arm AI

AI Virtual Tech Talks Series

EDGE IMPULSE

Efficient ML across Arm from Cortex-M to Web Assembly

Jan Jongboom

August 25, 2020

Jan Jongboom

jan@edgeimpulse.com

Some things are easy...

https://www.flickr.com/photos/free-stock/8425195959

-10°C

Some things were hard, but are now easy

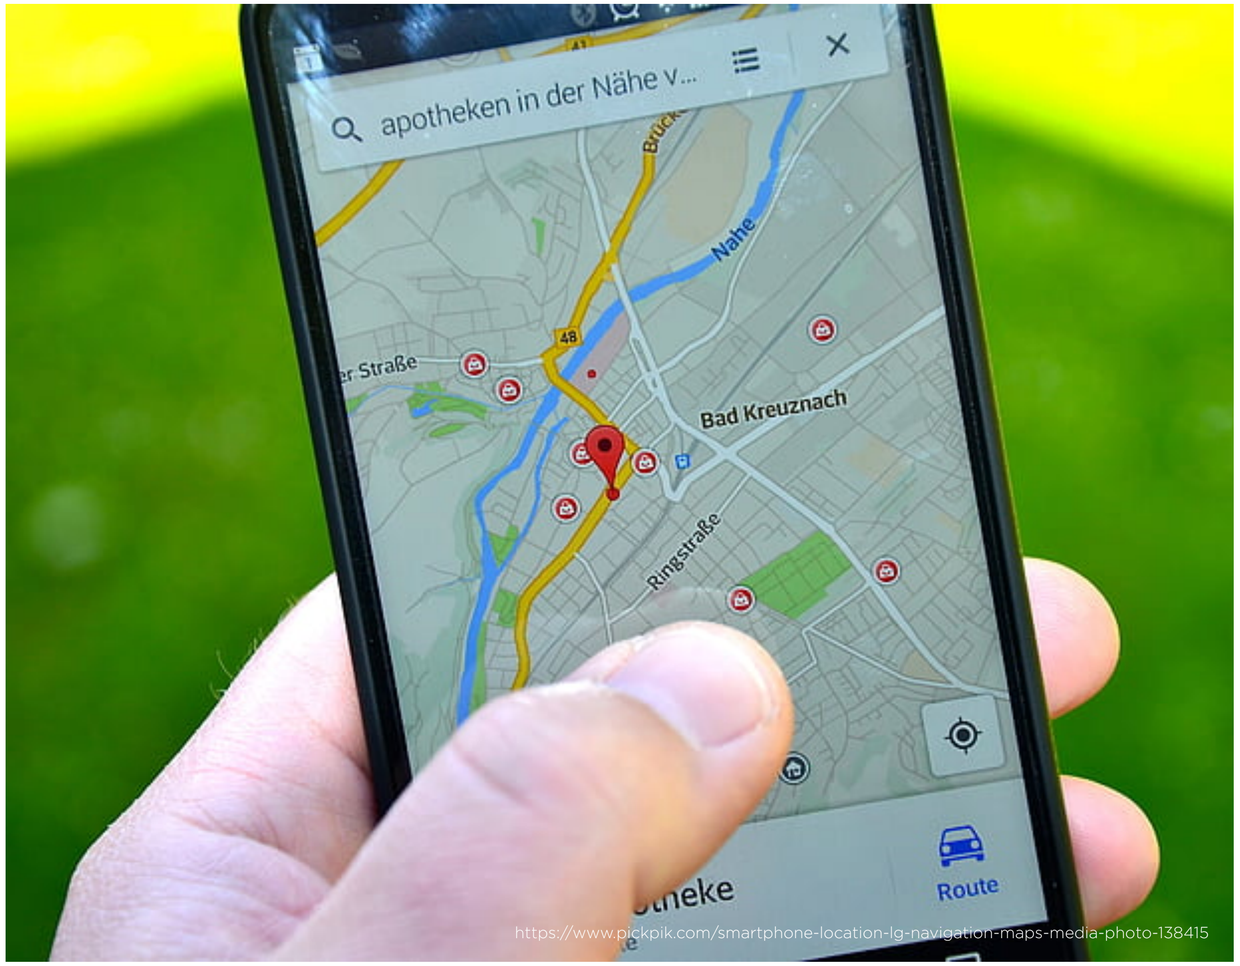

https://www.pickpik.com/smartphone-location-lg-navigation-maps-media-photo-138415

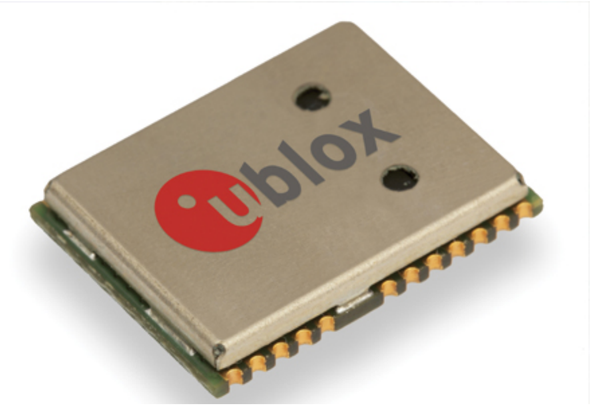

Some things are still hard...

Do I hear glass breaking?

Some things are still hard…

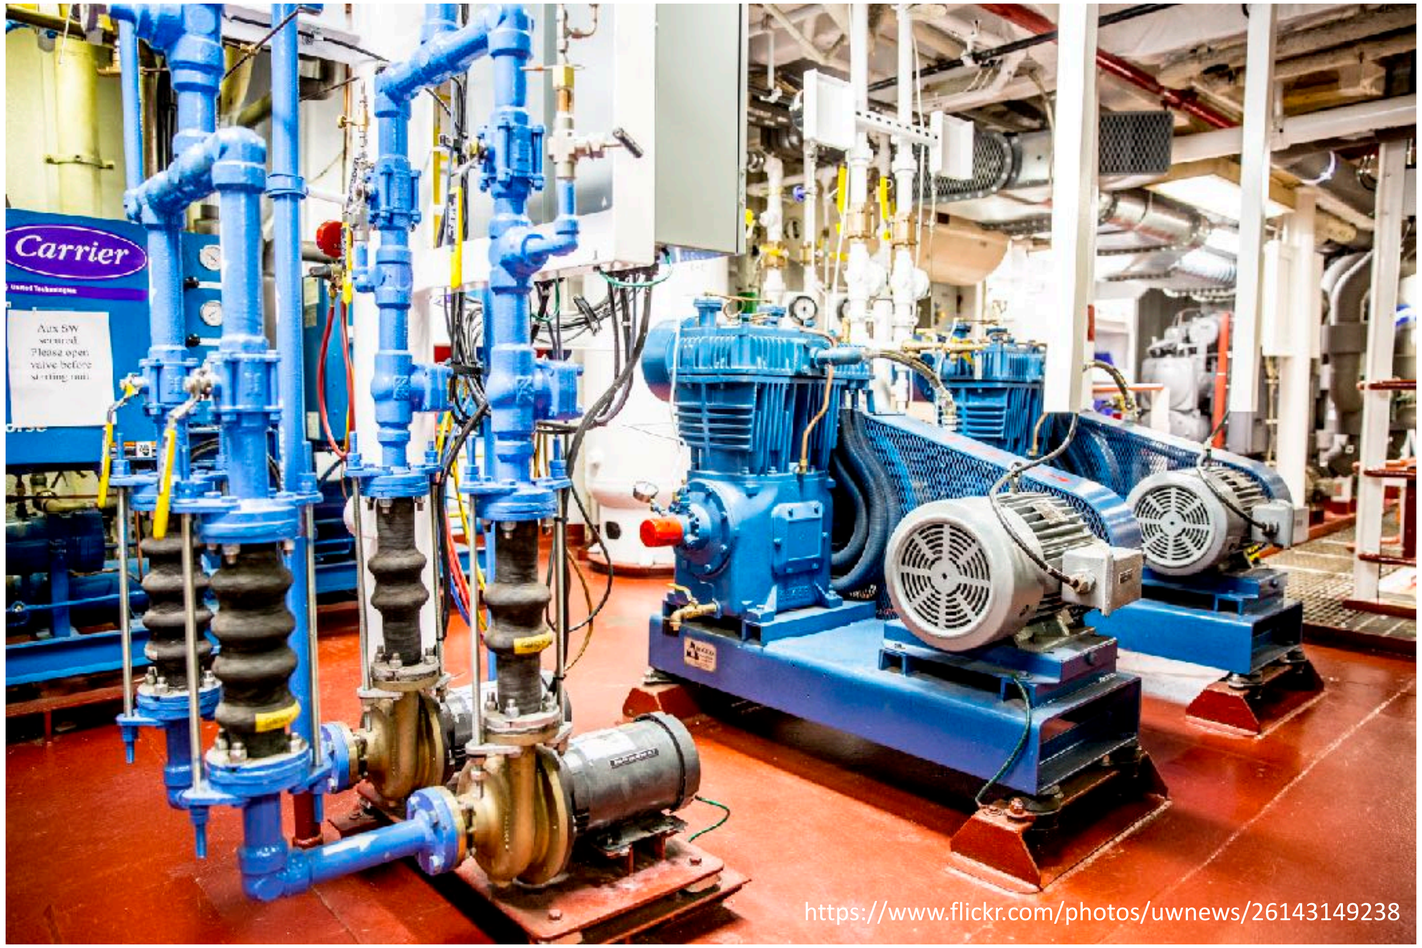

Is this machine vibrating differently?

But, no longer impossible...

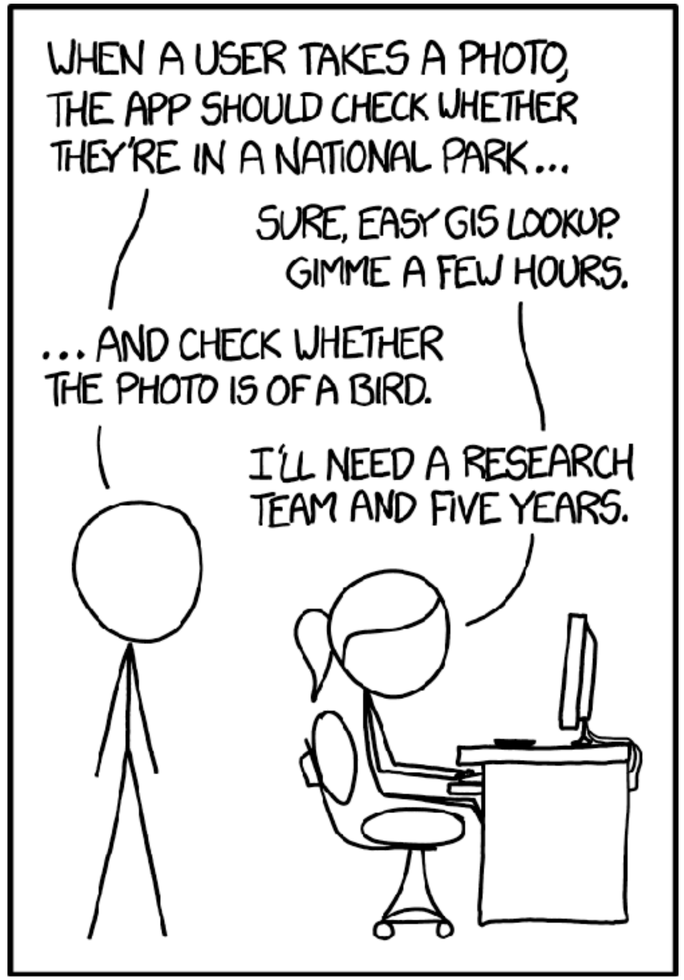

https://xkcd.com/1425/

Machine learning is great at finding patterns in messy data

(anything you can't reason about in Excel)

Machine learning?

ML models move to the edge

Focus on inferencing, not training

Machine learning model is just a mathematical function with lots of parameters

Accuracy vs. speed, reducing parameters, hardware-optimized paths

Leveraging silicon (either specialized, or vector extensions)

Pete Warden

Neil Tan

Much wider deployment options

Cloud

Gateway

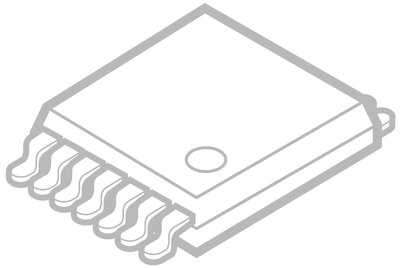

Sensor

Where to deploy?

Cloud - expensive to get data into place, compute is cheap.

Gateway - easy to retrofit, can be expensive to get raw data stream.

Edge device - requires new firmware and planning, but no radio required.

Or even hybrid... There's no right answer.

From 0 to model

EDGE IMPULSE

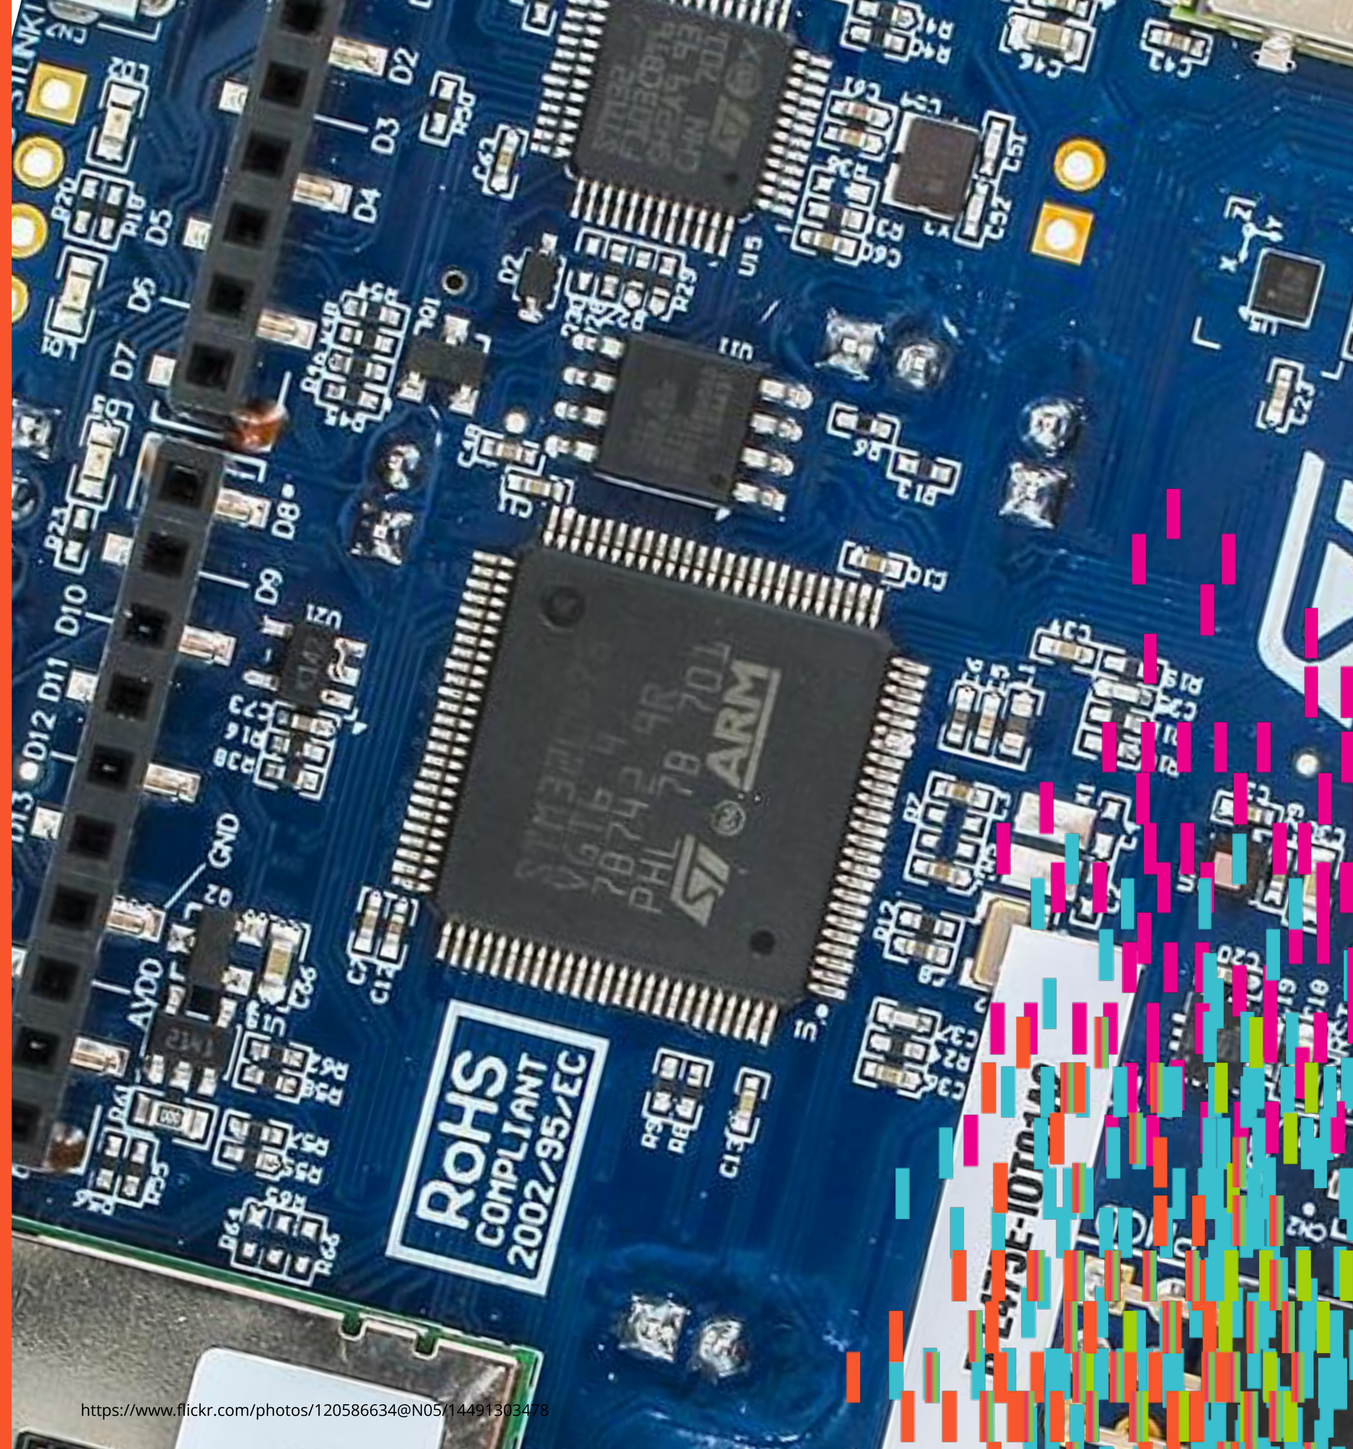

https://www.flickr.com/photos/120586634@N05/14491303478

1. Everything starts with raw data

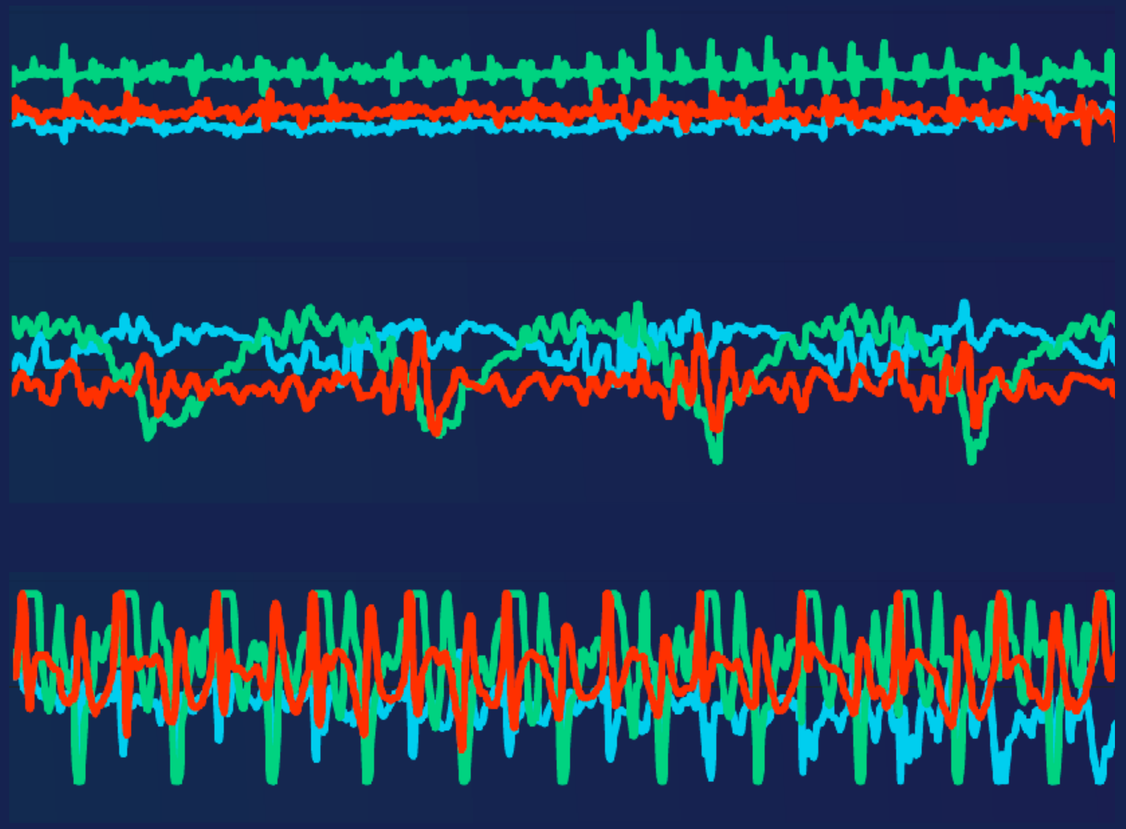

Get data at the highest resolution possible - e.g. using serial or directly over WiFi

2. Extracting meaningful features

Very dependent on your use case

Raw data can be notoriously hard to deal with

(3s. accelerometer data = 900 data points, 1s. audio data = 16,000 data points)

Use proven DSP algorithms

-1.1200, 0.5300, 10.6300, -0.5200, 1.1600, 13.1600, -0.1600, 1.4200, 13.5900, -0.2400, 1.3200, 10.8500, -0.3700, 0.5100, 7.7800, -0.1900, 0.0500, 8.8400, -0.1900, 0.0500, 8.8400, 0.2400, 0.7300, 11.1900, 0.5200, 1.6500, 12.5200, 0.6400, 1.5000, 10.8200, 0.2900, 0.3300, 7.4000, -0.3100, -0.5600, 8.5900, -0.8800, 0.6600, 11.9200, -0.8800, 0.6600, 11.9200, -0.4500, 1.3900, 11.9100,
-0.0800, 1.6000, 12.7800, -0.1100, 0.3600, 9.1700, -0.0200, 0.0700, 9.8200, 0.0600, 0.9600, 12.0000, 0.2800, 1.7700, 13.2000, 0.2800, 1.7700, 13.2000, 0.2300, 1.5100, 12.6000, 0.2700, 1.0800, 11.4300, 0.0900, 0.9300, 10.9400, 0.1700, 1.2100, 11.0400, 0.4100, 1.8500, 11.5000, 0.3900, 1.9600, 11.4300, 0.3900, 1.9600, 11.4300, 0.2400, 1.4400, 10.7200, 0.0300, 1.1900, 10.3600,
-0.0300, 1.3000, 10.6100, 0.4800, 1.7900, 11.7200, 1.0400, 2.6300, 13.3300, 1.0400, 2.6300, 13.3300, 1.0600, 2.3700, 13.5800, 0.3600, 1.9600, 13.3800, -0.1000, 2.1200, 13.9600, -0.3600, 2.0200, 14.5300, 0.0000, 2.1500, 14.6900, 0.0600, 2.1400, 14.3700, 0.0600, 2.1400, 14.3700, -0.3000, 1.8800, 13.7900, 0.0500, 1.7000, 13.5900, 0.1300, 1.6700, 13.1500, -0.0100, 1.7700, 12.9000,
0.4000, 1.8900, 12.2300, 0.5300, 2.3300, 12.2600, 0.5300, 2.3300, 12.2600, 0.0400, 1.9500, 11.8100, -0.2300, 1.9600, 11.2400, -0.0600, 2.1100, 10.2200, -0.1100, 2.4100, 9.7800, -0.3500, 2.7100, 9.7500, -0.7800, 3.1000, 10.1100, -0.7800, 3.1000, 10.1100, -1.0700, 3.1100, 9.8000, -1.2100, 2.9400, 9.1900, -1.1500, 3.2100, 8.6400, -0.7300, 3.6500, 8.4600, -0.5000, 3.9500, 8.6300,
-0.4300, 3.9100, 8.7400, -0.4300, 3.9100, 8.7400, -0.6400, 3.7800, 8.8700, -1.2000, 3.9200, 8.9300, -1.0800, 4.4400, 8.8100, -0.7800, 4.1900, 8.1200, -0.4400, 4.1000, 7.6400, -0.5400, 4.2000, 7.5600, -0.5400, 4.2000, 7.5600, -1.0700, 4.2600, 7.2700, -1.3000, 4.5100, 7.2300, -1.2600, 4.4600, 6.6900, -1.2800, 4.4100, 6.6000, -1.7000, 4.6800, 7.0800, -2.3400, 5.1100, 7.5900,
-2.3400, 5.1100, 7.5900, -2.8300, 4.8700, 6.8700, -2.9700, 4.7600, 6.2700, -3.2500, 4.6000, 6.1500, -3.4900, 4.5900, 6.2600, -3.3000, 4.9200, 6.3400, -2.7000, 4.9300, 5.8600, -2.7000, 4.9300, 5.8600, -2.9000, 4.5100, 4.9900, -3.5200, 4.3200, 4.9900, -4.1400, 4.2100, 5.7800, -3.7600, 4.1600, 5.7700, -3.0200, 4.2200, 5.4900, -3.0000, 3.8900, 3.9100, -3.0000, 3.8900, 3.9100,
-3.3500, 3.5800, 3.4400, -3.1100, 3.2000, 3.6000, -3.0900, 3.6000, 5.9400, -3.0800, 3.0900, 5.1800, -2.8000, 2.9400, 4.5600, -2.4000, 2.4900, 3.9200, -2.4000, 2.4900, 3.9200, -1.8500, 2.6300, 4.4900, -1.4300, 3.9900, 7.3400, -1.4900, 3.4900, 6.1600, -1.5300, 3.2100, 5.4500, -0.9900, 2.8600, 6.2400, -0.7900, 3.2200, 8.4700, -0.7900, 3.2200, 8.4700, -0.9300, 3.5700, 9.0300,
-1.6600, 2.9600, 7.0700, -1.7600, 2.1000, 6.9000, -1.4600, 2.1100, 9.0100, -1.3000, 2.5400, 10.3000, -1.3000, 2.5400, 10.3000, -1.4500, 2.6500, 9.7800, -1.5300, 1.8900, 8.4100, -1.1400, 1.3000, 9.4400, -0.7500, 1.6100, 10.3600, -0.9400, 1.5300, 9.7800, -1.0700, 1.0100, 9.2700, -1.0700, 1.0100, 9.2700, -0.8600, 1.0200, 10.1100, 0.3900, 1.6800, 11.3200, 1.2900, 1.8500, 11.4700,
1.0700, 1.3200, 11.2500, -0.2100, 1.5800, 12.4300, -1.8100, 1.3500, 12.3300, -1.8100, 1.3500, 12.3300, -2.1800, 1.0400, 11.1600, -1.5400, 0.3000, 10.3100, -0.4700, 0.2700, 11.0900, 0.7900, 1.4100, 12.9800, 1.1600, 1.7100, 12.4700, 0.5800, 0.8700, 9.6300, 0.5800, 0.8700, 9.6300, 0.1300, 0.3100, 9.5600, 0.2800, 0.3600, 10.5600, 0.7400, 0.8500, 11.4200, 0.9300, 1.0200, 11.4300,
0.6700, 0.5600, 10.5300, 0.8500, 0.3500, 9.4100, 0.8500, 0.3500, 9.4100, 1.6600, 1.2800, 10.9200, 2.0500, 0.9900, 9.7000, 2.1300, 0.8800, 10.1900, 2.0500, 0.9100, 11.3300, 1.7700, 1.4100, 12.2700, 1.4800, 1.7600, 12.1000, 1.4800, 1.7600, 12.1000, 0.9400, 1.1300, 10.8500, 0.2000, 0.8000, 10.1200, 0.2600, 1.1600, 10.5800, 0.5100, 1.6500, 10.7800, 0.4600, 1.2000, 9.9900, 0.9100,
0.8400, 9.6400, 0.9100, 0.8400, 9.6400, 1.4500, 0.7400, 10.2500, 2.0200, 1.3000, 11.4500, 1.8100, 1.8700, 12.1300, 1.0500, 1.5300, 12.0200, 0.6200, 0.6700, 11.3100, 0.7100, 0.8500, 12.0000, 0.7100, 0.8500, 12.0000, 0.6400, 1.2200, 13.1400, 1.1300, 2.0400, 14.6200, 0.8300, 2.0200, 15.5100, -0.1400, 1.4800, 15.6500, -0.6300, 1.5900, 16.0500, -1.3100, 1.7100, 16.3900, -1.3100,
1.7100, 16.3900, -1.7300, 1.5800, 16.6300, -1.1500, 1.4400, 16.0300, -0.5300, 1.1700, 15.1000, -0.1800, 0.9900, 14.4600, -0.3300, 1.0100, 13.5100, -0.3300, 1.0100, 13.5100, -0.4400, 0.9100, 12.6700, 0.0400, 1.2300, 12.5400, 0.6900, 2.0500, 13.1600, 0.3100, 1.7700, 12.8600, 0.0300, 1.3800, 11.1100, -0.4400, 1.2200, 9.4900, -0.4400, 1.2200, 9.4900, 0.1100, 1.1400, 7.3100, 0.8500,
2.2500, 8.4600, 0.8600, 3.3700, 11.2200, -0.1100, 2.2800, 8.4400, -1.3800, 1.5300, 7.1700, -1.0600, 1.5400, 6.9500, -1.0600, 1.5400, 6.9500, -0.5200, 2.8300, 8.7100, -0.2100, 2.3500, 8.1800, -0.3400, 2.7000, 8.9200, -0.3000, 2.3100, 8.7500, -0.4800, 1.4700, 7.8700, -0.3600, 0.9400, 6.9700, -0.3600, 0.9400, 6.9700, -0.2300, 1.4700, 7.6100, -0.3300, 2.2300, 8.5000, 0.3000, 1.9200,
7.8600, -0.2300, 1.5700, 6.8700, -1.4900, 1.5600, 6.3700, -2.8200, 1.6200, 7.2000, -2.8200, 1.6200, 7.2000, -3.1600, 1.8800, 7.1500, -2.7600, 2.2900, 6.8500, -2.6000, 2.2200, 6.2600, -2.9000, 1.9900, 5.8900, -3.3800, 2.2200, 6.2600, -3.9000, 2.1700, 6.0300, -3.9000, 2.1700, 6.0300, -3.8600, 2.3800, 5.6600, -3.5300, 2.5200, 5.6700, -3.2400, 2.3700, 5.8200, -3.2800, 2.1800,
5.5200, -3.1500, 2.1800, 5.6500, -3.0900, 2.0700, 5.1600, -3.0900, 2.0700, 5.1600, -2.4300, 2.1000, 5.3800, -2.0200, 2.3600, 6.0800, -2.0000, 2.5200, 6.4500, -2.2400, 2.4500, 6.0000, -2.0500, 1.8400, 4.6500, -1.3800, 1.3000, 4.6400, -1.3800, 1.3000, 4.6400, -1.2800, 1.8600, 6.9400, -1.3000, 2.5600, 9.0300, -1.5400, 2.7600, 8.5000, -1.7700, 1.6400, 6.1400, -1.6800, 1.4200,
7.5900, -1.3200, 2.0800, 9.8300, -1.3200, 2.0800, 9.8300, -0.8200, 2.1600, 10.3900, -0.7800, 1.7300, 9.8300, -1.1300, 1.3400, 9.7100, -1.3600, 1.6800, 10.2400, -1.5200, 1.6000, 9.3200, -1.8700, 1.4900, 9.1900, -1.8700, 1.4900, 9.1900, -1.9300, 1.0600, 9.9500, -1.3100, 0.8100, 10.6900, 0.0200, 2.0400, 11.0600, 0.2700, 2.5800, 9.3900, -0.0500, 2.2800, 7.3200, -0.3000, 0.4400,
7.6300, -0.3000, 0.4400, 7.6300, -1.4600, 1.0800, 12.3700, -1.9600, 1.7500, 15.3800, -0.7100, 2.1500, 14.0700, 0.7400, 1.7800, 10.4700, 0.6800, 0.8900, 9.9500, 0.0400, 1.5200, 12.0800, 0.0400, 1.5200, 12.0800, -0.4900, 1.7900, 12.7500

Example of a signal processing pipeline

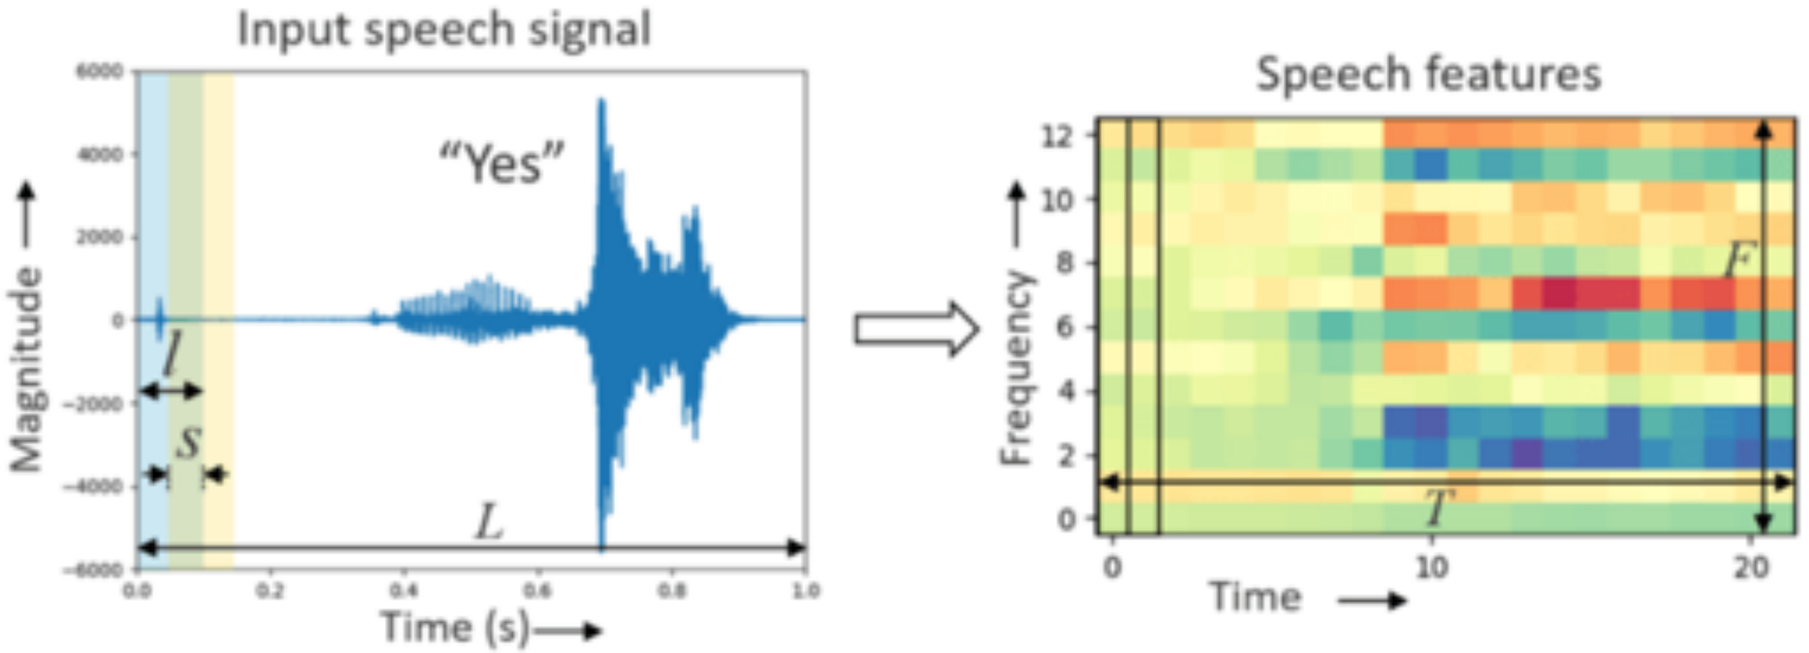

32,000 => 240

Before and after feature extraction

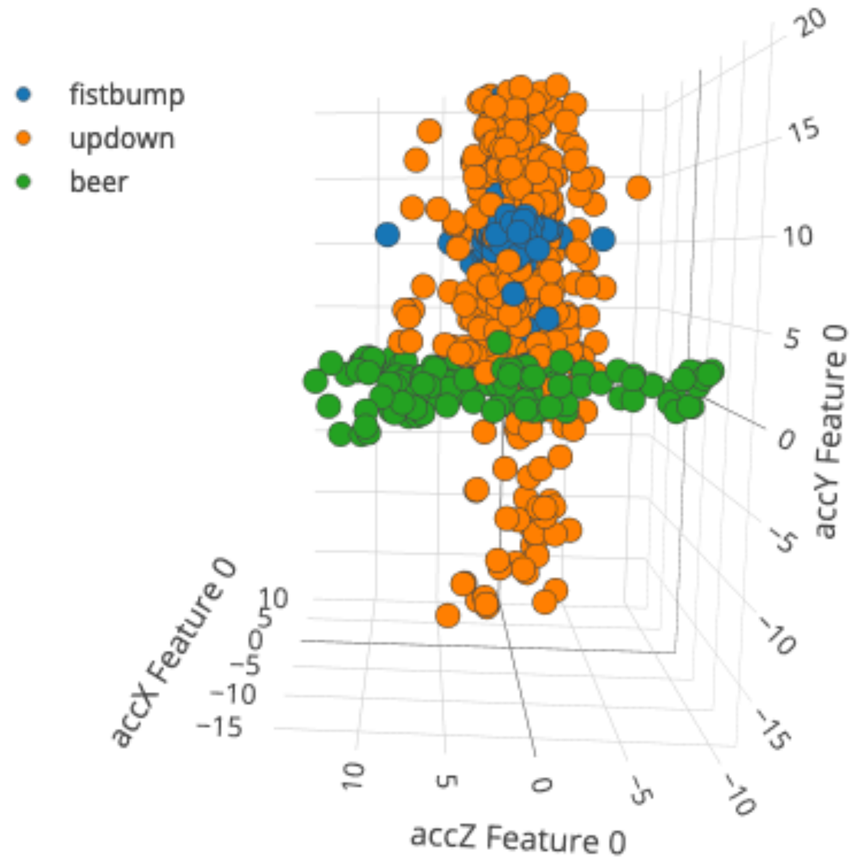

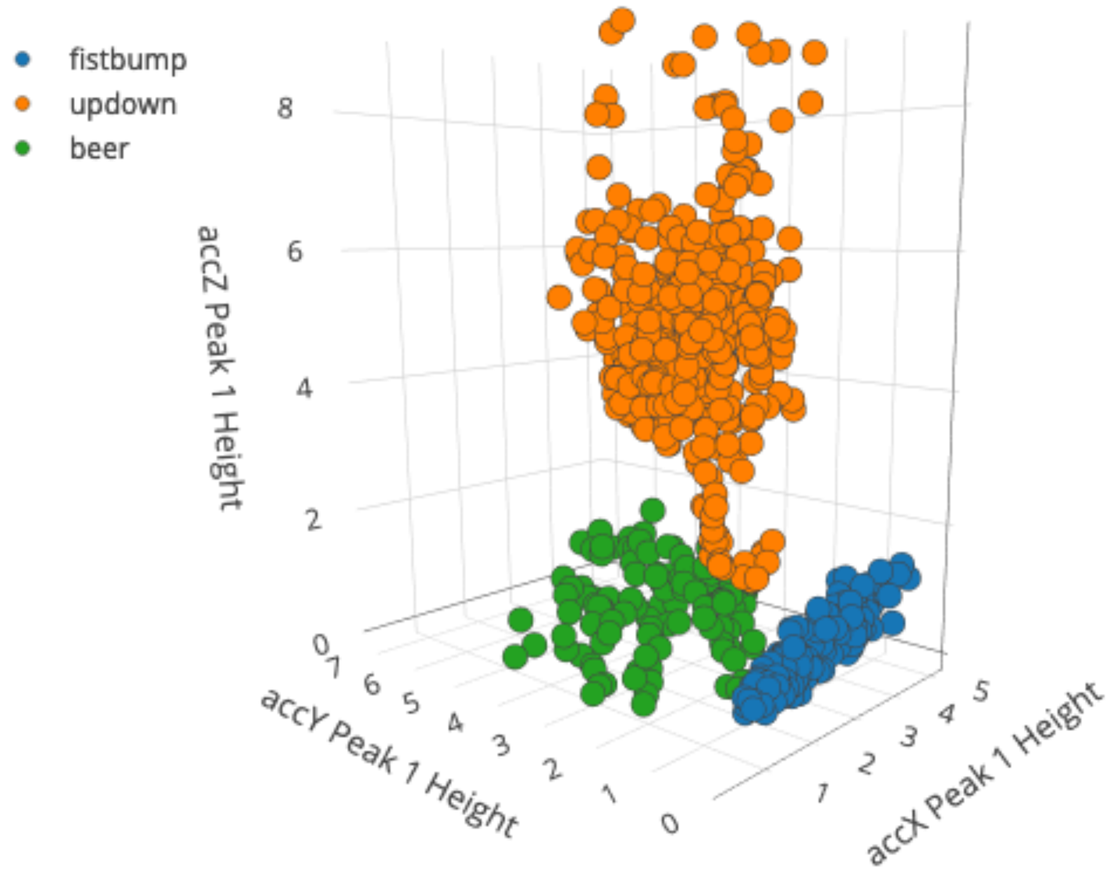

3. Letting the computers figure it out

Classification

What's happening right now?

Anomaly detection

Is this behavior out of the ordinary?

Forecasting

What will happen in the future?

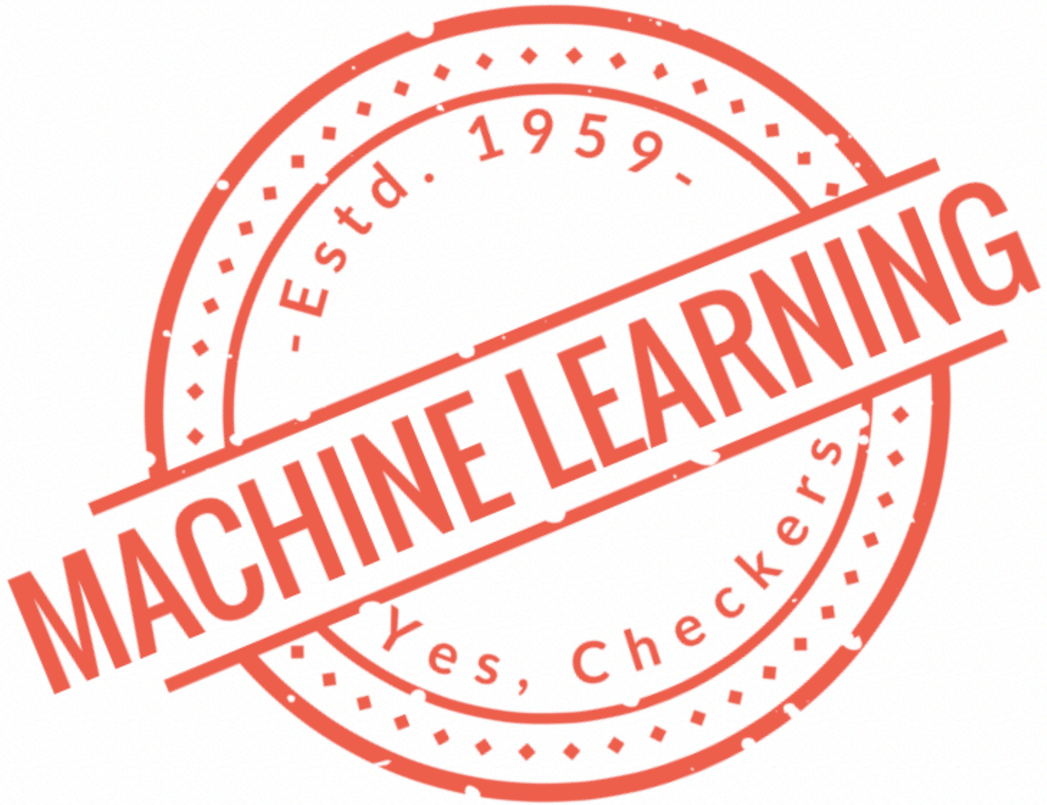

Picking the right algorithm

Classification

Neural network

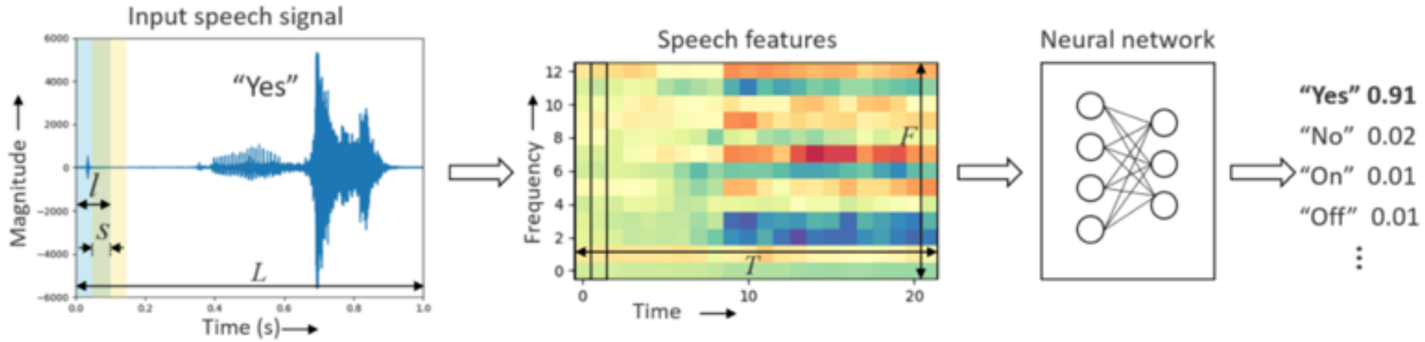

Anomaly detection

K-means clustering

Forecasting

Regression

4. Deploying

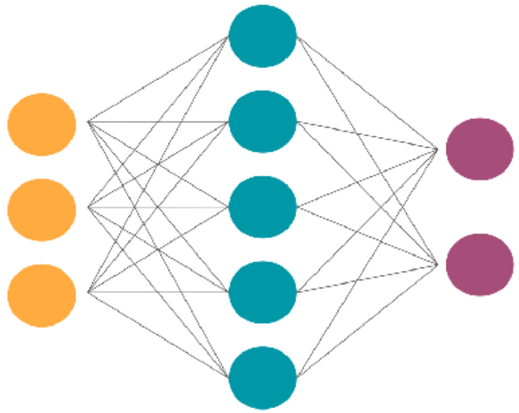

Signal processing, neural network and anomaly detection

Taking advantage of the hardware

Gateway or cloud - C++ to WebAssembly, adds isolation without VM. Trivial to deploy.

Cortex-M with FPU - use CMSIS-DSP and CMSIS-NN to leverage vector extensions.

Cortex-M with NN engine or external DSP - offload DSP and ML operations to co-processor.

Models driven by silicon capabilities

Cortex-M0+ - vibration and accelerometer classification, anomaly detection.

Cortex-M4F - realtime audio (continuous and discrete).

Cortex-M7 - basic image classification (10 frames/sec @ 96x96).

Cortex-A class - object detection, anomaly detection on a fleet of devices.

Cloud - anything (as long as you want to pay! $$$)

Co-processors can offset this, e.g. Eta Compute Cortex-M3 + external DSP can do realtime audio too.

Getting started 🚀

Get some hardware

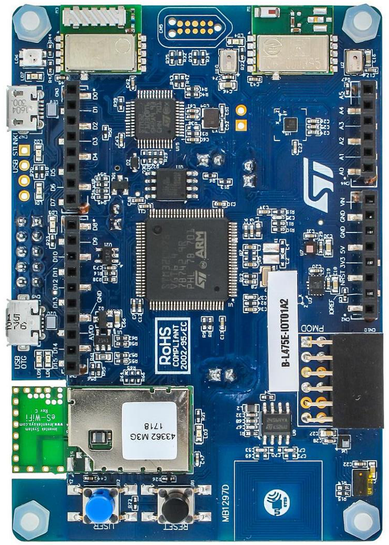

ST IoT Discovery Kit

Cortex-M4F 80MHz

Arduino Nano 33 BLE Sense

Cortex-M4F 64MHz

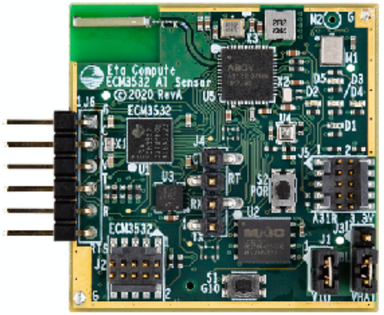

Eta Compute AI Sensor

Cortex-M3 + external DSP

OpenMV Cam H7+

Cortex-M7 480MHz

Any smartphone

Edge Impulse - TinyML as a service

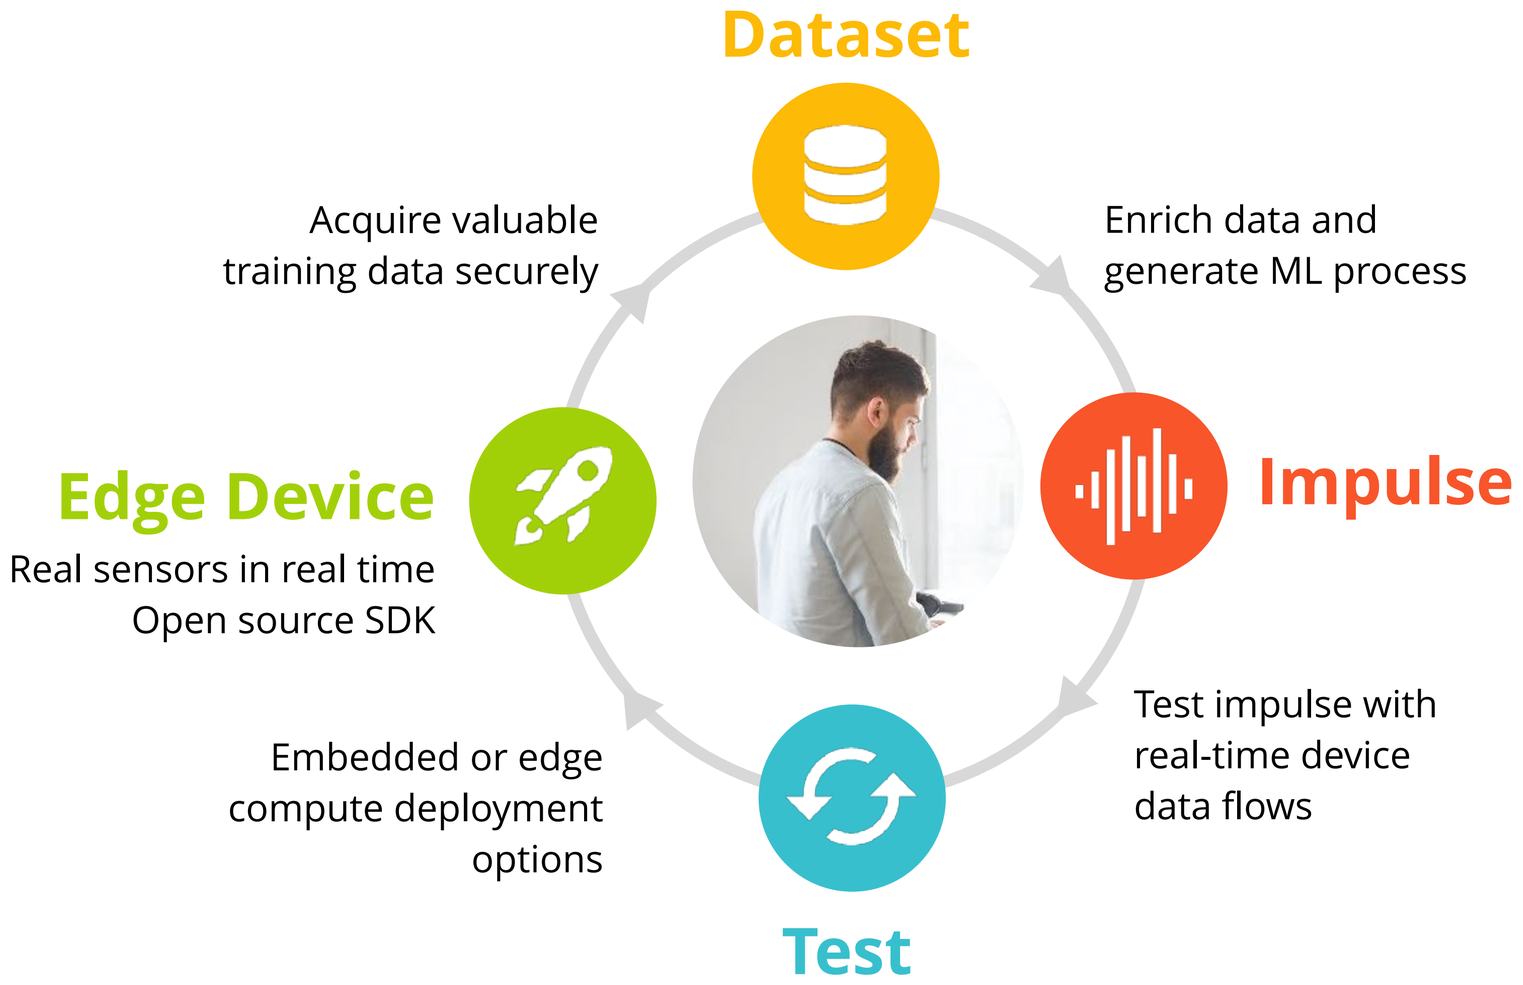

Free for developers, enterprise plan available: **edgeimpulse.com**

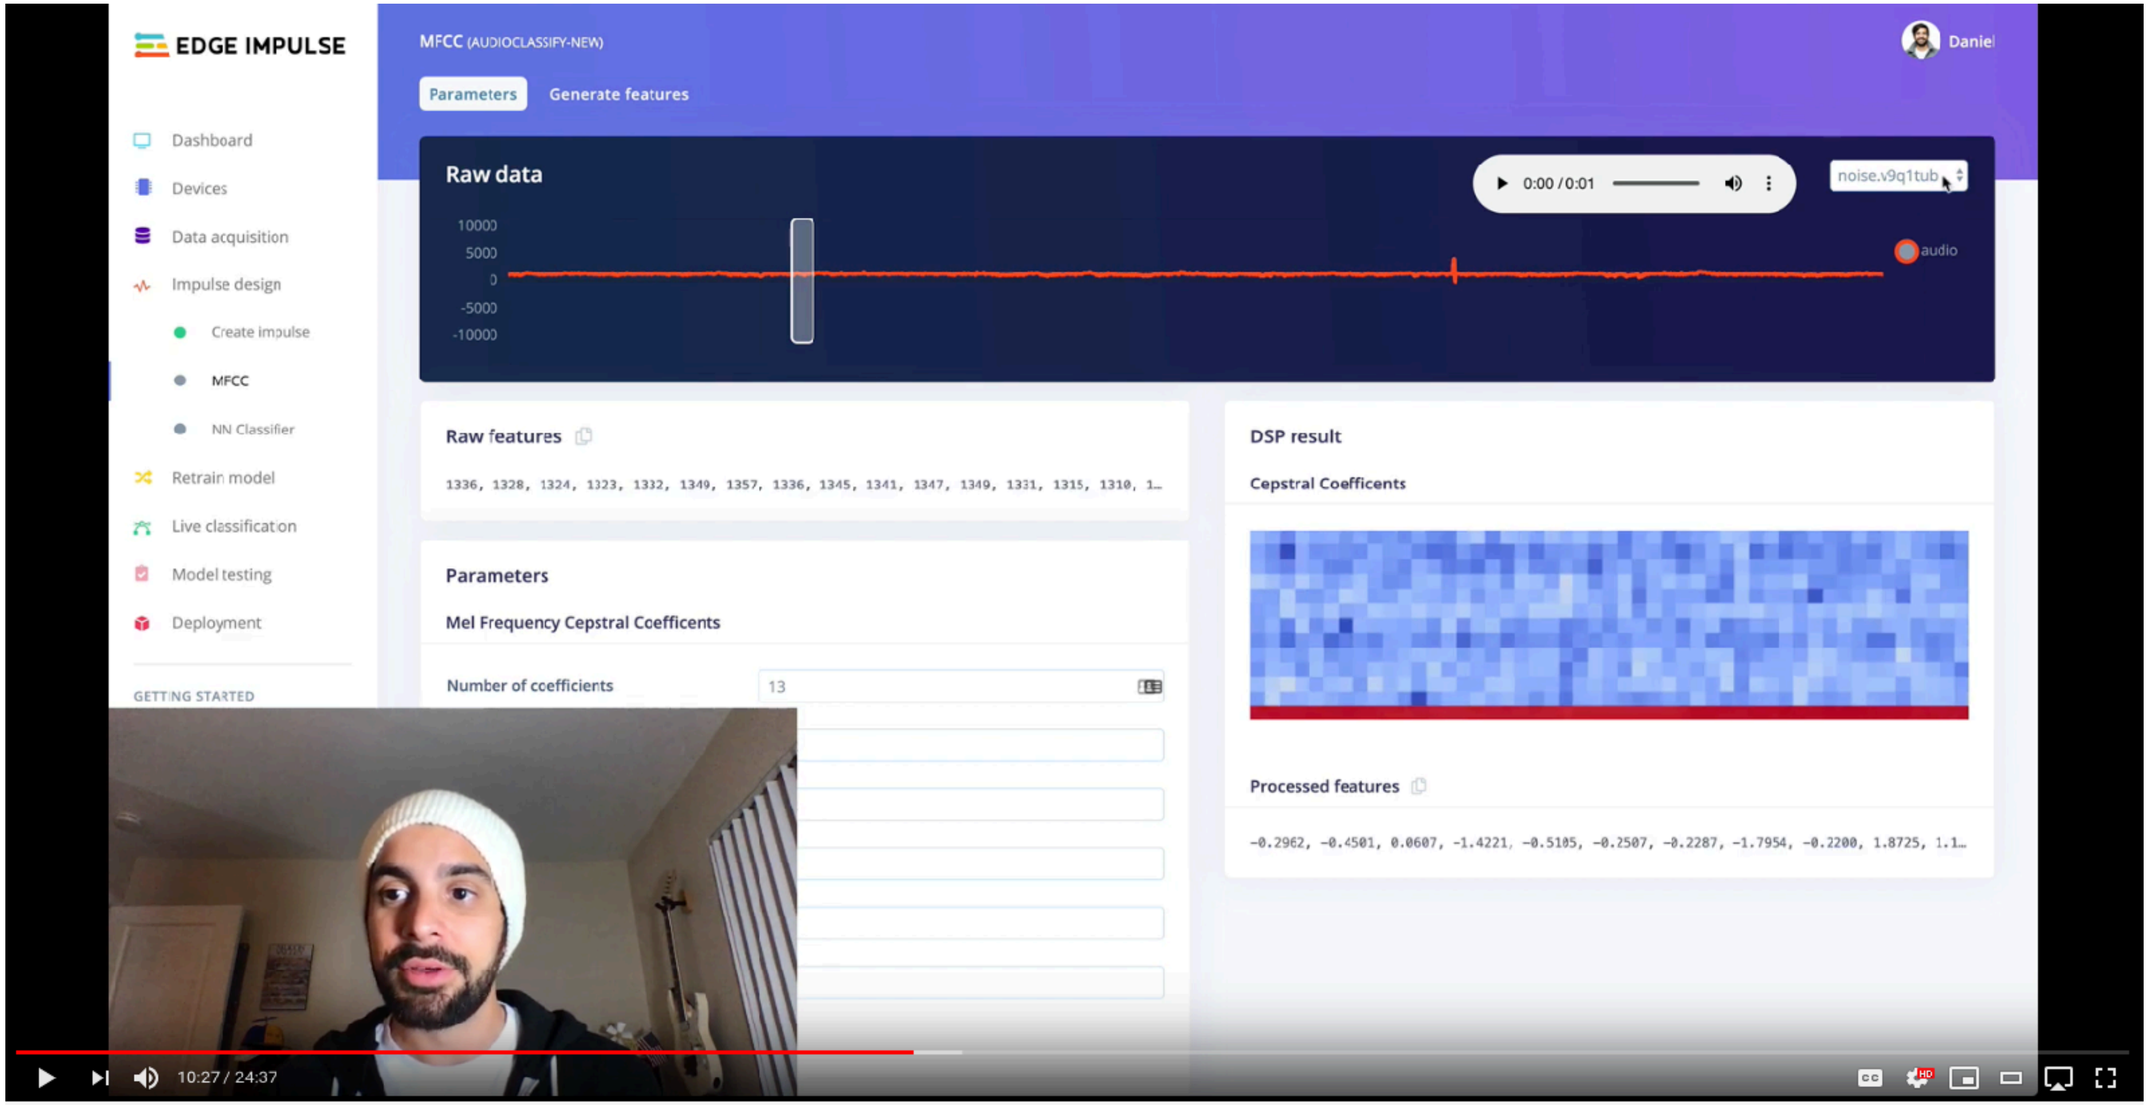

Building an audio classifier with embedded Machine Learning

697 views · Jan 26, 2020

21 0 SHARE SAVE

End-to-end tutorials on vibration, audio, and vision: docs.edgeimpulse.com

Demo 🚀

Recap

The ML hype is real

ML + sensors = perfect fit

Start doing the impossible!

edgeimpulse.com

AI Virtual Tech Talks Series

Date	Title	Host
September 8, 2020	Running Accelerated ML Applications on Mobile and Embedded Devices using Arm NN	Arm
September 22, 2020	How To Reduce AI Bias with Synthetic Data for Edge Applications	Dori Ai
October 20, 2020	Optimizing Power and Performance for ML At the Edge – Model Deployment Overview	Arm

Visit: developer.arm.com/solutions/machine-learning-on-arm/ai-virtual-tech-talks

Join us at Arm DevSummit

Oct 6 - 8 l Virtual Conference

Register here https://devsummit.arm.com/arm-ai-ml

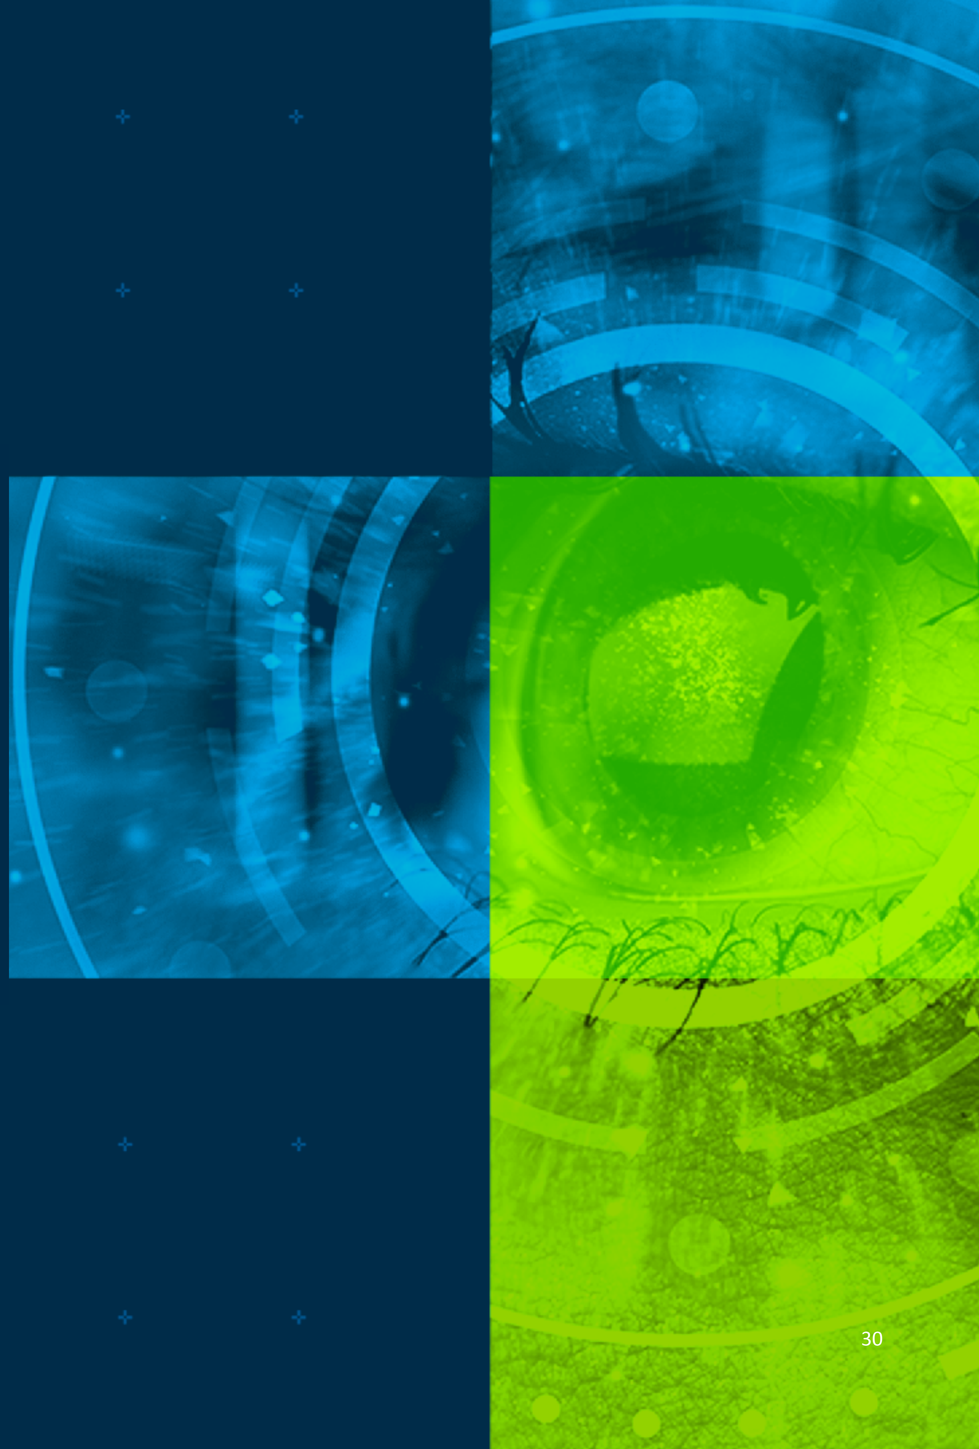

arm AI

AI Virtual Tech Talks Series

EDGE IMPULSE

Thank You
Danke
Merci
谢谢
ありがとう
Gracias
Kiitos
감사합니다
धन्यवाद
شكرًا
תודה