

Safer Roads, Powered by Machine Learning

...

The Future of Safer Roads Is Here: Automatic
Wheel Lug Nut Detection Using Machine
Learning at the Edge.

INTRODUCTION

Wheel lug nuts are such a tiny part of the overall automobile assembly that they're easy to overlook, yet serve a critical function in the safe operation of an automobile. In fact, it is not safe to drive even with one lug nut missing. A single missing lug nut will cause increased pressure on the wheel, in turn causing damage to the wheel bearings, studs, and make the other lug nuts fall off.

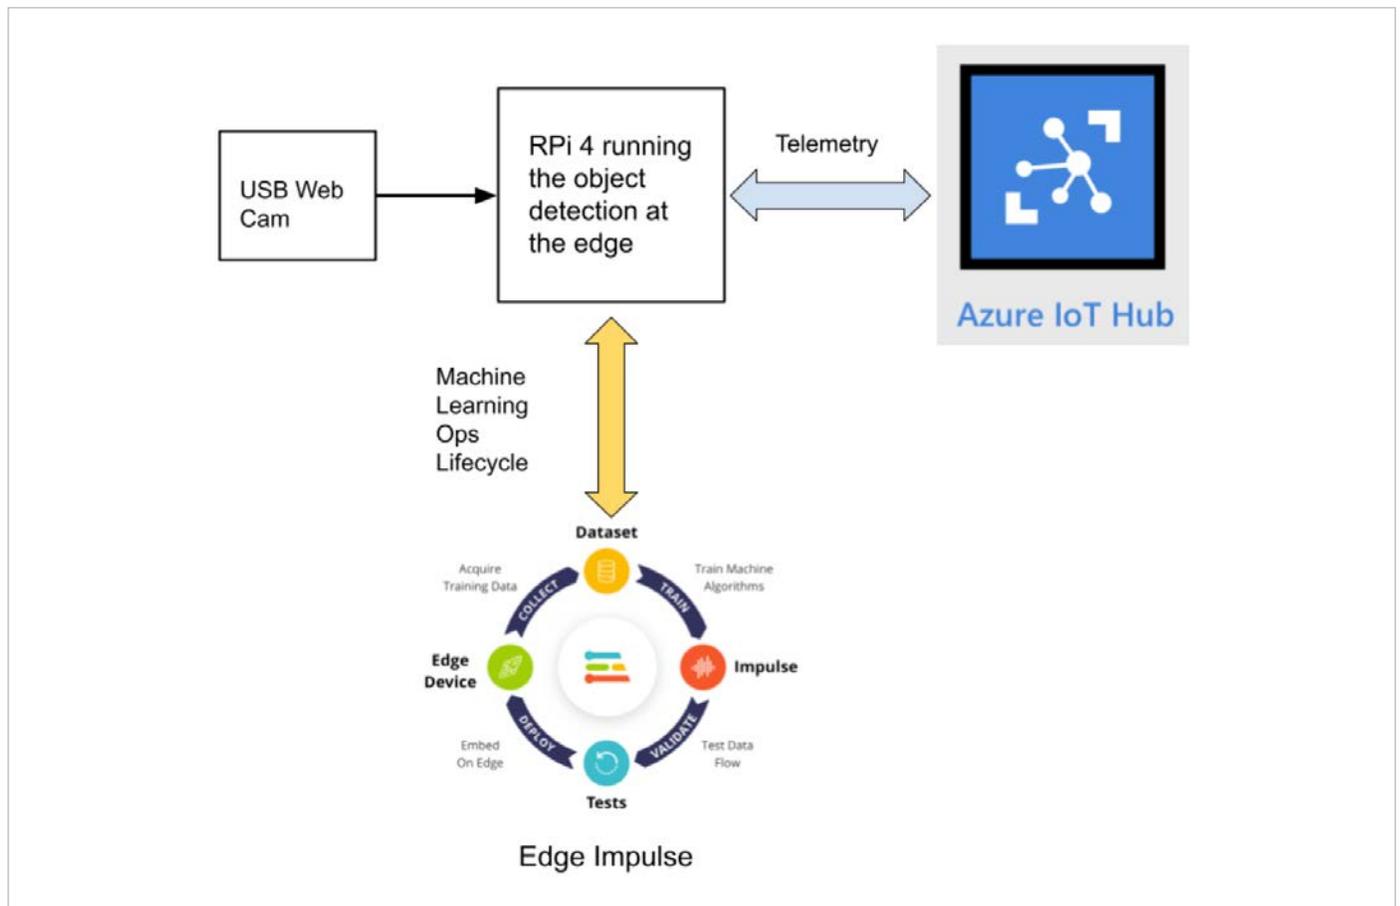
Over the years there have been a number of documented safety recalls and issues around wheel lug nuts. In some cases, it was only identified after the fact that the automobile manufacturer had installed incompatible lug

nut types with the wheel or had been inconsistent in installing the right type of lug nut. Even after delivery, after years of wear and tear, the lug nuts may become loose and may even fall off which would cause instability for an automobile to be in service. To reduce these incidents of quality control at manufacturing and maintenance in the field, there is a huge opportunity to leverage machine learning at the edge to automate wheel lug nut detection.

This motivated us to create a proof-of-concept reference project for automating wheel lug nut detection by easily putting together a USB webcam, Raspberry Pi 4, Microsoft Azure

IoT, and Edge Impulse, creating an end-to-end wheel lug nut detection system using object detection. This example use case and other derivatives will find a home in many industrial IoT scenarios where embedded machine learning can help improve the efficiency of factory automation and quality control processes including predictive maintenance.

This reference project will serve as a guide for quickly getting started with Edge Impulse on the Raspberry Pi 4 and Azure IoT, to train a model that detects lug nuts on a wheel and sends inference conclusions to Azure IoT as shown in the block diagram below:



DESIGN CONCEPT: EDGE IMPULSE AND AZURE IOT

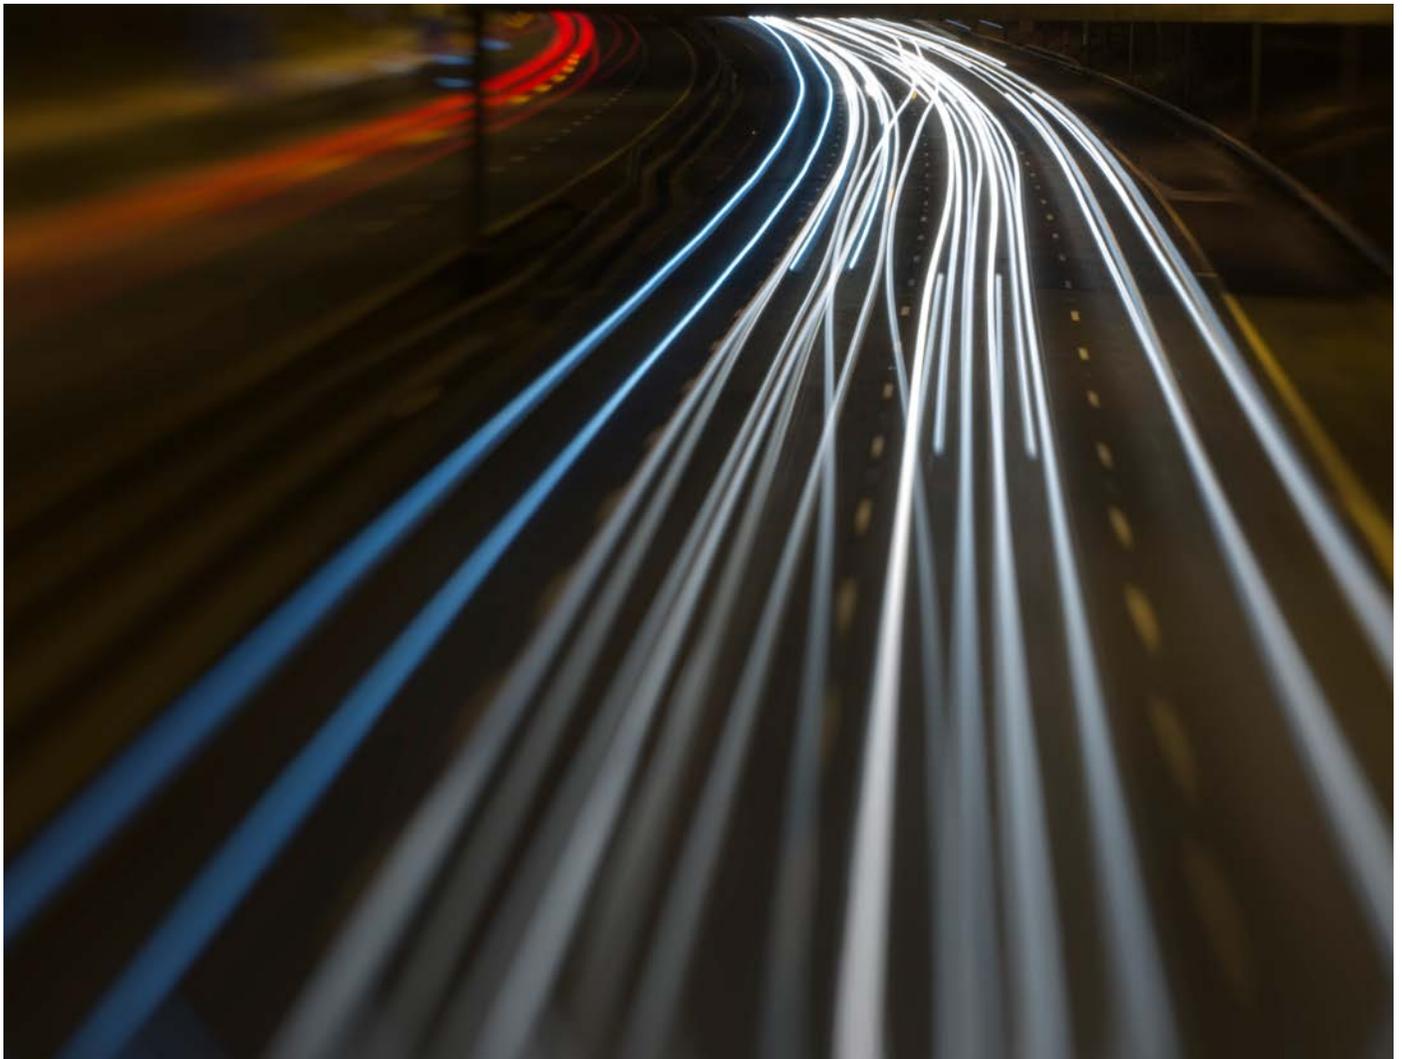
[Edge Impulse](#) is an embedded machine learning platform that allows you to manage the entire Machine Learning Ops (MLOps) lifecycle, which includes 1) Data acquisition, 2) Signal processing, 3) ML training, 4) Model testing, and 5) Creating a deployable model that can run efficiently on an edge device.

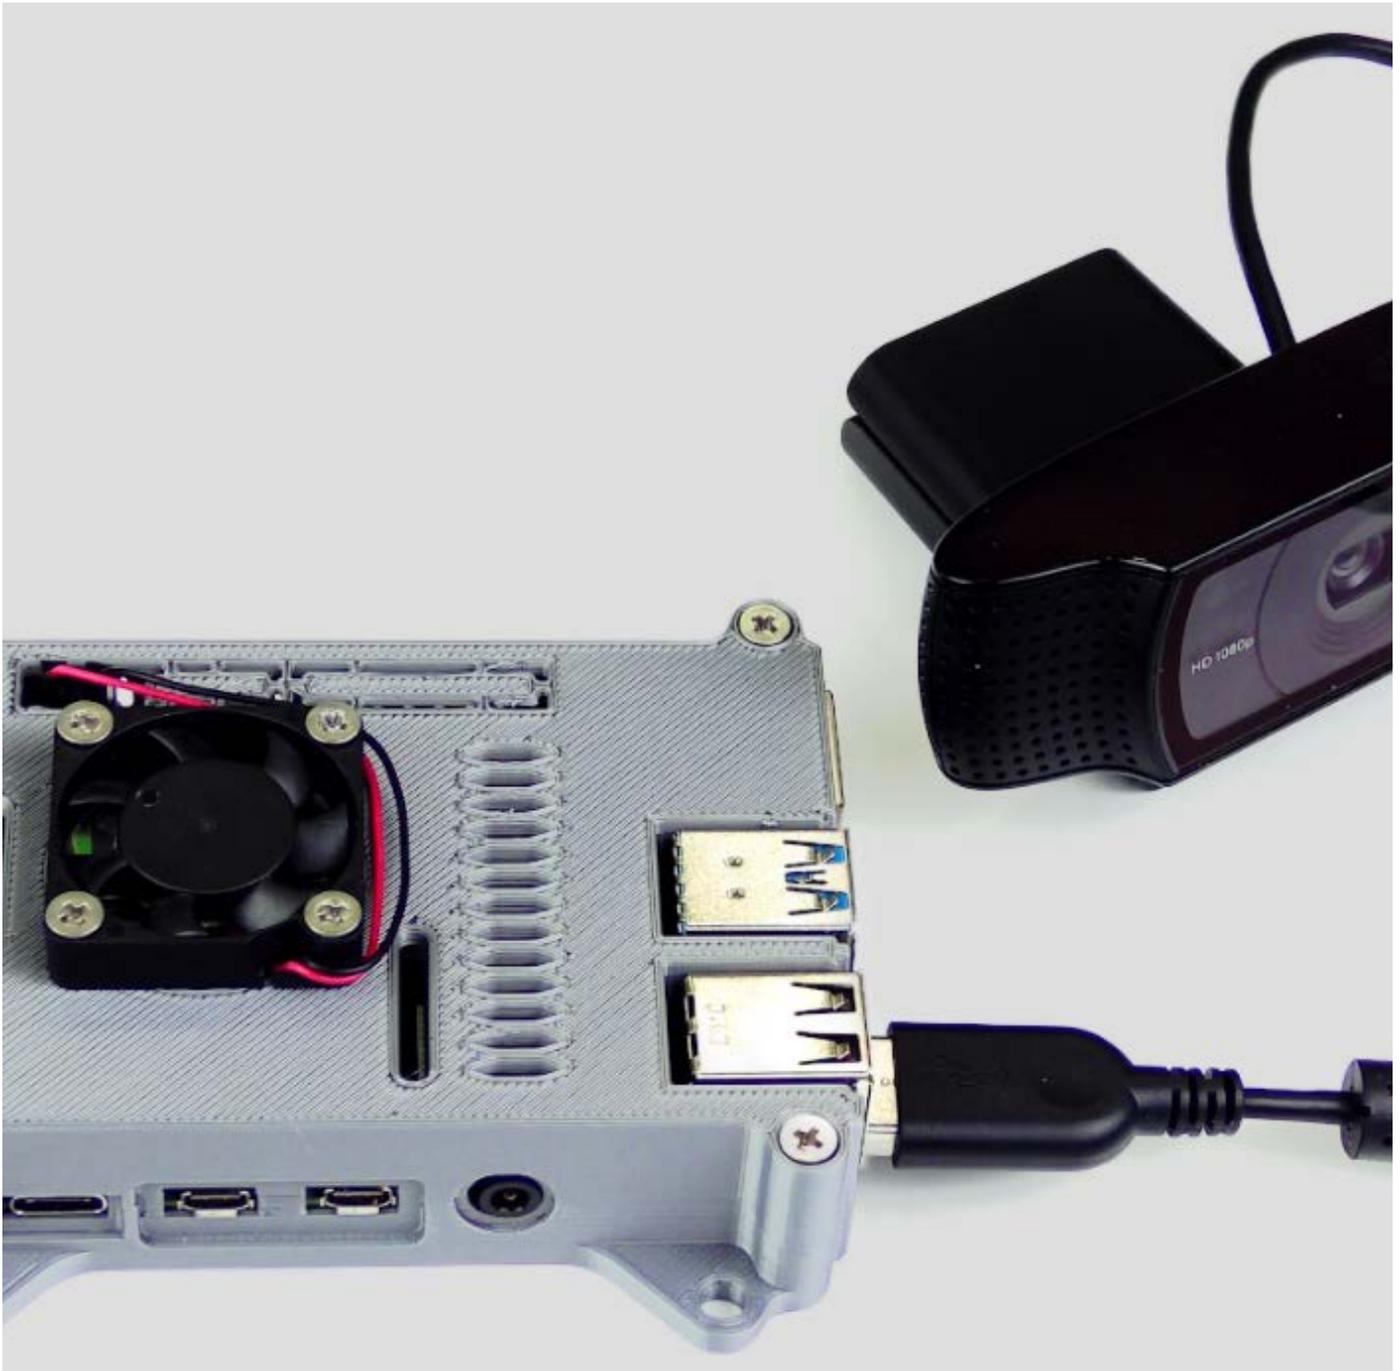
For the edge device, we chose to use the Raspberry Pi 4 due to its ubiquity and available processing power for efficiently running more sophisticated machine learning models such as object detection. By running the object detection model on the Raspberry Pi 4, we can optimize the network bandwidth connection

to Azure IoT for robustness and scalability by only sending the inference conclusions, i.e. “How many lug nuts are on the wheel?” Once the inference conclusions are available at the Azure IoT level, it becomes straightforward to feed these results into your business applications that can leverage other Azure services such as [Azure Stream Analytics](#) and [Power BI](#).

In the next sections we'll discuss how you can set this up yourself with the following items:

- [Raspberry Pi 4](#)
- USB webcam (such as a Logitech HD webcam)
- [Edge Impulse](#) account
- [Azure IoT Hub](#) instance





SETTING UP THE HARDWARE

We begin by setting up the Raspberry Pi 4 to connect to a Wi-Fi network for our network connection, configuring it for camera support, and installing the Edge Impulse Linux CLI (command line interface) tools on the Raspberry Pi 4. This

will allow the Raspberry Pi 4 to directly connect to Edge Impulse for data acquisition and finally, deployment of the wheel lug nut detection model.

For starters, you'll need a Raspberry Pi 4 with an up-to-date Raspberry Pi OS image that can be found [here](#). After flashing this image to an SD card and adding a file named 'wpa_supplicant.conf':

```
ctrl_interface=DIR=/var/run/wpa_supplicant GROUP=netdev
update_config=1
country=<Insert 2 letter ISO 3166-1 country code here>
network={
  ssid="<Name of your wireless LAN>"
  psk="<Password for your wireless LAN>"
}
```

along with an empty file named 'ssh' (both within the '/boot' directory), you can go ahead and power up the board. Once you've successfully SSH'd into the device with:

```
ssh pi@<IP_ADDRESS>
```

and the password 'raspberrypi', it's time to install the dependencies for the Edge Impulse Linux SDK. Simply run the next three commands to set up the NodeJS environment and everything else that's required for the edge-impulse-linux wizard:

```
curl -sL https://deb.nodesource.com/setup_12.x | sudo bash -
sudo apt install -y gcc g++ make build-essential nodejs sox gstreamer1.0-tools
gstreamer1.0-plugins-good gstreamer1.0-plugins-base gstreamer1.0-plugins-
base-apps
npm config set user root && sudo npm install edge-impulse-linux -g --un-
safe-perm
```

For more details on setting up the Raspberry Pi 4 with Edge Impulse, [visit this link](#).

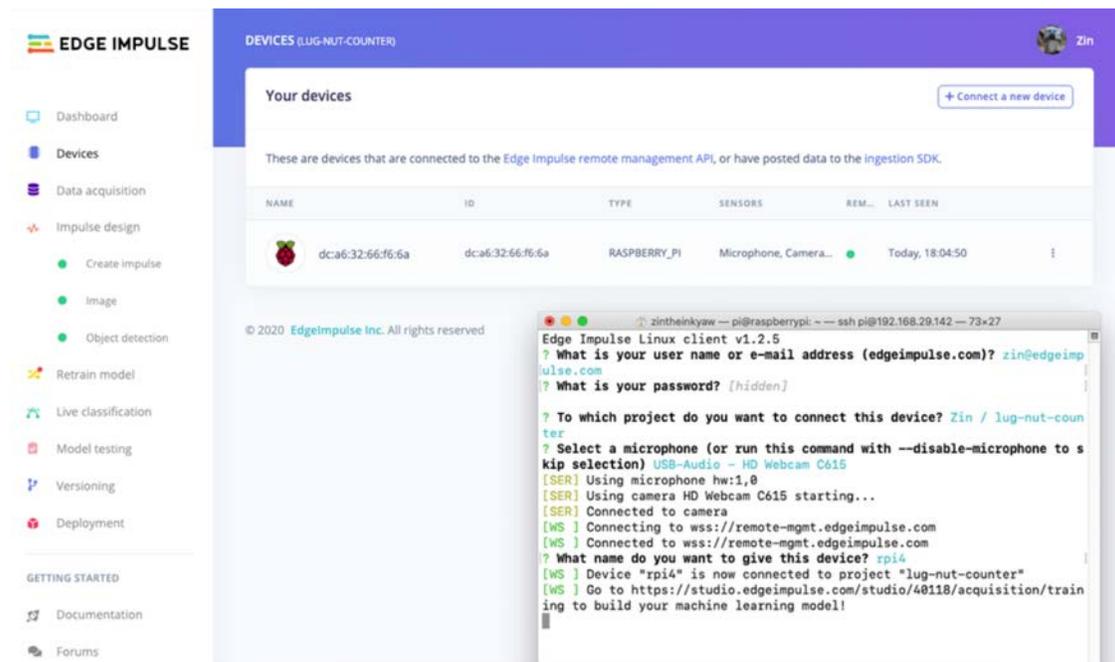
Since this project deals with images, we'll need some way to capture them. The wizard supports both the Pi camera modules and standard USB webcams, so make sure to enable the camera module first with:

```
sudo raspi-config
```

if you plan on using one. With that completed, go to Edge Impulse and create a new project, then run the wizard with:

```
edge-impulse-linux
```

and make sure your device appears within the Edge Impulse Studio's device section after logging in and selecting your project.



The screenshot displays the Edge Impulse Studio interface. On the left is a navigation sidebar with options like Dashboard, Devices, Data acquisition, Impulse design, Retrain model, Live classification, Model testing, Versioning, Deployment, GETTING STARTED, Documentation, and Forums. The main area shows the 'DEVICES (LUG-NUT-COUNTER)' section with a table of connected devices. One device is listed with the name 'lug-nut-counter', ID 'dca6:32:66:f6:6a', type 'RASPBERRY_PI', and sensors 'Microphone, Camera...'. Below the table, a terminal window shows the execution of the 'edge-impulse-linux' command, which prompts for user name, password, project name, and microphone selection, and finally reports that the device 'rpi4' is connected to the project 'lug-nut-counter'.

NAME	ID	TYPE	SENSORS	REM...	LAST SEEN
lug-nut-counter	dca6:32:66:f6:6a	RASPBERRY_PI	Microphone, Camera...		Today, 18:04:50

```
Edge Impulse Linux client v1.2.5
? What is your user name or e-mail address (edgeimpulse.com)? zin@edgeimp
ulse.com
? What is your password? [hidden]
? To which project do you want to connect this device? Zin / lug-nut-coun
ter
? Select a microphone (or run this command with --disable-microphone to s
kip selection) USB-Audio - HD Webcam C615
[SER] Using microphone hw:1,0
[SER] Using camera HD Webcam C615 starting...
[SER] Connected to camera
[WS ] Connecting to wss://remote-mgmt.edgeimpulse.com
[WS ] Connected to wss://remote-mgmt.edgeimpulse.com
? What name do you want to give this device? rpi4
[WS ] Device "rpi4" is now connected to project "lug-nut-counter"
[WS ] Go to https://studio.edgeimpulse.com/studio/40118/acquisition/train
ing to build your machine learning model!
```

DATA ACQUISITION

Training accurate production ready machine learning models requires feeding plenty of varied data, which means a lot of images are typically required. For this proof-of-concept, we captured around 145 images of a wheel that had lug nuts on it. The Edge Impulse Linux daemon allows you to directly connect the Raspberry Pi 4 to Edge Impulse and take snapshots using the USB webcam. Using the Labeling queue in the Data Acquisition page we then easily drew bounding boxes around each lug nut within every image, along with every wheel. To add some test data we went back to the main Dashboard page and clicked the 'Rebalance dataset' button that moves 20% of the training data to the test data bin.

EDGE IMPULSE DATA ACQUISITION (LUG-NUT-COUNTER) Zin

Training data | Test data | Labeling queue (3)

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

DATA COLLECTED: 145 items | LABELS: 2

SAMPLE NAME	LABELS	ADDED	LENGTH
tire.29joa5fa	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5g2	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5fi	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5d2	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5cq	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5ci	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5bp	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa5c2	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa53a	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa4ug	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa4v1	lug, lug, lug, ...	Jul 03 2021, ...	-
tire.29joa4ni	lug, lug, lug, ...	Jul 03 2021, ...	-

Record new data Connect using WebUSB

Device: rpi4

Label:

Sensor: Camera (640x480)

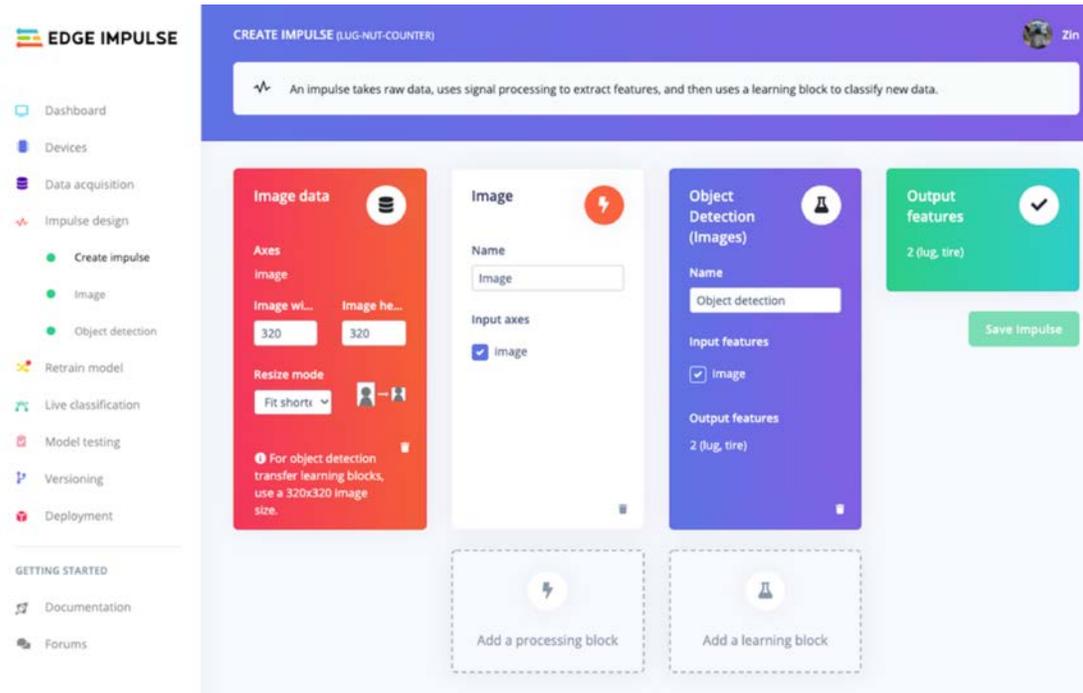
Start sampling

RAW DATA: tire.29joa4v1

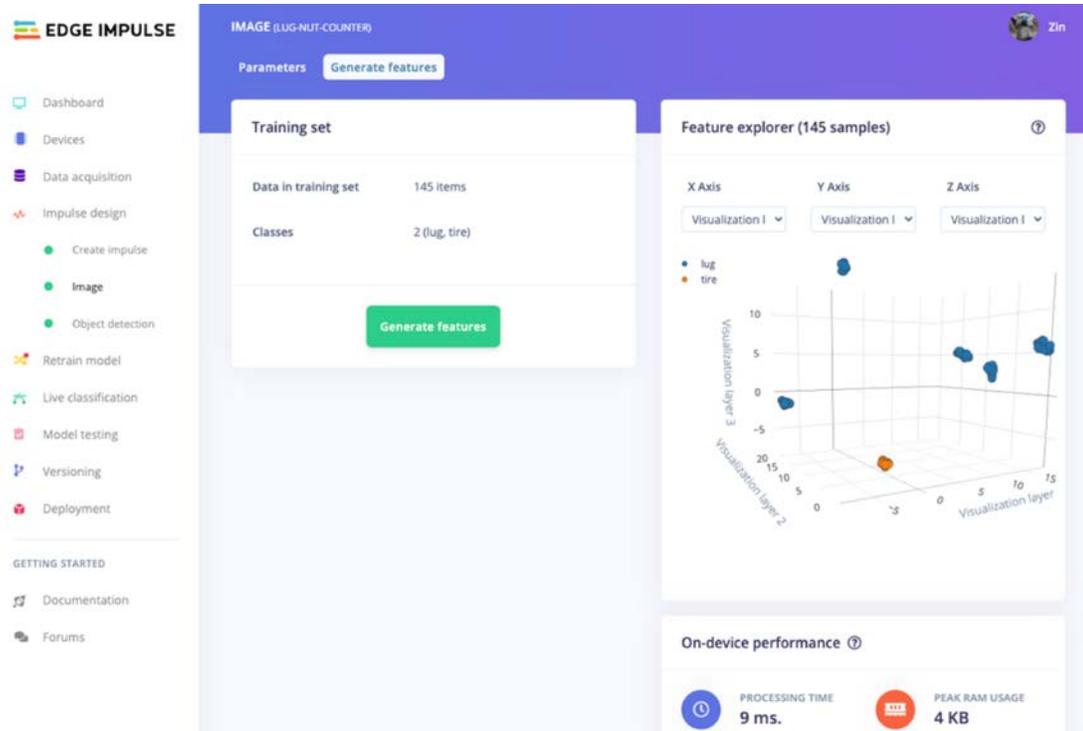
Image showing a wheel with bounding boxes for lug nuts.

IMPULSE DESIGN AND MODEL TRAINING

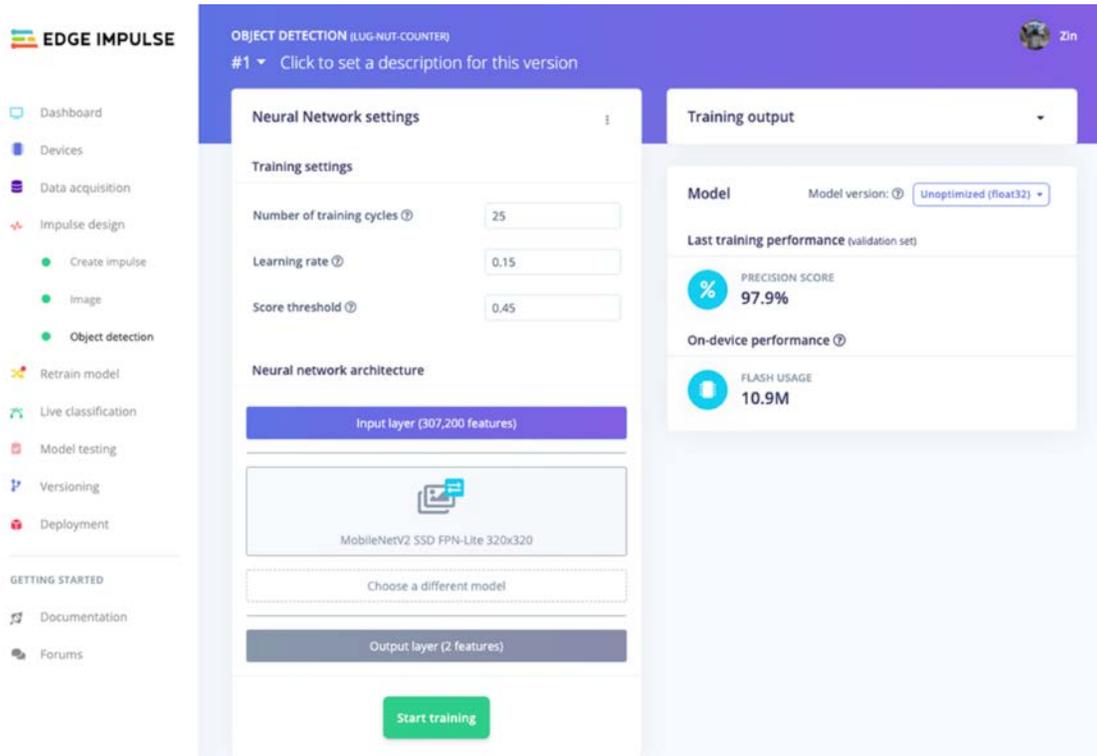
Now that we have plenty of training data, it's time to design and build our model. The first block in the Impulse Design is an Image Data block, and it scales each image to a size of '320' by '320' pixels.



Next, image data is fed to the Image processing block that takes the raw RGB data and derives features from it.



Finally, these features are used as inputs to the MobileNetV2 SSD FPN-Lite Transfer Learning Object Detection model that learns to recognize the lug nuts. The model is set to train for '25' cycles at a learning rate of '.15', but this can be adjusted to fine-tune for accuracy. As you can see from the screenshot below, the trained model indicates a precision score of 97.9%.



MODEL TESTING

If you'll recall from an earlier step we rebalanced the dataset to put 20% of the images we collected to be used for gauging how our trained model could perform in the real world. We use the model testing page to run a batch classification and see how we expect our model to perform. The 'Live Classification' tab will also allow you to acquire new data direct from the Raspberry Pi 4 and see how the model measures up against the immediate image sample.

The screenshot displays the Edge Impulse Model Testing interface. On the left is a navigation sidebar with options like Dashboard, Devices, Data acquisition, Impulse design, Retrain model, Live classification, Model testing (selected), Versioning, and Deployment. The main area is titled 'MODEL TESTING (LUG-NUT-COUNTER)' and contains a 'Test data' table and a 'Model testing output' log.

Test data table:

SAMPLE ...	EXPECTED O...	LE...	PRECISI...	RESULT
testing,...	lug, lug, lug...	-	94%	
testing,...	lug, lug, lug...	-	84%	
testing,...	lug, lug, lug...	-	87%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	

Model testing output log:

```

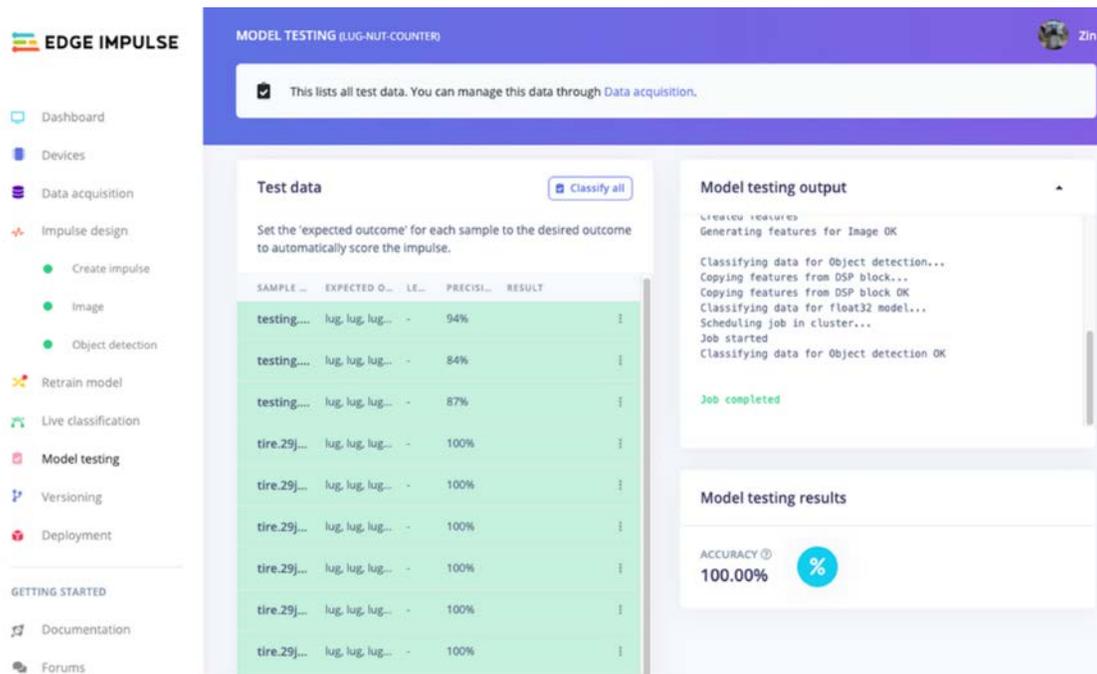
Generating features for Image OK
Classifying data for Object detection...
Copying features from DSP block...
Copying features from DSP block OK
Classifying data for float32 model...
Scheduling job in cluster...
Job started
Classifying data for Object detection OK
Job completed
    
```

Model testing results:

ACCURACY ① **100.00%** %

VERSIONING

An MLOps platform would not be complete without a way to archive your work as you iterate on your project. The 'Versioning' tab allows you to save your entire project including the entire dataset so you can always go back to a "known good version" as you experiment with different neural network parameters and project configurations. It's also a great way to share your efforts as you can designate any version as 'public' and other Edge Impulse users can clone your entire project and use it as a springboard to add their own enhancements.



The screenshot displays the Edge Impulse web interface for a project named 'MODEL TESTING (LUG-NUT-COUNTER)'. The left sidebar contains navigation options: Dashboard, Devices, Data acquisition, Impulse design, Create impulse, Image, Object detection, Retrain model, Live classification, Model testing, Versioning, Deployment, GETTING STARTED, Documentation, and Forums. The main content area is divided into three sections:

- Test data:** A table with columns 'SAMPLE', 'EXPECTED O...', 'LE...', 'PRECISI...', and 'RESULT'. It lists several test samples with their corresponding accuracy percentages.
- Model testing output:** A log showing the progress of the model testing process, including steps like 'Generating features for Image OK', 'Classifying data for Object detection...', and 'Job completed'.
- Model testing results:** A summary card showing 'ACCURACY' as '100.00%' with a blue percentage icon.

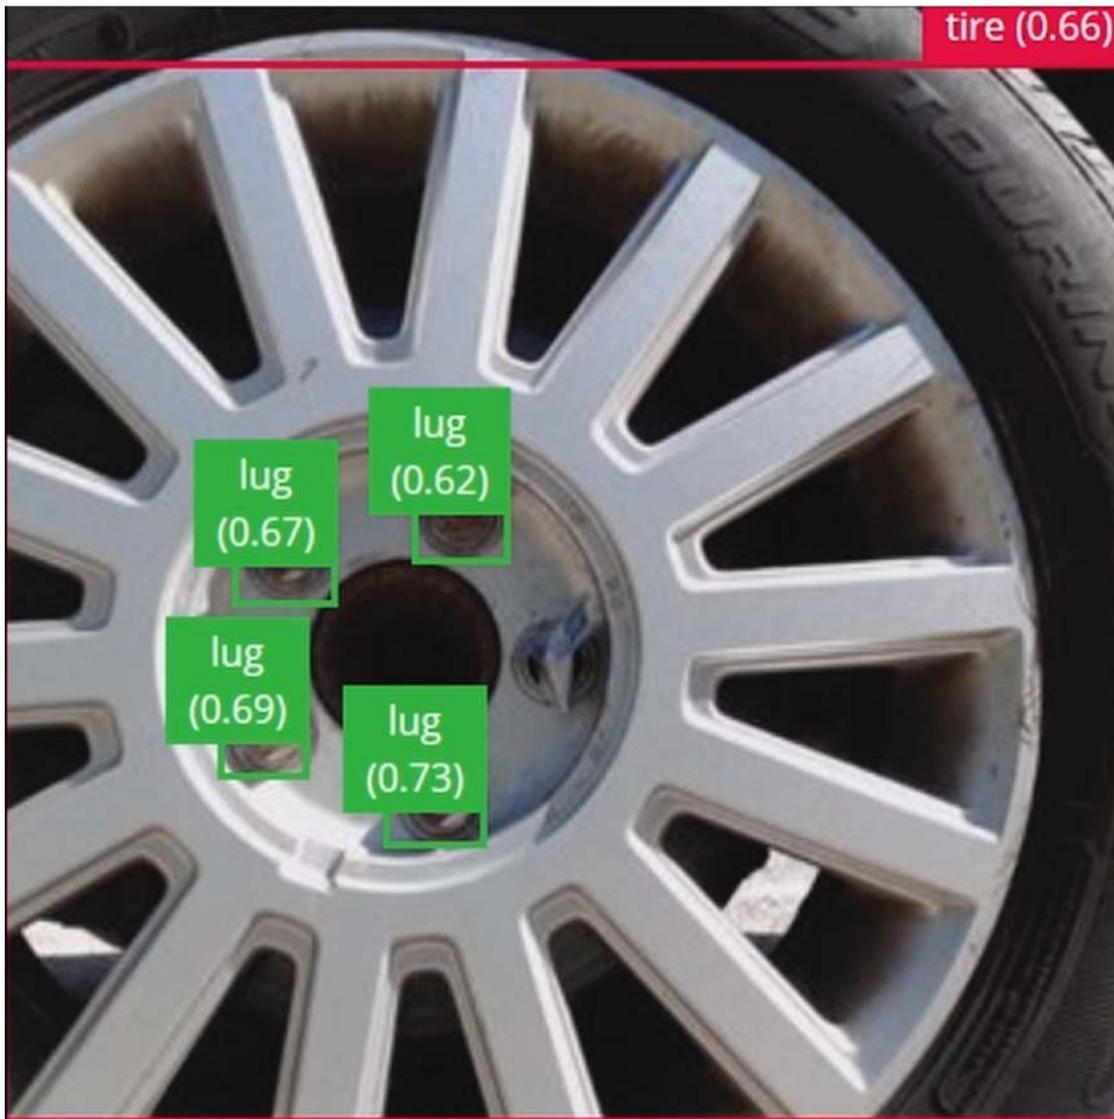
SAMPLE	EXPECTED O...	LE...	PRECISI...	RESULT
testing....	lug, lug, lug...	-	94%	
testing....	lug, lug, lug...	-	84%	
testing....	lug, lug, lug...	-	87%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	
tire.29j...	lug, lug, lug...	-	100%	

DEPLOYING MODELS

In order to verify that the model works correctly in the real world, we'll need to deploy it to the Raspberry Pi 4. This is a simple task thanks to the Edge Impulse CLI, as all we have to do is run:

```
edge-impulse-linux-runner
```

which downloads the model and creates a local webserver. From here, we can open a browser tab and visit the address listed after we run the command to see a live camera feed and any objects that are currently detected. Here's a sample of what the user will see in their browser tab:



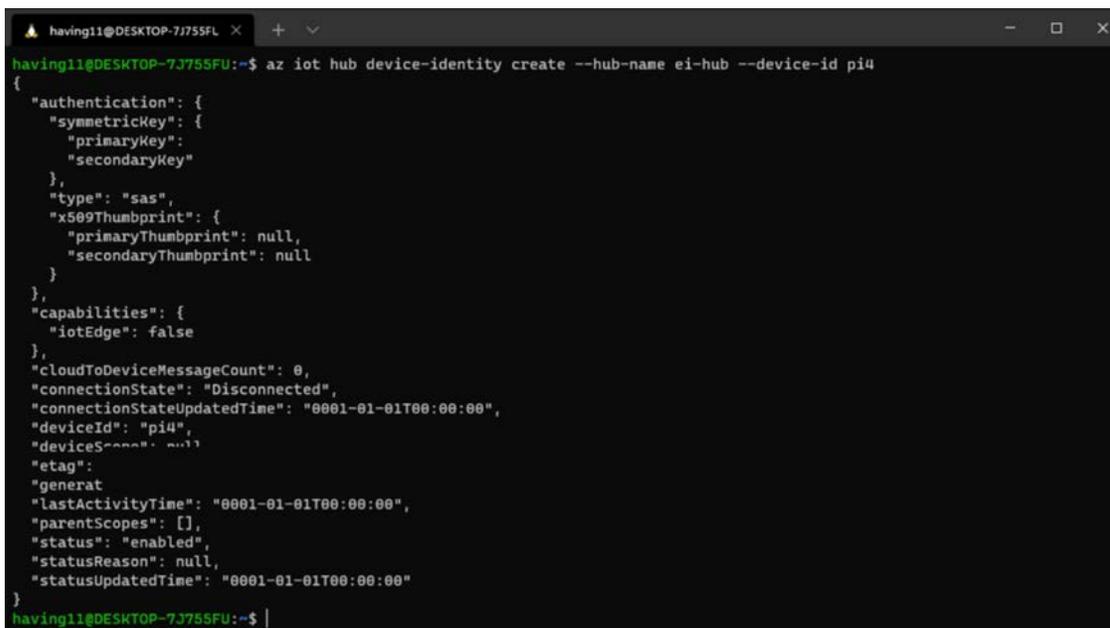
SENDING INFERENCE RESULTS TO AZURE IOT HUB

With the model working locally on the Raspberry Pi 4, let's see how we can send the inference results from the Raspberry Pi 4 to an Azure IoT Hub instance. As previously mentioned, these results will enable business applications to leverage other Azure services such as Azure Stream Analytics and Power BI. On your development machine, make sure you've installed the [Azure CLI](#) and have signed in using 'az login'. Then get the name of the resource group you'll be using for the project. If you don't have one, you can [follow this guide](#) on how to create a new resource group.

Name ↑	Type ↑	Resource group ↑	Location ↑
ei-hub	IoT Hub	ei-testing	Central US

After that, return to the terminal and run the following commands to create a new IoT Hub and register a new device ID:

```
az iot hub create --resource-group <your resource group> --name <your IoT
Hub name>
az extension add --name azure-iot
az iot hub device-identity create --hub-name <your IoT Hub name> --device-id
<your device id>
```



```
having11@DESKTOP-7J755FU:~$ az iot hub device-identity create --hub-name ei-hub --device-id pi4
{
  "authentication": {
    "symmetricKey": {
      "primaryKey":
      "secondaryKey"
    },
    "type": "sas",
    "x509Thumbprint": {
      "primaryThumbprint": null,
      "secondaryThumbprint": null
    }
  },
  "capabilities": {
    "iotEdge": false
  },
  "cloudToDeviceMessageCount": 0,
  "connectionState": "Disconnected",
  "connectionStateUpdatedTime": "0001-01-01T00:00:00",
  "deviceId": "pi4",
  "deviceStatus": null,
  "etag":
  "generation":
  "lastActivityTime": "0001-01-01T00:00:00",
  "parentScopes": [],
  "status": "enabled",
  "statusReason": null,
  "statusUpdatedTime": "0001-01-01T00:00:00"
}
having11@DESKTOP-7J755FU:~$ |
```

Retrieve the connection string the Raspberry Pi 4 will use to connect to Azure IoT with:

```
az iot hub device-identity connection-string show --device-id <your device id>
--hub-name <your IoT Hub name>
```

```
having11@DESKTOP-7J755FL x + v - □ x
},
"capabilities": {
  "iotEdge": false
},
},
"cloudToDeviceMessageCount": 0,
"connectionState": "Disconnected",
"connectionStateUpdateTime": "0001-01-01T00:00:00",
"deviceId": "pi4",
"deviceScope": null,
"etag": null,
"generated": null,
"lastActivityTime": "0001-01-01T00:00:00",
"parentScopes": [],
"status": "enabled",
"statusReason": null,
"statusUpdateTime": "0001-01-01T00:00:00"
}
having11@DESKTOP-7J755FU:~$ az iot hub device-identity show-connection-string --device-id pi4
This command has been deprecated and will be removed in a future release. Use 'az iot hub device-identity connection-string show' instead.
Please provide an IoT Hub entity name (via the '--hub-name' or '-n' parameter) or IoT Hub connection string via --login.
..
having11@DESKTOP-7J755FU:~$ az iot hub device-identity show-connection-string --device-id pi4 --hub-name ei-hub
This command has been deprecated and will be removed in a future release. Use 'az iot hub device-identity connection-string show' instead.
{
  "connectionString":
  ..
}
having11@DESKTOP-7J755FU:~$ |
```

Now it's time to SSH into the Raspberry Pi 4 and set the connection string as an environment variable with:

```
export IOTHUB_DEVICE_CONNECTION_STRING="<your connection string here>"
```

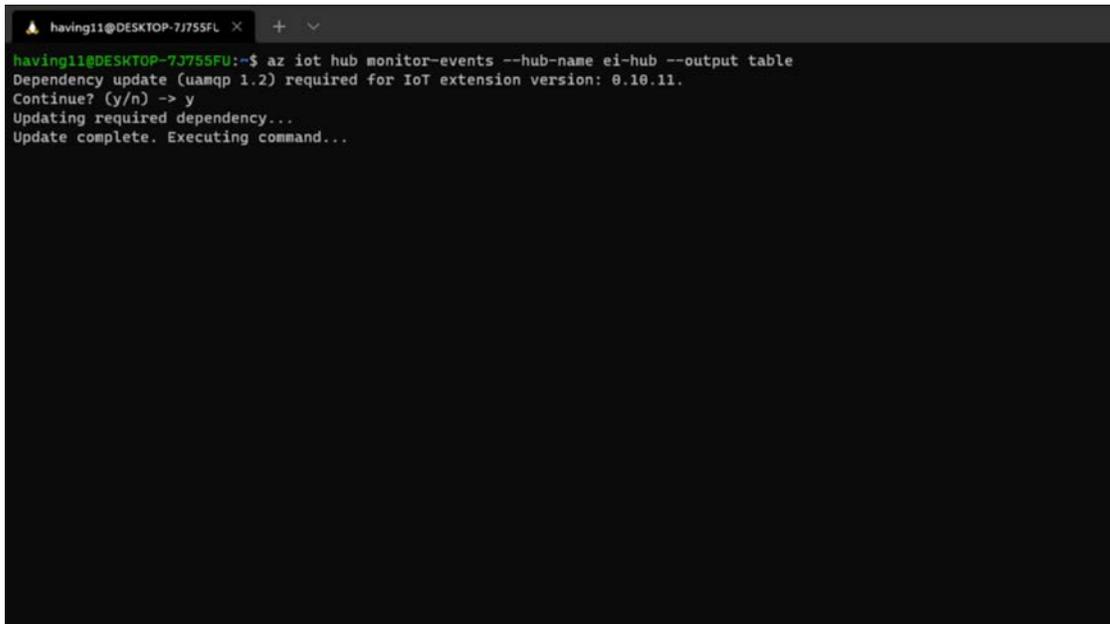
Then, add the necessary Azure IoT device libraries with:

```
pip install azure-iot-device
```

(Note: if you do not set the environment variable or pass it in as an argument the program will not work!) The connection string contains the information required for the Raspberry Pi 4 to establish a connection with the Azure IoT Hub service and communicate with it. You can then monitor output in the Azure IoT Hub with:

```
az iot hub monitor-events --hub-name <your IoT Hub name> --output table
```

or in the Azure Portal.



```
having11@DESKTOP-7J755FL x + v
having11@DESKTOP-7J755FU:~$ az iot hub monitor-events --hub-name ei-hub --output table
Dependency update (uamqp 1.2) required for IoT extension version: 0.10.11.
Continue? (y/n) -> y
Updating required dependency...
Update complete. Executing command...
```

To make sure it works, download and run this example on the Raspberry Pi 4 to make sure you can see the test message.

For the second half of deployment, we'll need a way to customize how our model is used within the code. Edge Impulse provides a Python SDK for this purpose. On the Raspberry Pi 4 install it with:

```
sudo apt-get install libatlas-base-dev libportaudio0 libportaudio2 libportaudiocpp0 portaudio19-dev
pip3 install edge_impulse_linux -i https://pypi.python.org/simple
```

We've made available a simple example on the Raspberry Pi 4 that sets up a connection to the Azure IoT Hub, runs the model, and sends the inference results to Azure IoT.

```

151 lines (124 sloc)  5.45 KB
7  import os
8  import sys
9  import cv2
10 import sys
11 import signal
12 import time
13 import uuid
14 from edge_impulse_linux_image import ImageImpulseRunner
15 from azure.iot.device.aio import IoTHubDeviceClient
16 from azure.iot.device import Message
17 import cli_parser
18
19
20 cli_args = cli_parser.parser.parse_args()
21
22 runner = None
23 WHEEL_COUNT = (200, 10, 0)
24 LUG_NUT_COUNT = (0, 250, 15)
25
26 def now():
27     return round(time.time() * 1000)
28
29 def get_camera() -> int:
30     if cli_args.port != None:
31         camera = cv2.VideoCapture(cli_args.port)
32         ret = camera.read()[0]
33         if ret:
34             print("Camera found on port (%). Resolution = (%) x (%).",
35                   name = "{}".format(cli_args.port, camera.get(3), camera.get(4),
36                   camera.get(cv2.CAP_PROP_FRAME_WIDTH)))
37             camera.release()
38             return cli_args.port
39         else:
40             raise Exception("Couldn't initialize camera")
41
42     ports = []
43     for port in range(1):
44         camera = cv2.VideoCapture(port)
45         if camera.isOpened():
46             ret = camera.read()[0]
47             if ret:
48                 print("Camera found on port (%). Resolution = (%) x (%).",
49                       name = "{}".format(port, camera.get(3), camera.get(4),
50                       camera.get(cv2.CAP_PROP_FRAME_WIDTH)))
51                 ports.append(port)
52                 camera.release()
53             if len(ports) > 1:
54                 raise Exception("More than one camera found! Use the -p option.")
55             if len(ports) == 0:
56                 raise Exception("No cameras found!")
57     return ports[0]
58
59 def signal_handler(sig, frame):
60     print("Interrupted")
61     if runner:
62         runner.stop()
63     sys.exit(0)
64
65 signal.signal(signal.SIGINT, signal_handler)
66
67 async def main():
68     # Start the connection stream from an independent coroutine on the line

```

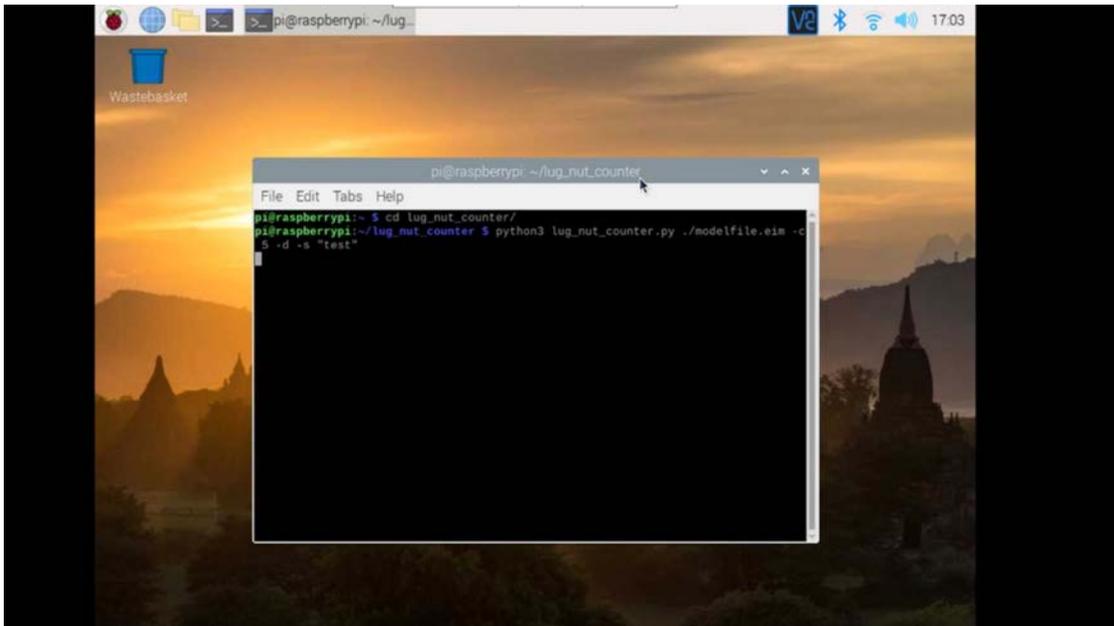
Once you've either downloaded the zip file or cloned the repo into a folder, get the model file by running:

```
edge-impulse-linux-runner --download modelfile.eim
```

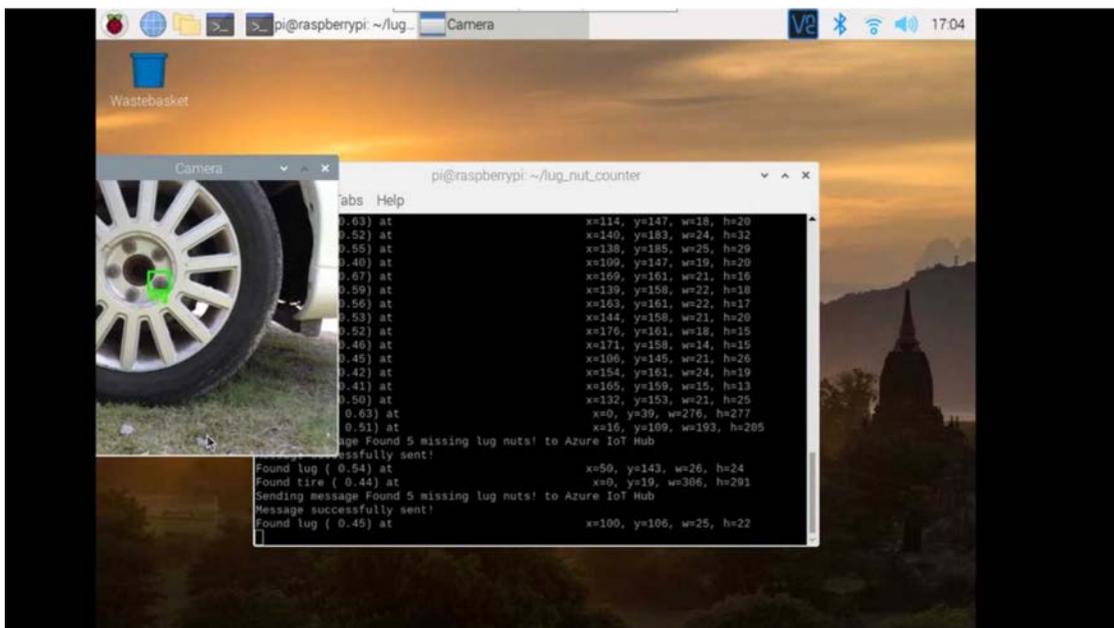
inside of the folder you just created from the cloning process. This will download a file called 'modelfile.eim'. Now, run the Python program with:

```
python lug_nut_counter.py ./modelfile.eim -c <LUG_NUT_COUNT>
```

where <LUG_NUT_COUNT> is the correct number of lug nuts that should be attached to the wheel (you might have to use 'python3' if both Python 2 and 3 are installed).



Now whenever a wheel is detected the number of lug nuts is calculated. If this number falls short of the target, a message is sent to the Azure IoT Hub.



And by only sending messages when there's something wrong, we can prevent an excess amount of bandwidth from being taken due to empty payloads.

```
having11@DESKTOP-7J755FL x + v
component: ''
interface: ''
module: ''
origin: pi4
payload: Found 5 missing lug nuts!

event:
component: ''
interface: ''
module: ''
origin: pi4
payload: Found 5 missing lug nuts!

event:
component: ''
interface: ''
module: ''
origin: pi4
payload: Found 5 missing lug nuts!

event:
component: ''
interface: ''
module: ''
origin: pi4
payload: Found 5 missing lug nuts!

event:
component: ''
interface: ''
module: ''
origin: pi4
payload: Found 4 missing lug nuts!

event:
component: ''
interface: ''
module: ''
origin: pi4
payload: Found 5 missing lug nuts!
```

CONCLUSION

We've just scratched the surface with wheel lug nut detection. Imagine utilizing object detection for other industrial applications in quality control, detecting ripe fruit amongst rows of crops, or identifying when machinery has malfunctioned with devices powered by machine learning.

With any hardware, Edge Impulse, and Microsoft Azure IoT, you can design comprehensive embedded

machine learning models, deploy them on any device, while authenticating each and every device with built-in security. You can set up individual identities and credentials for each of your connected devices to help retain the confidentiality of both cloud-to-device and device-to-cloud messages, revoke access rights for specific devices, upgrade device firmware remotely, and benefit from advanced analytics on devices running offline or with

intermittent connectivity.

The complete Edge Impulse project is available [here](#) for you to see how easy it is to start building your own embedded machine learning projects today using object detection. We look forward to your feedback at hello@edgeimpulse.com or on our [forum](#).



EDGE IMPULSE

hello@edgeimpulse.com

Edge Impulse
3031 Tisch Way
110 Plaza West
San Jose, CA 95128
USA