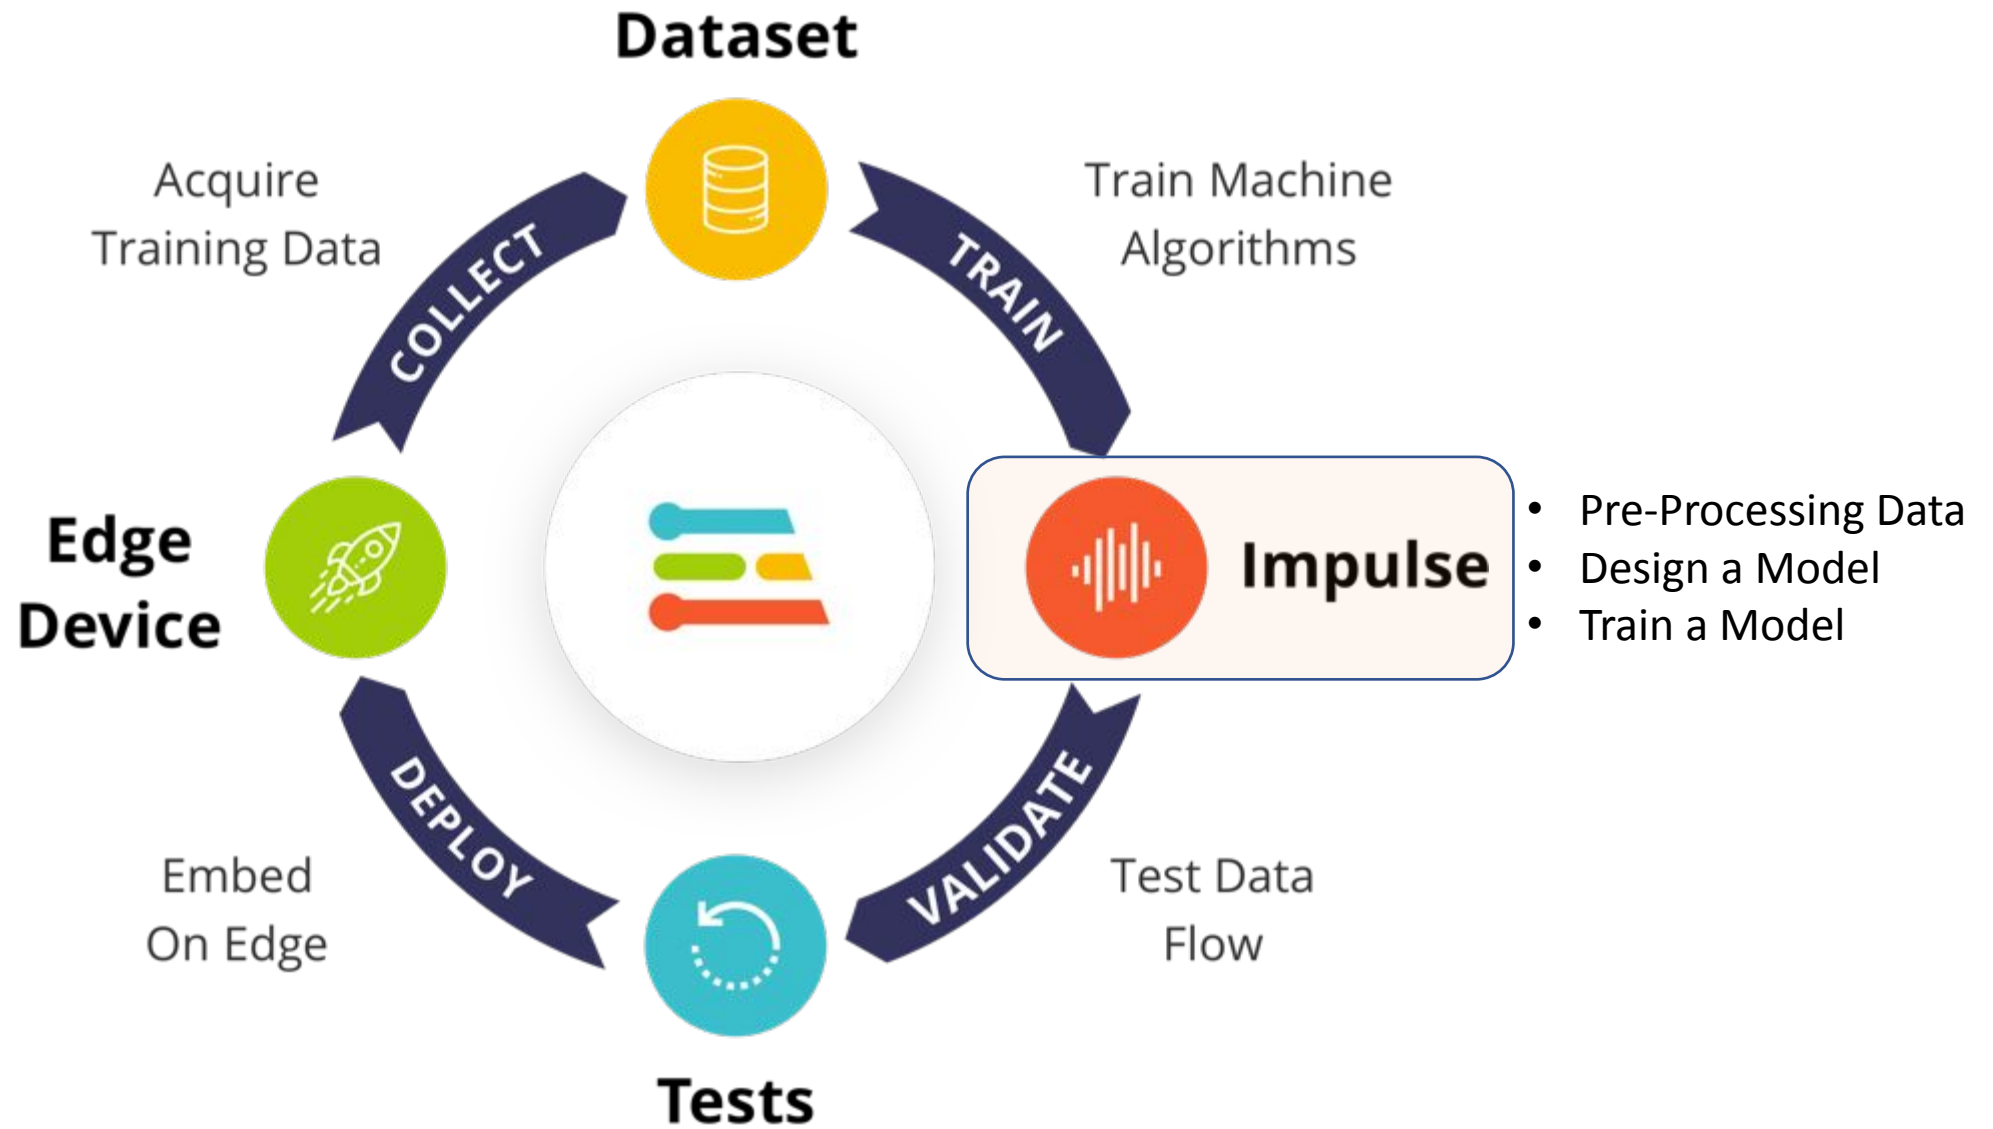
Record new data [Connect using WebUSB](#)

No devices connected to the remote management API.

RAW DATA: background.3600e161



SAMPLE NAME	LABEL	ADDED	LENGTH
background.3600e7fj	background	Today, 10:48:52	-
background.3600e161	background	Today, 10:48:46	-
background.3600dqt1	background	Today, 10:48:39	-
background.3600dj23	background	Today, 10:48:31	-
background.3600dco9	background	Today, 10:48:25	-
background.3600d7o5	background	Today, 10:48:20	-
background.3600d38u	background	Today, 10:48:15	-
background.3600cjr	background	Today, 10:47:59	-
background.3600cdb7	background	Today, 10:47:53	-
background.3600c5e5	background	Today, 10:47:45	-
background.3600bv6a	background	Today, 10:47:38	-
background.3600bnb4	background	Today, 10:47:30	-

Navigation: < 1 2 3 4 5 6 >



Data Pre-processing

- Image Pre-Process *:
 - Convert to Grayscale
 - Re-scale 96 x 96



- * During Inference, the OV7675 captures the raw image as:
 - QQVGA (160 x 120), crop it and resize it to 96x96
 - RGB565 and convert it first to RGB888 and after it to Grayscale
 - 1FPS

studio.edgeimpulse.com/studio/114253/create-impulse

EDGE IMPULSE

CREATE IMPULSE (IESTI01 - IMAGE CLASSIFICATION) 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

Input axes
image

Image width: 96 Image height: 96

Resize mode: Squash

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: 2 (background, medicine)

Output features

2 (background, medicine)

Save Impulse

Dashboard
Devices
Data sources
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums

© 2022 EdgeImpulse Inc. All rights reserved

Image - IEST101 - Image Classi x +

studio.edgeimpulse.com/studio/114253/dsp/image/3


EDGE IMPULSE


IMAGE (IEST101 - IMAGE CLASSIFICATION) Marcelo Rovai

#1 Click to set a description for this version

Parameters Generate features

Raw data background.3600e7fj (background)



Raw features 

0x7b8376, 0x798174, 0x798174, 0x7a8072, 0x7b8072, 0x7e8274, 0x7d8173...

Parameters

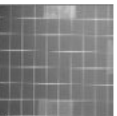
Image


Color depth Grayscale

Save parameters


DSP result

Image



Processed features 

0.4985, 0.4907, 0.4907, 0.4887, 0.4898, 0.4989, 0.4949, 0.4942, 0.50...

On-device performance 

PROCESSING TIME **11 ms.** PEAK RAM USAGE **4 KB**

Dashboard
Devices
Data sources
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums

Image - IEST101 - Image Classi x +

studio.edgeimpulse.com/studio/114253/dsp/image/3/generate-features

EDGE IMPULSE

IMAGE (IEST101 - IMAGE CLASSIFICATION)

#1 Click to set a description for this version

Parameters **Generate features**

Training set

Data in training set	68 items
Classes	2 (background, medicine)

Generate features

Feature explorer

● background
● medicine

On-device performance ?

PROCESSING TIME	PEAK RAM USAGE
11 ms.	4 KB

© 2022 EdgImpulse Inc. All rights reserved

- Dashboard
- Devices
- Data sources
- Data acquisition
- Impulse design
 - Create impulse
 - Image
 - Transfer learning
- EON Tuner
- Retrain model
- Live classification
- Model testing
- Versioning
- Deployment

GETTING STARTED

- Documentation
- Forums

Model Design

MobileNetV1 96x96 0.25

A pre-trained multi-layer convolutional network designed to efficiently classify images. Uses around 105.9K RAM and 301.6K ROM with default settings and optimizations. Works best with 96x96 input size. Supports both RGB and grayscale.

Model

Image Size

MobileNetV1 96x96 0.2 Alpha

Uses around 83.1K RAM and 218.3K ROM with default settings and optimizations. Works best with 96x96 input size. Supports both RGB and grayscale.

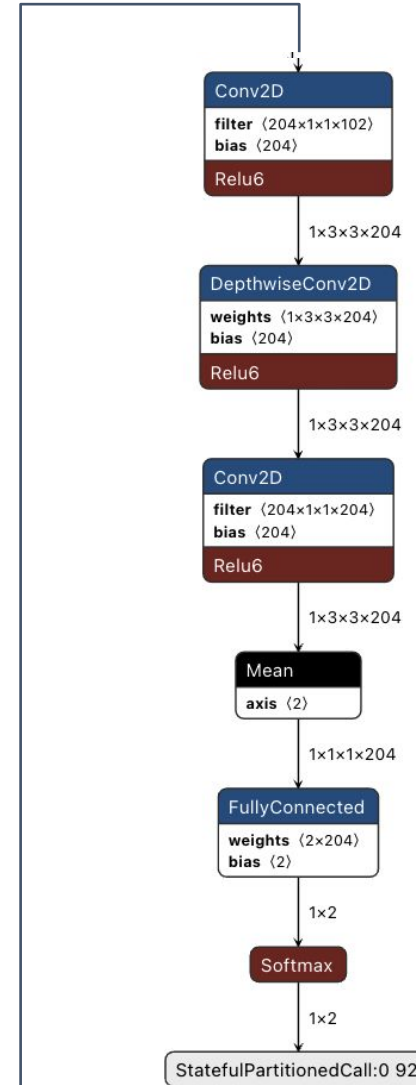


MobileNetV1 96x96 0.1

Uses around 53.2K RAM and 101K ROM with default settings and optimizations. Works best with 96x96 input size. Supports both RGB and grayscale.

Model Design

Input layer
(9,216 features
96 x 96)

MobileNetV1.0_2
.96x96
.grayscale
.bsize_96
.lr_0_05
.epoch_204
.val_loss_3.87
.val_accuracy_0.23



Output layer
(2 classes)

Train

The screenshot displays the Edge Impulse Studio interface for a transfer learning project. The main area is titled "TRANSFER LEARNING (IESTI01 - IMAGE CLASSIFICATION)".

Neural Network settings:

- Training settings:
 - Number of training cycles: 50
 - Learning rate: 0.0005
 - Validation set size: 20 %
 - Auto-balance dataset:
 - Data augmentation:
- Neural network architecture:
 - Input layer (9,216 features)
 - MobileNetV1 96x96 0.2 (no final dense layer, 0.1 dropout)
 - Choose a different model
 - Output layer (2 classes)

Training output: Very Good Accuracy

Model: Model version: Quantized (int8)

Last training performance (validation set):

- ACCURACY: 100.0%
- LOSS: 0.28

Confusion matrix (validation set):

	BACKGROUND	MEDICINE
BACKGROUND	100%	0%
MEDICINE	0%	100%
F1 SCORE	1.00	1.00

Data explorer (full training set):

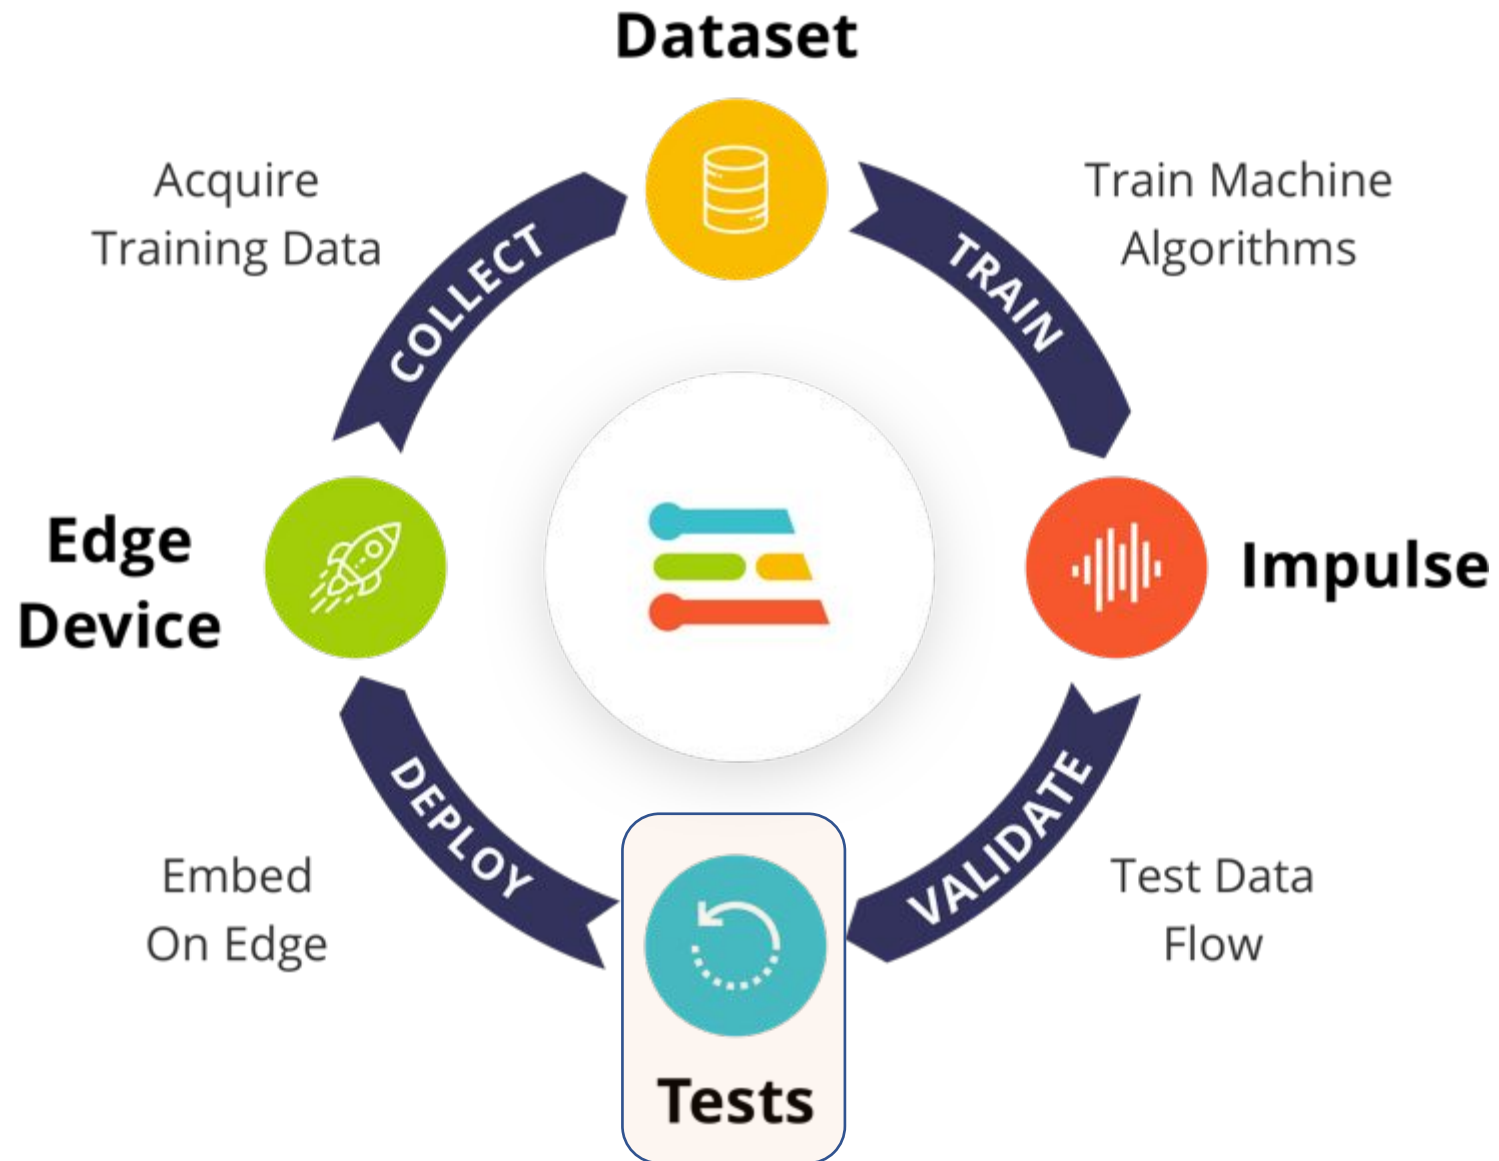
- background - correct (green dots)
- medicine - correct (red dots)
- medicine - incorrect (yellow dots)

On-device performance:

- INFERENCING: 3,592 ms.
- PEAK RAM USAGE: 106.2K
- FLASH USAGE: 225.3K

Annotations on the image:

- A green arrow points to the "Model" section.
- An orange box highlights the "Model version" dropdown and the "Last training performance" metrics.
- An orange box highlights the "MobileNetV1" model selection.
- A blue arrow points to the "Start training" button with the text "High Latency".
- A pink arrow points to the "Data explorer" scatter plot with the text "Memory OK".
- An orange box highlights the "On-device performance" metrics.



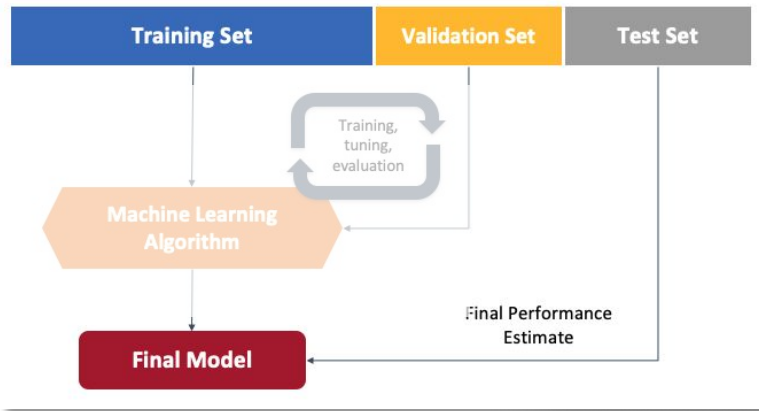
- Pre-Processing Data
- Design a Model
- Train a Model

Test

The screenshot displays the Edge Impulse Studio interface for model testing. The left sidebar contains navigation options: Dashboard, Devices, Data sources, Data acquisition, Impulse design, EON Tuner, Retrain model, Live classification, Model testing, Versioning, and Deployment. The main content area is titled 'Test data' and includes a 'Classify all' button. Below this is a table with columns for SAMPLE ID, EXPECTED OUTCOME, LENGTH, ACCURACY, and RESULT. The table shows 12 samples, all with 100% accuracy and correct classifications. The right sidebar shows 'Model testing output' with a log of training progress and a 'Model testing results' section. The results section displays an accuracy of 100.00% and a confusion matrix table.

	BACKGROUND	MEDICINE	UNCERTAIN
BACKGROUND	100%	0%	0%
MEDICINE	0%	100%	0%
F1 SCORE	1.00	1.00	

The 'Feature explorer' section shows a scatter plot with two data series: 'background - correct' (yellow dots) and 'medicine - correct' (green dots). The plot shows a clear separation between the two classes in the feature space.



Live classification - IESTIO

studio.edgeimpulse.com/studio/114253/classification#load-sample-10929...

Classification result

Summary

Name: testing.3601u5mo

Expected outcome: testing


CATEGORY	COUNT
background	1
medicine	0
uncertain	0

Detailed result Show only unknowns

BACKGROUND	MEDICINE
0.91	0.09

RAW DATA

testing.3601u5mo



Raw features [📄](#)

0x92917f, 0x92917f, 0x93907f, 0x969080, 0x969080, 0x979080, 0x958e7e, 0x948e80, 0x948f83, 0x949085, 0x939085, 0x8...

Live classification - IESTIO

studio.edgeimpulse.com/studio/114253/classification#load-sample-10929...

Classification result

Summary

Name: testing.36020g3v

Expected outcome: testing


CATEGORY	COUNT
background	0
medicine	1
uncertain	0

Detailed result Show only unknowns

BACKGROUND	MEDICINE
0.03	0.97


RAW DATA

testing.36020g3v



Raw features [📄](#)

0x1d2a23, 0x1d2a23, 0x1c2922, 0x1d2824, 0x1e2825, 0x1e2925, 0x1f2a26, 0x1d2926, 0x1e2a28, 0x1c2b28, 0x1b2a28, 0x1...



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

Deployment - IESTI01 - Image | x +

studio.edgeimpulse.com/studio/114253/deployment

Computer Mobile phone

Select optimizations (optional)

Model optimizations can increase on-device performance but may reduce accuracy. Click below to analyze optimizations and see the recommended choices for your target. Or, just click Build to use the currently selected options.

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

Available optimization

Quantized (int8)
Currently selected

Unoptimized (float)
Click to select


Estimate for Arduino Nano

Build output

```

Creating ... (ID: 3231036)
Writing ...
Writing ... OK
Schedul... in cluster...
Job ...
Copy... pulse SDK...
Copy... pulse SDK OK
  
```

ei-iesti01---imag....zip ^ Show All x



Built Arduino library

Add this library through the Arduino IDE via:
Sketch > Include Library > Add .ZIP Library...

Examples can be found under:
**File > Examples > IESTI01_--
Image_classification_inferencing**

Select a zip file or a folder containing the library you'd like to add

Downloads

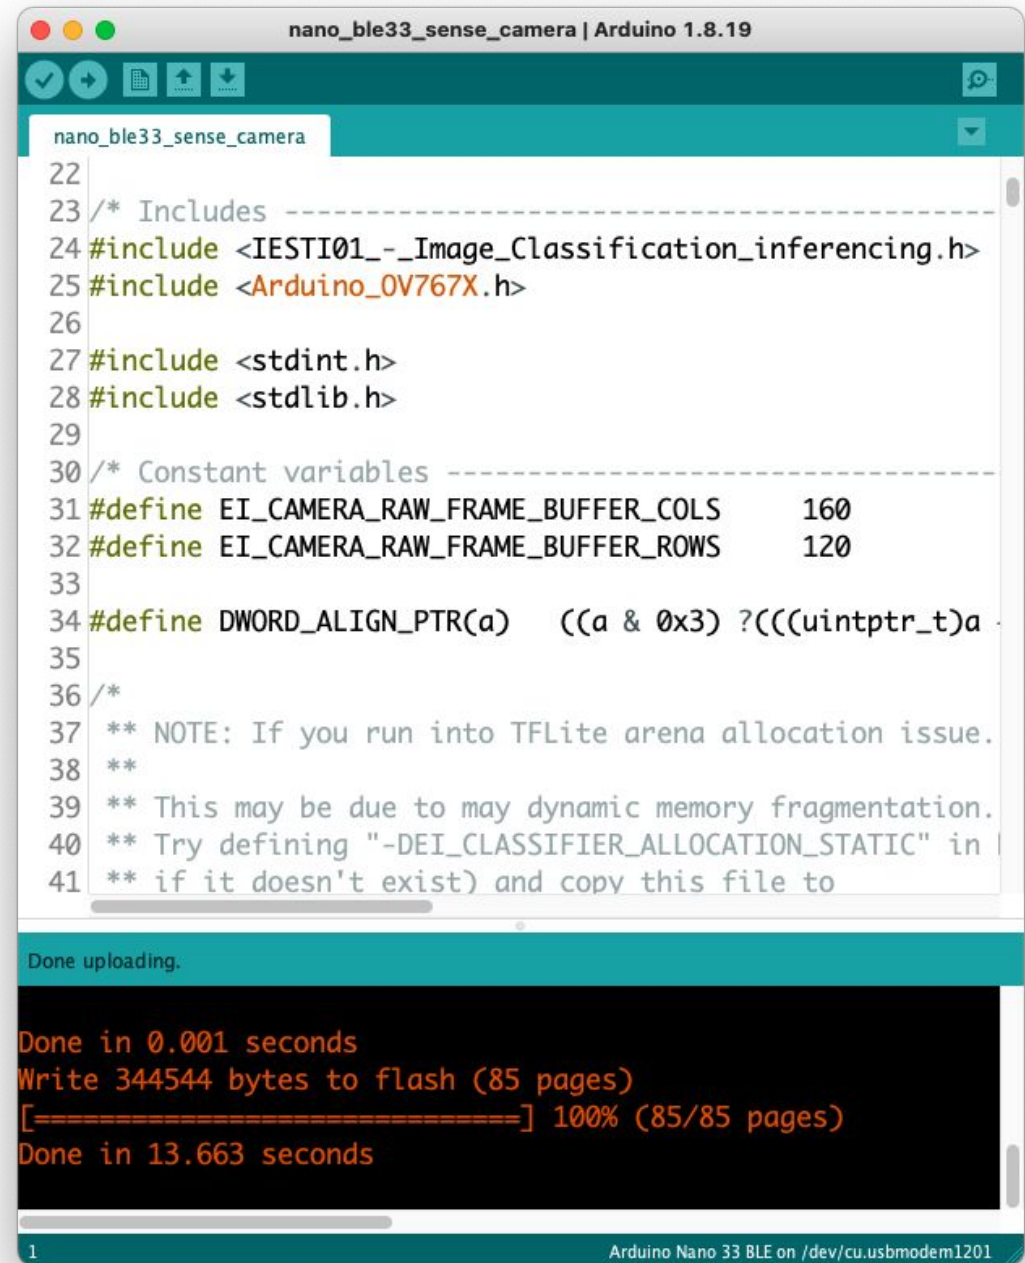
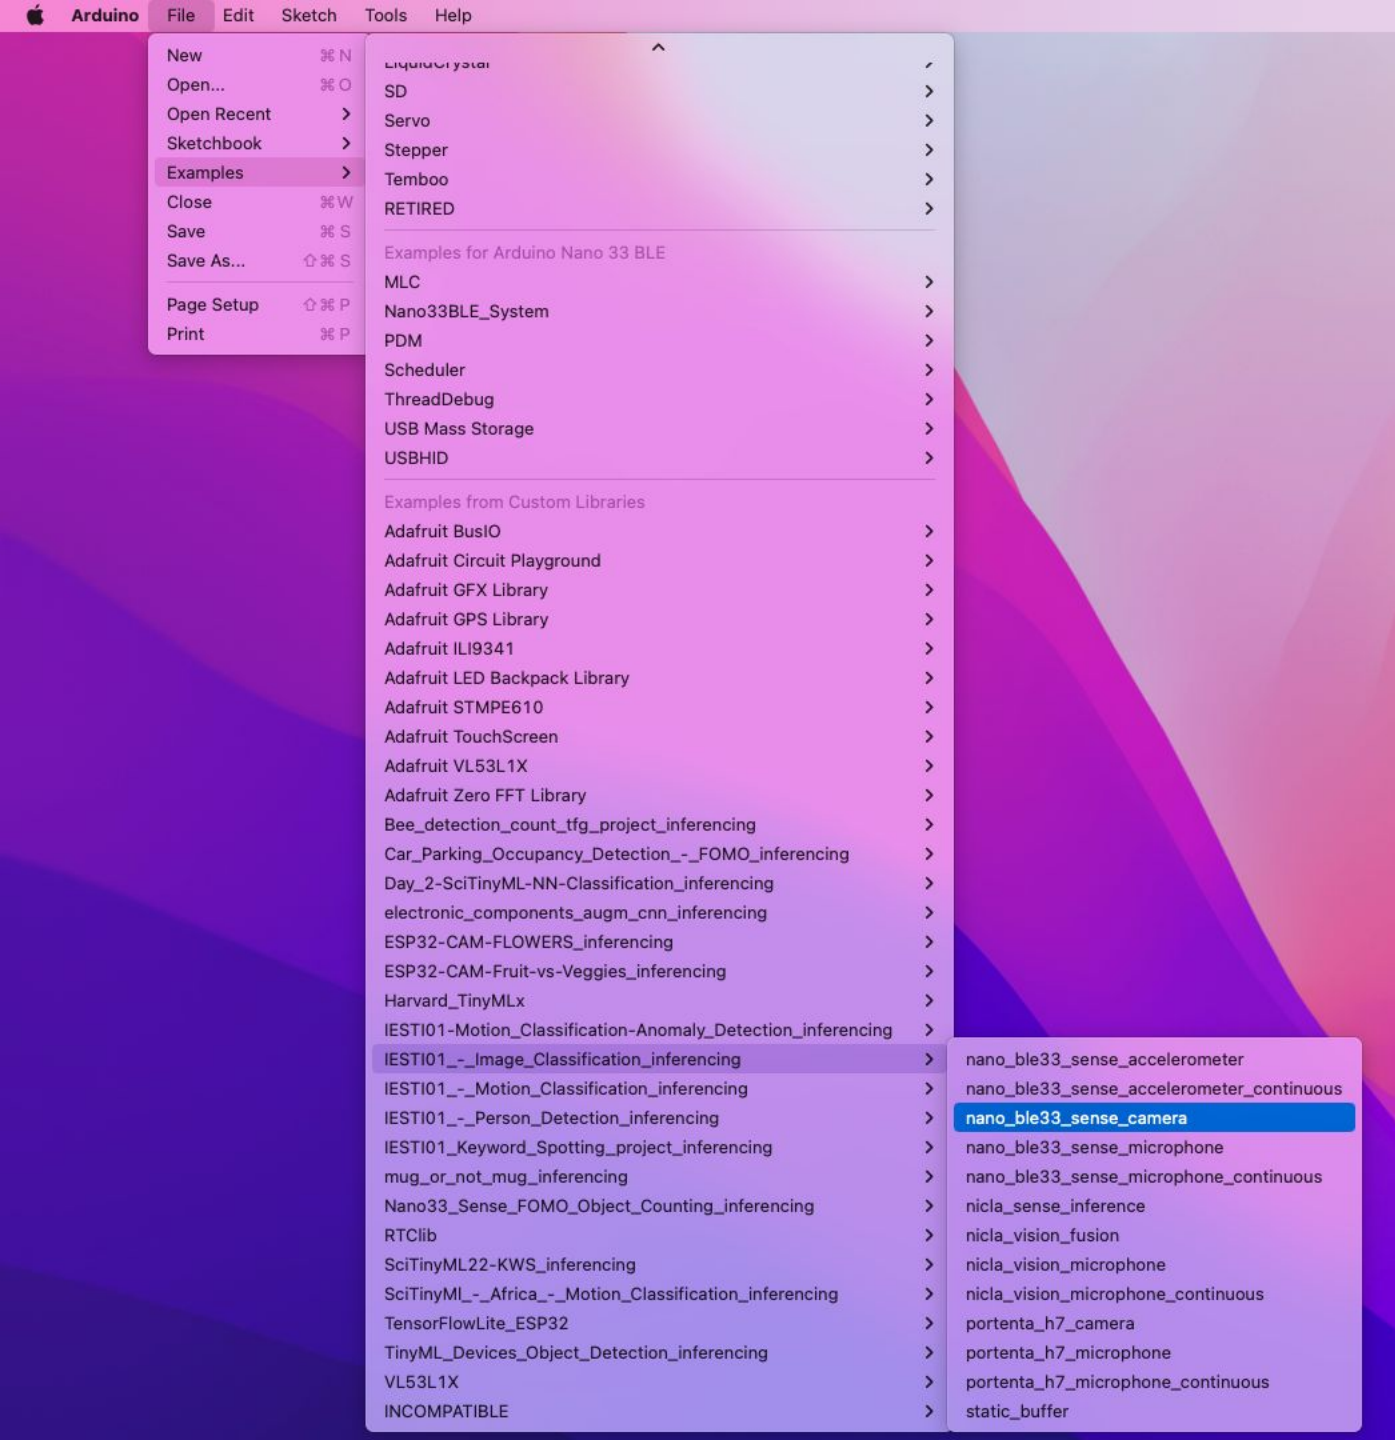
Name	Date Modified
ei-iesti01---image-classification-arduino-1.0.1.zip	Tuesday, June 21, 2022 11:17 AM
arduino-cli_0.23.0_macOS_64bit	Monday, June 20, 2022 4:29 PM
arduino-cli_0.23.0_macOS_64bit.tar.gz	Monday, June 20, 2022 4:29 PM
arduino-nano-33-ble-sense (8)	Monday, June 20, 2022 3:59 PM
arduino-nano-33-ble-sense (8).zip	Monday, June 20, 2022 3:59 PM
IESTI01_TinyML_class_25a.pptx (1).pdf	Monday, June 20, 2022 3:34 PM
IESTI01_TinyML_class_25a.pptx.pdf	Monday, June 20, 2022 3:23 PM
IESTI01_TinyML_class_25a (1).pptx	Monday, June 20, 2022 3:22 PM
IESTI01_TinyML_class_25a.pptx	Monday, June 20, 2022 3:18 PM
IESTI01_TinyML_class_25.pptx.pdf	Monday, June 20, 2022 3:07 PM
IESTI01_TinyML_class_25 (1).pptx	Monday, June 20, 2022 3:06 PM
IESTI01_TinyML_class_25.pptx	Monday, June 20, 2022 2:44 PM
Vittorio_Corbo_registro_2206.pdf	Monday, June 20, 2022 11:00 AM
IESTI01_TinyML_class_27.pptx	Monday, June 20, 2022 10:49 AM
IESTI01_TinyML_class_26a.pptx	Monday, June 20, 2022 10:49 AM
IESTI01_TinyML_class_26.pptx	Monday, June 20, 2022 10:07 AM
JEC757015_preprint.pdf	Sunday, June 19, 2022 12:22 PM
Deep-learning-diagram.png	Sunday, June 19, 2022 12:21 PM
recording (8).wav	Friday, June 17, 2022 5:42 PM
recording (7).wav	Friday, June 17, 2022 5:32 PM
recording (6).wav	Friday, June 17, 2022 5:27 PM
recording (5).wav	Friday, June 17, 2022 5:26 PM
recording (4).wav	Friday, June 17, 2022 5:26 PM

File Format: ZIP files or folders

Cancel Choose

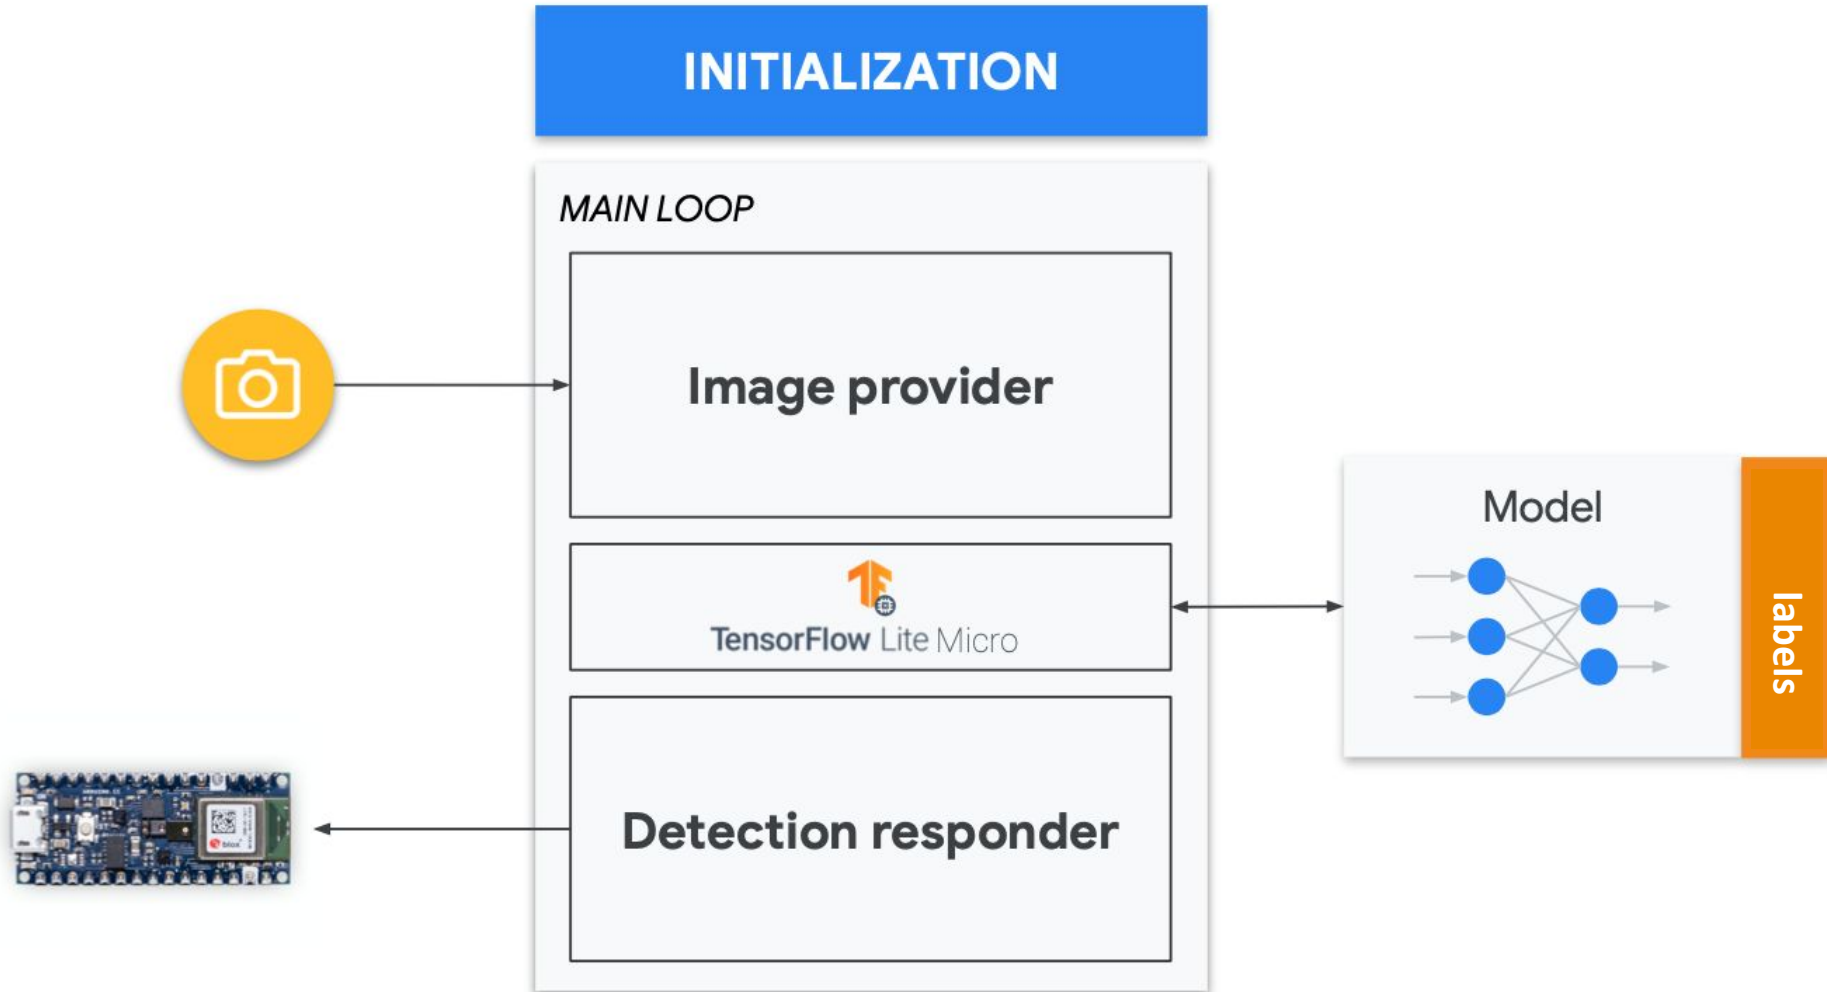


Model Inference

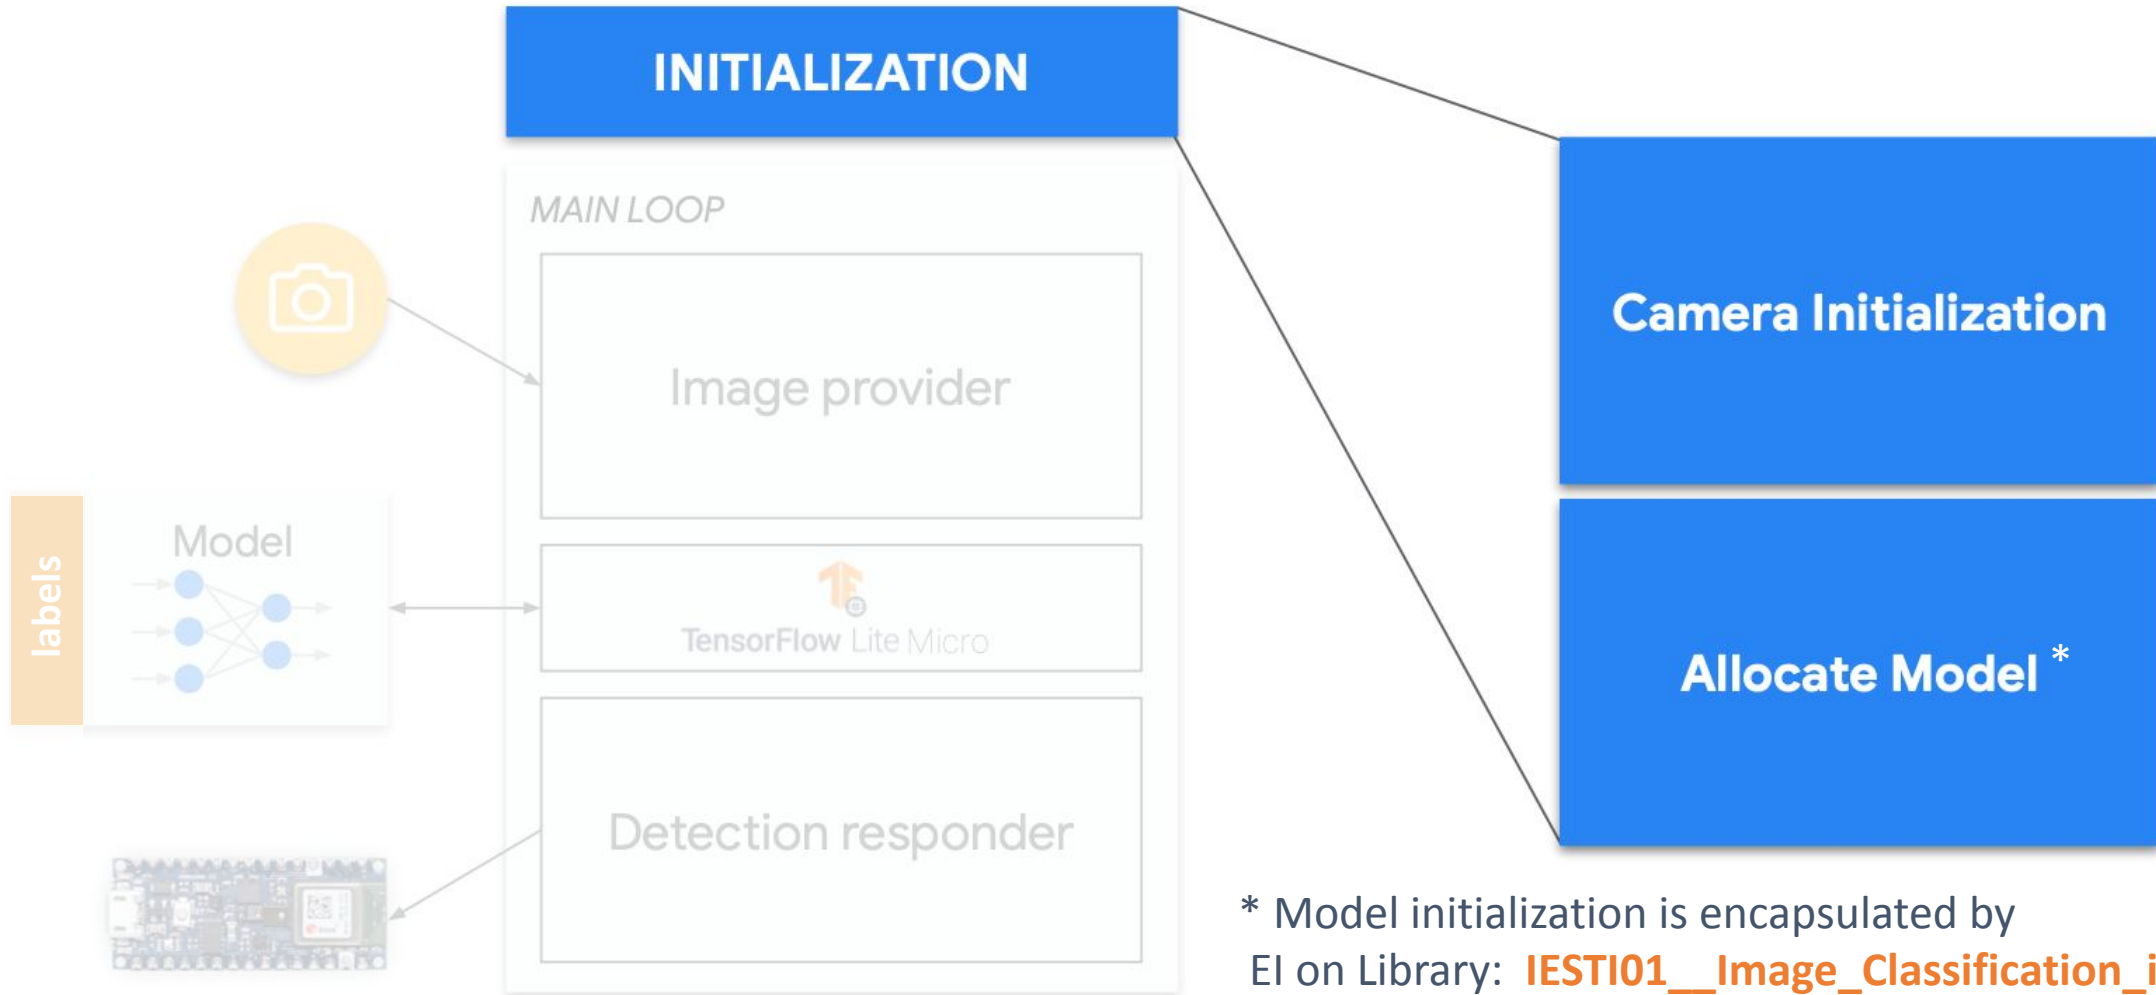


Img. Classification Components

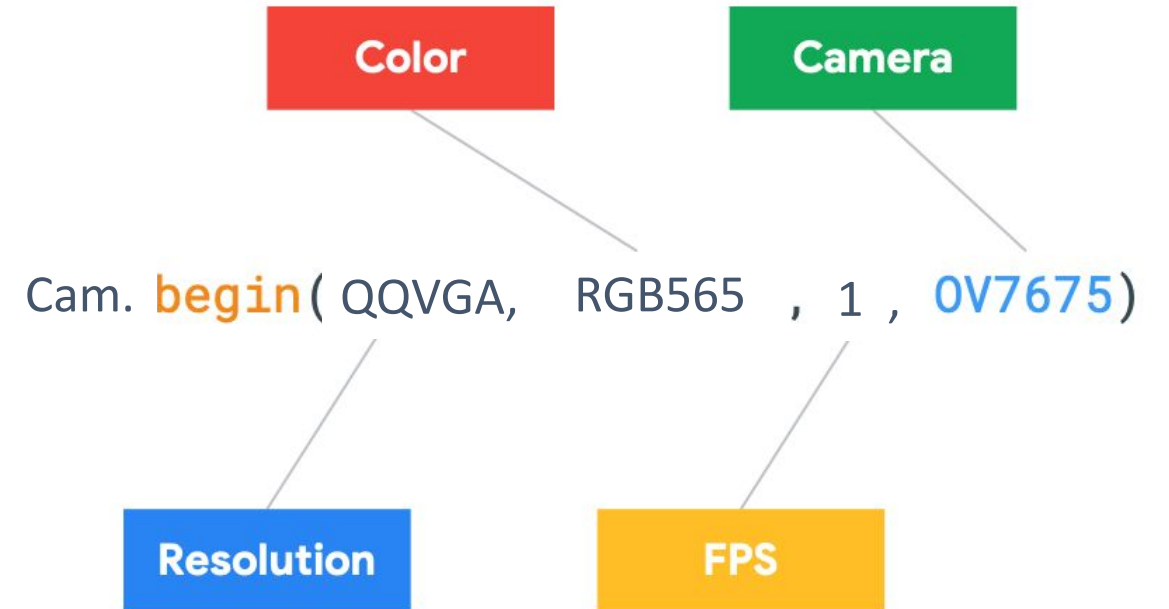
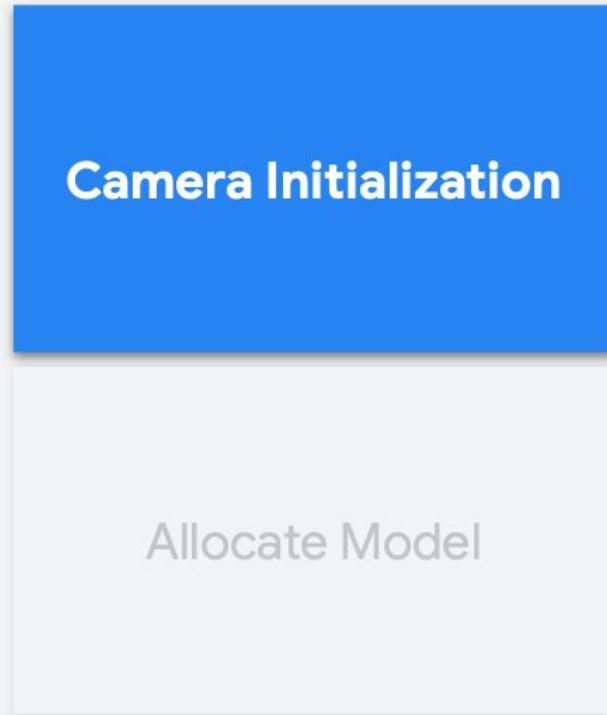
(nano_ble33_sense_camera.ino)



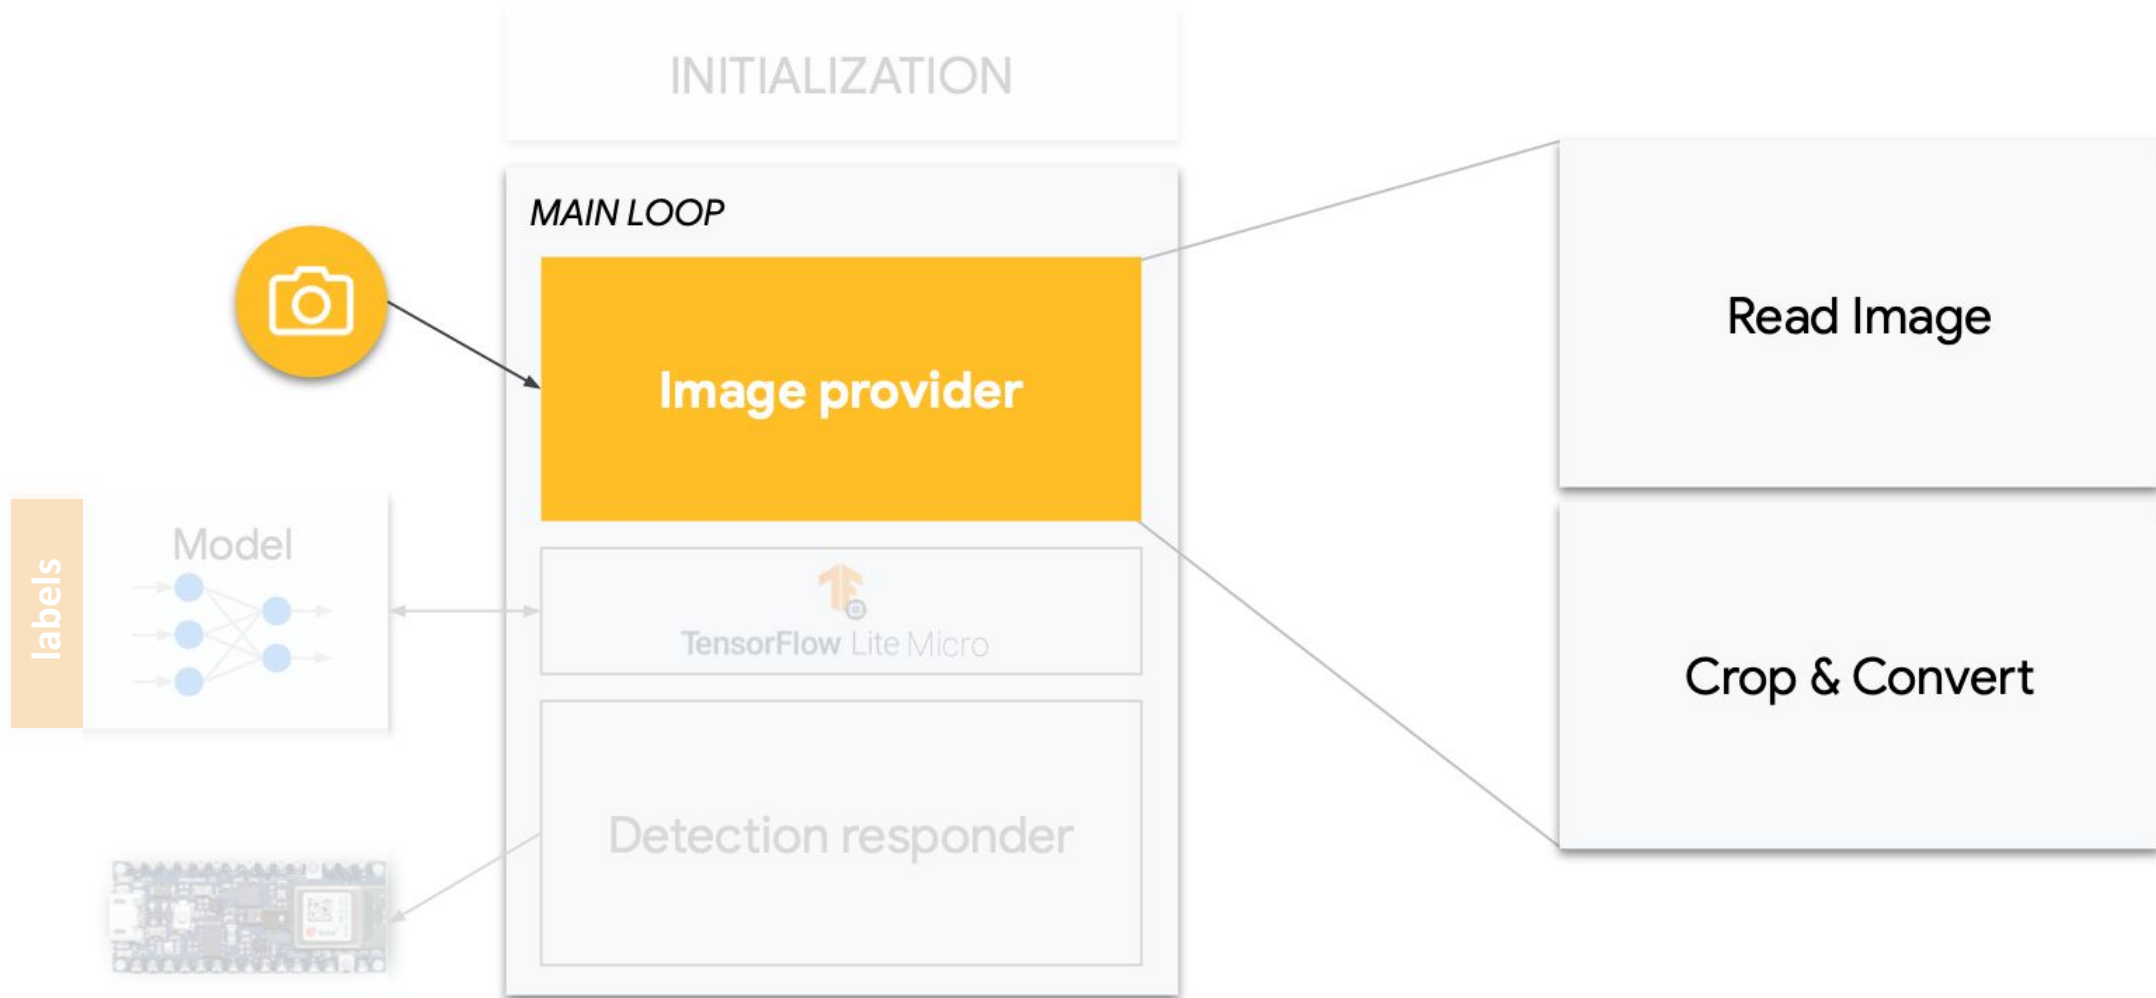
Initialization



Initialization



Pre-processing



Pre-processing

Read Image

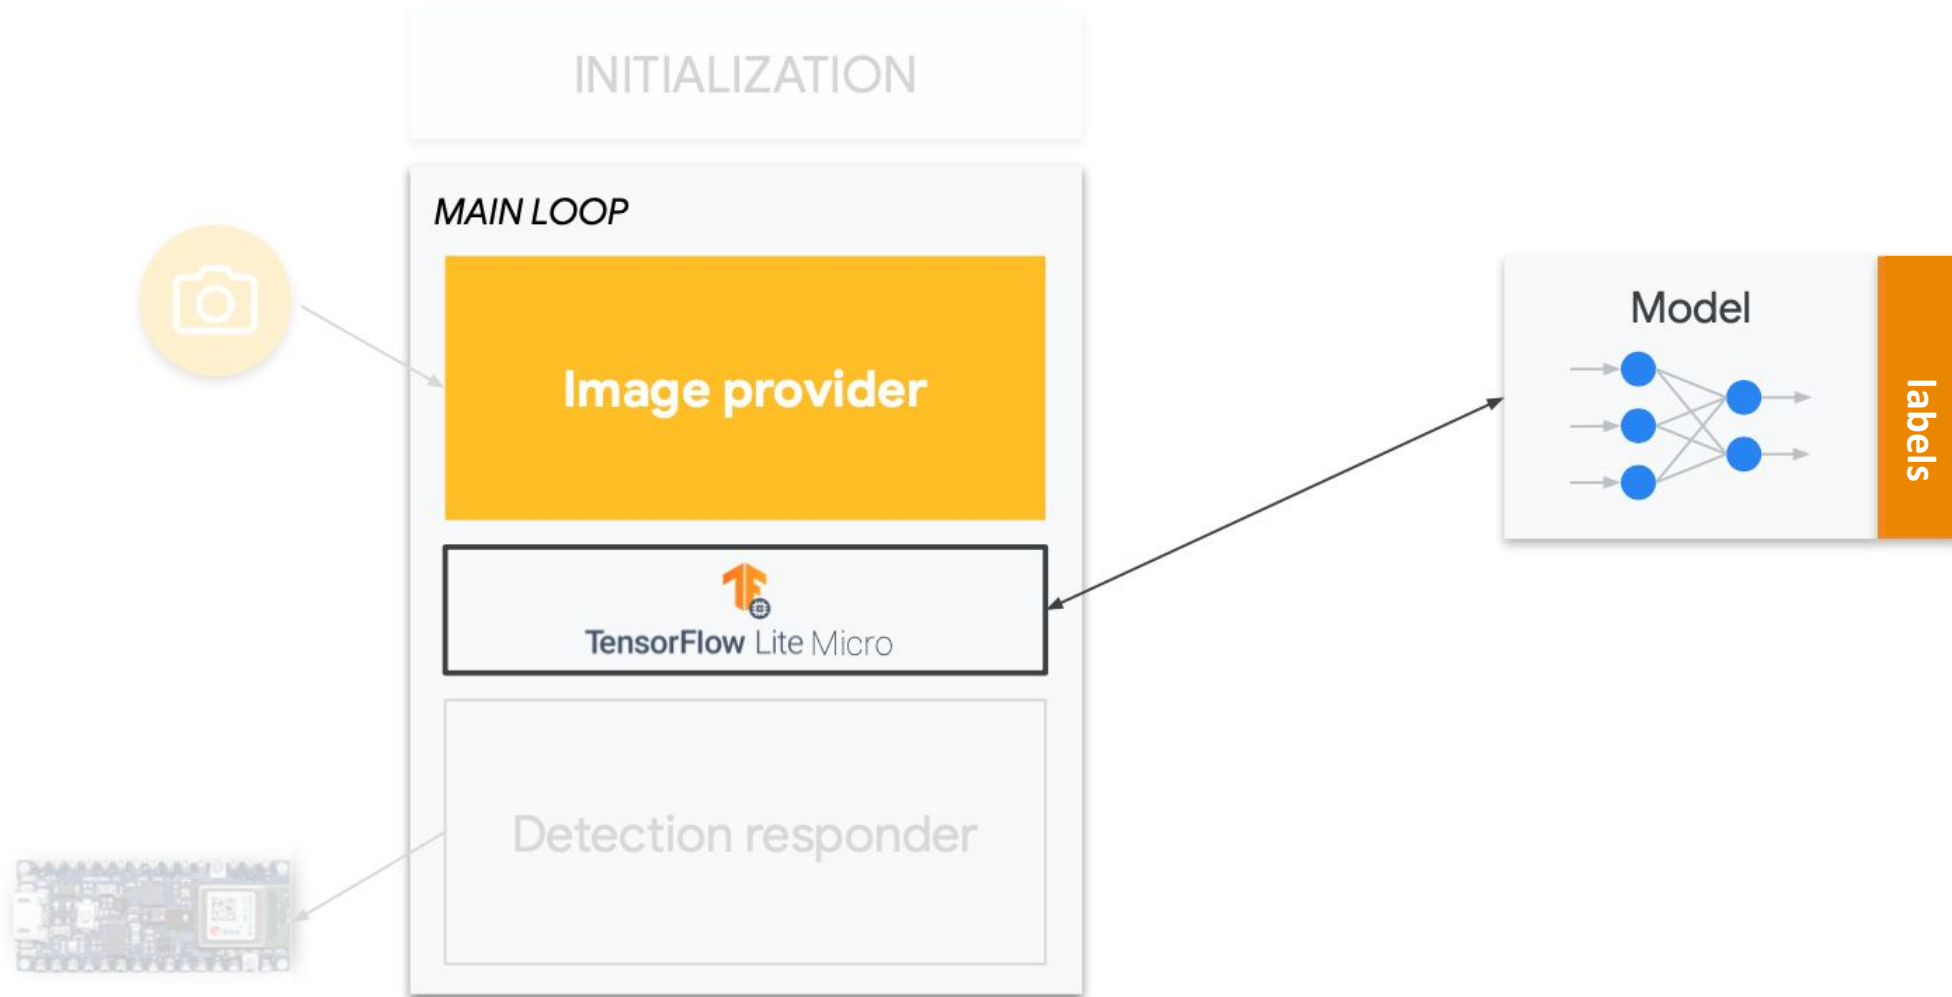
Crop & Convert

```
void *snapshot_mem = NULL;
uint8_t *snapshot_buf = NULL;
snapshot_mem = ei_malloc(resize_col_sz*resize_row_sz*2);
if(snapshot_mem == NULL) {
    ei_printf("failed to create snapshot_mem\r\n");
    break;
}
snapshot_buf = (uint8_t *)DWORD_ALIGN_PTR((uintptr_t)snapshot_mem);


if (ei_camera_capture(EI_CLASSIFIER_INPUT_WIDTH,
    EI_CLASSIFIER_INPUT_HEIGHT, snapshot_buf) == false) {
    ei_printf("Failed to capture image\r\n");
    if (snapshot_mem) ei_free(snapshot_mem);
    break;
}

ei::signal_t signal;
signal.total_length = EI_CLASSIFIER_INPUT_WIDTH * EI_CLASSIFIER_INPUT_HEIGHT;
signal.get_data = &ei_camera_cutout_get_data;
```

Interpreter + Model



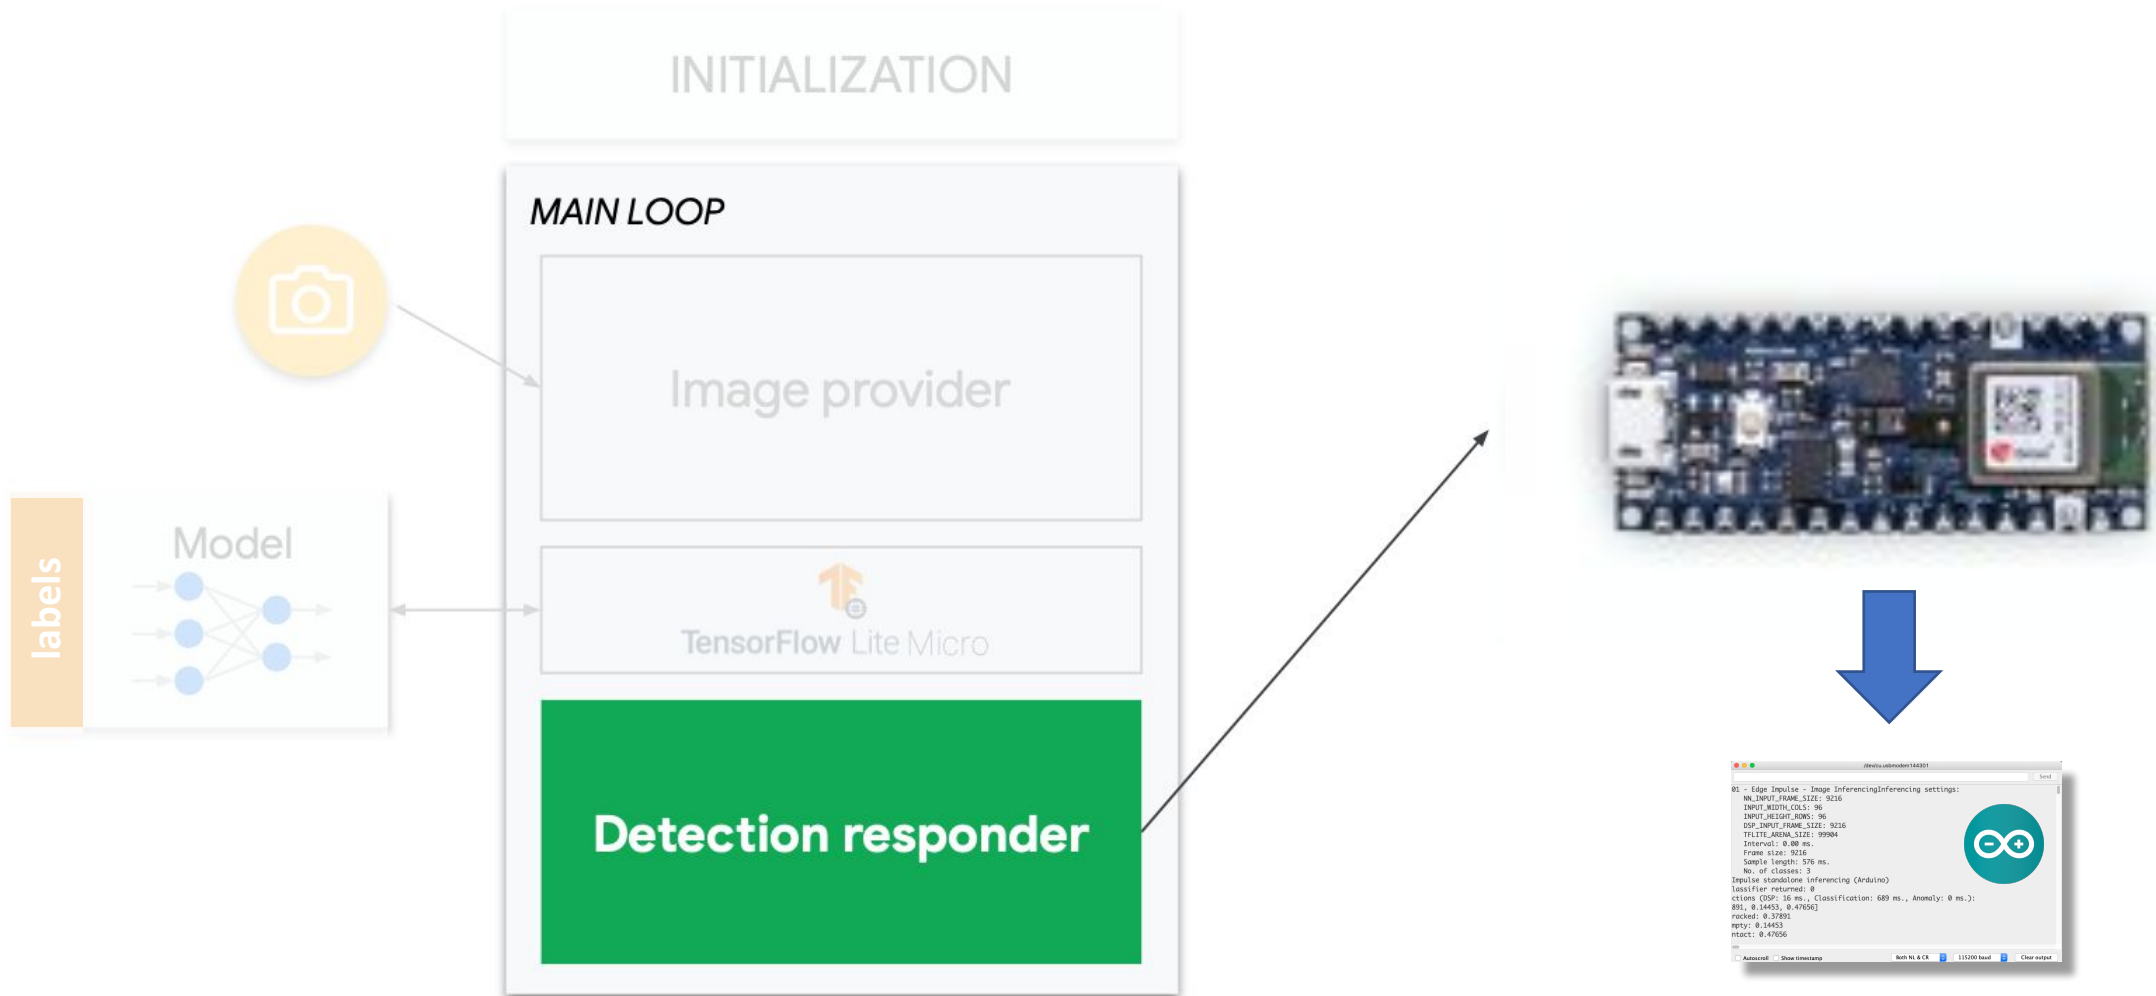
Interpreter + Model



```
// run the impulse: DSP, neural network and the Anomaly algorithm
ei_impulse_result_t result = { 0 };
```

```
EI_IMPULSE_ERROR ei_error = run_classifier(&signal, &result, debug_nn);
if (ei_error != EI_IMPULSE_OK) {
    ei_printf("Failed to run impulse (%d)\n", ei_error);
    ei_free(snapshot_mem);
    break;
}
```

Post-processing



Detection responder

```
// print the predictions
ei_printf("Predictions (DSP: %d ms., Classification: %d ms., Anomaly: %d ms.): \n",
         result.timing.dsp, result.timing.classification, result.timing.anomaly);
#if EI_CLASSIFIER_OBJECT_DETECTION == 1
bool bb_found = result.bounding_boxes[0].value > 0;
for (size_t ix = 0; ix < EI_CLASSIFIER_OBJECT_DETECTION_COUNT; ix++) {
    auto bb = result.bounding_boxes[ix];
    if (bb.value == 0) {
        continue;
    }

    ei_printf("    %s (%f) [ x: %u, y: %u, width: %u, height: %u ]\n",
             bb.label, bb.value, bb.x, bb.y, bb.width, bb.height);
}

if (!bb_found) {
    ei_printf("    No objects found\n");
}
#else
for (size_t ix = 0; ix < EI_CLASSIFIER_LABEL_COUNT; ix++) {
    ei_printf("    %s: %.5f\n", result.classification[ix].label,
             result.classification[ix].value);
}
#endif
```

```

/dev/cu.usbmodem1201
Send
Starting inferencing in 2 seconds...
Taking photo...
Predictions (DSP: 149 ms., Classification: 677 ms., Anomaly: 0 ms.):
  background: 0.10938
  medicine: 0.89062

Starting inferencing in 2 seconds...
Taking photo...
Predictions (DSP: 149 ms., Classification: 677 ms., Anomaly: 0 ms.):
  background: 0.10547
  medicine: 0.89453

Starting inferencing in 2 seconds...
Taking photo...

 Autoscroll  Show timestamp
Both NL & CR 115200 baud Clear

```




```

/dev/cu.usbmodem1201


Starting inferencing in 2 seconds...
Taking photo...
Predictions (DSP: 149 ms., Classification: 678 ms., Anomaly: 0 ms.):
  background: 0.69922
  medicine: 0.30078

Starting inferencing in 2 seconds...
Taking photo...
Predictions (DSP: 149 ms., Classification: 678 ms., Anomaly: 0 ms.):
  background: 0.71484
  medicine: 0.28516

Starting inferencing in 2 seconds...

 Autoscroll  Show timestamp
Both NL & CR 115200 baud

```



Main References:

- UNIFEI - IESTI01 TinyML - Machine Learning for Embedding Devices
- 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

I want to thank [Shawn Hymel](#) and Edge Impulse, [Pete Warden](#) and [Laurence Moroney](#) from Google, Professor [Vijay Janapa Reddi](#) and [Brian Plancher](#) from Harvard, and the rest of the [TinyMLedu](#) team for preparing the excellent material on TinyML that is the basis of this Workshop.

The SciTinyML workshops are part of the [TinyML4D](#), an initiative to make TinyML education available to everyone globally.

Thanks



UNIFEI