Data Pre-Processing for Hands-on Keyword Spotting

Brian Plancher
Harvard John A. Paulson School of Engineering and Applied Sciences
Barnard College, Columbia University
[*brianplancher.com*](https://brianplancher.com)

Quick Disclaimer:

Today will be **both** **too fast**

and **too slow**!

Quick Disclaimer:

Today will be **both too fast**

and **too slow**!

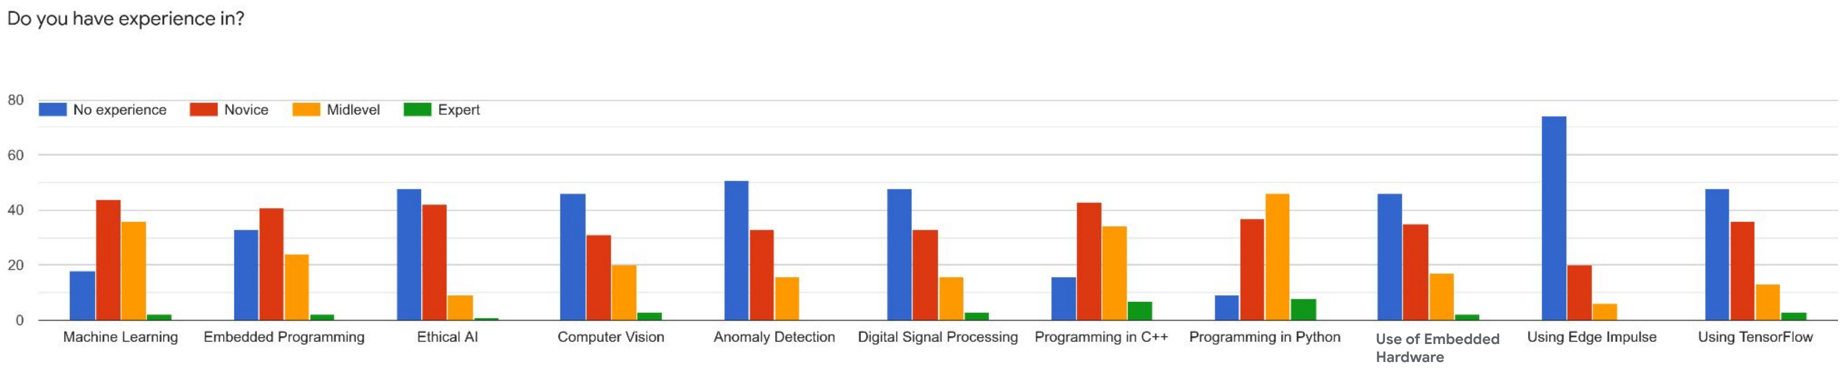

Feedback from yesterday:

The depth of material covered today was

15 responses

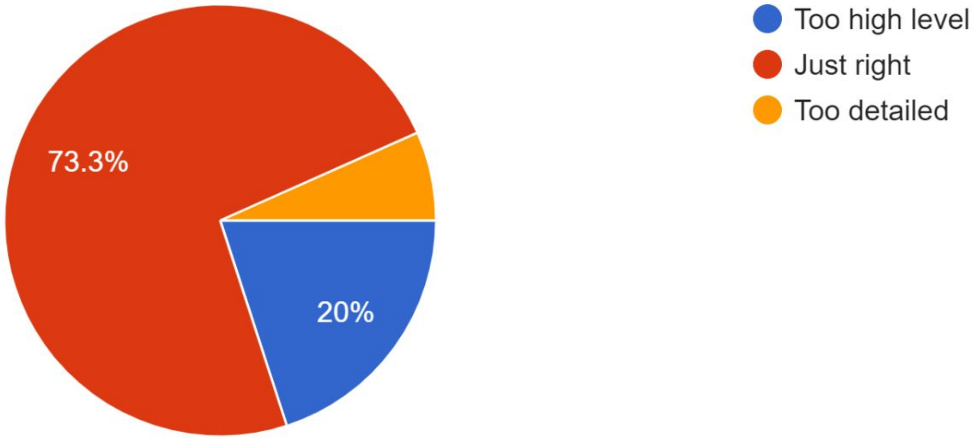

Feedback from yesterday:

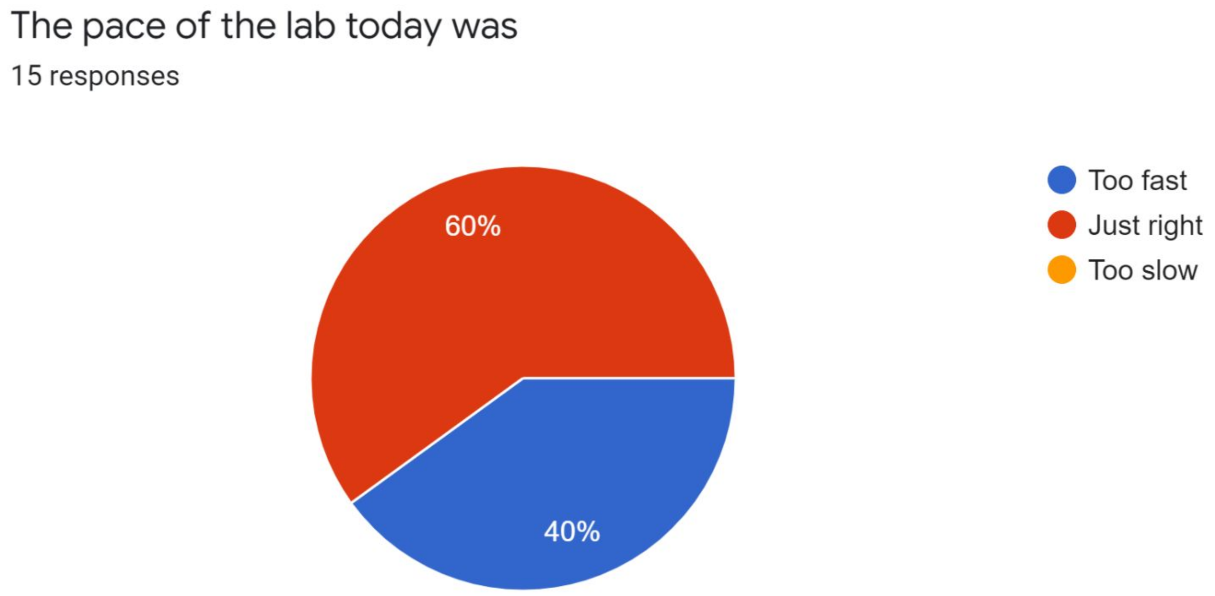

Feedback from yesterday:

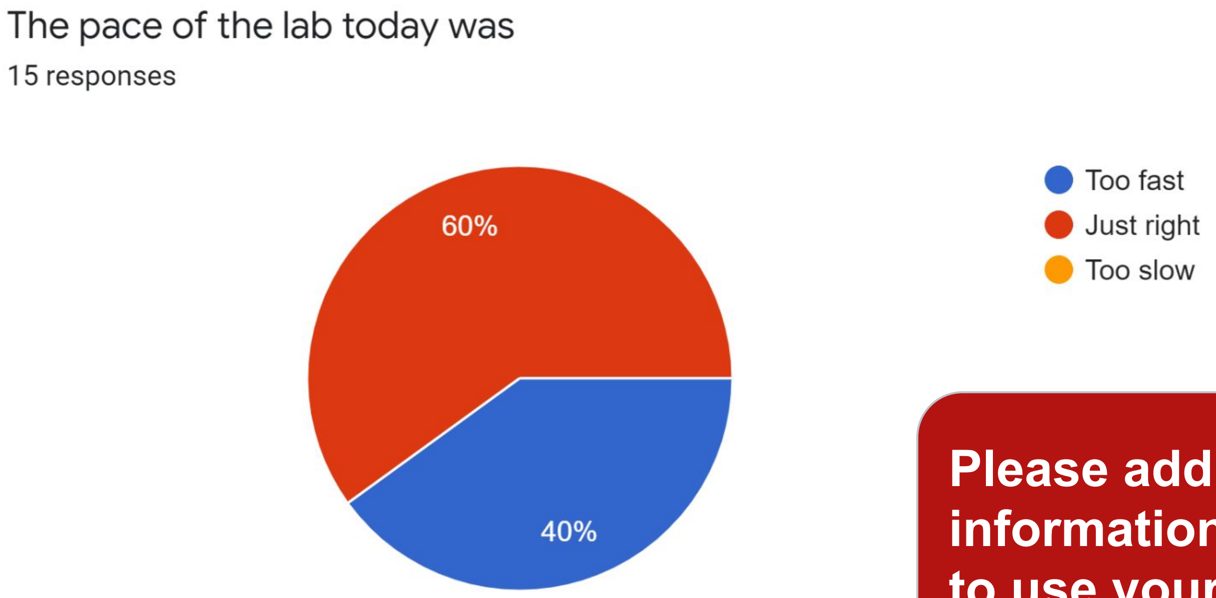

Please add more information about how to use your cell phone!

Keyword Spotting in One Slide

If we **pick a simple task** to only identifying a **few key words** we can then use a **small model** and train it with **little data** and fit it onto an **embedded device**

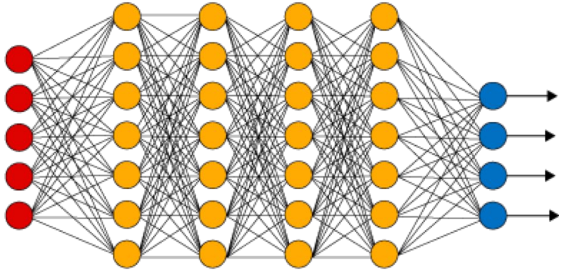

By the end of today: **Hands-on Keyword Spotting (KWS)**

We will explore the **science** behind KWS and **collect data** and **train** our own custom model to recognize “yes” vs. “no” using **Edge Impulse**

Today’s Agenda

- A Quick Review of What We’ve Learned
- Data Engineering for KWS
- Hands-on KWS Data Collection with Edge Impulse
- Training our Model using Transfer Learning
- Deploying our Model onto our Arduino
- Summary

Today’s Agenda

- **A Quick Review of What We’ve Learned**
- Data Engineering for KWS
- Hands-on KWS Data Collection with Edge Impulse
- (Hands-on) Data Preprocessing for KWS
- Deploying our Model onto our Arduino
- Summary

Machine Learning

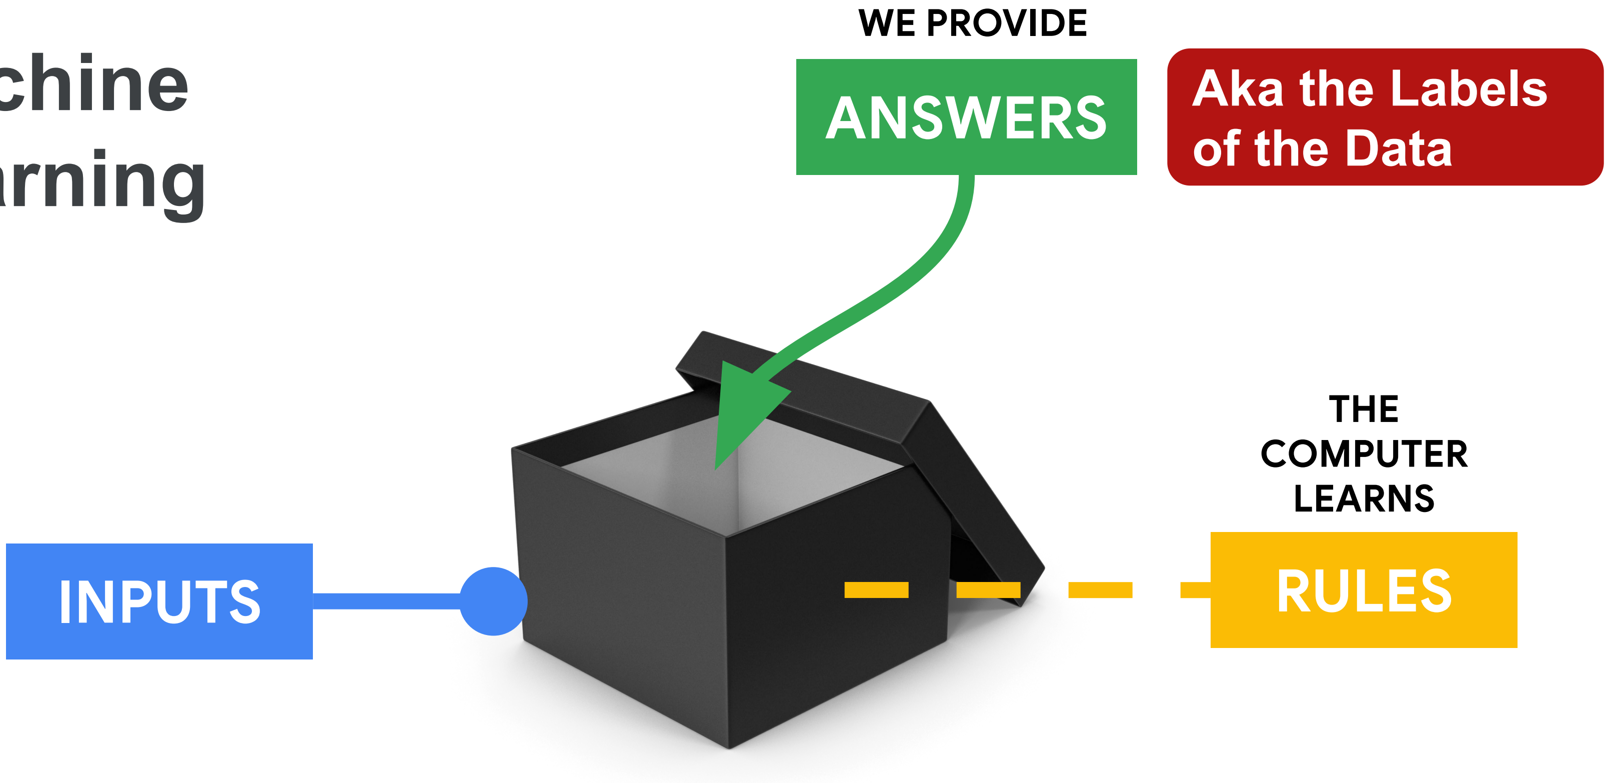

Deep Learning with **Neural Networks**

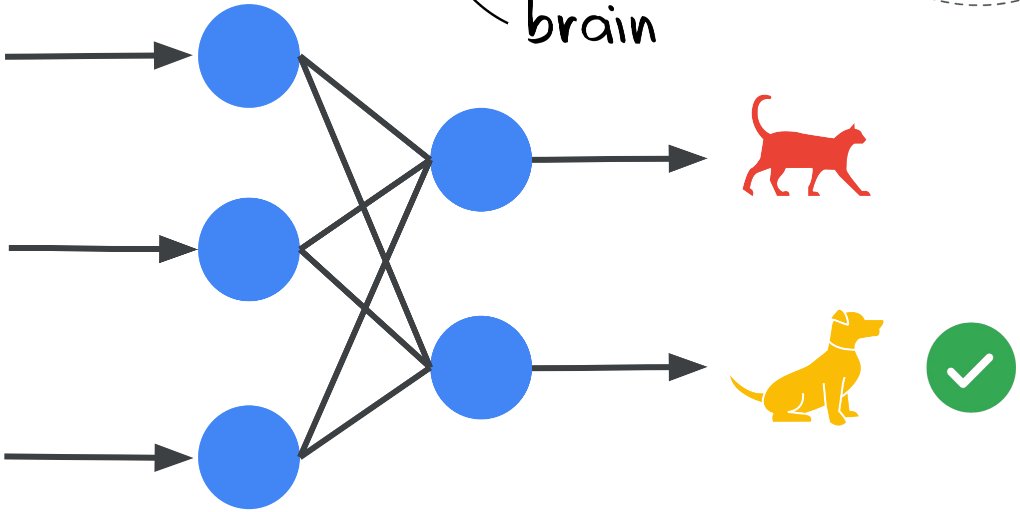

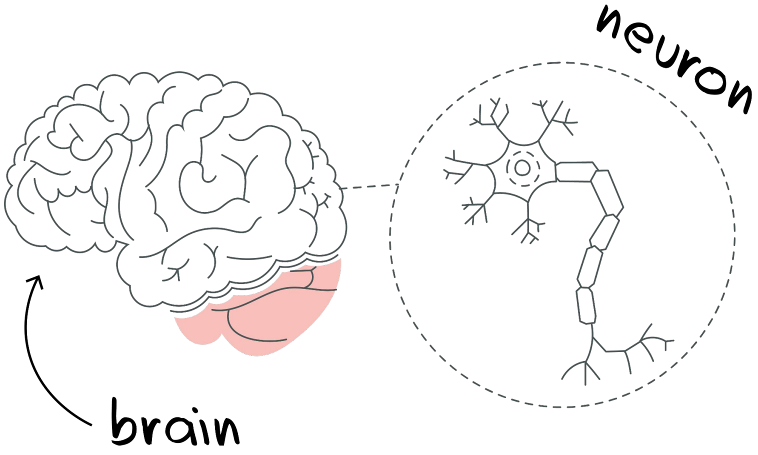

[Colab Link]

Features can be found with **Convolutions**

-1	0	1
-2	0	2
-1	0	1

Features

AlexNet Paper

Convolutional Neural Networks

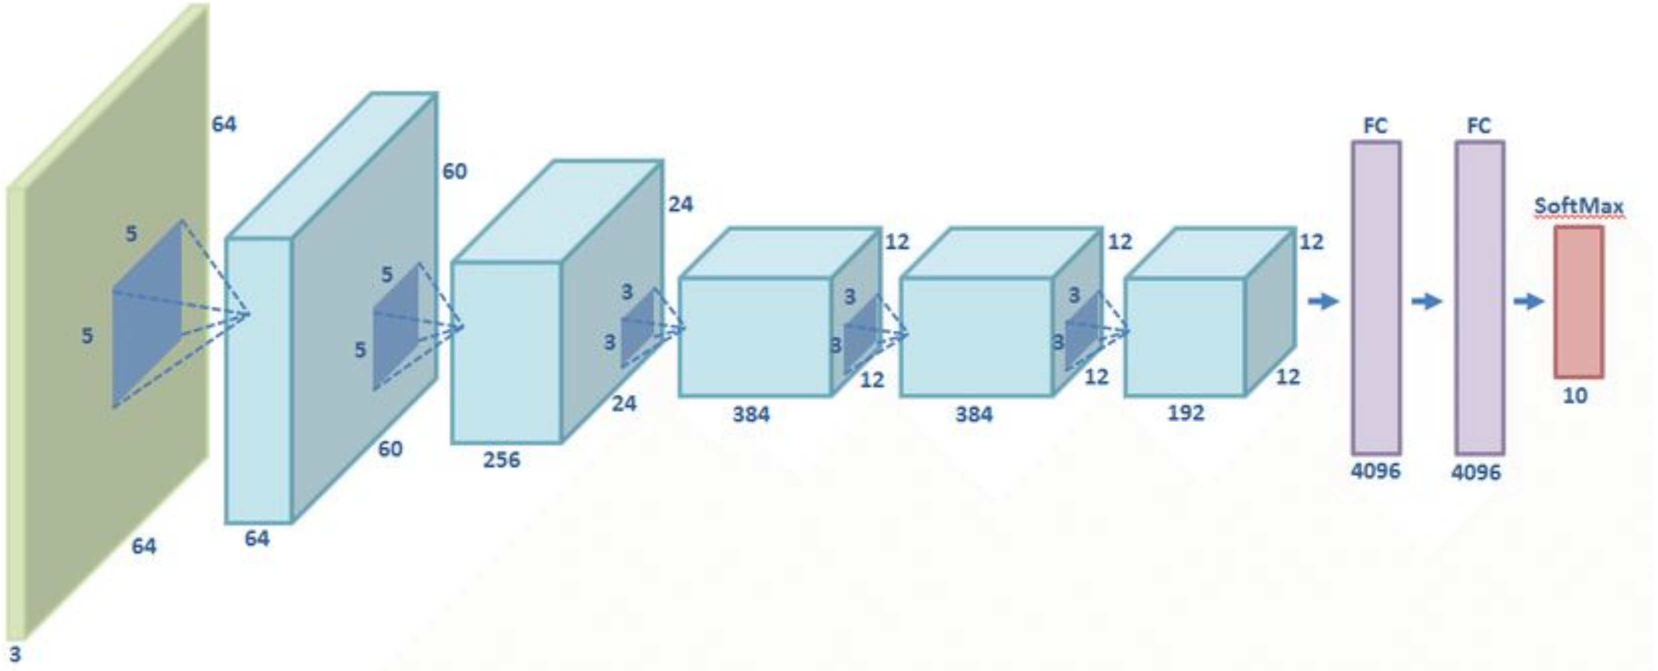

The **TinyML** Workflow

Camera feed

```
Starting inferencing in 2 seconds...
Taking photo...
Predictions (DSP: 9 ms., Classification:
    car: 0.07812
    truck: 0.92188
```

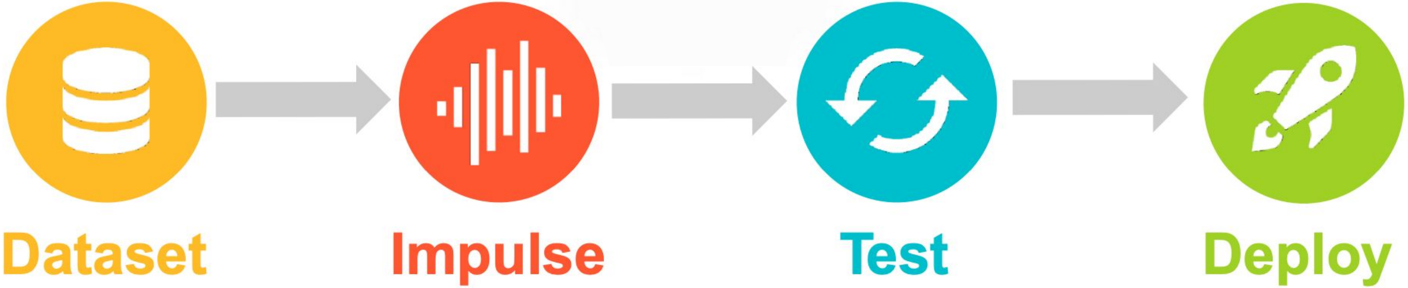

Today’s Agenda

- A Quick Review of What We’ve Learned
- **Data Engineering for KWS**
- Hands-on KWS Data Collection with Edge Impulse
- (Hands-on) Data Preprocessing for KWS
- Deploying our Model onto our Arduino
- Summary

Data Engineering for KWS

(How to collect good data)

Data Engineering for KWS

(How to collect good data)

Who will use your ML model?

- What **languages** will they speak?
- What **accents** will they have?
- Will they use **slang** or formal diction?

Data Engineering for KWS

(How to collect good data)

Who will use your ML model?

- What **languages** will they speak?
- What **accents** will they have?
- Will they use **slang** or formal diction?

Where will your ML model be used?

- Will there be **background noise**?
- **How far** will users be from the microphone?
- Will there be **echos**?

Data Engineering for KWS
(How to collect good data)

Who will use your ML model?

- What **languages** will they speak?
- What **accents** will they have?
- Will they use **slang** or formal diction?

Where will your ML model be used?

- Will there be **background noise**?
- **How far** will users be from the microphone?
- Will there be **echos**?

Why will your ML model be used?
Why those Keywords?

- What **tone of voice** will be used?
- Are your **keywords commonly** used? (aka will you get a lot of false positives)
- What about false negatives?

Data Engineering for KWS

(How to collect good data)

There are a lot more things to consider to **eliminate bias** and **protect privacy** when collecting data that we will talk about in future sessions!

ML model be used?

Why those Keywords?

... aka will you get a lot of false positives)

- What about false negatives?

Tips and Tricks for Custom KWS

- Pick **uncommon words** for Keywords
- Record lots of “**other words**”
- Record in the **location** you are going to be **deploying**
- Get **your end users** to help you build a dataset
- Record with the same **hardware** you will **deploy**
- Always **test** and then **improve** your dataset and model

Tips and Tricks for Custom KWS

Today we are just working on a demo so to give our demo the the best chance of working we will:

1. **Stay in one spot**
2. **Only record ourselves**
3. **Use common words (yes, no)**
4. **Only test ourselves**

(we’re cheating)

Today’s Agenda

- A Quick Review of What We’ve Learned
- Data Engineering for KWS
- **Hands-on KWS Data Collection with Edge Impulse**
- (Hands-on) Data Preprocessing for KWS
- Deploying our Model onto our Arduino
- Summary

The TinyML Workflow using Edge Impulse

Today we'll also collect all of our data using Edge Impulse...

...and deploy to your cell phone as well

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

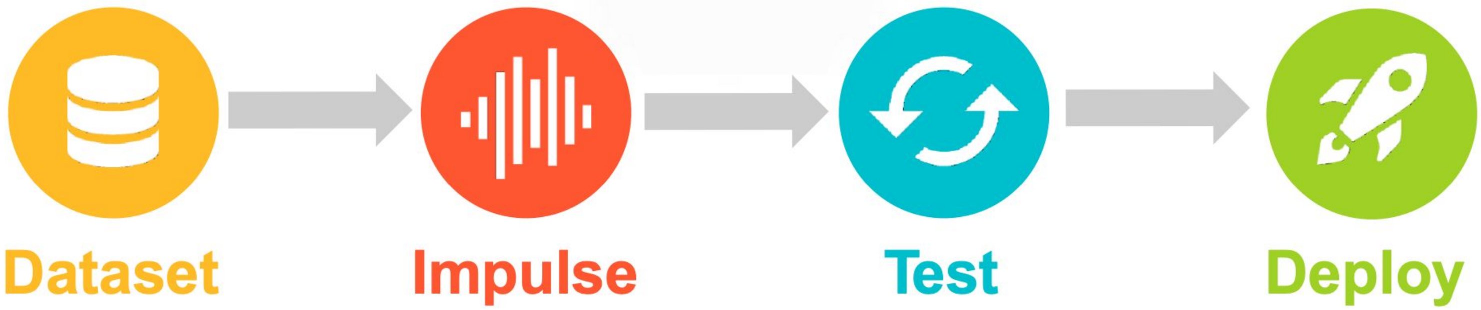

Edge Impulse Project Dashboard

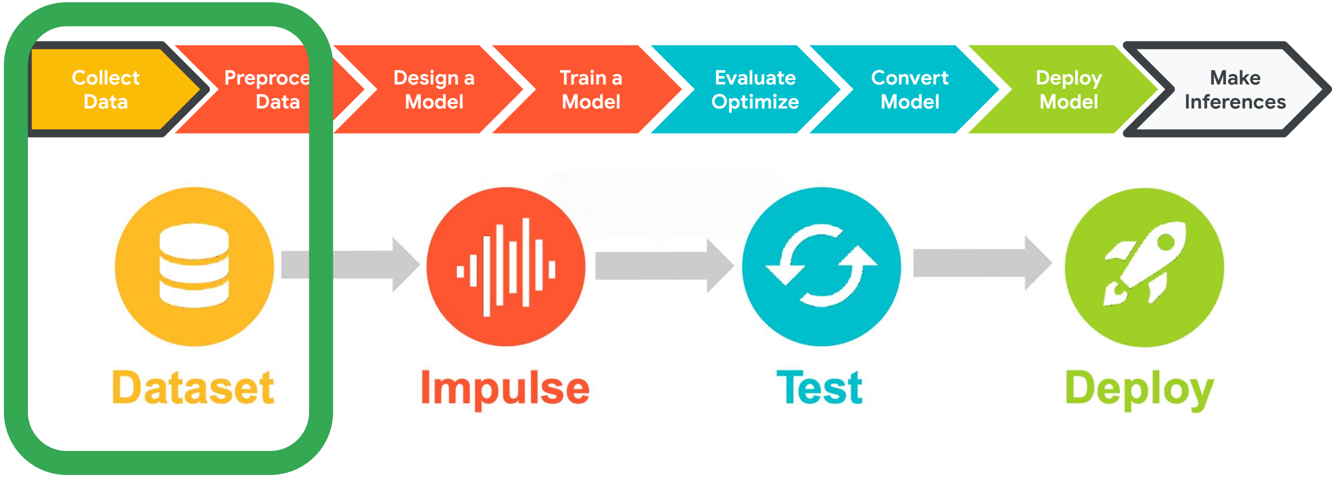

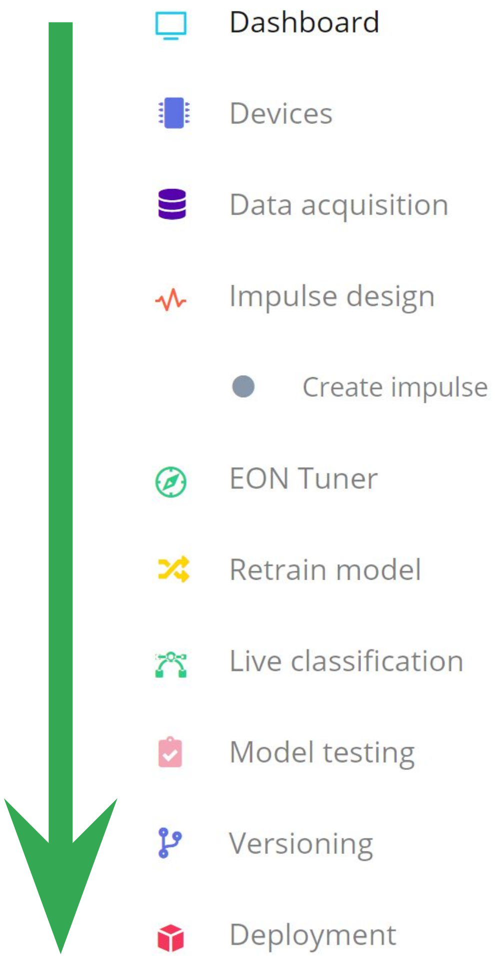

Create an Edge Impulse Account

1. Create an Edge Impulse account: https://studio.edgeimpulse.com/signup

2. Validate your email by clicking the link in the email sent to your account's email address

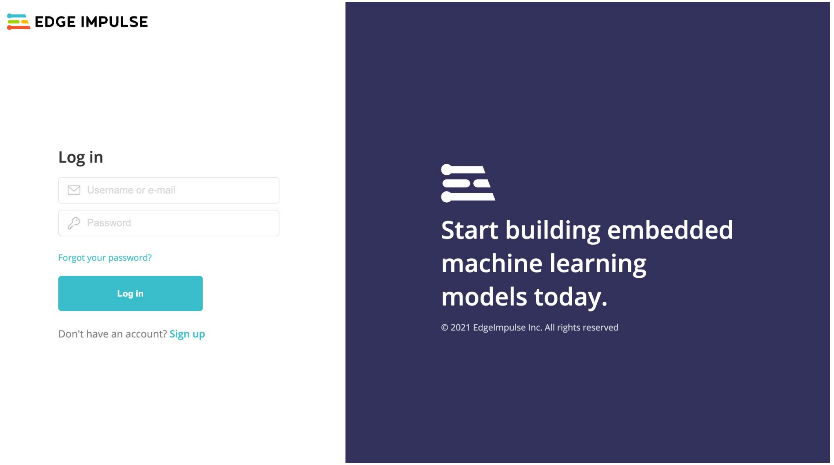

Activity: Create a Keyword Spotting Dataset

Collect **~30 samples each** of the following classes of data:

- **Keyword #1 “yes”** (label: yes) (length: 1 seconds)
- **Keyword #2 “no”** (label: no) (length: 1 seconds)
- **“Unknown” words** that are not the keyword **and background noise** (label: unknown) (length: 1 seconds)

https://docs.edgeimpulse.com/docs/pre-built-datasets/keyword-spotting

Activity: Create a Keyword Spotting Dataset

Collect **~30 samples each** of the following classes of data:

- **Keyword #1 "yes"** (label: yes) (length: 1 seconds)
- **Keyword #2 "no"** (label: no) (length: 1 seconds)
- ~~**"Unknown" words** that are not the keyword **and** **bac**~~ ~~nds)~~

I've pre-loaded in a bunch of background noise and unknown words!

https://www.edgeimpulse.com/blog/public-projects-launch

Clone my starter KWS project: https://bit.ly/SciTinyML22-KWS

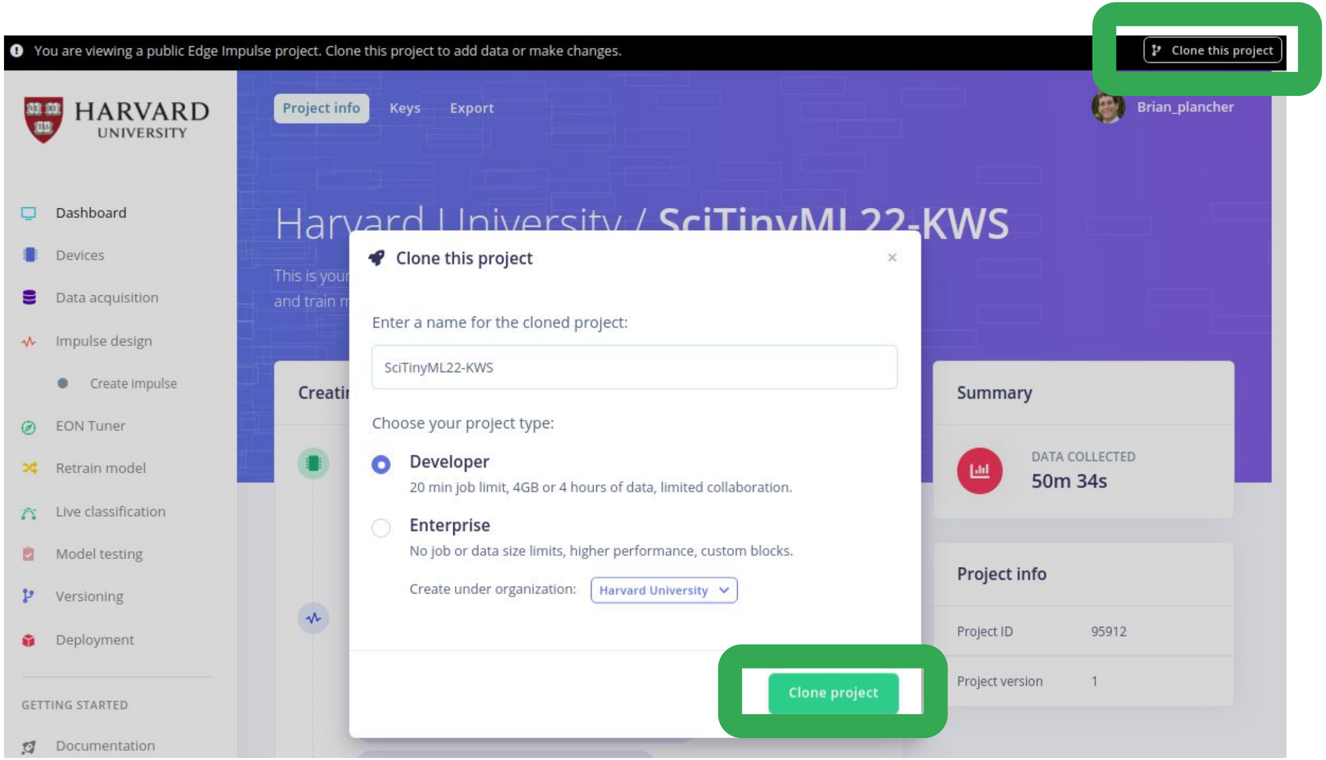

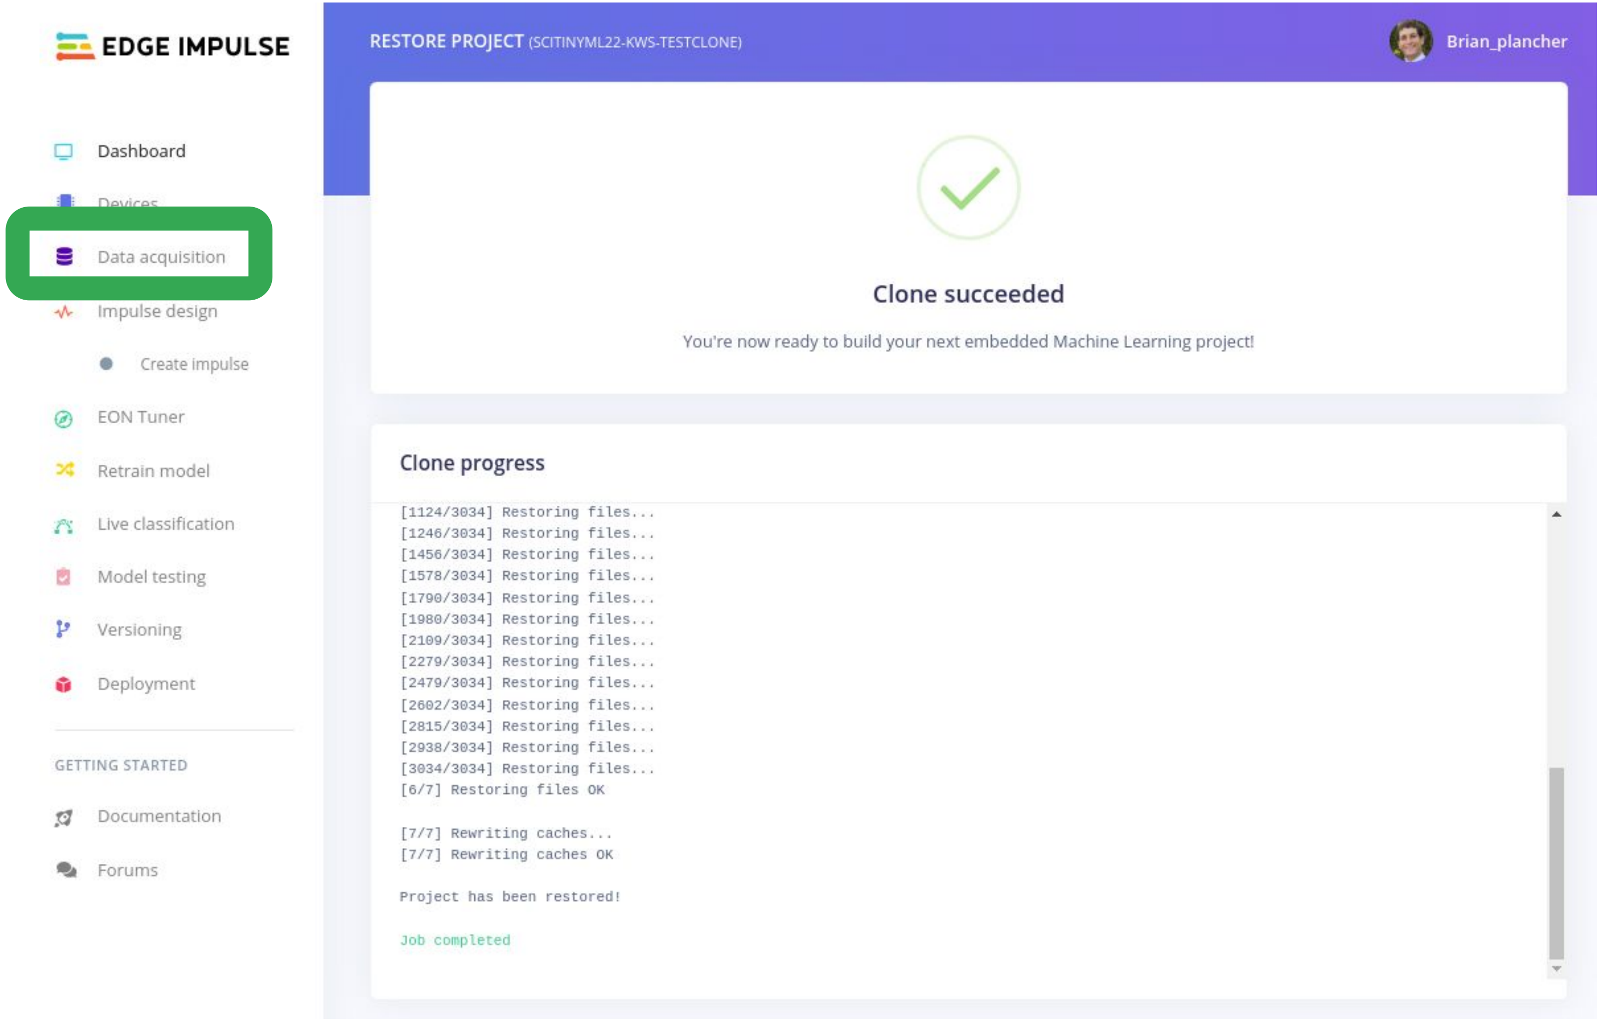
EDGE IMPULSE
RESTORE PROJECT (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Clone succeeded
You're now ready to build your next embedded Machine Learning project!
Clone progress
[1124/3034] Restoring files...
[1246/3034] Restoring files...
[1456/3034] Restoring files...
[1578/3034] Restoring files...
[1790/3034] Restoring files...
[1980/3034] Restoring files...
[2109/3034] Restoring files...
[2279/3034] Restoring files...
[2479/3034] Restoring files...
[2602/3034] Restoring files...
[2815/3034] Restoring files...
[2938/3034] Restoring files...
[3034/3034] Restoring files...
[6/7] Restoring files OK
[7/7] Rewriting caches...
[7/7] Rewriting caches OK
Project has been restored!
Job completed

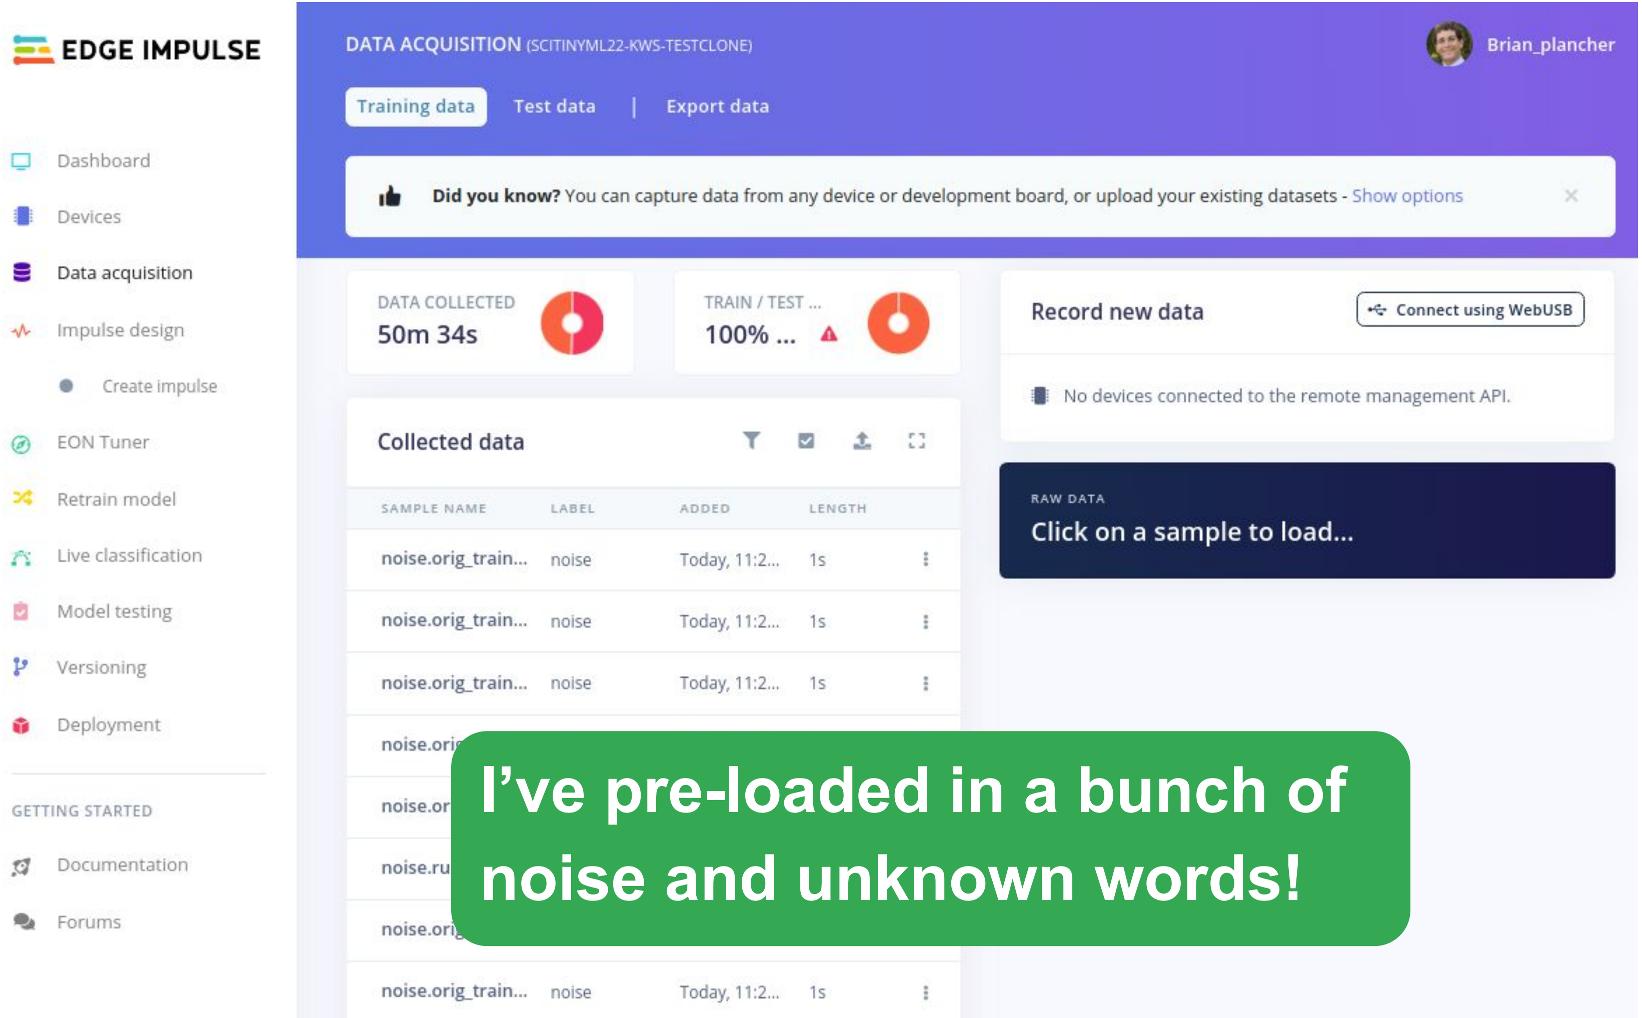

I've pre-loaded in a bunch of noise and unknown words!

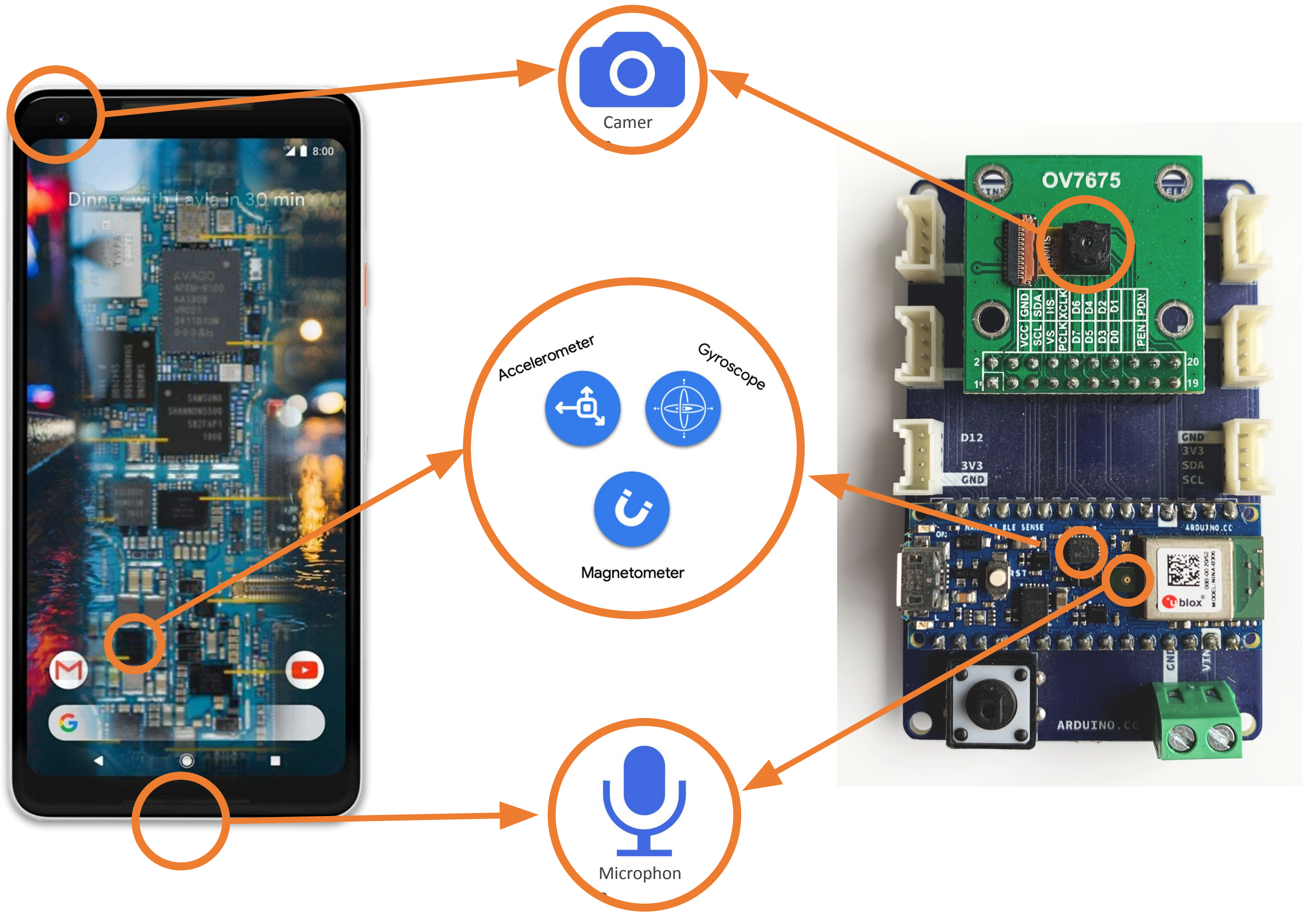
Camer
8:00
Dinner with Layla in 30 min
Accelerometer
Gyroscope
Magnetometer
Microphon
OV7675
D12
3V3
GND
SDA
SCL
ARDUINO.CC

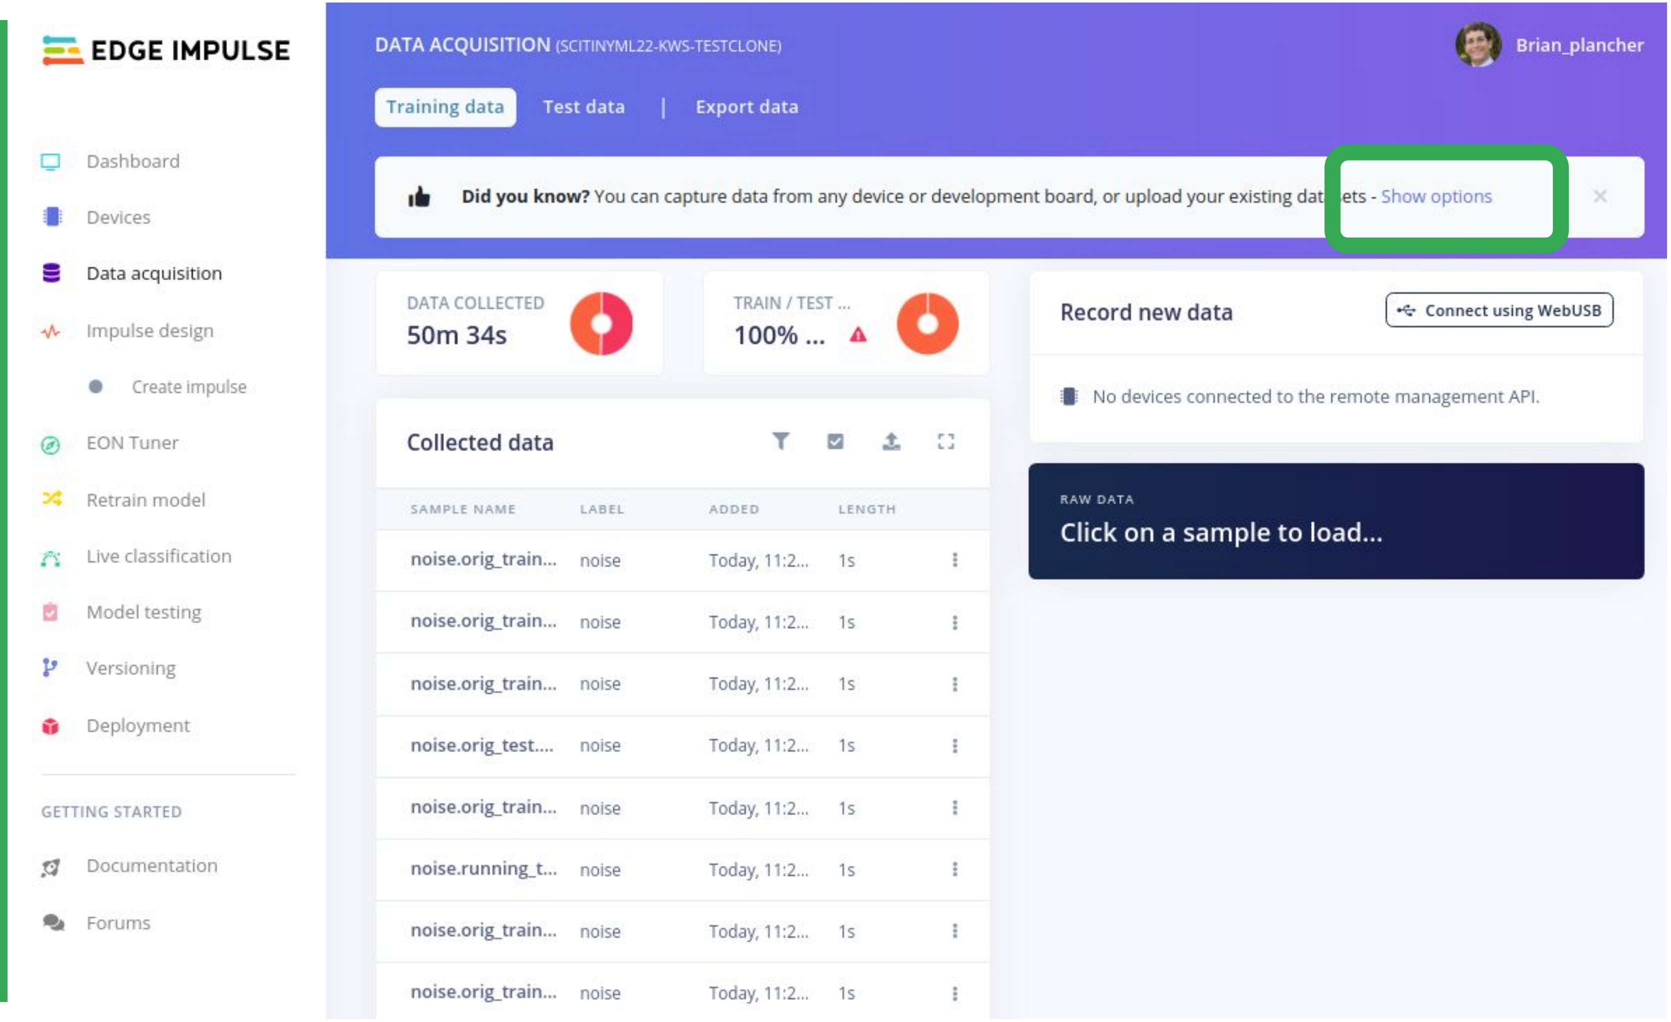
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
DATA ACQUISITION (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
Training data
Test data
Export data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
DATA COLLECTED
50m 34s
TRAIN / TEST ...
100% ...
Record new data
Connect using WebUSB
No devices connected to the remote management API.
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
noise.orig_train... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.orig_test.... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.running_t... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
RAW DATA
Click on a sample to load...

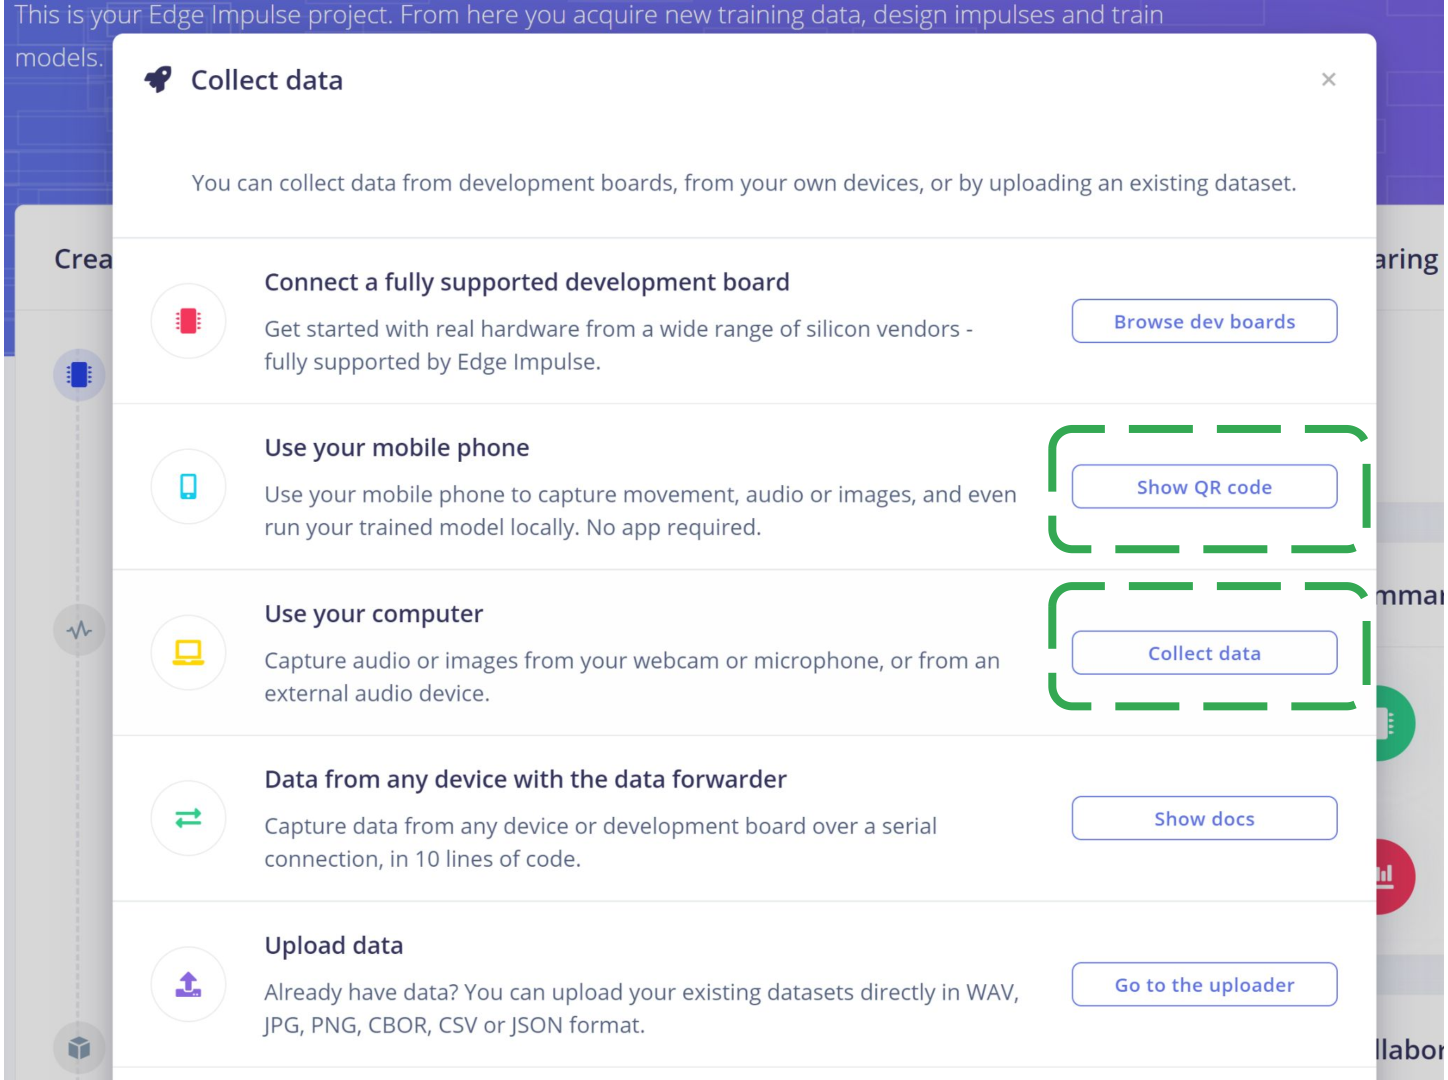
This is your Edge Impulse project. From here you acquire new training data, design impulses and train models.
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
Use your mobile phone
Use your mobile phone to capture movement, audio or images, and even run your trained model locally. No app required.
Show QR code
Use your computer
Capture audio or images from your webcam or microphone, or from an external audio device.
Collect data
Data from any device with the data forwarder
Capture data from any device or development board over a serial connection, in 10 lines of code.
Show docs
Upload data
Already have data? You can upload your existing datasets directly in WAV, JPG, PNG, CBOR, CSV or JSON format.
Go to the uploader

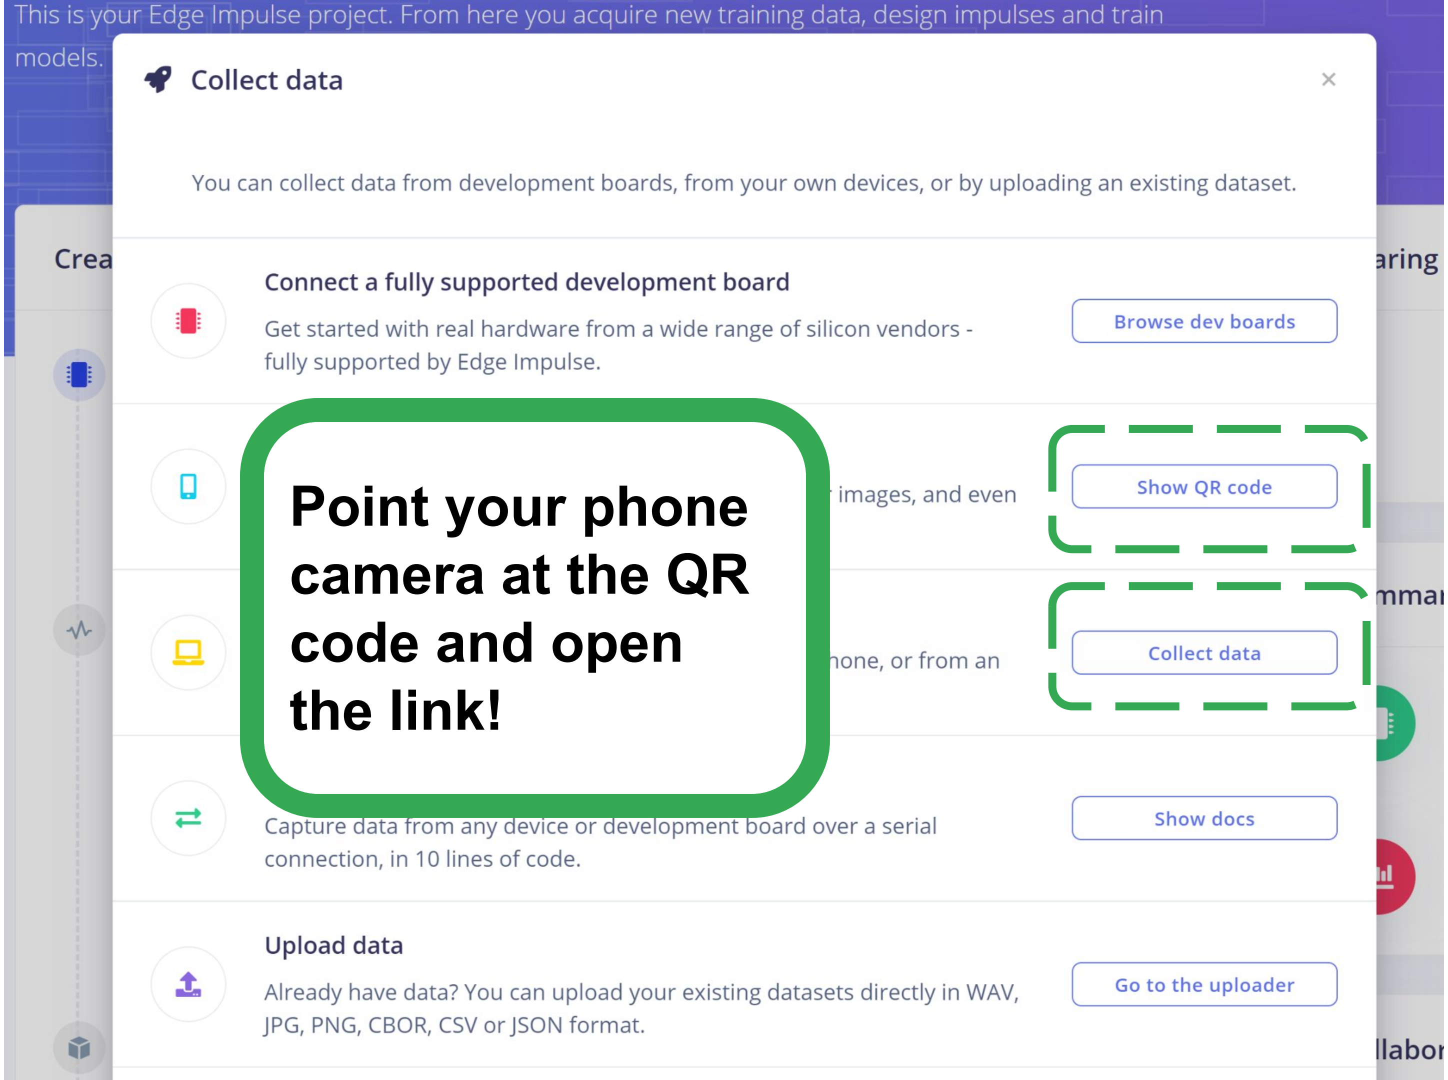
This is your Edge Impulse project. From here you acquire new training data, design impulses and train models.
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
images, and even
Show QR code
hone, or from an
Collect data
Capture data from any device or development board over a serial connection, in 10 lines of code.
Show docs
Upload data
Already have data? You can upload your existing datasets directly in WAV, JPG, PNG, CBOR, CSV or JSON format.
Go to the uploader
Point your phone camera at the QR code and open the link!

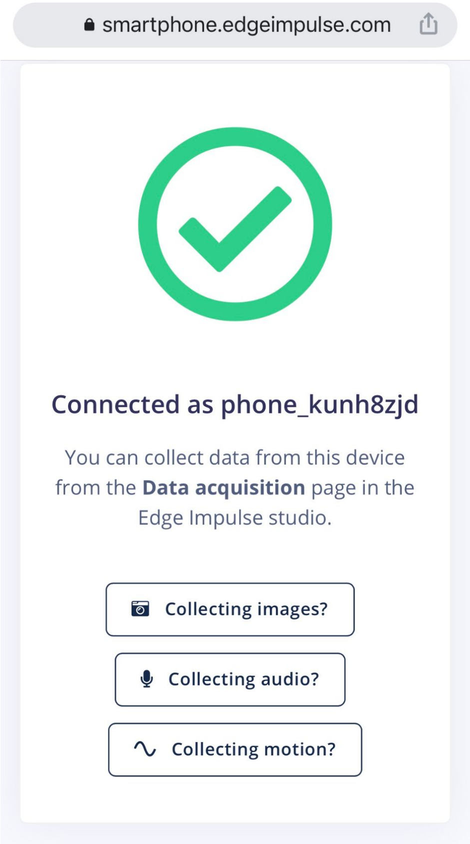
smartphone.edgeimpulse.com
Connected as phone_kunh8zjd
You can collect data from this device from the **Data acquisition** page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?

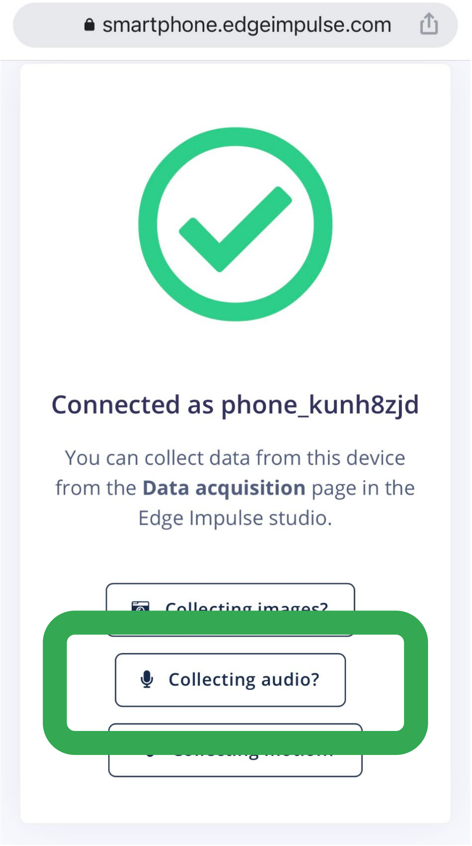
smartphone.edgeimpulse.com
Connected as phone_kunh8zjd
You can collect data from this device from the **Data acquisition** page in the Edge Impulse studio.
Collecting audio?

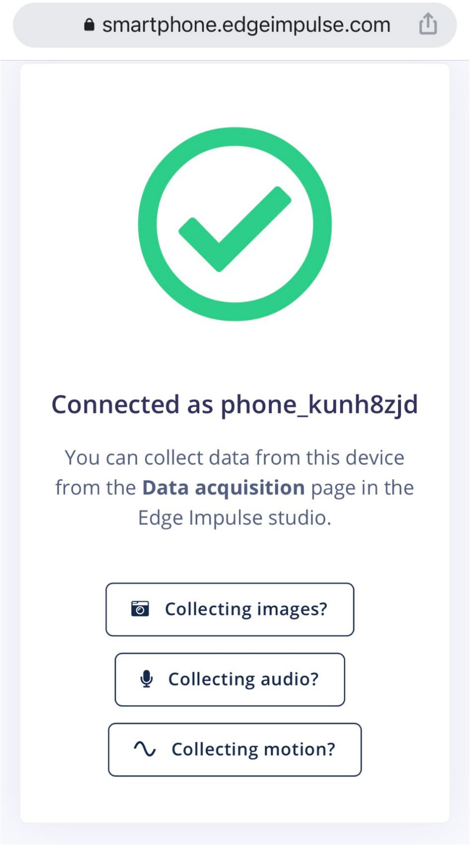
smartphone.edgeimpulse.com
Connected as phone_kunh8zjd
You can collect data from this device from the Data acquisition page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?

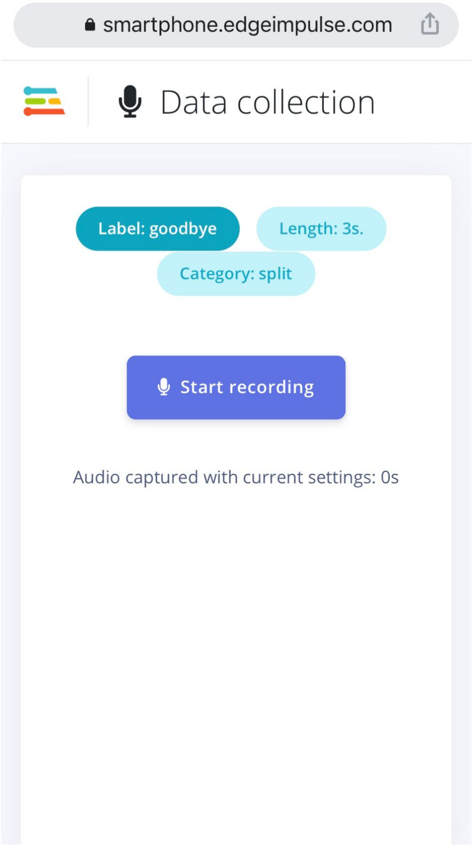
smartphone.edgeimpulse.com
Data collection
Label: goodbye
Length: 3s.
Category: split
Start recording
Audio captured with current settings: 0s

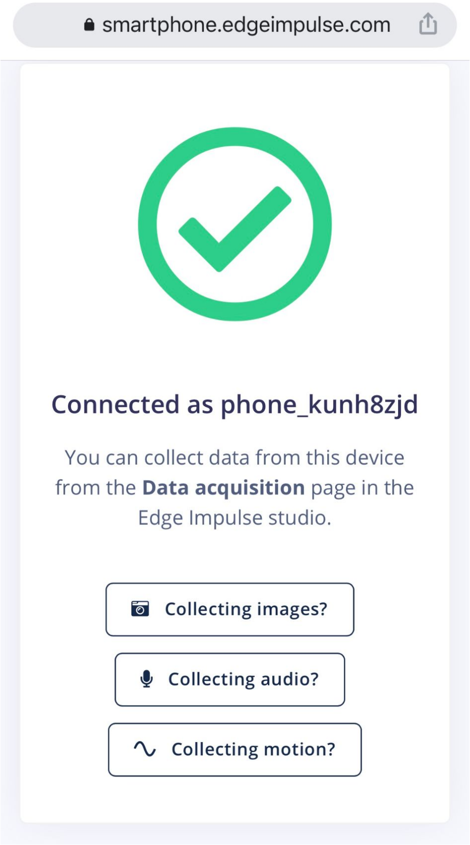
smartphone.edgeimpulse.com
Connected as phone_kunh8zjd
You can collect data from this device from the Data acquisition page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?

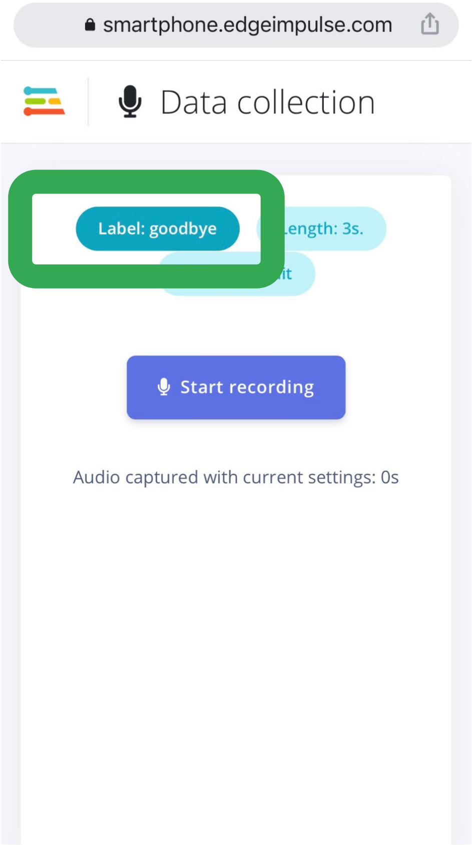
smartphone.edgeimpulse.com
Data collection
Label: goodbye
Start recording
Audio captured with current settings: 0s

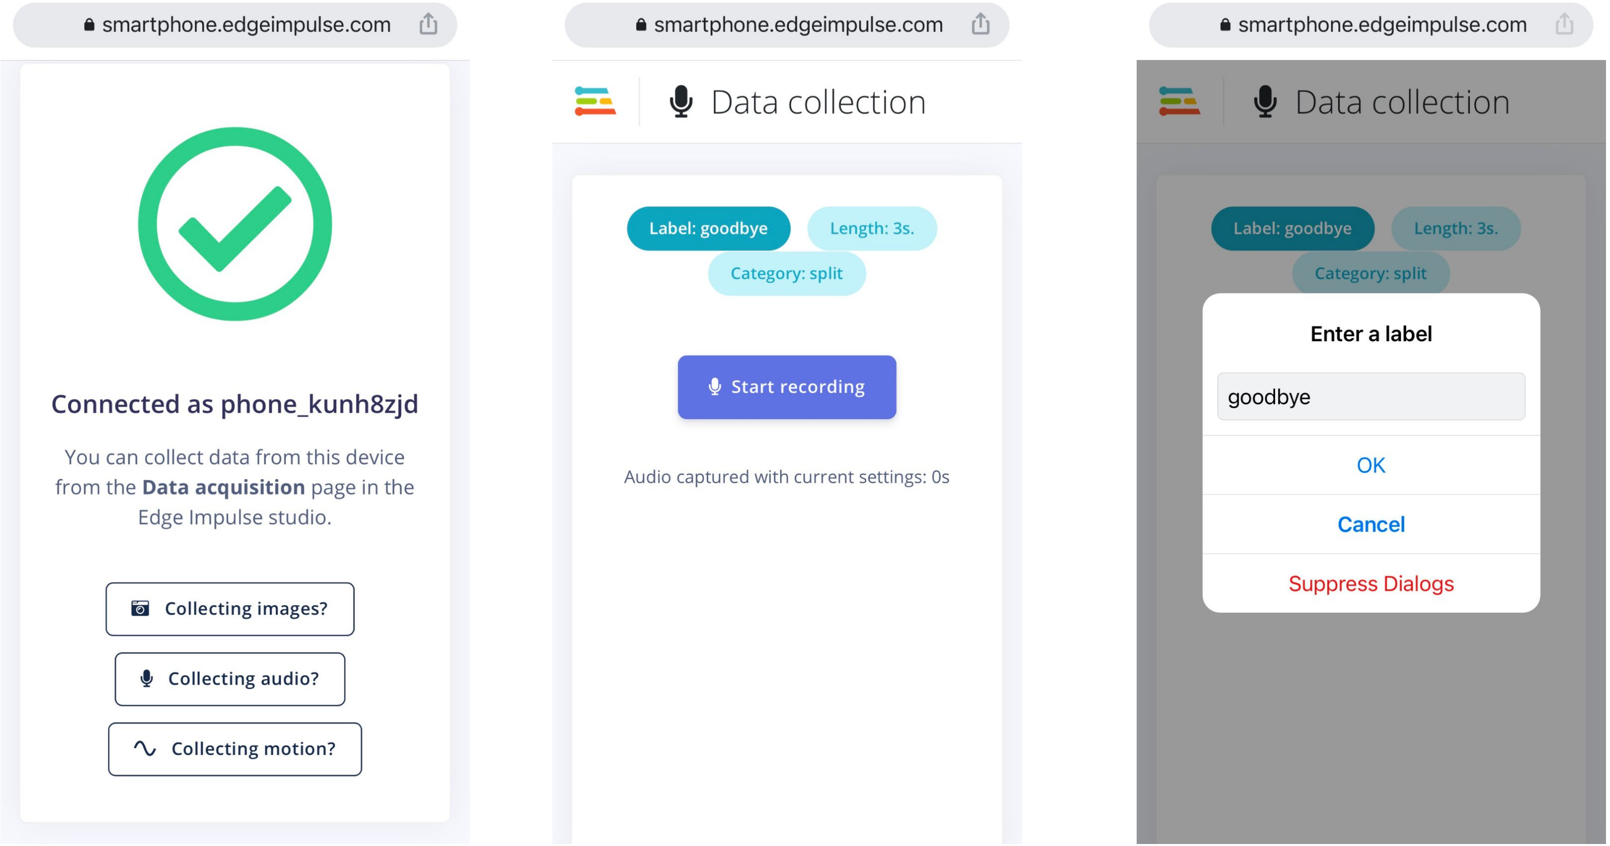
smartphone.edgeimpulse.com
Connected as phone_kunh8zjd
You can collect data from this device from the Data acquisition page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?
smartphone.edgeimpulse.com
Data collection
Label: goodbye
Length: 3s.
Category: split
Start recording
Audio captured with current settings: 0s
smartphone.edgeimpulse.com
Data collection
Label: goodbye
Length: 3s.
Category: split
Enter a label
goodbye
OK
Cancel
Suppress Dialogs

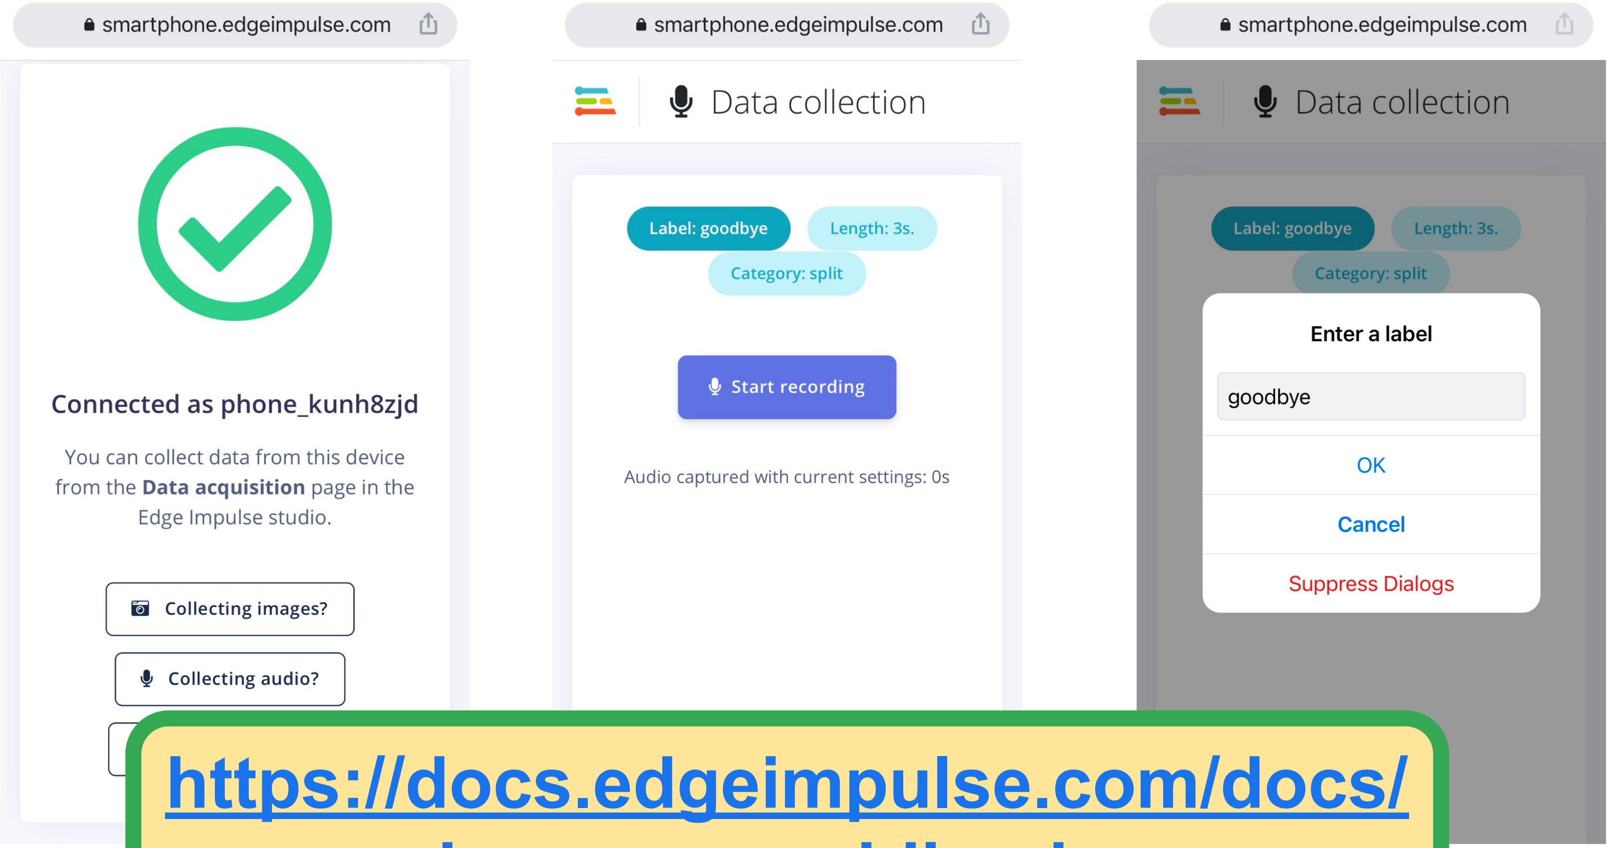

https://docs.edgeimpulse.com/docs/using-your-mobile-phone

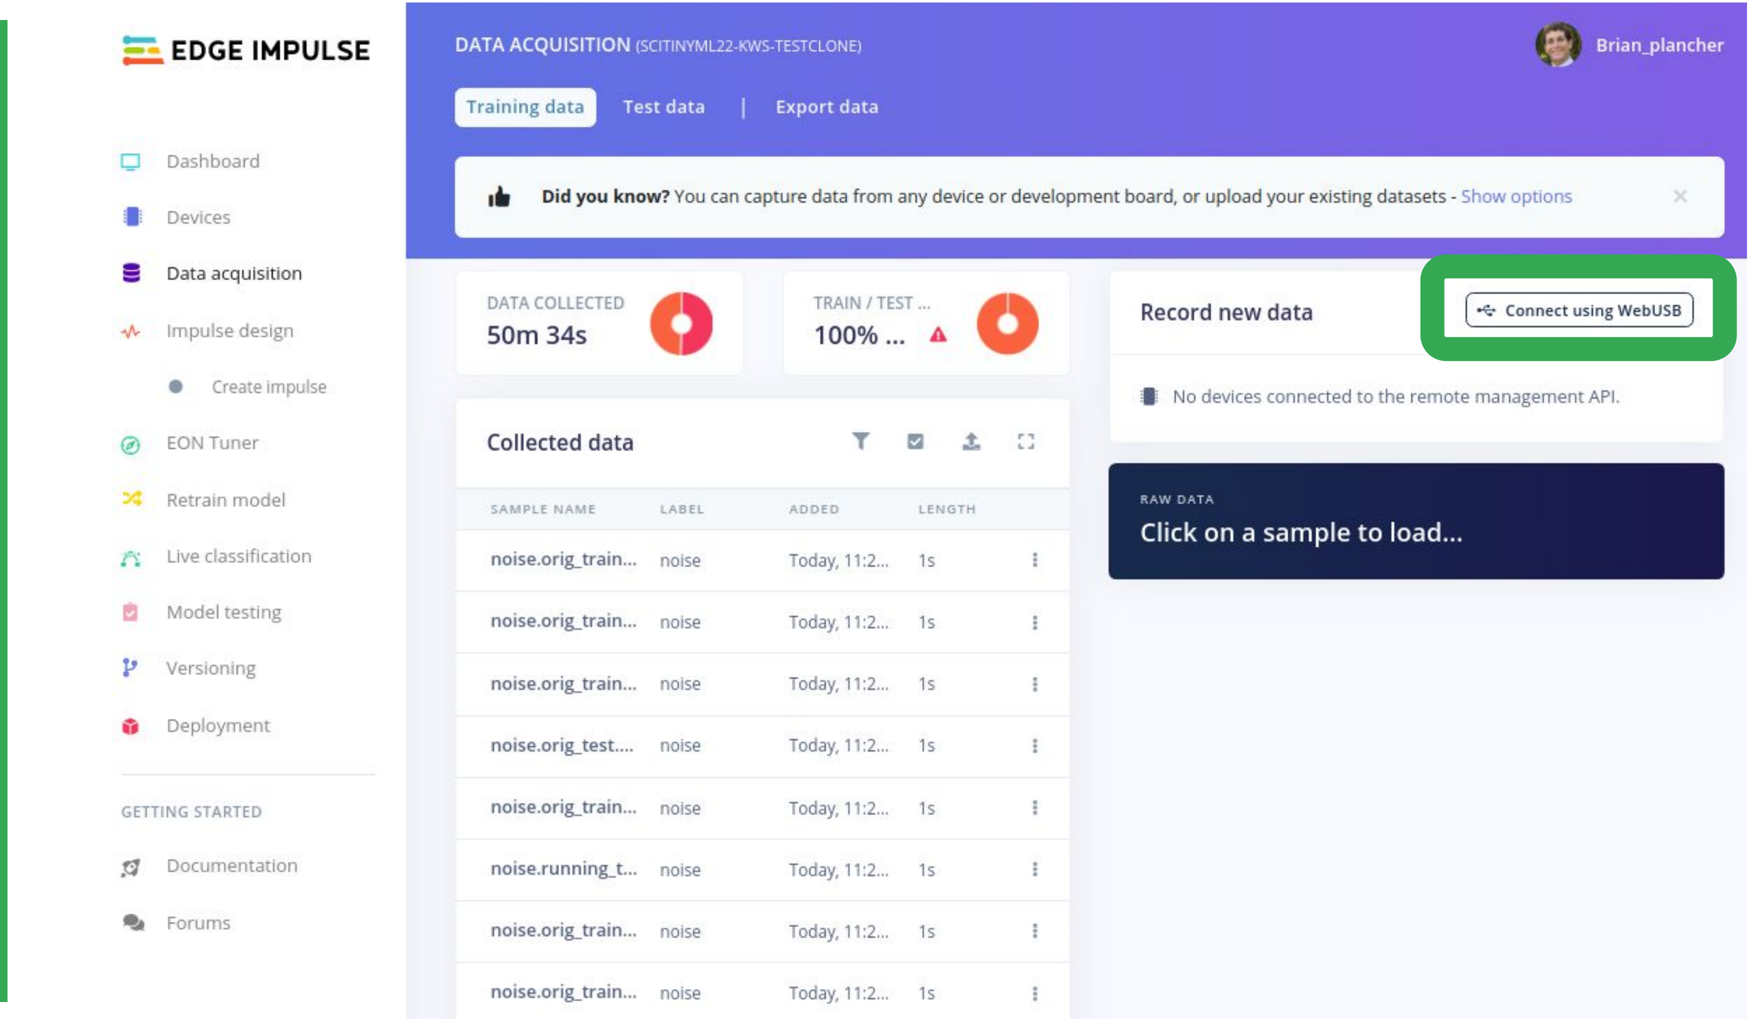
EDGE IMPULSE
DATA ACQUISITION (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
Training data
Test data
Export data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
DATA COLLECTED
50m 34s
TRAIN / TEST ...
100% ...
Record new data
Connect using WebUSB
No devices connected to the remote management API.
RAW DATA
Click on a sample to load...
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
noise.orig_train...
noise
Today, 11:2...
1s
noise.orig_train...
noise
Today, 11:2...
1s
noise.orig_train...
noise
Today, 11:2...
1s
noise.orig_test....
noise
Today, 11:2...
1s
noise.orig_train...
noise
Today, 11:2...
1s
noise.running_t...
noise
Today, 11:2...
1s
noise.orig_train...
noise
Today, 11:2...
1s
noise.orig_train...
noise
Today, 11:2...
1s

You may need to re-flash the EI Firmware!

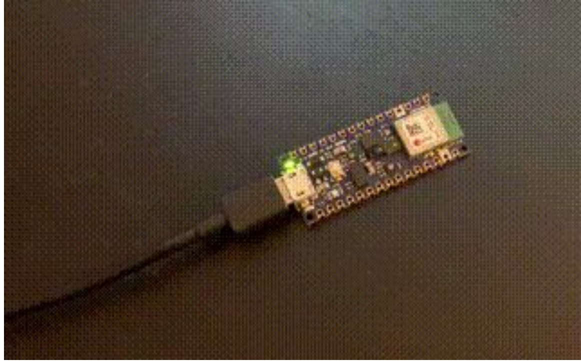

1. Double tap RESET to enter bootloader mode

2. Download the firmware: bit.ly/EI-Nano33-Firmware

3. Run the flash script for your operating system (`flash_windows.bat`, `flash_mac.command` or `flash_linux.sh`).

4. Wait until flashing is complete, and press the RESET button once to launch the new firmware.

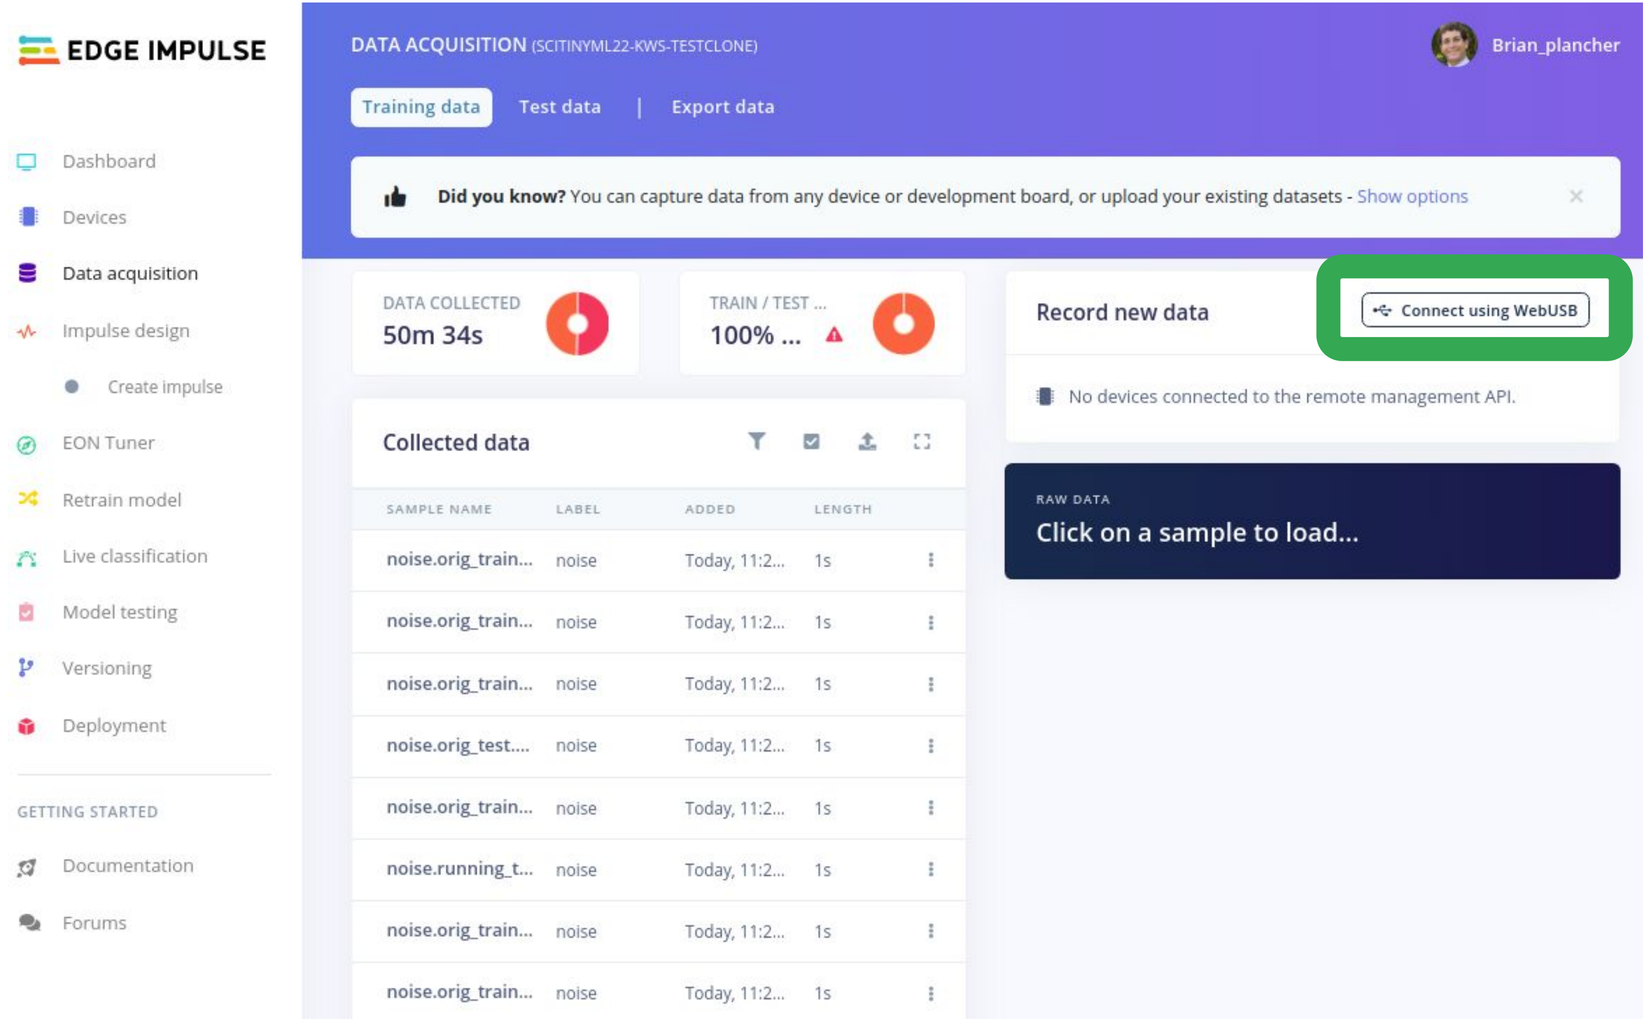
EDGE IMPULSE
DATA ACQUISITION (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
Training data
Test data
Export data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
DATA COLLECTED
50m 34s
TRAIN / TEST ...
100% ...
Record new data
Connect using WebUSB
No devices connected to the remote management API.
RAW DATA
Click on a sample to load...
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
noise.orig_train... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.orig_test.... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.running_t... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s
noise.orig_train... noise Today, 11:2... 1s

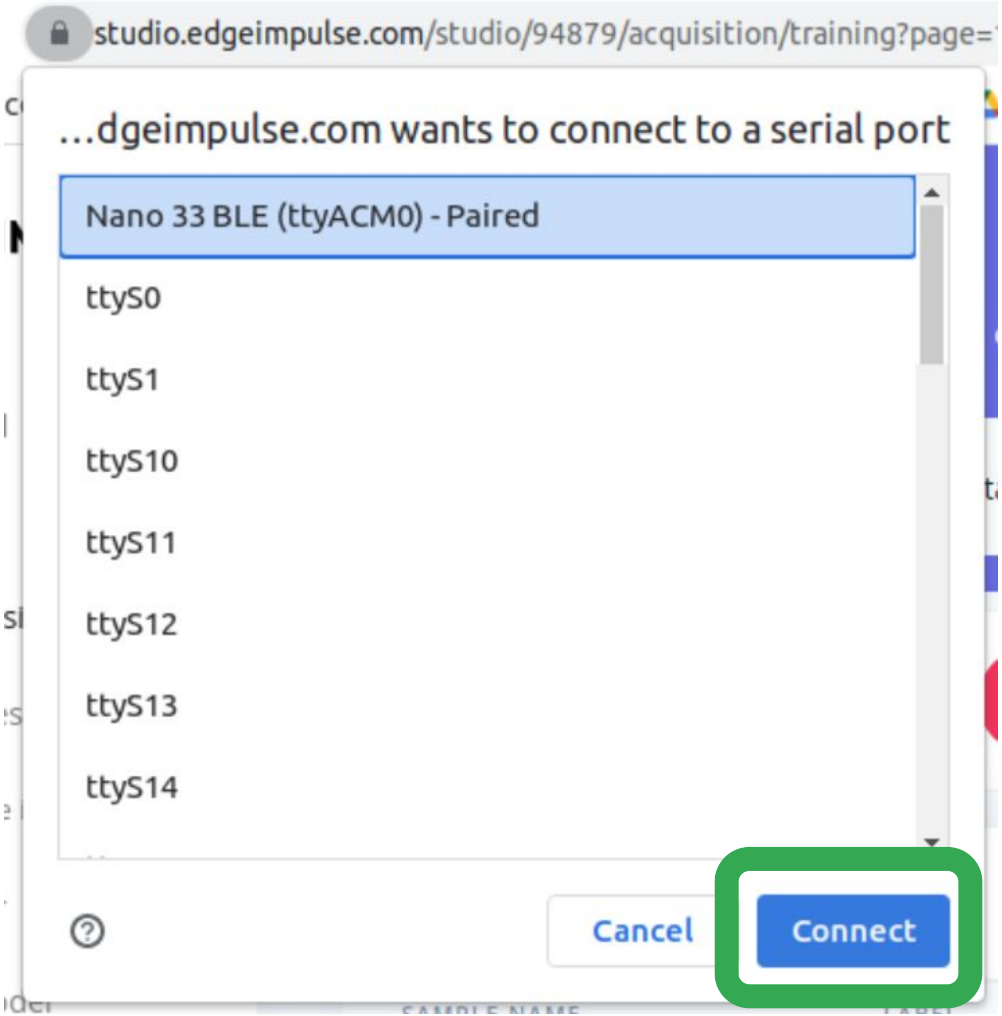
studio.edgeimpulse.com/studio/94879/acquisition/training?page=
...dgeimpulse.com wants to connect to a serial port
Nano 33 BLE (ttyACM0) - Paired
ttyS0
ttyS1
ttyS10
ttyS11
ttyS12
ttyS13
ttyS14
Cancel
Connect

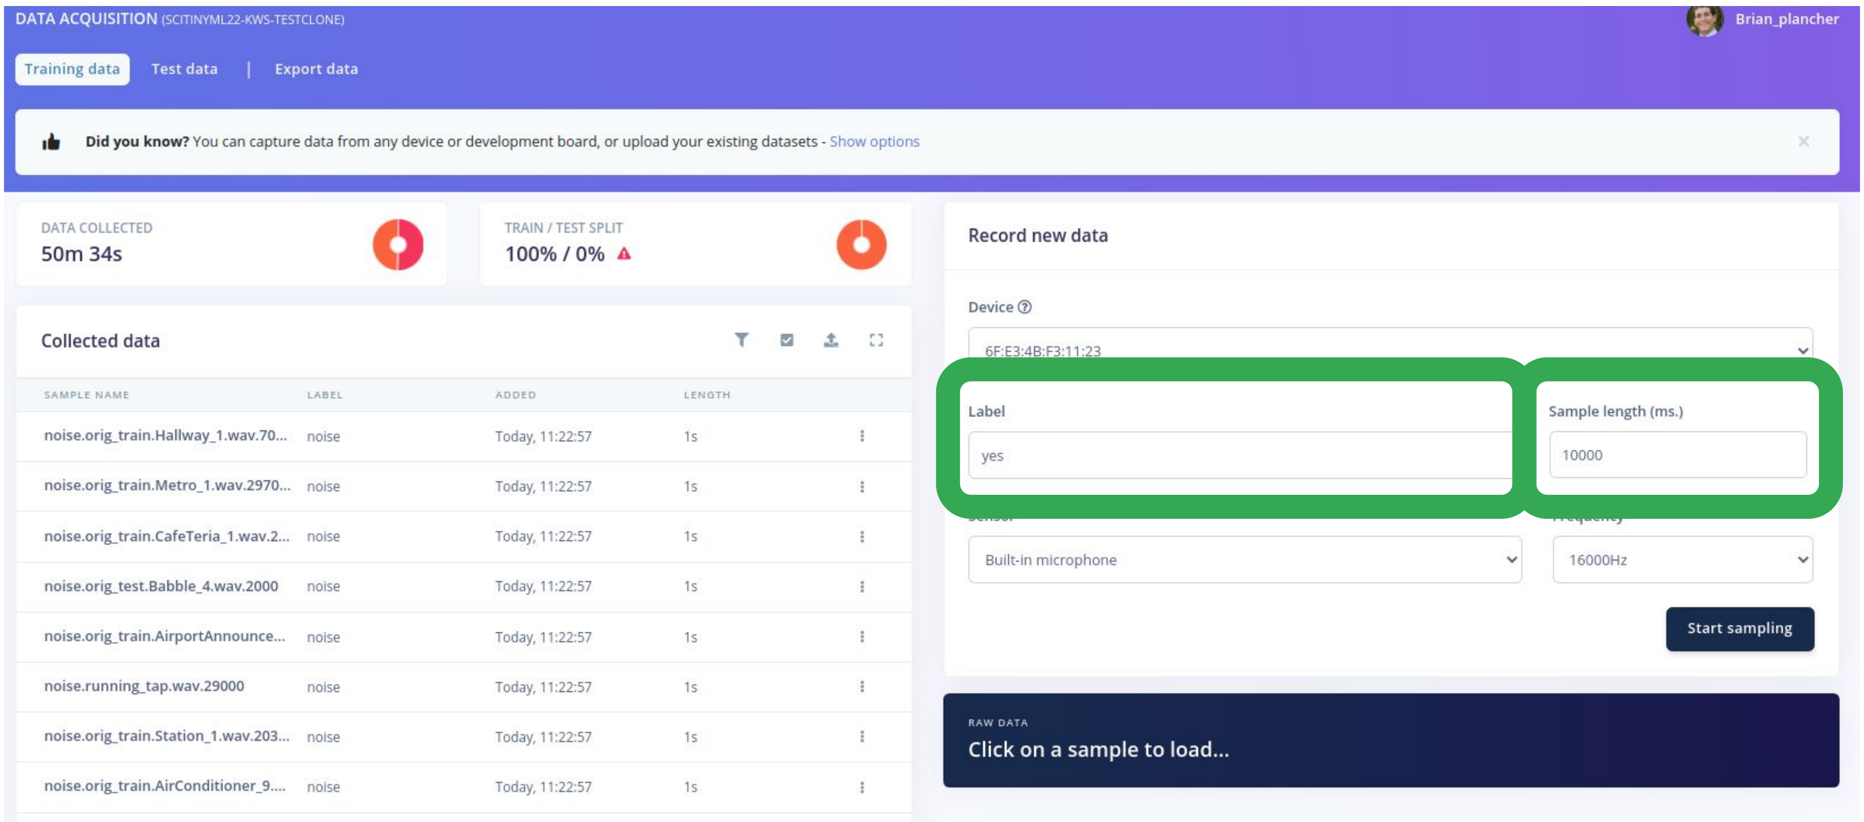
DATA ACQUISITION (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
Training data
Test data
Export data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
DATA COLLECTED
50m 34s
TRAIN / TEST SPLIT
100% / 0%
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
noise.orig_train.Hallway_1.wav.70... noise Today, 11:22:57 1s
noise.orig_train.Metro_1.wav.2970... noise Today, 11:22:57 1s
noise.orig_train.CafeTeria_1.wav.2... noise Today, 11:22:57 1s
noise.orig_test.Babble_4.wav.2000 noise Today, 11:22:57 1s
noise.orig_train.AirportAnnounce... noise Today, 11:22:57 1s
noise.running_tap.wav.29000 noise Today, 11:22:57 1s
noise.orig_train.Station_1.wav.203... noise Today, 11:22:57 1s
noise.orig_train.AirConditioner_9.... noise Today, 11:22:57 1s
Record new data
Device
6F:E3:4B:F3:11:23
Label
yes
Sample length (ms.)
10000
Built-in microphone
16000Hz
Start sampling
RAW DATA
Click on a sample to load...

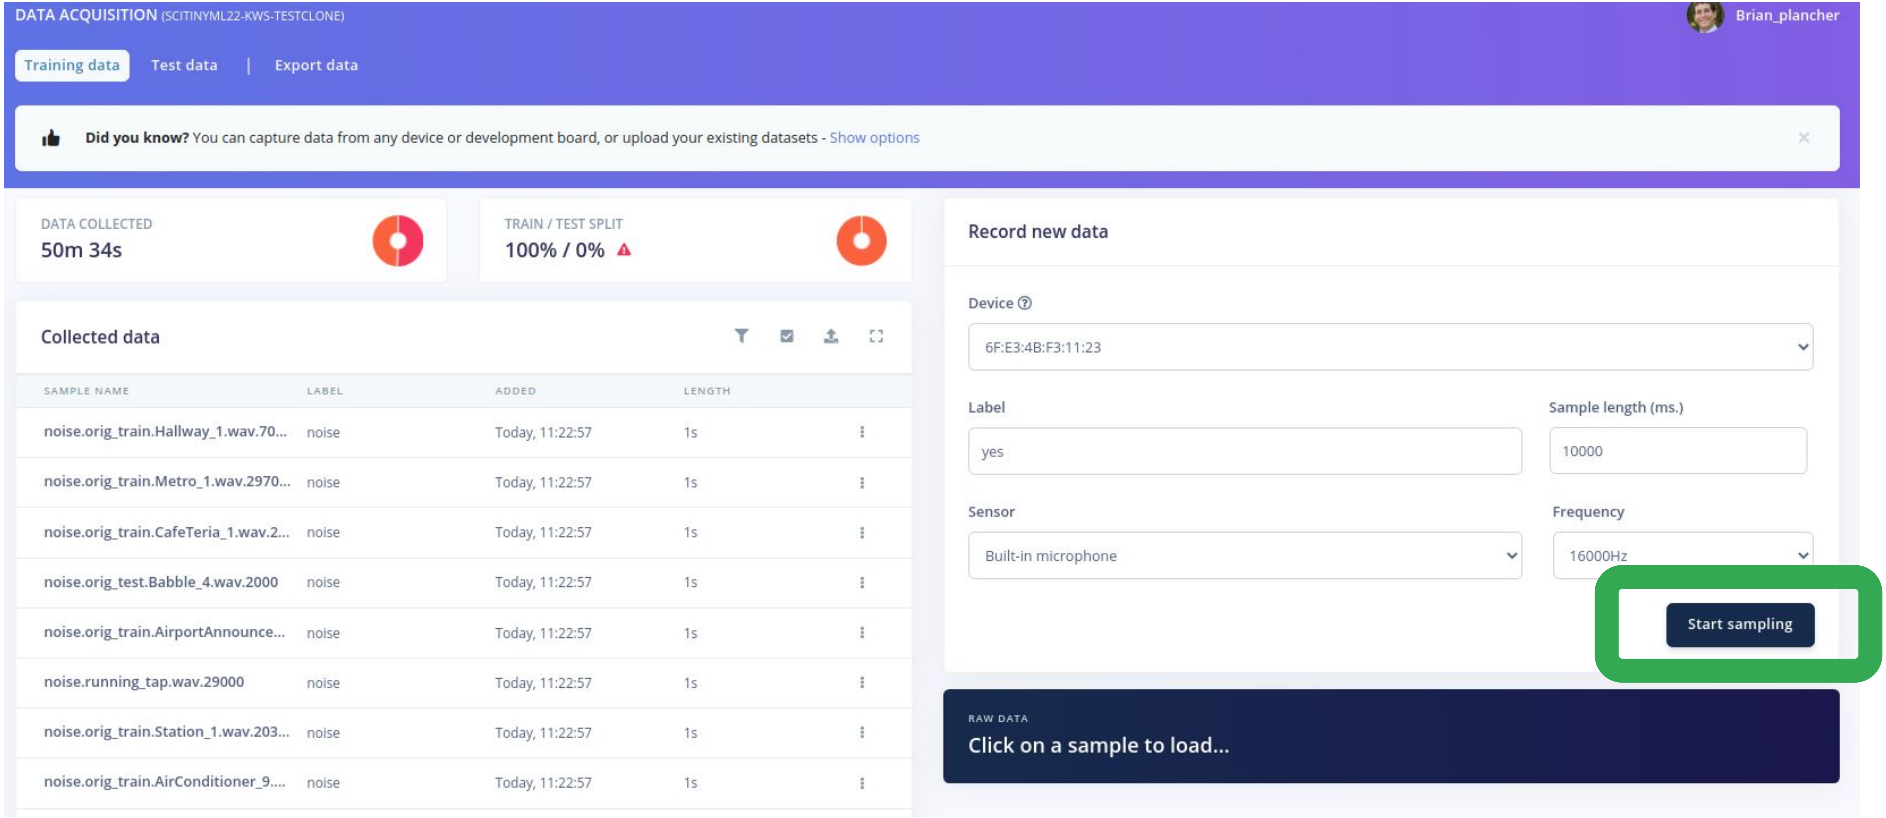
DATA ACQUISITION (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
Training data
Test data
Export data
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
DATA COLLECTED
50m 34s
TRAIN / TEST SPLIT
100% / 0%
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
noise.orig_train.Hallway_1.wav.70...
noise
Today, 11:22:57
1s
noise.orig_train.Metro_1.wav.2970...
noise.orig_train.CafeTeria_1.wav.2...
noise.orig_test.Babble_4.wav.2000
noise.orig_train.AirportAnnounce...
noise.running_tap.wav.29000
noise.orig_train.Station_1.wav.203...
noise.orig_train.AirConditioner_9....
Record new data
Device
6F:E3:4B:F3:11:23
Label
yes
Sample length (ms.)
10000
Sensor
Built-in microphone
Frequency
16000Hz
Start sampling
RAW DATA
Click on a sample to load...

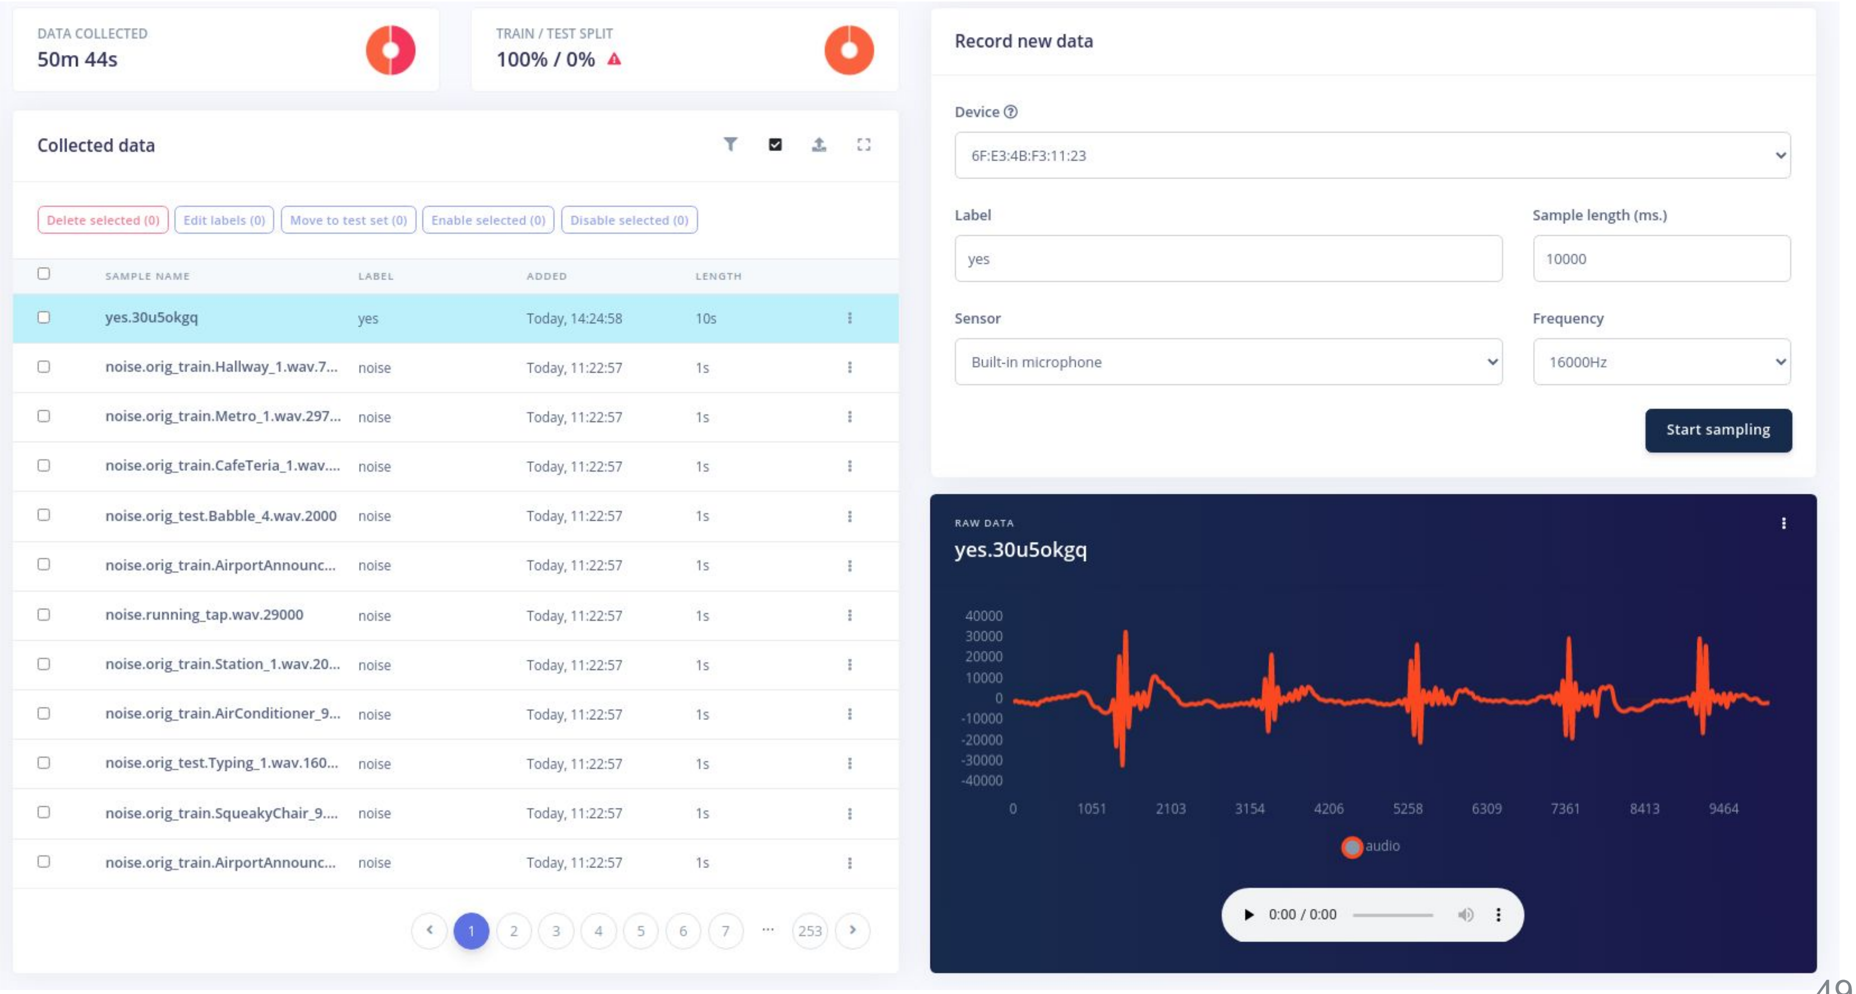
DATA COLLECTED
50m 44s
TRAIN / TEST SPLIT
100% / 0%
Collected data
Delete selected (0)
Edit labels (0)
Move to test set (0)
Enable selected (0)
Disable selected (0)
SAMPLE NAME
LABEL
ADDED
LENGTH
yes.30u5okgq yes Today, 14:24:58 10s
noise.orig_train.Hallway_1.wav.7... noise Today, 11:22:57 1s
noise.orig_train.Metro_1.wav.297... noise Today, 11:22:57 1s
noise.orig_train.CafeTeria_1.wav.... noise Today, 11:22:57 1s
noise.orig_test.Babble_4.wav.2000 noise Today, 11:22:57 1s
noise.orig_train.AirportAnnounc... noise Today, 11:22:57 1s
noise.running_tap.wav.29000 noise Today, 11:22:57 1s
noise.orig_train.Station_1.wav.20... noise Today, 11:22:57 1s
noise.orig_train.AirConditioner_9... noise Today, 11:22:57 1s
noise.orig_test.Typing_1.wav.160... noise Today, 11:22:57 1s
noise.orig_train.SqueakyChair_9.... noise Today, 11:22:57 1s
noise.orig_train.AirportAnnounc... noise Today, 11:22:57 1s
1 2 3 4 5 6 7 ... 253
Record new data
Device
6F:E3:4B:F3:11:23
Label
yes
Sample length (ms.)
10000
Sensor
Built-in microphone
Frequency
16000Hz
Start sampling
RAW DATA
yes.30u5okgq
40000 30000 20000 10000 0 -10000 -20000 -30000 -40000
0 1051 2103 3154 4206 5258 6309 7361 8413 9464
audio
0:00 / 0:00

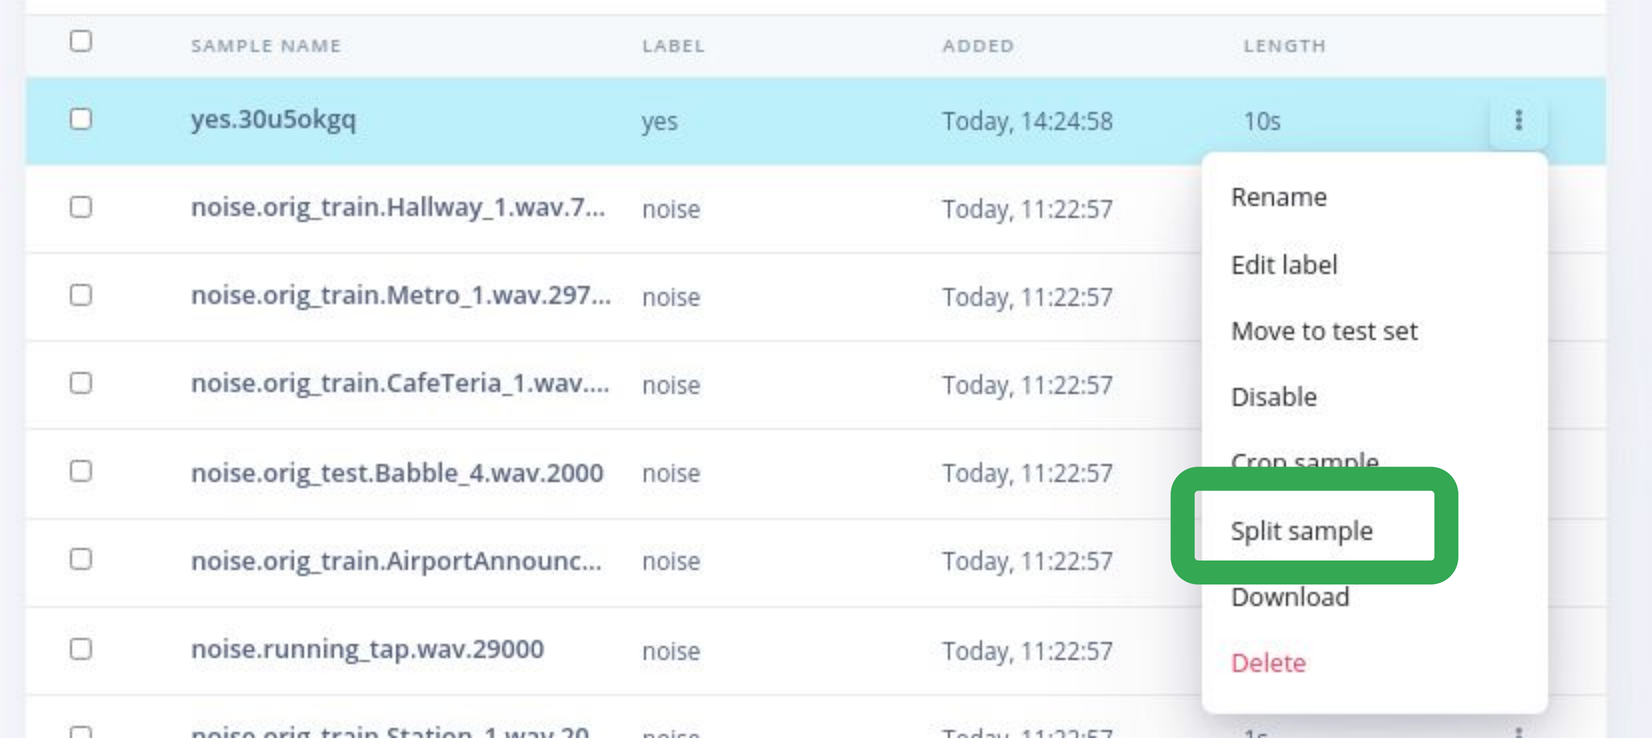
SAMPLE NAME
LABEL
ADDED
LENGTH
yes.30u5okgq
yes
Today, 14:24:58
10s
noise.orig_train.Hallway_1.wav.7...
noise
Today, 11:22:57
noise.orig_train.Metro_1.wav.297...
noise
Today, 11:22:57
noise.orig_train.CafeTeria_1.wav....
noise
Today, 11:22:57
noise.orig_test.Babble_4.wav.2000
noise
Today, 11:22:57
noise.orig_train.AirportAnnounc...
noise
Today, 11:22:57
noise.running_tap.wav.29000
noise
Today, 11:22:57
Rename
Edit label
Move to test set
Disable
Split sample
Download
Delete


Split sample 'yes.30u5okgq'
Zoom
+ Add Segment
Set segment length (ms.):
1000
Apply
40000
30000
20000
10000
0
-10000
-20000
-30000
-40000
0
1013
2027
3041
4054
5068
6082
7096
8109
9123
audio
Cancel
Shift samples
Split


Split sample 'yes.30u5okgq'
Zoom
+ Add Segment
Set segment length (ms.):
1000
Apply
40000
30000
20000
10000
0
-10000
-20000
-30000
-40000
0
1013
2027
3041
4054
5068
6082
7096
8109
9123
audio
Cancel
Shift samples
Split

Split sample 'yes.30u5okgq'
Zoom
+ Add Segment
Set segment length (ms.):
1000
Apply
40000
30000
20000
10000
0
-10000
-20000
-30000
-40000
0
1013
2027
3041
4054
5068
6082
9123
audio
Cancel
Shift samples
Split

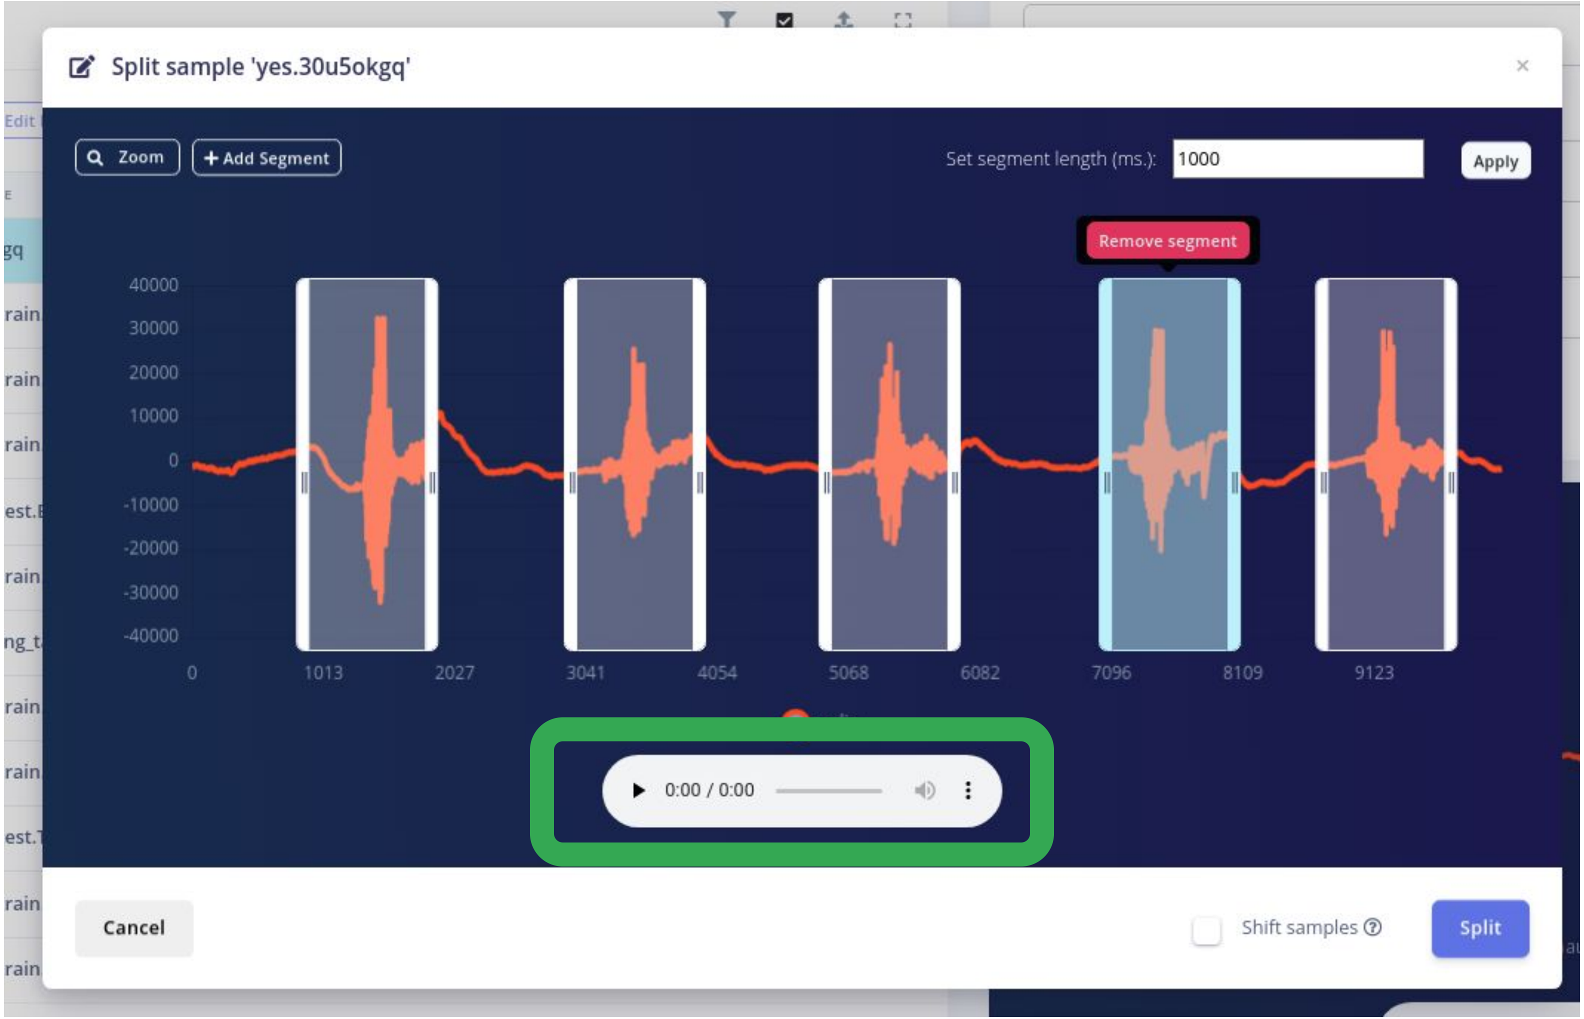
Split sample 'yes.30u5okgq'
Zoom
+ Add Segment
Set segment length (ms.):
1000
Apply
Remove segment
40000
30000
20000
10000
0
-10000
-20000
-30000
-40000
0
1013
2027
3041
4054
5068
6082
7096
8109
9123
0:00 / 0:00
Cancel
Shift samples
Split

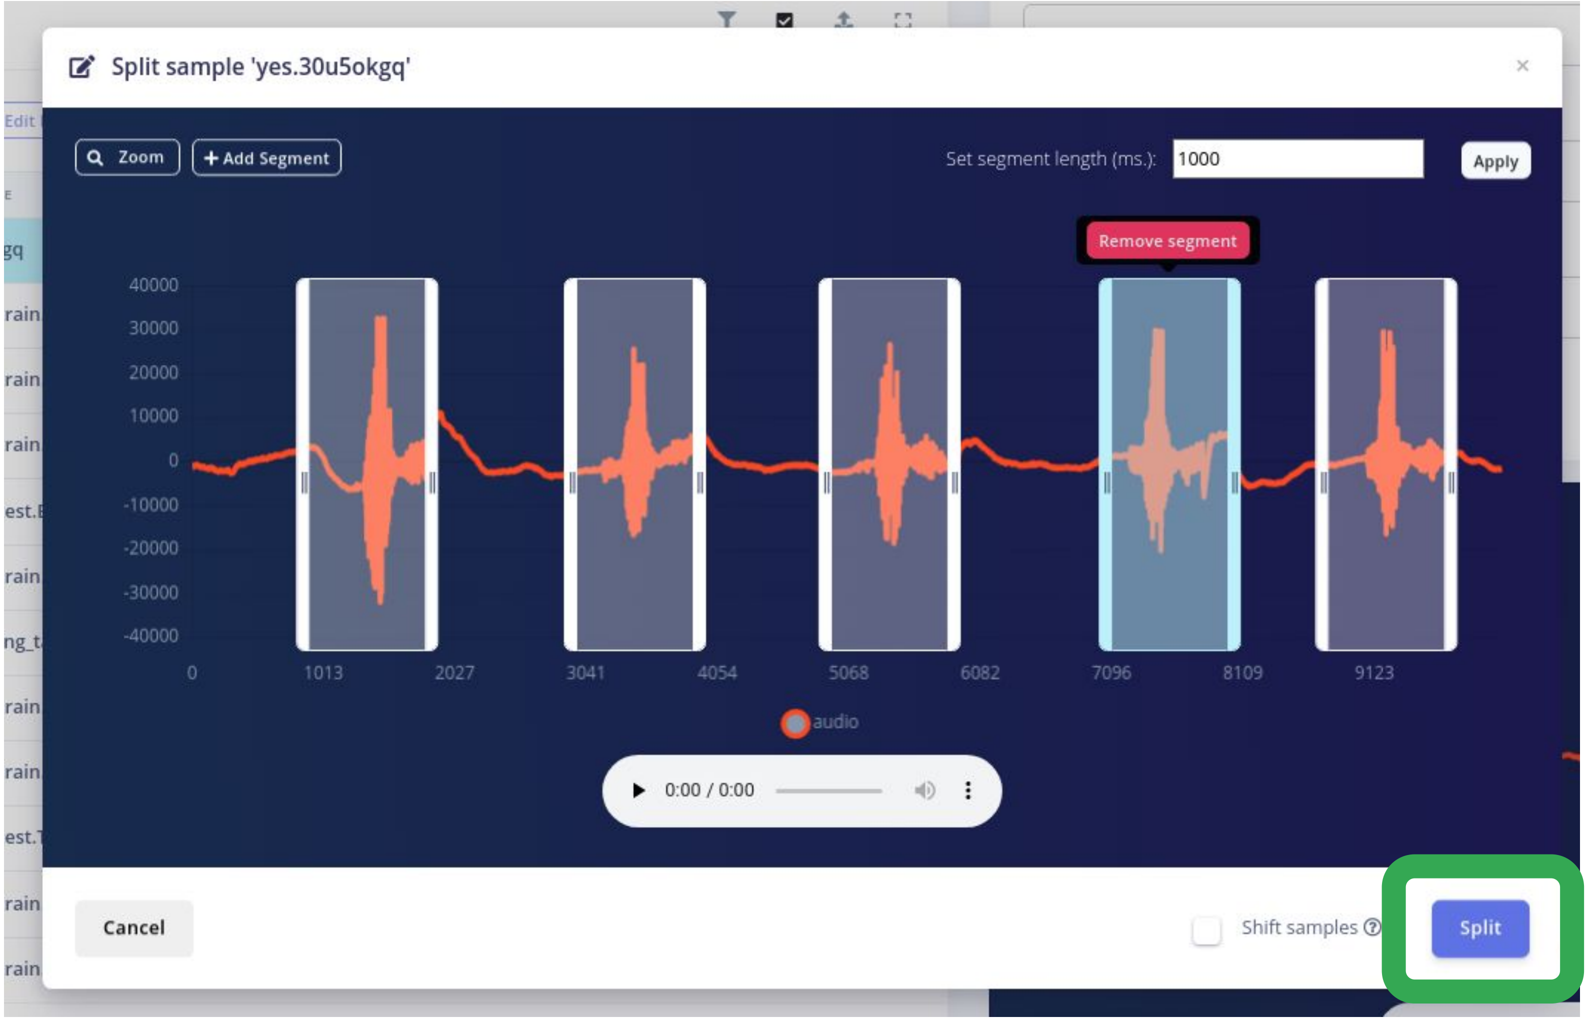

Split sample 'yes.30u5okgq'
Zoom
+ Add Segment
Set segment length (ms.):
1000
Apply
Remove segment
40000
30000
20000
10000
0
-10000
-20000
-30000
-40000
0
1013
2027
3041
4054
5068
6082
7096
8109
9123
audio
0:00 / 0:00
Cancel
Shift samples
Split

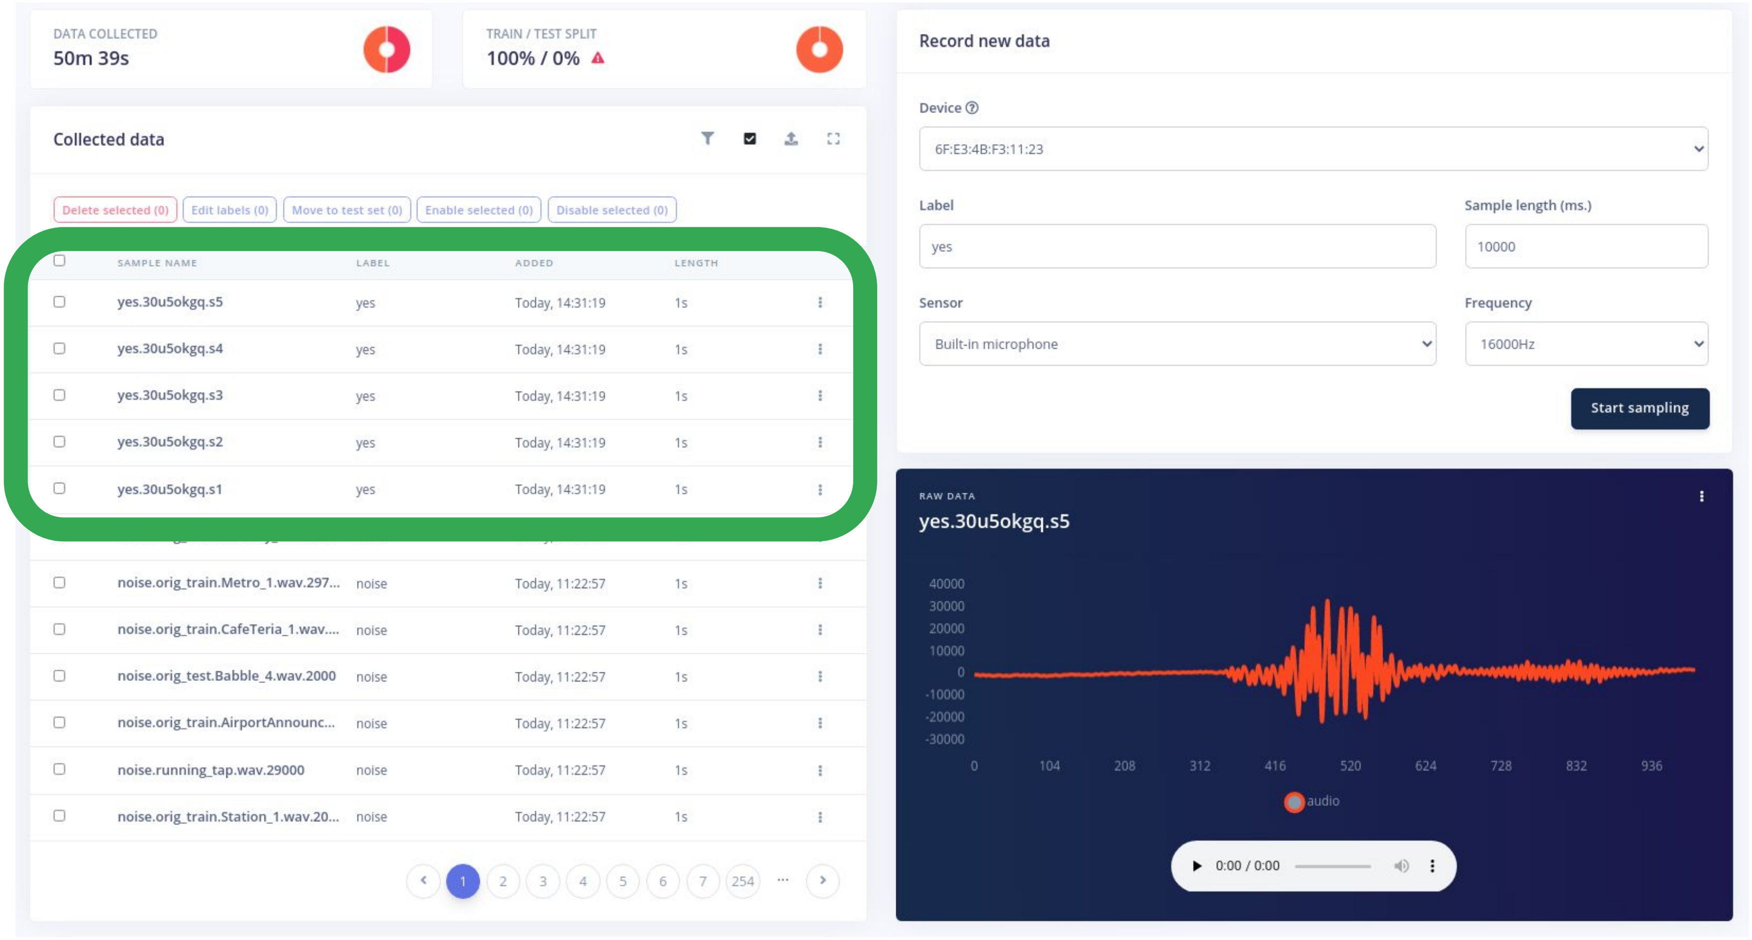
DATA COLLECTED
50m 39s
TRAIN / TEST SPLIT
100% / 0%
Collected data
Delete selected (0)
Edit labels (0)
Move to test set (0)
Enable selected (0)
Disable selected (0)
SAMPLE NAME
LABEL
ADDED
LENGTH
yes.30u5okgq.s5 yes Today, 14:31:19 1s
yes.30u5okgq.s4 yes Today, 14:31:19 1s
yes.30u5okgq.s3 yes Today, 14:31:19 1s
yes.30u5okgq.s2 yes Today, 14:31:19 1s
yes.30u5okgq.s1 yes Today, 14:31:19 1s
noise.orig_train.Metro_1.wav.297... noise Today, 11:22:57 1s
noise.orig_train.CafeTeria_1.wav.... noise Today, 11:22:57 1s
noise.orig_test.Babble_4.wav.2000 noise Today, 11:22:57 1s
noise.orig_train.AirportAnnounc... noise Today, 11:22:57 1s
noise.running_tap.wav.29000 noise Today, 11:22:57 1s
noise.orig_train.Station_1.wav.20... noise Today, 11:22:57 1s
1 2 3 4 5 6 7 254 ...
Record new data
Device
6F:E3:4B:F3:11:23
Label
yes
Sample length (ms.)
10000
Sensor
Built-in microphone
Frequency
16000Hz
Start sampling
RAW DATA
yes.30u5okgq.s5
40000
30000
20000
10000
0
-10000
-20000
-30000
0
104
208
312
416
520
624
728
832
936
audio
0:00 / 0:00

Activity: Create a Keyword Spotting Dataset

Collect **~30 samples each** of the following classes of data:

- **Keyword #1 "yes"** (label: yes) (length: 1 seconds)
- **Keyword #2 "no"** (label: no) (length: 1 seconds)

We'll resume in 10 minutes!

Scroll Down to the Bottom

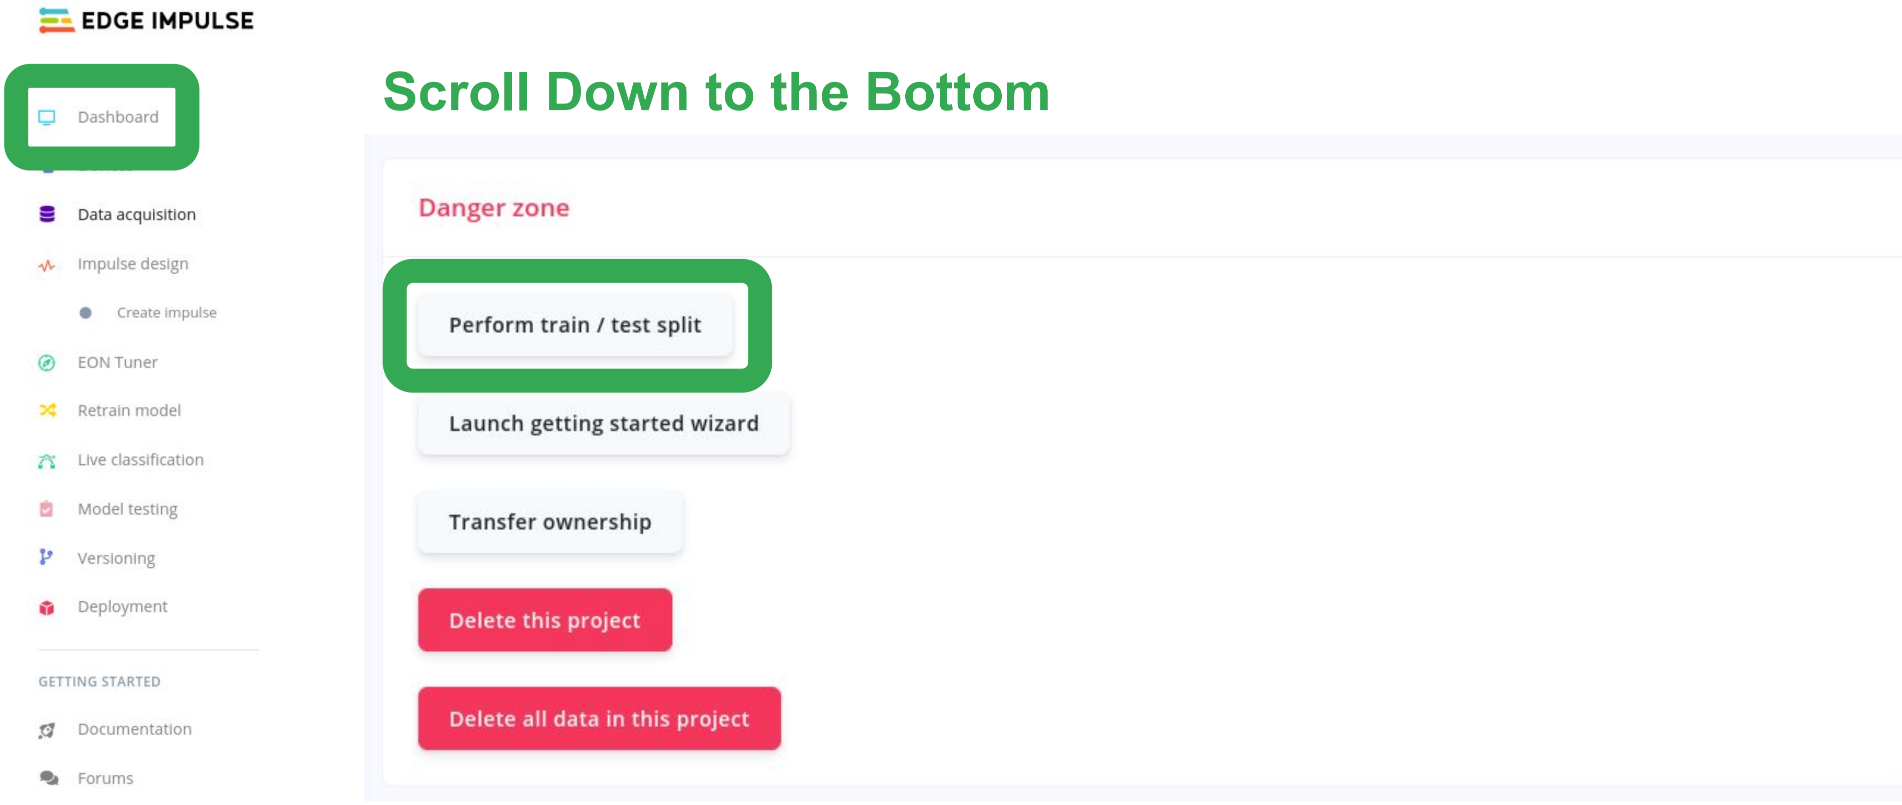

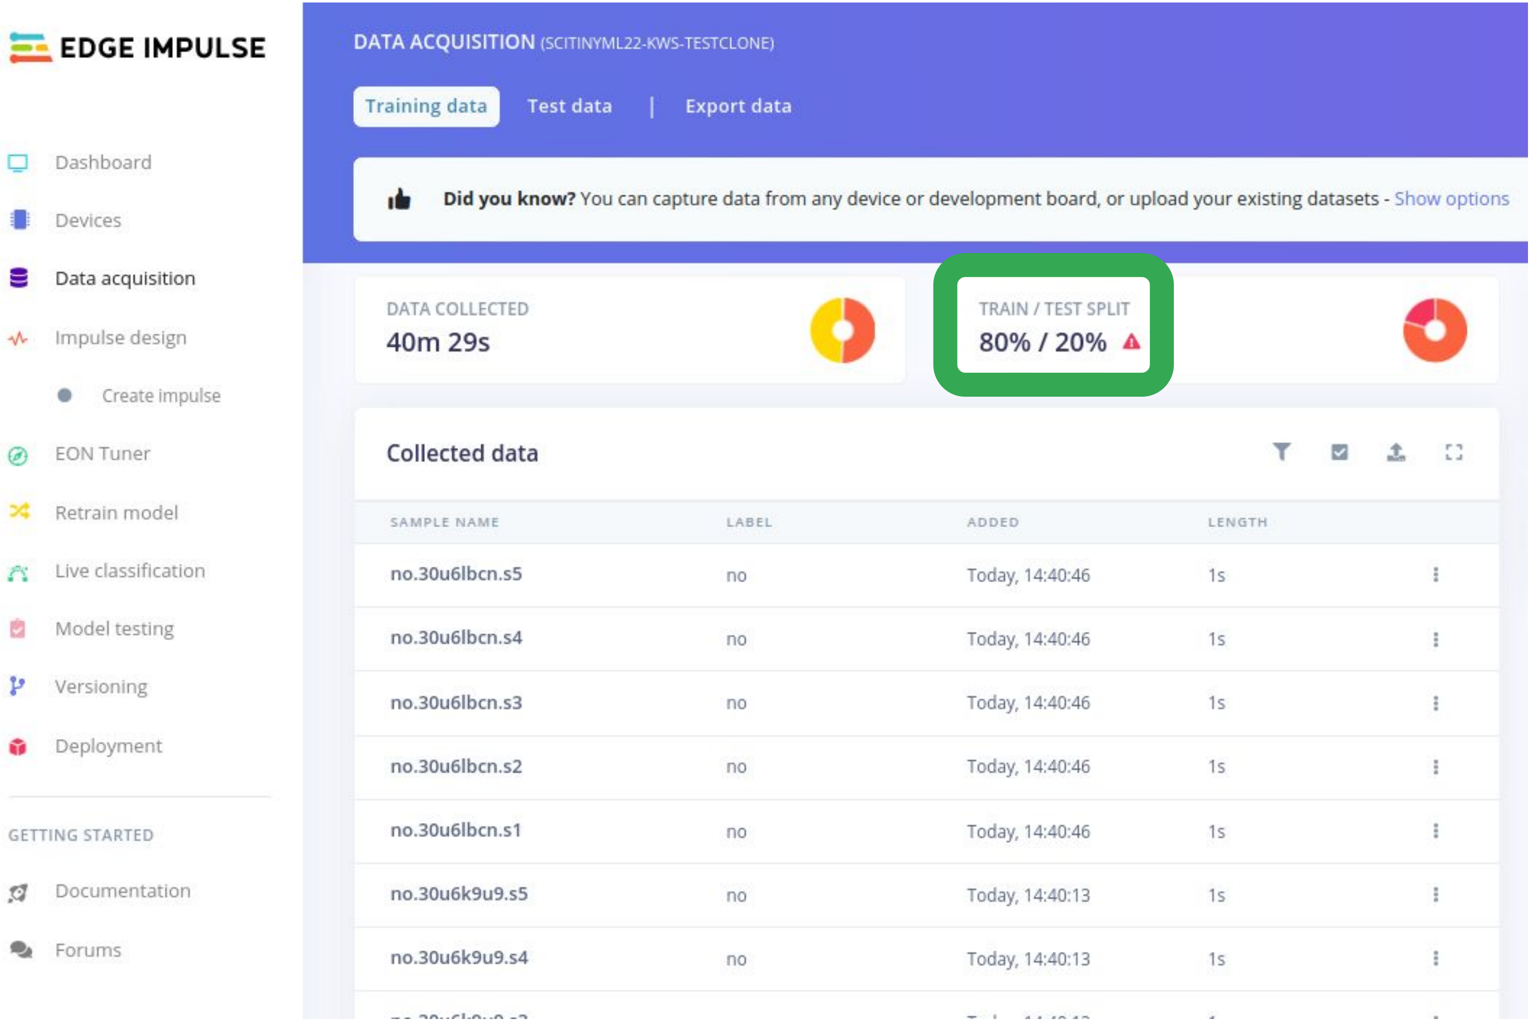
EDGE IMPULSE
DATA ACQUISITION (SCITINYML22-KWS-TESTCLONE)
Training data
Test data
Export data
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Did you know? You can capture data from any device or development board, or upload your existing datasets - Show options
DATA COLLECTED
40m 29s
TRAIN / TEST SPLIT
80% / 20%
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
no.30u6lbcn.s5 no Today, 14:40:46 1s
no.30u6lbcn.s4 no Today, 14:40:46 1s
no.30u6lbcn.s3 no Today, 14:40:46 1s
no.30u6lbcn.s2 no Today, 14:40:46 1s
no.30u6lbcn.s1 no Today, 14:40:46 1s
no.30u6k9u9.s5 no Today, 14:40:13 1s
no.30u6k9u9.s4 no Today, 14:40:13 1s

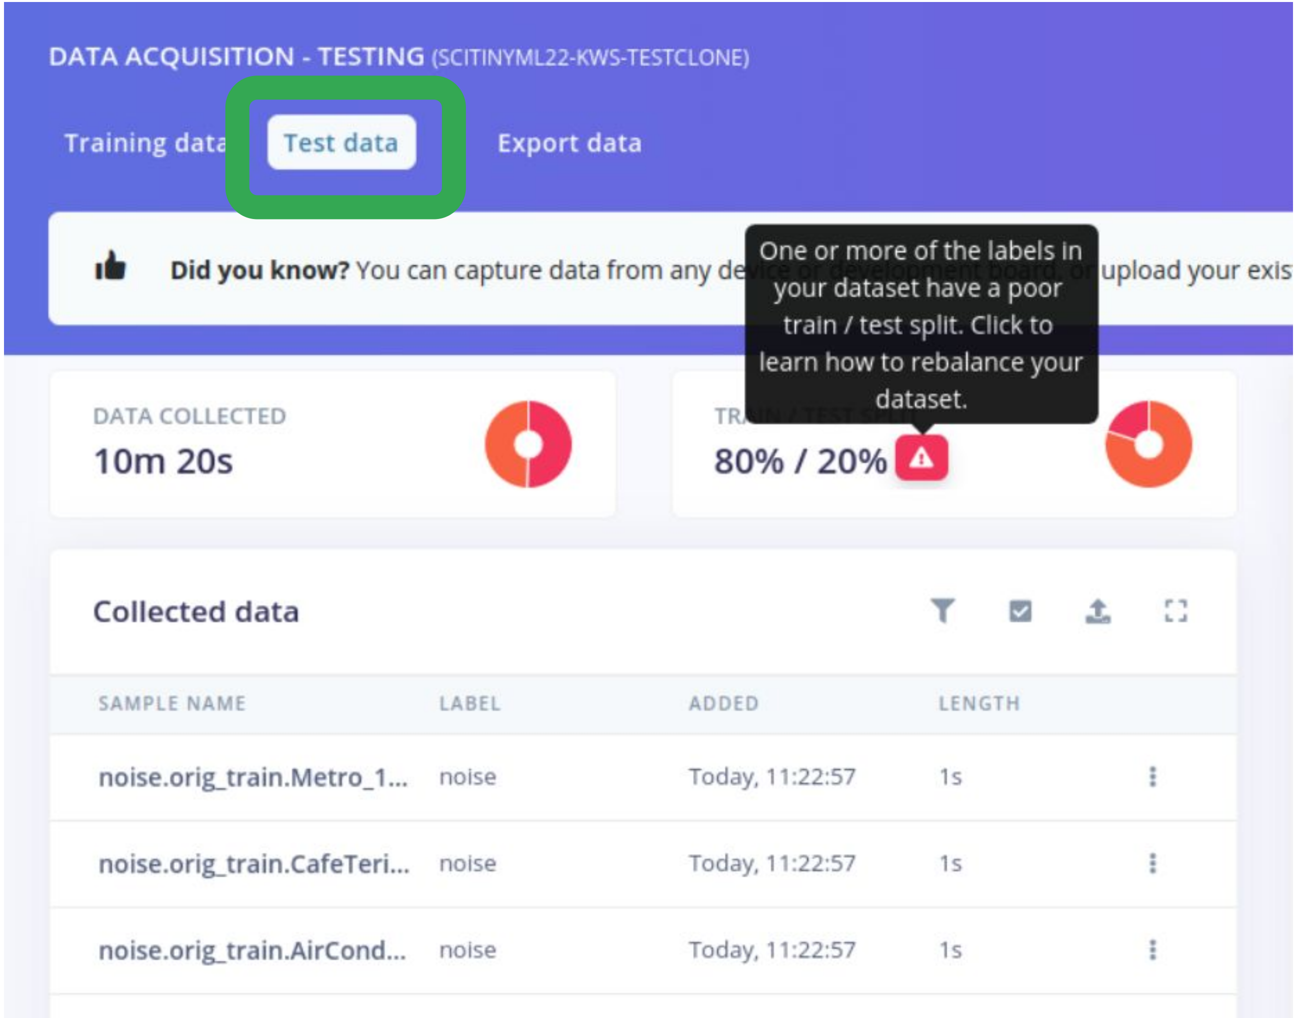
DATA ACQUISITION - TESTING (SCITINYML22-KWS-TESTCLONE)
Training data
Test data
Export data
Did you know? You can capture data from any de
upload your exis
One or more of the labels in your dataset have a poor train / test split. Click to learn how to rebalance your dataset.
DATA COLLECTED
10m 20s
80% / 20%
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
noise.orig_train.Metro_1...
noise
Today, 11:22:57
1s
noise.orig_train.CafeTeri...
noise
Today, 11:22:57
1s
noise.orig_train.AirCond...
noise
Today, 11:22:57
1s

Dataset train / test split ratio

×

Training data is used to train your model, and **testing data** is used to test your model's accuracy after training. We recommend an approximate 80/20 train/test split ratio for your data for every class (or label) in your dataset, although especially large datasets may require less testing data.

SUGGESTED TRAIN / TEST SPLIT 80% / 20%

Labels in your dataset ?

The 'no' class has a poor train/test split ratio. To fix this, add or move samples to the training or testing data.

NO 100% / 0% (27s / 0s)

NOISE 80% / 20% (20m 22s / 5m 13s)

UNKNOWN 80% / 20% (19m 52s / 5m 7s)

YES 81% / 19% (22s / 5s)

Perform train / test split

Use this option to rebalance your data, automatically splitting items between training and testing datasets.
Warning: this action cannot be undone.

Perform train / test split

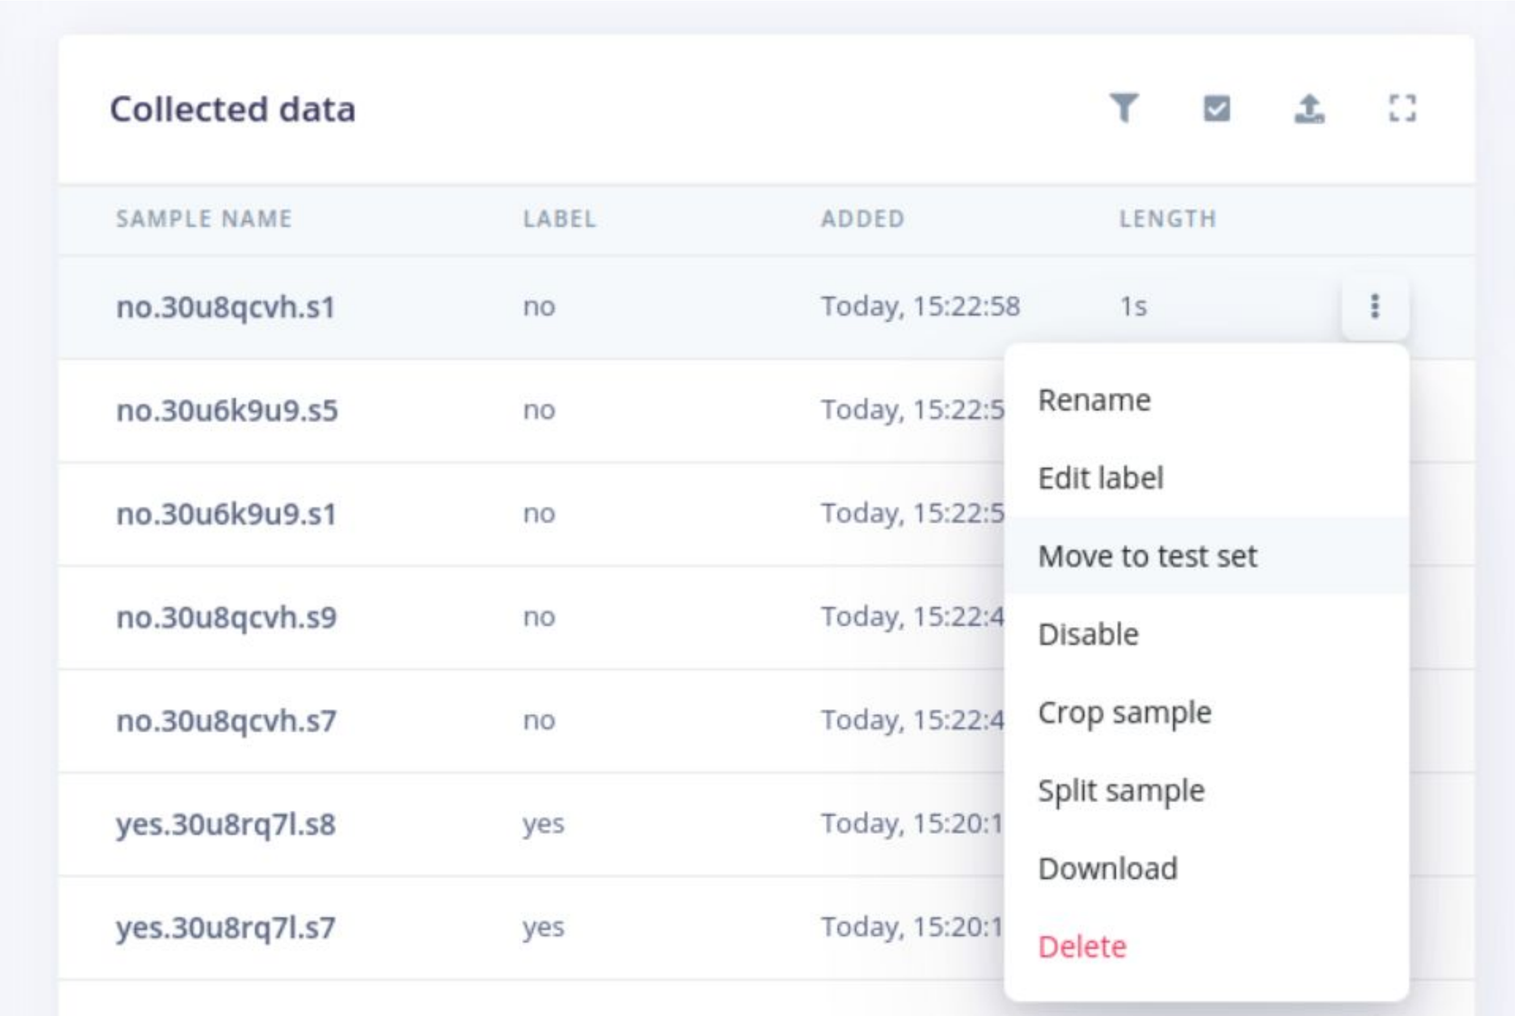
Collected data
SAMPLE NAME
LABEL
ADDED
LENGTH
no.30u8qcvh.s1
no
Today, 15:22:58
1s
no.30u6k9u9.s5
no
Today, 15:22:5
no.30u6k9u9.s1
no
Today, 15:22:5
no.30u8qcvh.s9
no
Today, 15:22:4
no.30u8qcvh.s7
no
Today, 15:22:4
yes.30u8rq7l.s8
yes
Today, 15:20:1
yes.30u8rq7l.s7
yes
Today, 15:20:1
Rename
Edit label
Move to test set
Disable
Crop sample
Split sample
Download
Delete

Dataset train / test split ratio

Training data is used to train your model, and **testing data** is used to test your model's accuracy after training. We recommend an approximate 80/20 train/test split ratio for your data for every class (or label) in your dataset, although especially large datasets may require less testing data.

SUGGESTED TRAIN / TEST SPLIT — 80% / 20%

Labels in your dataset

The 'no' class has a poor train/test split ratio. To fix this, add or move samples to the training or testing data.

NO — 81% / 19% (22s / 5s)

NOISE — 80% / 20% (20m 22s / 5m 13s)

UNKNOWN — 80% / 20% (19m 52s / 5m 7s)

YES — 81% / 19% (22s / 5s)

Today's Agenda

- A Quick Review of What We've Learned
- Data Engineering for KWS
- Hands-on KWS Data Collection with Edge Impulse
- **(Hands-on) Data Preprocessing for KWS**
- Deploying our Model onto our Arduino
- Summary

Edge Impulse Project Dashboard

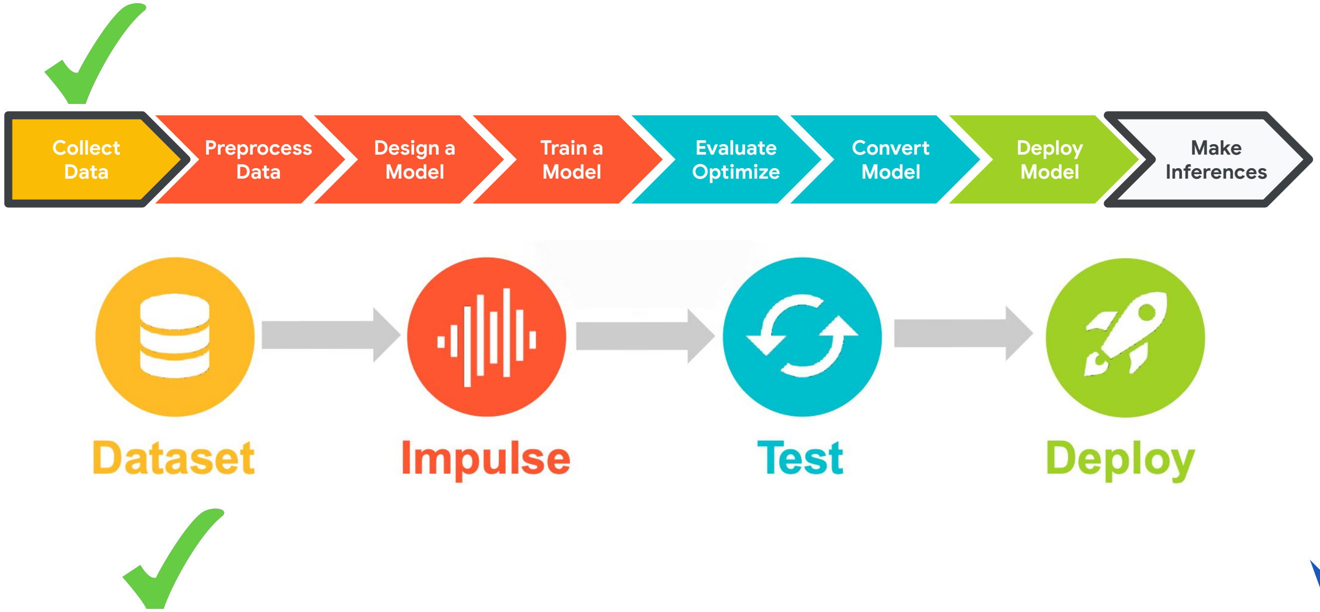

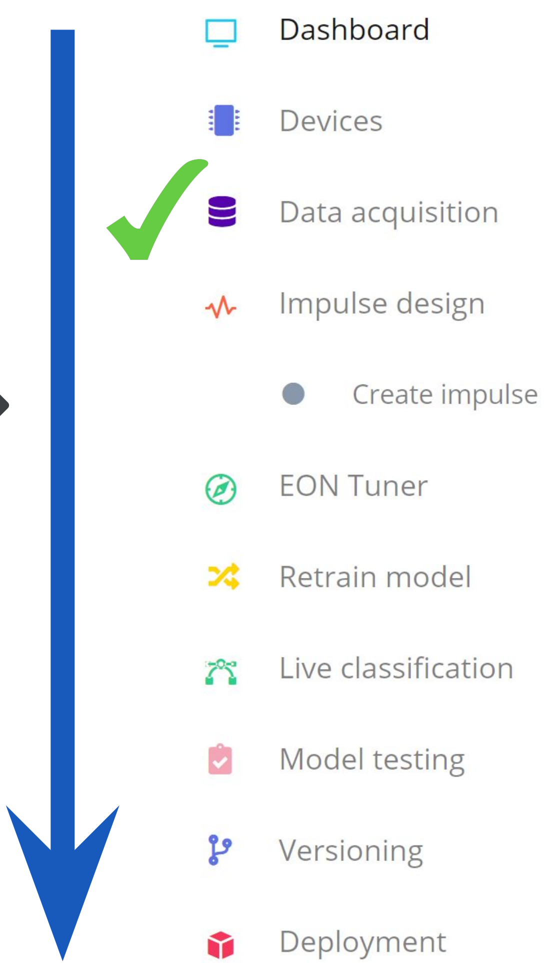

Why might we want to **preprocess** data and not send the raw data to the neural network?

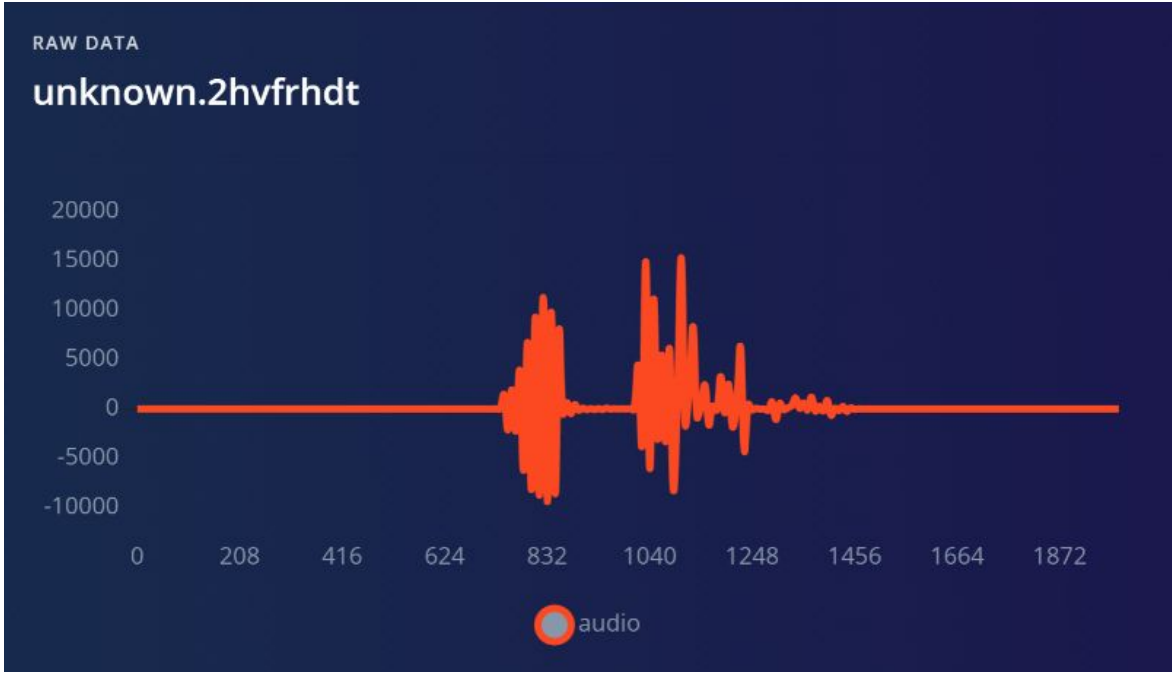

Can you tell these two signals apart?

“Yes” *(spoken loudly)*

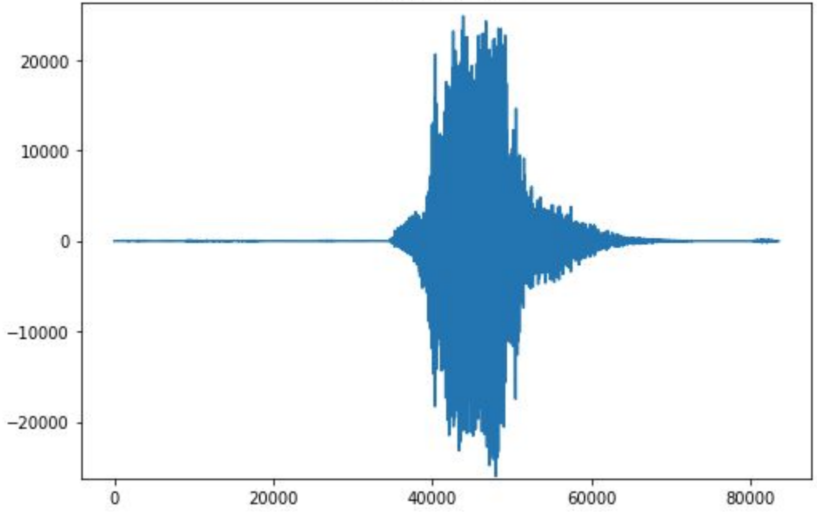

“No” *(spoken loudly)*

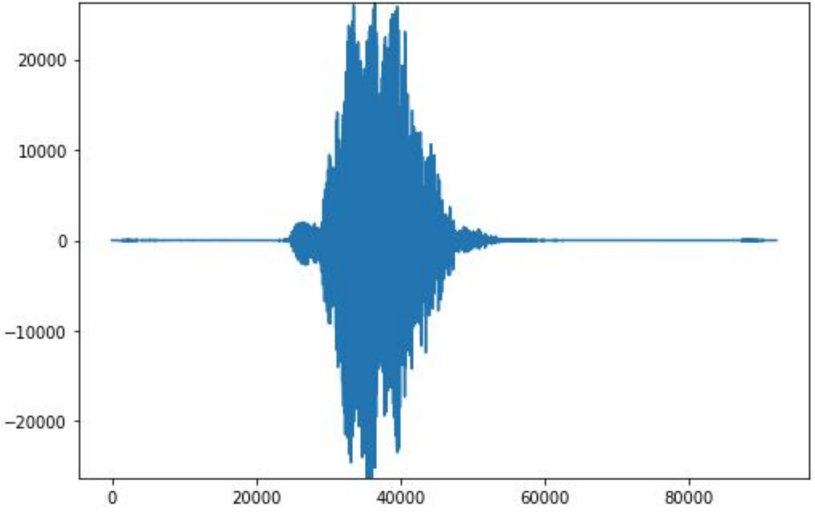

Signal **Components**?

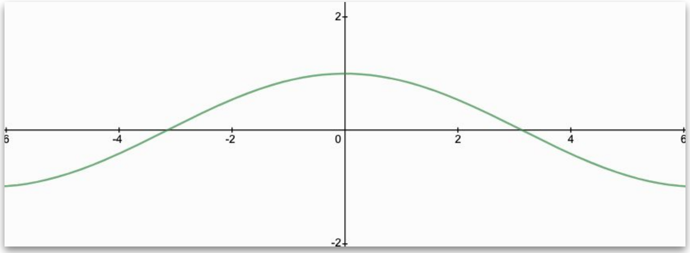

+

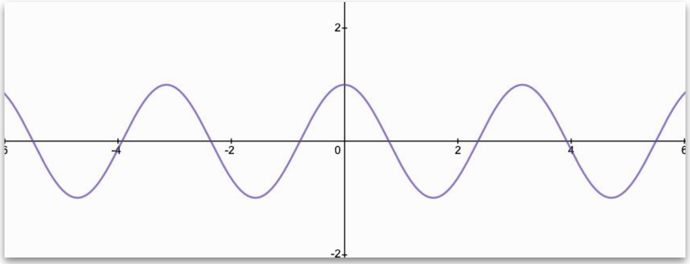

=

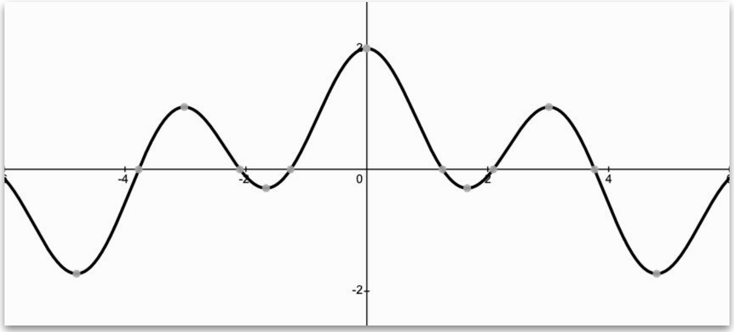

Signal **Components**?

?

+

?

=

"No" *(spoken loudly)*

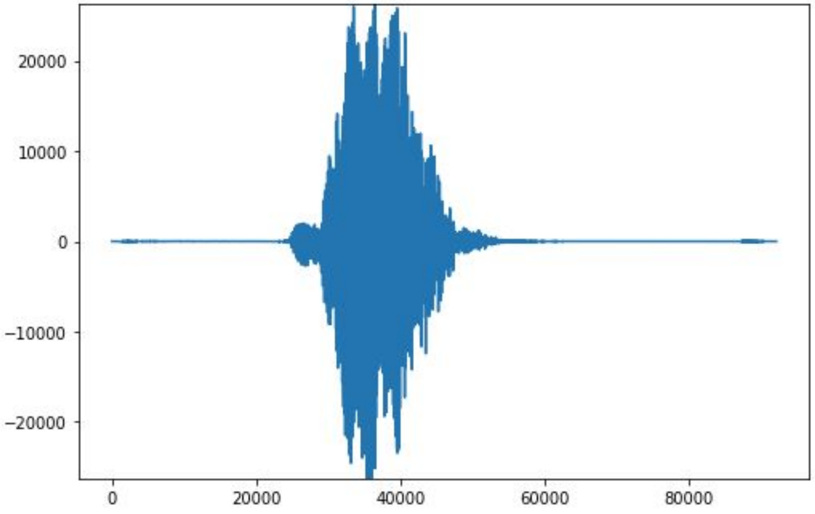

Fast Fourier Transform:
extract the frequencies from a signal

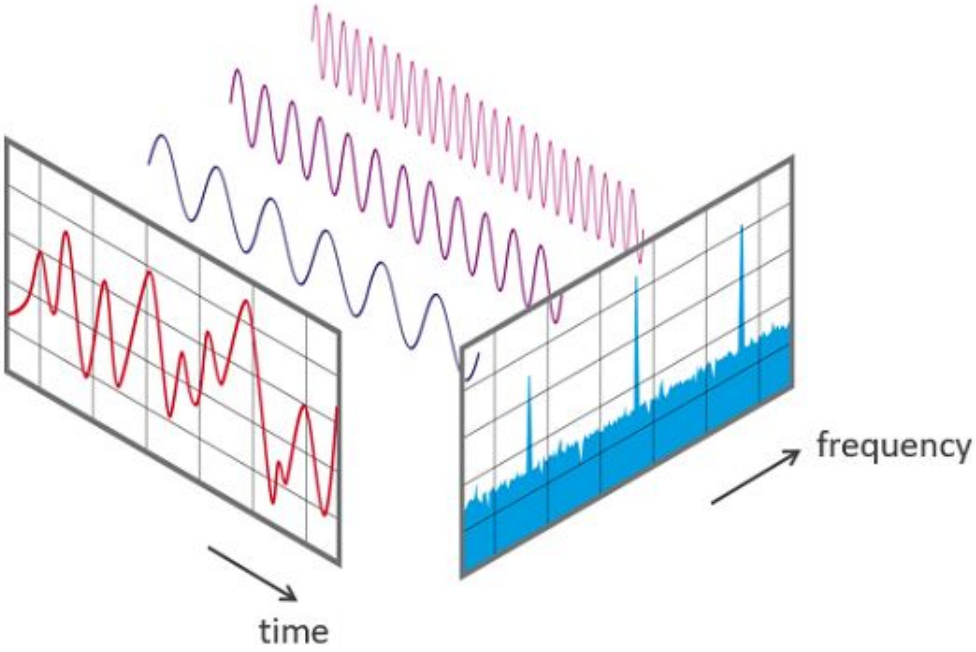

Fast Fourier Transform

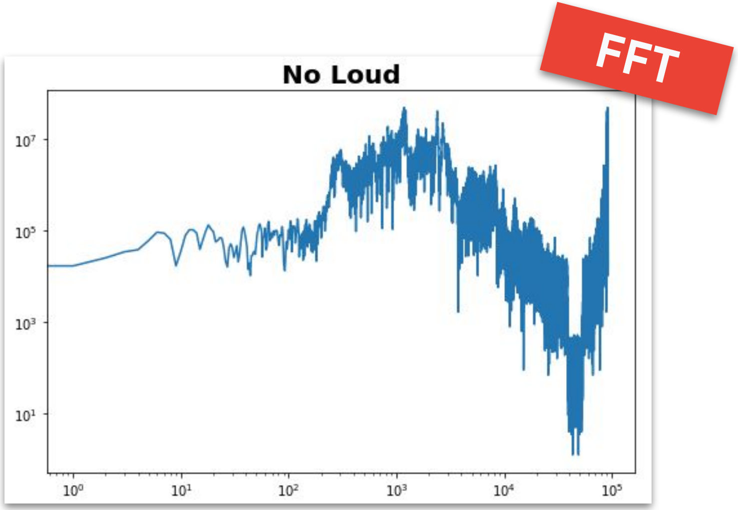

Frequency

=

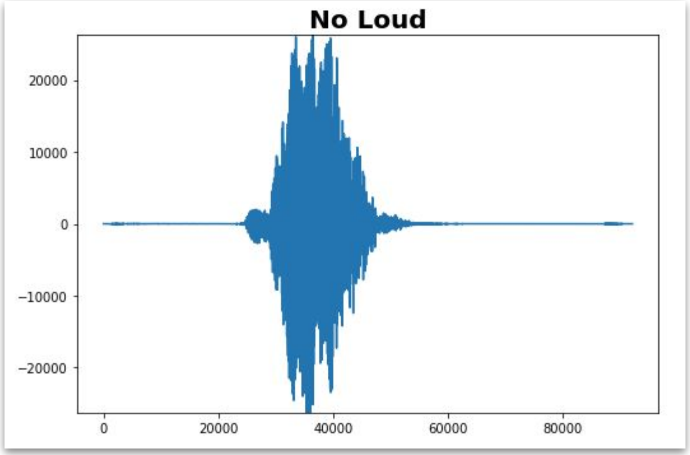

Time

Building a **Spectrogram** using FFTs

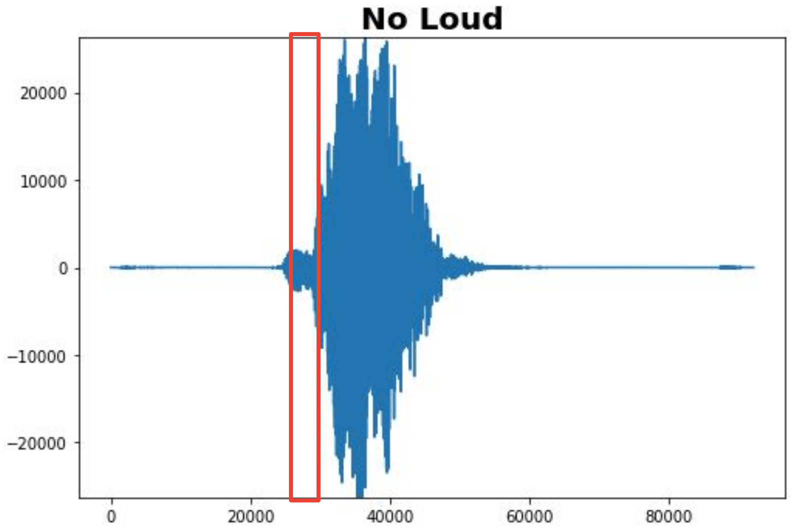

Building a **Spectrogram** using FFTs

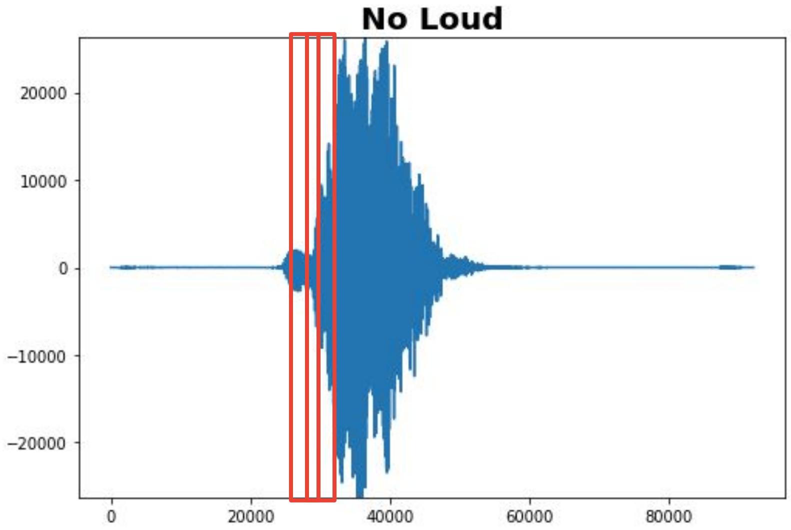

Building a **Spectrogram** using FFTs

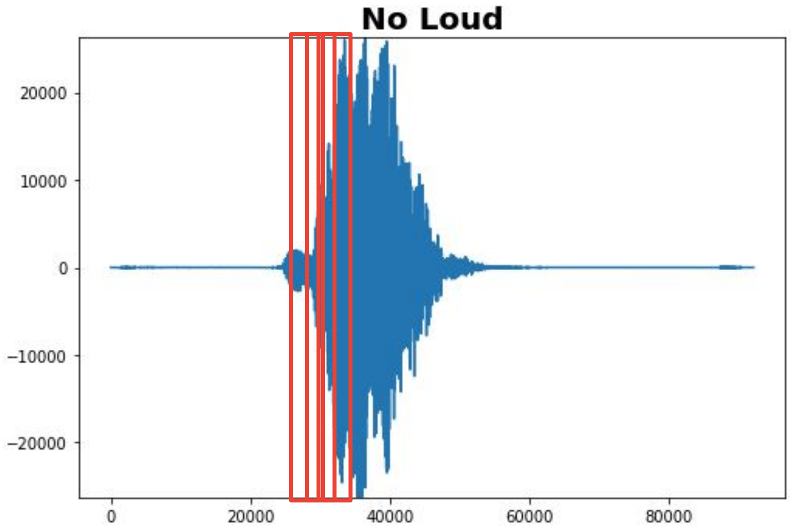

Building a **Spectrogram** using FFTs

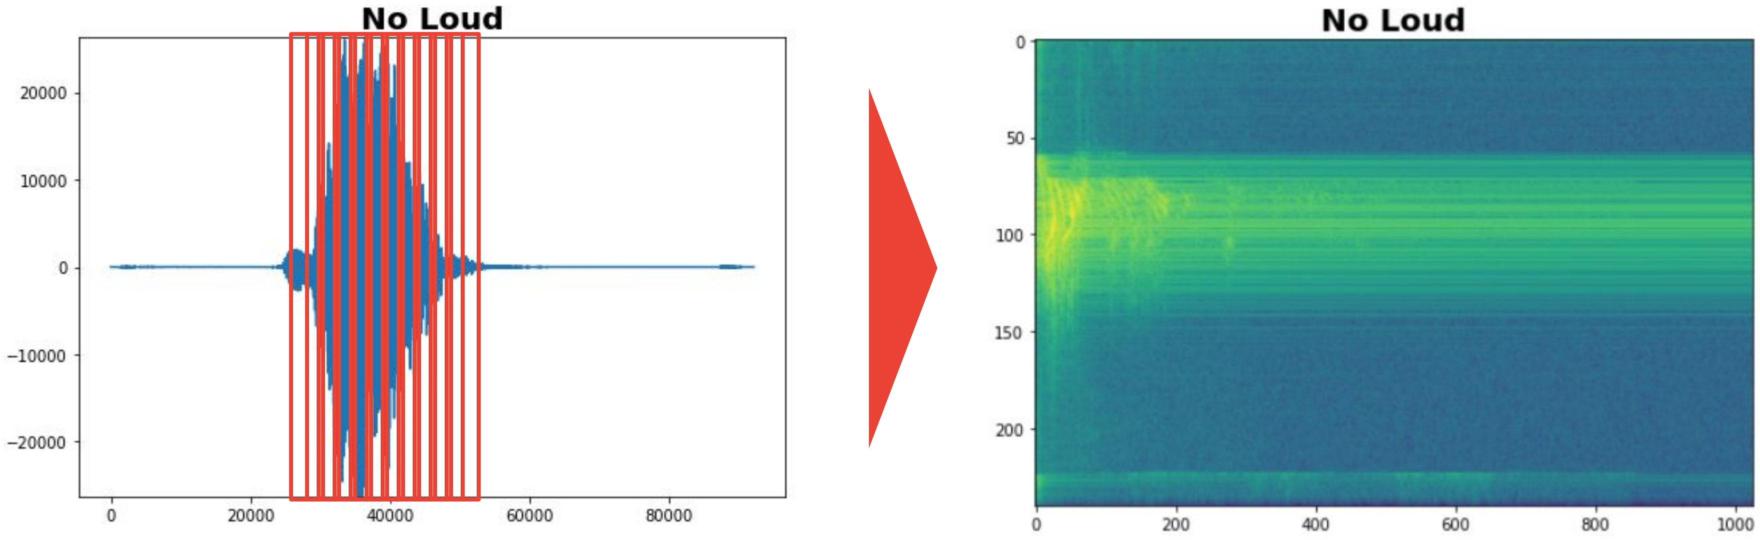

Essentially if you **stack up all the FFTs in a row** then you get the **Spectrogram** (time vs. frequency with color indicating intensity)

Spectrograms help differentiate the data

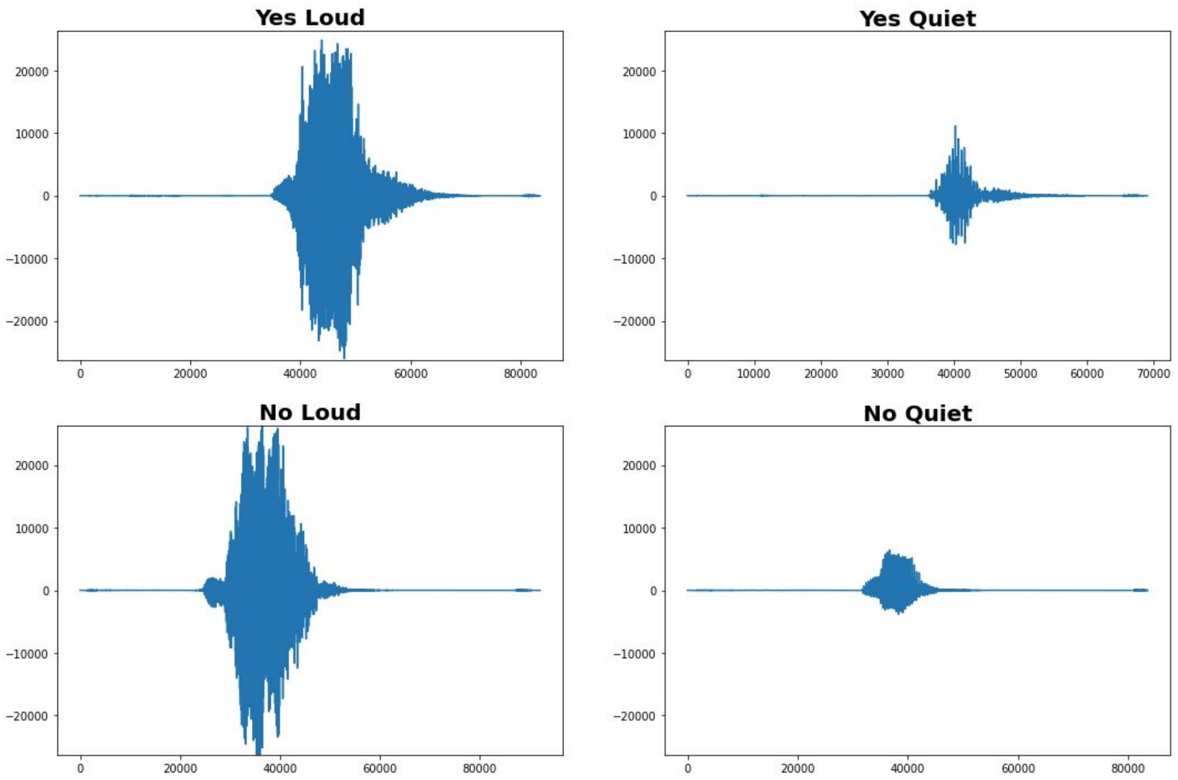

Spectrograms help differentiate the data

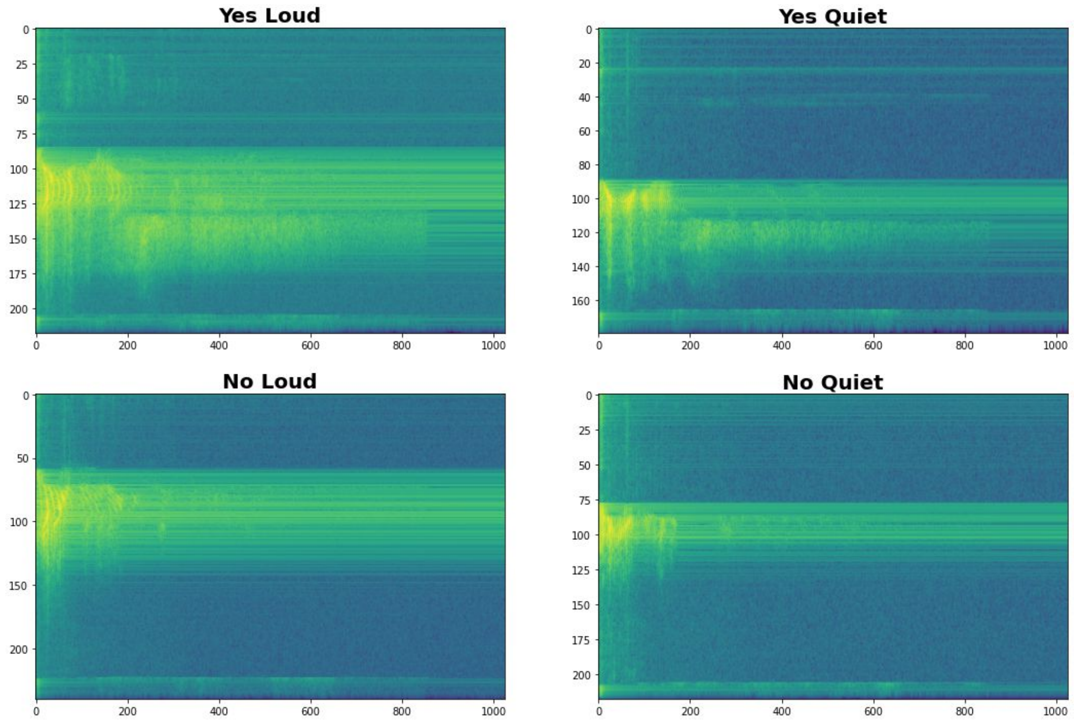

Spectrograms help differentiate the data

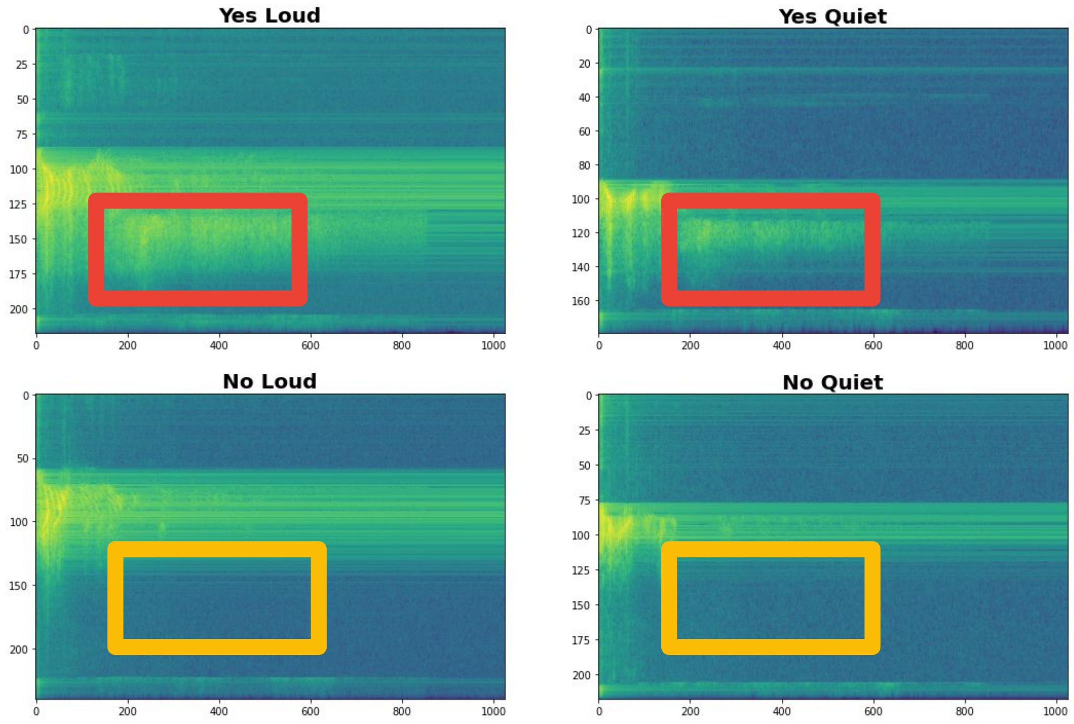

Data Preprocessing: **Spectrograms**

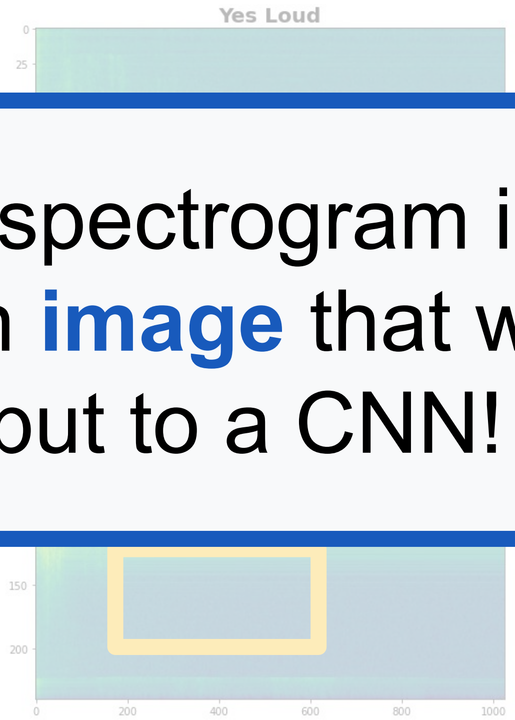

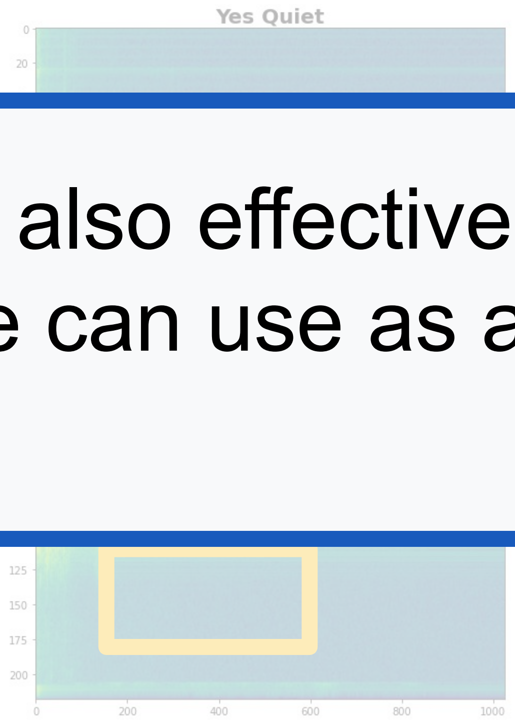

A spectrogram is also effectively an **image** that we can use as an input to a CNN!

Can we do **better** than a spectrogram?

Can we take **domain knowledge** into account?

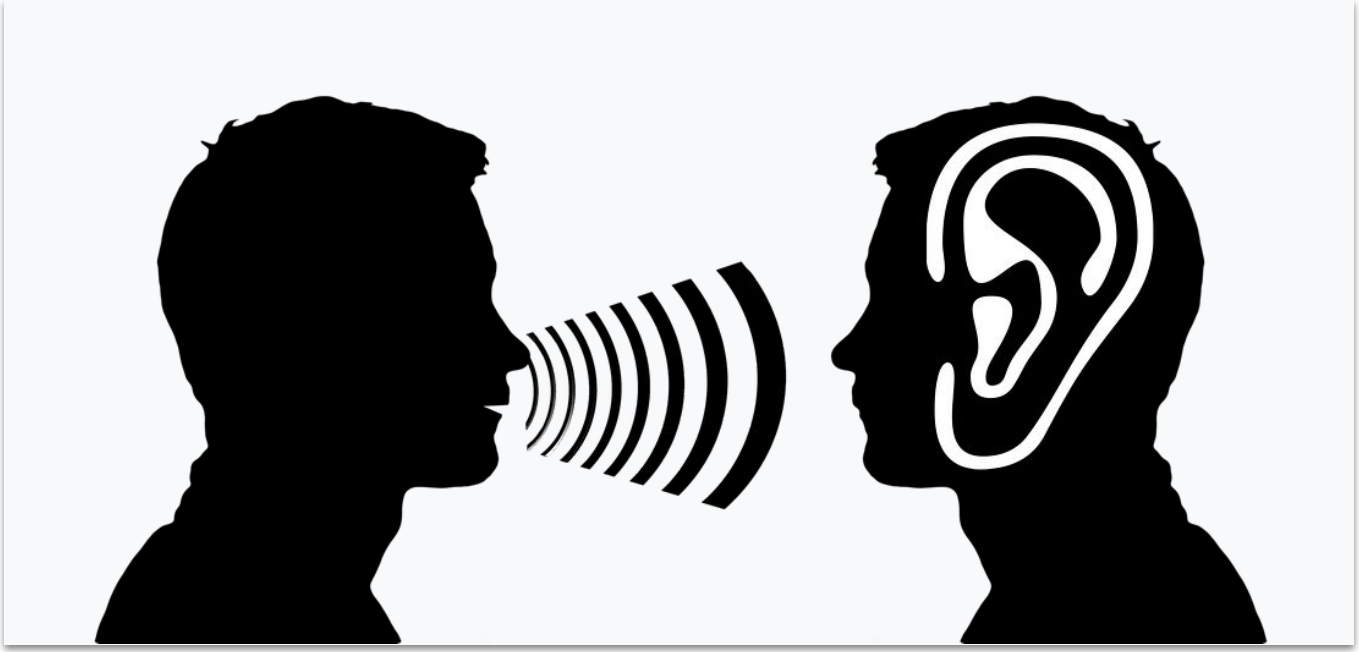

Mel Filterbanks

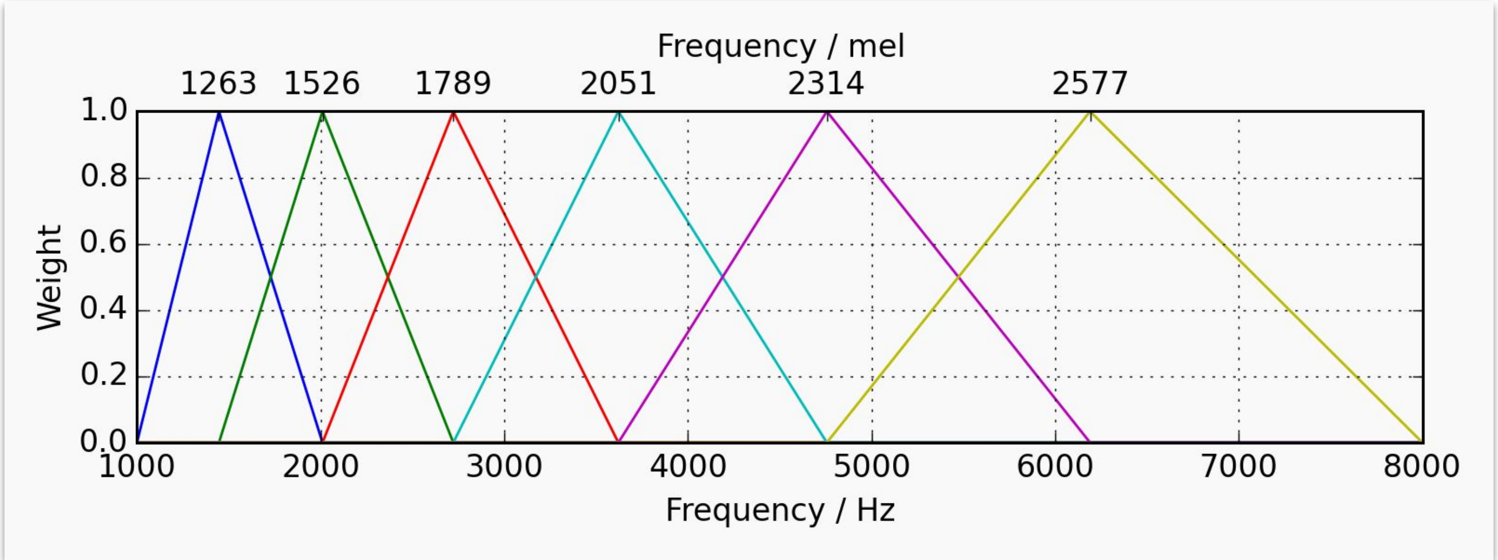

Spectrograms v. MFCCs

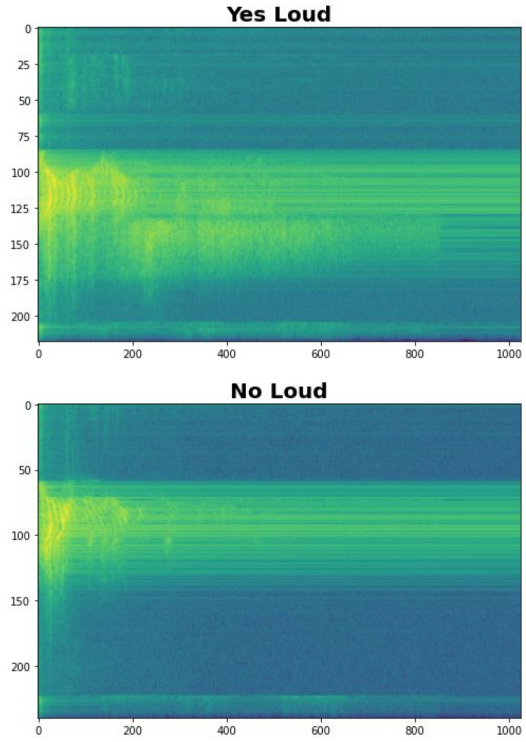

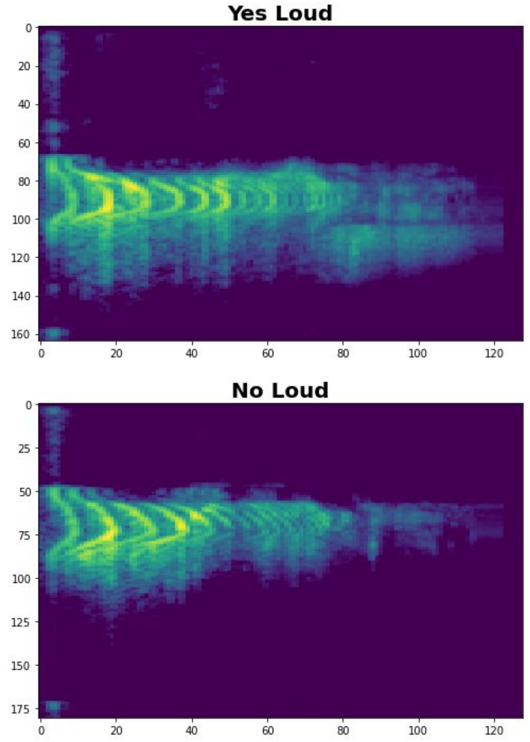

Colab Link

Spectrograms v. MFCCs

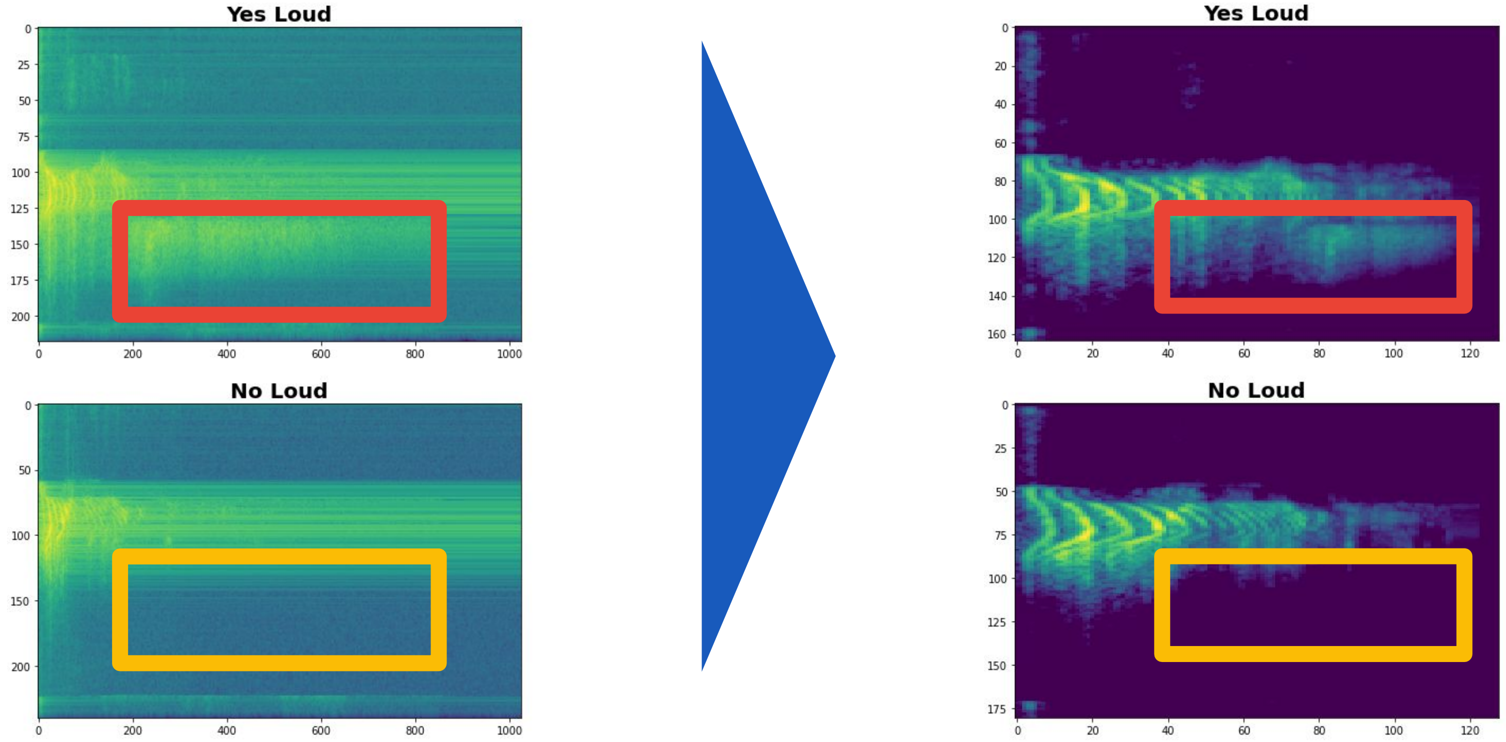

Additional **Feature Engineering**

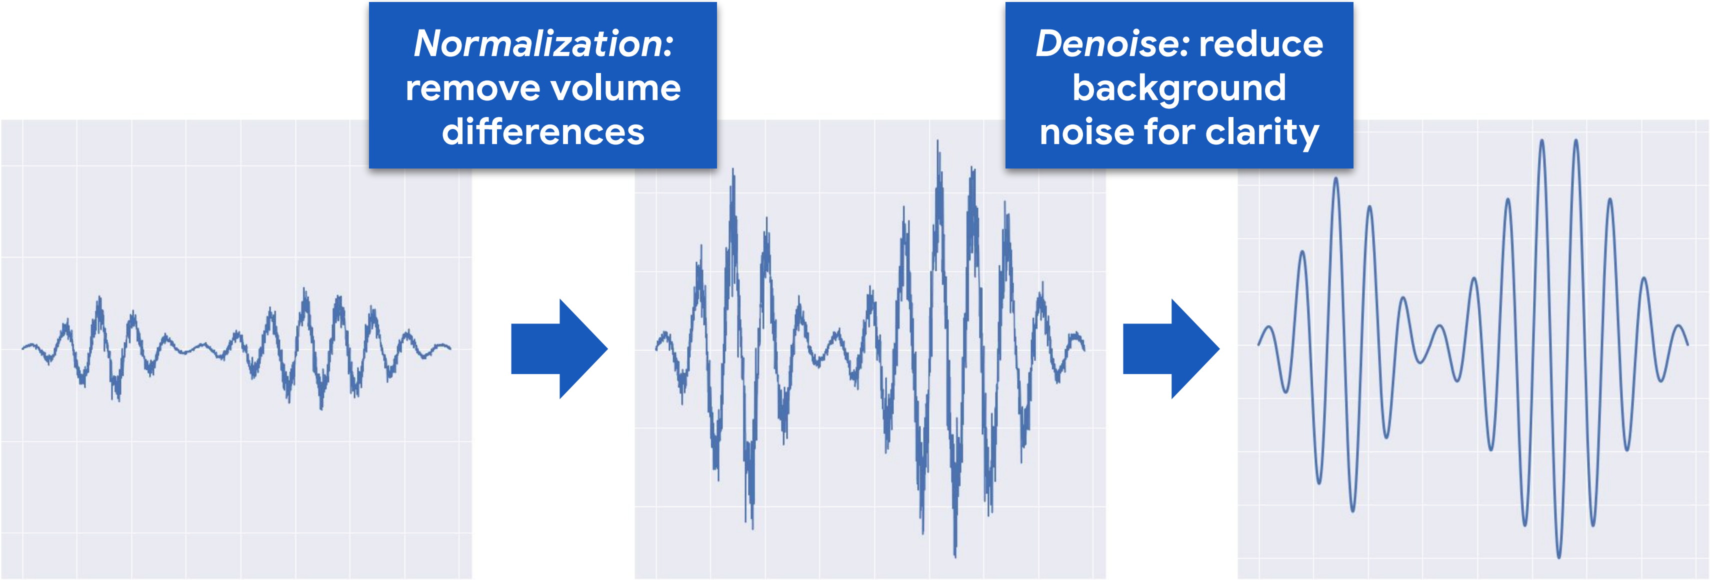

Additional **Feature Engineering**

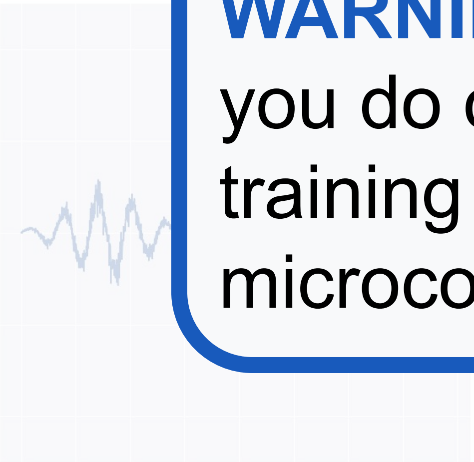
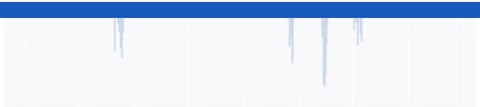


WARNING: Whatever preprocessing you do on the computer in python for training you need to do in C++ on the microcontroller!

Today's Agenda

- Deep ML Background
- Hands-on Computer Vision: Thing Translator
- The Tiny Machine Learning Workflow
- Keyword Spotting (KWS) Data Collection
- KWS Preprocessing and Training
    - Preprocessing (for KWS)
    - **Hands-on Preprocessing and Training with Edge Impulse**
- Deployment Challenges and Opportunities for Embedded ML
- Summary

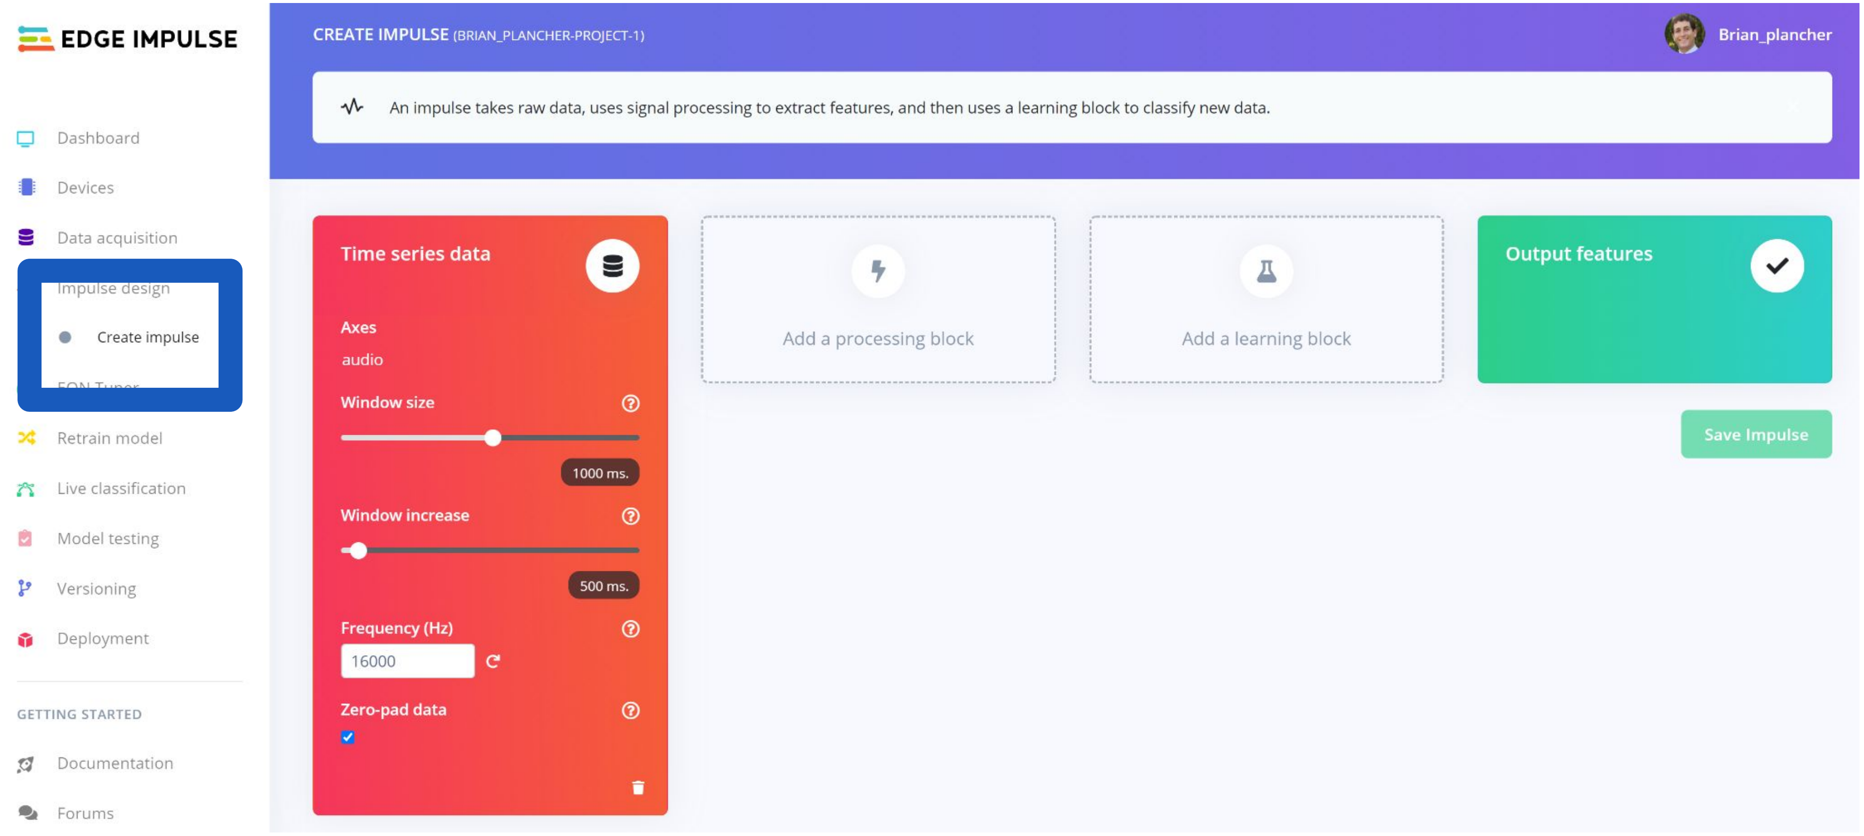
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
CREATE IMPULSE (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Time series data
Axes
audio
Window size
1000 ms.
Window increase
500 ms.
Frequency (Hz)
16000
Zero-pad data
Add a processing block
Add a learning block
Output features
Save Impulse

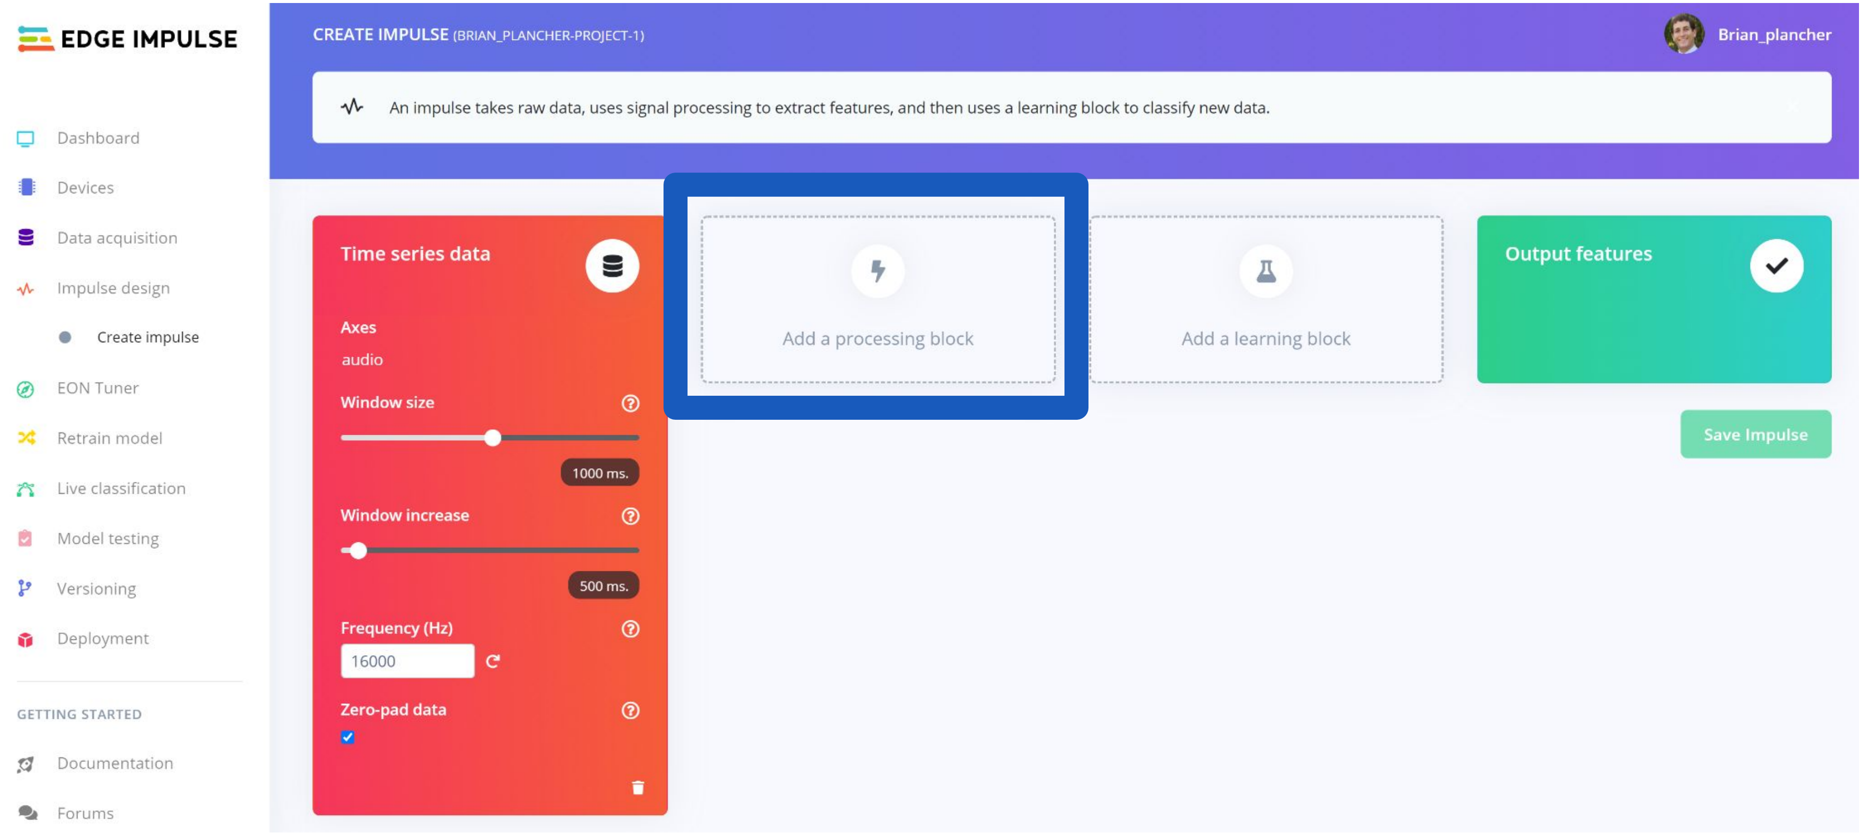
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
CREATE IMPULSE (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Time series data
Axes
audio
Window size
1000 ms.
Window increase
500 ms.
Frequency (Hz)
16000
Zero-pad data
Add a processing block
Add a learning block
Output features
Save Impulse

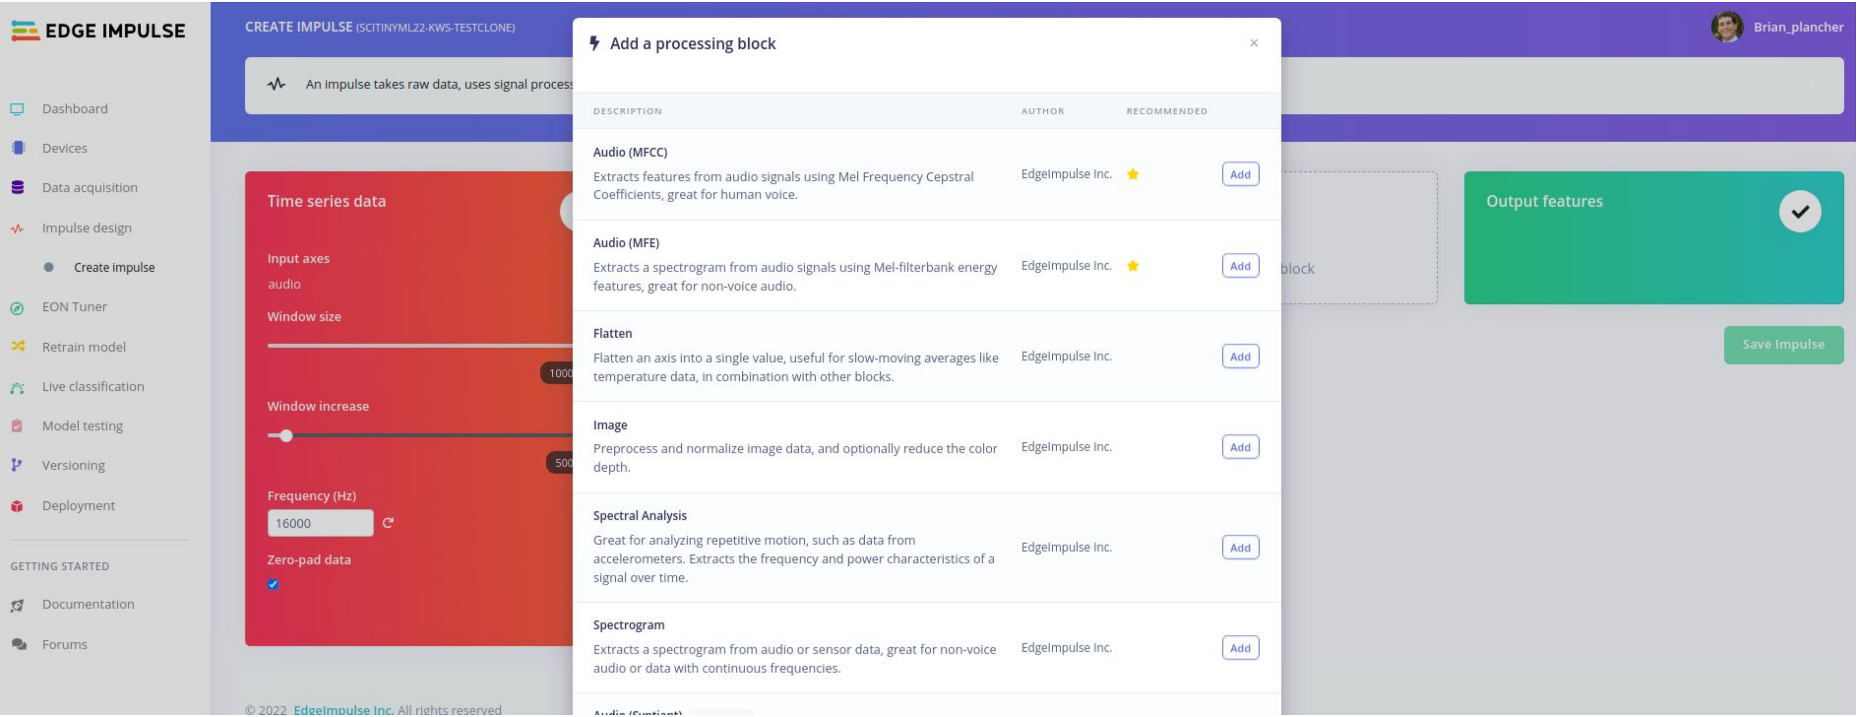
EDGE IMPULSE
CREATE IMPULSE (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
An impulse takes raw data, uses signal process
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Time series data
Input axes
audio
Window size
Window increase
Frequency (Hz)
16000
Zero-pad data
Add a processing block
DESCRIPTION
AUTHOR
RECOMMENDED
Audio (MFCC)
Extracts features from audio signals using Mel Frequency Cepstral Coefficients, great for human voice.
EdgeImpulse Inc.
Add
Audio (MFE)
Extracts a spectrogram from audio signals using Mel-filterbank energy features, great for non-voice audio.
EdgeImpulse Inc.
Add
Flatten
Flatten an axis into a single value, useful for slow-moving averages like temperature data, in combination with other blocks.
EdgeImpulse Inc.
Add
Image
Preprocess and normalize image data, and optionally reduce the color depth.
EdgeImpulse Inc.
Add
Spectral Analysis
Great for analyzing repetitive motion, such as data from accelerometers. Extracts the frequency and power characteristics of a signal over time.
EdgeImpulse Inc.
Add
Spectrogram
Extracts a spectrogram from audio or sensor data, great for non-voice audio or data with continuous frequencies.
EdgeImpulse Inc.
Add
block
Output features
Save Impulse
© 2022 EdgeImpulse Inc. All rights reserved

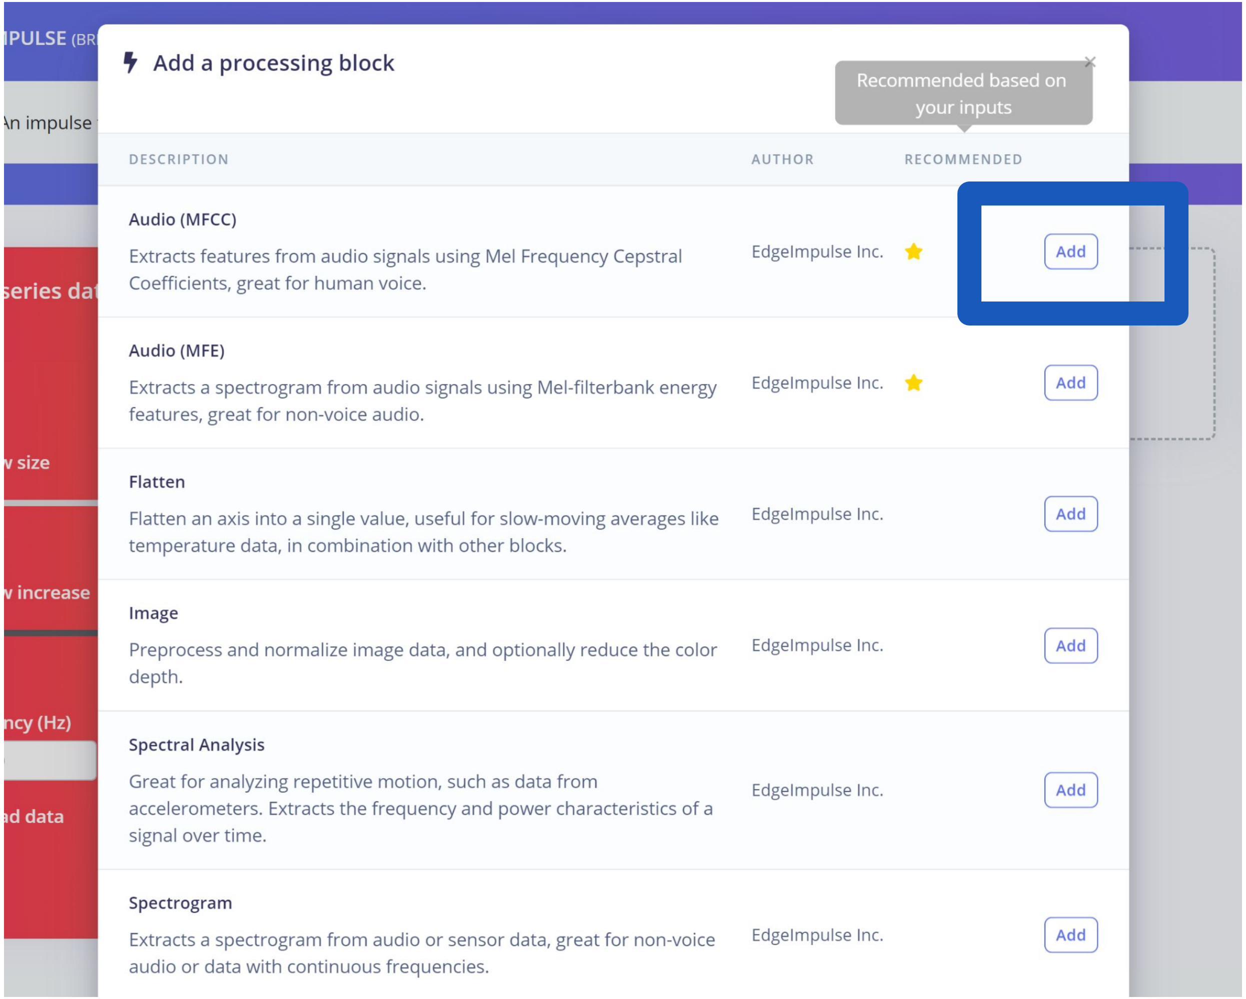

We'll keep things simple today and just add an MFCC but/and in future projects you can:

- **create your own blocks**
- **use multiple blocks**

https://docs.edgeimpulse.com/docs/custom-blocks

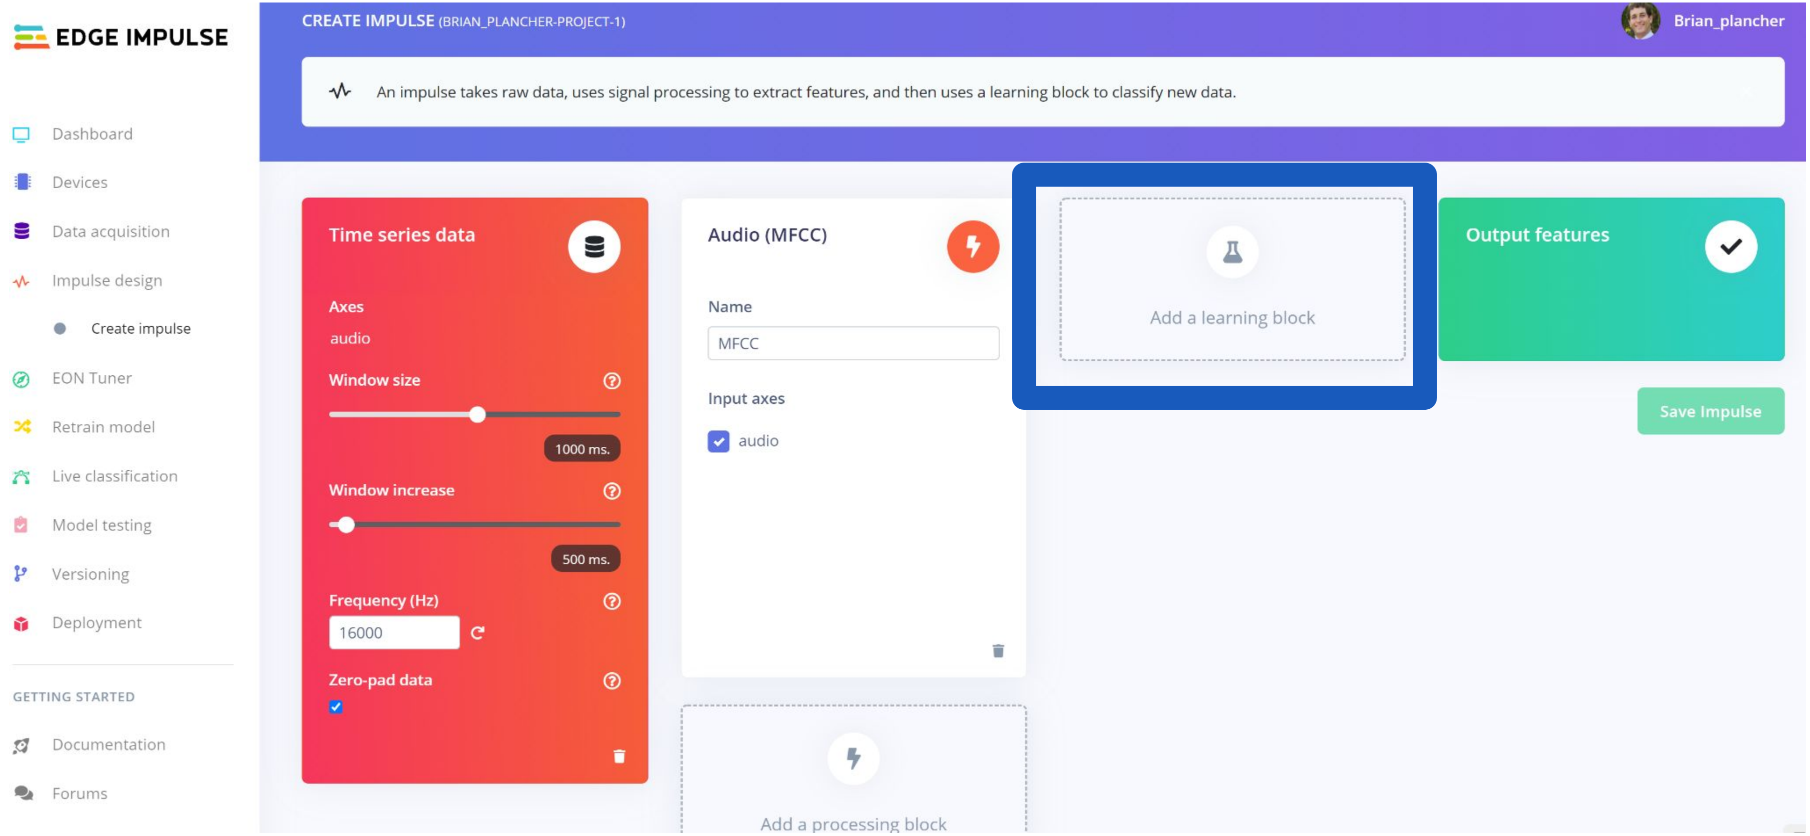
EDGE IMPULSE
CREATE IMPULSE (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
An impulse takes raw data, uses signal processing to extract features, and then uses a learning block to classify new data.
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Time series data
Axes
audio
Window size
1000 ms.
Window increase
500 ms.
Frequency (Hz)
16000
Zero-pad data
Audio (MFCC)
Name
MFCC
Input axes
audio
Add a processing block
Add a learning block
Output features
Save Impulse

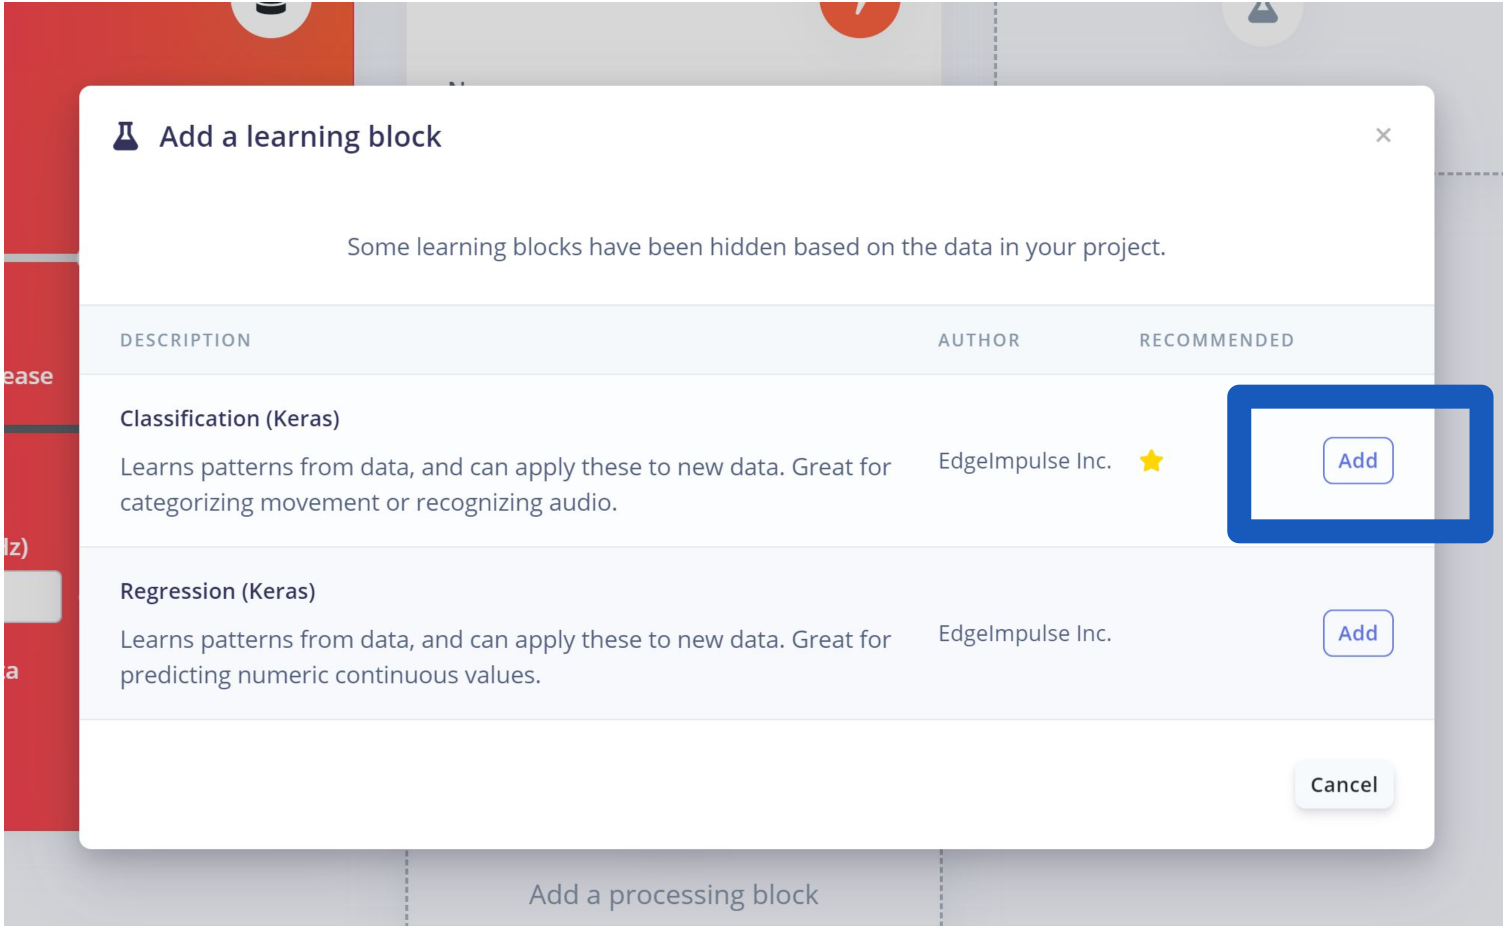
Add a learning block
Some learning blocks have been hidden based on the data in your project.
DESCRIPTION
AUTHOR
RECOMMENDED
Classification (Keras)
Learns patterns from data, and can apply these to new data. Great for categorizing movement or recognizing audio.
EdgeImpulse Inc.
Add
Regression (Keras)
Learns patterns from data, and can apply these to new data. Great for predicting numeric continuous values.
EdgeImpulse Inc.
Add
Cancel
Add a processing block

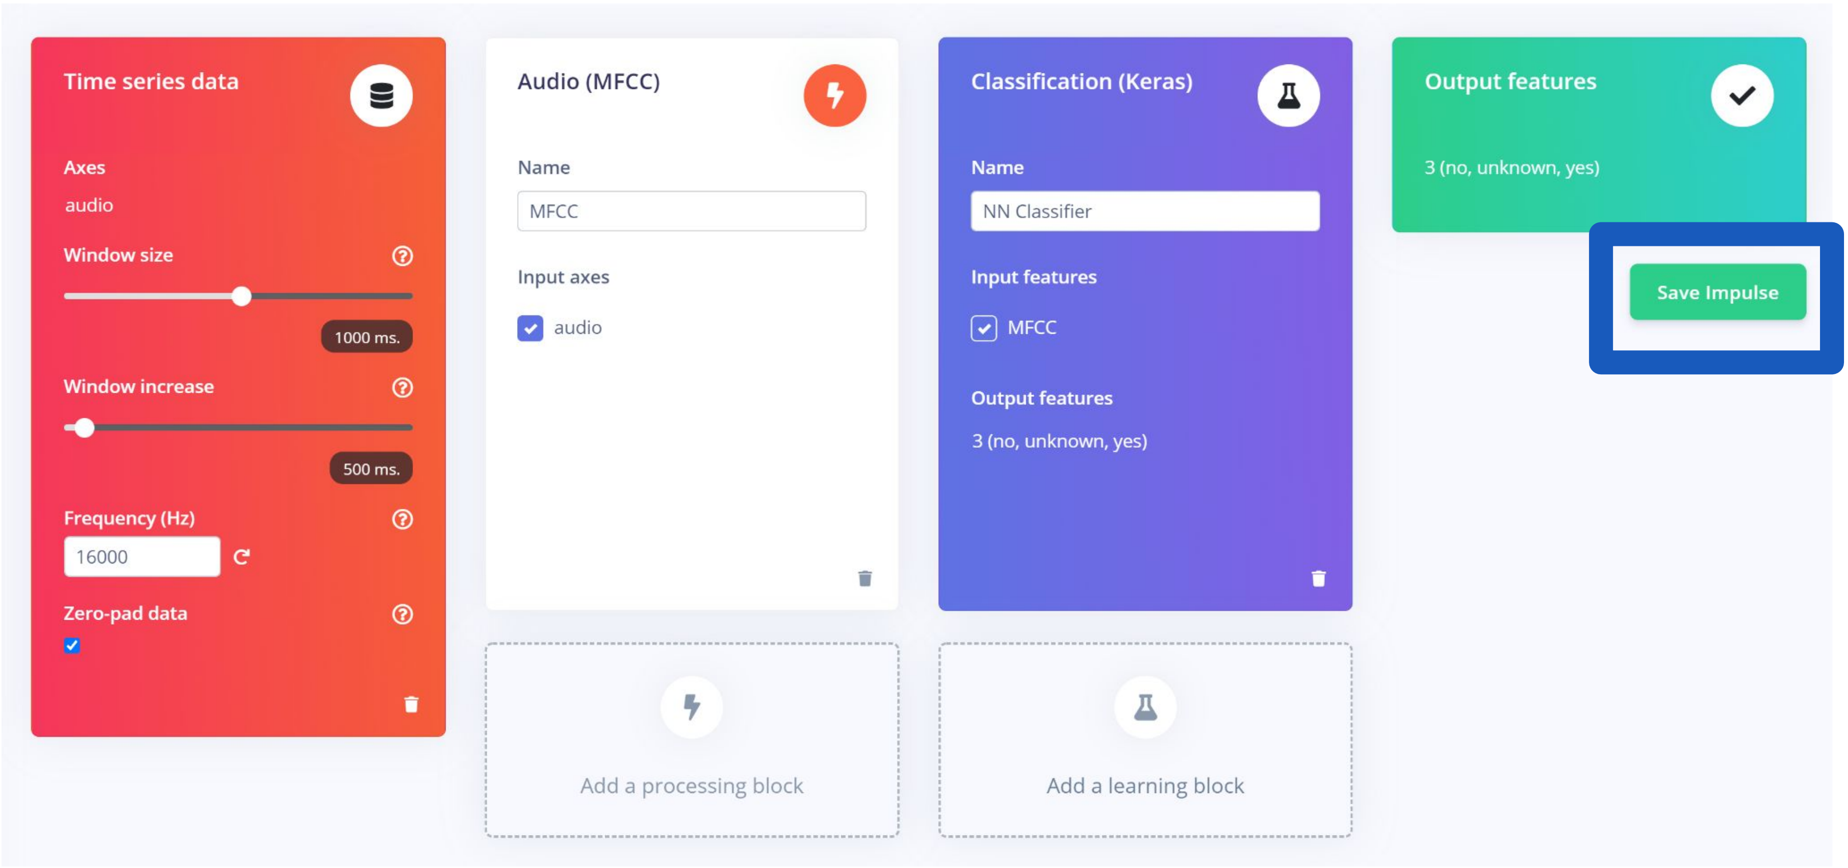

Time series data
Axes
audio
Window size
1000 ms.
Window increase
500 ms.
Frequency (Hz)
16000
Zero-pad data
Audio (MFCC)
Name
MFCC
Input axes
audio
Add a processing block
Classification (Keras)
Name
NN Classifier
Input features
MFCC
Output features
3 (no, unknown, yes)
Add a learning block
Output features
3 (no, unknown, yes)
Save Impulse

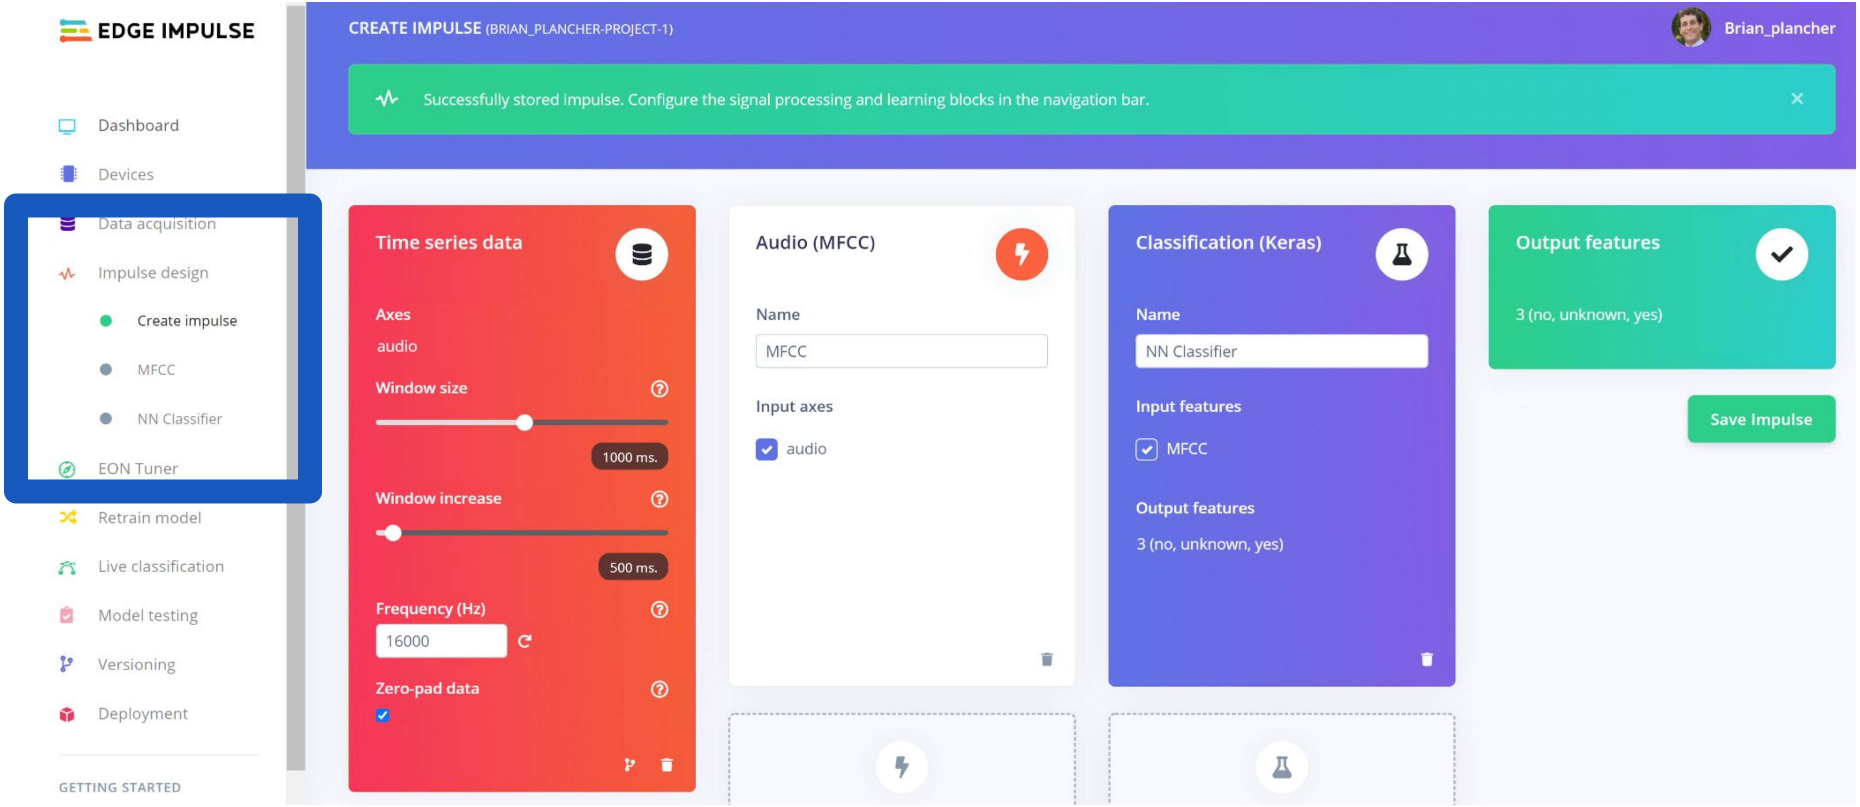

EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
MFCC
NN Classifier
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
CREATE IMPULSE (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
Successfully stored impulse. Configure the signal processing and learning blocks in the navigation bar.
Time series data
Axes
audio
Window size
1000 ms.
Window increase
500 ms.
Frequency (Hz)
16000
Zero-pad data
Audio (MFCC)
Name
MFCC
Input axes
audio
Classification (Keras)
Name
NN Classifier
Input features
MFCC
Output features
3 (no, unknown, yes)
Output features
3 (no, unknown, yes)
Save Impulse

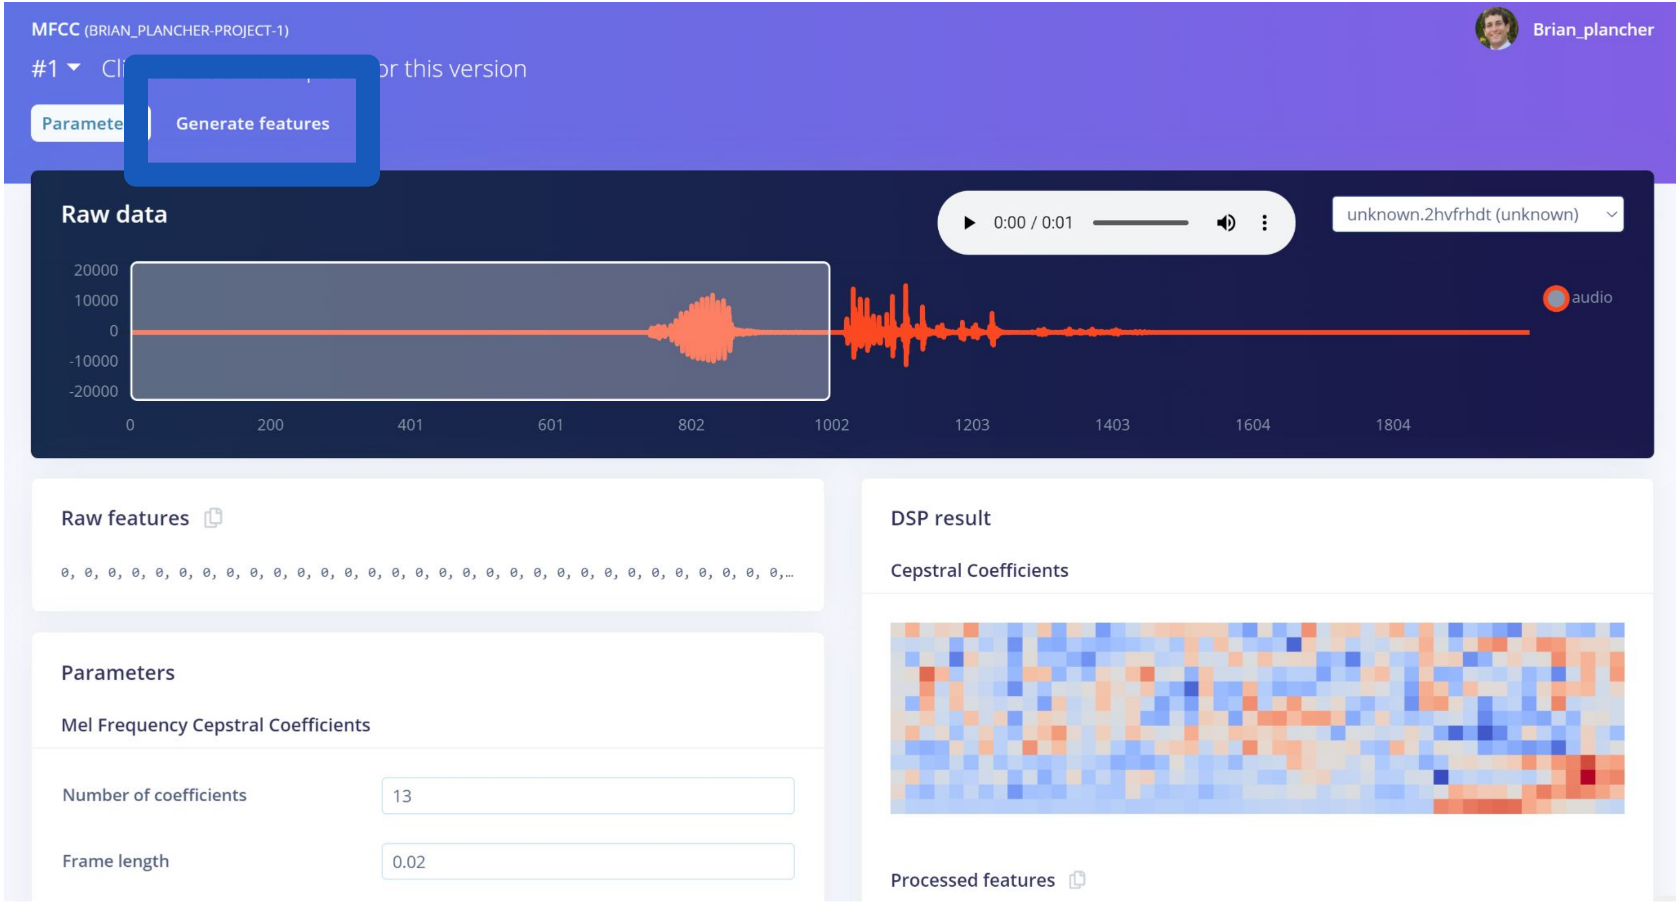
MFCC (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
#1
or this version
Paramete
Generate features
Raw data
0:00 / 0:01
unknown.2hvfrhdt (unknown)
audio
20000
10000
0
-10000
-20000
0
200
401
601
802
1002
1203
1403
1604
1804
Raw features
0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0,…
Parameters
Mel Frequency Cepstral Coefficients
Number of coefficients
13
Frame length
0.02
DSP result
Cepstral Coefficients
Processed features

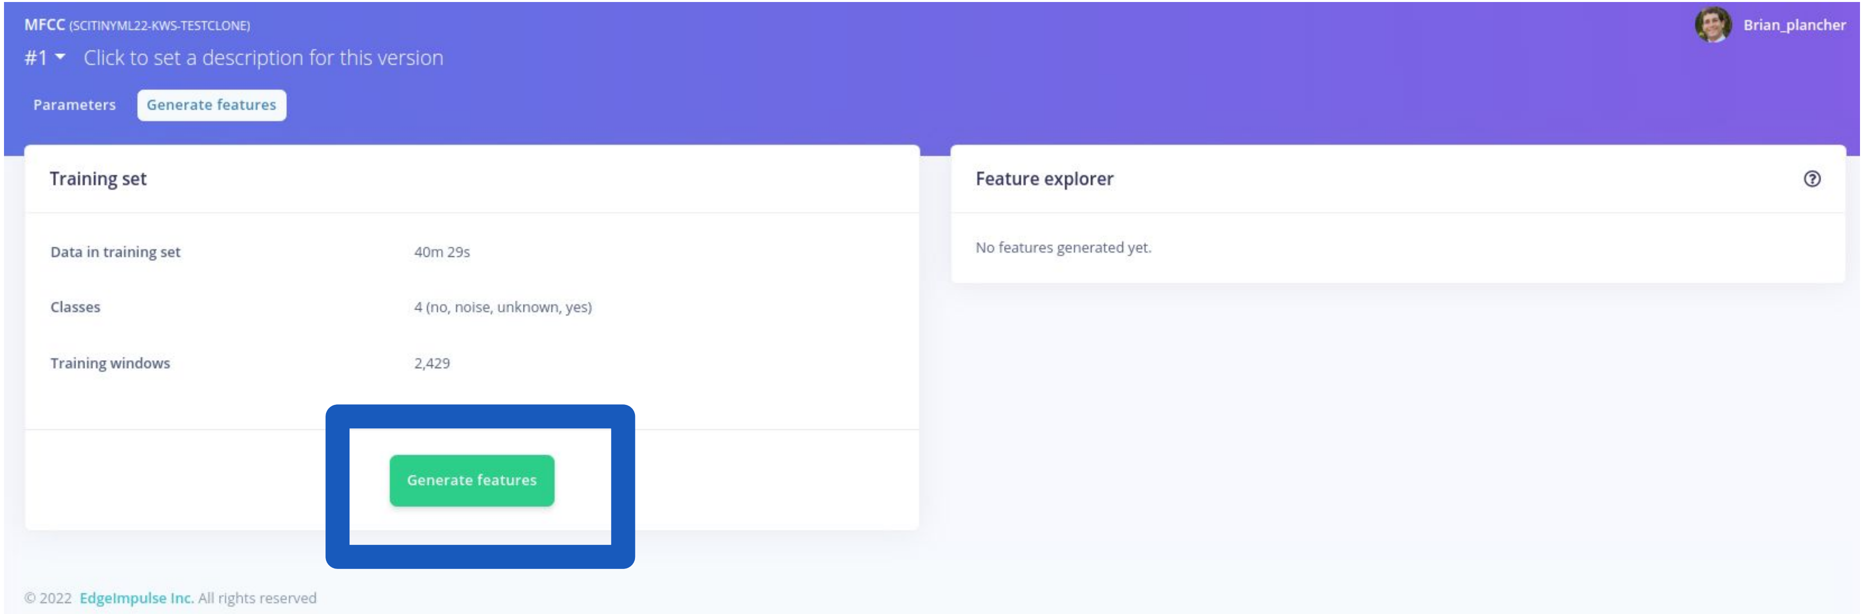
MFCC (SCITINYML22-KWS-TESTCLONE)
Brian_plancher
#1 Click to set a description for this version
Parameters
Generate features
Training set
Data in training set
40m 29s
Classes
4 (no, noise, unknown, yes)
Training windows
2,429
Generate features
Feature explorer
No features generated yet.
© 2022 EdgeImpulse Inc. All rights reserved

MFCC (SCITINYML22-KWS-TESTCLONE)

Brian_plancher

#1 Click to set a description for this version

Parameters | Generate features

Training set

Data in training set	40m 29s
Classes	4 (no, noise, unknown, yes)
Training windows	2,429

Generating features...

Feature explorer

No features generated yet.

Feature generation output

Cancel

```
Creating job... OK (ID: 2596741)

Scheduling job in cluster...
Job started
Creating windows from 2429 files...
[2/3] Pre-caching files...
[3/3] Pre-caching files...
Pre-caching files OK

[   1/2429] Creating windows from files...
```

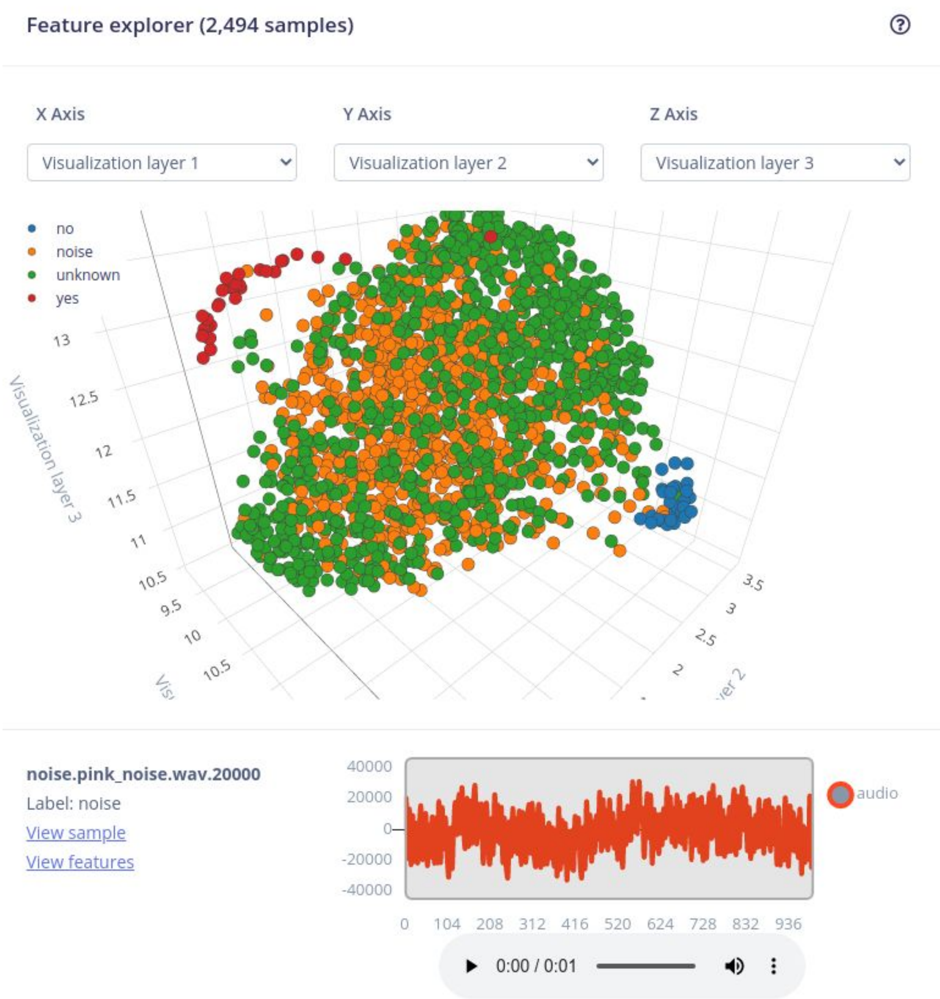

If you can visually see the clustering of the data then it is easier for the ML model to learn!
(But its not required and provides no guarantees)

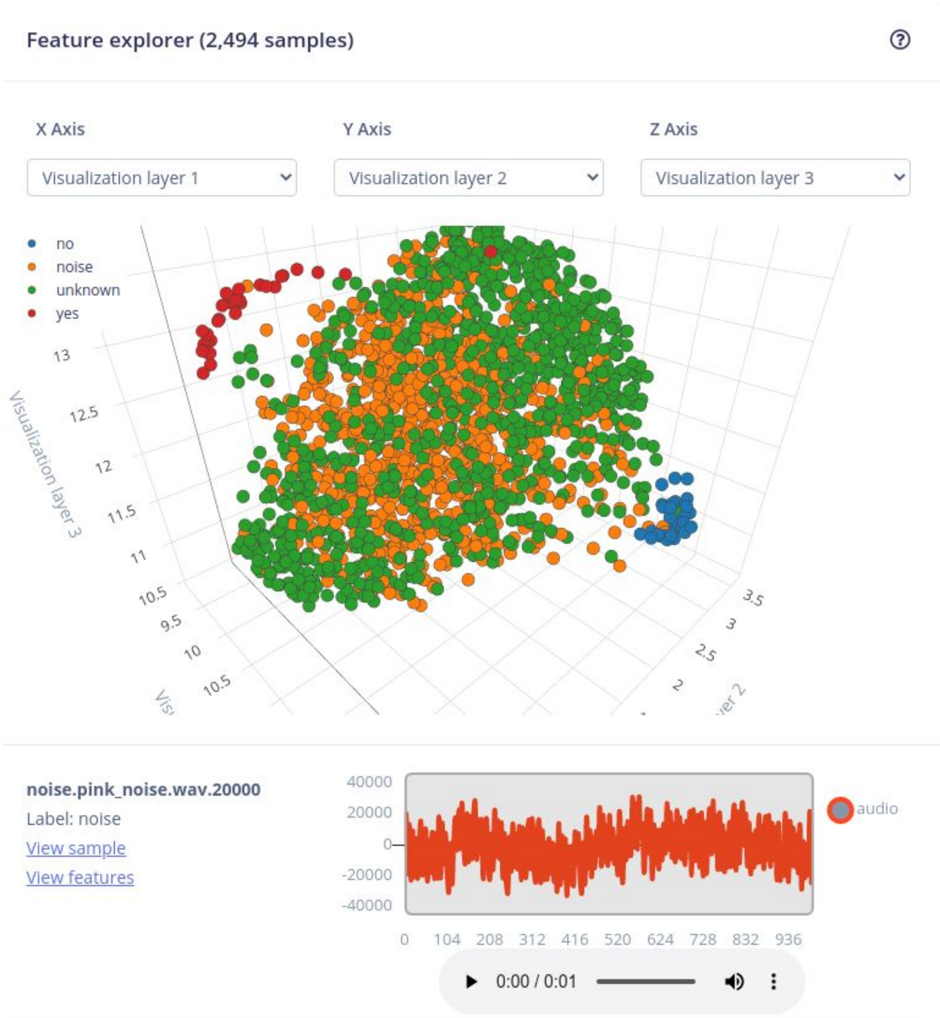
Feature explorer (2,494 samples)
X Axis
Y Axis
Z Axis
Visualization layer 1
Visualization layer 2
Visualization layer 3
no
noise
unknown
yes
Visualization layer 3
noise.pink_noise.wav.20000
Label: noise
View sample
View features
audio
0:00 / 0:01

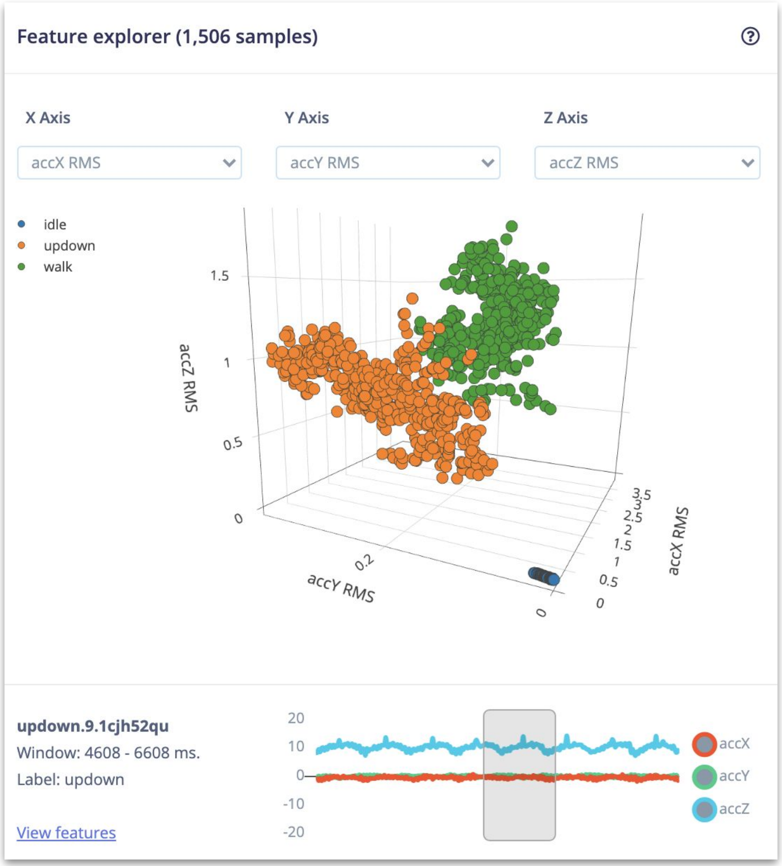
Feature explorer (1,506 samples)
X Axis
Y Axis
Z Axis
accX RMS
accY RMS
accZ RMS
idle
updown
walk
accZ RMS
accY RMS
accX RMS
updown.9.1cjh52qu
Window: 4608 - 6608 ms.
Label: updown
View features
accX
accY
accZ

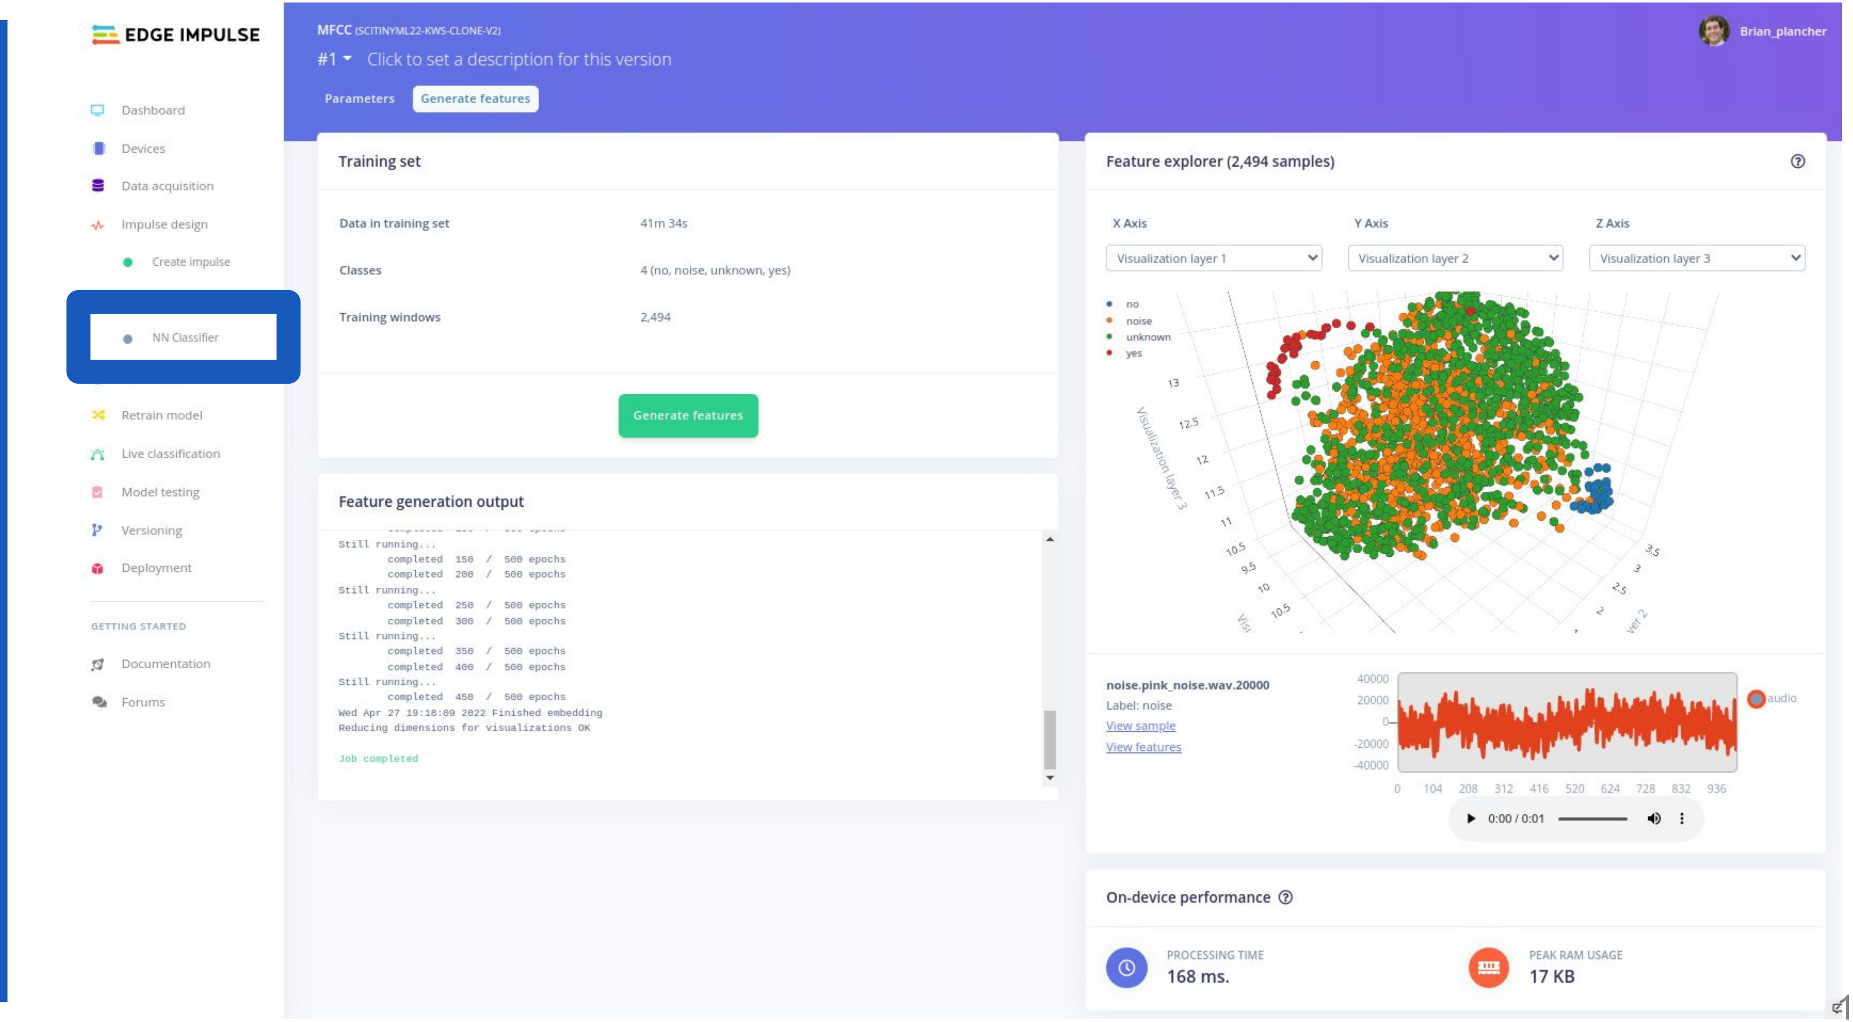
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
NN Classifier
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
MFCC (SCITINYML22-KWS-CLONE-V2)
#1 Click to set a description for this version
Brian_plancher
Parameters
Generate features
Training set
Data in training set 41m 34s
Classes 4 (no, noise, unknown, yes)
Training windows 2,494
Generate features
Feature generation output
Still running...
completed 150 / 500 epochs
completed 200 / 500 epochs
Still running...
completed 250 / 500 epochs
completed 300 / 500 epochs
Still running...
completed 350 / 500 epochs
completed 400 / 500 epochs
Still running...
completed 450 / 500 epochs
Wed Apr 27 19:18:09 2022 Finished embedding
Reducing dimensions for visualizations OK
Job completed
Feature explorer (2,494 samples)
X Axis
Y Axis
Z Axis
Visualization layer 1
Visualization layer 2
Visualization layer 3
no
noise
unknown
yes
Visualization layer 3
noise.pink_noise.wav.20000
Label: noise
View sample
View features
audio
0:00 / 0:01
On-device performance
PROCESSING TIME
168 ms.
PEAK RAM USAGE
17 KB

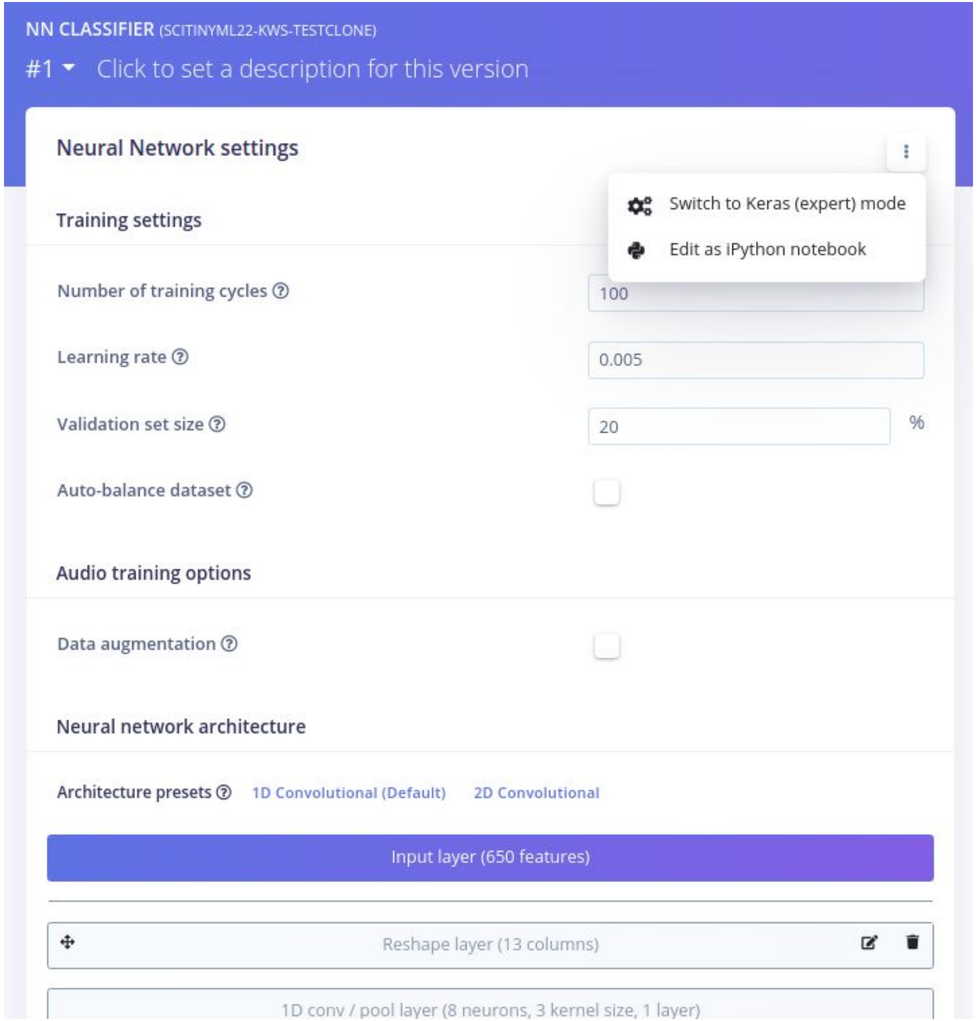
NN CLASSIFIER (SCITINYML22-KWS-TESTCLONE)
#1 Click to set a description for this version
Neural Network settings
Switch to Keras (expert) mode
Edit as iPython notebook
Training settings
Number of training cycles
100
Learning rate
0.005
Validation set size
20
%
Auto-balance dataset
Audio training options
Data augmentation
Neural network architecture
Architecture presets
1D Convolutional (Default)
2D Convolutional
Input layer (650 features)
Reshape layer (13 columns)
1D conv / pool layer (8 neurons, 3 kernel size, 1 layer)

Model Design with Edge Impulse

Pre-made neural network "blocks" that you can add!

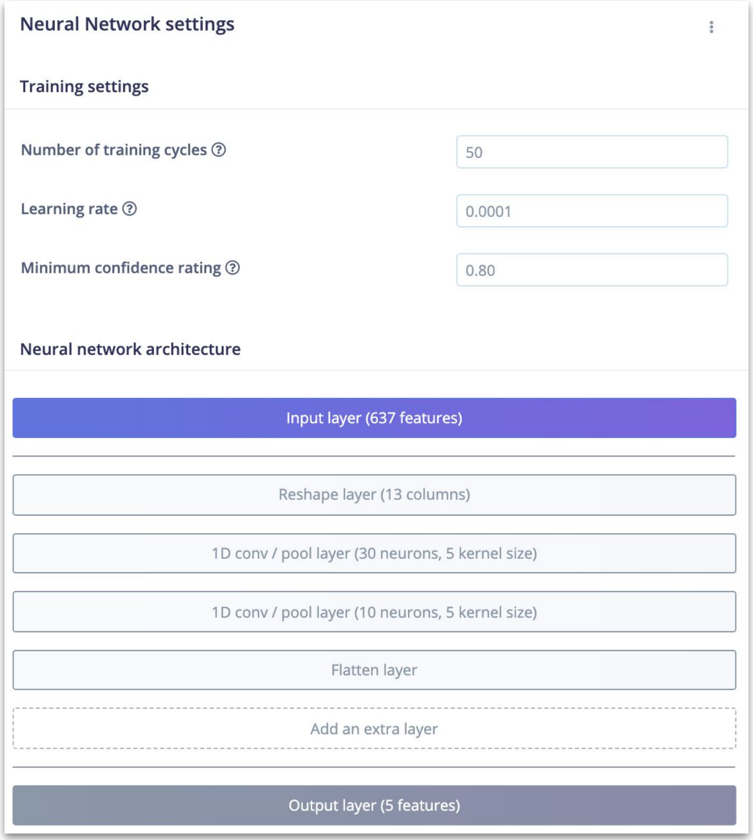

Model Design with Edge Impulse

“**Expert**” mode to write your own TensorFlow code

Neural network architecture

```
import tensorflow as tf
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense, InputLayer,
    Dropout, Conv1D, Conv2D, Flatten, Reshape, MaxPooling1D,
    MaxPooling2D, BatchNormalization
from tensorflow.keras.optimizers import Adam
sys.path.append('./resources/libraries')
import ei_tensorflow.training

# model architecture
model = Sequential()
channels = 1
columns = 13
rows = int(input_length / (columns * channels))
model.add(Reshape((rows, columns, channels), input_shape
    =(input_length, )))
model.add(Conv2D(8, kernel_size=3, activation='relu',
    kernel_constraint=tf.keras.constraints.MaxNorm(1),
    padding='same'))
model.add(MaxPooling2D(pool_size=2, strides=2, padding
    ='same'))
model.add(Dropout(0.25))
model.add(Conv2D(16, kernel_size=3, activation='relu',
    kernel_constraint=tf.keras.constraints.MaxNorm(1),
    padding='same'))
model.add(MaxPooling2D(pool_size=2, strides=2, padding
    ='same'))
model.add(Dropout(0.25))
model.add(Flatten())
model.add(Dense(classes, activation='softmax', name='y_pred'
    ))
```

Start training

Neural network architecture

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

Input layer (650 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 (3 features)

Start training

Neural network architecture

```
import tensorflow as tf
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense, InputLayer, Dropout, Conv1D, Conv2D,
    Flatten, Reshape, MaxPooling1D, MaxPooling2D, BatchNormalization,
    TimeDistributed
from tensorflow.keras.optimizers import Adam

# model architecture

model.add(Reshape((int(input_length / 13), 13), input_shape=(input_length, )))
model.add(Conv1D(8, kernel_size=3, activation='relu', padding='same'))
model.add(MaxPooling1D(pool_size=2, strides=2, padding='same'))

model.add(Conv1D(16, kernel_size=3, activation='relu', padding='same'))
model.add(MaxPooling1D(pool_size=2, strides=2, padding='same'))
model.add(Dropout(0.25))
model.add(Flatten())
model.add(Dense(classes, activation='softmax', name='y_pred'))

# this controls the learning rate
opt = Adam(lr=0.005, beta_1=0.9, beta_2=0.999)
# this controls the batch size, or you can manipulate the tf.data.Dataset objects
    yourself
BATCH_SIZE = 32
train_dataset = train_dataset.batch(BATCH_SIZE, drop_remainder=False)
validation_dataset = validation_dataset.batch(BATCH_SIZE, drop_remainder=False)
callbacks.append(BatchLoggerCallback(BATCH_SIZE, train_sample_count))

# train the neural network
model.compile(loss='categorical_crossentropy', optimizer=opt, metrics=['accuracy'])
model.fit(train_dataset, epochs=100, validation_data=validation_dataset, verbose=2,
    callbacks=callbacks)
```

Neural network architecture

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

Input layer (650 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 (3 features)

Start training

Neural network architecture

```
import tensorflow as tf
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense, InputLayer, Dropout, Conv1D, Conv2D,
    Flatten, Reshape, MaxPooling1D, MaxPooling2D, BatchNormalization,
    TimeDistributed
from tensorflow.keras.optimizers import Adam

# model architecture
model = Sequential()
model.add(Reshape((int(input_length / 13), 13), input_shape=(input_length, )))
model.add(Conv1D(8, kernel_size=3, activation='relu', padding='same'))
model.add(MaxPooling1D(pool_size=2, strides=2, padding='same'))
model.add(Dropout(0.25))
model.add(Conv1D(16, kernel_size=3, activation='relu', padding='same'))
model.add(MaxPooling1D(pool_size=2, strides=2, padding='same'))
model.add(Dropout(0.25))
model.add(Flatten())
model.add(Dense(classes, activation='softmax', name='y_pred'))

# this controls the learning rate
opt = Adam(lr=0.005, beta_1=0.9, beta_2=0.999)
# this controls the batch size, or you can manipulate the tf.data.Dataset objects
    yourself
BATCH_SIZE = 32
train_dataset = train_dataset.batch(BATCH_SIZE, drop_remainder=False)
validation_dataset = validation_dataset.batch(BATCH_SIZE, drop_remainder=False)
```

For now just stick with the defaults but/and you can easily design **any model** you want and use **any optimizer** you want using **TensorFlow**!

WARNING: if you want to deploy to a microcontroller make sure you only use Ops supported by TensorFlow Lite Micro! https://github.com/tensorflow/tflite-micro/blob/main/tensorflow/lite/micro/all_ops_resolver.cc#L22

the … can easily design **any model** you want and use **any optimizer** you want using **TensorFlow**!

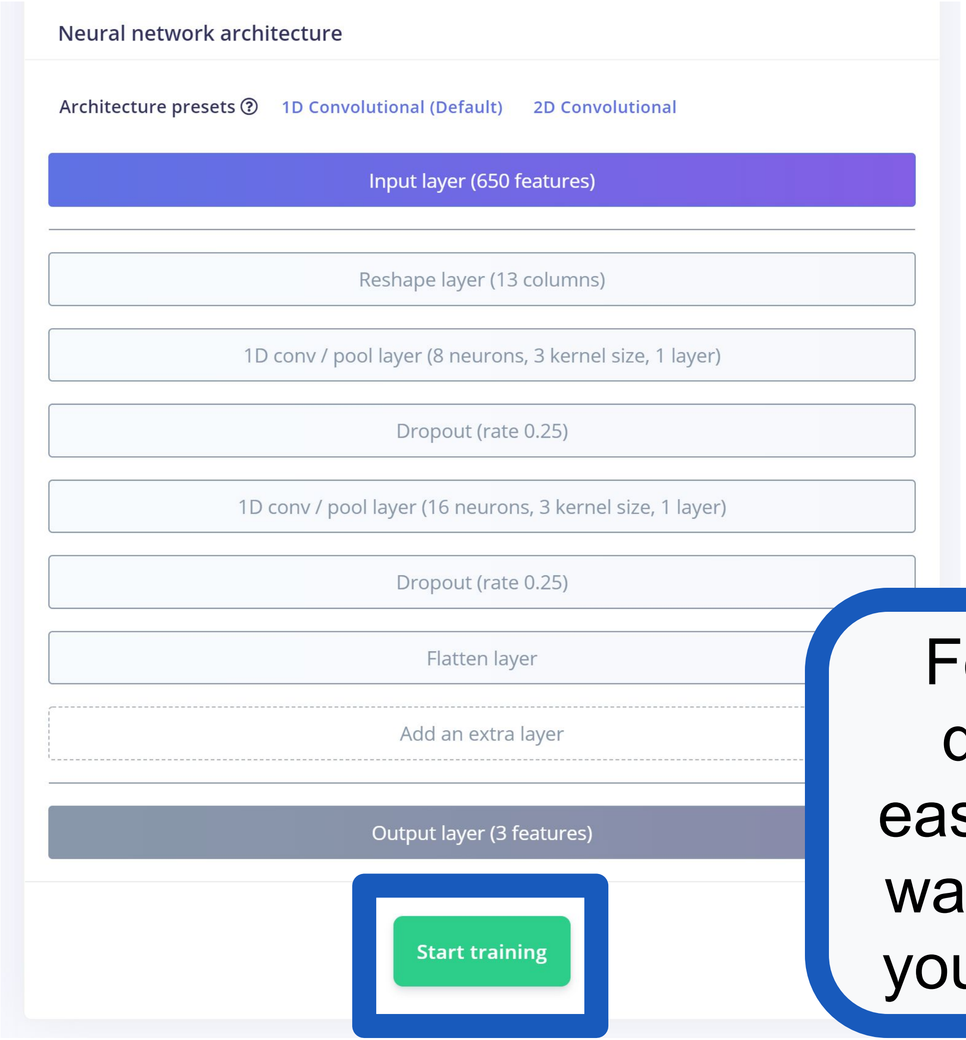

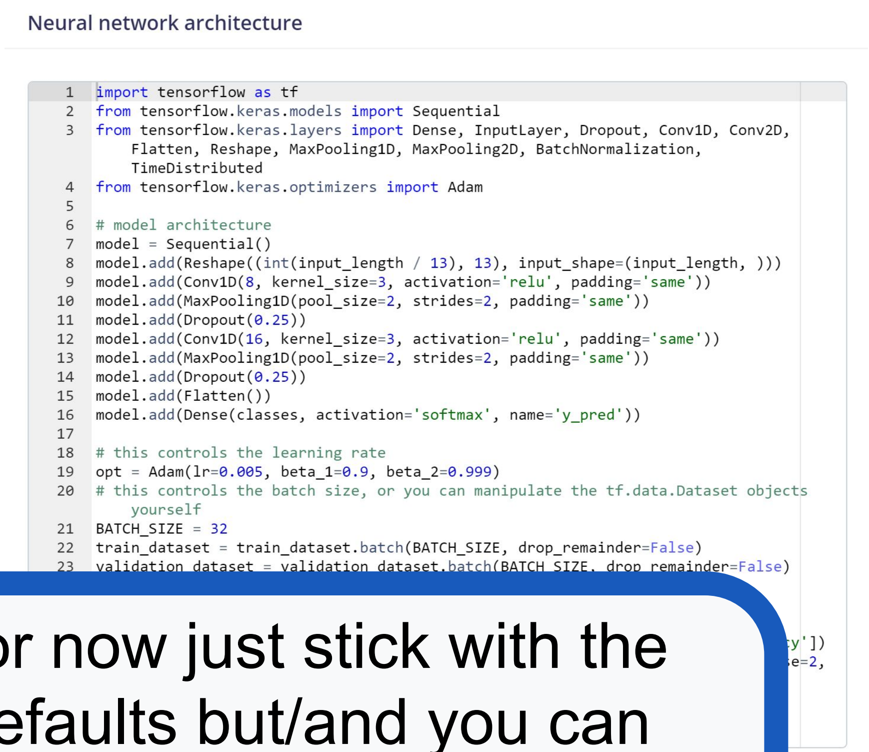

For now just stick with the defaults but/and you can easily design **any model** you want and use **any optimizer** you want using **TensorFlow**!

Training output

```
Epoch 95/100
4/4 - 0s - loss: 0.1044 - accuracy: 0.9500 - val_loss: 0.2934 - val_accuracy: 0.9231
Epoch 96/100
4/4 - 0s - loss: 0.0256 - accuracy: 1.0000 - val_loss: 0.3830 - val_accuracy: 0.8846
Epoch 97/100
4/4 - 0s - loss: 0.0523 - accuracy: 0.9800 - val_loss: 0.4366 - val_accuracy: 0.8462
Epoch 98/100
4/4 - 0s - loss: 0.0451 - accuracy: 0.9800 - val_loss: 0.4265 - val_accuracy: 0.8846
Epoch 99/100
4/4 - 0s - loss: 0.0514 - accuracy: 0.9900 - val_loss: 0.3926 - val_accuracy: 0.8846
Epoch 100/100
4/4 - 0s - loss: 0.0348 - accuracy: 0.9900 - val_loss: 0.3571 - val_accuracy: 0.9231
Finished training
```

Training Set

Validation Set

Final Accuracy

Model

Model version: ⓘ Quantized (int8)

Last training performance (validation set)

ACCURACY
96.6%

LOSS
0.09

Confusion matrix (validation set)

	NO	NOISE	UNKNOWN	YES
NO	100%	0%	0%	0%
NOISE	0%	96.4%	3.6%	0%
UNKNOWN	0%	2.5%	97.1%	0.4%
YES	0%	0%	20%	80%
F1 SCORE	1.00	0.97	0.96	0.80

Feature explorer (full training set) ⓘ

- no - correct
- noise - correct
- unknown - correct
- yes - correct
- noise - incorrect
- unknown - incorrect
- yes - incorrect

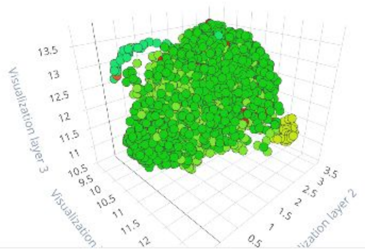

noise.orig_test.Neighbor_6.wav.8000
Label: noise
Predicted: noise
View sample
View features

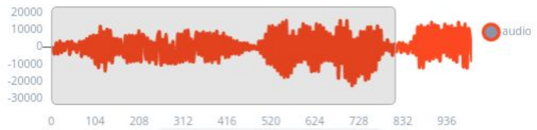

0:00 / 0:01

On-device performance ⓘ

INFERENCING TIME
11 ms.

PEAK RAM USAGE
5.0K

FLASH USAGE
34.8K

Final Accuracy

Accuracy Breakdown

Model

Model version: ⓘ Quantized (int8)

Last training performance (validation set)

ACCURACY
96.6%

LOSS
0.09

Confusion matrix (validation set)

	NO	NOISE	UNKNOWN	YES
NO	100%	0%	0%	0%
NOISE	0%	96.4%	3.6%	0%
UNKNOWN	0%	2.5%	97.1%	0.4%
YES	0%	0%	20%	80%
F1 SCORE	1.00	0.97	0.96	0.80

Feature explorer (full training set) ⓘ

- no - correct
- noise - correct
- unknown - correct
- yes - correct
- noise - incorrect
- unknown - incorrect
- yes - incorrect

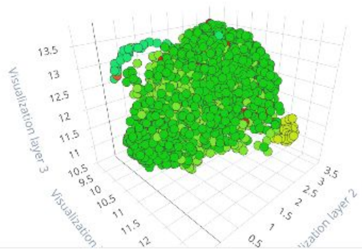

noise.orig_test.Neighbor_6.wav.8000
Label: noise
Predicted: noise
View sample
View features

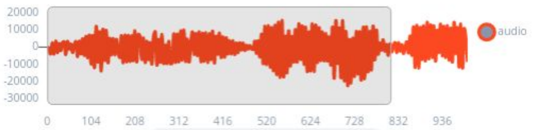

0:00 / 0:01

On-device performance ⓘ

INFERENCING TIME
11 ms.

PEAK RAM USAGE
5.0K

FLASH USAGE
34.8K

Confusion Matrix

	Actual Output = Yes	Actual Output = No
Predicted Output = Yes	**# of True Positive**	# of False Positive ***Type 1 Error***
Predicted Output = No	# of False Negative ***Type 2 Error***	**# of True Negative**

Final Accuracy

Accuracy Breakdown

Feature Explorer

Individual Data Points

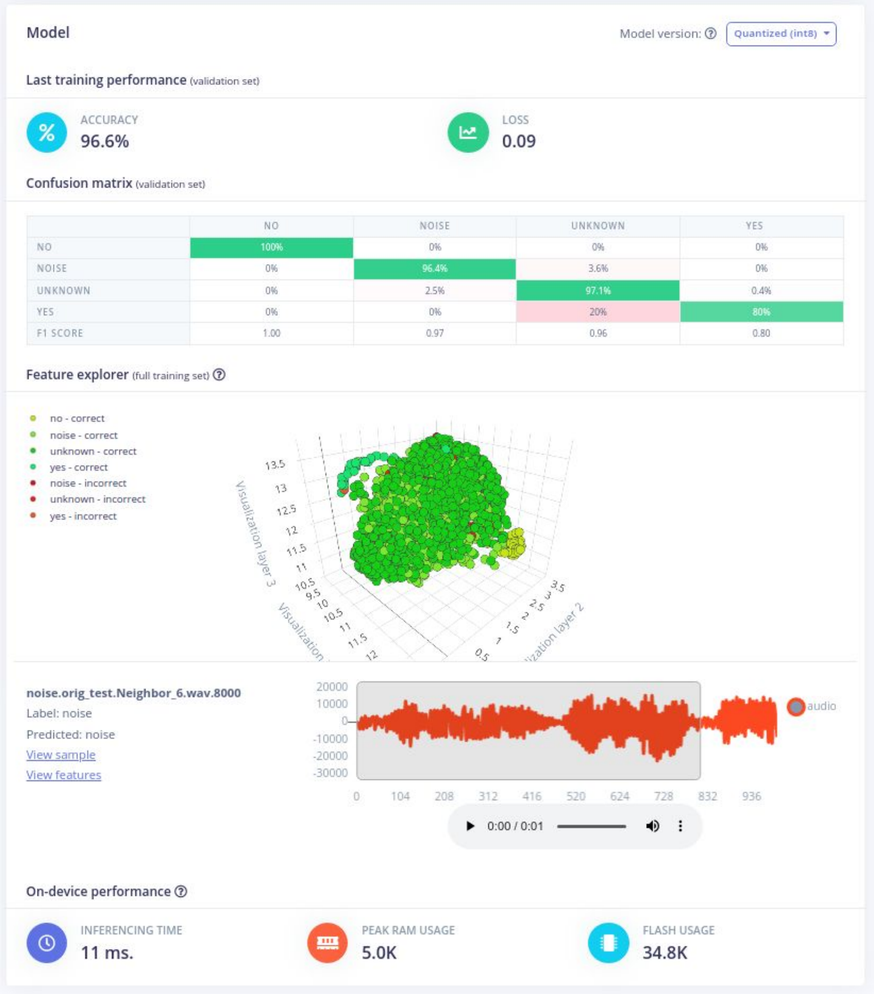

Final Accuracy

Accuracy Breakdown

Feature Explorer

Individual Data Points

Expected runtime/memory

Model

Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY
96.6%

LOSS
0.09

Confusion matrix (validation set)

	NO	NOISE	UNKNOWN	YES
NO	100%	0%	0%	0%
NOISE	0%	96.4%	3.6%	0%
UNKNOWN	0%	2.5%	97.1%	0.4%
YES	0%	0%	20%	80%
F1 SCORE	1.00	0.97	0.96	0.80

Feature explorer (full training set)

- no - correct
- noise - correct
- unknown - correct
- yes - correct
- noise - incorrect
- unknown - incorrect
- yes - incorrect

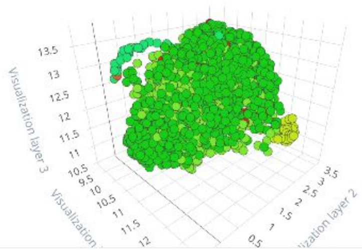

noise.orig_test.Neighbor_6.wav.8000
Label: noise
Predicted: noise
View sample
View features

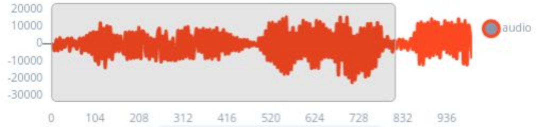

0:00 / 0:01

On-device performance

INFERENCING TIME
11 ms.

PEAK RAM USAGE
5.0K

FLASH USAGE
34.8K

Edge Impulse Project Dashboard

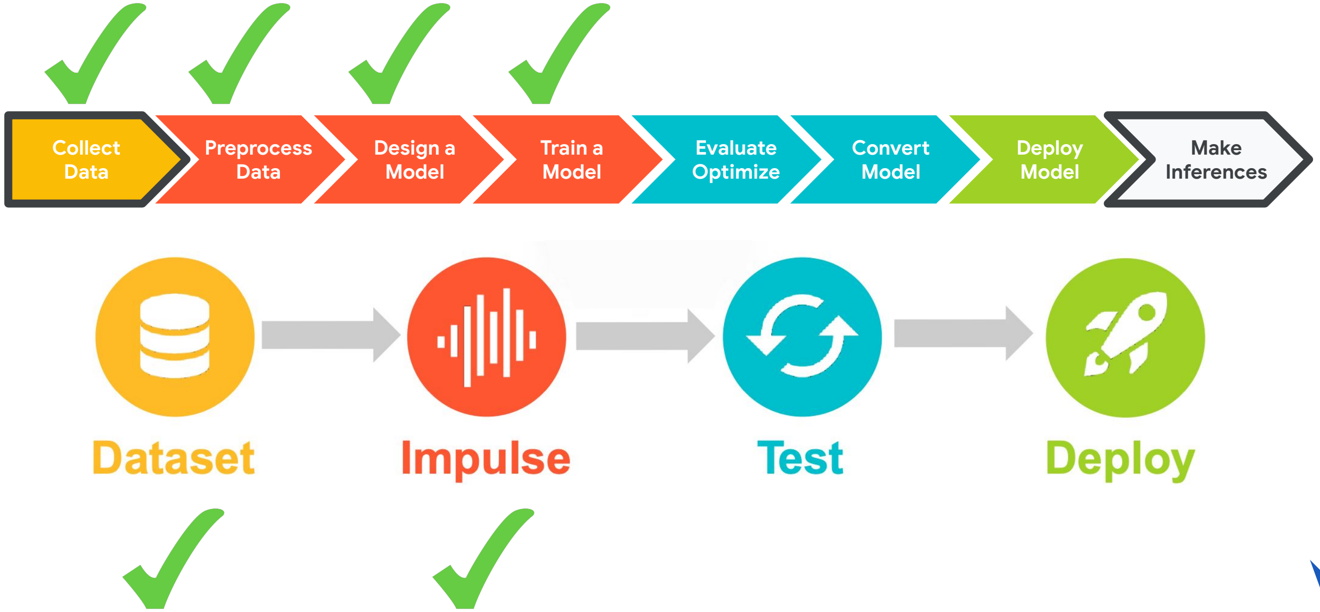

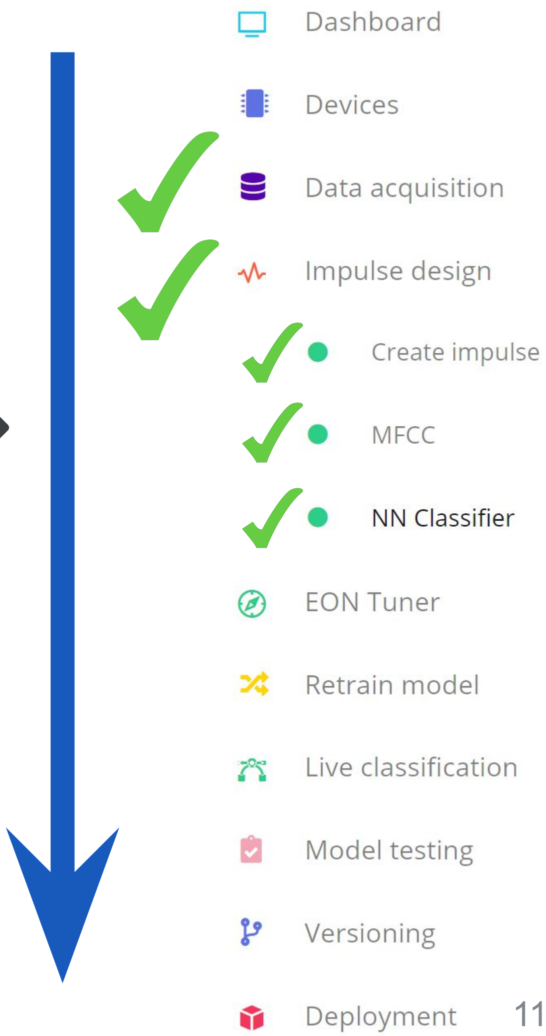

Today’s Agenda

- A Quick Review of What We’ve Learned
- Data Engineering for KWS
- Hands-on KWS Data Collection with Edge Impulse
- (Hands-on) Data Preprocessing for KWS
- **Deploying our Model onto our Arduino**
- Summary

Edge Impulse Project Dashboard

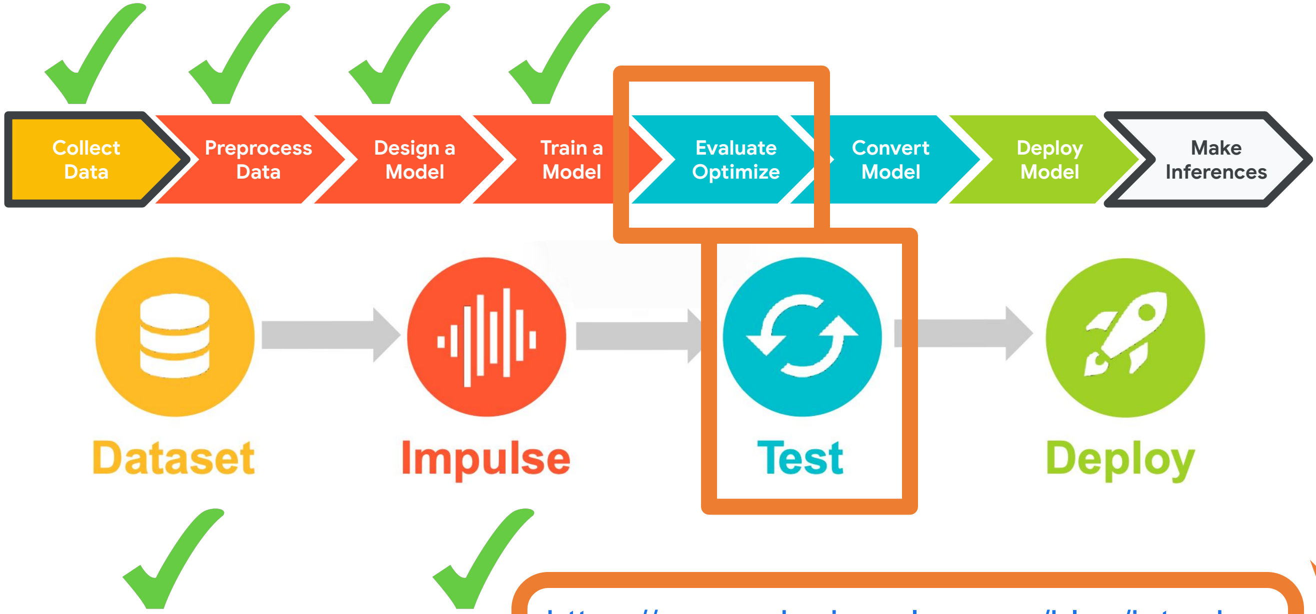

https://www.edgeimpulse.com/blog/introducing-the-eon-tuner-edge-impulses-new-automl-tool-for-embedded-machine-learning

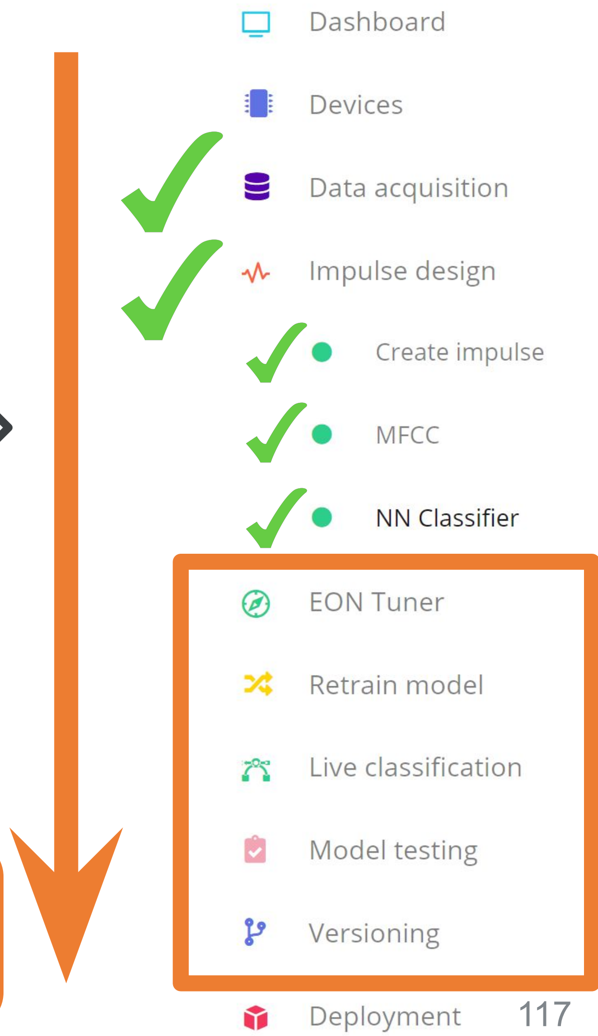

Edge Impulse Project Dashboard

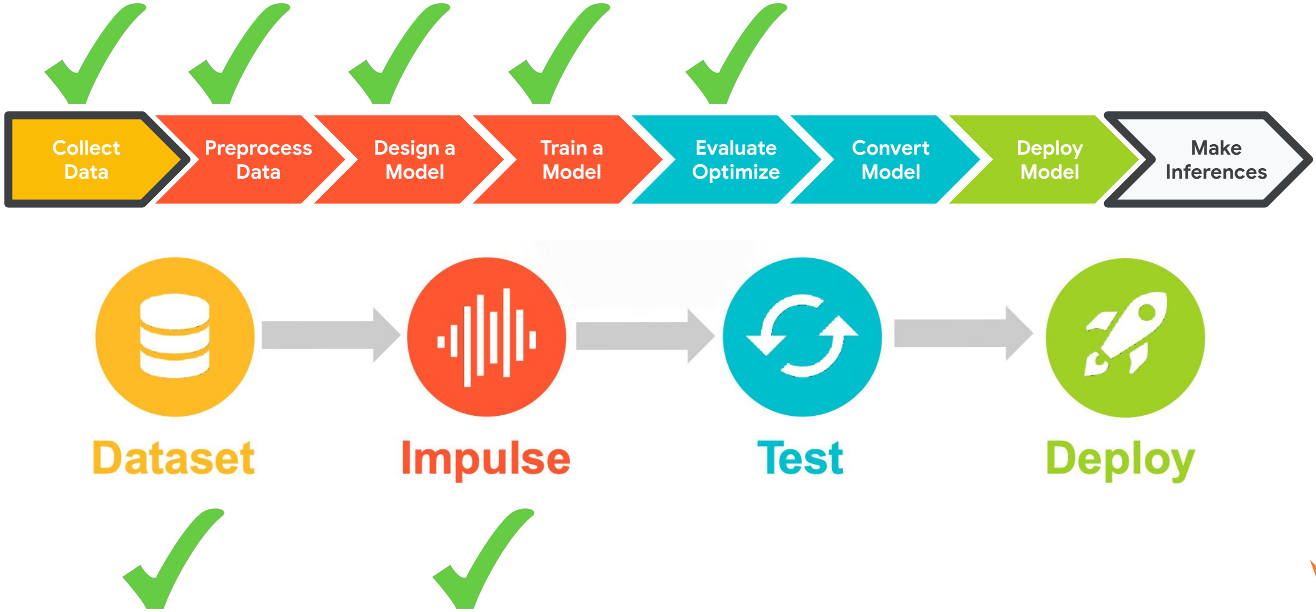

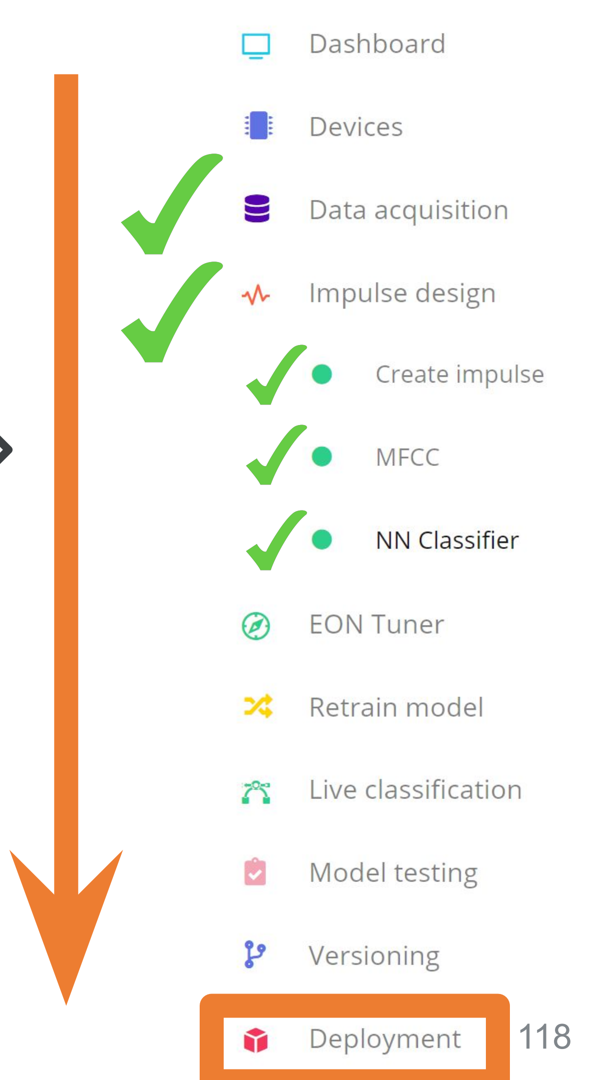

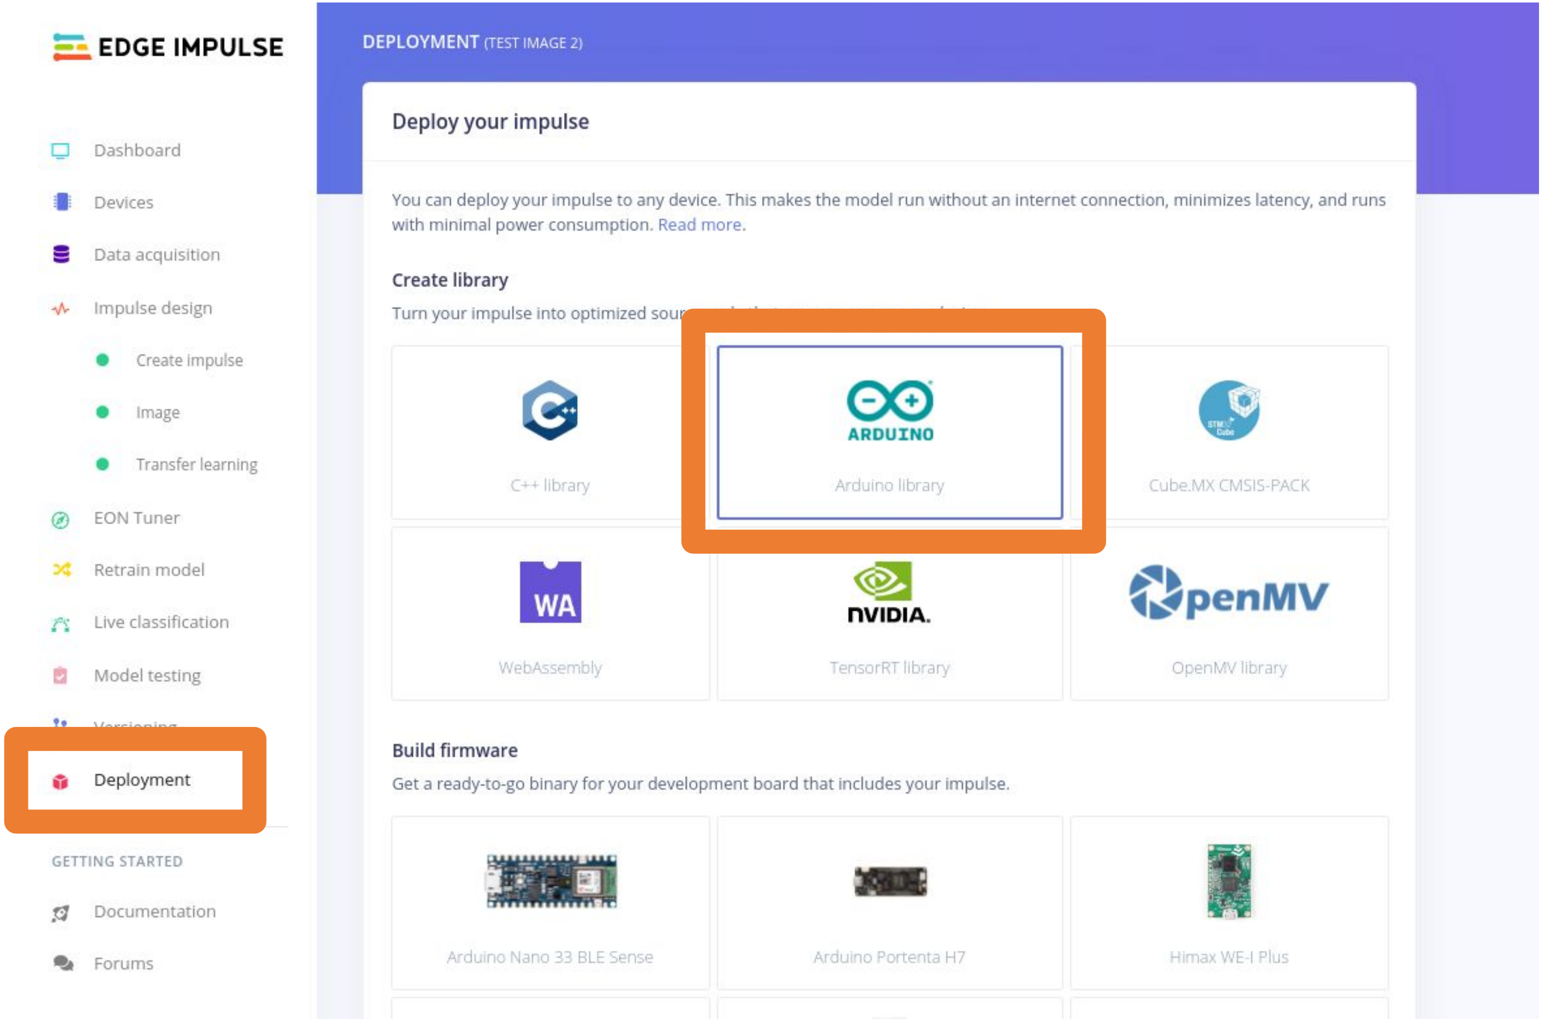
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Deployment
GETTING STARTED
Documentation
Forums
DEPLOYMENT (TEST IMAGE 2)
Deploy your impulse
You can deploy your impulse to any device. This makes the model run without an internet connection, minimizes latency, and runs with minimal power consumption. Read more.
Create library
Turn your impulse into optimized sour
C++ library
ARDUINO
Arduino library
Cube.MX CMSIS-PACK
WA
WebAssembly
NVIDIA.
TensorRT library
OpenMV
OpenMV library
Build firmware
Get a ready-to-go binary for your development board that includes your impulse.
Arduino Nano 33 BLE Sense
Arduino Portenta H7
Himax WE-I Plus

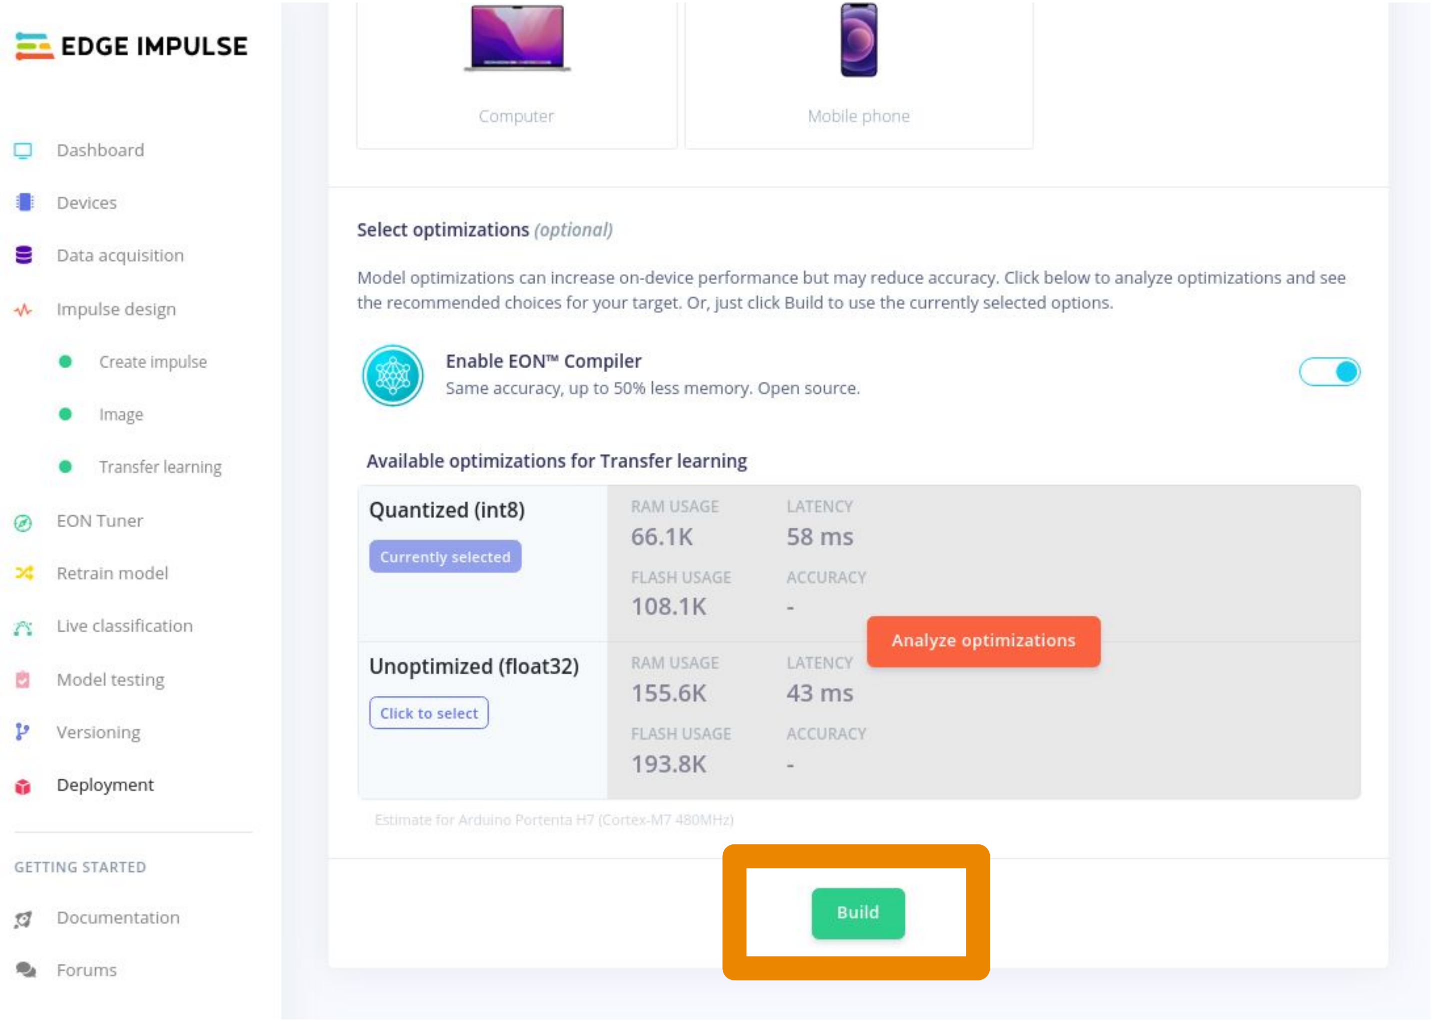
EDGE IMPULSE
Dashboard
Devices
Data acquisition
Impulse design
Create impulse
Image
Transfer learning
EON Tuner
Retrain model
Live classification
Model testing
Versioning
Deployment
GETTING STARTED
Documentation
Forums
Computer
Mobile phone
Select optimizations (optional)
Model optimizations can increase 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 optimizations for Transfer learning
Quantized (int8)
Currently selected
RAM USAGE
66.1K
LATENCY
58 ms
FLASH USAGE
108.1K
ACCURACY
-
Analyze optimizations
Unoptimized (float32)
Click to select
RAM USAGE
155.6K
LATENCY
43 ms
FLASH USAGE
193.8K
ACCURACY
-
Estimate for Arduino Portenta H7 (Cortex-M7 480MHz)
Build

Edge Impulse Project Dashboard

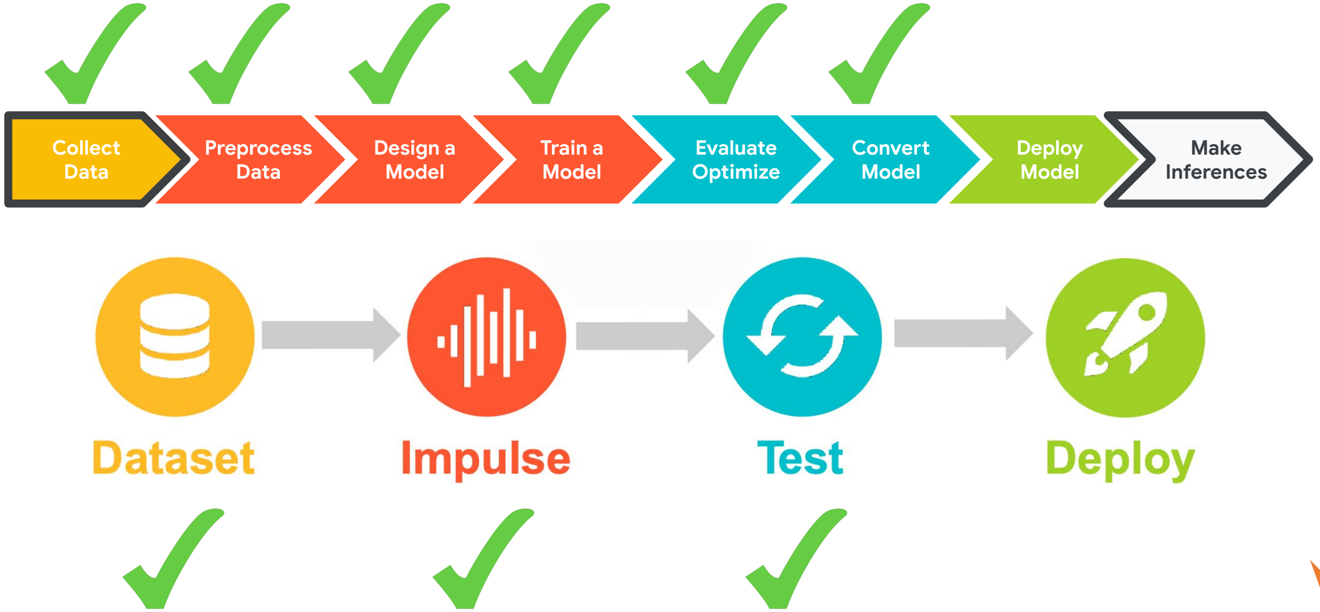

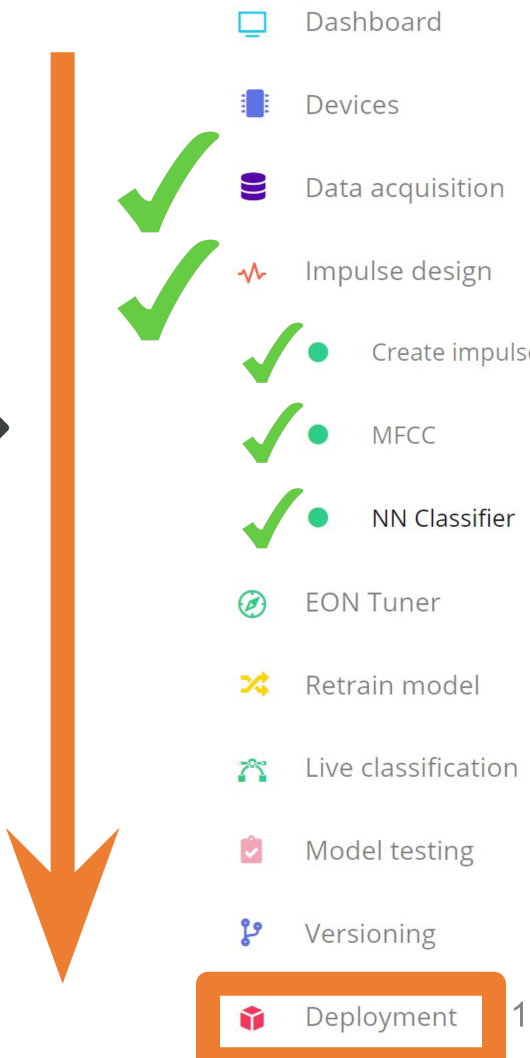

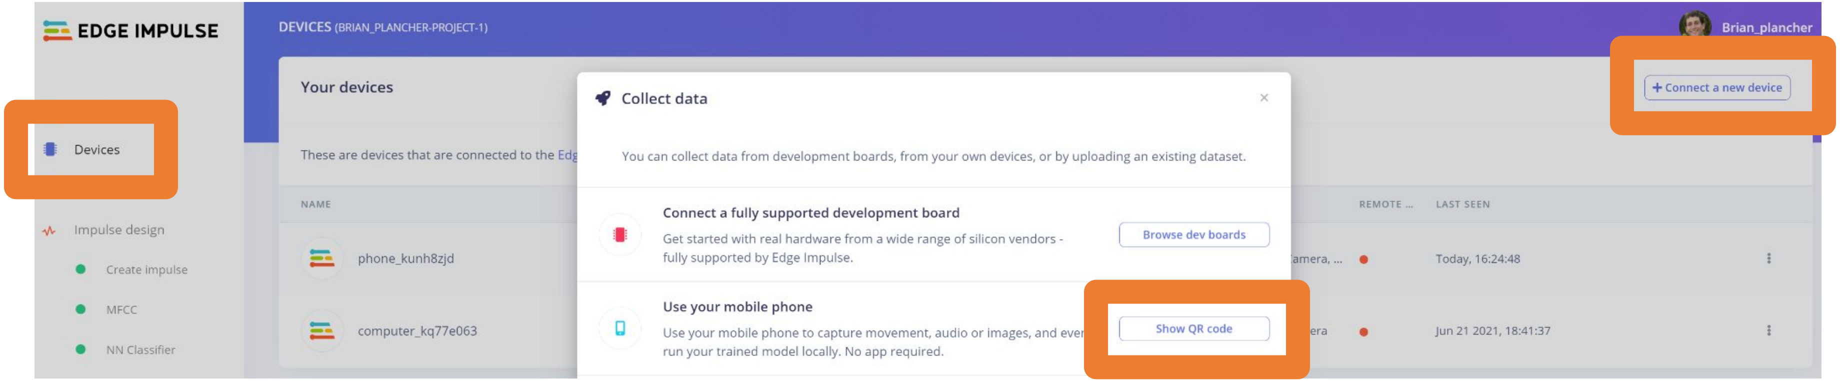
EDGE IMPULSE
DEVICES (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
Your devices
Devices
These are devices that are connected to the Edg
NAME
Impulse design
Create impulse
MFCC
NN Classifier
phone_kunh8zjd
computer_kq77e063
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
Use your mobile phone
Use your mobile phone to capture movement, audio or images, and ever run your trained model locally. No app required.
Show QR code
+ Connect a new device
REMOTE ...
LAST SEEN
Today, 16:24:48
Jun 21 2021, 18:41:37

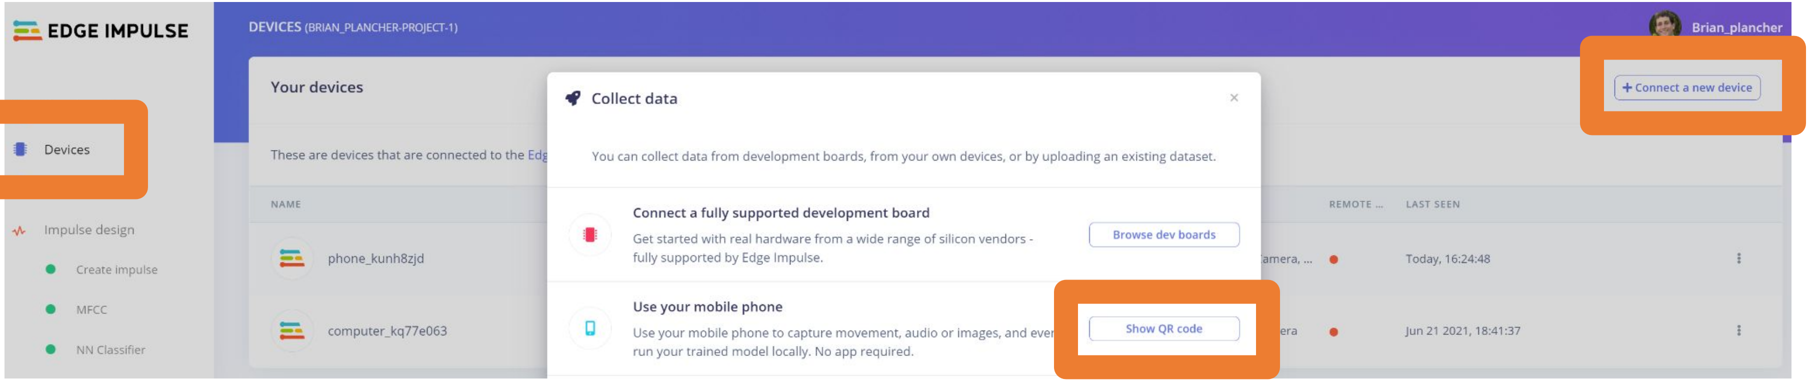
EDGE IMPULSE
DEVICES (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
Devices
Impulse design
Create impulse
MFCC
NN Classifier
Your devices
These are devices that are connected to the Edg
NAME
phone_kunh8zjd
computer_kq77e063
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
Use your mobile phone
Use your mobile phone to capture movement, audio or images, and ever run your trained model locally. No app required.
Show QR code
+ Connect a new device
REMOTE ...
LAST SEEN
amera, ...
Today, 16:24:48
era
Jun 21 2021, 18:41:37

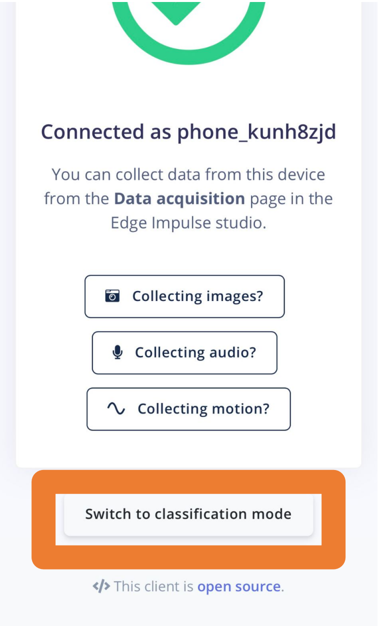
Connected as phone_kunh8zjd
You can collect data from this device from the Data acquisition page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?
Switch to classification mode
This client is open source.

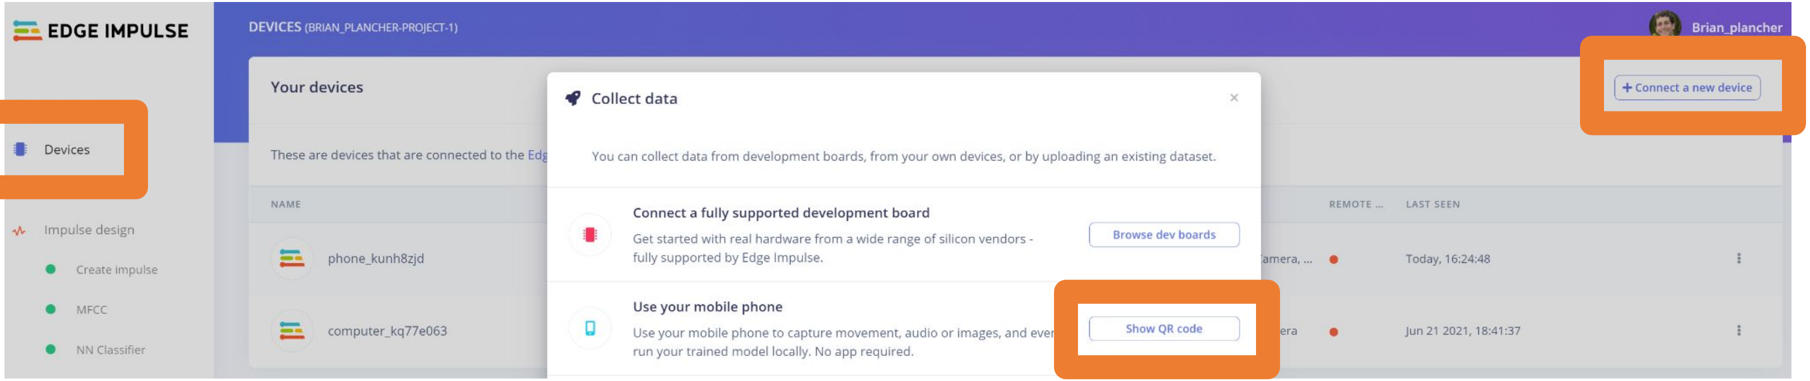
EDGE IMPULSE
DEVICES (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
Devices
Impulse design
Create impulse
MFCC
NN Classifier
Your devices
These are devices that are connected to the Edg
NAME
phone_kunh8zjd
computer_kq77e063
+ Connect a new device
REMOTE ...
LAST SEEN
Today, 16:24:48
Jun 21 2021, 18:41:37
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
Use your mobile phone
Use your mobile phone to capture movement, audio or images, and ever run your trained model locally. No app required.
Show QR code

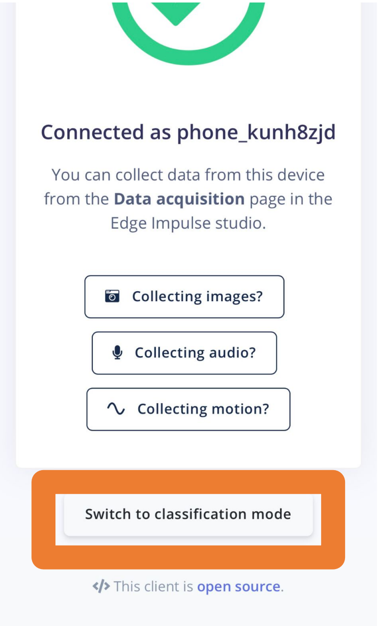
Connected as phone_kunh8zjd
You can collect data from this device from the **Data acquisition** page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?
Switch to classification mode
This client is open source.

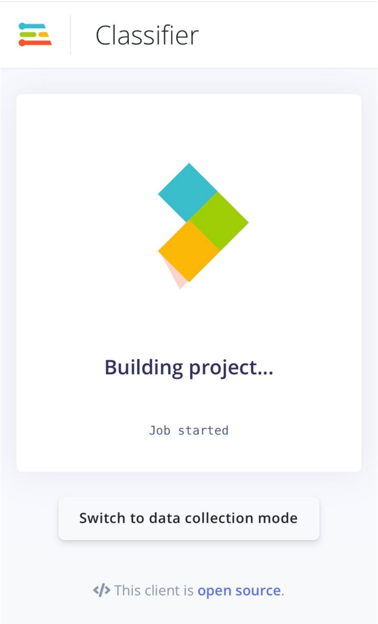
Classifier
Building project...
Job started
Switch to data collection mode
This client is open source.

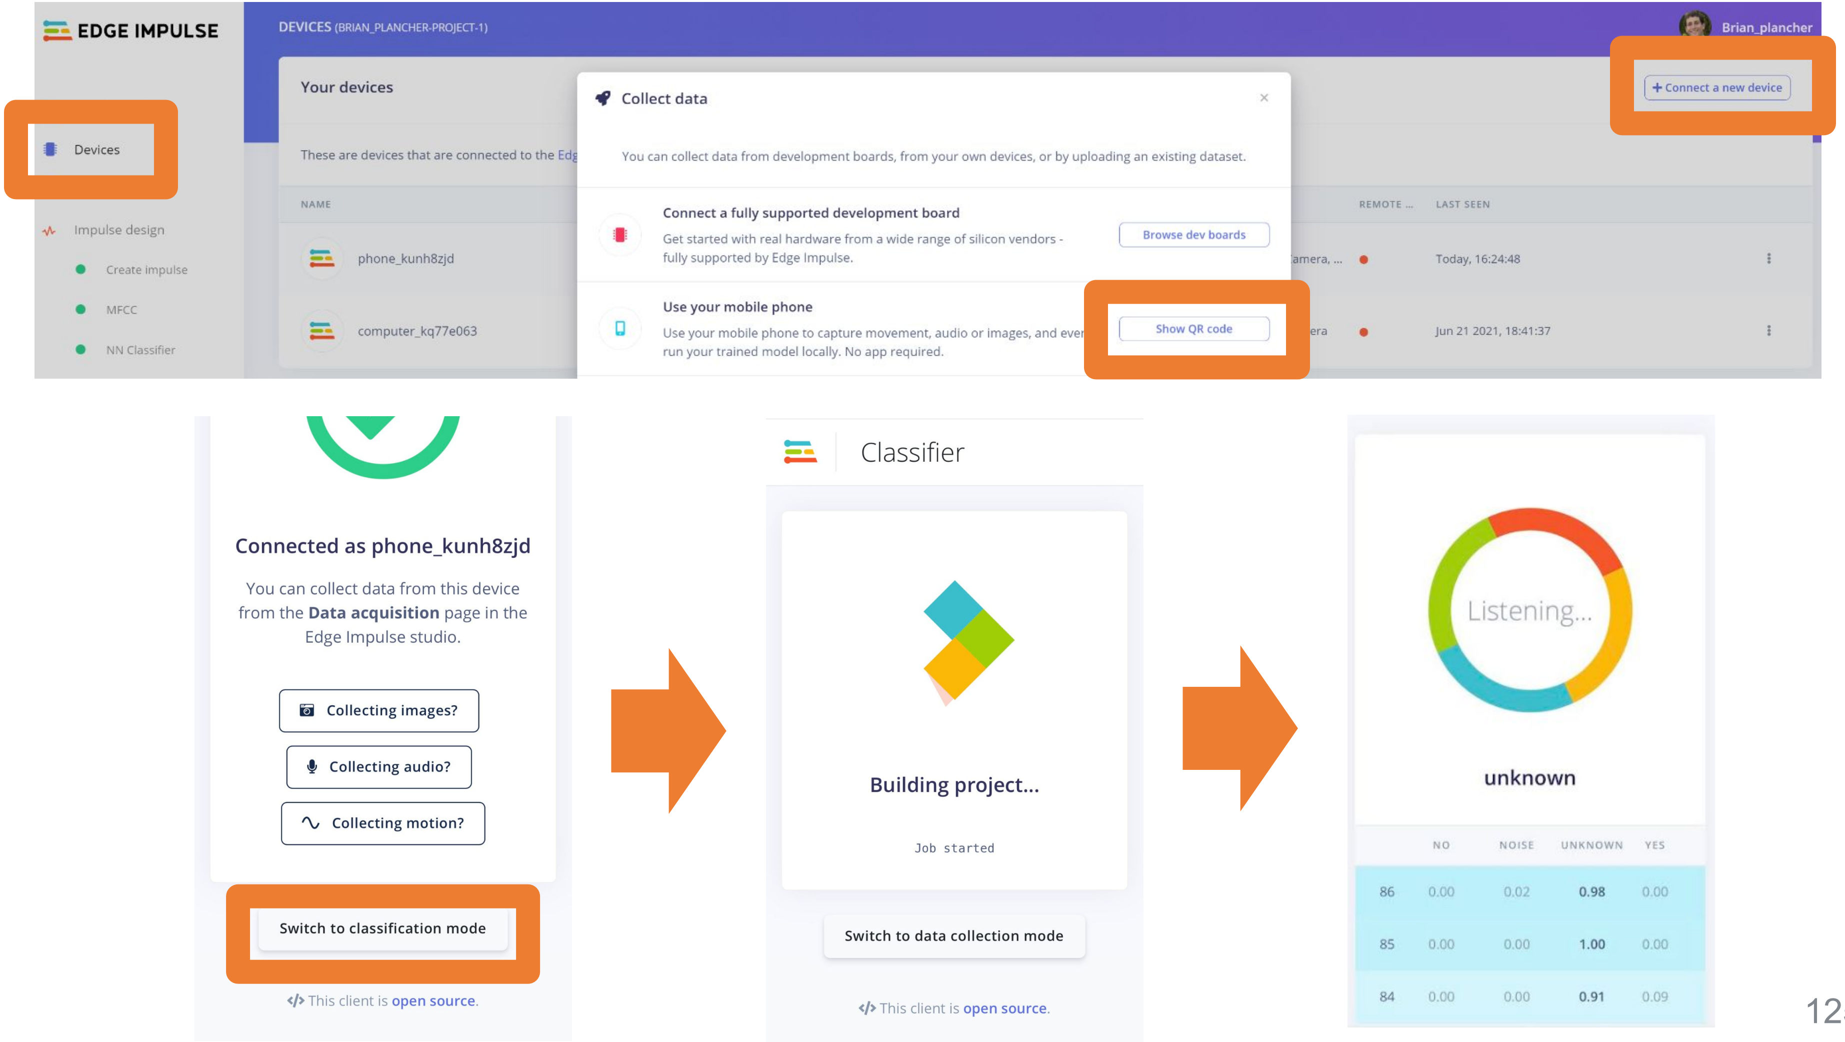
EDGE IMPULSE
DEVICES (BRIAN_PLANCHER-PROJECT-1)
Brian_plancher
Devices
Impulse design
Create impulse
MFCC
NN Classifier
Your devices
These are devices that are connected to the Edg
NAME
phone_kunh8zjd
computer_kq77e063
+ Connect a new device
REMOTE ...
LAST SEEN
Today, 16:24:48
Jun 21 2021, 18:41:37
Collect data
You can collect data from development boards, from your own devices, or by uploading an existing dataset.
Connect a fully supported development board
Get started with real hardware from a wide range of silicon vendors - fully supported by Edge Impulse.
Browse dev boards
Use your mobile phone
Use your mobile phone to capture movement, audio or images, and ever run your trained model locally. No app required.
Show QR code
Connected as phone_kunh8zjd
You can collect data from this device from the Data acquisition page in the Edge Impulse studio.
Collecting images?
Collecting audio?
Collecting motion?
Switch to classification mode
This client is open source.
Classifier
Building project...
Job started
Switch to data collection mode
This client is open source.
Listening...
unknown
NO NOISE UNKNOWN YES
86 0.00 0.02 0.98 0.00
85 0.00 0.00 1.00 0.00
84 0.00 0.00 0.91 0.09

Deploy and Test your Model

Shows the **score** for **(confidence that the current sounds is)** each of the various keywords and unknown and bolds the highest score.

unknown

	NO	NOISE	UNKNOWN	YES
86	0.00	0.02	**0.98**	0.00
85	0.00	0.00	**1.00**	0.00
84	0.00	0.00	**0.91**	0.09

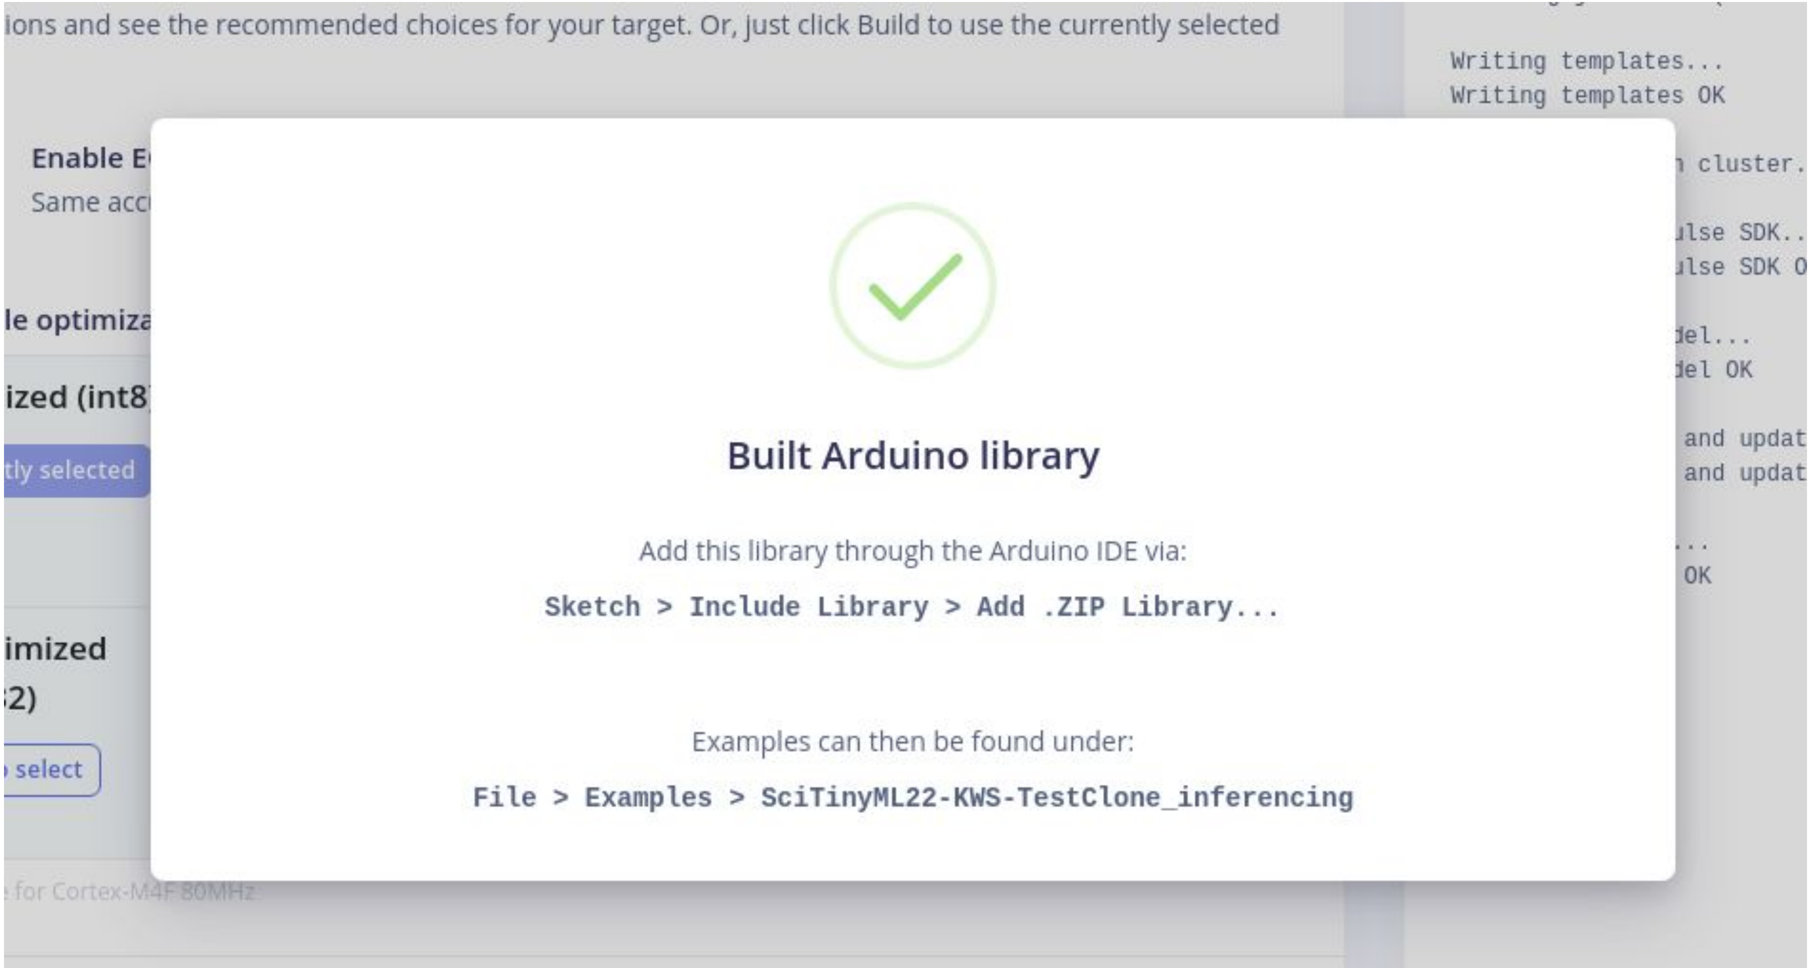
ions and see the recommended choices for your target. Or, just click Build to use the currently selected
Writing templates...
Writing templates OK
Built Arduino library
Add this library through the Arduino IDE via:
Sketch > Include Library > Add .ZIP Library...
Examples can then be found under:
File > Examples > SciTinyML22-KWS-TestClone_inferencing

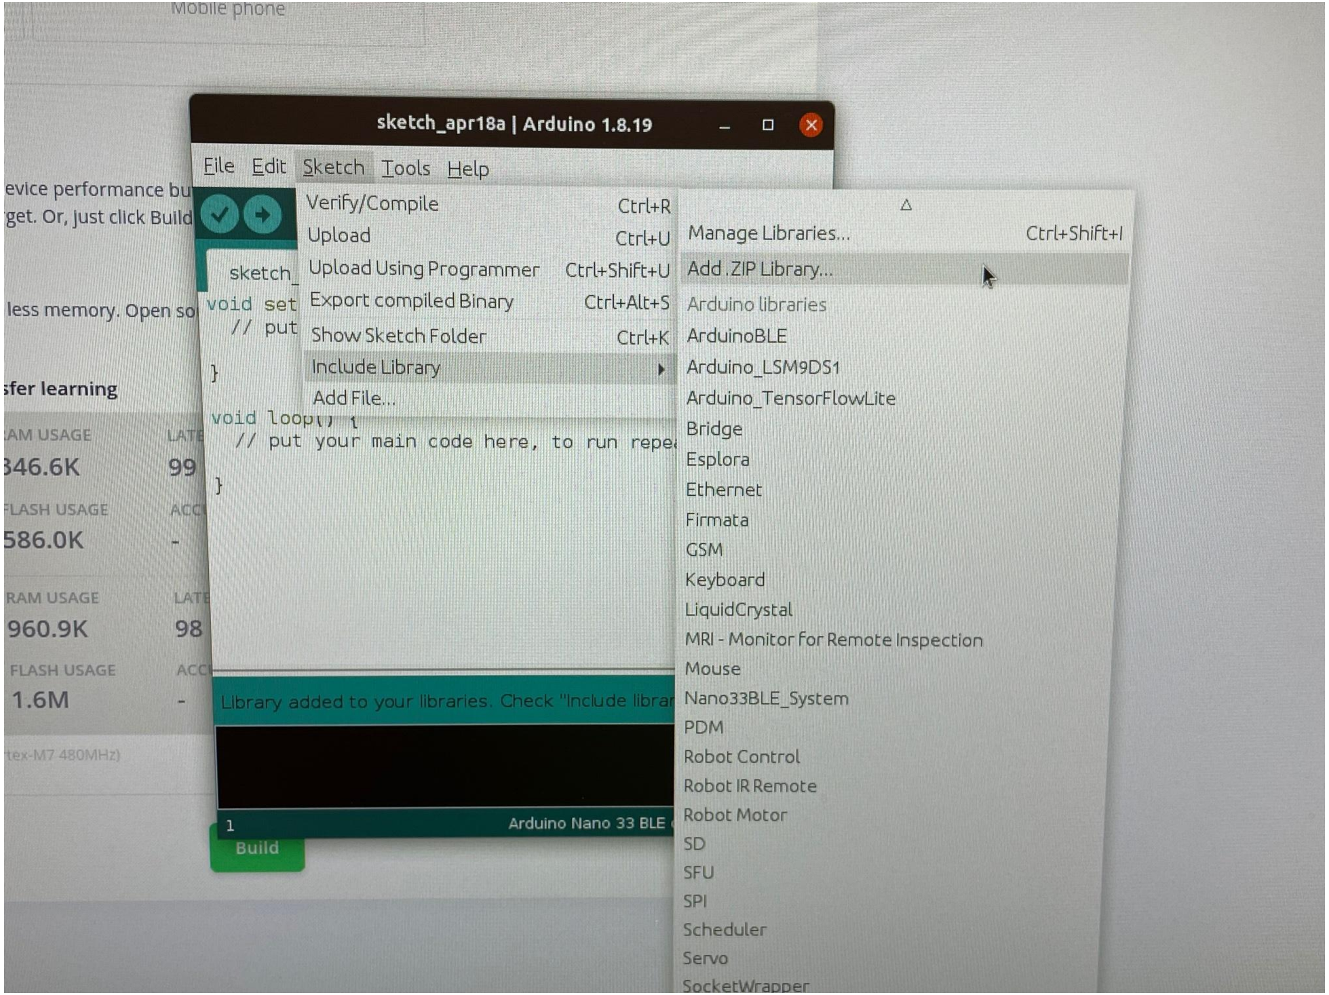
sketch_apr18a | Arduino 1.8.19
File Edit Sketch Tools Help
Verify/Compile Ctrl+R
Upload Ctrl+U
Upload Using Programmer Ctrl+Shift+U
Export compiled Binary Ctrl+Alt+S
Show Sketch Folder Ctrl+K
Include Library
Add File...
Manage Libraries... Ctrl+Shift+I
Add .ZIP Library...
Arduino libraries
ArduinoBLE
Arduino_LSM9DS1
Arduino_TensorFlowLite
Bridge
Esplora
Ethernet
Firmata
GSM
Keyboard
LiquidCrystal
MRI - Monitor for Remote Inspection
Mouse
Nano33BLE_System
PDM
Robot Control
Robot IR Remote
Robot Motor
SD
SFU
SPI
Scheduler
Servo
SocketWrapper
Library added to your libraries. Check "Include librar
Arduino Nano 33 BLE
Build

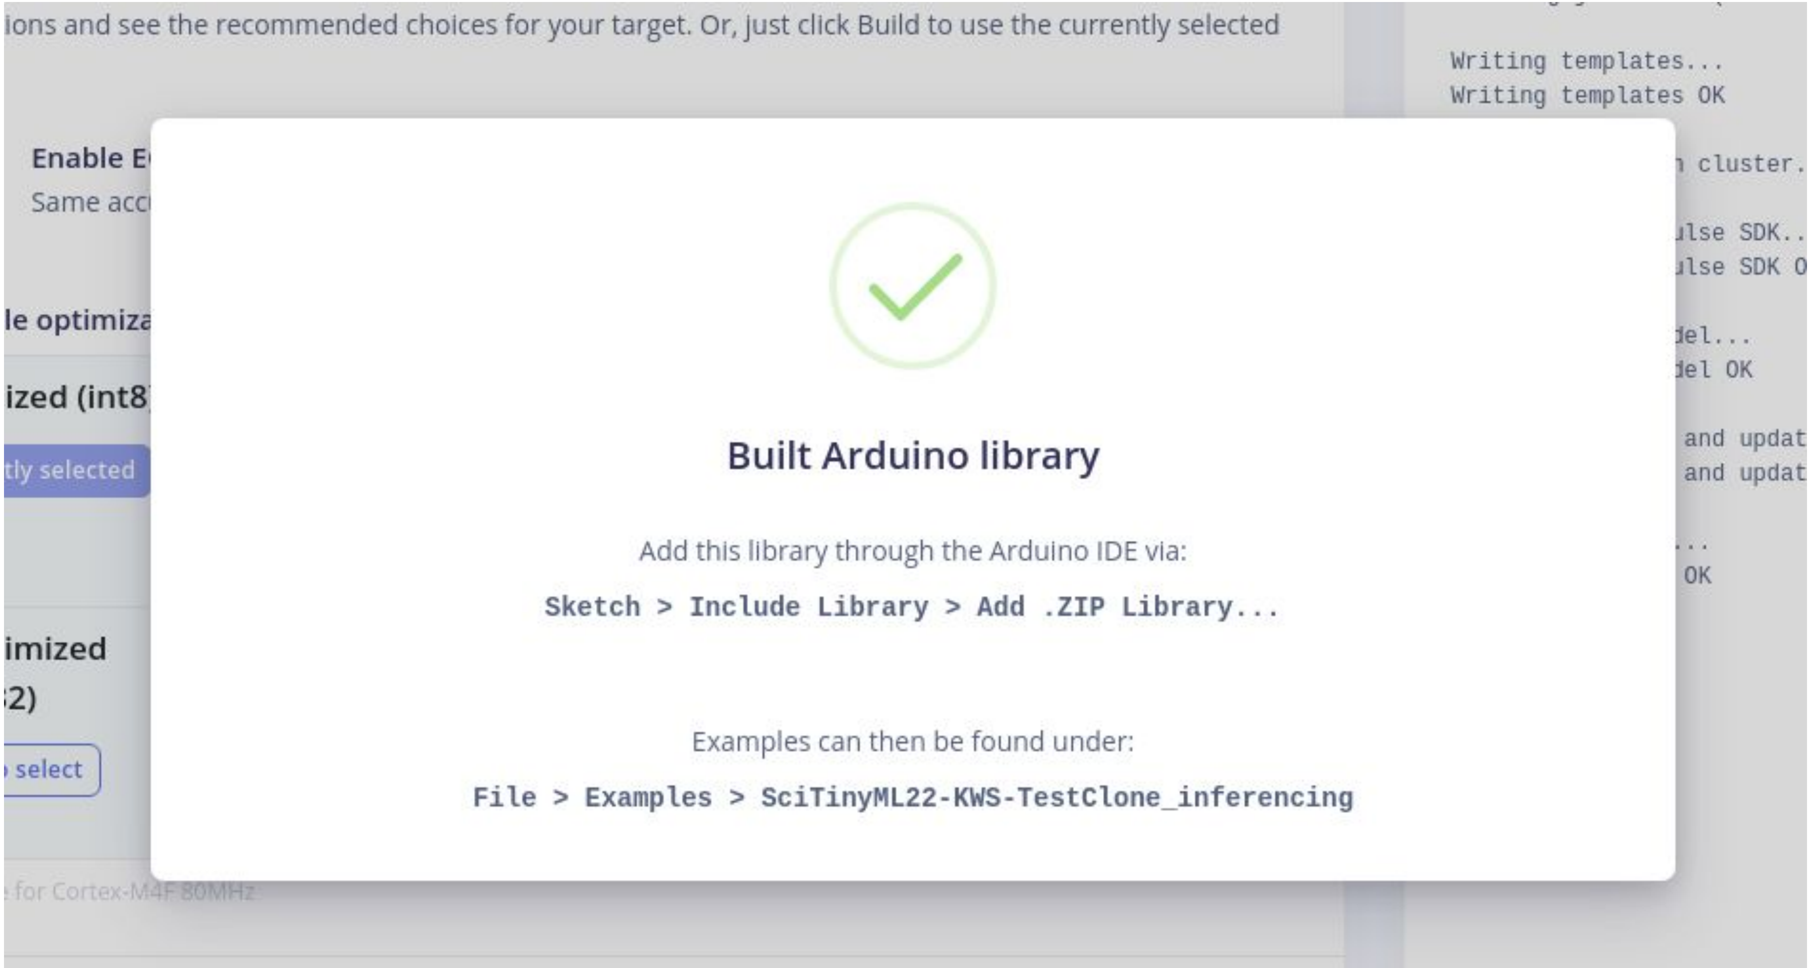
Writing templates...
Writing templates OK
Built Arduino library
Add this library through the Arduino IDE via:
Sketch > Include Library > Add .ZIP Library...
Examples can then be found under:
File > Examples > SciTinyML22-KWS-TestClone_inferencing

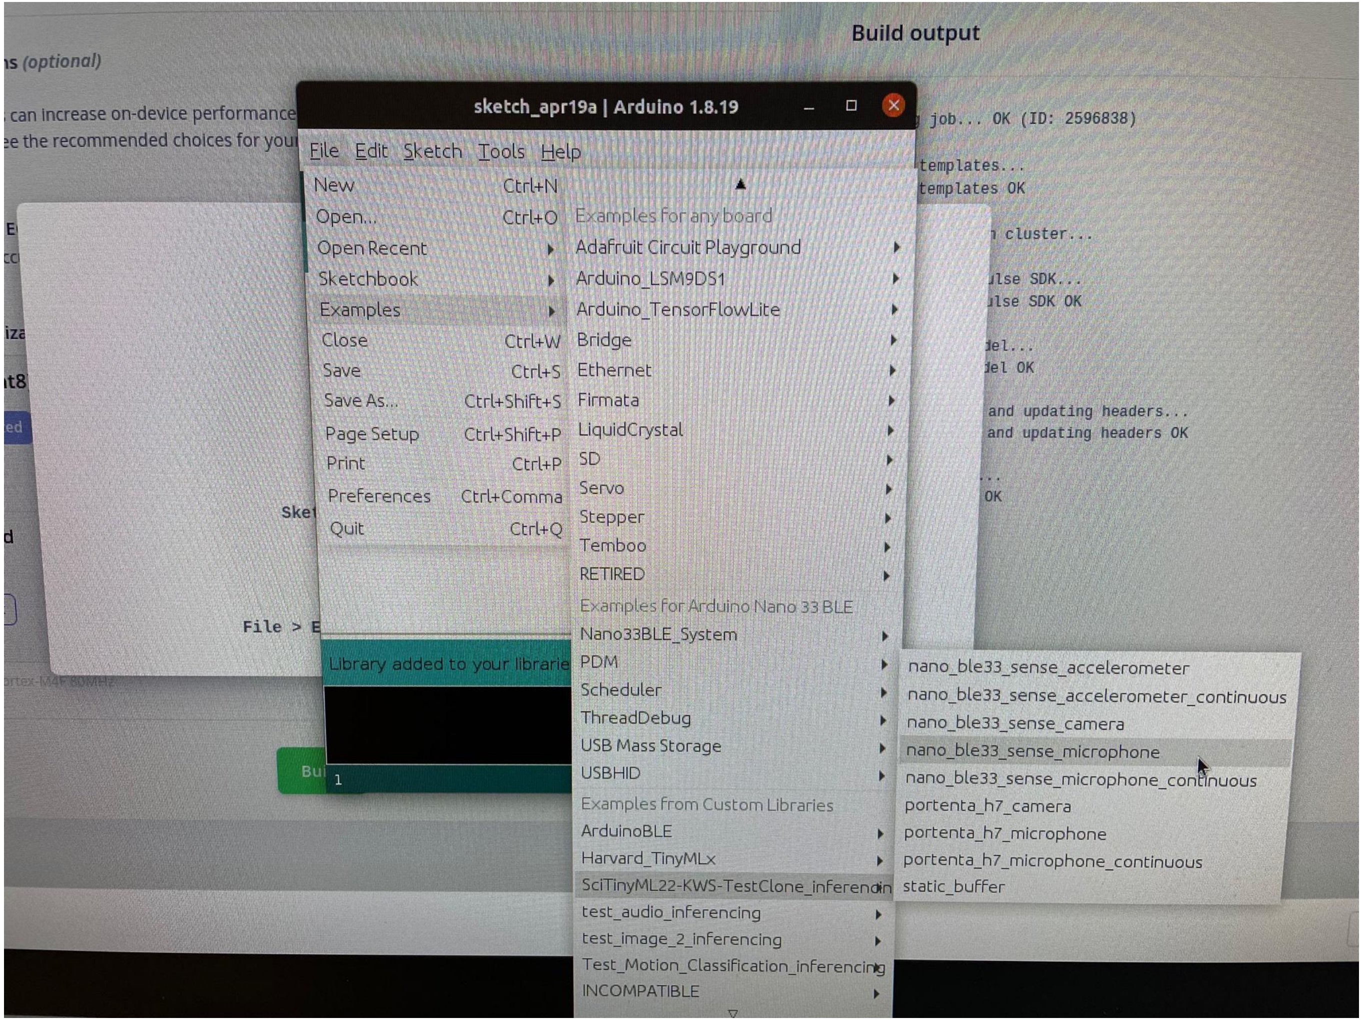

Build output
sketch_apr19a | Arduino 1.8.19
File
Edit
Sketch
Tools
Help
New Ctrl+N
Open... Ctrl+O
Open Recent
Sketchbook
Examples
Close Ctrl+W
Save Ctrl+S
Save As... Ctrl+Shift+S
Page Setup Ctrl+Shift+P
Print Ctrl+P
Preferences Ctrl+Comma
Quit Ctrl+Q
Examples for any board
Adafruit Circuit Playground
Arduino_LSM9DS1
Arduino_TensorFlowLite
Bridge
Ethernet
Firmata
LiquidCrystal
SD
Servo
Stepper
Temboo
RETIRED
Examples for Arduino Nano 33 BLE
Nano33BLE_System
PDM
Scheduler
ThreadDebug
USB Mass Storage
USBHID
Examples from Custom Libraries
ArduinoBLE
Harvard_TinyMLx
SciTinyML22-KWS-TestClone_inferencin
test_audio_inferencing
test_image_2_inferencing
Test_Motion_Classification_inferencing
INCOMPATIBLE
nano_ble33_sense_accelerometer
nano_ble33_sense_accelerometer_continuous
nano_ble33_sense_camera
nano_ble33_sense_microphone
nano_ble33_sense_microphone_continuous
portenta_h7_camera
portenta_h7_microphone
portenta_h7_microphone_continuous
static_buffer
Library added to your librarie
job... OK (ID: 2596838)
templates...
templates OK
and updating headers...
and updating headers OK

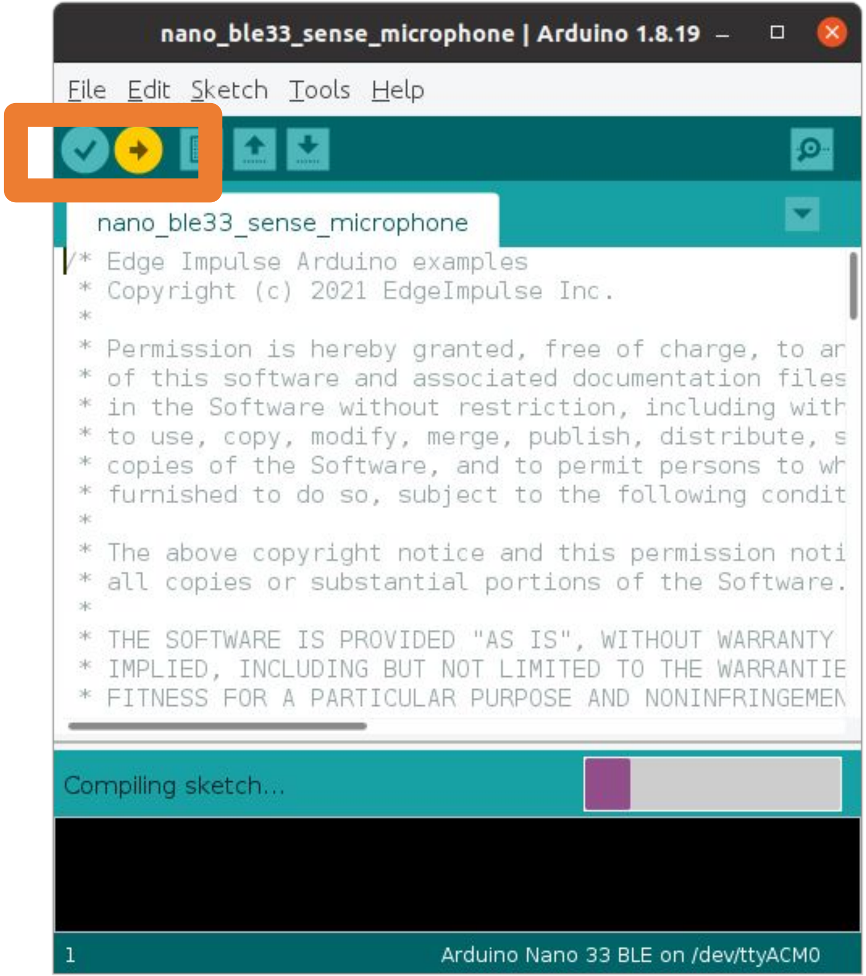
nano_ble33_sense_microphone | Arduino 1.8.19
File Edit Sketch Tools Help
nano_ble33_sense_microphone
/* Edge Impulse Arduino examples
* Copyright (c) 2021 EdgeImpulse Inc.
*
* Permission is hereby granted, free of charge, to an
* of this software and associated documentation files
* in the Software without restriction, including with
* to use, copy, modify, merge, publish, distribute, s
* copies of the Software, and to permit persons to wh
* furnished to do so, subject to the following condit
*
* The above copyright notice and this permission noti
* all copies or substantial portions of the Software.
*
* THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY
* IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIE
* FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMEN
Compiling sketch...
1
Arduino Nano 33 BLE on /dev/ttyACM0

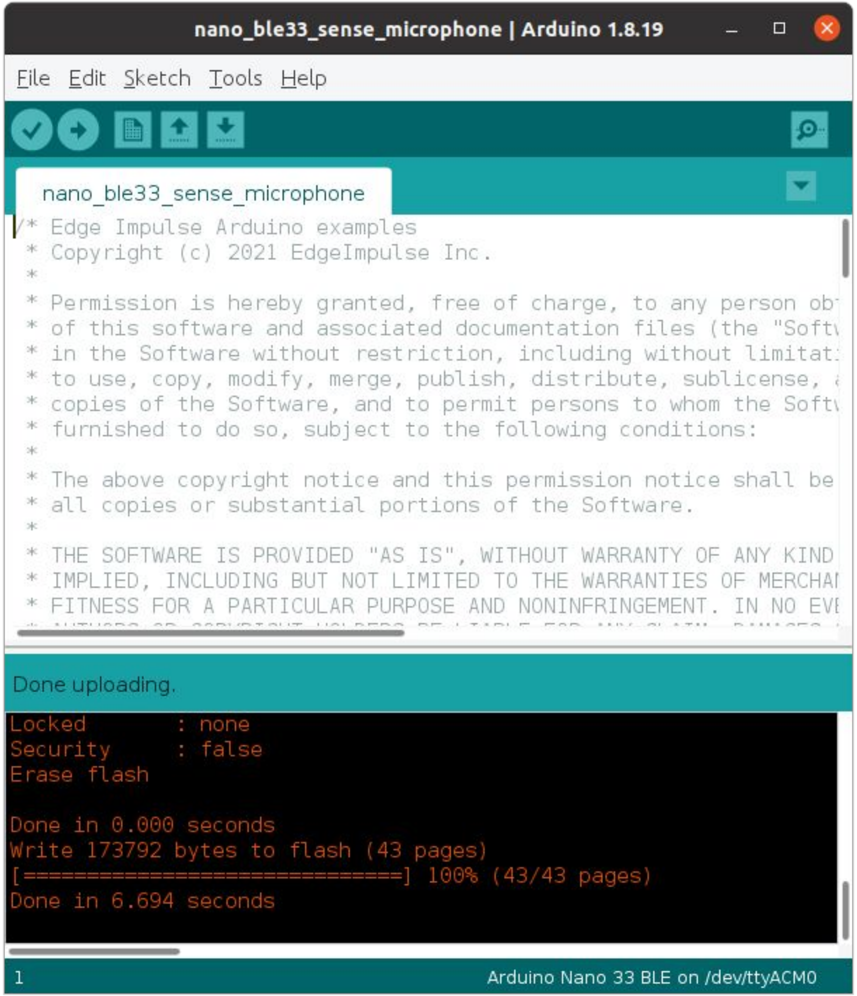
nano_ble33_sense_microphone | Arduino 1.8.19
File Edit Sketch Tools Help
nano_ble33_sense_microphone
/* Edge Impulse Arduino examples
* Copyright (c) 2021 EdgeImpulse Inc.
Done uploading.
Locked : none
Security : false
Erase flash
Done in 0.000 seconds
Write 173792 bytes to flash (43 pages)
[==============================] 100% (43/43 pages)
Done in 6.694 seconds
Arduino Nano 33 BLE on /dev/ttyACM0

An error occurred while uploading the sketch

```
/home/plancher/Arduino/libraries/test_image_2_inferencing/src/edge-impulse-sdk/CMSIS/NN/Source/PoolingFunctions/arm_pool_q7_HWC
         *__SIMD32(pCnt)++ = __QADD16(vo2, in);
         ^
/home/plancher/Arduino/libraries/test_image_2_inferencing/src/edge-impulse-sdk/tensorflow/lite/core/api/op_resolver.cpp: In fun
/home/plancher/Arduino/libraries/test_image_2_inferencing/src/edge-impulse-sdk/tensorflow/lite/core/api/op_resolver.cpp:34:20:
       builtin_code < BuiltinOperator_MIN) {
       ~~~~~~~~~~~~~^~~~~~~~~~~~~~~~~~~~~
Sketch uses 224024 bytes (22%) of program storage space. Maximum is 983040 bytes.
Global variables use 58672 bytes (22%) of dynamic memory, leaving 203472 bytes for local variables. Maximum is 262144 bytes.
An error occurred while uploading the sketch

Device unsupported
```

Double Tap Reset for Bootloader Mode!

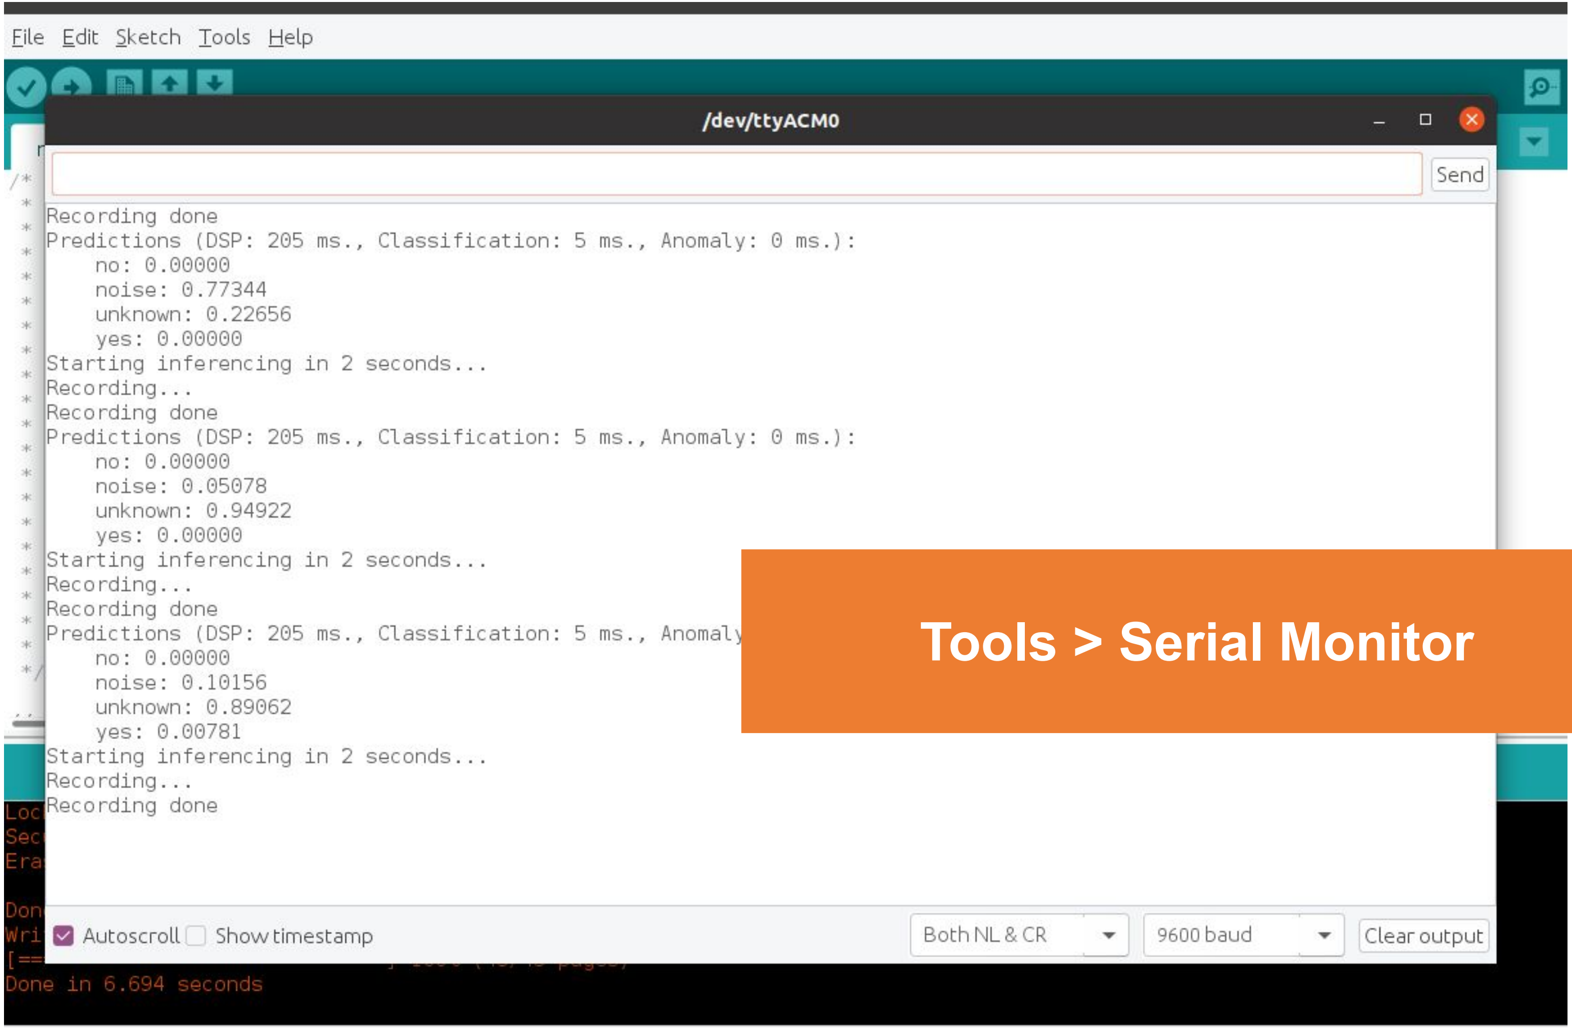

Tools > Serial Monitor

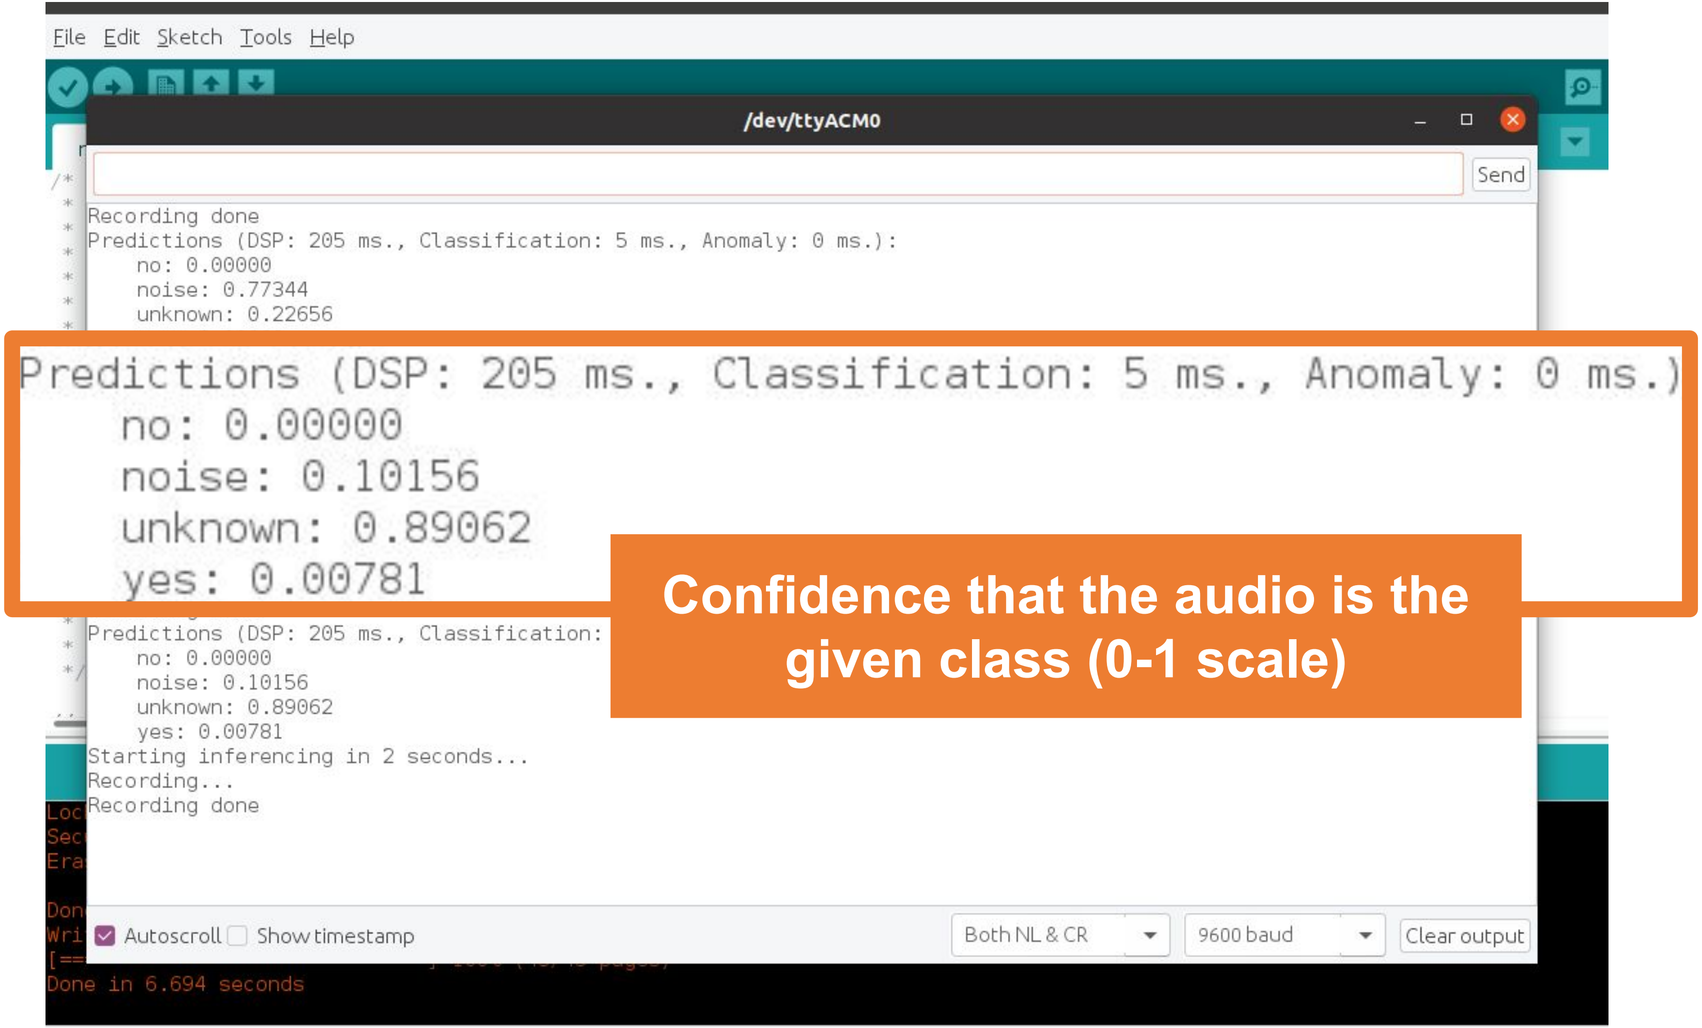
File Edit Sketch Tools Help
/dev/ttyACM0
Send
Recording done
Predictions (DSP: 205 ms., Classification: 5 ms., Anomaly: 0 ms.):
no: 0.00000
noise: 0.77344
unknown: 0.22656
Predictions (DSP: 205 ms., Classification: 5 ms., Anomaly: 0 ms.)
no: 0.00000
noise: 0.10156
unknown: 0.89062
yes: 0.00781
Confidence that the audio is the given class (0-1 scale)
Predictions (DSP: 205 ms., Classification:
no: 0.00000
noise: 0.10156
unknown: 0.89062
yes: 0.00781
Starting inferencing in 2 seconds...
Recording...
Recording done
Autoscroll
Show timestamp
Both NL & CR
9600 baud
Clear output
Done in 6.694 seconds

Today’s Agenda

- A Quick Review of What We’ve Learned
- Data Engineering for KWS
- Hands-on KWS Data Collection with Edge Impulse
- (Hands-on) Data Preprocessing for KWS
- Deploying our Model onto our Arduino
- **Summary**

Machine Learning

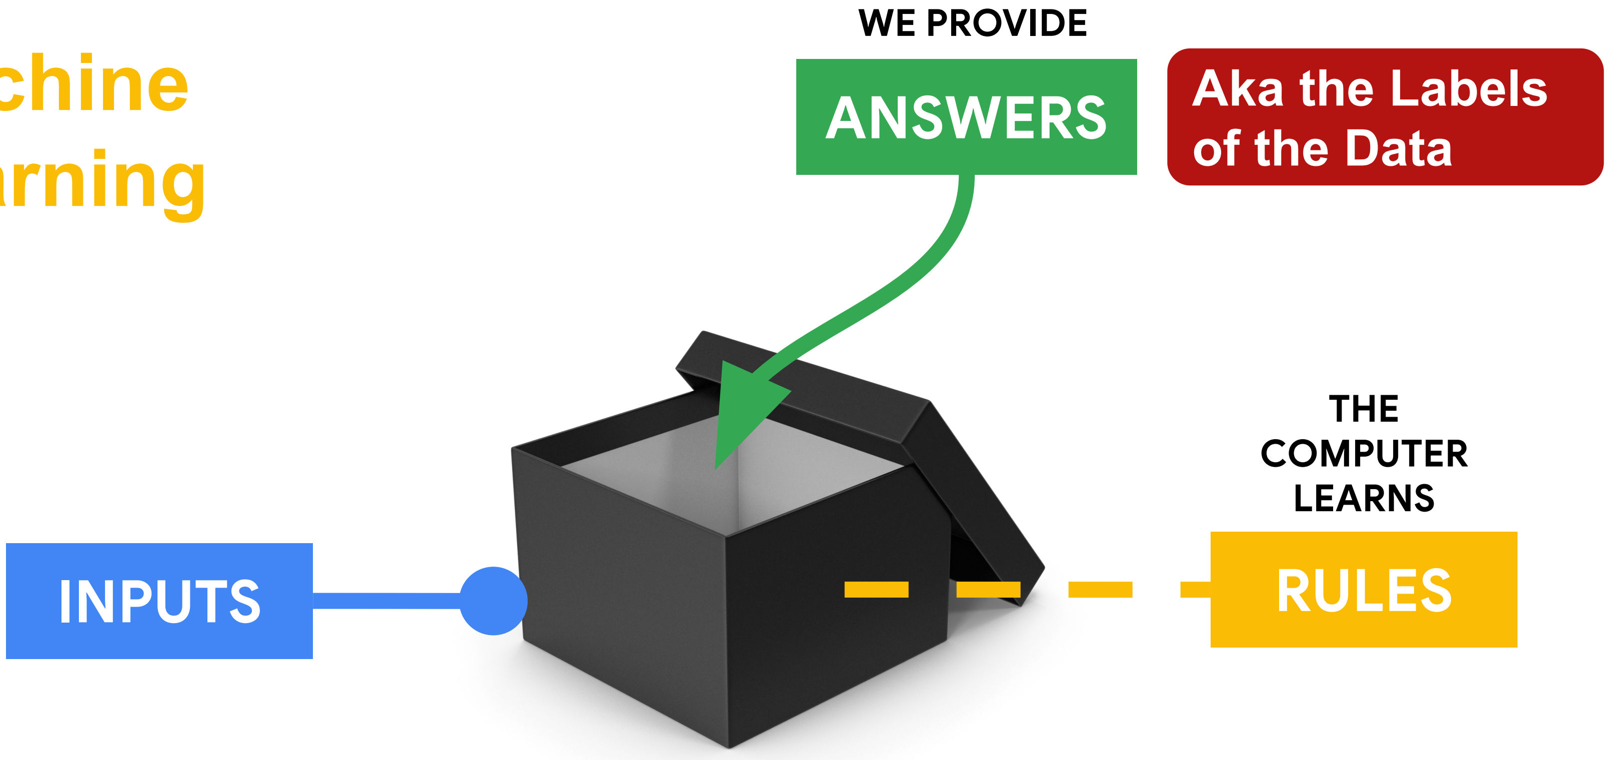

Deep Learning
with **Neural Networks**

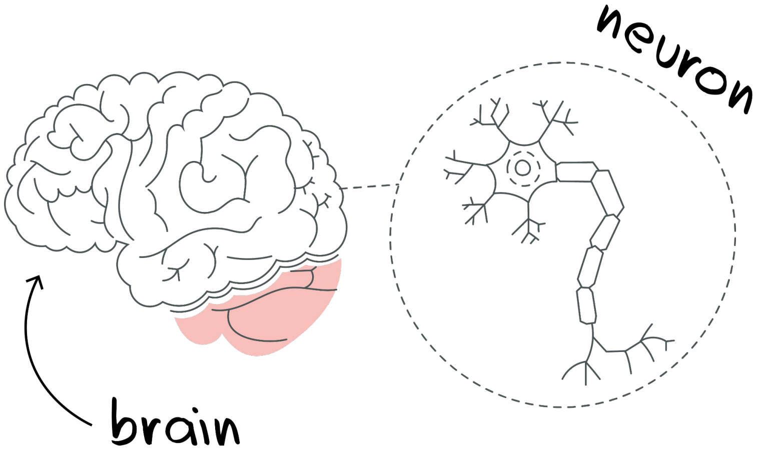

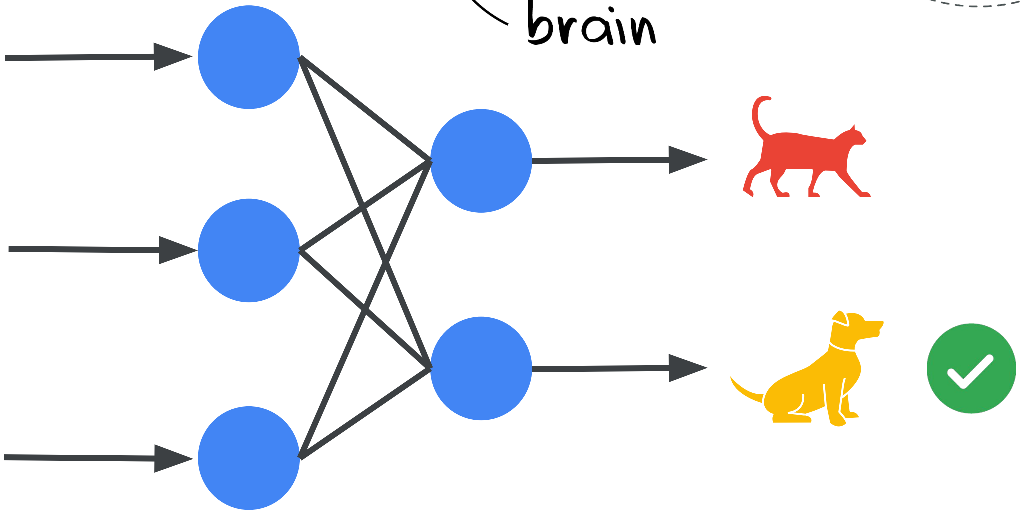

Colab Link

Features can be found with **Convolutions**

-1	0	1
-2	0	2
-1	0	1

Features

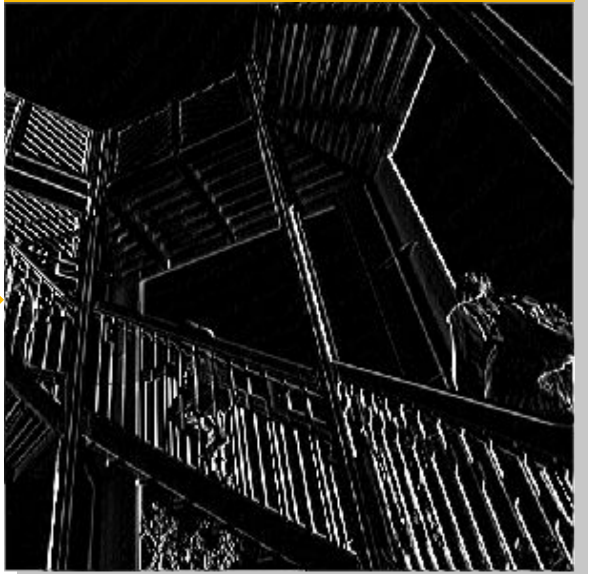

The TinyML Workflow

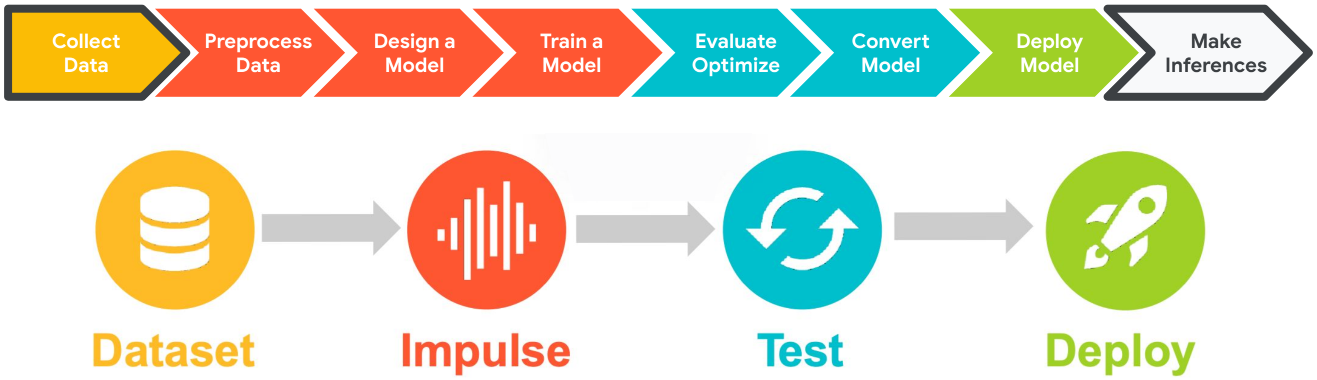

The TinyML Workflow

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

Who will use your ML model?

Where will your ML model be used?

Why will your ML model be used?
Why those Keywords?

Training Set	Validation Set	Test Set

The TinyML Workflow

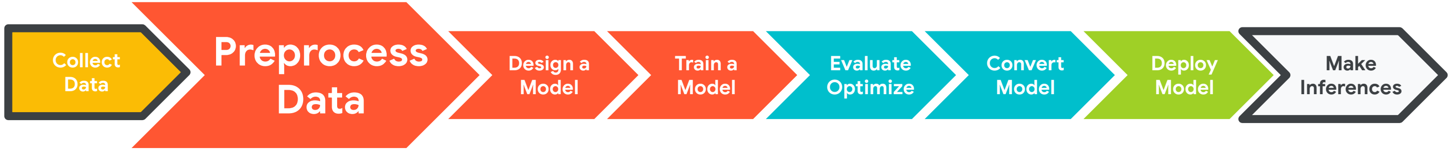

FFT, Spectrogram, MFCC

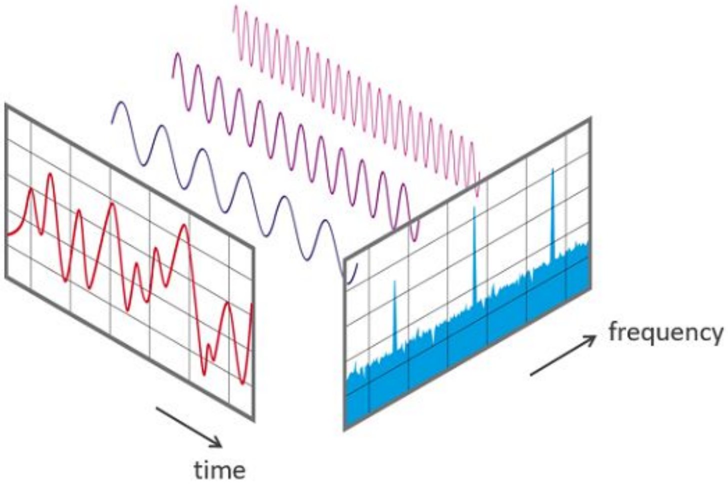

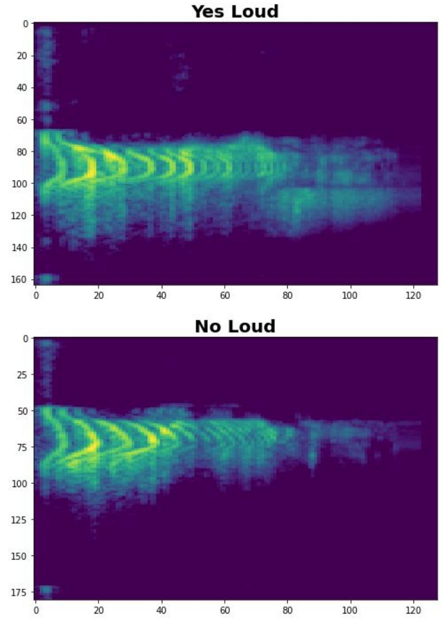

The TinyML Workflow

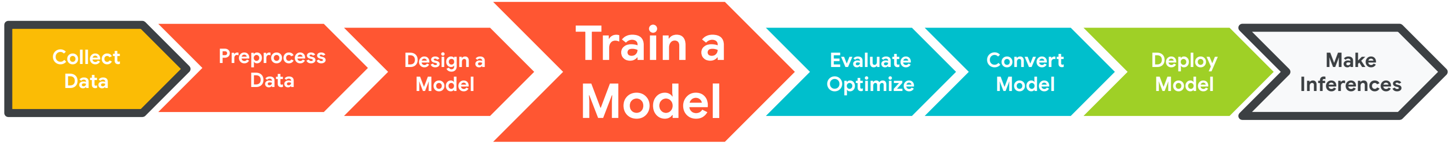

Confusion Matrix	Actual Output = Yes	Actual Output = No
Predicted Output = Yes	# of True Positive	# of False Positive *Type 1 Error*
Predicted Output = No	# of False Negative *Type 2 Error*	# of True Negative

The TinyML Workflow

Quantization

Reduces the precision of numbers used in a model which results in:

- **smaller model size**
- **faster computation**

The TinyML Workflow

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

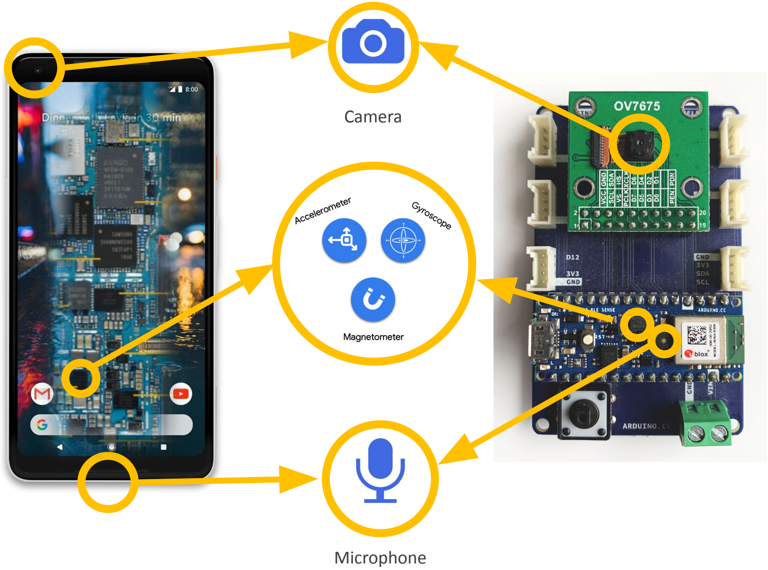

Edge Impulse Simplifies Deployment

Better Data = Better Models!

African Regional Workshop on SciTinyML: Scientific Use of Machine Learning on Low-Power Devices

25-29 April 2022
Online

Further information:
http://indico.ictp.it/event/9792/
smr3708@ictp.it

Data Pre-Processing for Hands-on Keyword Spotting

Brian Plancher
Harvard John A. Paulson School of Engineering and Applied Sciences
Barnard College, Columbia University
brianplancher.com

https://arduino.github.io/arduino-cli/0.21/installation/

Edge Impulse CLI Notes:

1. Install the [Arduino CLI](https://arduino.github.io/arduino-cli/0.21/installation/)
   a. On linux:
   ```
   curl -fsSL https://raw.githubusercontent.com/arduino/arduino-cli/master/install.sh | sh
   ```
   b. On mac:
   ```
   brew update
   brew install arduino-cli
   ```
   c. Or view the link for binaries
2. Add to your `.bashrc`:
   ```
   # Arduino (CLI)
   export PATH="ARDUINO_INSTALL_LOCATION/bin:$PATH"
   ```
   Where `ARDUINO_INSTALL_LOCATION` is e.g.,: `$HOME/Documents/arduino-1.8.19`

https://docs.edgeimpulse.com/docs/edge-impulse-cli/cli-installation

Edge Impulse CLI Notes:

1. Install the Edge Impulse CLI
   a. Install Node.js by following the link or on Linux:

      ```
      curl -sL https://deb.nodesource.com/setup_14.x | sudo -E bash -
      sudo apt-get install -y nodejs
      ```
   b. Run: `npm install -g edge-impulse-cli --force`
   c. Add to your `.bashrc`:

      ```
      # EI (CLI)
      export PATH="$HOME/.npm-global/bin:$PATH"
      ```
2. Run `edge-impulse-daemon --clean` to start the daemon and then follow the instructions in the terminal to add it to your current project using your edge impulse account!

https://docs.edgeimpulse.com/docs/edge-impulse-cli/cli-installation

Edge Impulse CLI Notes:

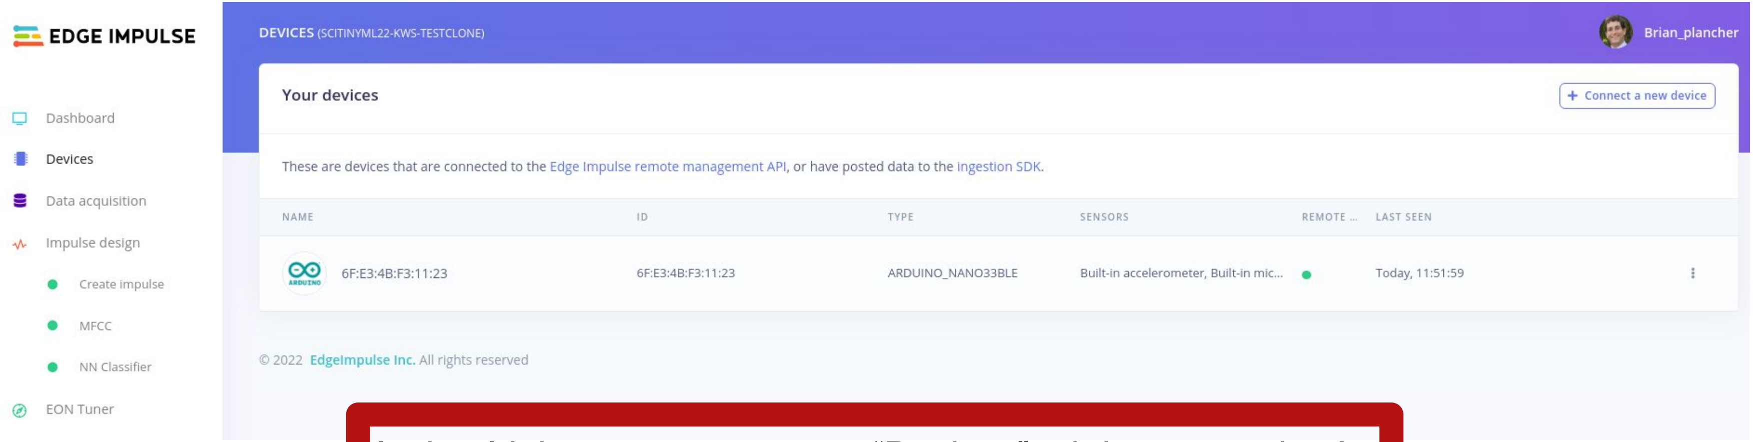

It should then appear on your "Devices" tab in your project!

And then if you go to "Data Acquisition" you should be able to proceed as you would with the standard instructions!