# tinyML® EMEA

*Enabling Ultra-low Power Machine Learning at the Edge*

# tinyML EMEA Technical Forum 2021 Proceedings

June 7 – 10, 2021

Virtual Event

www.tinyML.org

Photo by Braydon Anderson on Unsplash

# SECOND SENSE

## LESSONS LEARNED FROM BUILDING AN ARTIFICIAL NOSE

Benjamin Cabé
@kartben


COFFEE
PREMIUM QUALITY
ESPRESSO

# BENJAMIN CABÉ

Principal Program Manager,
Azure IoT – Microsoft

Open Source &
Community Advocate

**@kartben**


MIS EN BOUTEILLE AU CHATEAU
CHATEAU DE COSTIS
2010
GRAND VIN DE BORDEAUX
BORDEAUX
APPELLATION BORDEAUX CONTRÔLÉE


Francine
Francine
Francine

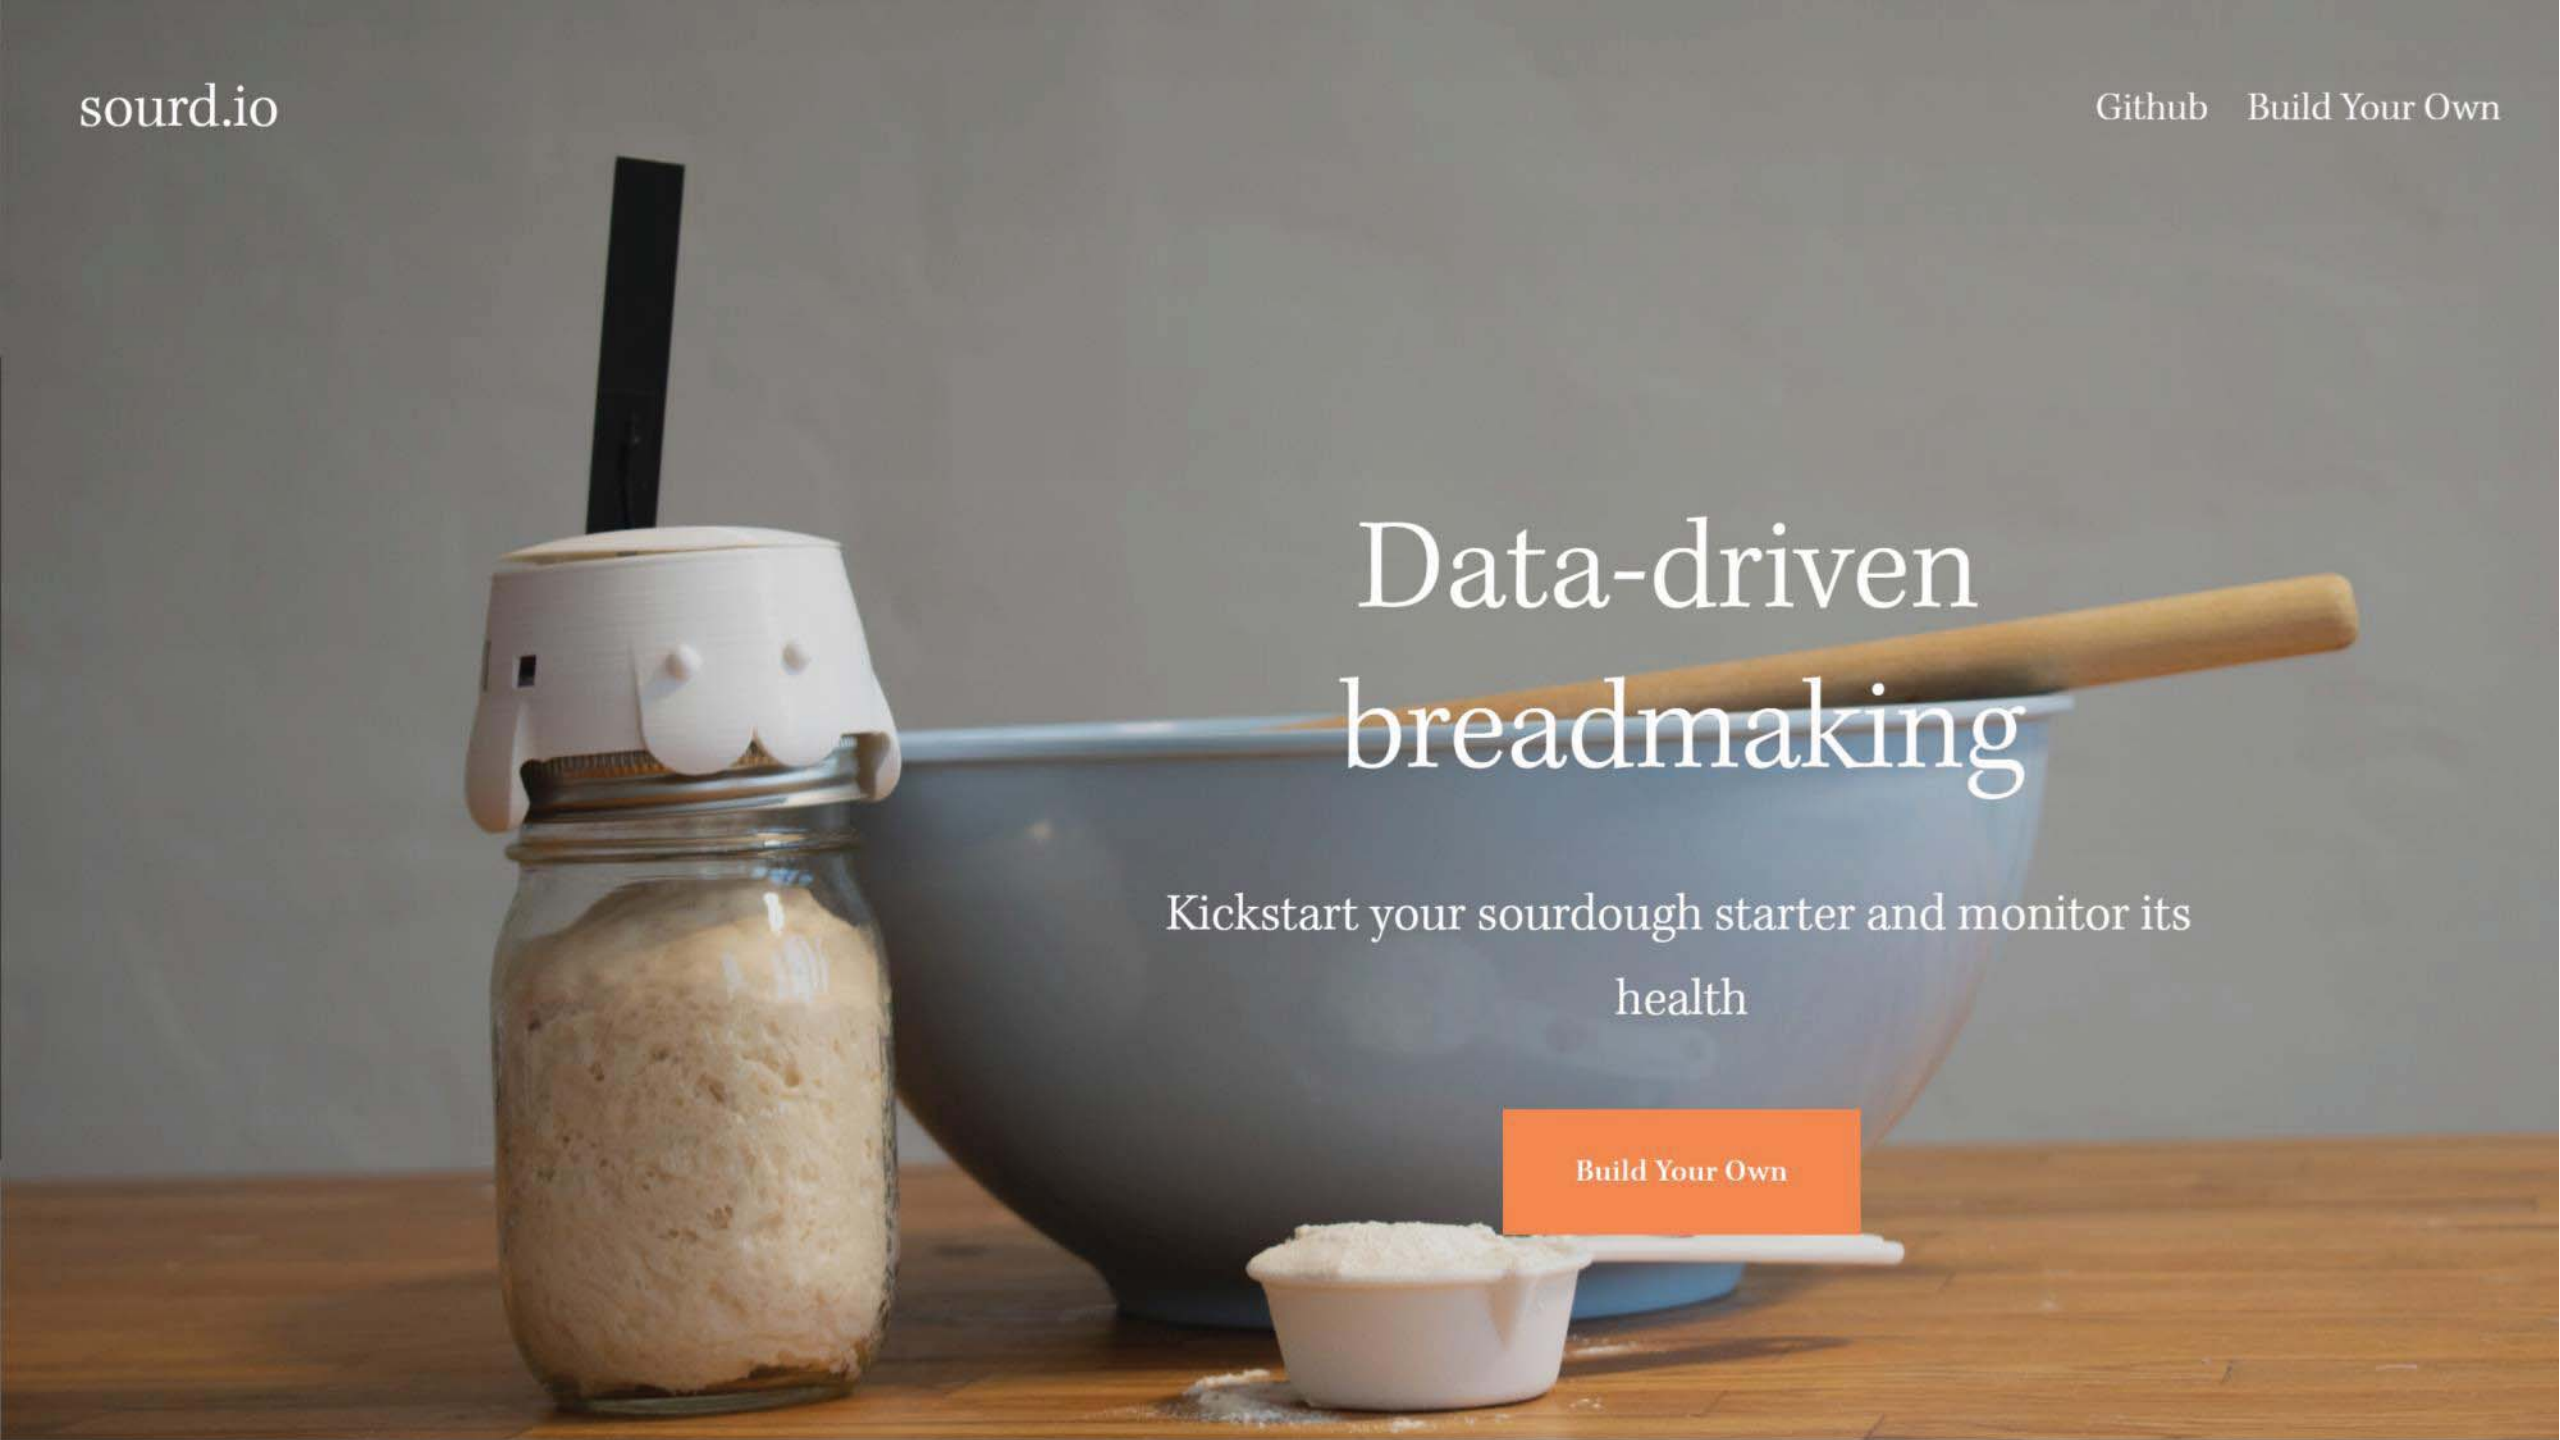
sourd.io
Github
Build Your Own
Data-driven breadmaking
Kickstart your sourdough starter and monitor its health
Build Your Own

# TINYML?

> The ability to run a neural network model at an energy cost of below 1 mW.

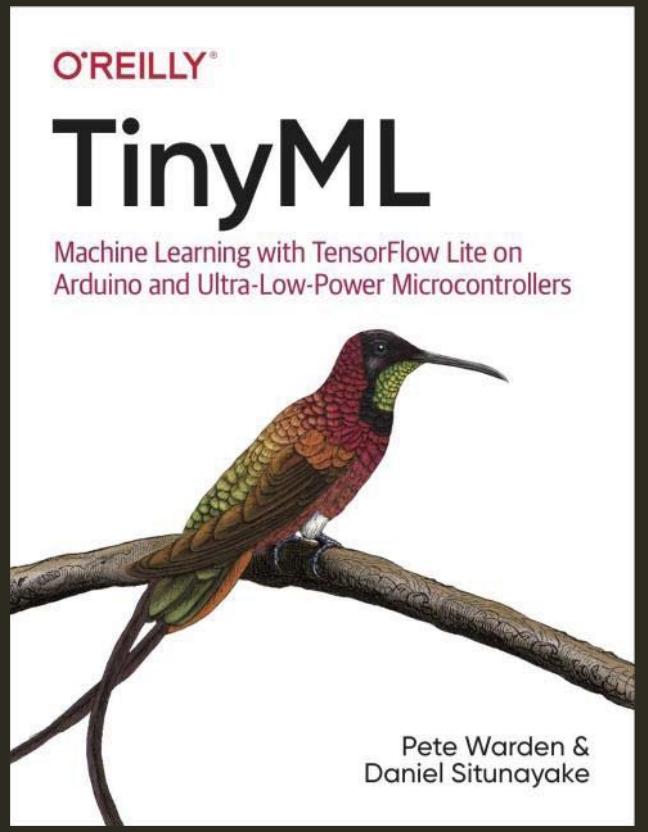

DATASET 

\+ HARDWARE 

\+ MODEL 

+FRAMEWORK 

---

# DEEP LEARNING APPLICATION

# ASSEMBLING A DATASET

Wait… do I really want to bake dozens of baguettes to **gather enough data** to train my model?

➔ Let's instead build a model that "just" recognizes smells.

My device will be rather **constrained**, how can I efficiently capture my training data?

➔ Edge Impulse to the rescue ☺


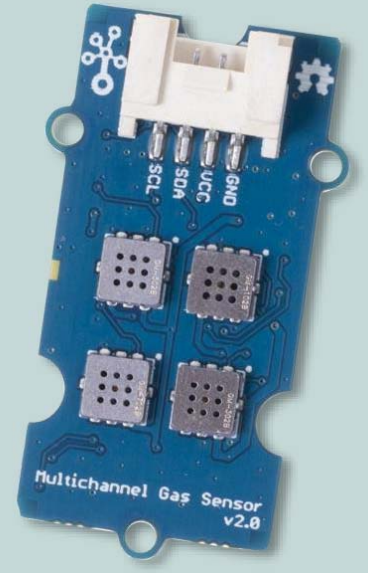



# WIO TERMINAL + GAS SENSOR

Carbon monoxide (CO)
Ethyl alcohol($C_2H_5OH$)
Nitrogen dioxide ($NO_2$)
Volatile Organic Compounds (VOC)

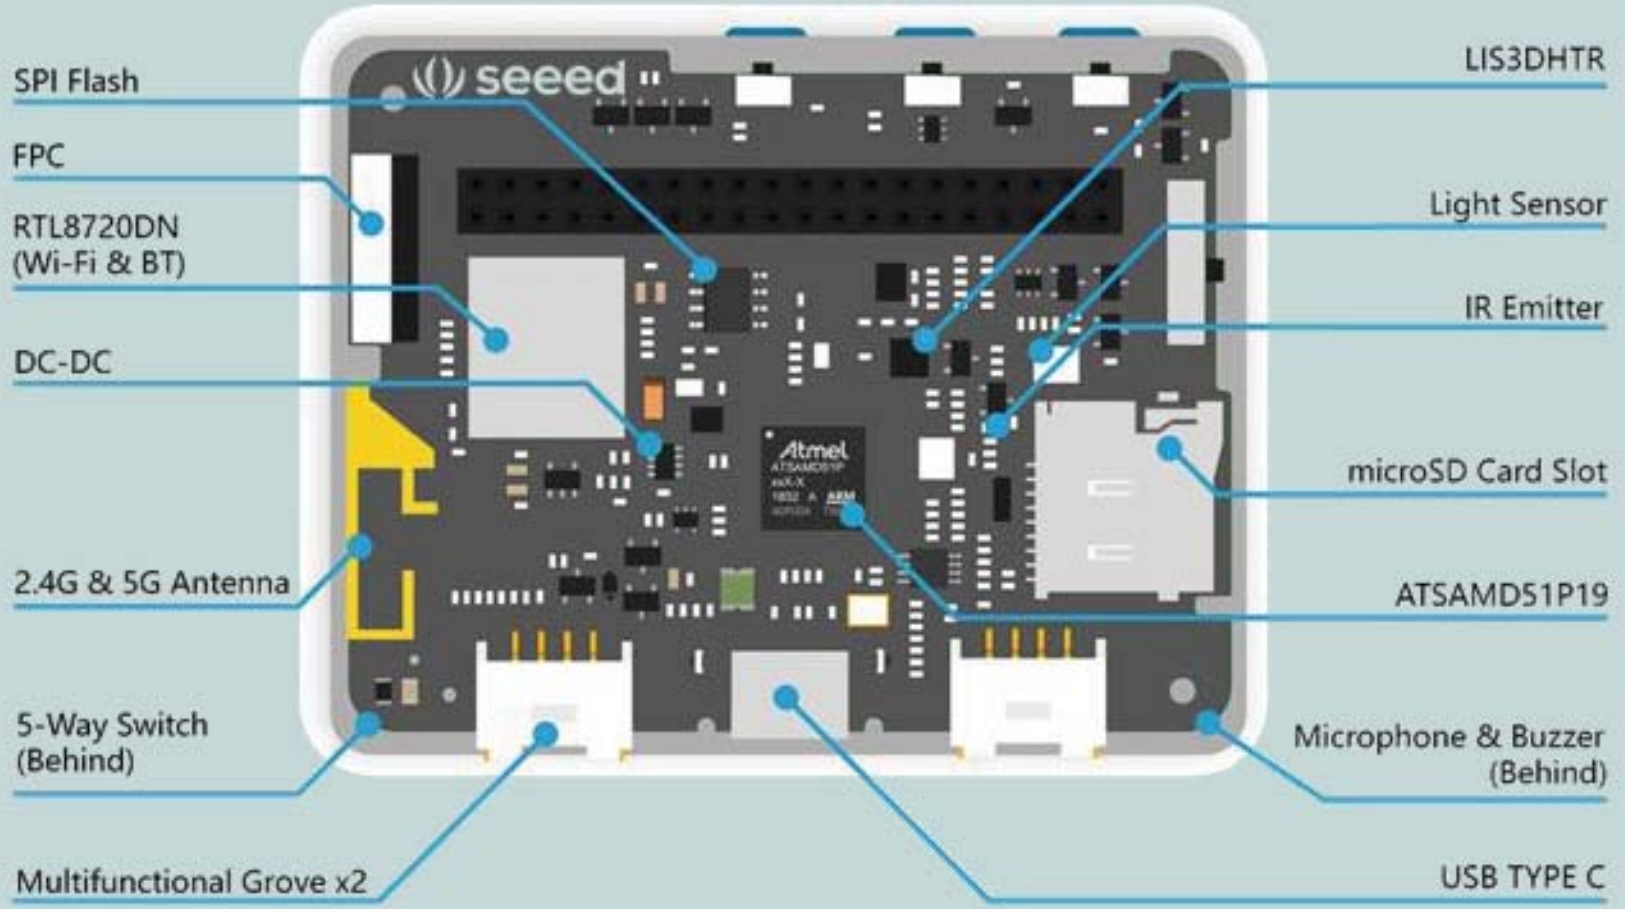


# A CLOSER LOOK AT THE WIO TERMINAL

Arm Cortex-M4F

192KB of RAM, 512KB of Flash, 4MB External Flash

# COST CONSIDERATIONS

**Wio Terminal**
~$30

Arm Cortex-M4
512K of Flash
192K of RAM


**ATSAMD51P19A**
~$5 (when ordering 3000+ units)

Arm Cortex-M4F
512K of Flash
192K of RAM

# BUT HOW CAN A MACHINE SMELL ANYWAYS?

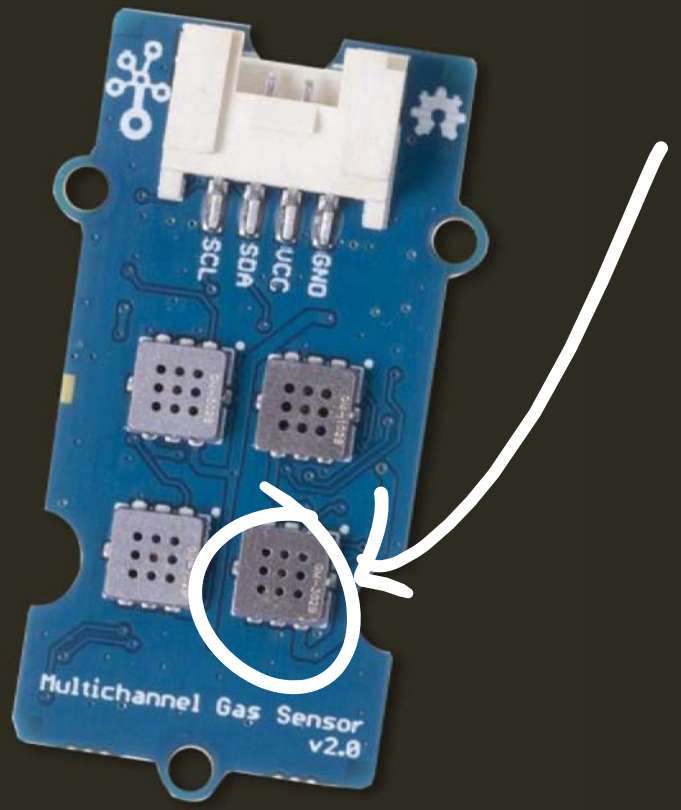

# 10 ppm

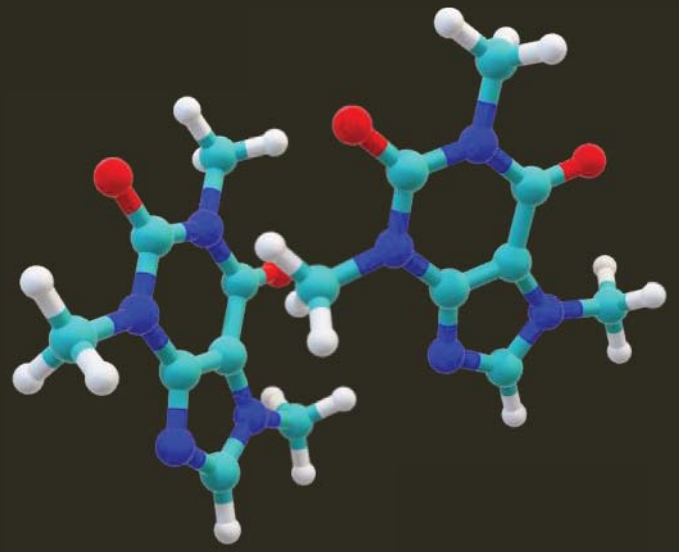

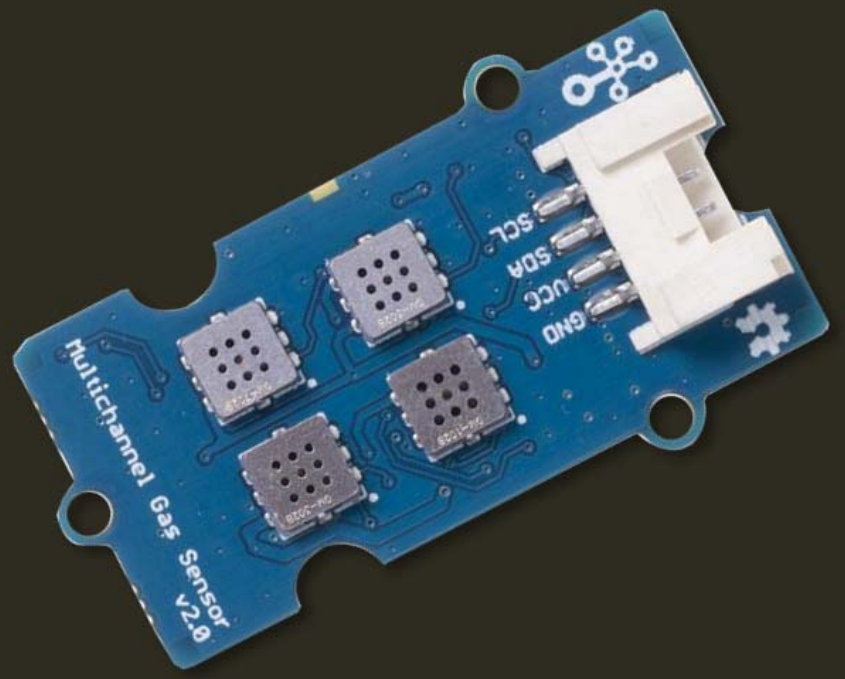


# VOC?

# *100 Ω!*

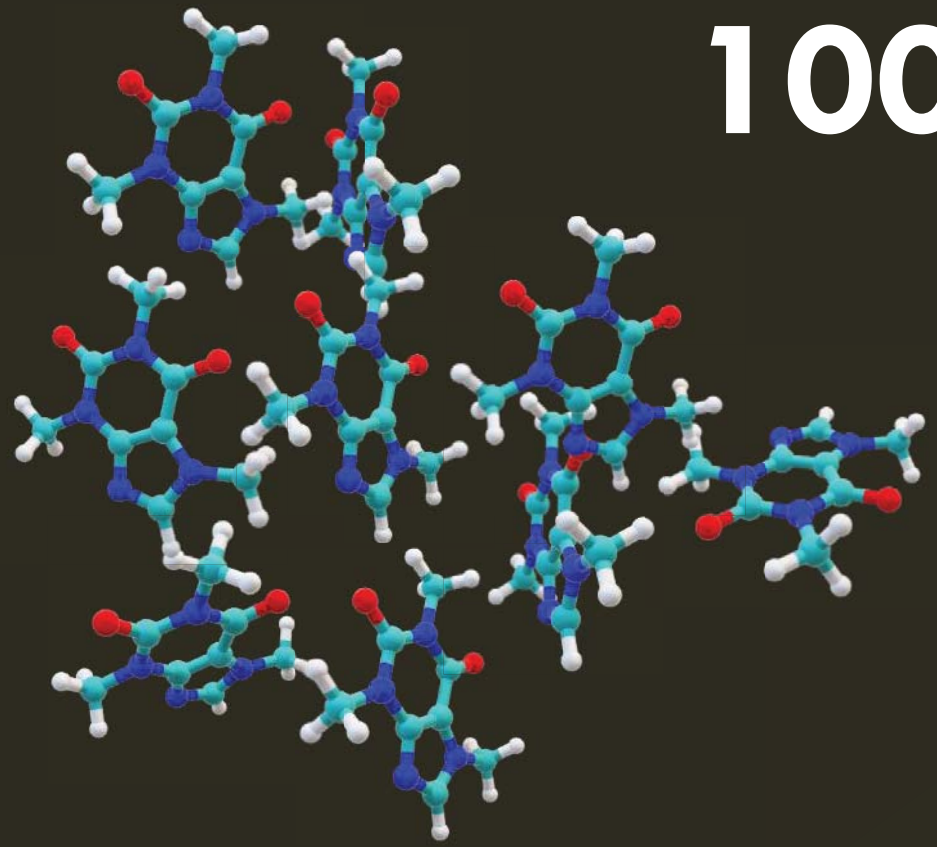

# 1000 ppm

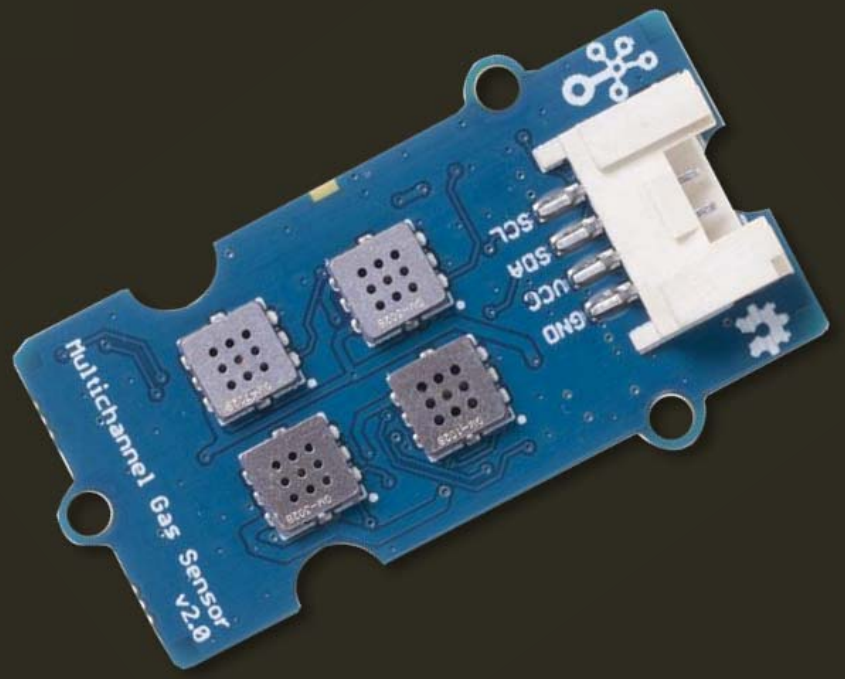


# VOC?

# *200 Ω!*

# Raw sensor data

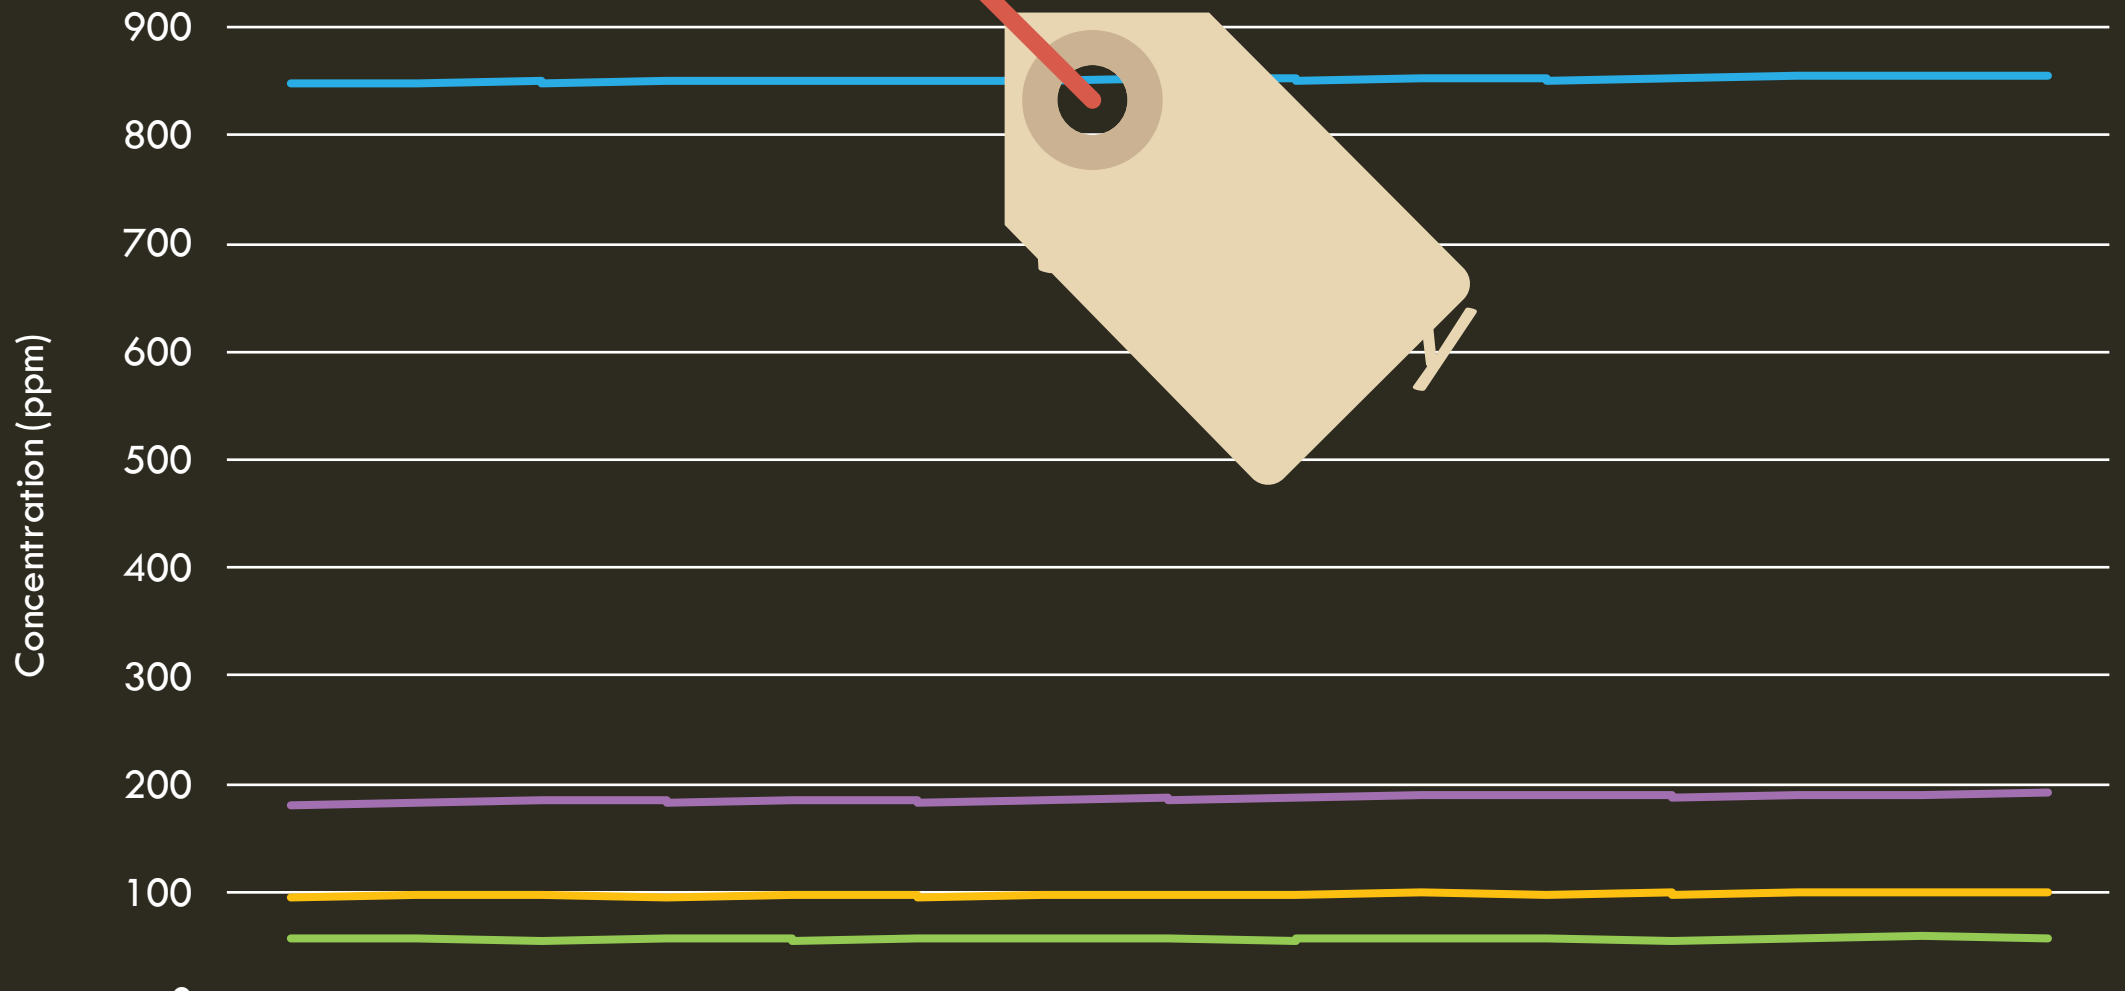


| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Nitrogen dioxide | 849 | 849 | 850 | 851 | 851 | 851 | 851 | 852 | 852 | 853 | 853 | 854 | 855 | 855 | 855 |
| Carbon monoxide | 57 | 57 | 56 | 57 | 57 | 58 | 58 | 58 | 57 | 57 | 57 | 56 | 57 | 59 | 58 |
| Ethyl alcohol | 95 | 96 | 96 | 95 | 97 | 97 | 98 | 98 | 98 | 99 | 98 | 99 | 100 | 100 | 100 |
| Volatile organic compounds | 182 | 183 | 184 | 184 | 185 | 185 | 186 | 187 | 188 | 189 | 189 | 189 | 190 | 190 | 191 |

# BUT HOW CAN A MACHINE SMELL ANYWAYS?

Raw features (whiskey) – 1.5 s of sensor data, 10 Hz

[ 849, 57, 95, 182, 849, 57, 96, 183, 850, 56, 96, 184, 851, 57, 95, 184, 851, 57, 97, 185, 851, 58, 97, 185, 851, 58, 98, 186, 852, 58, 98, 187, 852, 57, 98, 188, 853, 57, 99, 189, 853, 57, 98, 189, 854, 56, 99, 189, 855, 57, 100, 190, 855, 59, 100, 190, 855, 58, 100, 191 ]

# BUT HOW CAN A MACHINE SMELL ANYWAYS?

really

Flattened features (whiskey) – step 1:

```
[
 [ 849, 57,  95, 182 ],
 [ 849, 57,  96, 183 ],
 …
 [ 855, 58, 100, 191 ]
]
```

# BUT HOW CAN A MACHINE SMELL ANYWAYS?

really

Flattened features (whiskey) – step 2 (scale axes):

```
[
 [ 0.849, 0.057, 0.095, 0.182 ],
 [ 0.849, 0.057, 0.096, 0.183 ],
 ...
 [ 0.855, 0.058, 0.100, 0.191 ]
]
```

# BUT HOW CAN A MACHINE SMELL ANYWAYS?

really

Flattened features (whiskey) – step 3 (DSP):

[ 0.8520, 0.849, 0.855, 0.8250, 0.0019,
0.0572, 0.056, 0.059, 0.0554, 0.0007,
0.0977, 0.095, 0.100, 0.0946, 0.0016,
0.1868, 0.182, 0.191, 0.1808, 0.0027 ]

$NO_2$

CO

$C_2H_5OH$

VOC

avera ge

min

max

root-mean

standar d devia on

# MODEL?

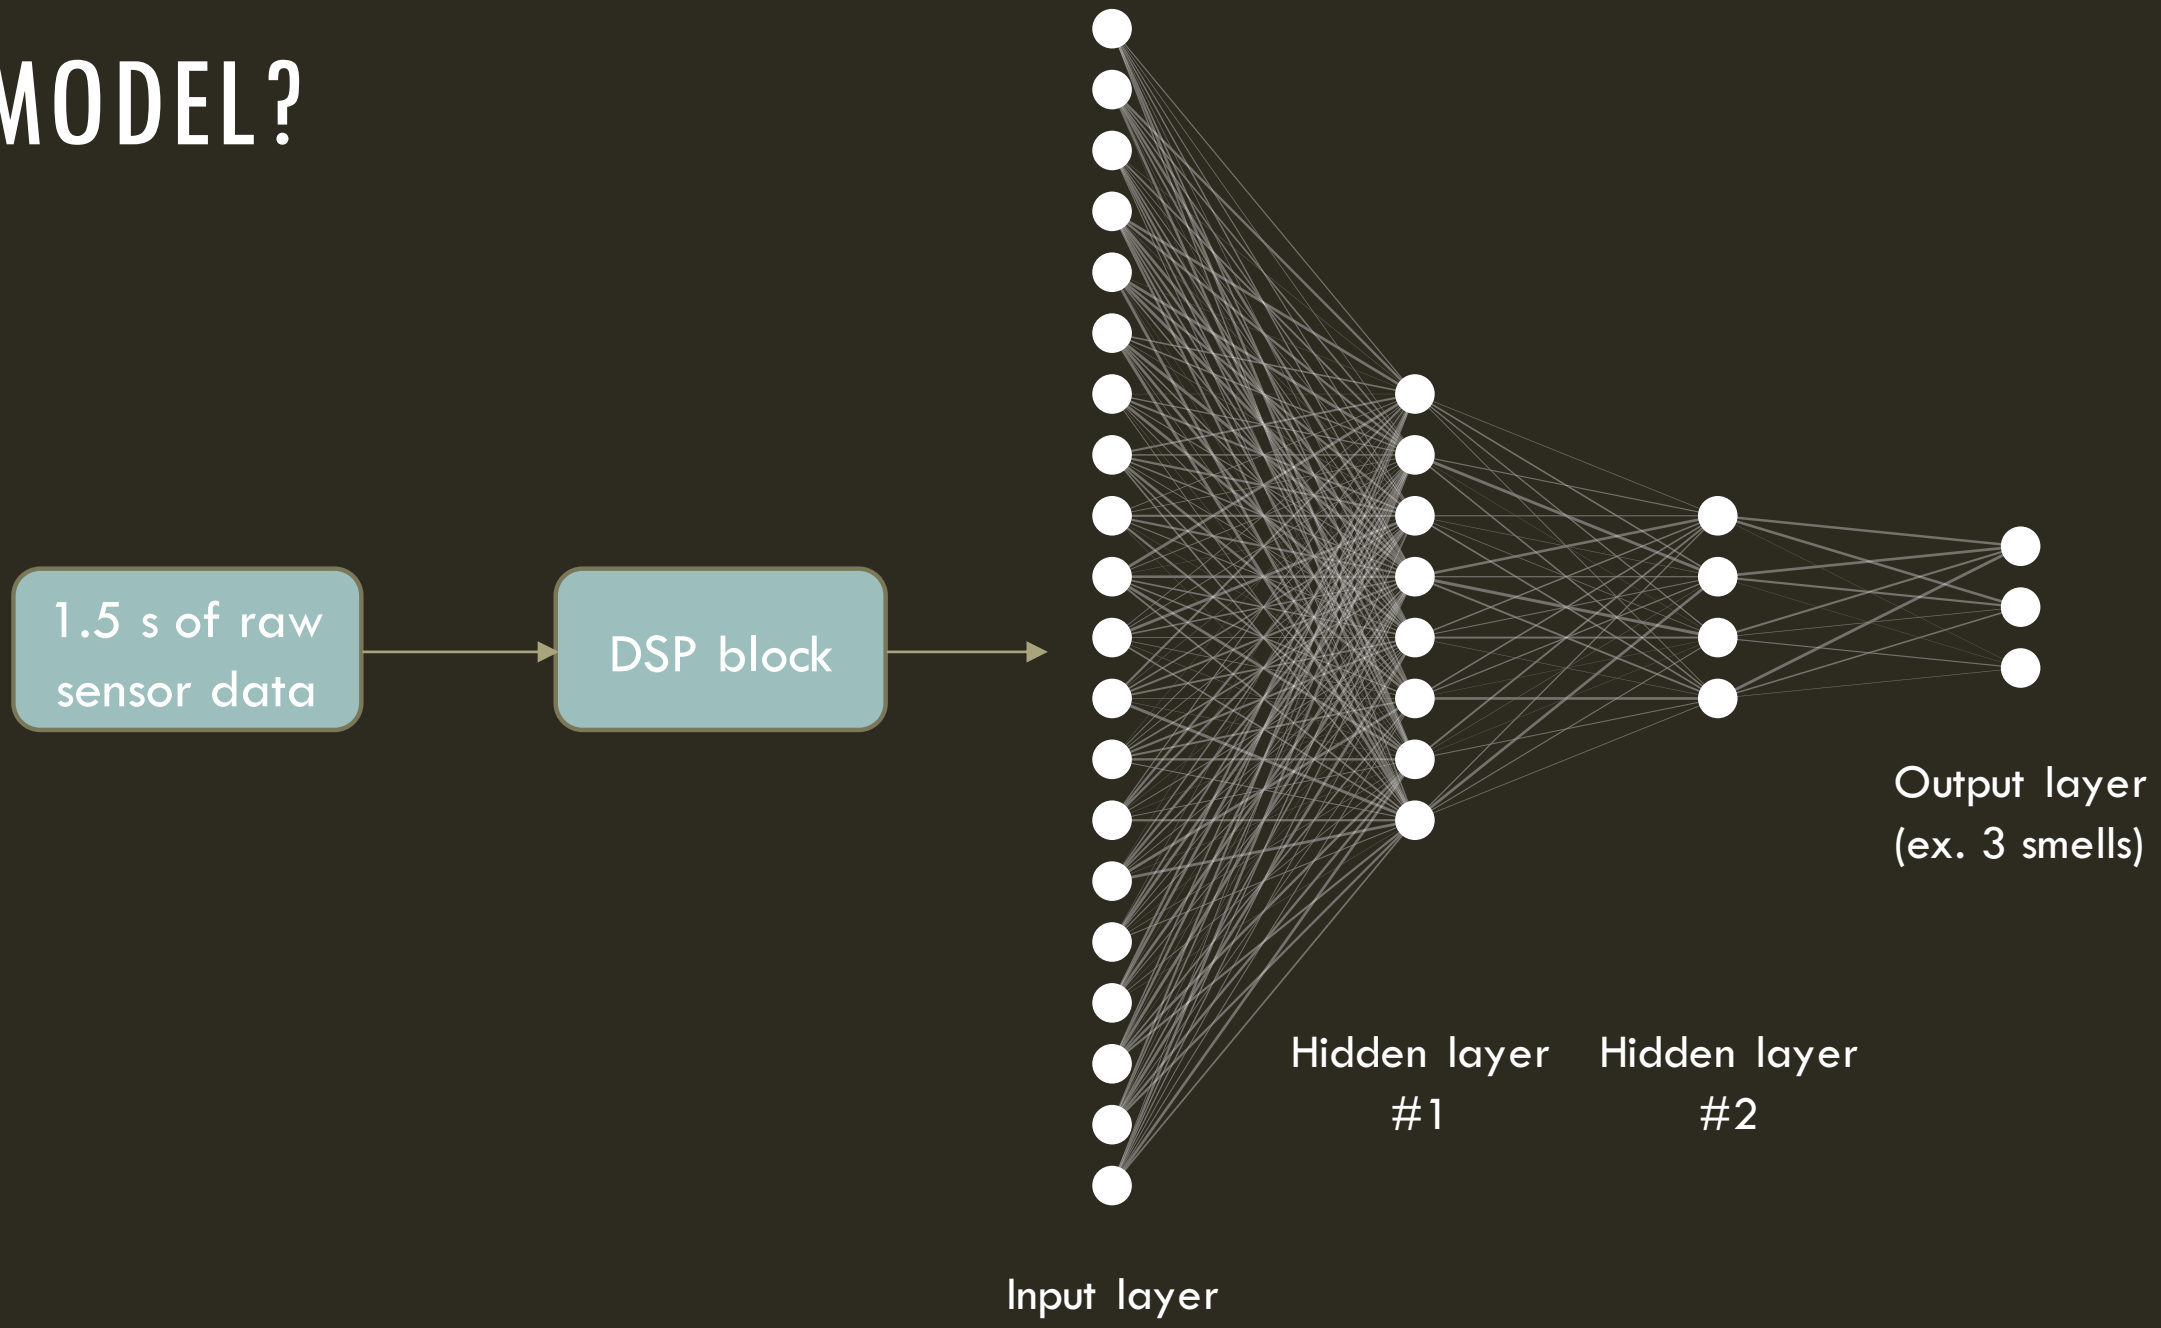
1.5 s of raw sensor data
DSP block
Input layer
Hidden layer #1
Hidden layer #2
Output layer (ex. 3 smells)

# TENSORFLOW LITE FOR MICROCONTROLLERS

**Optimized for on-device machine learning**

- **latency** – there's no round-trip to a server
- **privacy** – no personal data leaves the device
- **connectivity** – Internet connectivity is not required
- **size** – reduced model and binary size
- **power consumption** – efficient inference & a lack of network connections

**High performance** (hardware acceleration and model optimization)

Available as **Arduino library**

# ON PERFORMANCE AND CODE SIZE

Classifying 3-5 smells:

- ~4KB of RAM, ~27KB of ROM (the actual TFLite model is ~3KB)
- Inference is ~1ms on an 80MHz 32-bit MCU

**Quantization**

- Reducing the precision of the numbers used to represent a model's parameters. Think float32 → int8, i.e 4x size improvement

**EON compiler** – up to 50% less memory!

**CMSIS-DSP** & **CMSIS-NN**

# OH, A FRIENDLY REMINDER…

**Classification** using neural networks <u>always</u> gives you a result!

 51% WHISKEY!

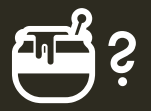 78% COFFEE!

→ **Anomaly detection** helps flag input data that is too different from data seen during training


# “GREAT… SO… YOU’VE BEEN TINKERING, EH?”

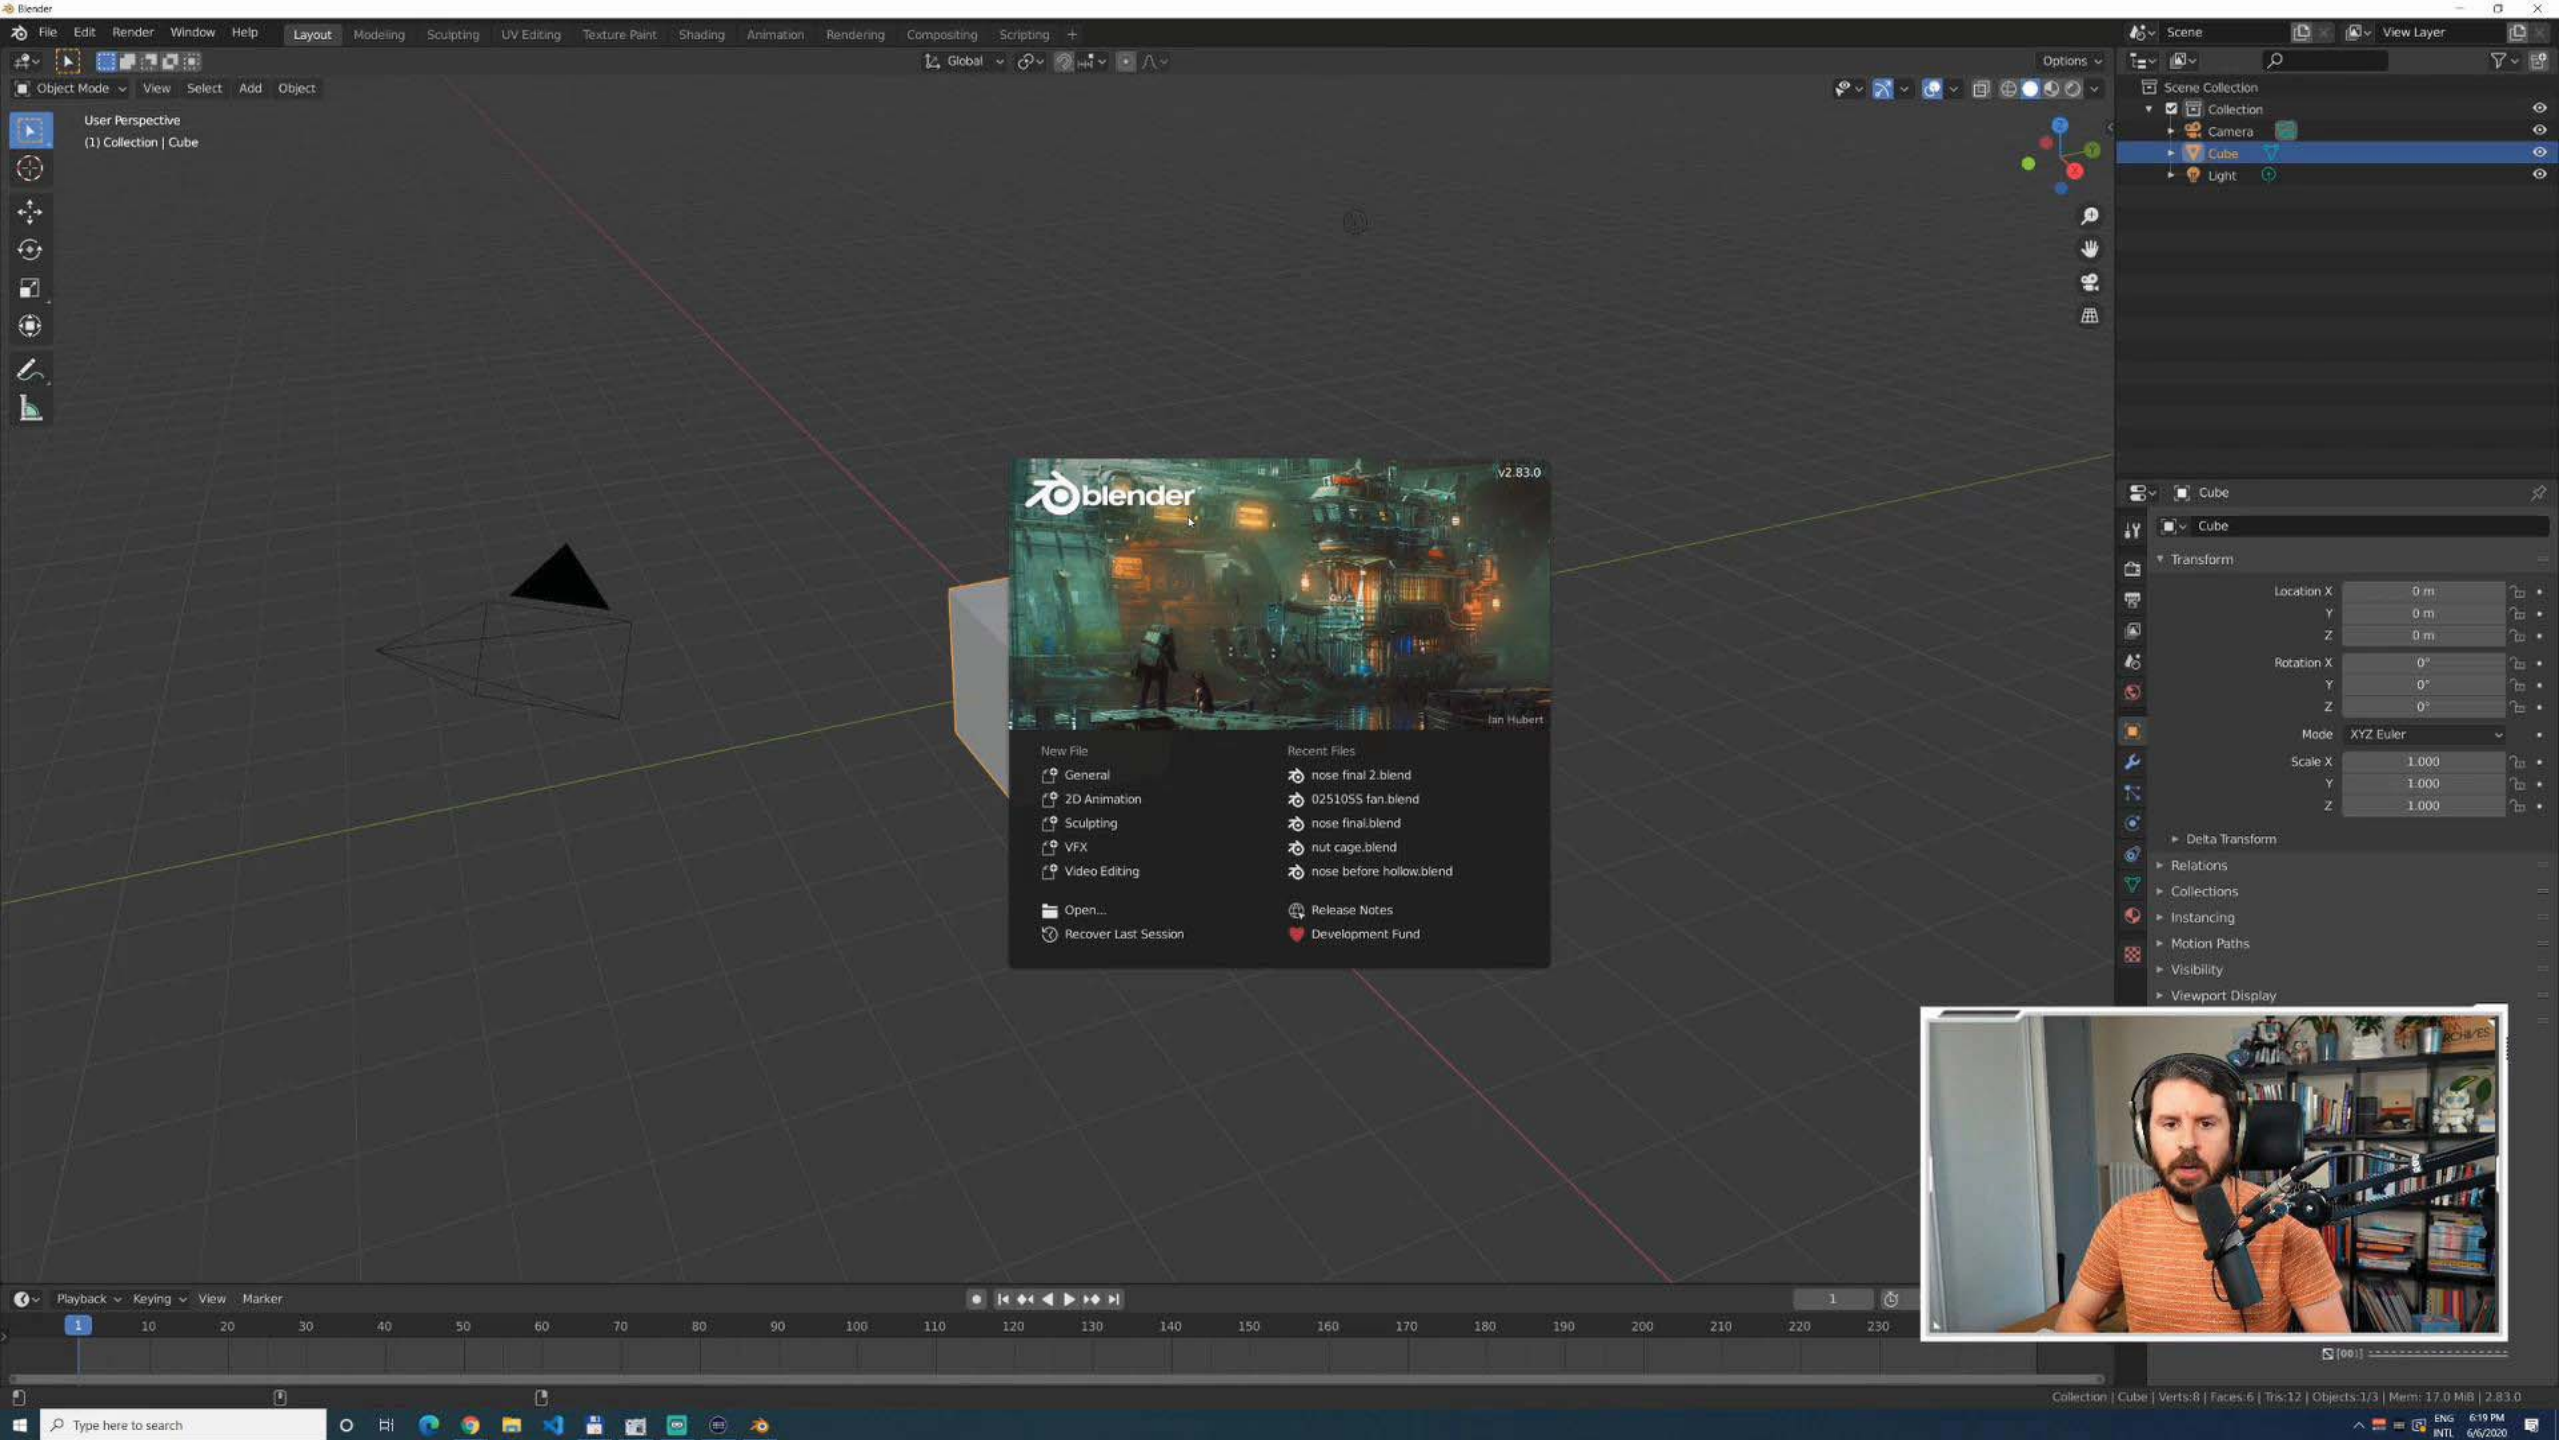
File Edit Render Window Help
Layout Modeling Sculpting UV Editing Texture Paint Shading Animation Rendering Compositing Scripting
Object Mode View Select Add Object
User Perspective
(1) Collection | Cube
Global
Options
Scene
View Layer
Scene Collection
Collection
Camera
Cube
Light
Transform
Location X 0 m
Y 0 m
Z 0 m
Rotation X 0°
Y 0°
Z 0°
Mode XYZ Euler
Scale X 1.000
Y 1.000
Z 1.000
Delta Transform
Relations
Collections
Instancing
Motion Paths
Visibility
Viewport Display
blender
v2.83.0
Ian Hubert
New File
General
2D Animation
Sculpting
VFX
Video Editing
Open...
Recover Last Session
Recent Files
nose final 2.blend
02510SS fan.blend
nose final.blend
nut cage.blend
nose before hollow.blend
Release Notes
Development Fund
Playback Keying View Marker
Collection | Cube | Verts:8 | Faces:6 | Tris:12 | Objects:1/3 | Mem: 17.0 MiB | 2.83.0
Type here to search





THINGIVERSE.COM/THING:4493907

**Benjamin Cabé** @kartben · Feb 12

I realized I **never quite published** the instructions to replicate my #TinyML and #IoT artificial **nose** project, powered by awesome tech from @EdgeImpulse and hardware from @seeedstudio. Working on getting this fixed asap while sipping my espresso! 👃☕

8 22 97

# Make:

MOTORIZED
**SELF SOLVING**
RUBIK'S CUBE

# MACHINE LEARNING

**Swear Bear... Trash Sorter... Easy AI Trainers**
Teach your project to think for itself

**BENJAMIN CABÉ'S "NOSE" KNOWS!**
Build this smell-identifying AI sniffer


## 23 PROJECTS!

- **Raspberry Pi Meteor Camera**
- Digital String Art Portraits
- **Animated LED Skirt**
- Arduino Borealis Lights

## SKILL BUILDERS

- **Digital Mobile Radio**
- Hack a Knitting Machine

makezine.com | makercampus.com | makerfaire.com

VOLUME 77

# SECOND SENSE

MAKE A SMART SNIFFER THAT CAN SORT COFFEE FROM TEA, CHOOSE YOUR FAVORITE BOOZE, OR WHATEVER ELSE YOU TRAIN IT TO SMELL!

Written and photographed by Benjamin Cabé


**BENJAMIN CABÉ** is a software engineer and IoT jack of all trades at Microsoft, living in Toulouse, France. Find him on Twitter @kartben.



It was a long weekend of May 2020. Like many of my human siblings stuck at home with time on their hands due to an ongoing pandemic, I was busy trying to perfect my bread recipe. In fact, just a few days before, I had ordered a gas sensor (Figure A) that I thought would be ideal to help me monitor my sourdough starter and bake my bread at just the right moment.

And then I thought about it some more. "Surely, this is the perfect excuse for me to finally start learning this *machine learning* thing that everyone's talking about. But … do I really want to bake dozens of baguettes before I have a training set large enough to teach an AI the relationship between the olfactory fingerprint of the sourdough starter and the yumminess of the final loaf? Plus, flour is pretty scarce these days!"

That's how, over the course of the next few days, I ended up building a DIY, general-purpose, artificial nose — one that can smell virtually anything you teach it to recognize! The artificial nose is powered by artificial intelligence — a TinyML neural network that I trained using the free online tool Edge Impulse and then uploaded onto an Arduino-compatible microcontroller.

I learned a lot along the way, and not just about machine learning. From designing my first 3D enclosure to rudimentary fluid dynamics (the airflow within the nose is not exactly optimal), it was the first time I built my own "thing" from scratch, so I'm excited to share it with the *Make:* community. Here are the steps for replicating the build for yourself.

**TIME REQUIRED:** A Few Hours

**DIFFICULTY:** Intermediate

**COST:** $80–$100

## MATERIALS

- **Wio Terminal microcontroller board with LCD display** Seeed Studio part #102991299
- **Grove Multichannel Gas Sensor v2** Seeed 101020820, measures nitrogen dioxide ($NO_2$), carbon monoxide (CO), ethyl alcohol ($C_2H_5OH$), and volatile organic compounds (VOCs)
- **3D printed nose enclosure** Download the 3D files for free at thingiverse.com/thing:4493907.
- **Grove MOSFET board** Seeed 103020008
- **Cable, 4-pin Grove connector to male jumper wires** Seeed 110990210
- **Fan, 5V DC, 25×25×10mm** such as NMB Technologies 02510SS
- **Fan finger guard, 25×25mm** Sunon FG-2
- **USB-C right-angle cable (optional)**
- **Wio Terminal Battery Chassis (optional)** Seeed 103990564
- **M2 and M3 screws, nuts, and washers**
- **Breadboard jumper wires, 100mm (2)**

## TOOLS

- **3D printer**
- **Computer with internet connection** for setup only; not needed for nose operation

## BUILD YOUR ARTIFICIAL NOSE

### 1. GET YOUR PARTS READY

You can 3D print the nose enclosure from thingiverse.com/thing:4493907 (Figure B). Alternatively, grab your Miniaturizer 3000™, fly to Easter Island, and capture your own 1:100 copy of a moai!

Note that you don't need the handsome enclosure to build the artificial nose, but you certainly need all the electronic components (Figure C on the following page). They're easy to put together; there's no soldering involved at all, just plugging in some jumper wires and connectors.


A

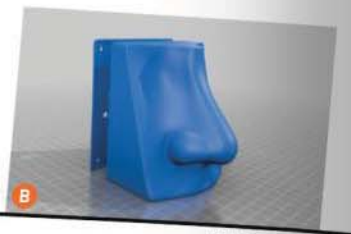
B

Photo by Wilhelm Gunkel on Unsplash

# “INTELLIGENCE AT THE EDGE” + INTERNET = ♡*

* a.k.a. AIoT

Contoso Virtual Scents Inc.

Search

Devices > Artificial Nose > **nose001**

Connect | Block | Attach to gate... | Rename | Edit template | Delete

# nose001

About | **Overview** | Settings | Raw data

Last data received: 4/30/2021, 4:45:56 PM | Status: Provisioned

## VOC (ppm), $NO_2$ (ppm), $C_2H_5OH$ (ppm), CO (ppm)

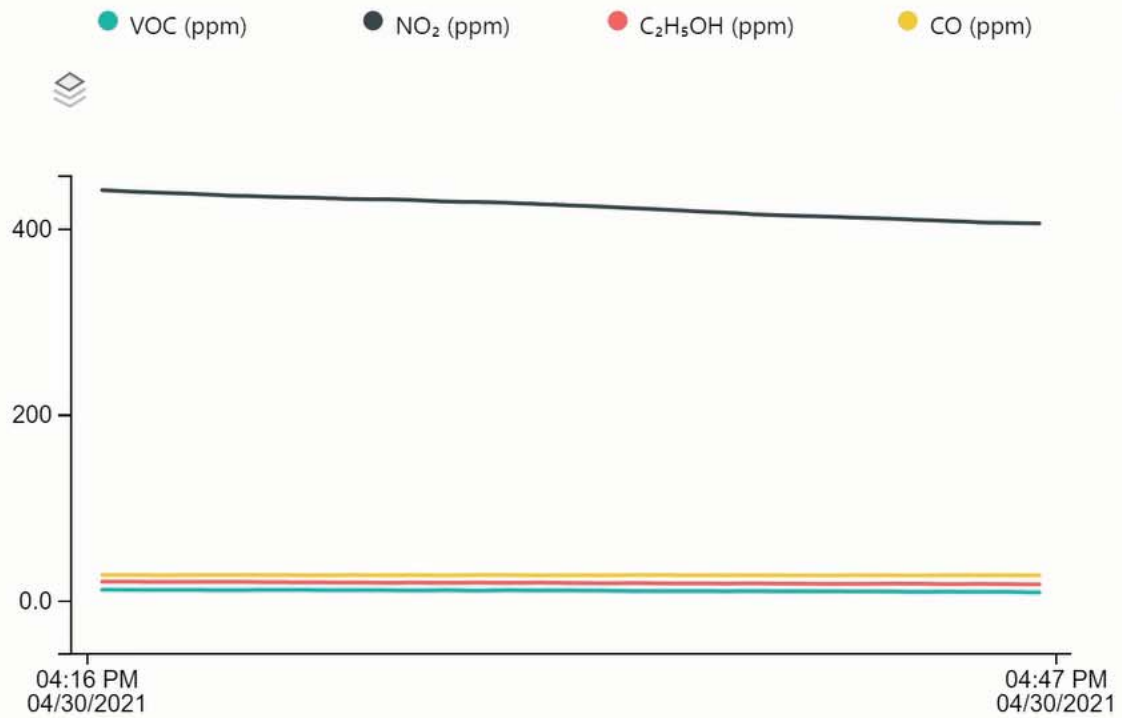


## VOC (ppm)

12.35

Average, Past 12 hours

## $NO_2$ (ppm)

442.91

Average, Past 12 hours

## $C_2H_5OH$ (ppm)

20.08

Average, Past 12 hours

## CO (ppm)

27.01

Average, Past 12 hours

# FROM CONNECTED THINGS TO CONNECTED ENVIRONMENTS*


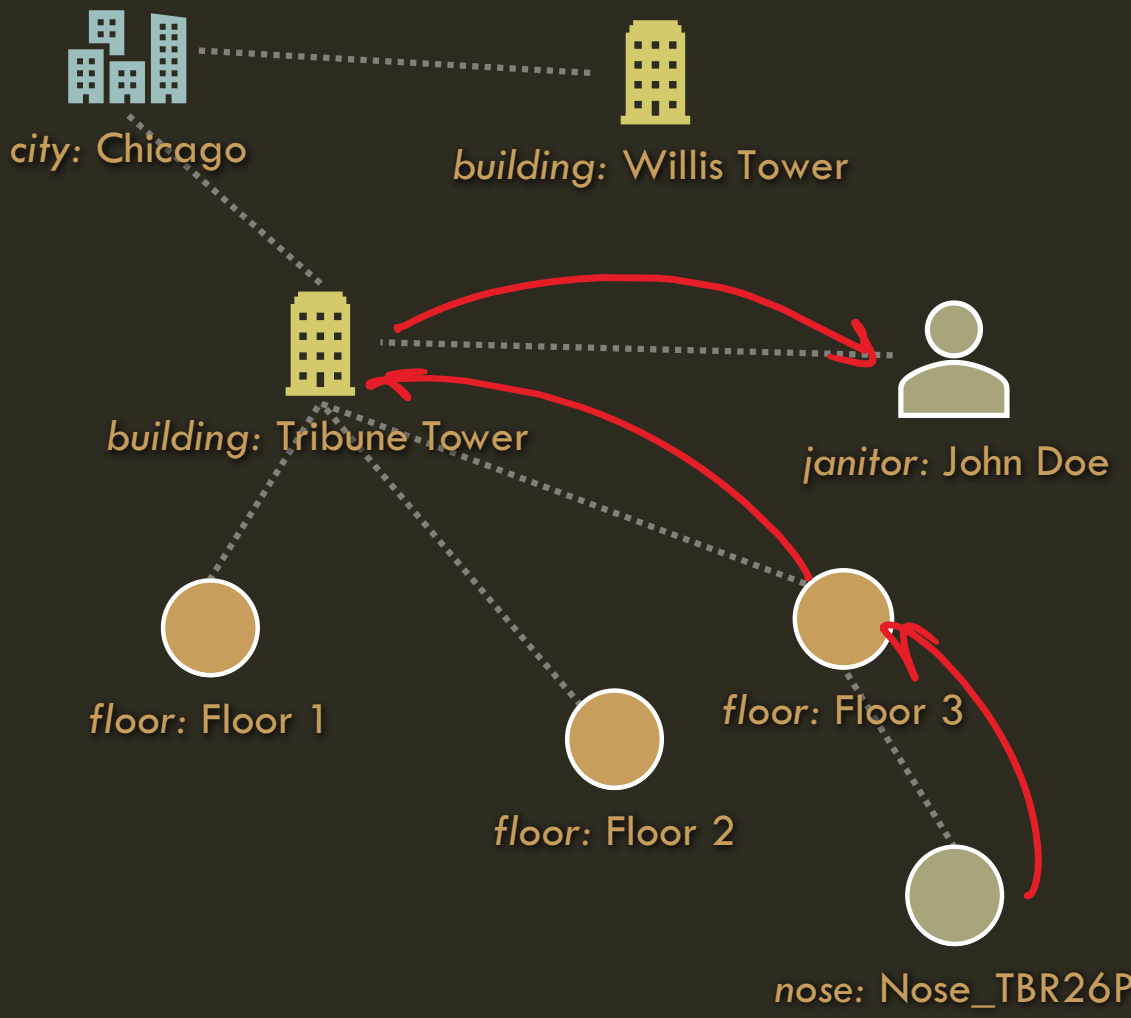


* a.k.a. Digital Twins


# OPEN SOURCE // OPEN HARDWARE

Bill of materials, source code, schemaitcs, …

@kartben

https://blog.benjamin-cabe.com

# THANKS!

TINY ML EMEA
Virtual
JUNE 7-10,2021

# Premier Sponsor

Intelligent Agent
Neuton

# Automated TinyML

Zero-code SaaS solution

**Create tiny models, ready for embedding, in just a few clicks!**

Compare the benchmarks of our compact models to those of TensorFlow and other leading neural network frameworks.

***Build Fast. Build Once. Never Compromise.***

TINY ML EMEA
Virtual
JUNE 7-10,2021

# Executive Sponsors

# Arm: The Software and Hardware Foundation for tinyML

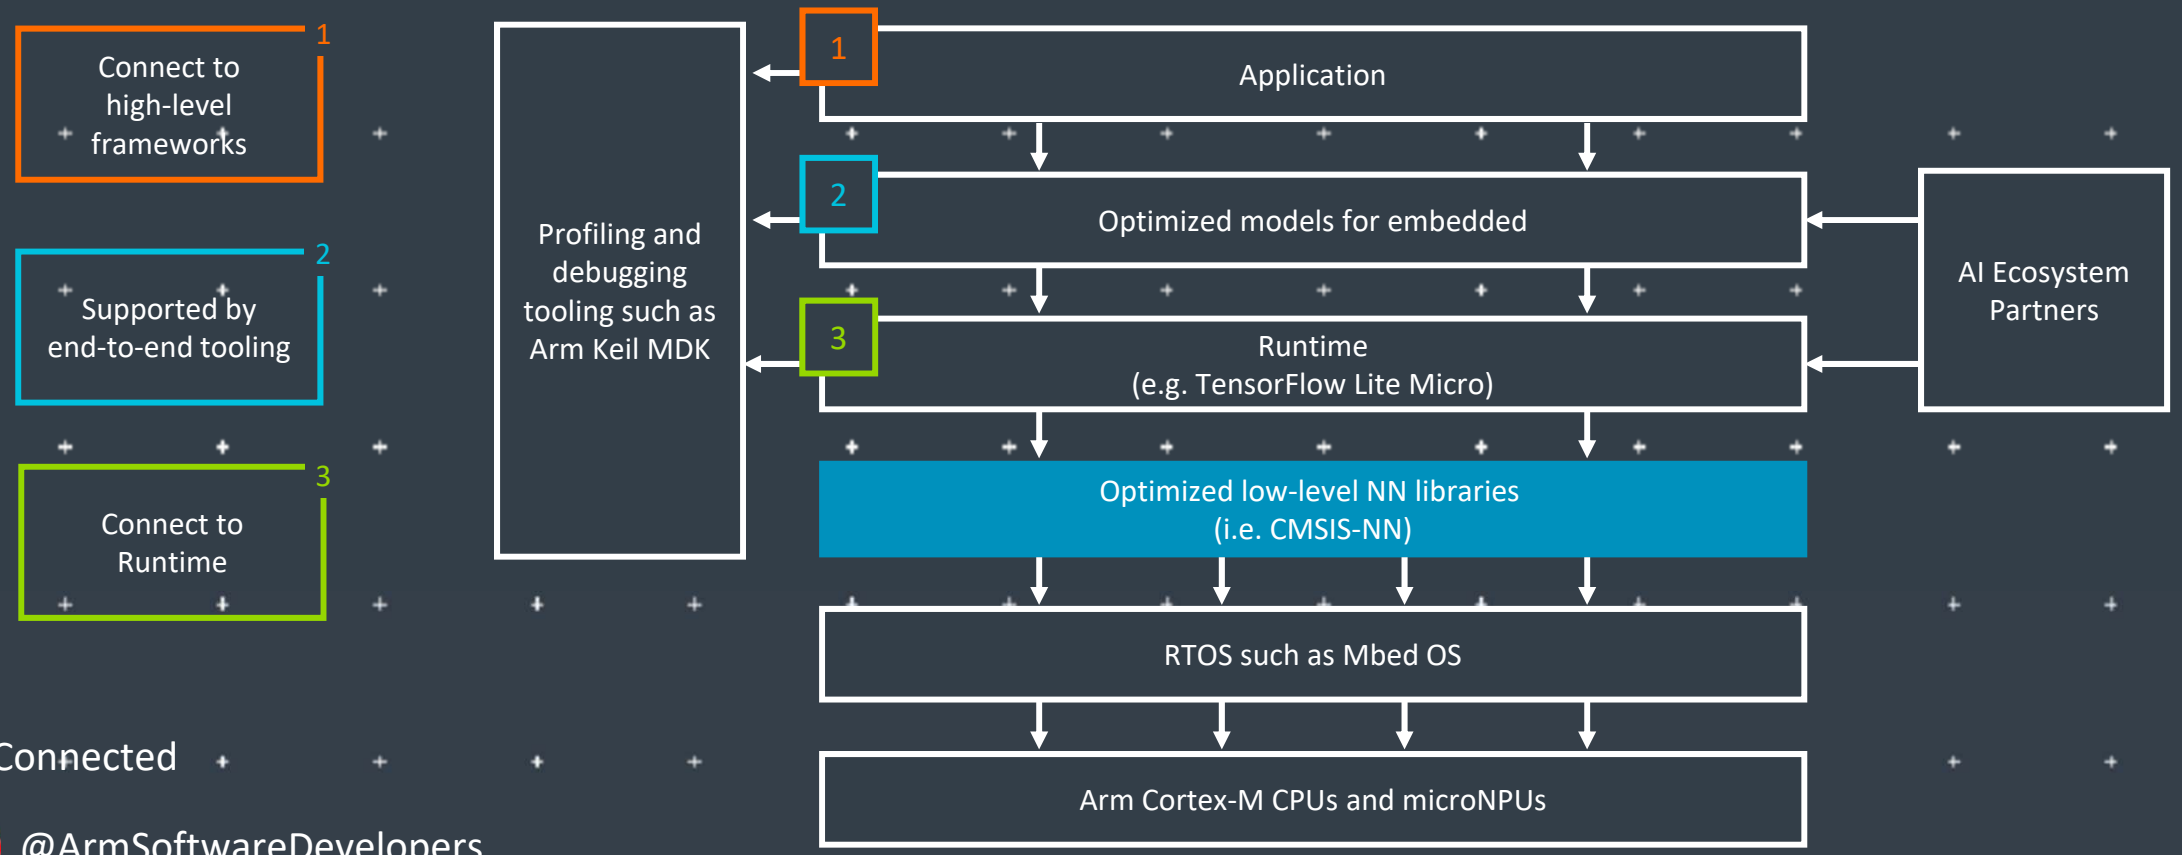


Stay Connected

@ArmSoftwareDevelopers

@ArmSoftwareDev

Resources: developer.arm.com/solutions/machine-learning-on-arm




arm

# TinyML for all developers

C++ library

Arduino library

WebAssembly

Dataset

Acquire valuable training data securely

Enrich data and train ML algorithms

Edge Device

Real sensors in real time

Open source SDK

Impulse

Embedded and edge compute deployment options

Test impulse with real-time device data flows

Test

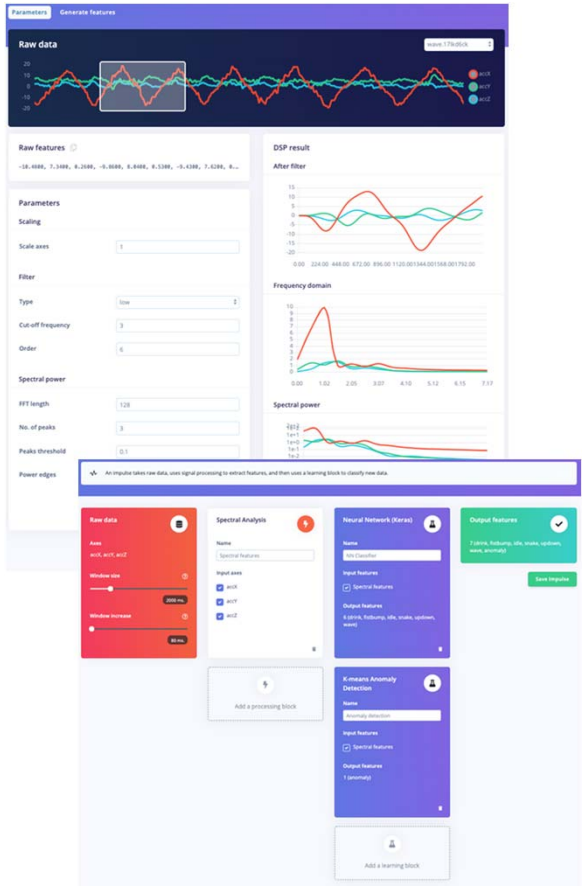

www.edgeimpulse.com

Qualcomm
AI research

# Advancing AI research to make efficient AI ubiquitous

### Power efficiency

Model design, compression, quantization, algorithms, efficient hardware, software tool

### Personalization

Continuous learning, contextual, always-on, privacy-preserved, distributed learning

### Efficient learning

Robust learning through minimal data, unsupervised learning, on-device learning

## A platform to scale AI across the industry

### Perception

Object detection, speech recognition, contextual fusion

### Reasoning

Scene understanding, language understanding, behavior prediction

### Action

Reinforcement learning for decision making

Edge cloud

Cloud

IoT/IIoT

Automotive

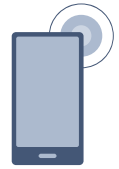

Mobile

# SYNTIANT

Syntiant Corp. is moving artificial intelligence and machine learning from the cloud to edge devices. Syntiant's chip solutions merge deep learning with semiconductor design to produce ultra-low-power, high performance, deep neural network processors. These network processors enable always-on applications in battery-powered devices, such as smartphones, smart speakers, earbuds, hearing aids, and laptops. Syntiant's Neural Decision Processors™ offer wake word, command word, and event detection in a chip for always-on voice and sensor applications.

Founded in 2017 and headquartered in Irvine, California, the company is backed by Amazon, Applied Materials, Atlantic Bridge Capital, Bosch, Intel Capital, Microsoft, Motorola, and others. Syntiant was recently named a CES® 2021 Best of Innovation Awards Honoree, shipped over 10M units worldwide, and unveiled the NDP120 part of the NDP10x family of inference engines for low-power applications.

www.syntiant.com

@Syntiantcorp

TINY ML EMEA
Virtual
JUNE 7-10,2021

# Platinum Sponsors


Part of your life. Part of tomorrow.

www.infineon.com

# Reality AI®

# Add Advanced Sensing to your Product with Edge AI / TinyML

https://reality.ai info@reality.ai @SensorAI Reality AI

## Pre-built Edge AI sensing modules, plus tools to build your own

### Reality AI solutions

Prebuilt sound recognition models for indoor and outdoor use cases

Solution for industrial anomaly detection

Pre-built automotive solution that lets cars "see with sound"

### Reality AI Tools® software

Build prototypes, then turn them into real products

Explain ML models and relate the function to the physics

Optimize the hardware, including sensor selection and placement

TINY ML EMEA
Virtual
JUNE 7-10,2021

# Gold Sponsors

LatentAI

# Adaptive AI for the Intelligent Edge

Latentai.com

SensiML™

# Build Smart IoT Sensor Devices From Data

SensiML pioneered TinyML software tools that auto generate AI code for the intelligent edge.

- End-to-end AI workflow
- Multi-user auto-labeling of time-series data
- Code transparency and customization at each step in the pipeline

We enable the creation of production-grade smart sensor devices.

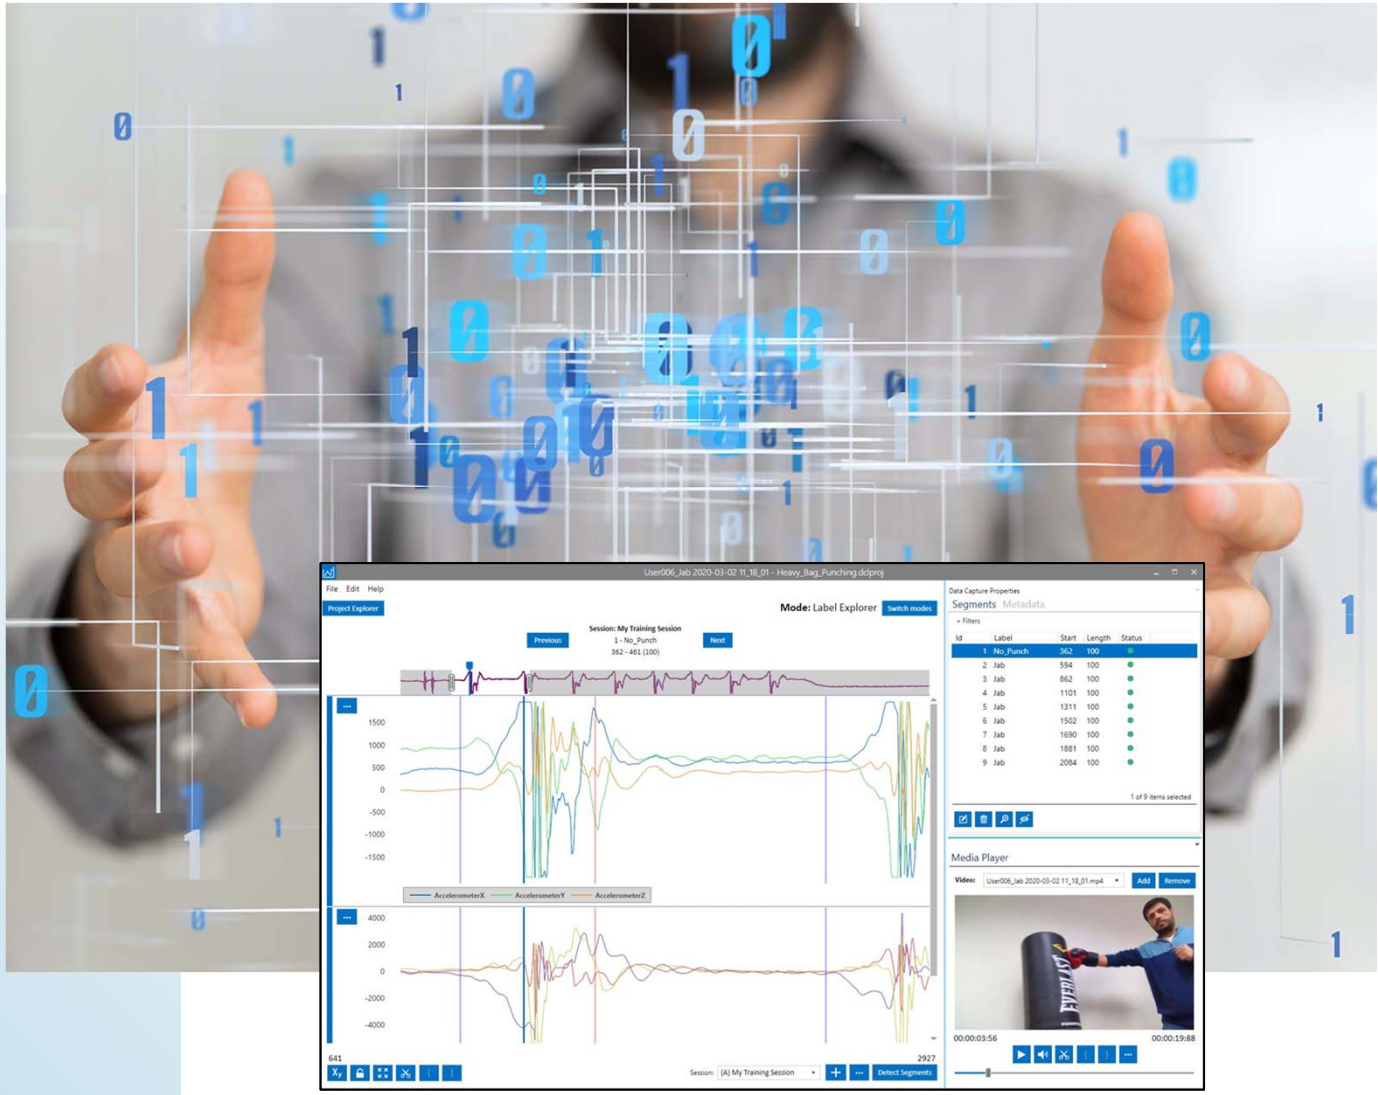


sensiml.com

TINY ML EMEA Virtual JUNE 7-10,2021

# Silver Sponsors




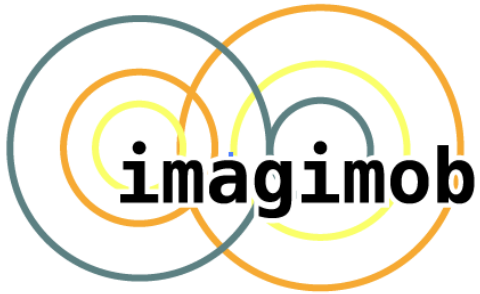




# Copyright Notice