Harvard

Yá'át'ééh 👋

EASI-22

Edge AI Summer Institute 2022

with Navajo Tech

IEEE COMPUTER SOCIETY

Our website!

tinyMLedu.org/EASI-22

home base for **all information**!

Workshop **Agenda**

Day 1	Introduction to AI and (Tiny)ML	Cloud ML
Day 2	Keyword Spotting for the Navajo Language	Mobile ML
Day 3	Bringing AI/ML from the Cloud to the Edge	Embedded ML

Workshop **Agenda**

Day 1	Introduction to AI and (Tiny)ML	Cloud ML
Day 2	Keyword Spotting for the Navajo Language	Mobile ML
Day 3	Bringing AI/ML from the Cloud to the Edge	Embedded ML

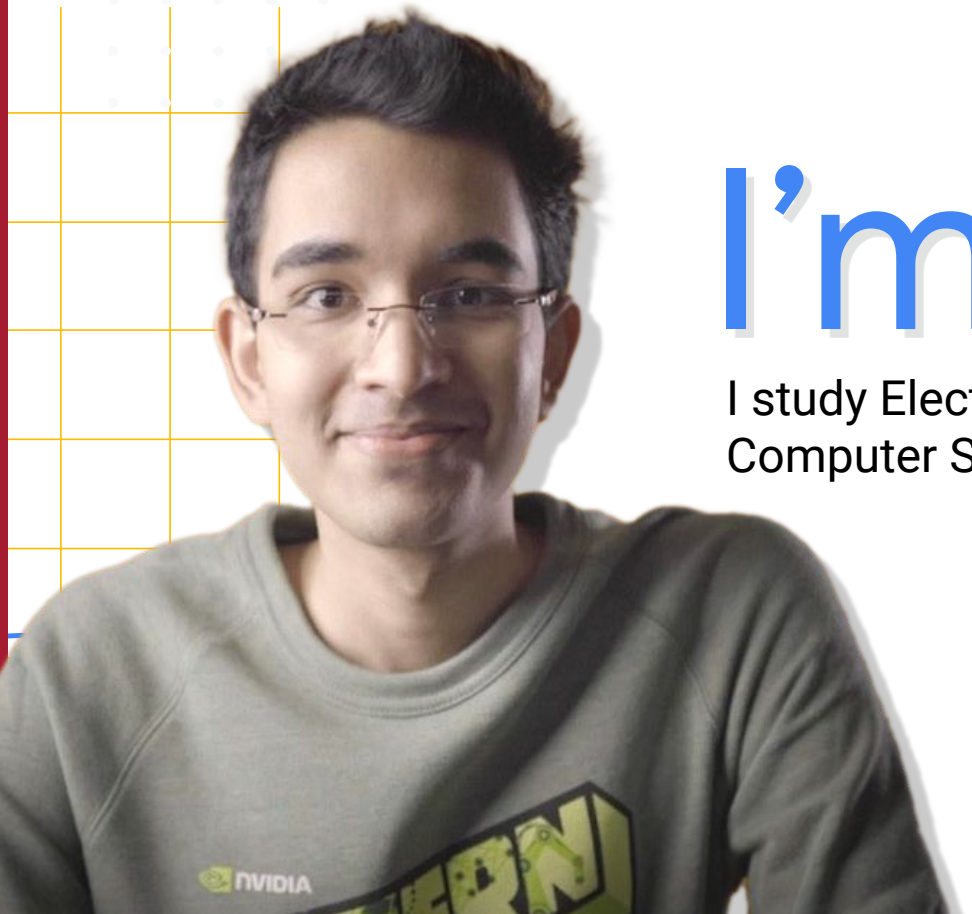

I'm Dhilan!

I study Electrical Engineering and Computer Science at **Harvard**.

Harvard

BARNA

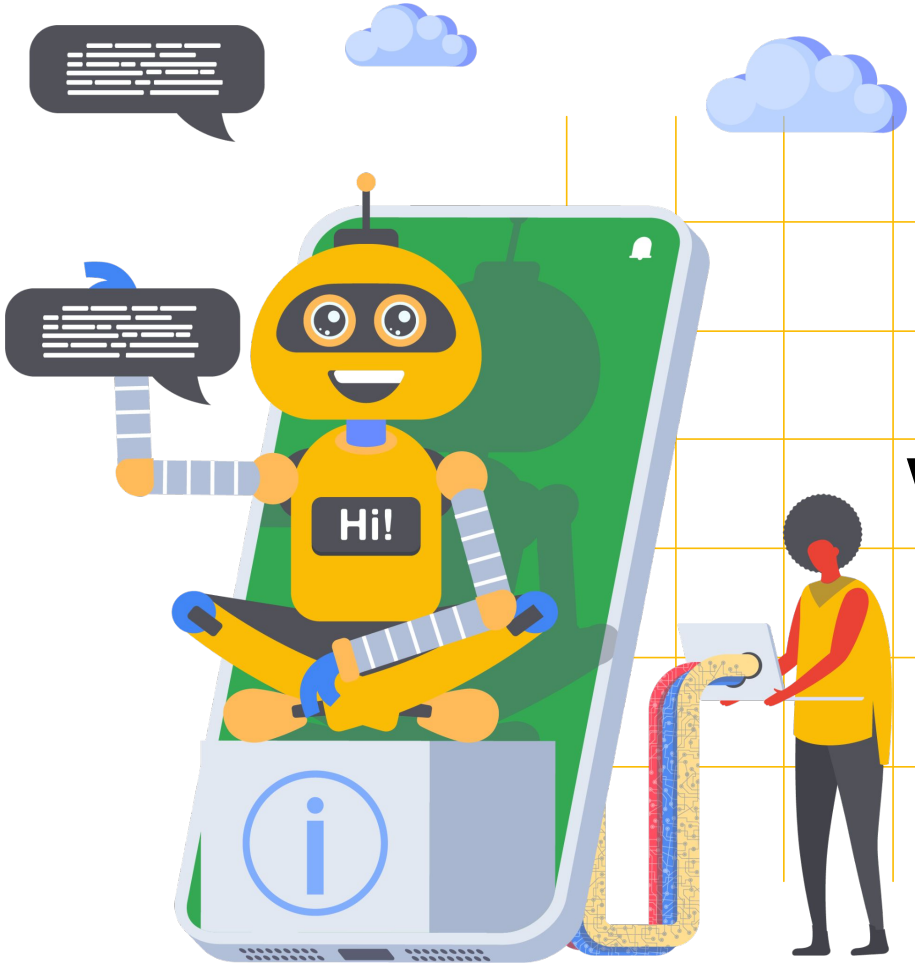

How ML works?

with Brian

Harvard

BARNA

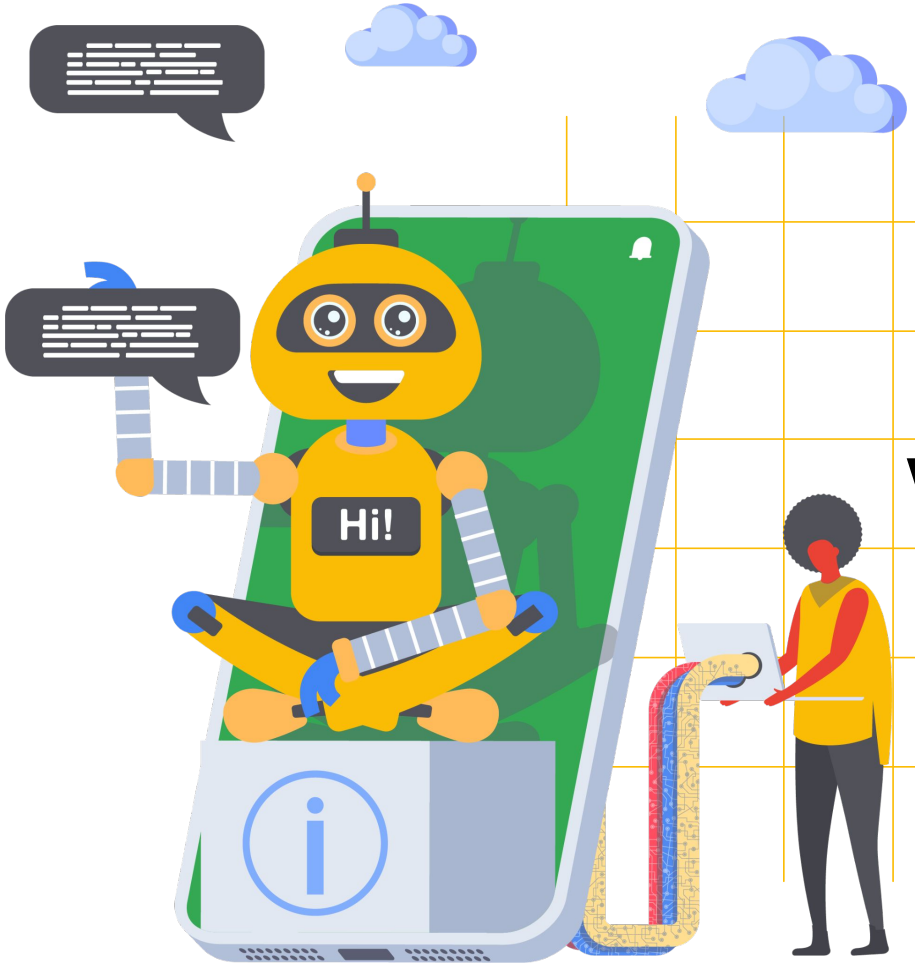

Where ML works?

with Dhilan

Last time: **Phone**

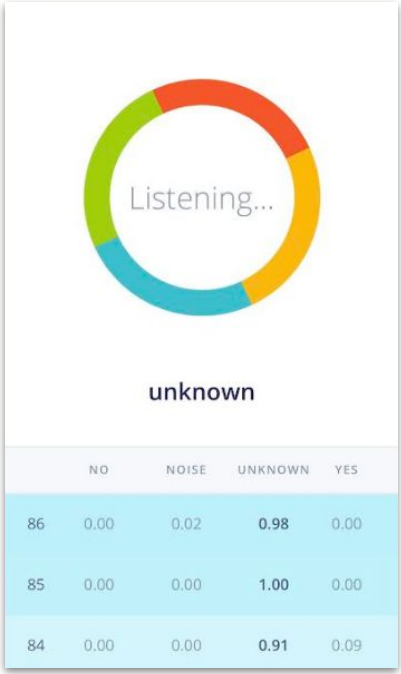

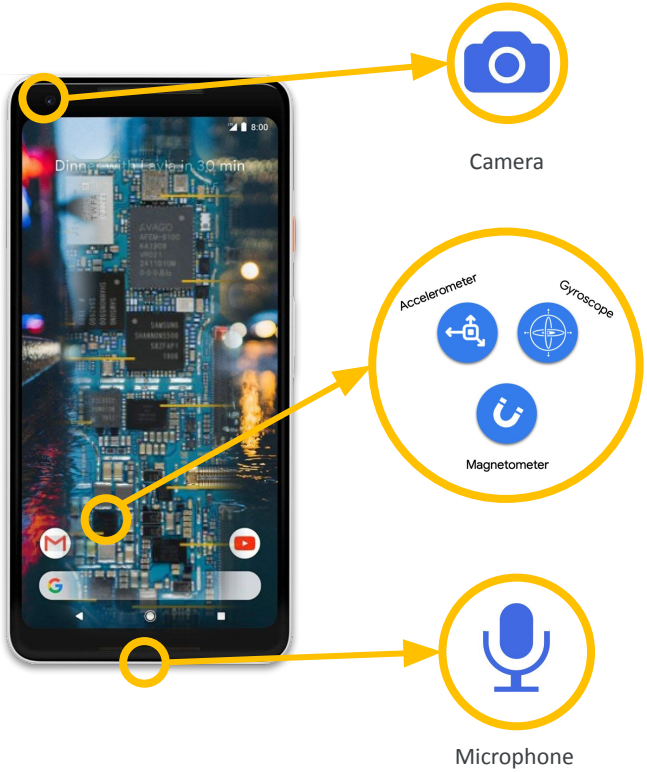

Today: **TinyML Kit**

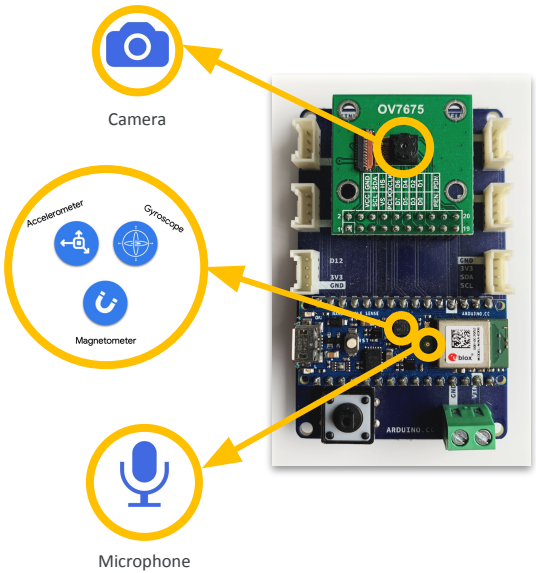

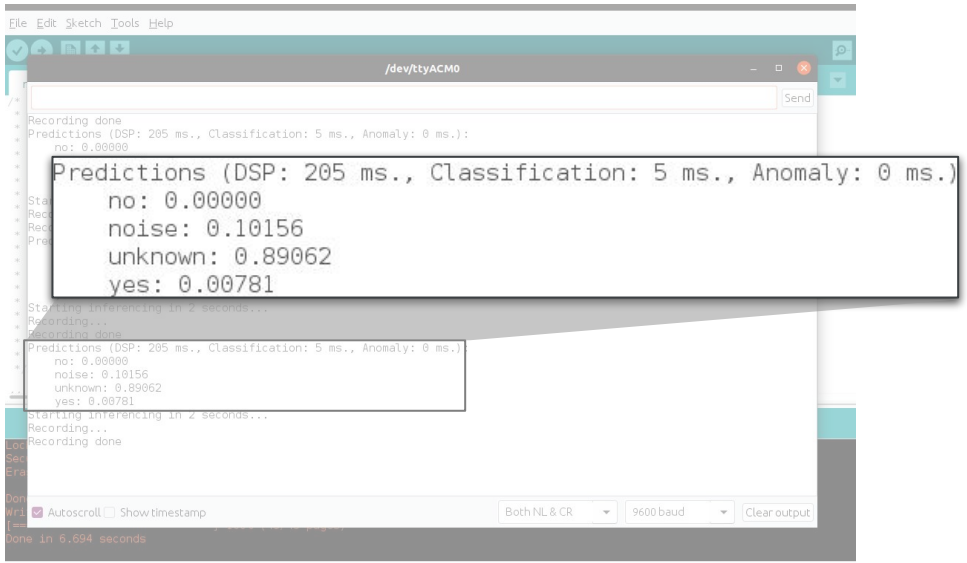

Today's Agenda

- Review + Why Tiny?
- Hardware Basics
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino
- Summary & Next Steps

Today's Agenda

Review + Why Tiny?

- Hardware Basics
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino
- Summary & Next Steps

Machine Learning **Workflow**

Machine Learning **Workflow**

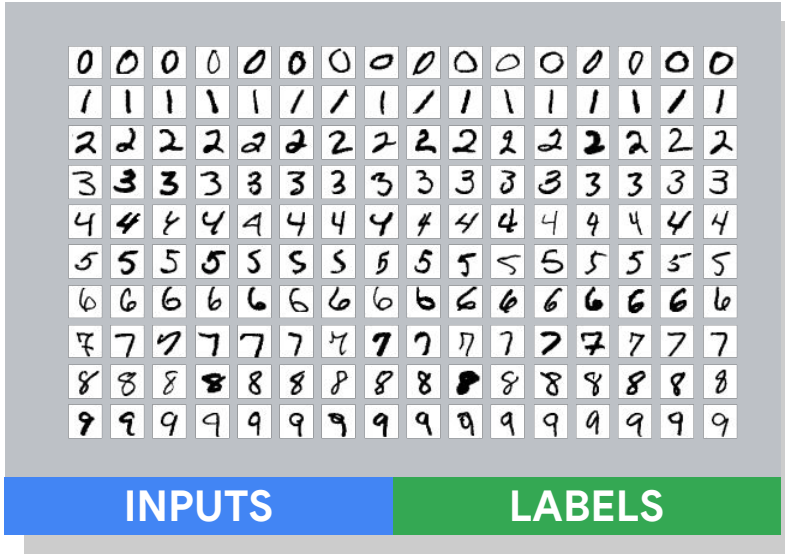

Collect & Transform Data → Design & Train a Model → Deploy Model → Make Inferences

Machine Learning **Workflow**

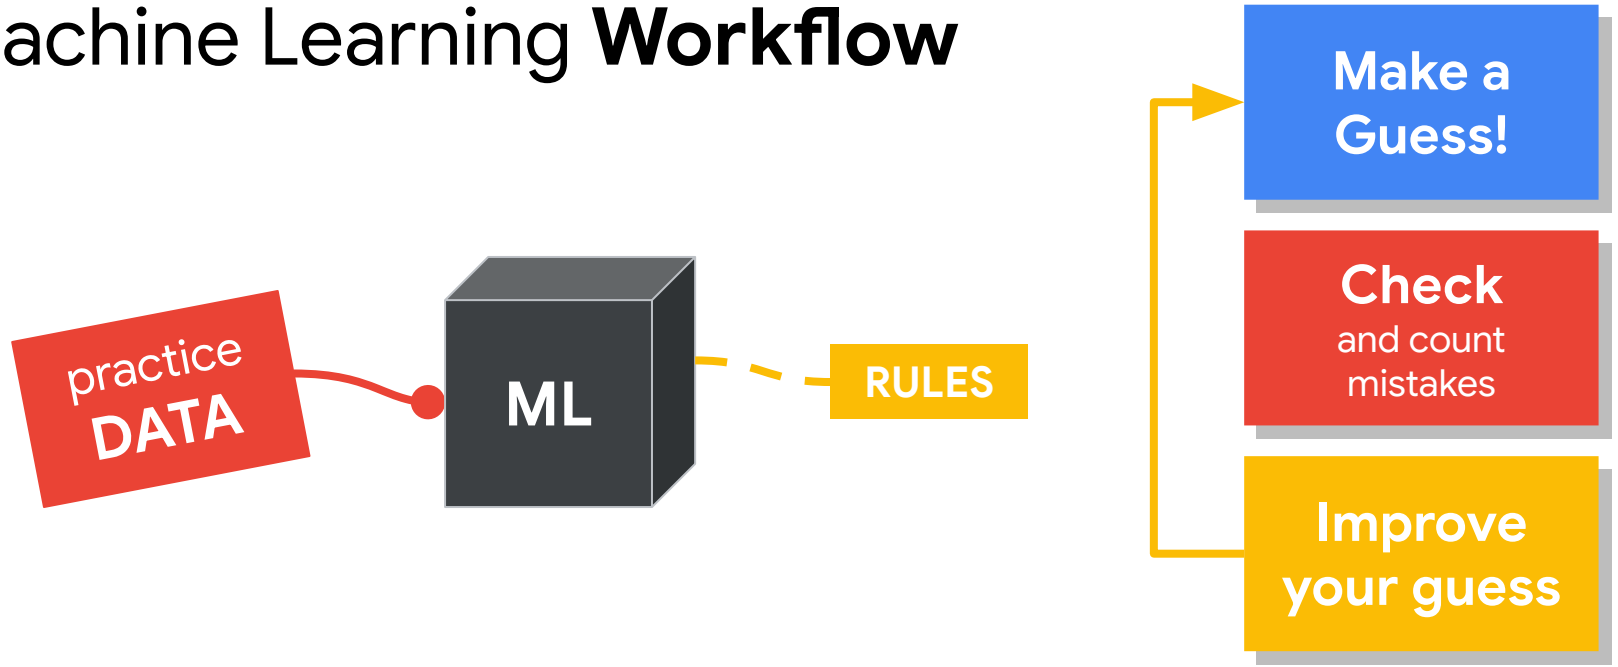

Machine Learning **Workflow**

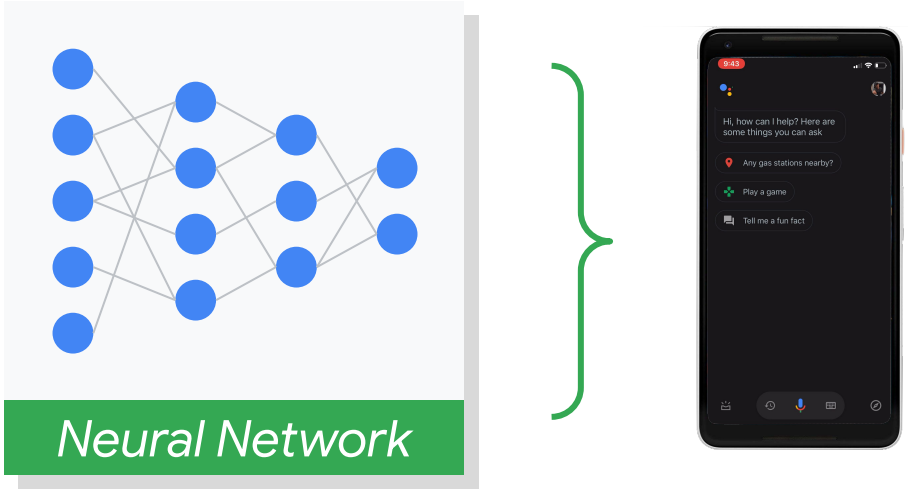

Collect & Transform Data → Design & Train a Model → **Deploy Model** → Make Inferences

Machine Learning **Workflow**

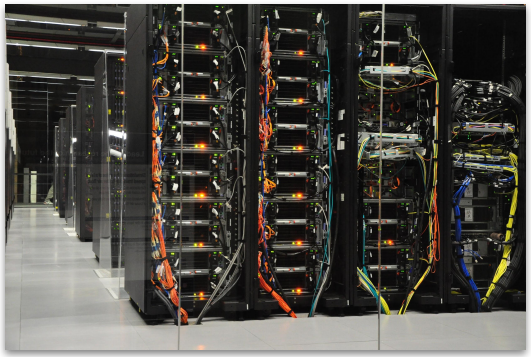

Collect & Transform Data → Design & Train a Model → **Deploy Model** → Make Inferences

Machine Learning **Workflow**

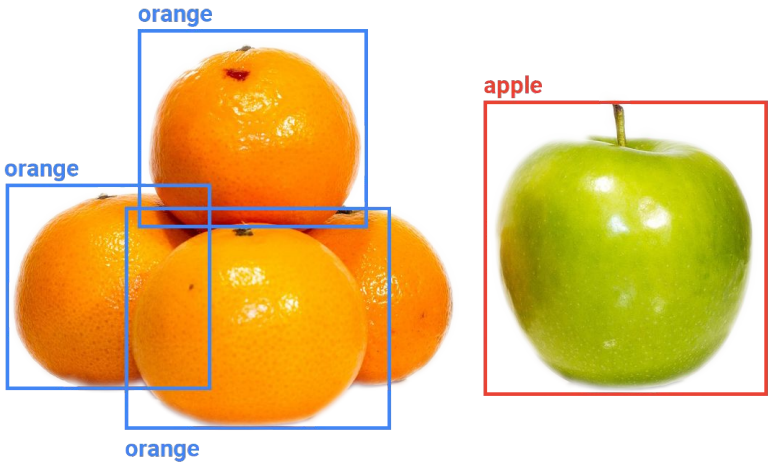

Collect & Transform Data → Design & Train a Model → Deploy Model → **Make Inferences**

Machine Learning **Workflow**

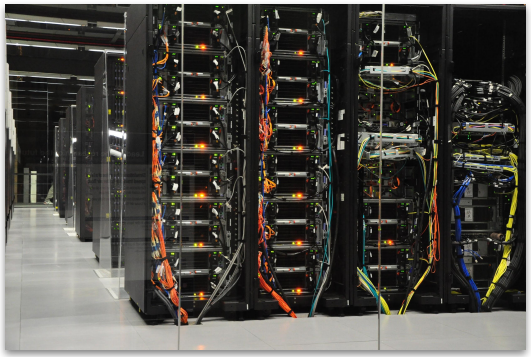

Collect & Transform Data → Design & Train a Model → **Deploy Model** → Make Inferences

Machine Learning **Workflow**

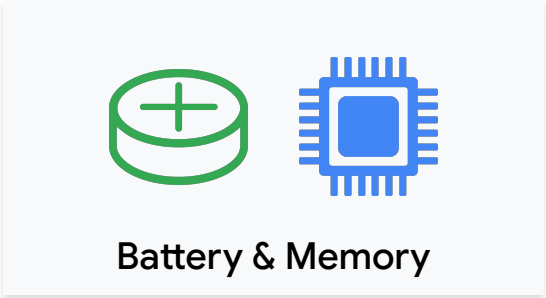

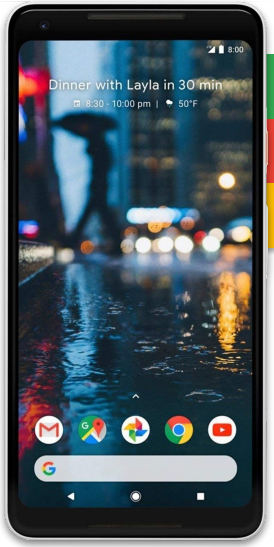

Less memory

Less compute power

Only focused on *inference*

Collect & Transform Data → Design & Train a Model → **Deploy Model** → Make Inferences

deploy

to your tiny **devices**!

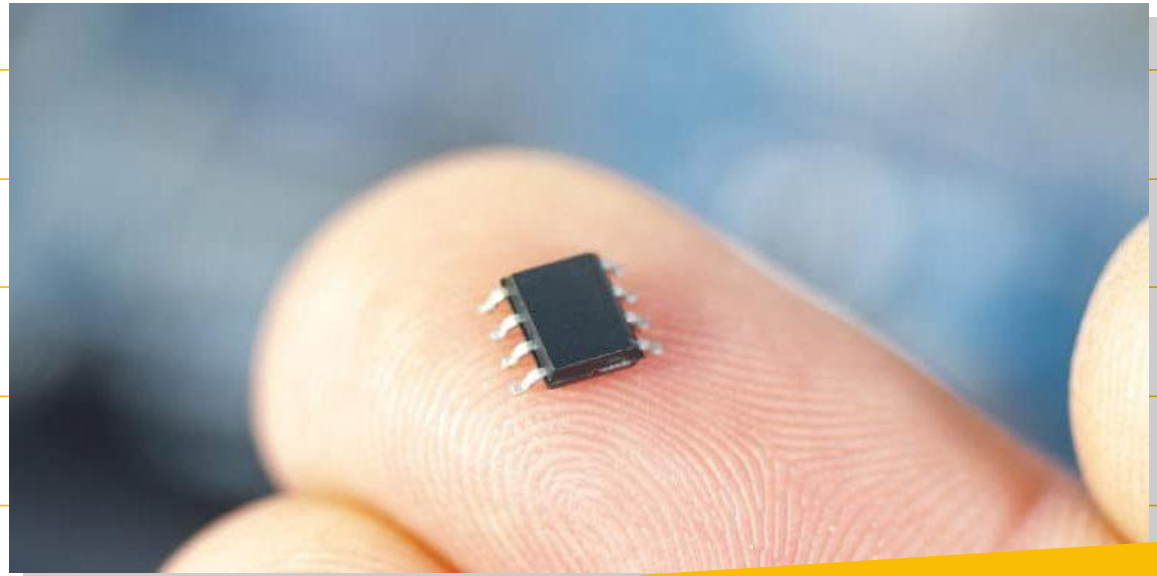

TinyML

Datacenter

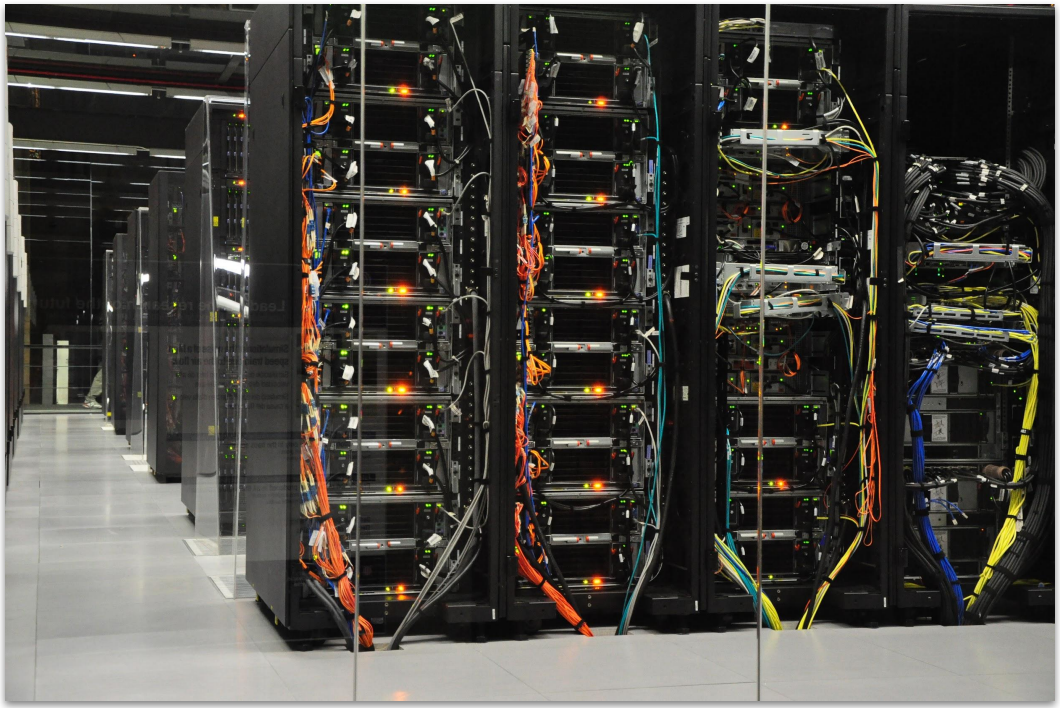

TPUs/GPUs

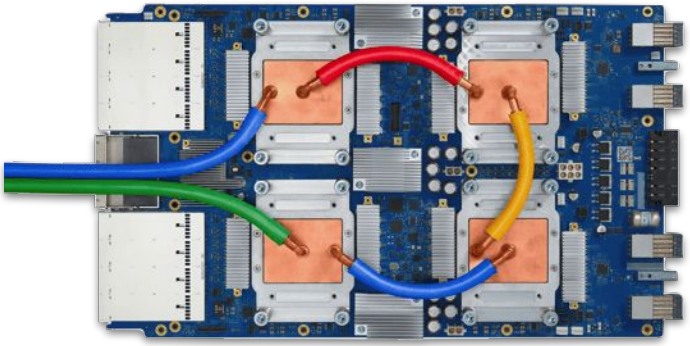

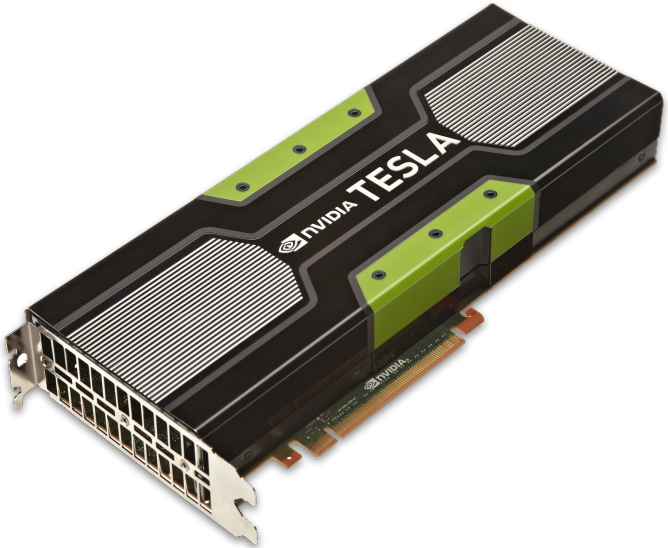

Bigger Is Not Always Better.

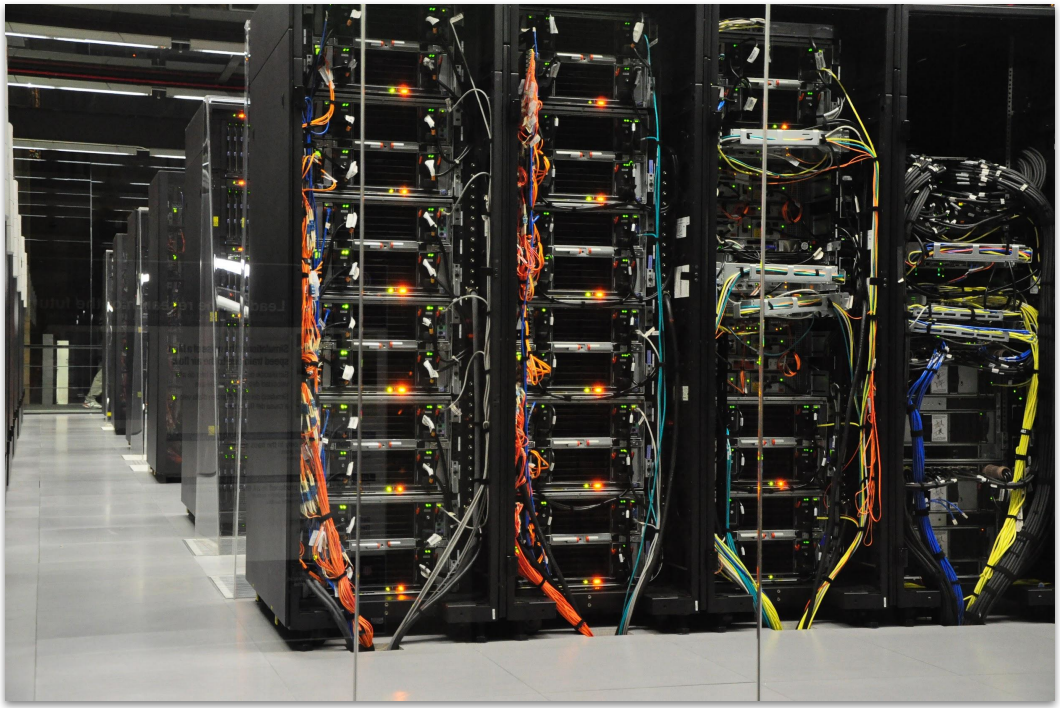

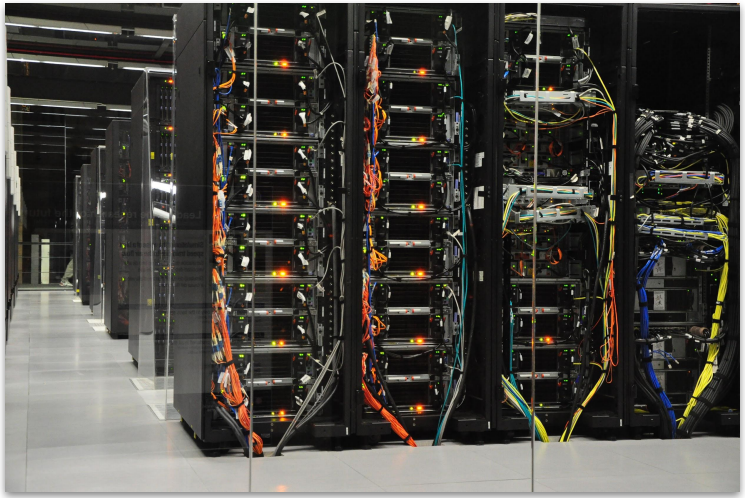

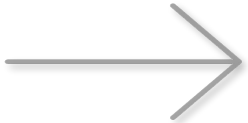

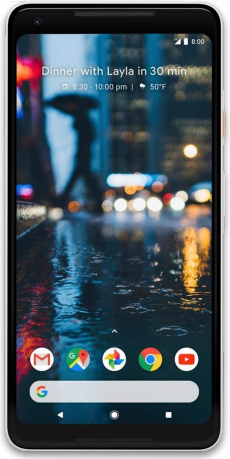

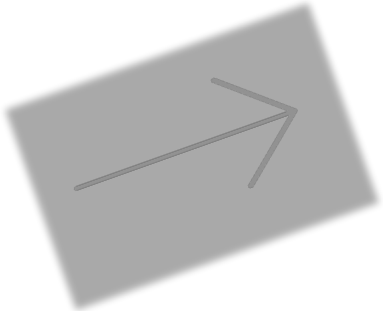

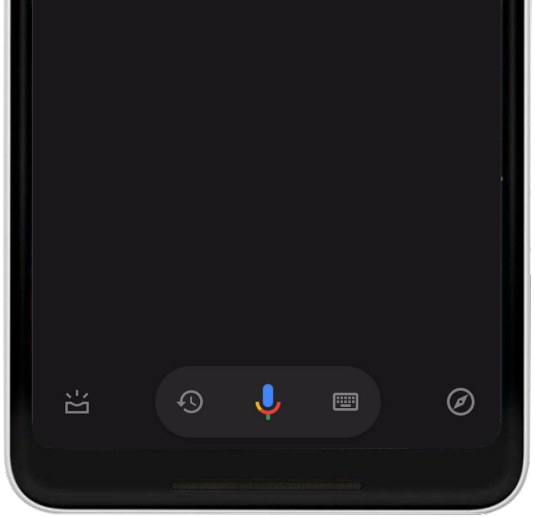

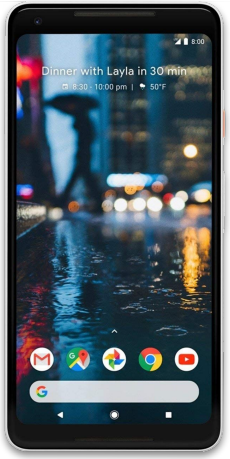
Dinner with Layla in 30 min

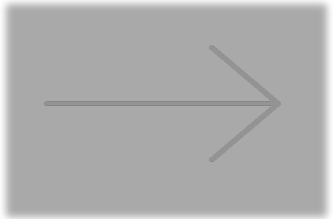

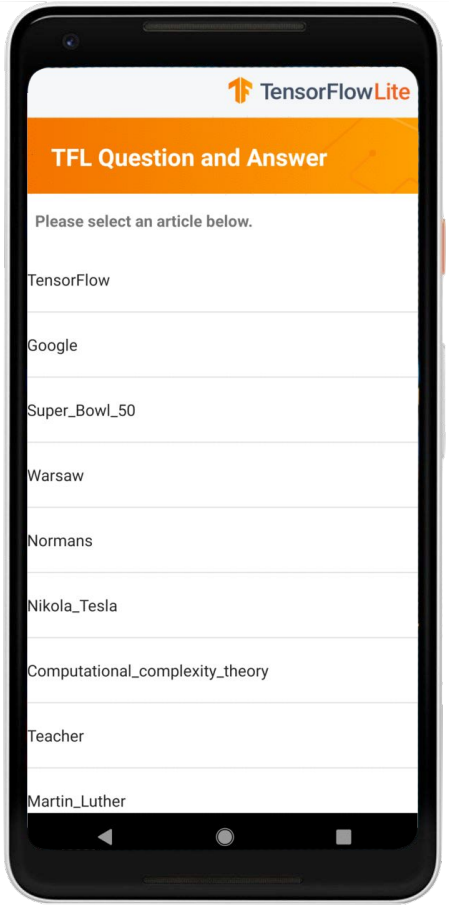
TensorFlowLite
TFL Question and Answer
Please select an article below.
TensorFlow
Google
Super_Bowl_50
Warsaw
Normans
Nikola_Tesla
Computational_complexity_theory
Teacher
Martin_Luther

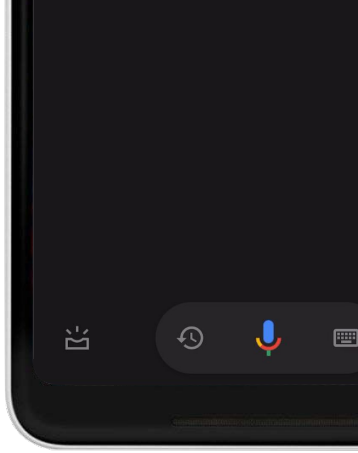

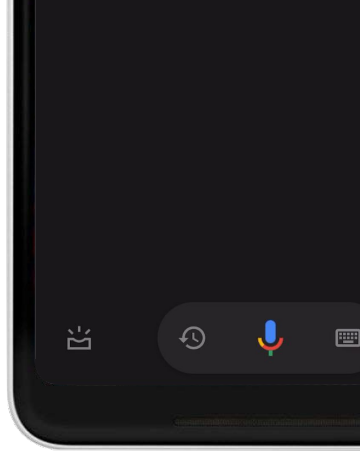

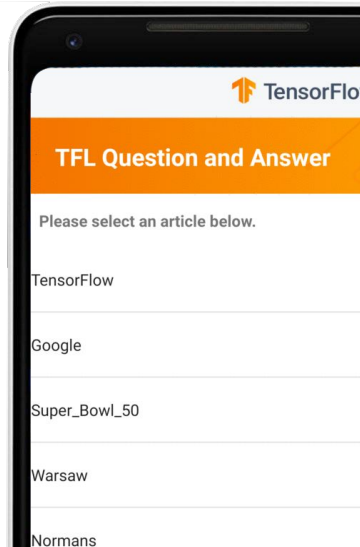
TensorFlo
TFL Question and Answer
Please select an article below.
TensorFlow
Google
Super_Bowl_50
Warsaw
Normans

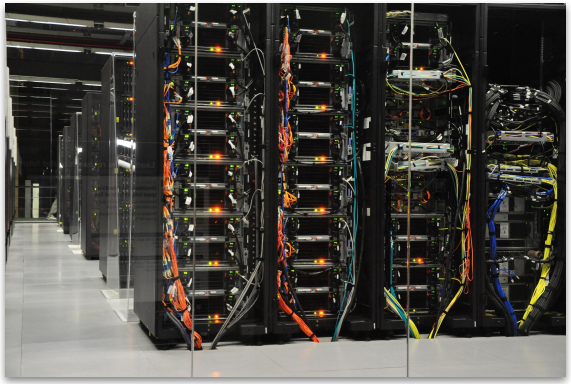

High power
High bandwidth
High latency

Low power
Low bandwidth
Low latency

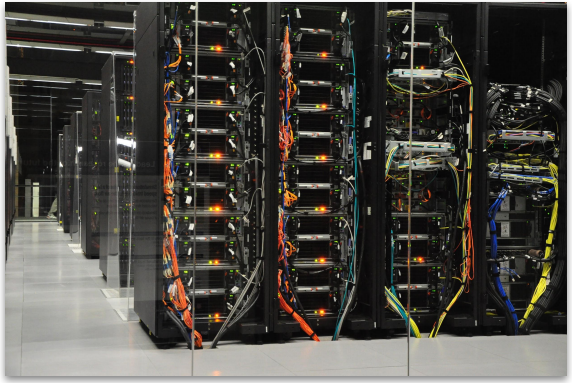

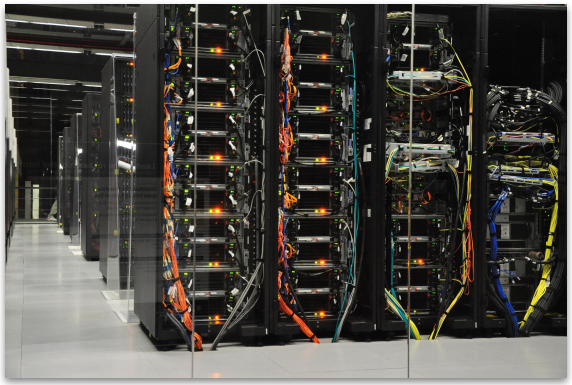

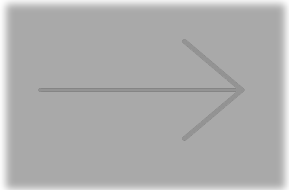

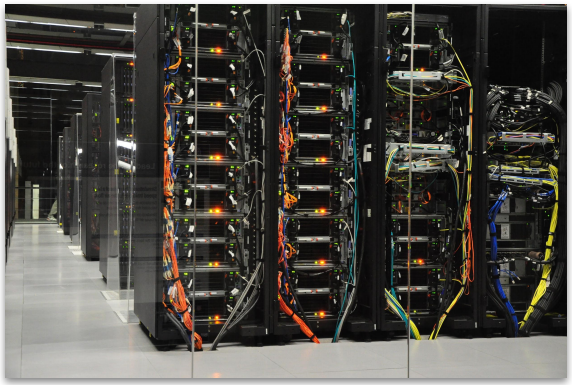

Endpoint Devices

Google Assistant

Endpoint Devices

Google Assistant

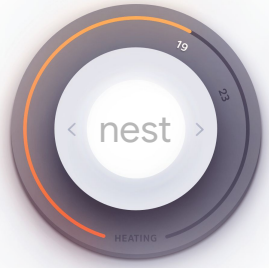

Endpoints Have **Sensors**, Tons of Sensors

Motion Sensors
Gyroscope, radar, magnetometer, accelerator

Acoustic Sensors
Ultrasonic, Microphones, Geophones, Vibrometers

Environmental Sensors
Temperature, Humidity, Pressure, IR, etc.

Touchscreen Sensors
Capacitive, IR

Image Sensors
Thermal, Image

Biometric Sensors
Fingerprint, Heart rate, etc.

Force Sensors
Pressure, Strain

Rotation Sensors
Encoders

Endpoints Have **Sensors**, Tons of Sensors

Motion Sensors
Gyroscope, radar, magnetometer, accelerator

Acoustic Sensors
Ultrasonic, Microphones, Geophones, Vibrometers

Environmental Sensors
Temperature, Humidity, Pressure, IR, etc.

Touchscreen Sensors
Capacitive, IR

Image Sensors
Thermal, Image

Biometric Sensors
Fingerprint, Heart rate, etc.

Force Sensors
Pressure, Strain

Rotation Sensors
Encoders

Biometric Sensors

Non-invasive Glucose Monitoring

Fingerprint + Photoplethysmography (**PPG**)

Source: Jacobs School of Engineering/UC San Diego

Endpoints Have **Sensors**, Tons of Sensors

Motion Sensors
Gyroscope, radar, magnetometer, accelerator

Acoustic Sensors
Ultrasonic, Microphones, Geophones, Vibrometers

Environmental Sensors
Temperature, Humidity, Pressure, IR, etc.

Touchscreen Sensors
Capacitive, IR

Image Sensors
Thermal, Image

Biometric Sensors
Fingerprint, Heart rate, etc.

Force Sensors
Pressure, Strain

Rotation Sensors
Encoders

Thinking **Big**

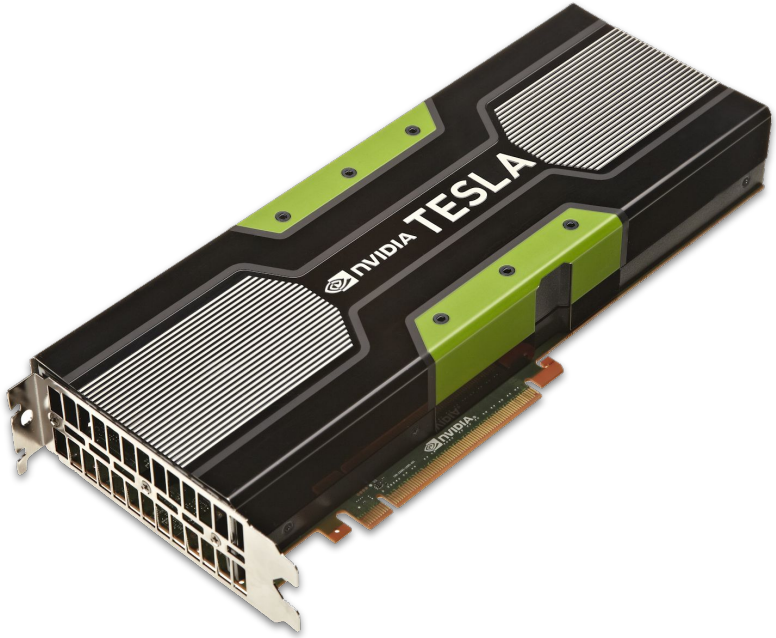

Thinking Big

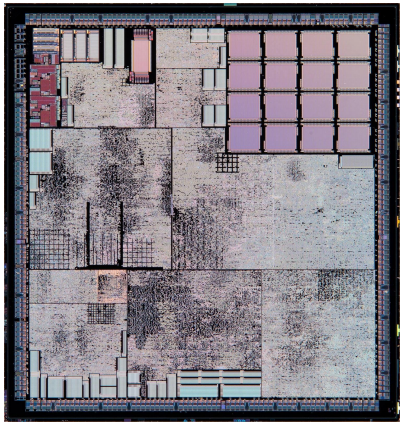

Thinking Big

BIG

GPU / CPU

561mm^2

Thinking **Small**

BIG

GPU / CPU

561mm^2

Thinking Small

BIG

GPU / CPU

561mm^2

Thinking Small

Thinking Tiny

BIG
GPU / CPU
561mm^2

SMALL
Mobile SoC
83mm^2

Thinking **Tiny**

BIG
GPU / CPU
561mm²

SMALL
Mobile SoC
83mm²

Thinking **Tiny**

BIG
GPU / CPU
561mm^2

SMALL
Mobile SoC
83mm^2

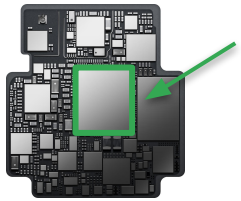

Thinking **Tiny**

We're just getting started.

Thinking **Record-breaking**

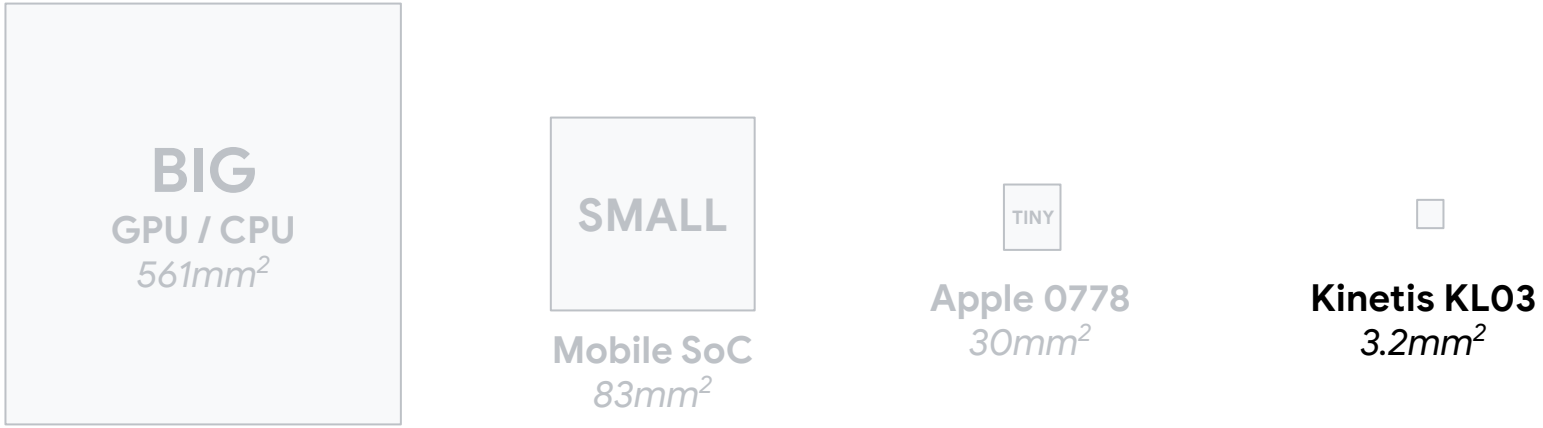

Thinking **Record-breaking**

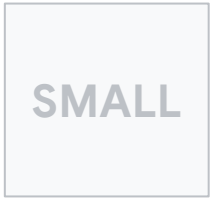

Mobile SoC
83mm²

Apple 0778
30mm²

world's smallest ARM-Powered MCU

48MHz, 32KB flash, 20-pin

Kinetis KL03
3.2mm²


2.0 mm
1.6 mm

250 Billion

today

Challenges

Latency & Bandwidth

Accuracy & Personalization

Security & Privacy

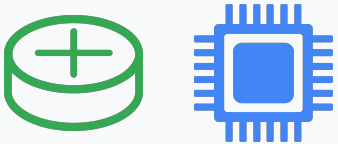

Battery & Memory

Source: Google

Source: Google

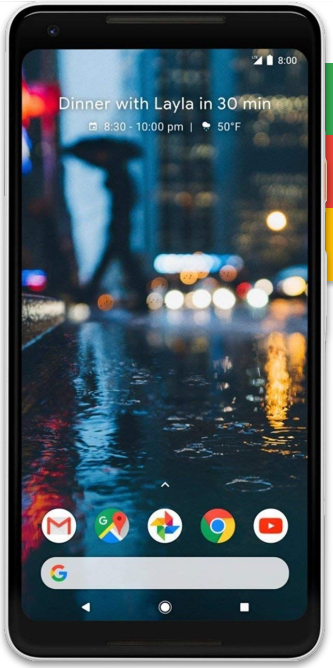

Less memory

Less compute power

Only focused on *inference*

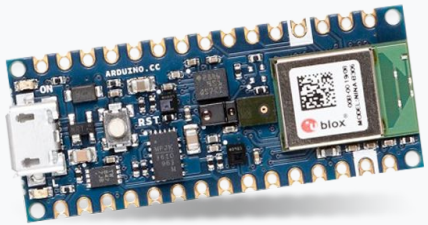

Even less memory

Even less compute power

Also, only focused on *inference*

What is Tiny Machine Learning (**TinyML**)?

TinyML

Fastest-growing field of **ML**

Algorithms, hardware, software

On-device sensor analytics

Low power consumption

Battery-operated

Always-on ML

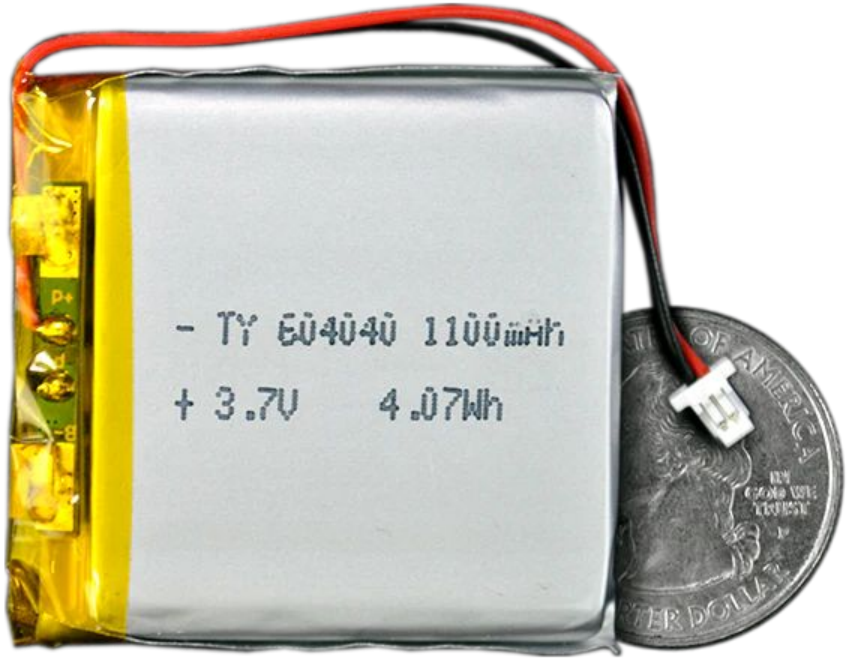

Measuring capacity by watt-hours lets us compare any type of battery

1
12 V car battery

≈

133
9 V batteries

≈

67
3.6 V smartphones

1 battery
× (12 V x 50 Ah)
600 Wh

133 batteries
× (9 V x .5 Ah)
600 Wh

67 smartphones
× (3.6 V x 2.5 Ah)
600 Wh

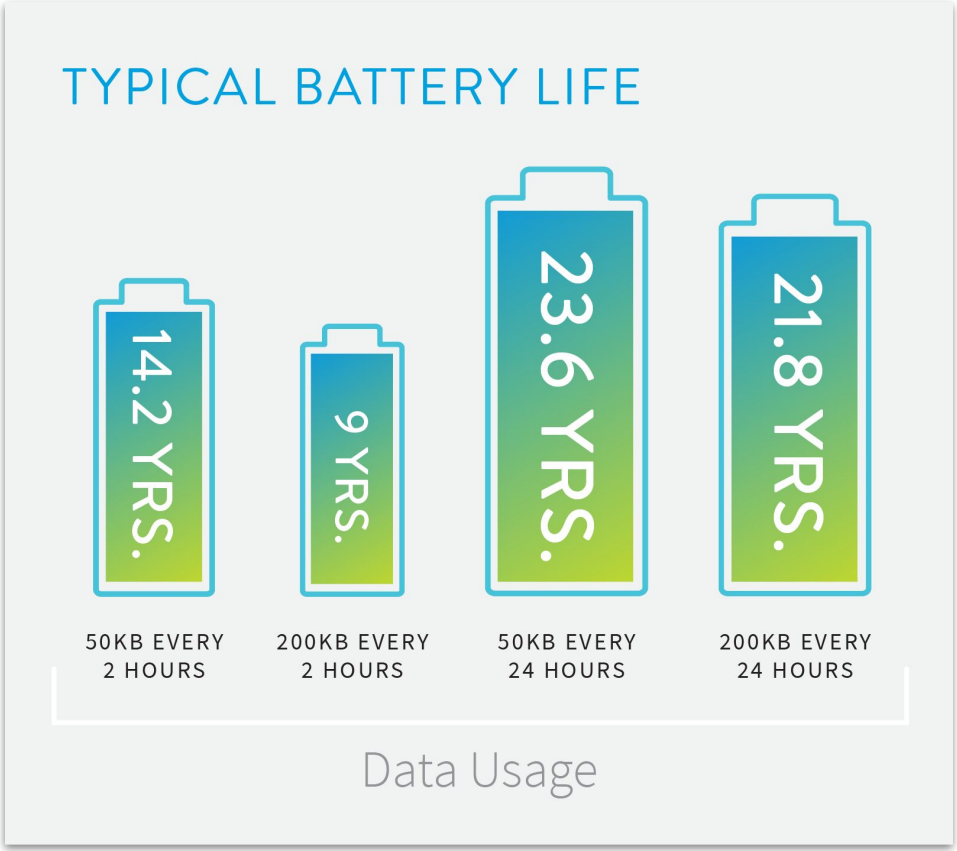

transmit less data!

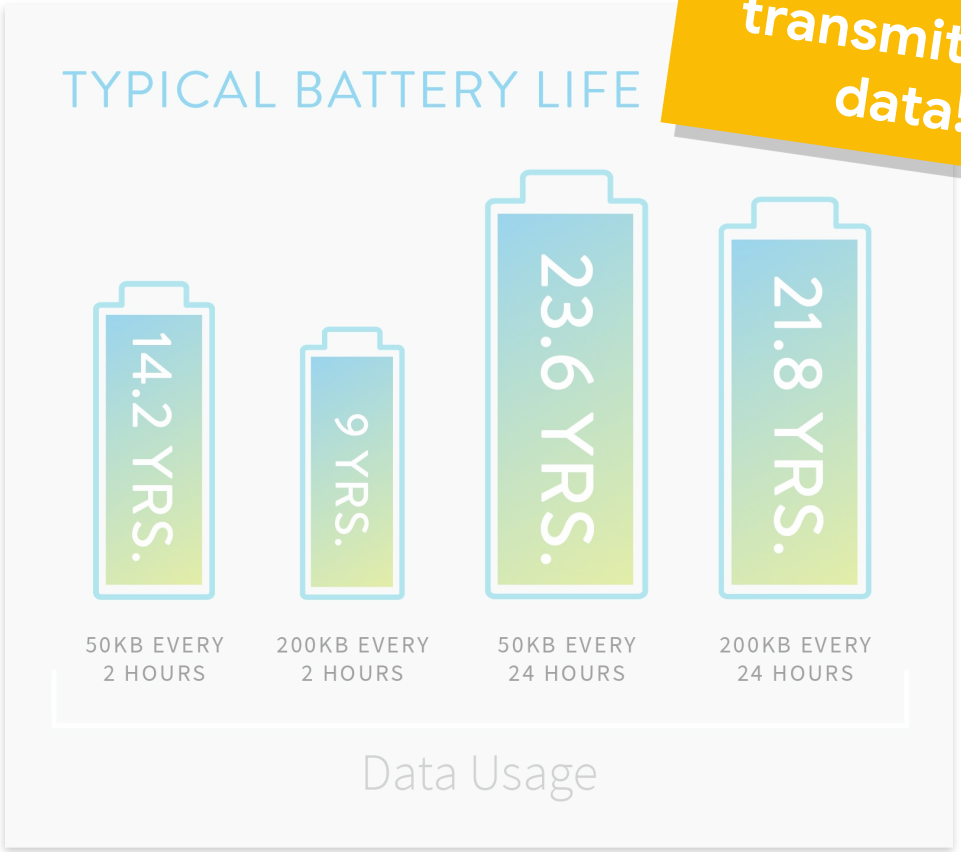

Today's Agenda

- Review + Why Tiny?
- Hardware Basics
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino
- Summary & Next Steps

Today's **Agenda**

- Review + Why Tiny?
- **Hardware Basics**
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino
- Summary & Next Steps

language of computers

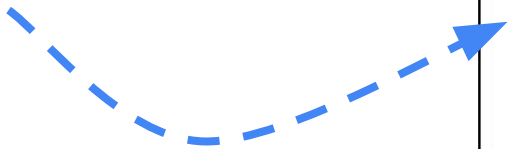

Binary	Hex	Decimal
0000	0	0
0001	1	1
0010	2	2
0011	3	3
0100	4	4
0101	5	5
0110	6	6
0111	7	7
1000	8	8
1001	9	9
1010	A	10
1011	B	11
1100	C	12
1101	D	13
1110	E	14
1111	F	15

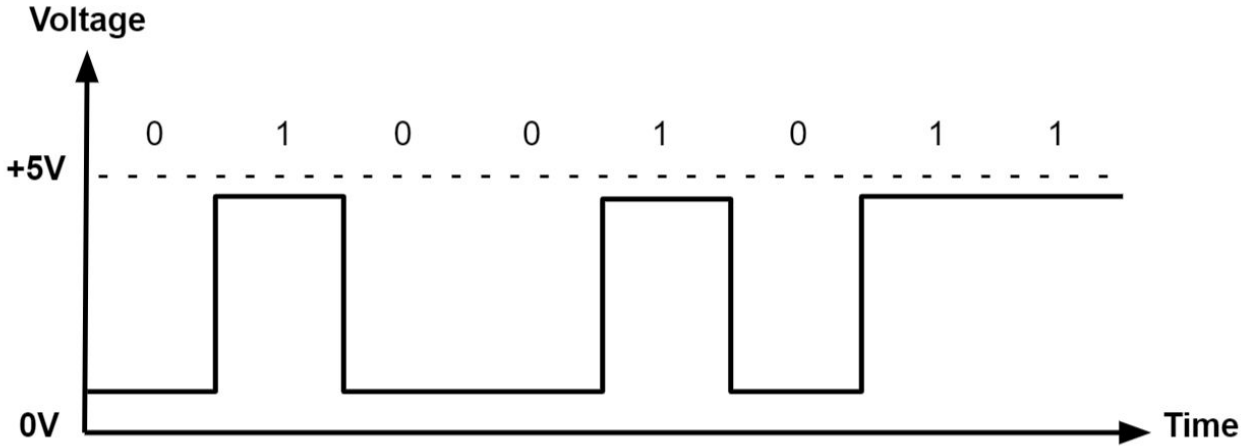

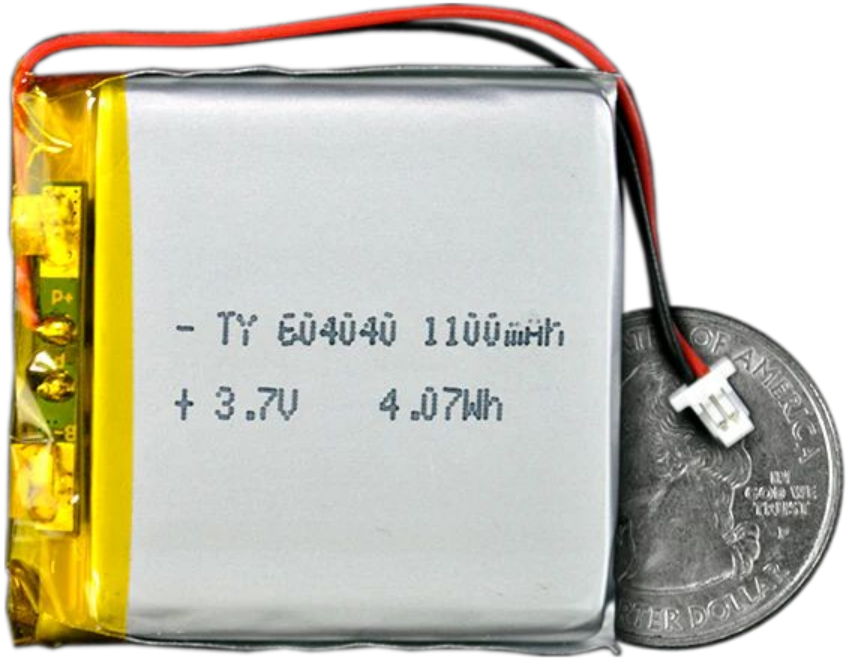

Microcontroller

CPU	Read-Only Memory (ROM)	Read-Write Memory
Timer	I/O Port	Serial Interface

Microcontroller: a **complete package**

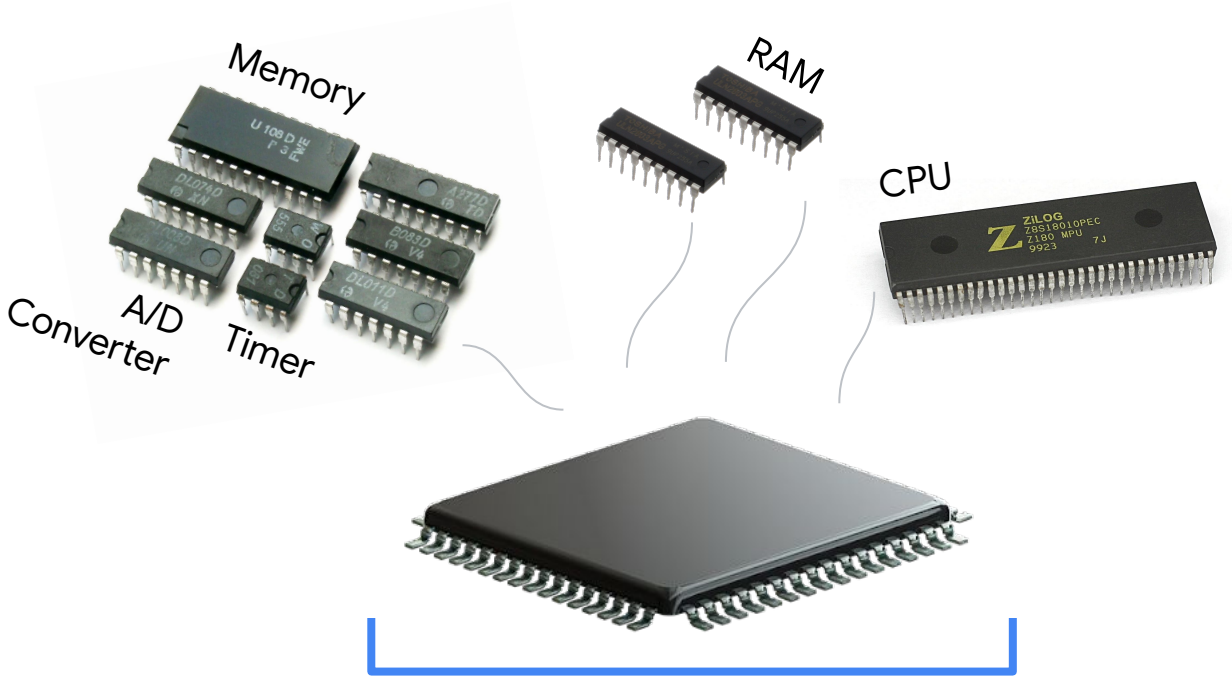

Embedded Systems

	Board	MCU / ASIC	Clock	Memory	Sensors	Radio
	Himax WE-I Plus EVB	HX6537-A 32-bit EM9D DSP	400 MHz	2MB flash 2MB RAM	Accelerometer, Mic, Camera	None
	Arduino Nano 33 BLE Sense	32-bit nRF52840	64 MHz	1MB flash 256kB RAM	Mic, IMU, Temp, Humidity, Gesture, Pressure, Proximity, Brightness, Color	BLE
	SparkFun Edge 2	32-bit ArtemisV1	48 MHz	1MB flash 384kB RAM	Accelerometer, Mic, Camera	BLE
	Espressif EYE	32-bit ESP32-D0WD	240 MHz	4MB flash 520kB RAM	Mic, Camera	WiFi, BLE

Embedded Systems

	Board	MCU / ASIC	Clock	Memory	Sensors	Radio
	Himax WE-I Plus EVB	HX6537-A 32-bit EM9D DSP	400 MHz	2MB flash 2MB RAM	Accelerometer, Mic, Camera	None
	Arduino Nano 33 BLE Sense	32-bit nRF52840	64 MHz	1MB flash 256kB RAM	Mic, IMU, Temp, Humidity, Gesture, Pressure, Proximity, Brightness, Color	BLE
	SparkFun Edge 2	32-bit ArtemisV1	48 MHz	1MB flash 384kB RAM	Accelerometer, Mic, Camera	BLE
	Espressif EYE	32-bit ESP32-D0WD	240 MHz	4MB flash 520kB RAM	Mic, Camera	WiFi, BLE

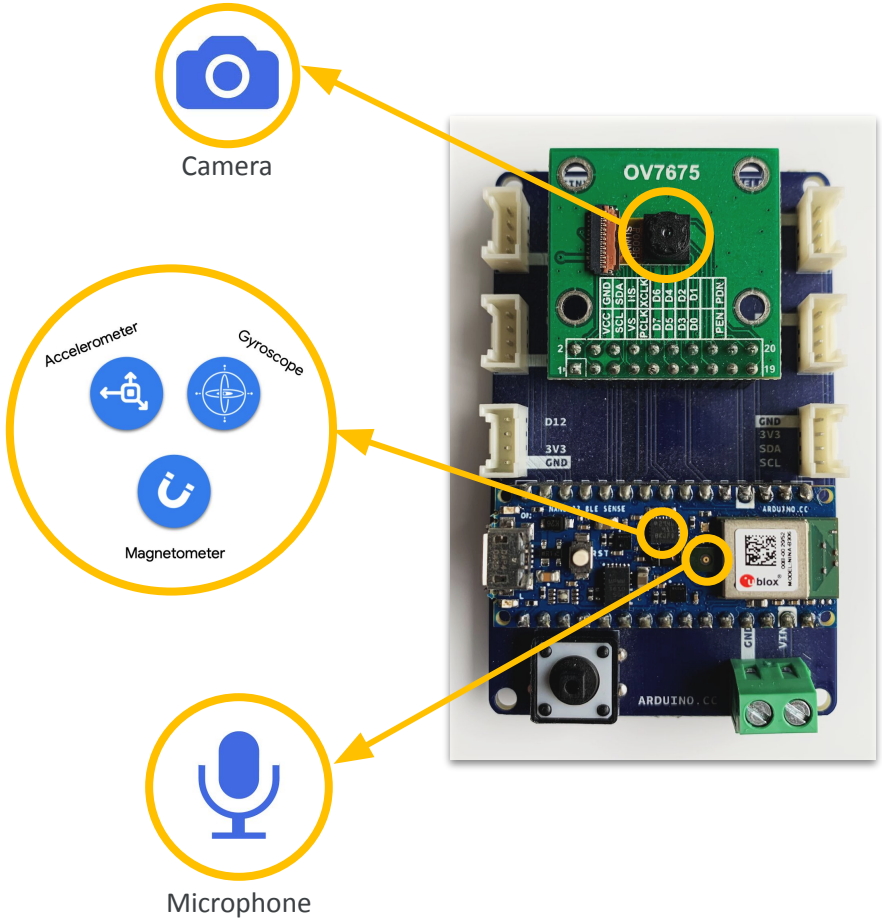

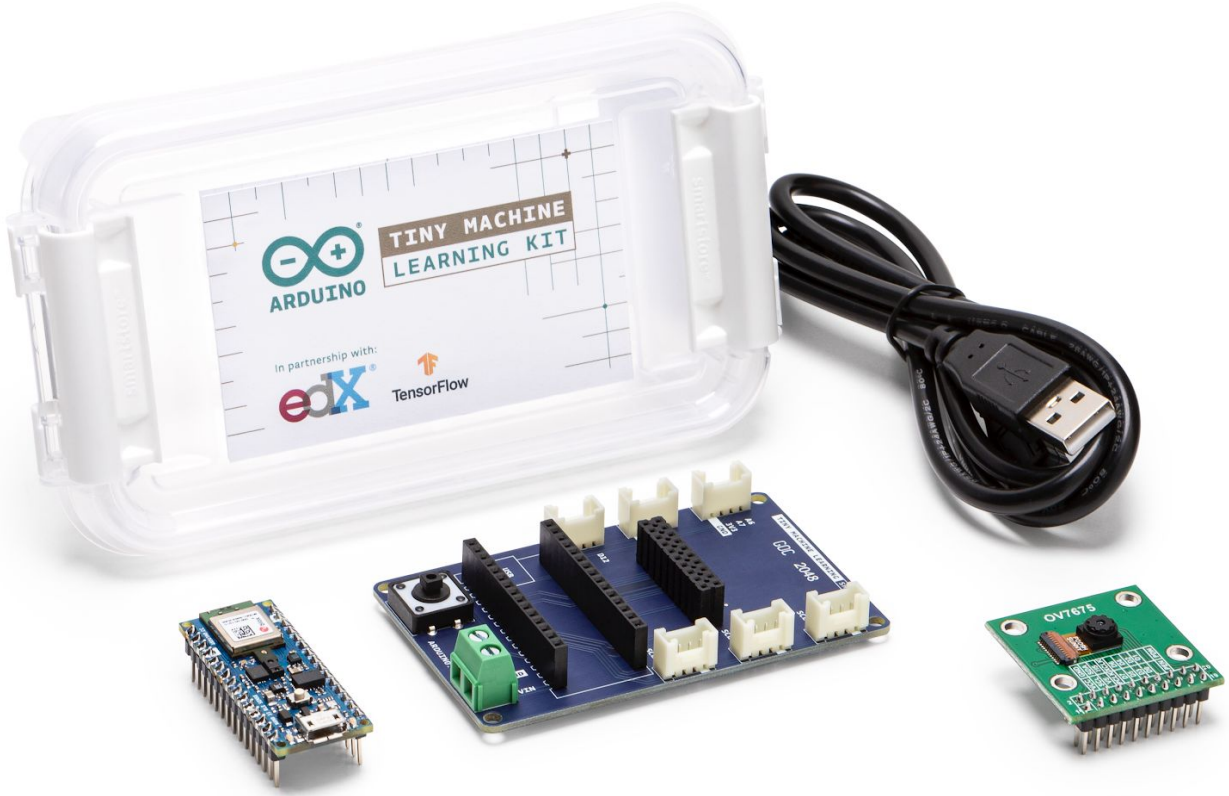

Nano 33 BLE Sense (+ USB cable)

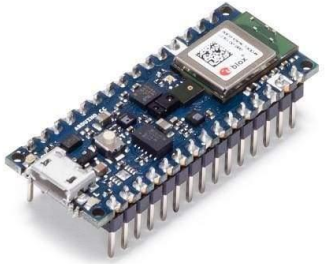

Purpose

AI-enabled developmental **microcontroller board** with USB-A to microB cable

Specifications

- nRF52840 MCU (ARM Cortex-M4): 3.3V, 64MHz, 1MB flash, 256 kB RAM
- Sensors on board: microphone, IMU, color, light, proximity, temperature, humidity, and more!
- BLE module with application-adjacent protocol layers (GAP, GATT) covered by **ArduinoBLE** library

OV 7675 Camera Module

Purpose

Breakout PCB for ***tiny*** camera, for use in person-detection exercises

Specifications

- Low-voltage, 0.3 MP CMOS VGA (can step down to QVGA, QQVGA) image sensor
- Serial Camera Control Bus (SCCB) + Camera Parallel Interface (CPI) / Digital Video Port (DVP) interface
- Breaks ribbon cable out to 2x10 pin array

Tiny Machine Learning Shield

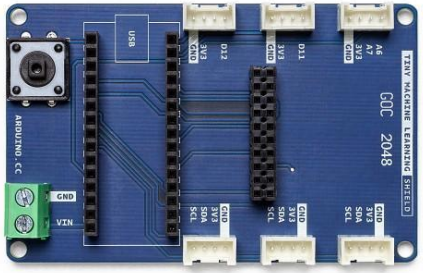

Purpose

A daughter PCB designed to **breakout the I/O** from the Nano 33 BLE sense to permit easy, reliable **communication with** other local, **off-board elements**

Specifications

- Grove connectors (3.3V I2C and simple digital / analog - see pinouts)
- 2x10 pin array for OV7675 camera module
- Voltage input terminal block, accepts 4.5 to 21V (down regulated to 3.3V on Nano 33)

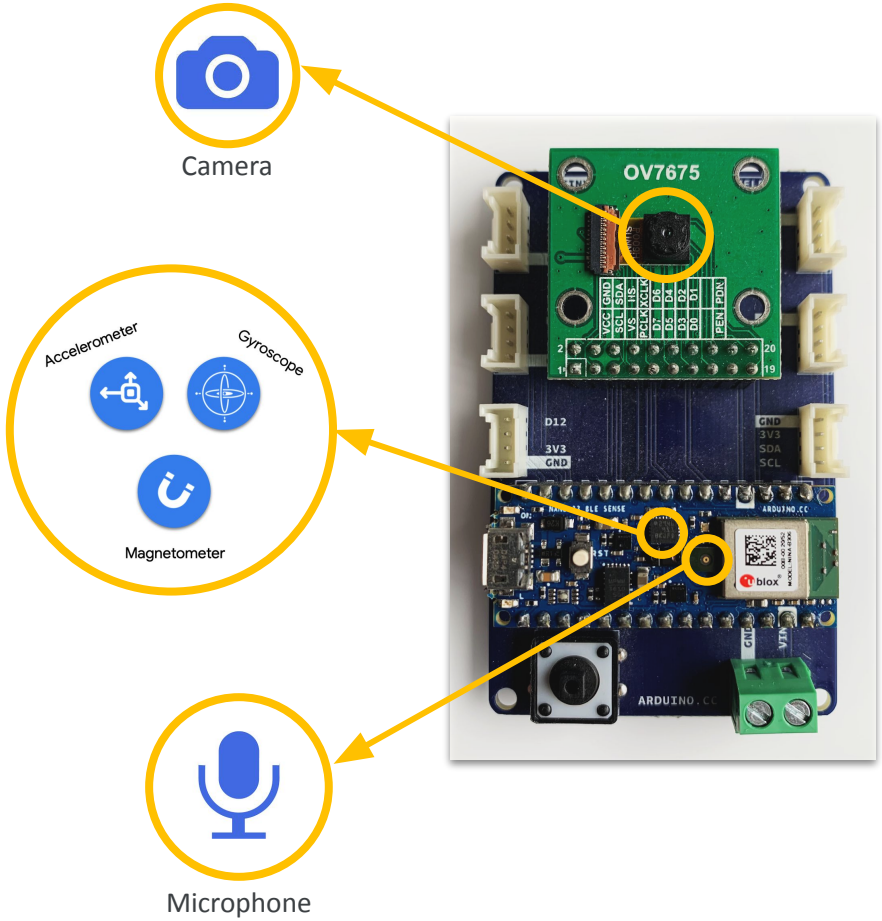

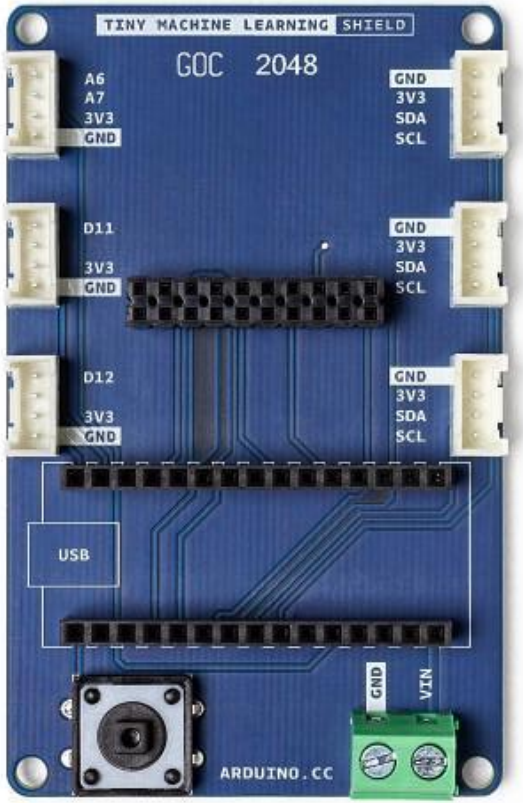

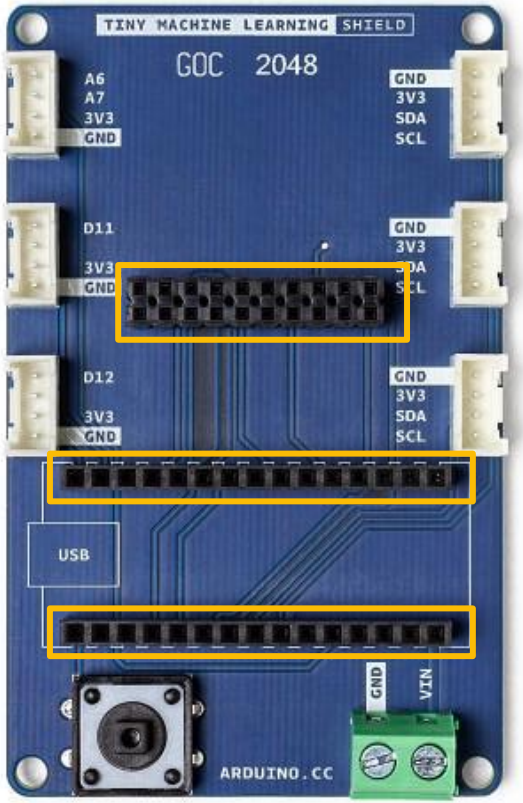

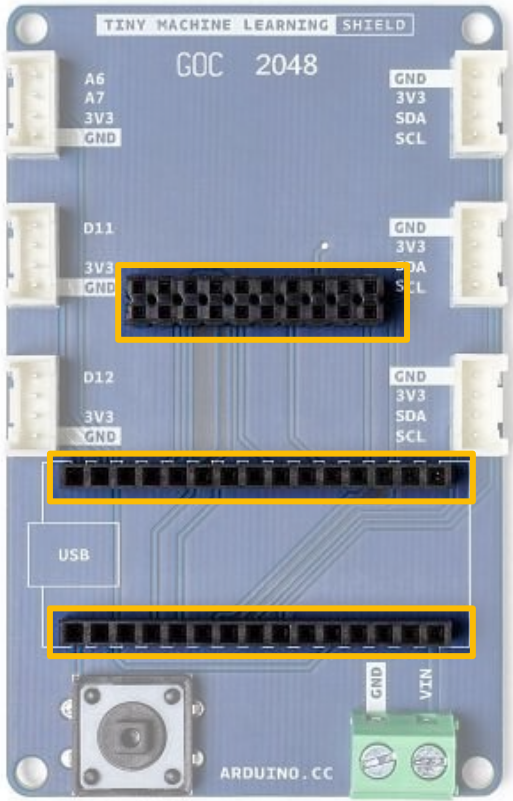

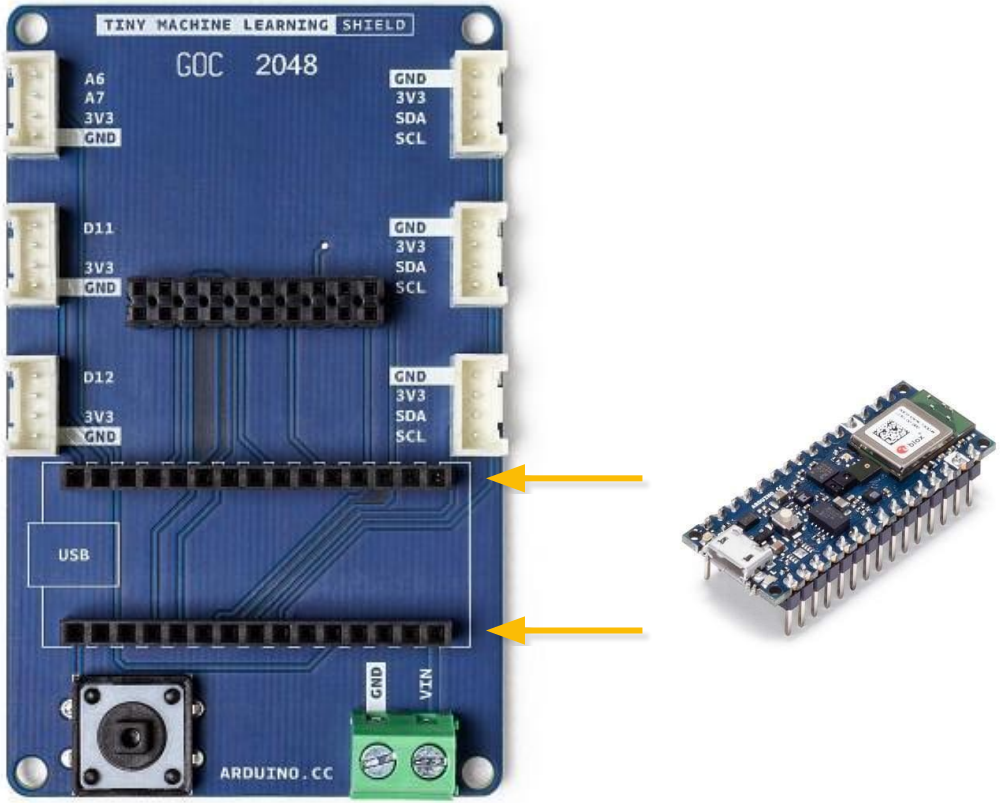

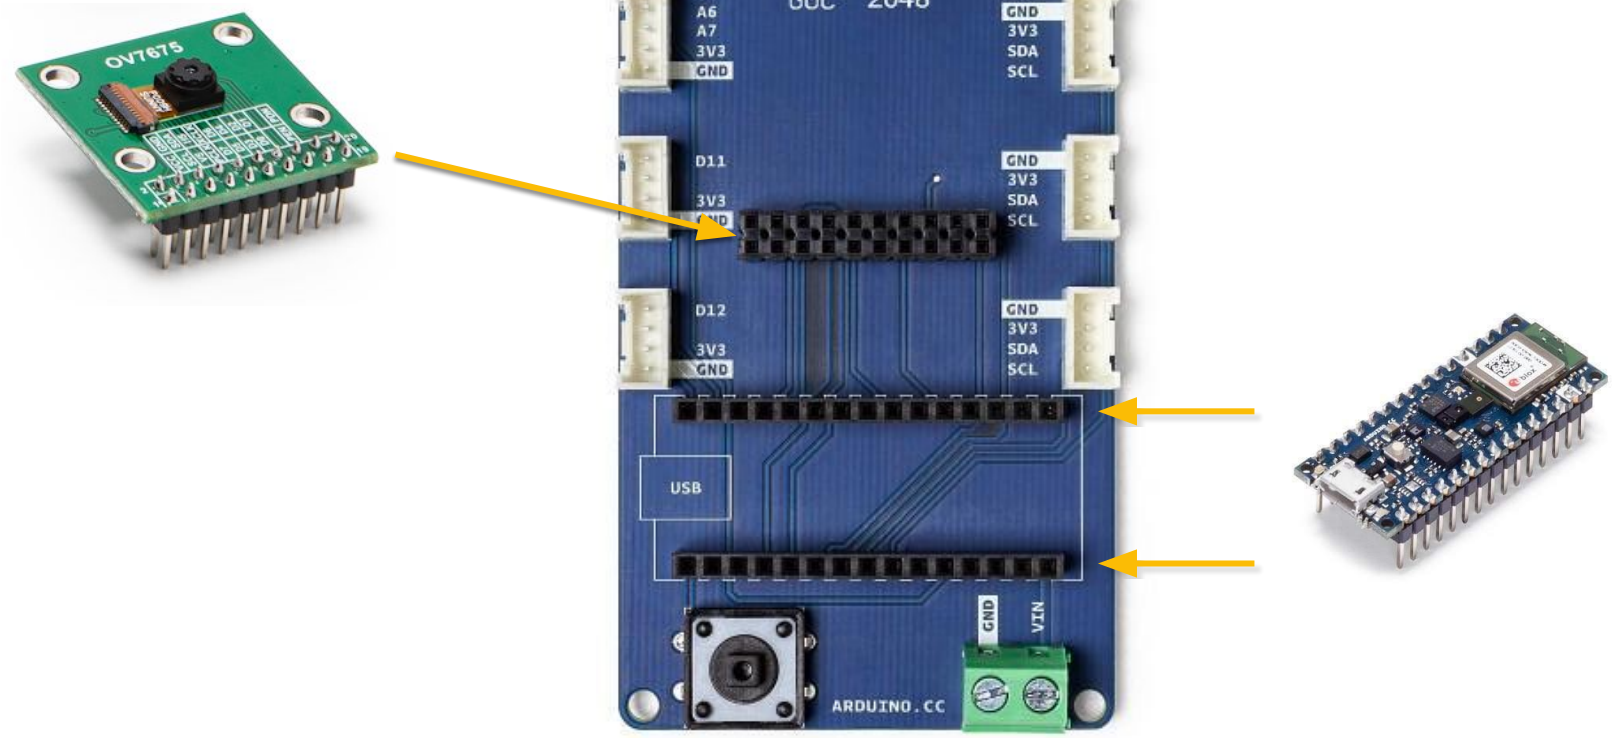

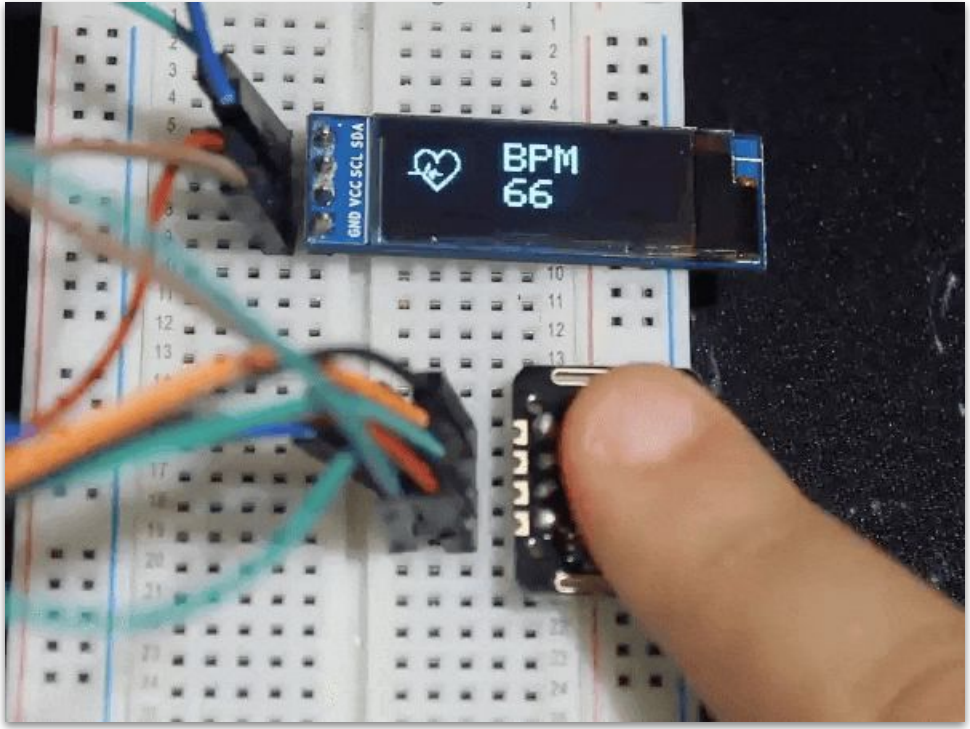

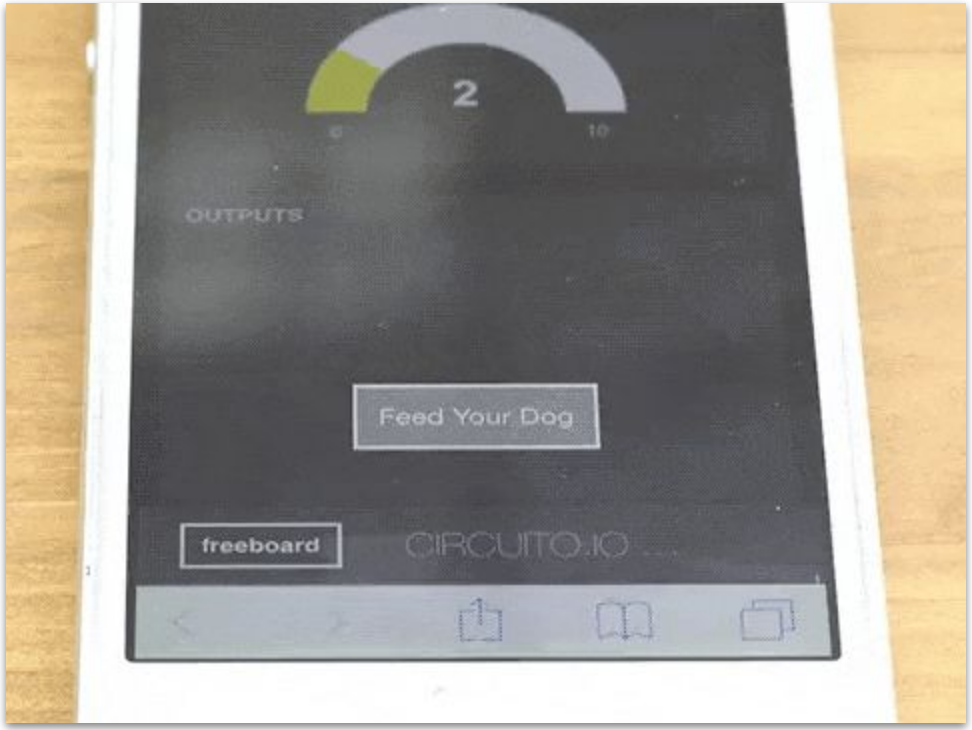

Today's Agenda

- Review + Why Tiny?
- Hardware Basics
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino
- Summary & Next Steps

Today's Agenda

- Review + Why Tiny?
- Hardware Basics
- **Installing and Starting the Arduino IDE**
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino
- Summary & Next Steps

www.arduino.cc/en/software

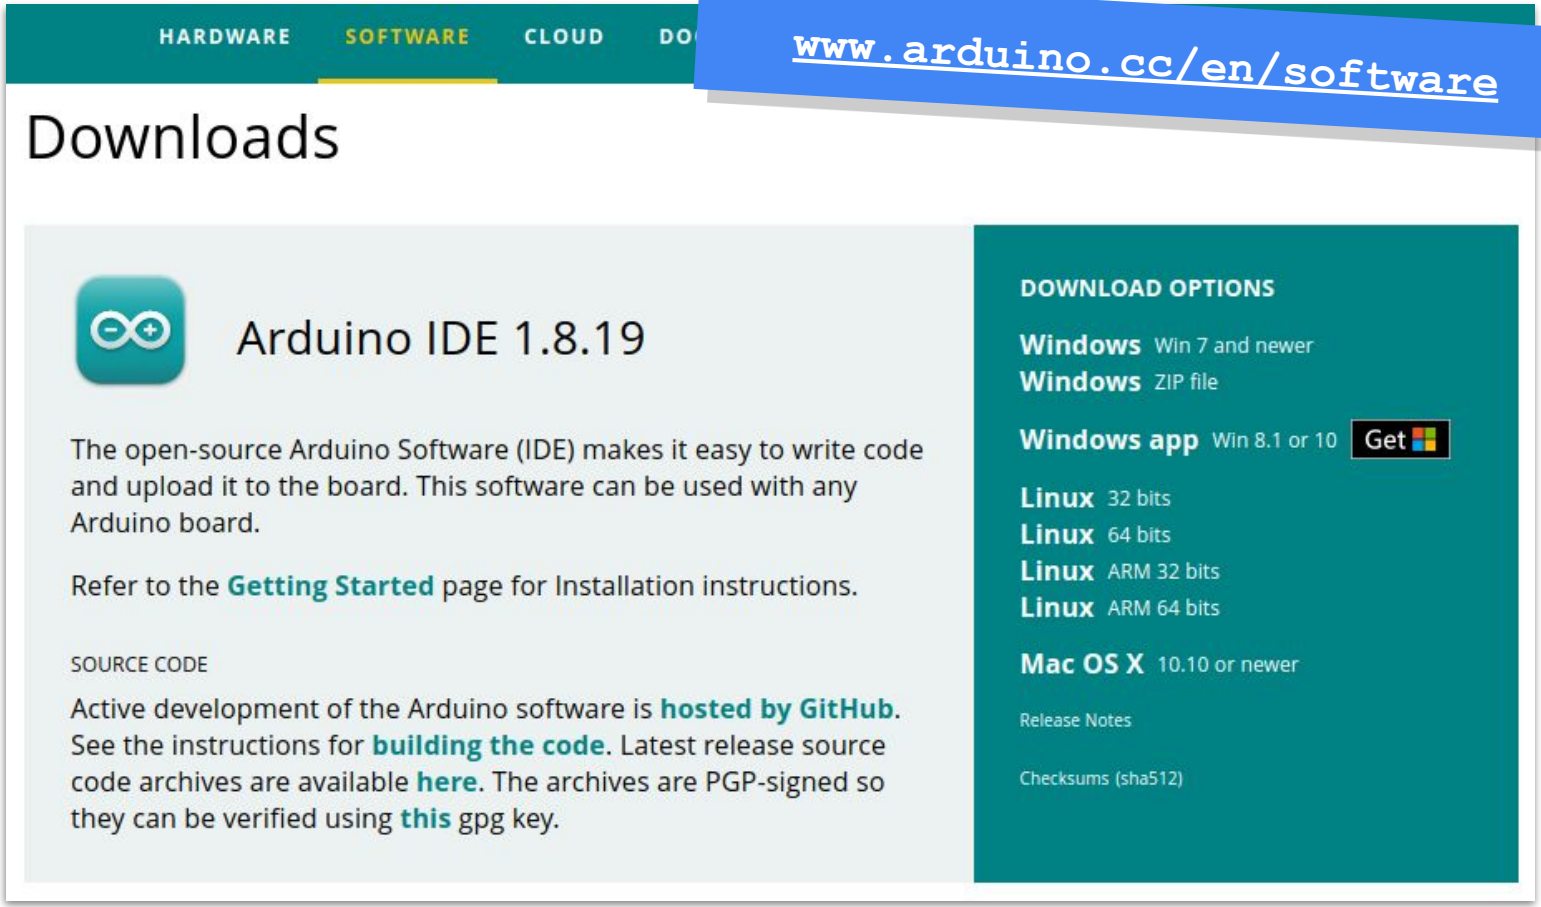

www.arduino.cc/en/software

www.arduino.cc/en/software

Getting Started with Arduino products

Welcome to Arduino! Before you start controlling the world around you, you'll need to set up the software to program your board.

The Arduino Software (IDE) allows you to write programs and upload them to your board. In the Arduino Software page you will find two options:

1. If you have a reliable Internet connection, you should use the online IDE (Arduino Web Editor). It will allow you to save your sketches in the cloud, having them available from any device and backed up. You will always have the most up-to-date version of the IDE without the need to install updates or community generated libraries.
2. If you would rather work offline, you should use the latest version of the desktop IDE.

Code online on the Arduino Web Editor

To use the online IDE simply follow these instructions. Remember that boards work out-of-the-box on the Web Editor, you only need to install Arduino Create Agent to get started.

Install the Arduino Desktop IDE

To get step-by-step instructions select one of the following link accordingly to your operating system.

- Windows
- Mac OS
- Linux
- Portable IDE (Windows and Linux)
- ChromeOS (Arduino Create App) in the Chrome Web Store

Choose your board in the list here on the right to learn how to get started with it and how to use it on the Desktop IDE.

Learn Arduino

- Read an introduction on what is Arduino and why you'd want to use it.
- What is the Arduino Software (IDE) and how do I change the default language?

Arduino IoT Cloud:

- Getting Started with Arduino IoT Cloud

Instructions for our boards:

- Due
- MEGA2560
- MKR1000
- MKR WiFi 1010
- MKR FOX 1200
- MKR WAN 1300
- MKR WAN 1310
- MKR GSM 1400
- MKR NB 1500
- MKR Vidor 4000
- MKRZERO
- Nano
- NANO 33 IoT
- NANO 33 BLE
- NANO 33 BLE Sense

Getting Started with Arduino products

Welcome to Arduino! Before you start controlling the world around you, you'll need to set up the software to program your board.

The Arduino Software (IDE) allows you to write programs and upload them to your board. In the Arduino Software page you will find two options:

1. If you have a reliable Internet connection, you should use the online IDE (Arduino Web Editor). It will allow you to save your sketches in the cloud, having them available from any device and backed up. You will always have the most up-to-date version of the IDE without the need to install updates or community generated libraries.
2. If you would rather work offline, you should use the latest version of the desktop IDE.

Code online on the Arduino Web Editor

To use the online IDE simply follow these instructions. Remember that boards work out-of-the-box on the Web Editor, you only need to install Arduino Create Agent to get started.

Install the Arduino Desktop IDE

To get step-by-step instructions select one of the following link accordingly to your operating system.

- Windows
- Mac OS
- Linux

- Portable IDE (Windows and Linux)
- ChromeOS (Arduino Create App) in the Chrome Web Store

Choose your board in the list here on the right to learn how to get started with it and how to use it on the Desktop IDE.

Learn Arduino

- Read an introduction on what is Arduino and why you'd want to use it.
- What is the Arduino Software (IDE) and how do I change the default language?

Arduino IoT Cloud:

- Getting Started with Arduino IoT Cloud

Instructions for our boards:

- Due
- MEGA2560
- MKR1000
- MKR WiFi 1010
- MKR FOX 1200
- MKR WAN 1300
- MKR WAN 1310
- MKR GSM 1400
- MKR NB 1500
- MKR Vidor 4000
- MKRZERO
- Nano
- NANO 33 IoT
- NANO 33 BLE
- NANO 33 BLE Sense

head to Arduino's software **website**

www.arduino.cc/en/software

then **download** and **install** *right* **now!**

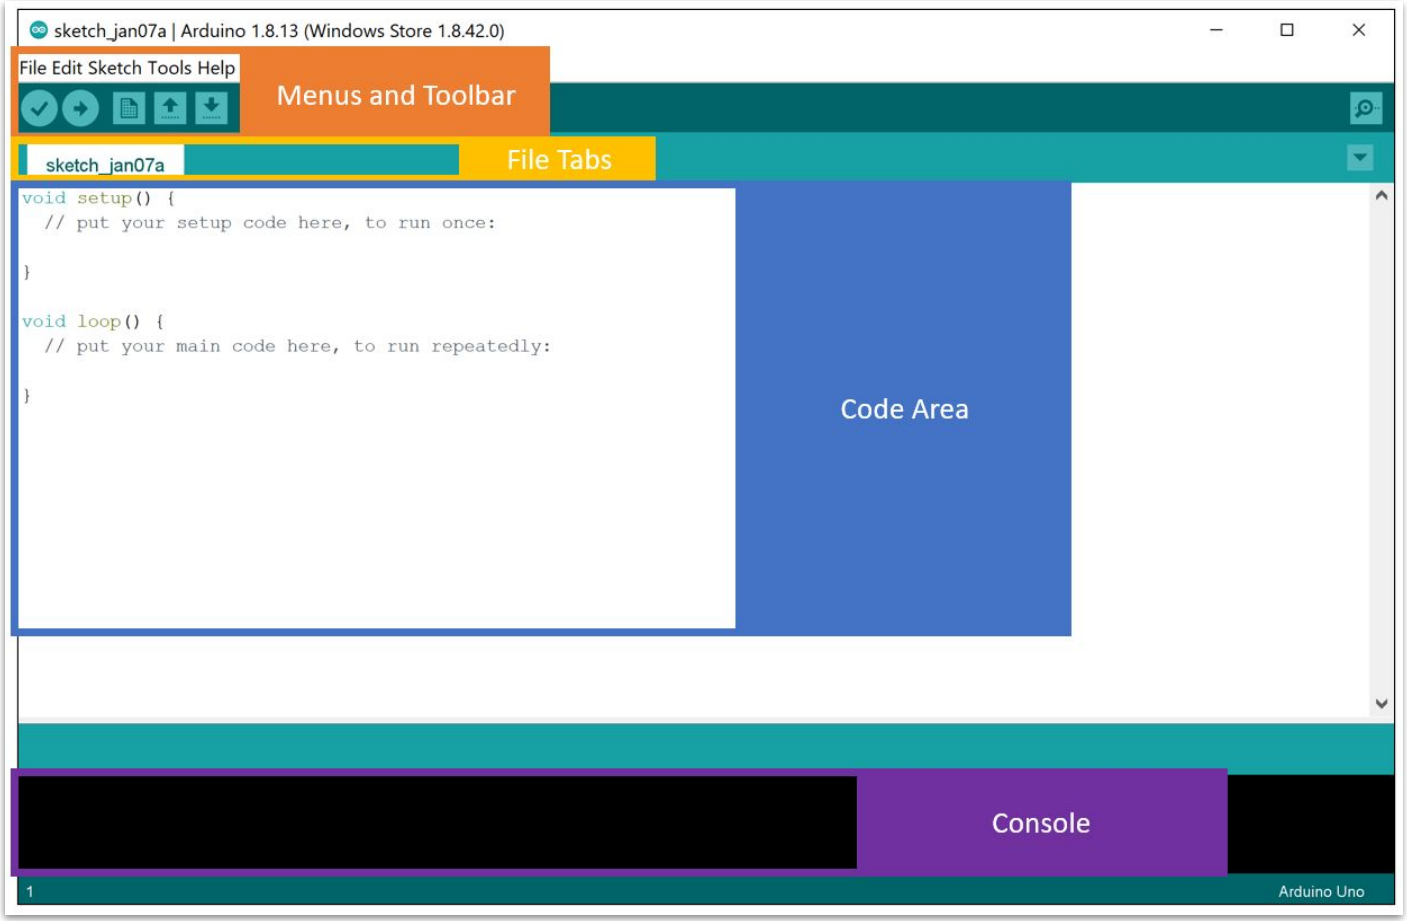

Install Extras

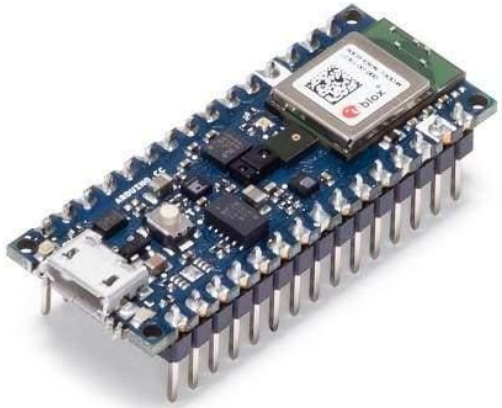

Tools → Board → Boards Manager

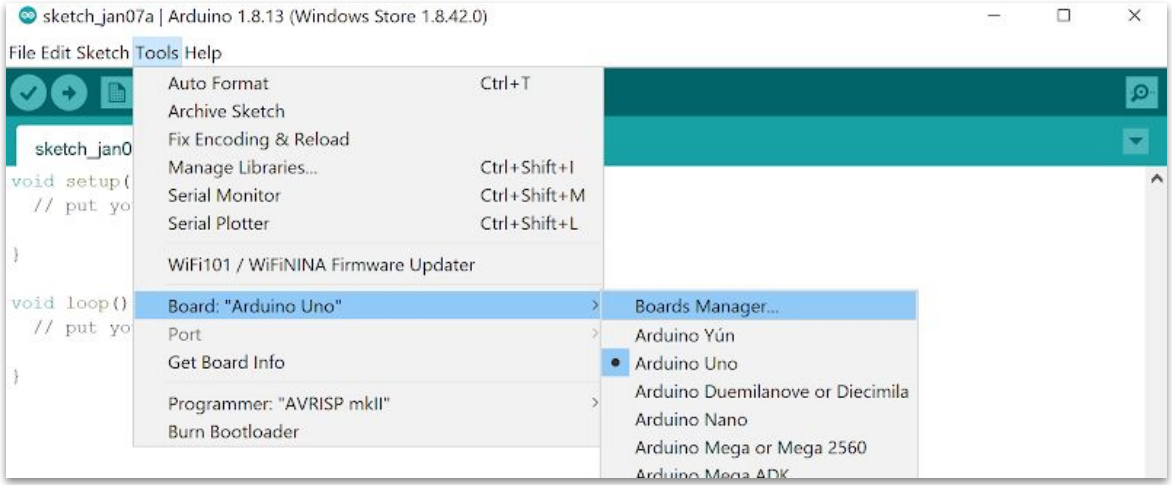

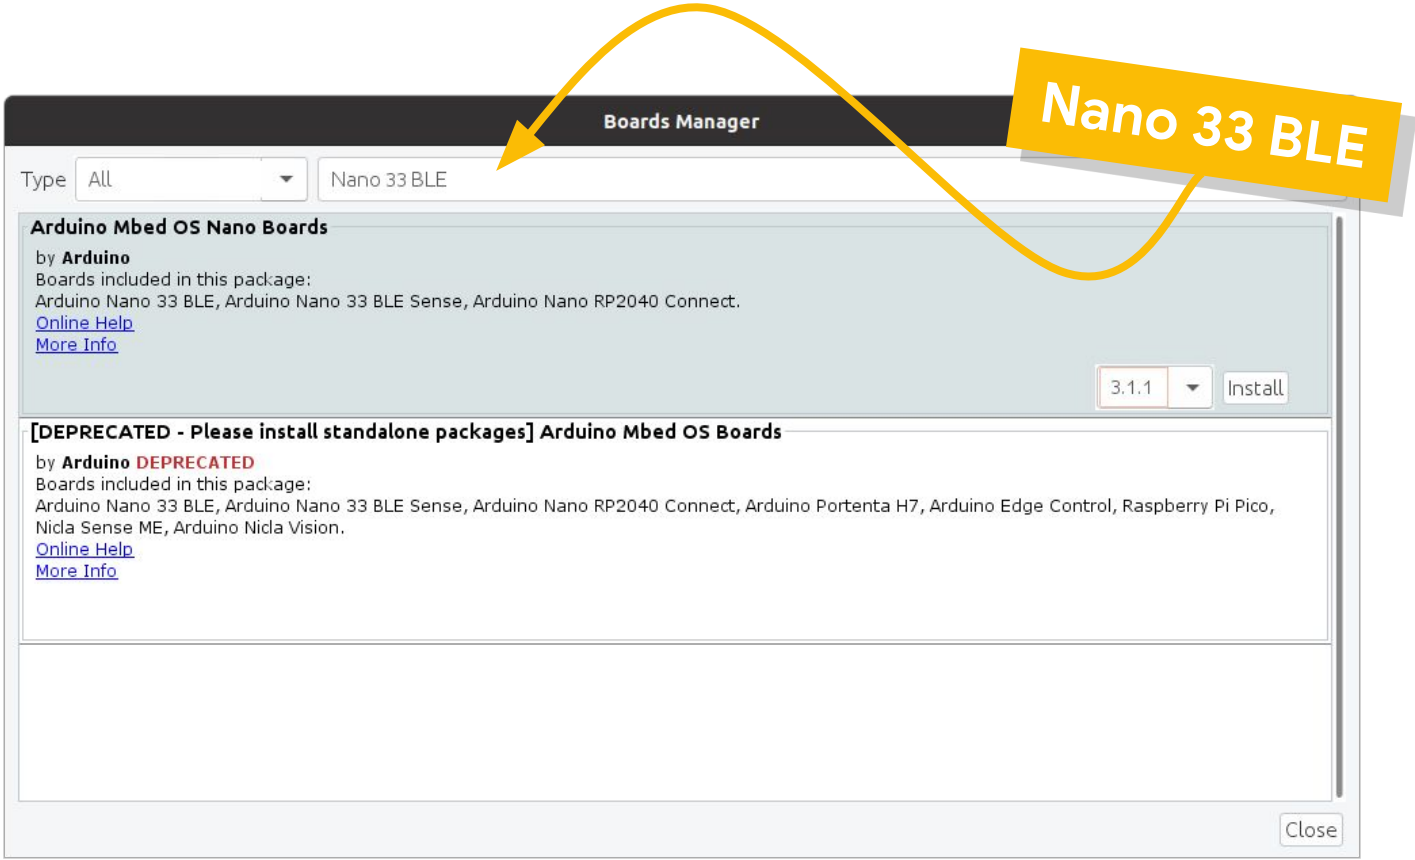

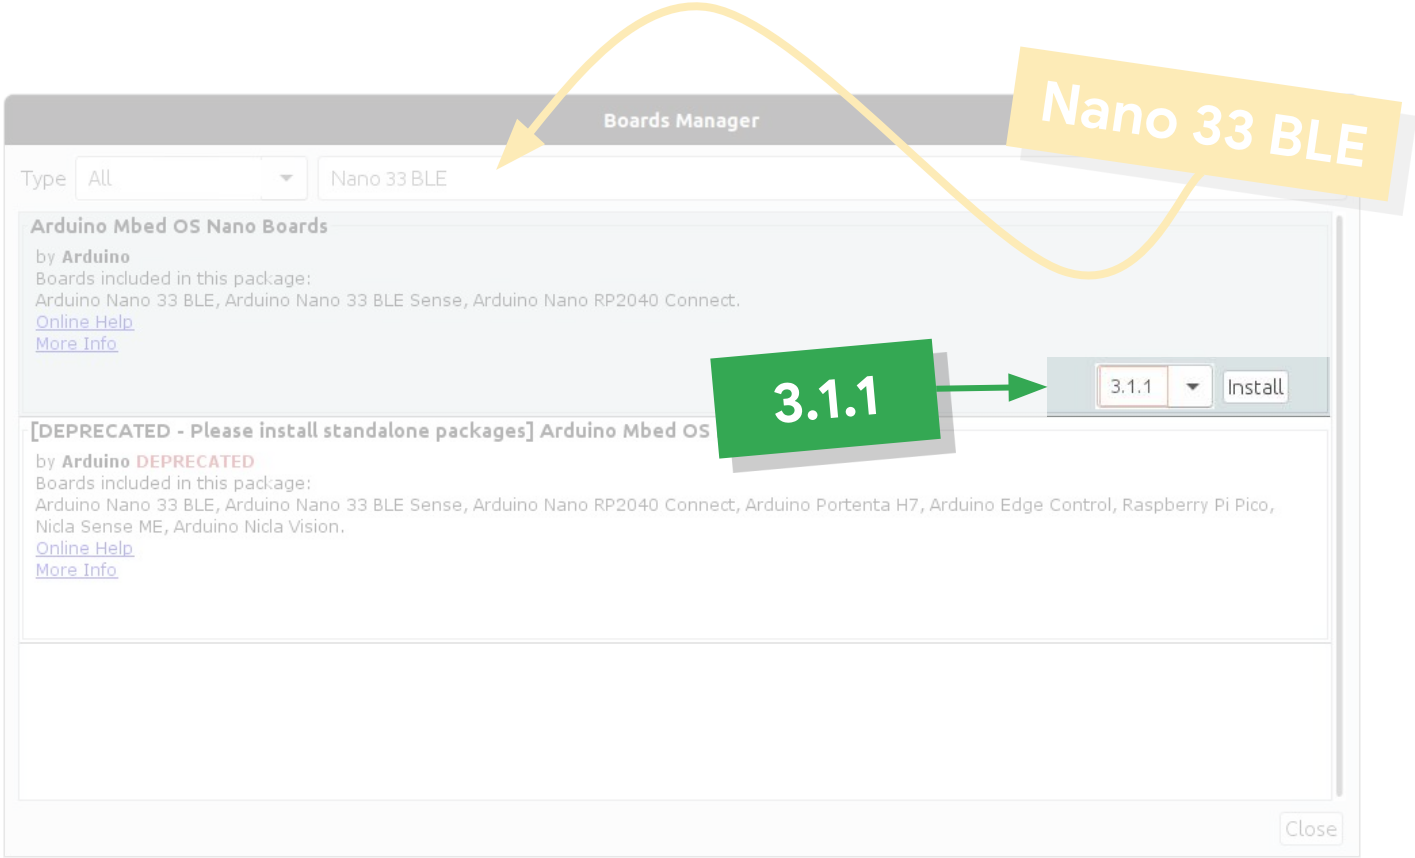

Tools → Manage Libraries...

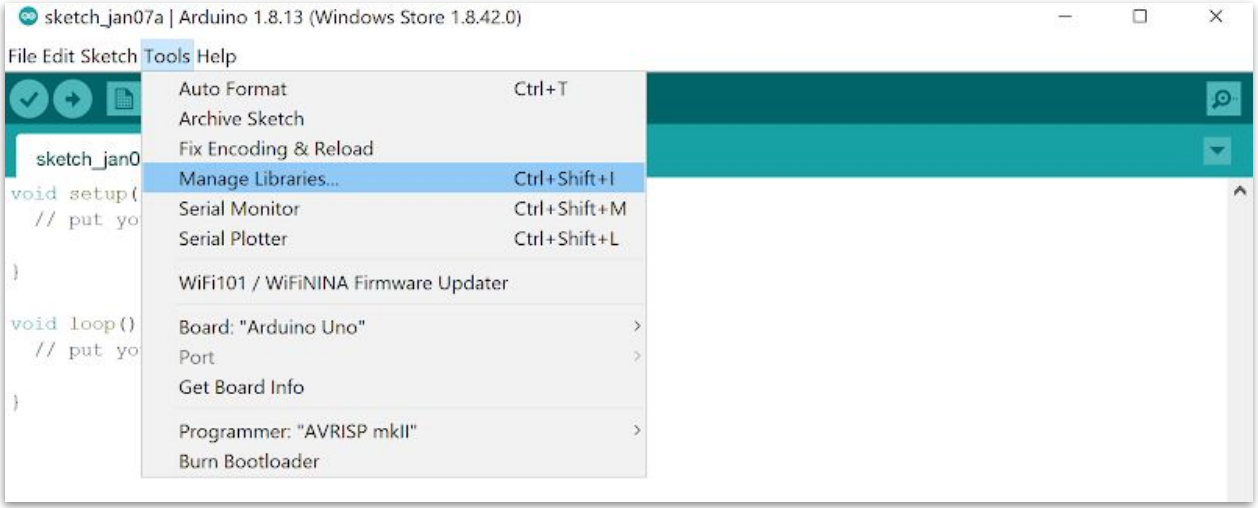

Tools → Manage Libraries...

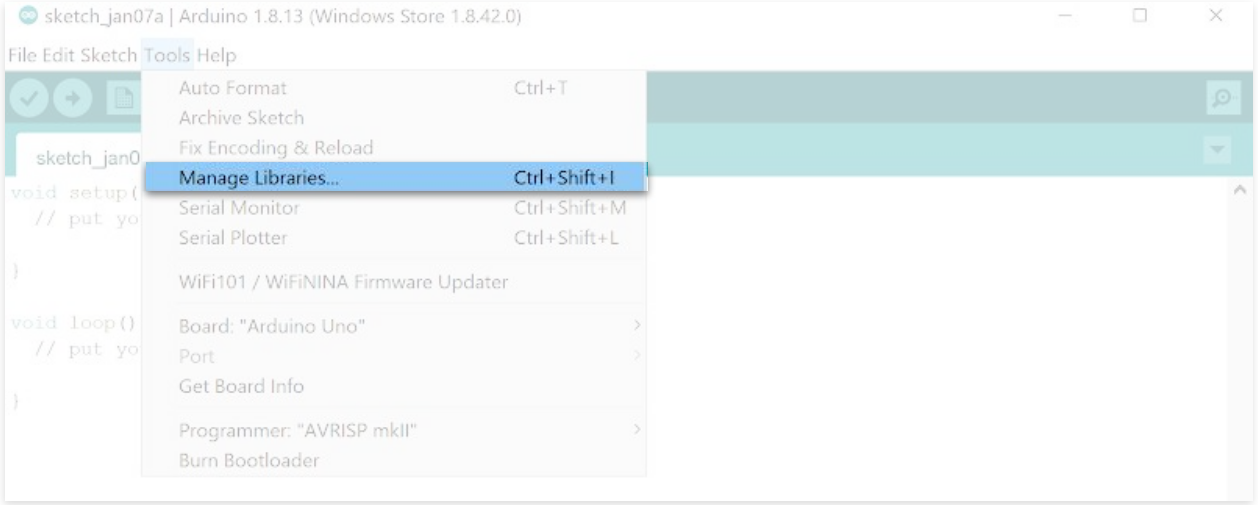

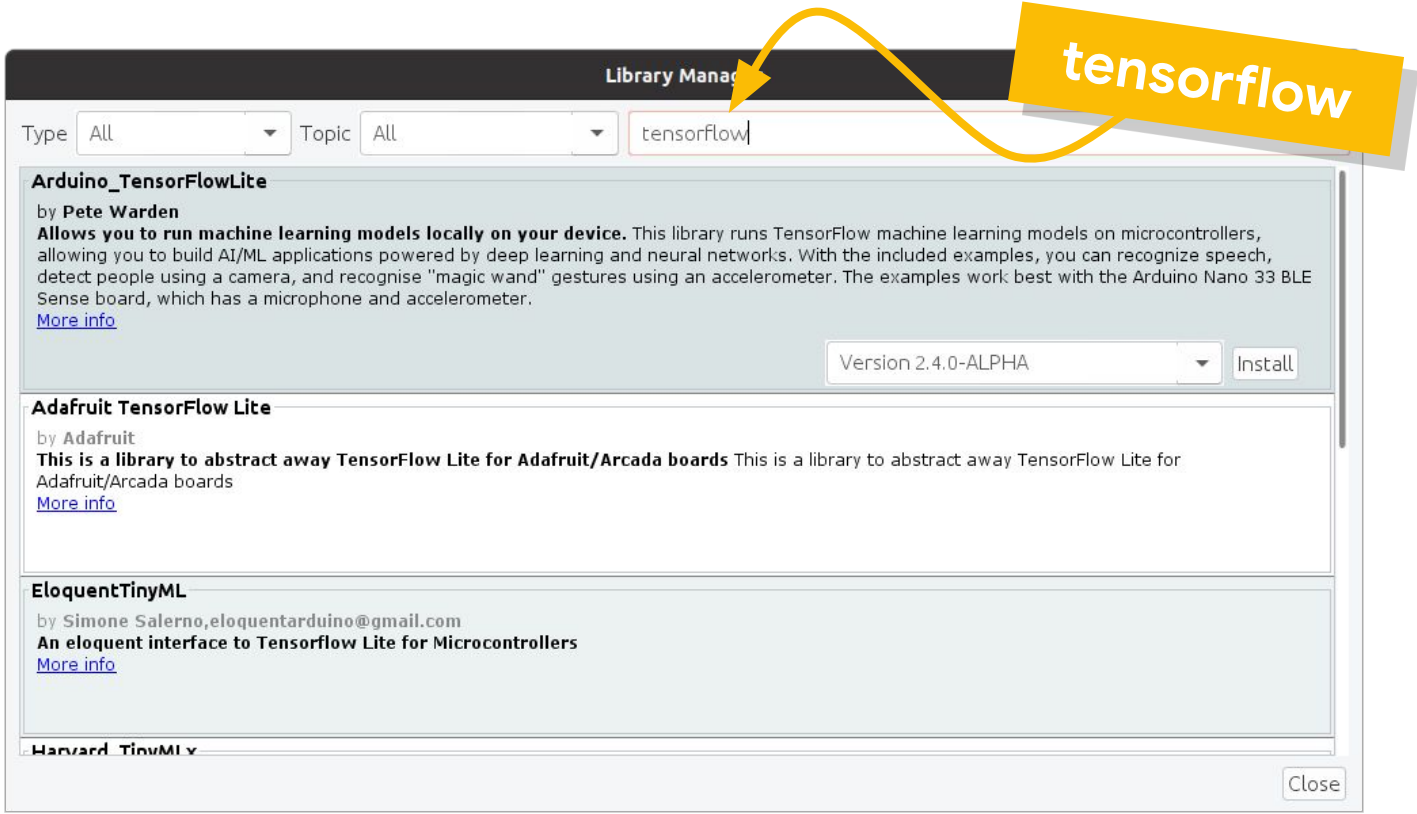

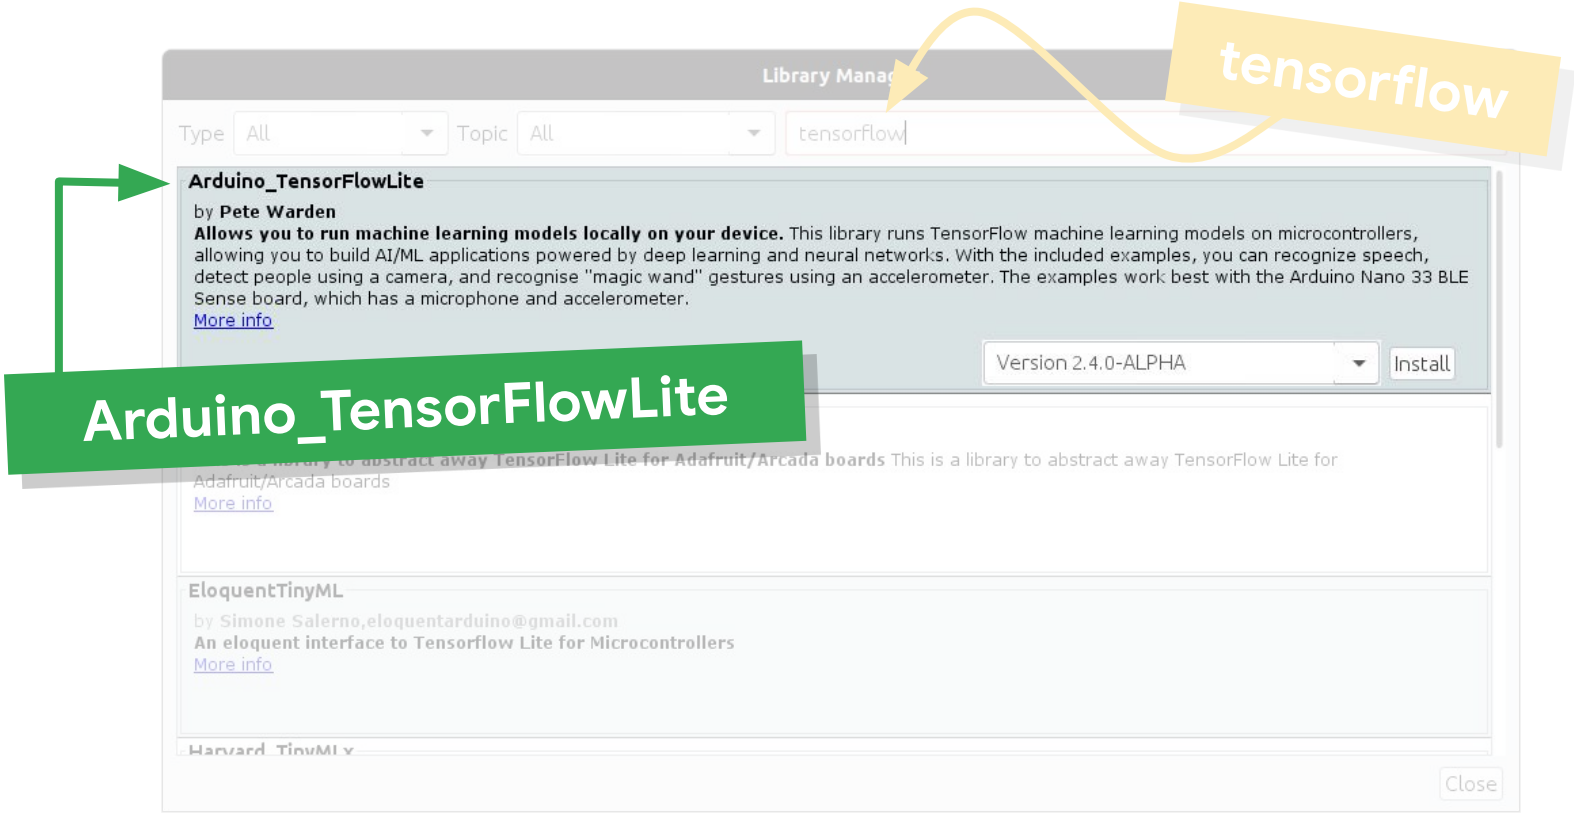

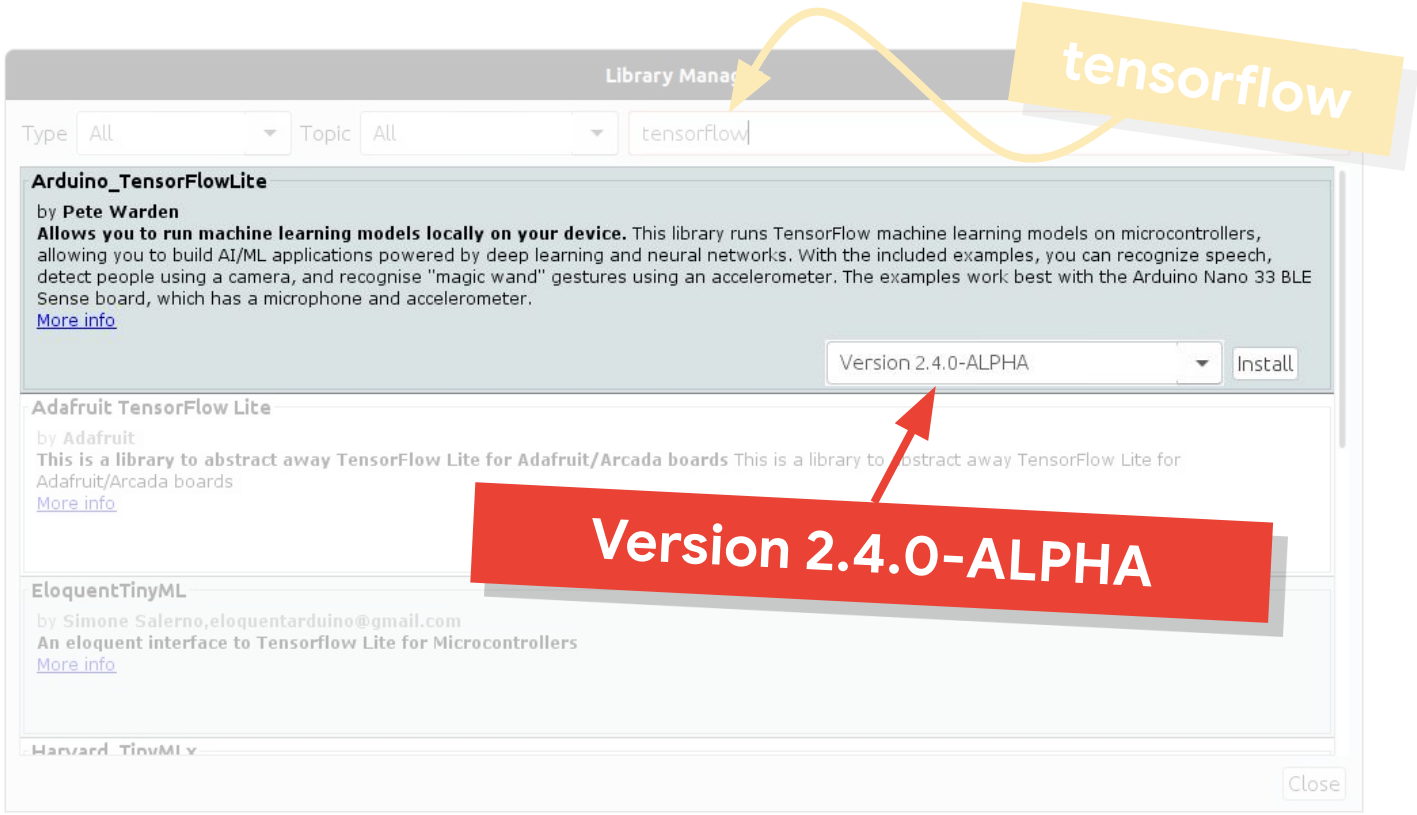

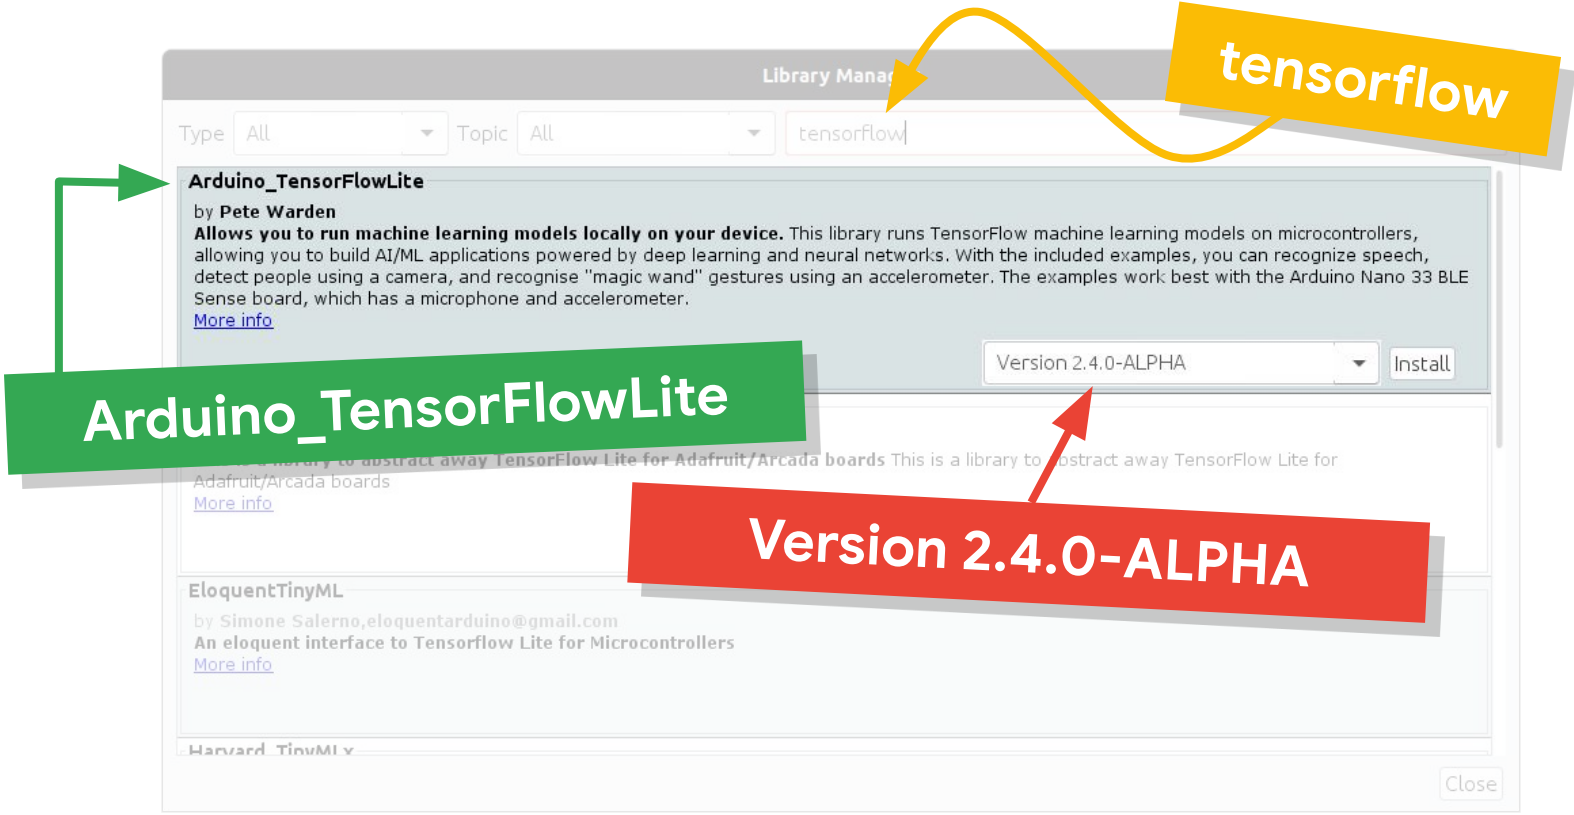

Install Extras

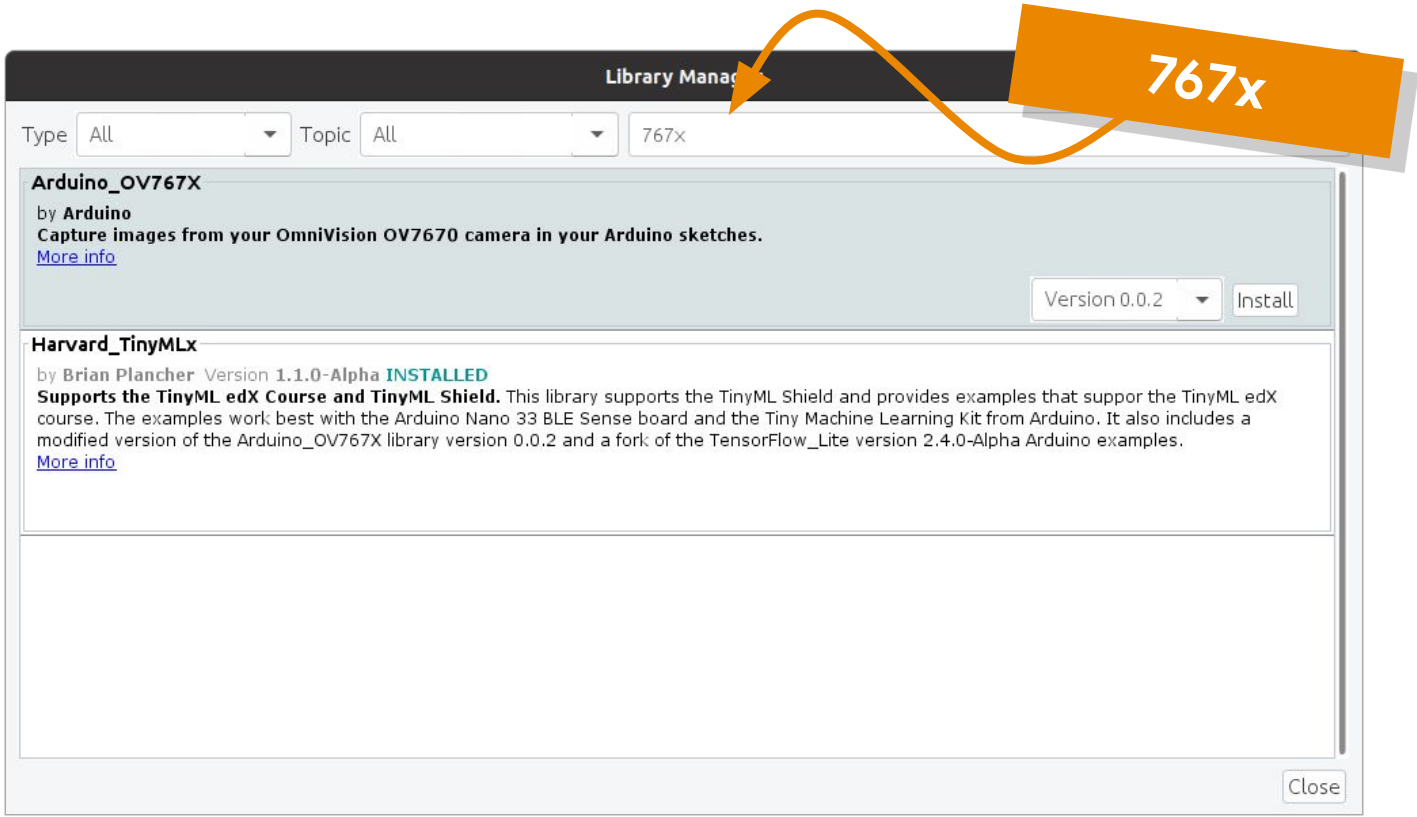

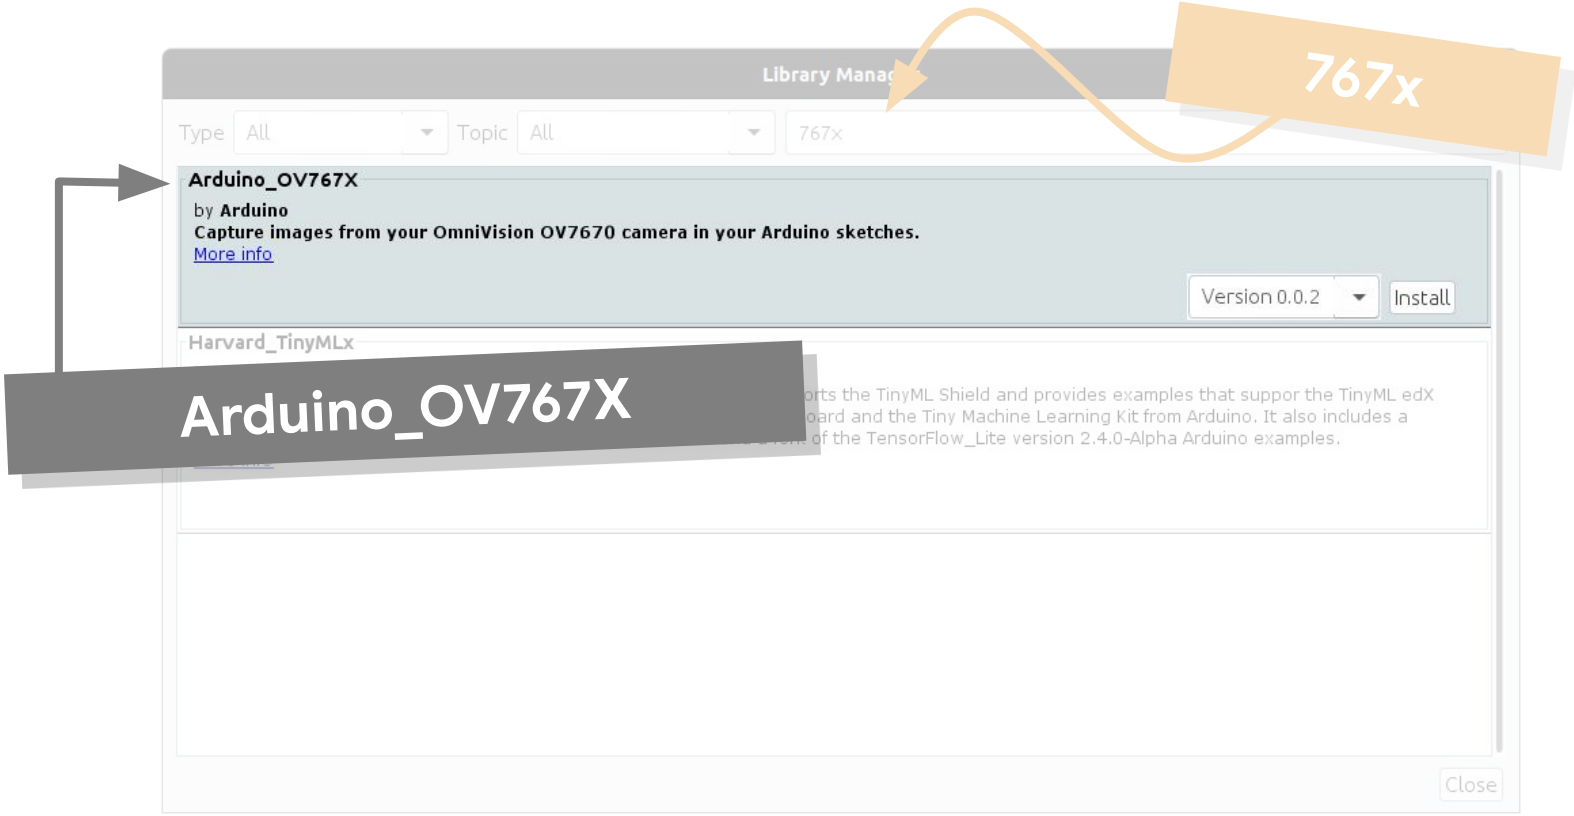

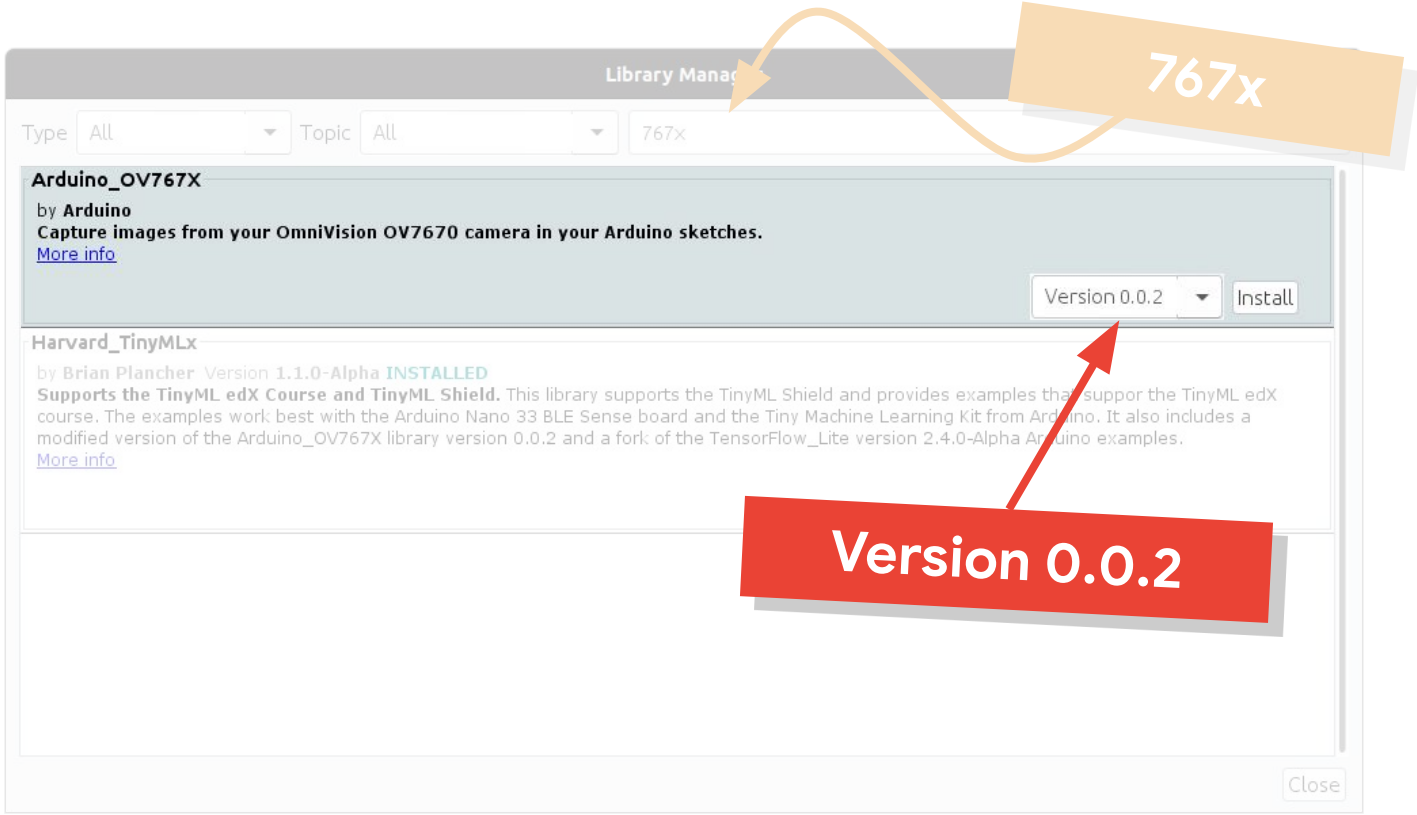

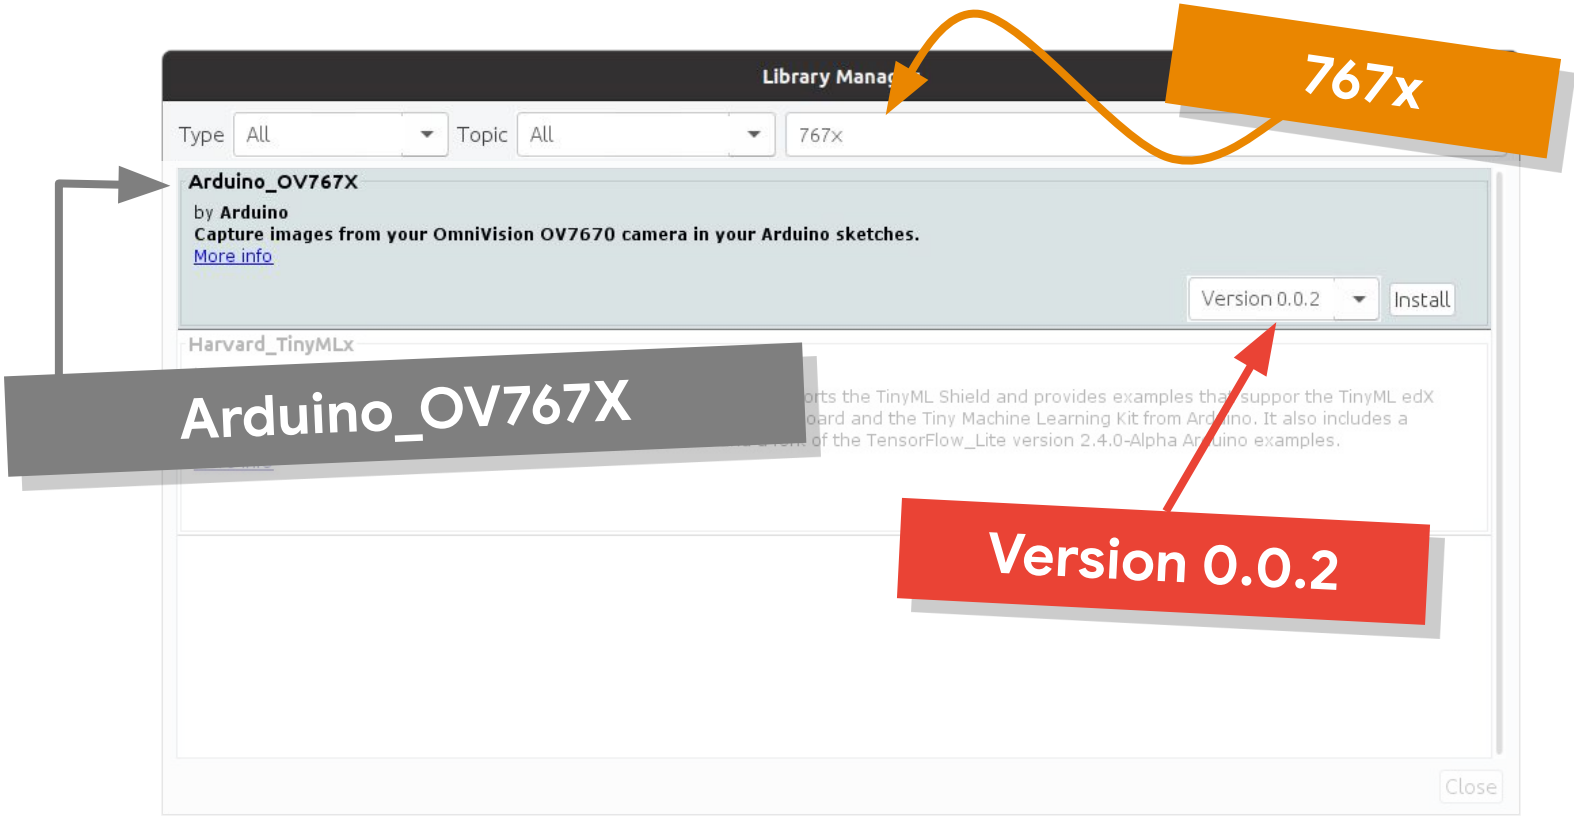

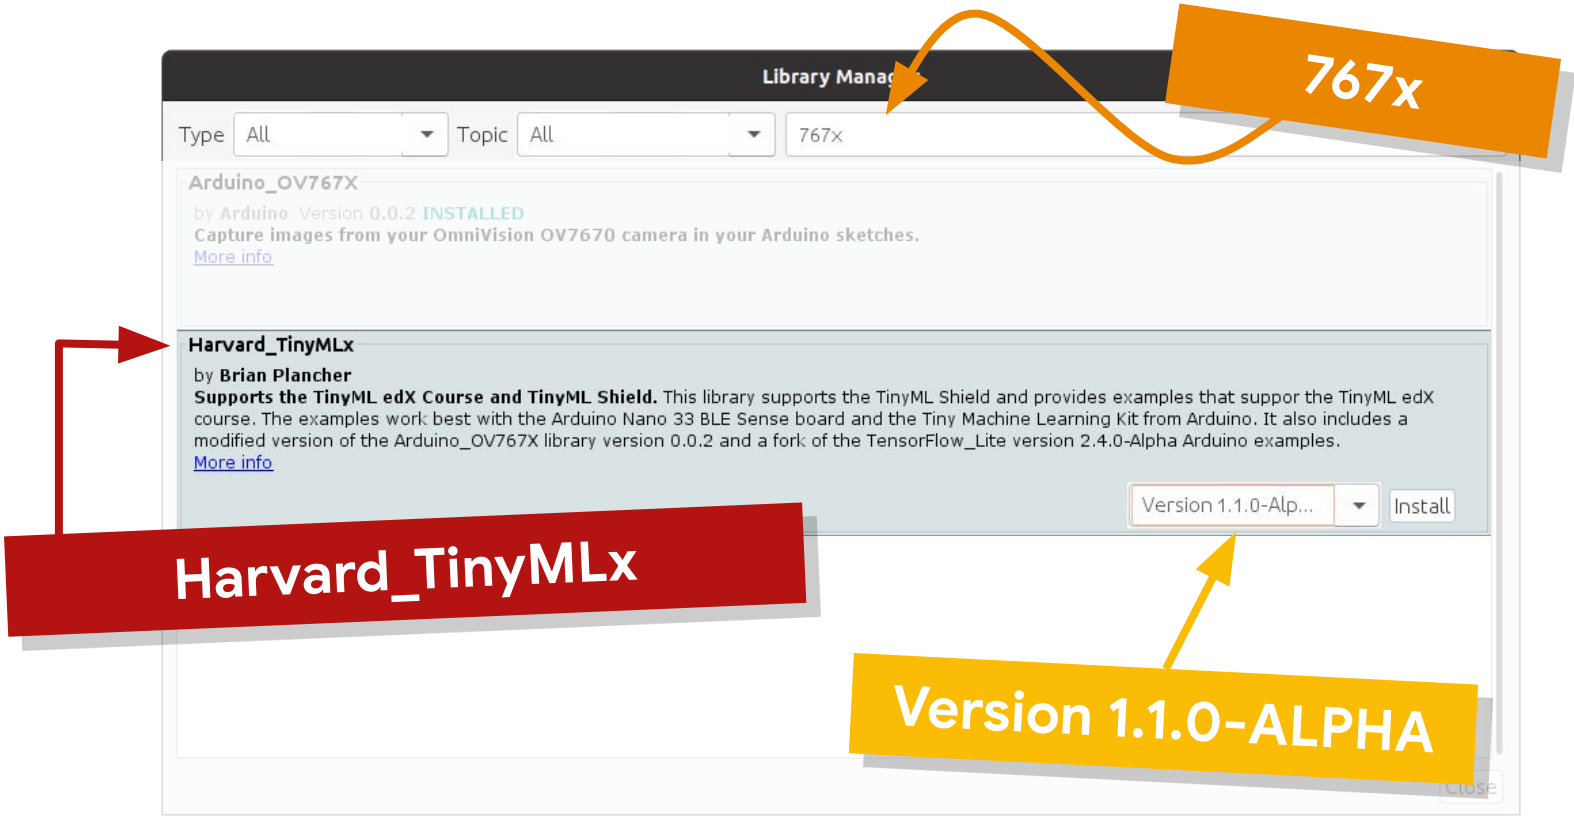

Install Extras

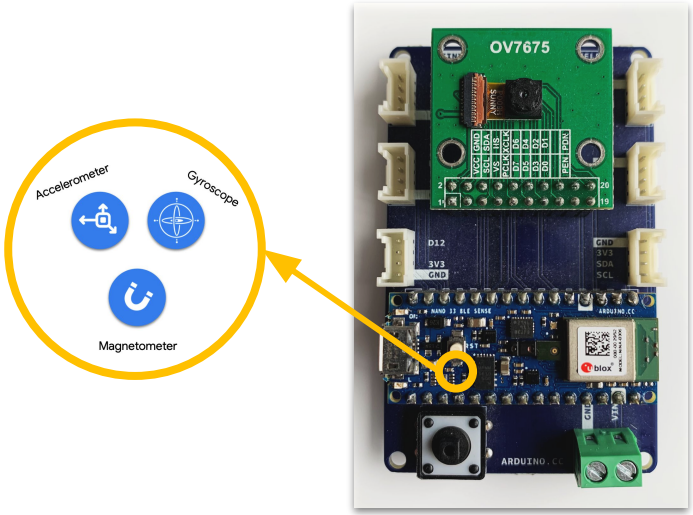

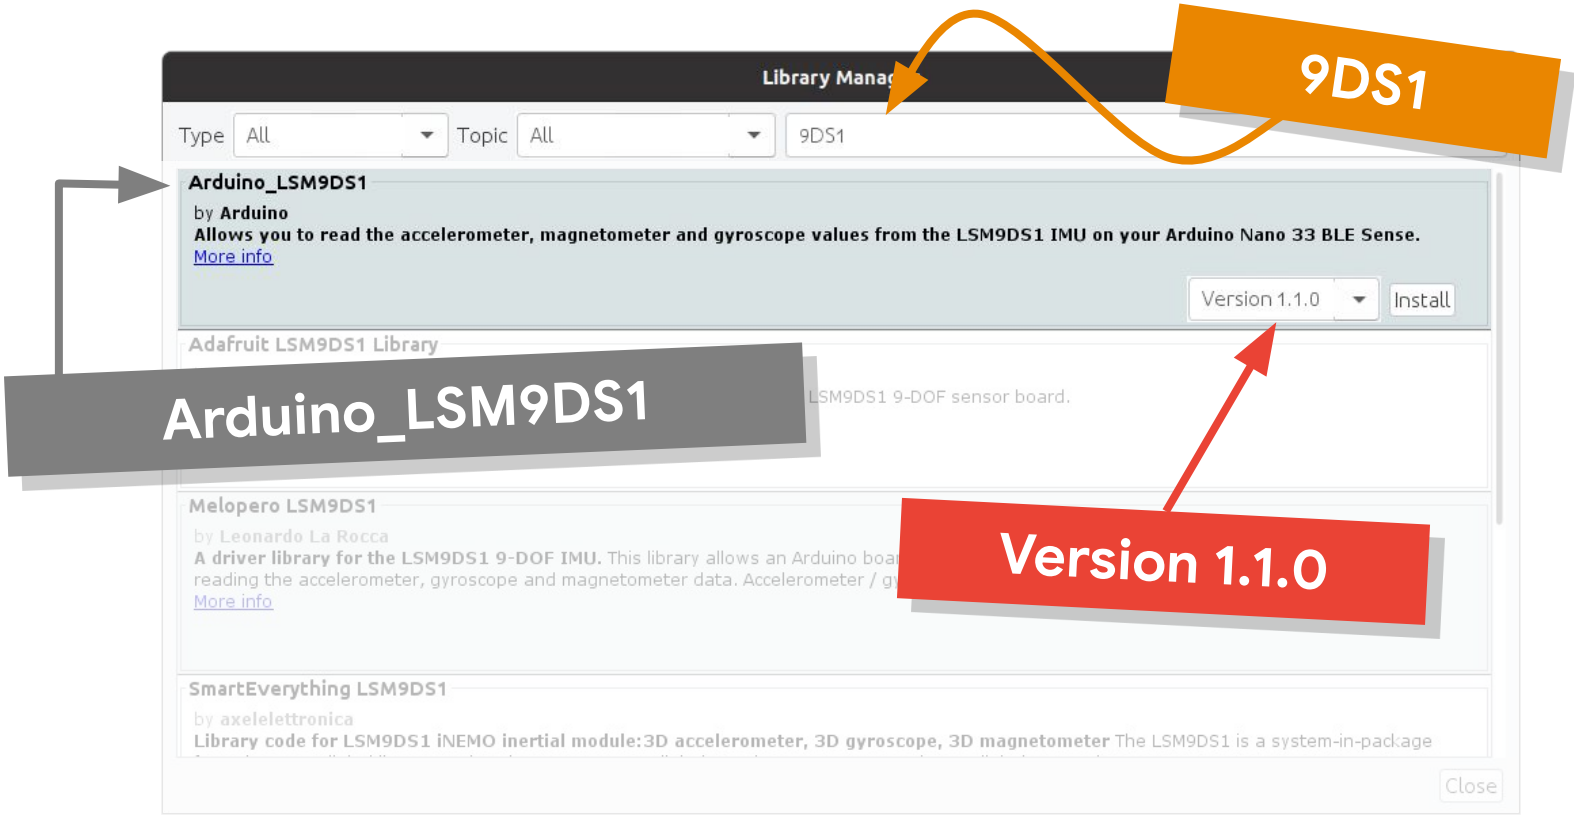

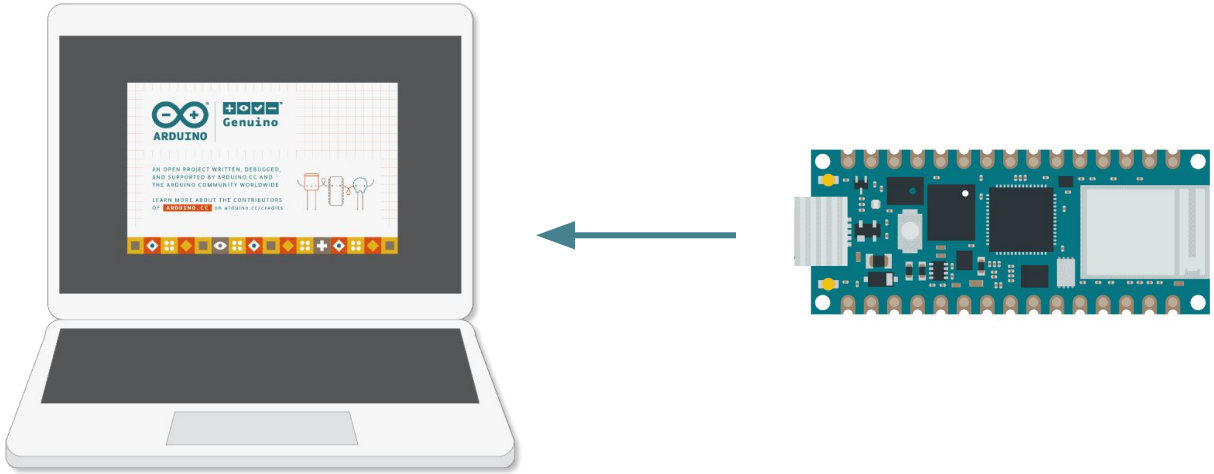

Tools → **Board: “[something]”**

→ Arduino Mbed OS Nano Boards

→ Arduino Nano 33 BLE

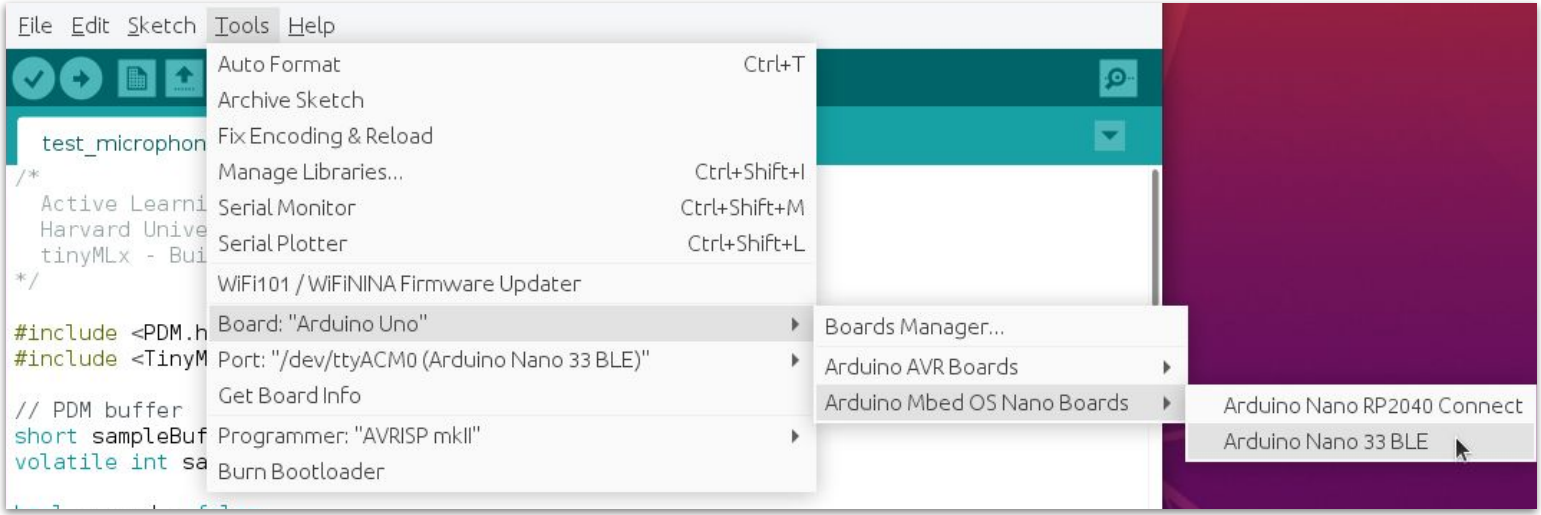

MICRO USB CABLE

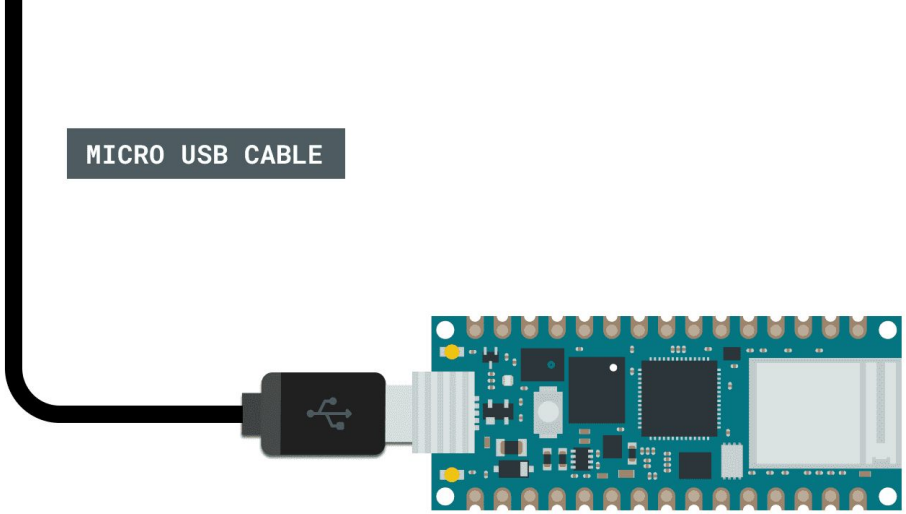

PLUG IN via MicroUSB

to your computer *right* ***now!***

set **board** and **port** for later

Tools → **Port: "[???] (Arduino Nano 33 BLE)"**

→ **<???> (Arduino Nano 33 BLE)**

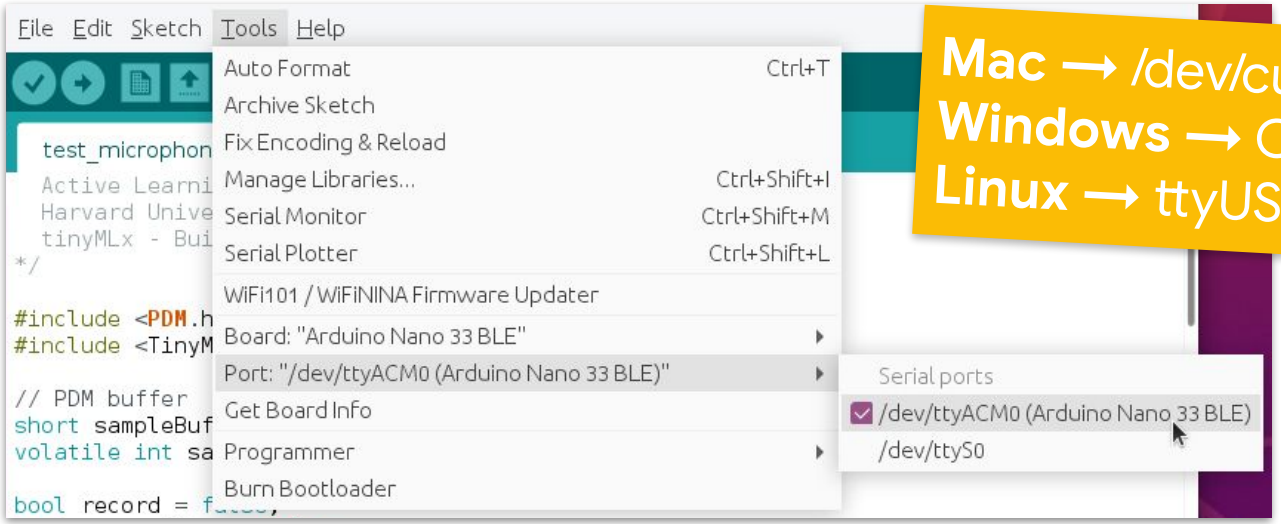

Mac → /dev/cu.usbmodem<#>
Windows → COM<#>
Linux → ttyUSB<#> or ttyACM<#>

Today's Agenda

- Review + Why Tiny?
- Hardware Basics
- Installing and Starting the Arduino IDE
- **Testing Your TinyML Kit**
- Deploying KWS model onto Arduino
- Summary & Next Steps

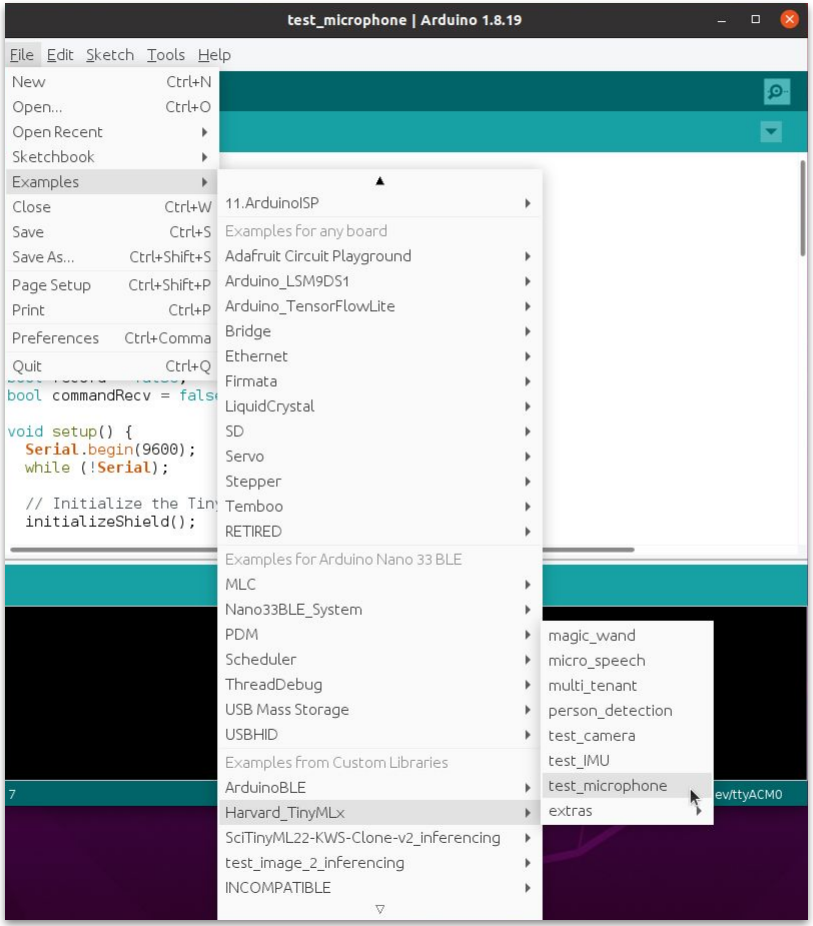

File → **Examples**

→ **Harvard_TinyMLx**

→ **test_microphone**

upload

done?

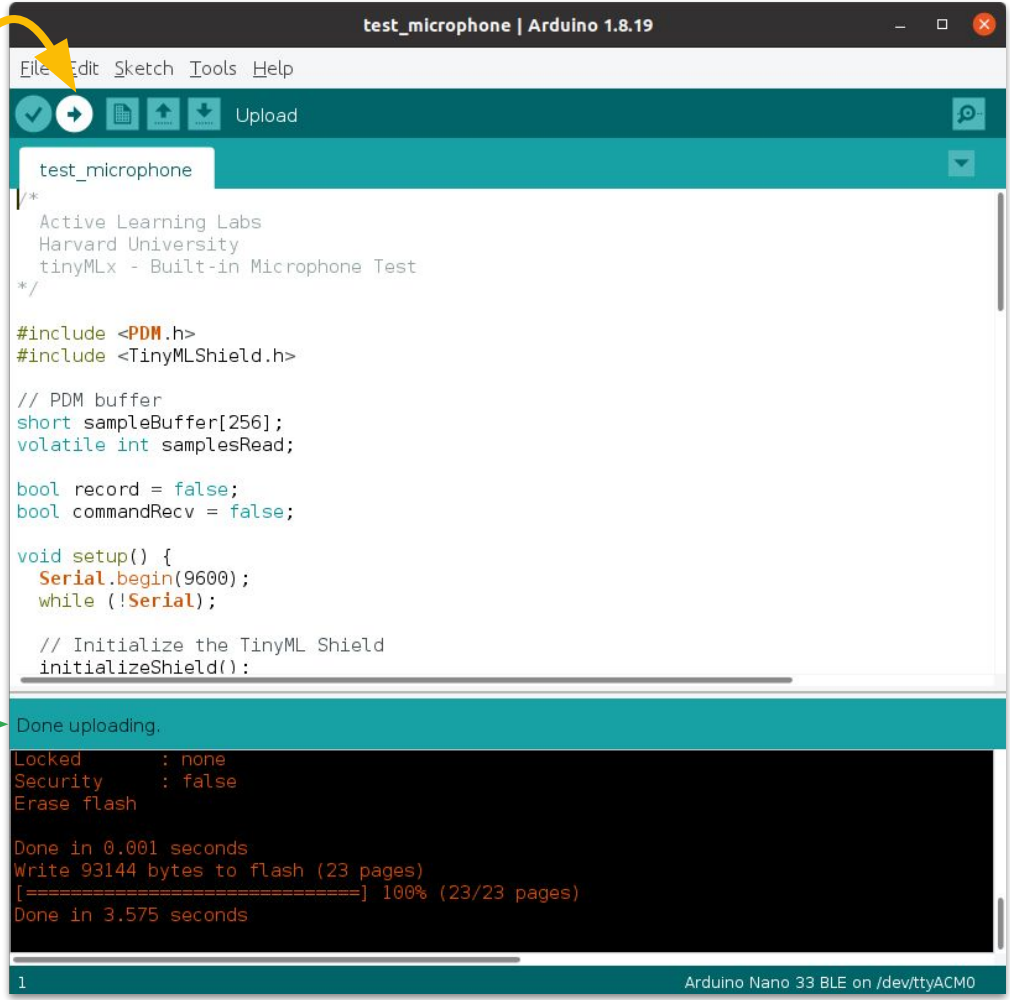

An error occurred while uploading the sketch

```
/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!

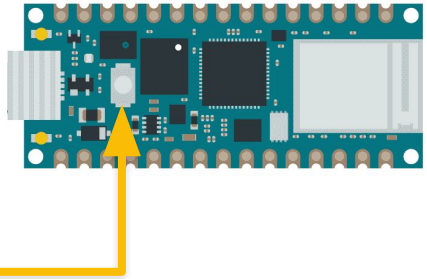

Tools → Serial Monitor

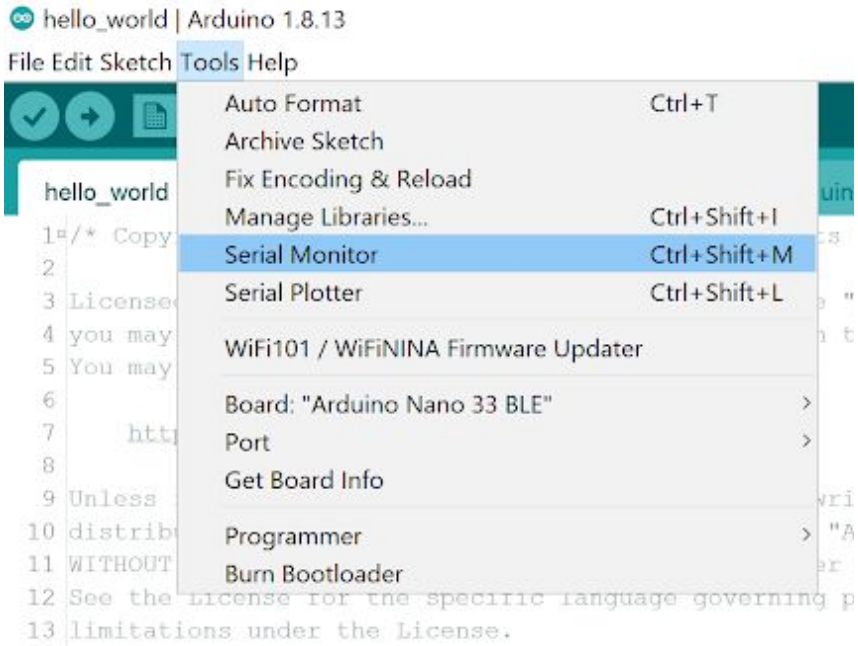

Tools → Serial Monitor

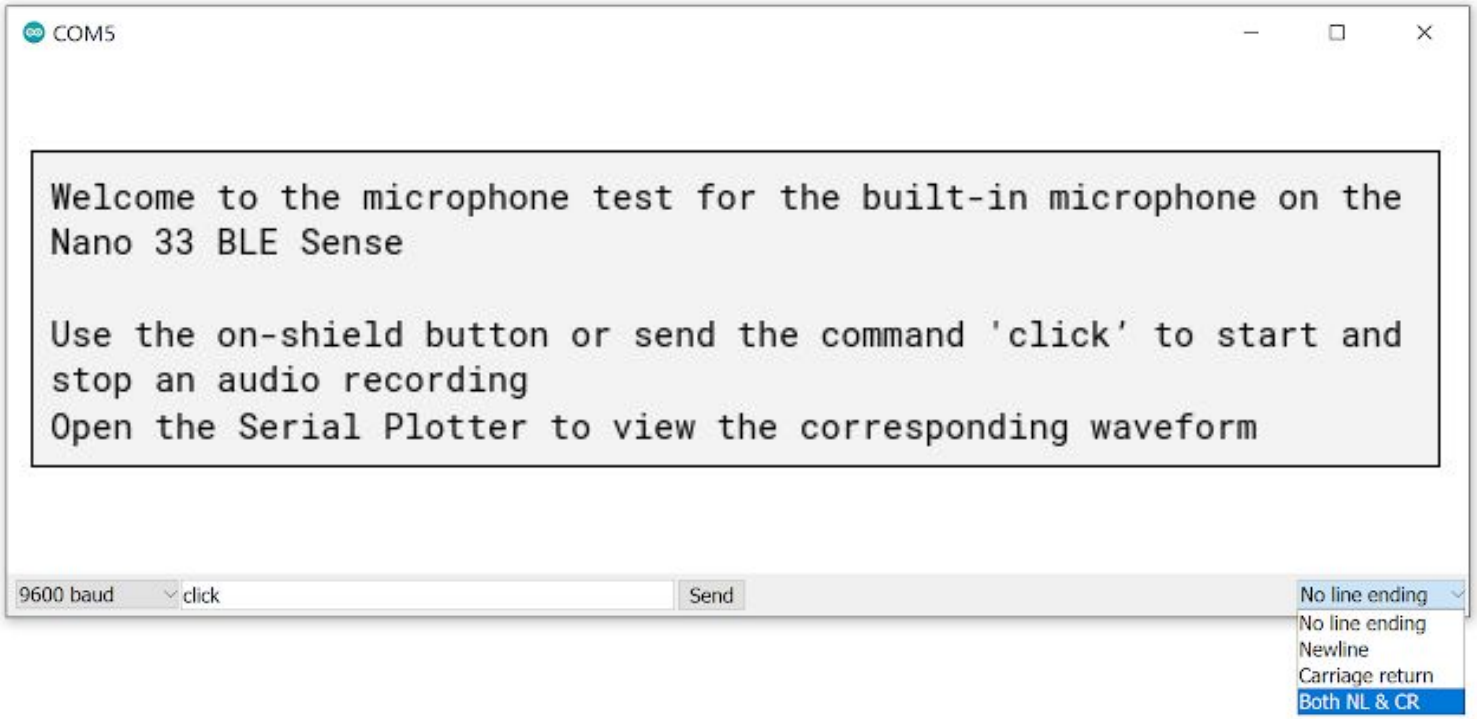

Tools → Serial Monitor

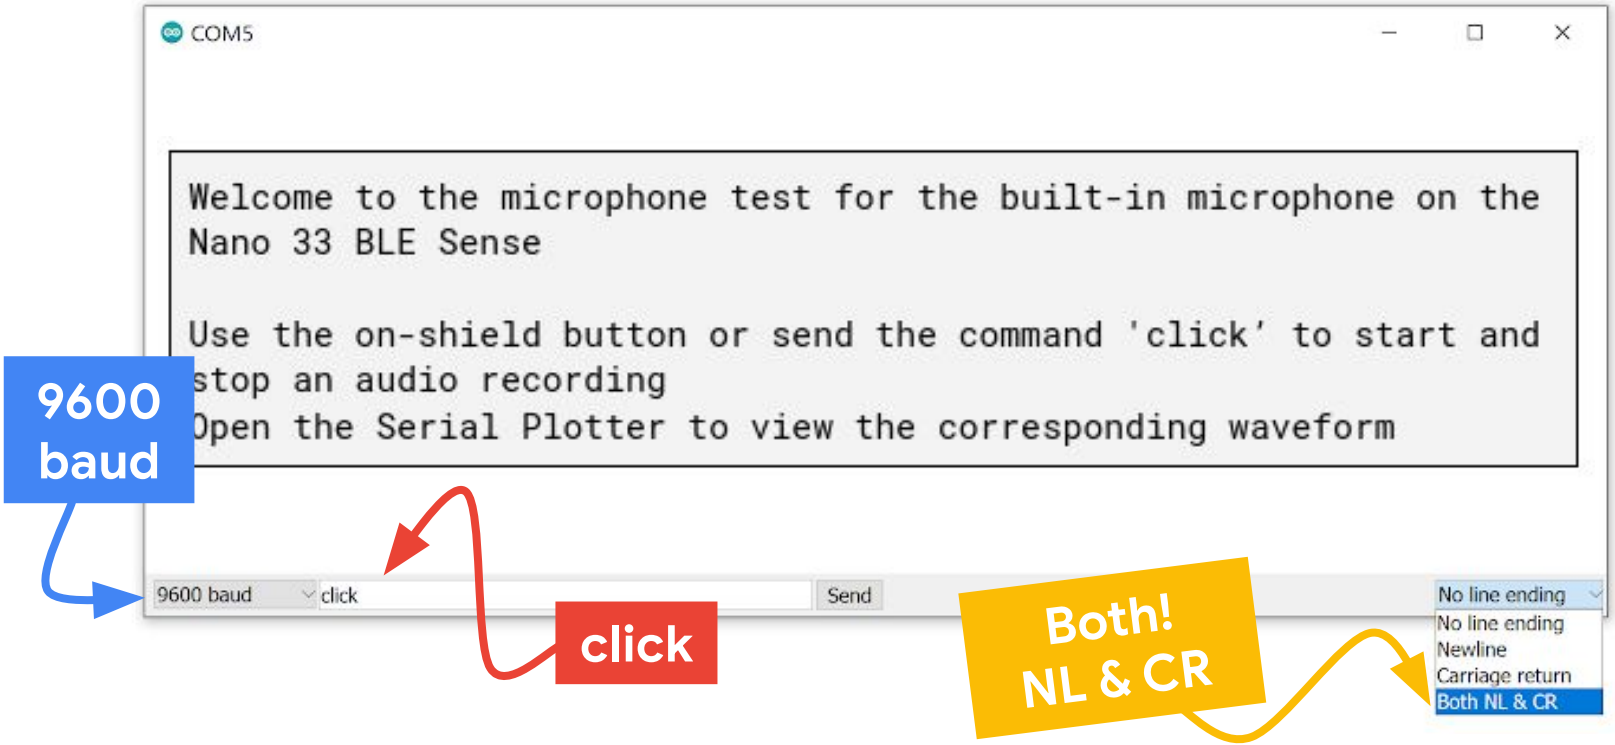

Tools → Serial Monitor

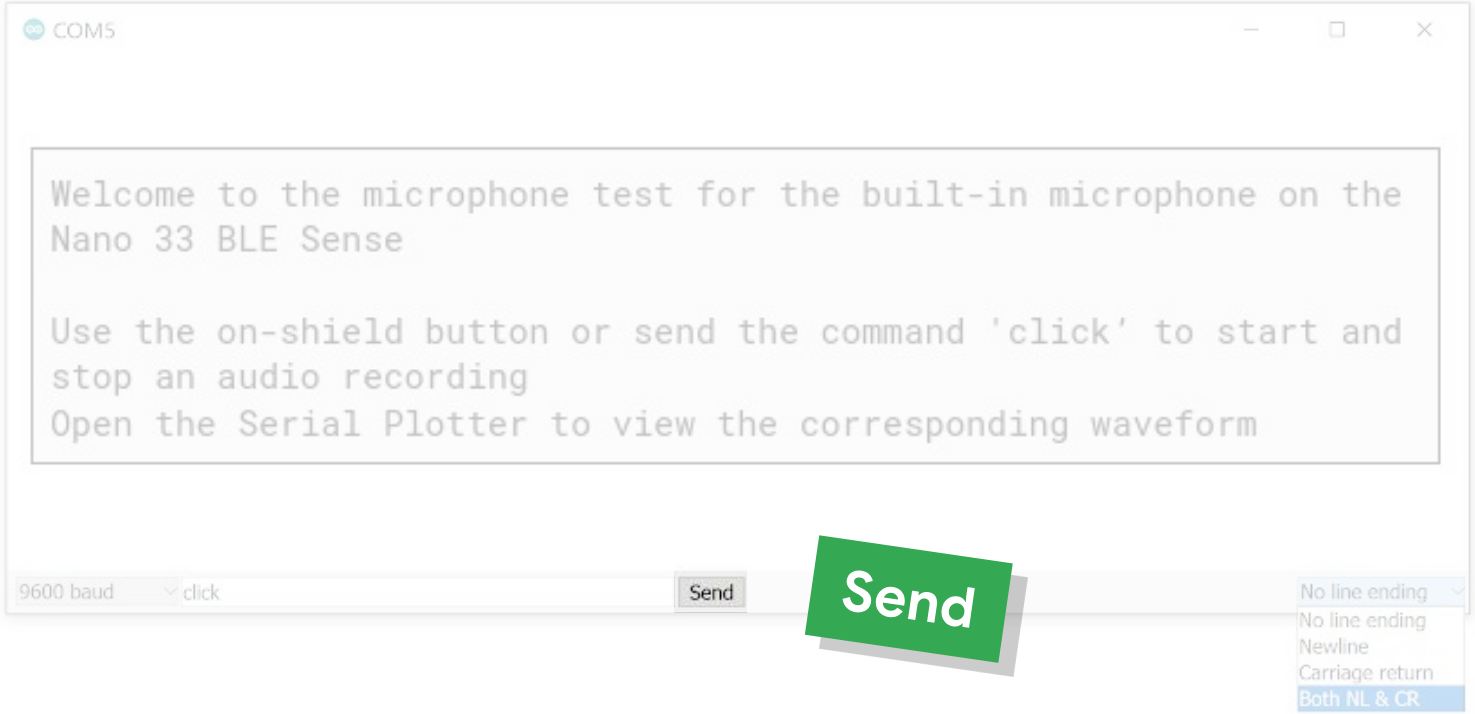

Tools → Serial Monitor

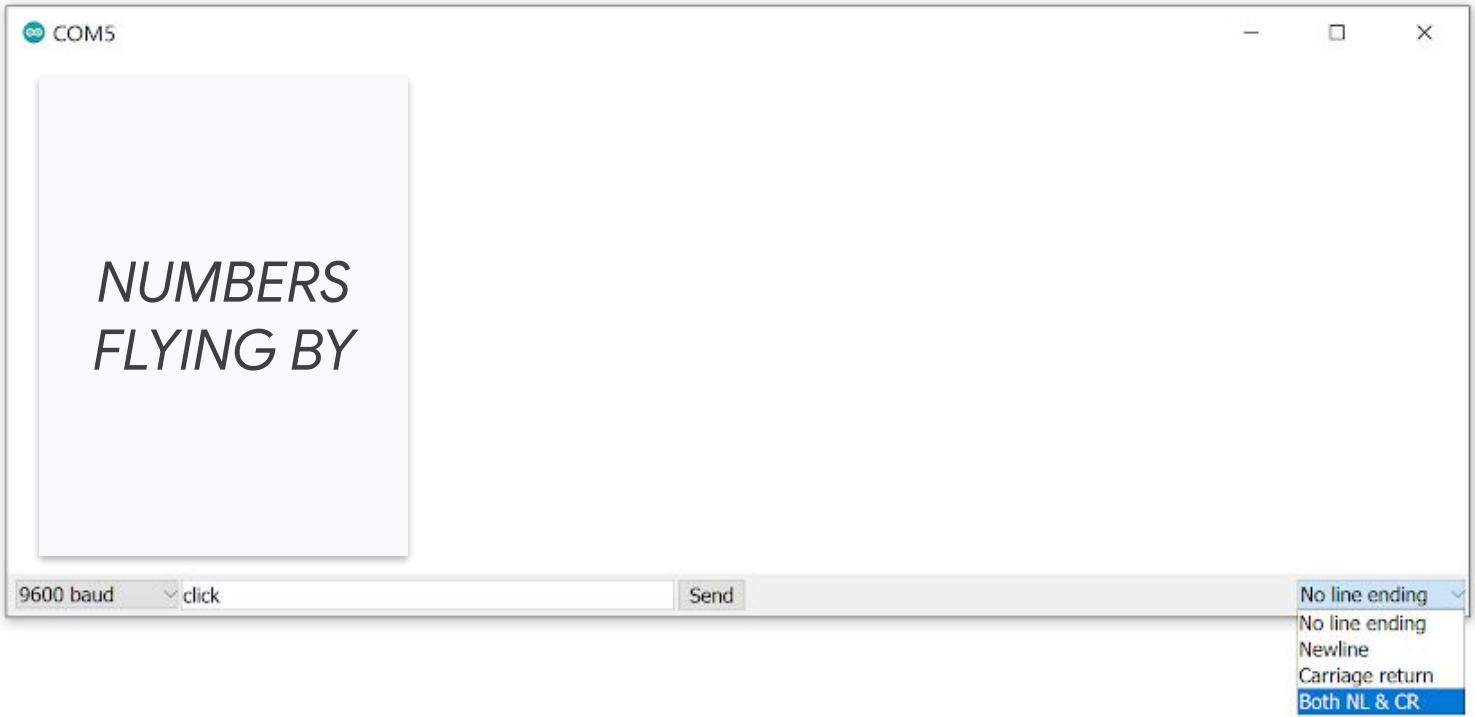

Tools → Serial *Plotter*

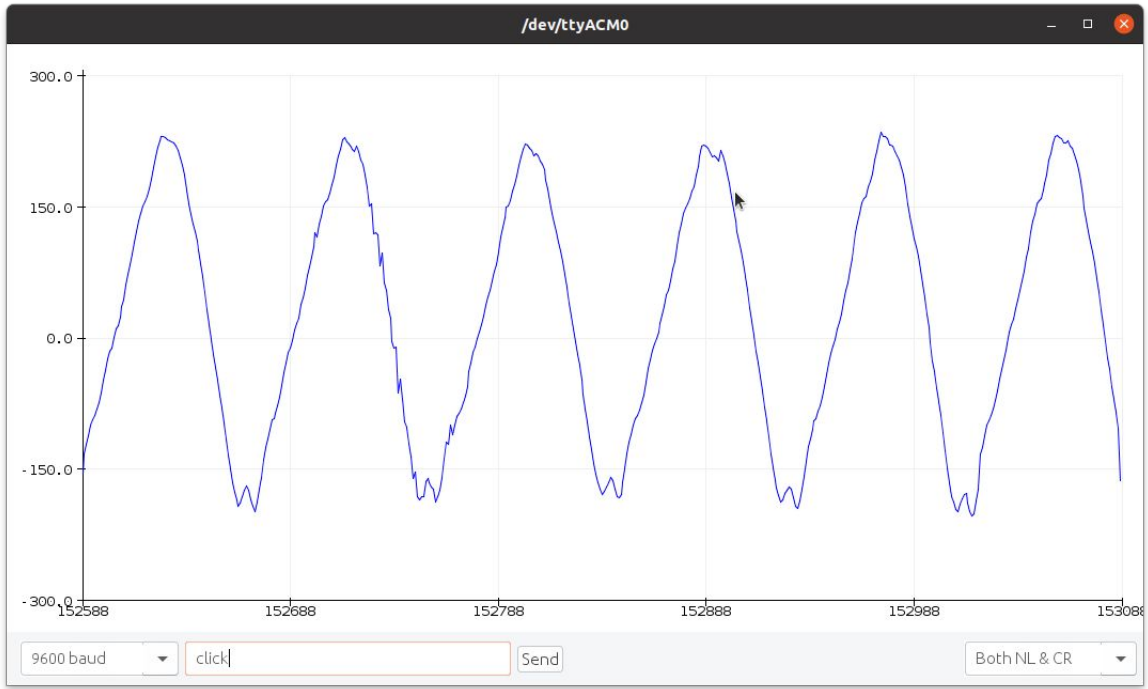

Tools → Serial *Plotter*

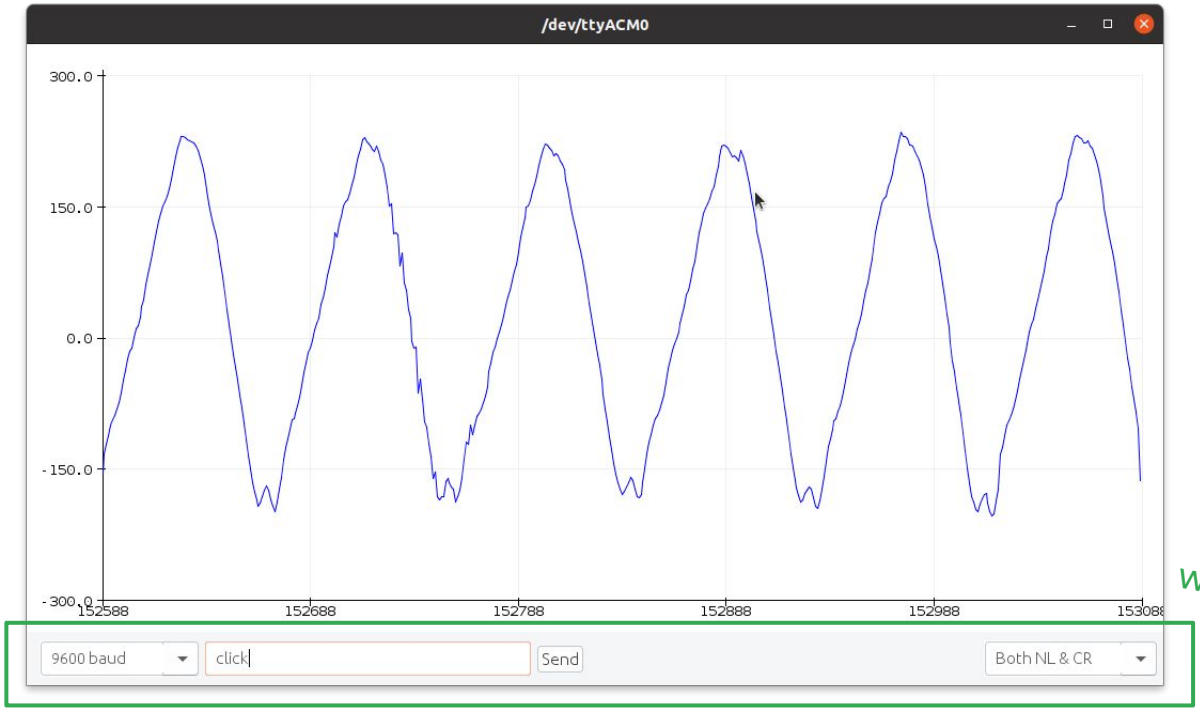

How clean of a wave can you get?

Today's Agenda

- Review + Why Tiny?
- Hardware Basics
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- **Deploying KWS model onto Arduino**
- Summary & Next Steps

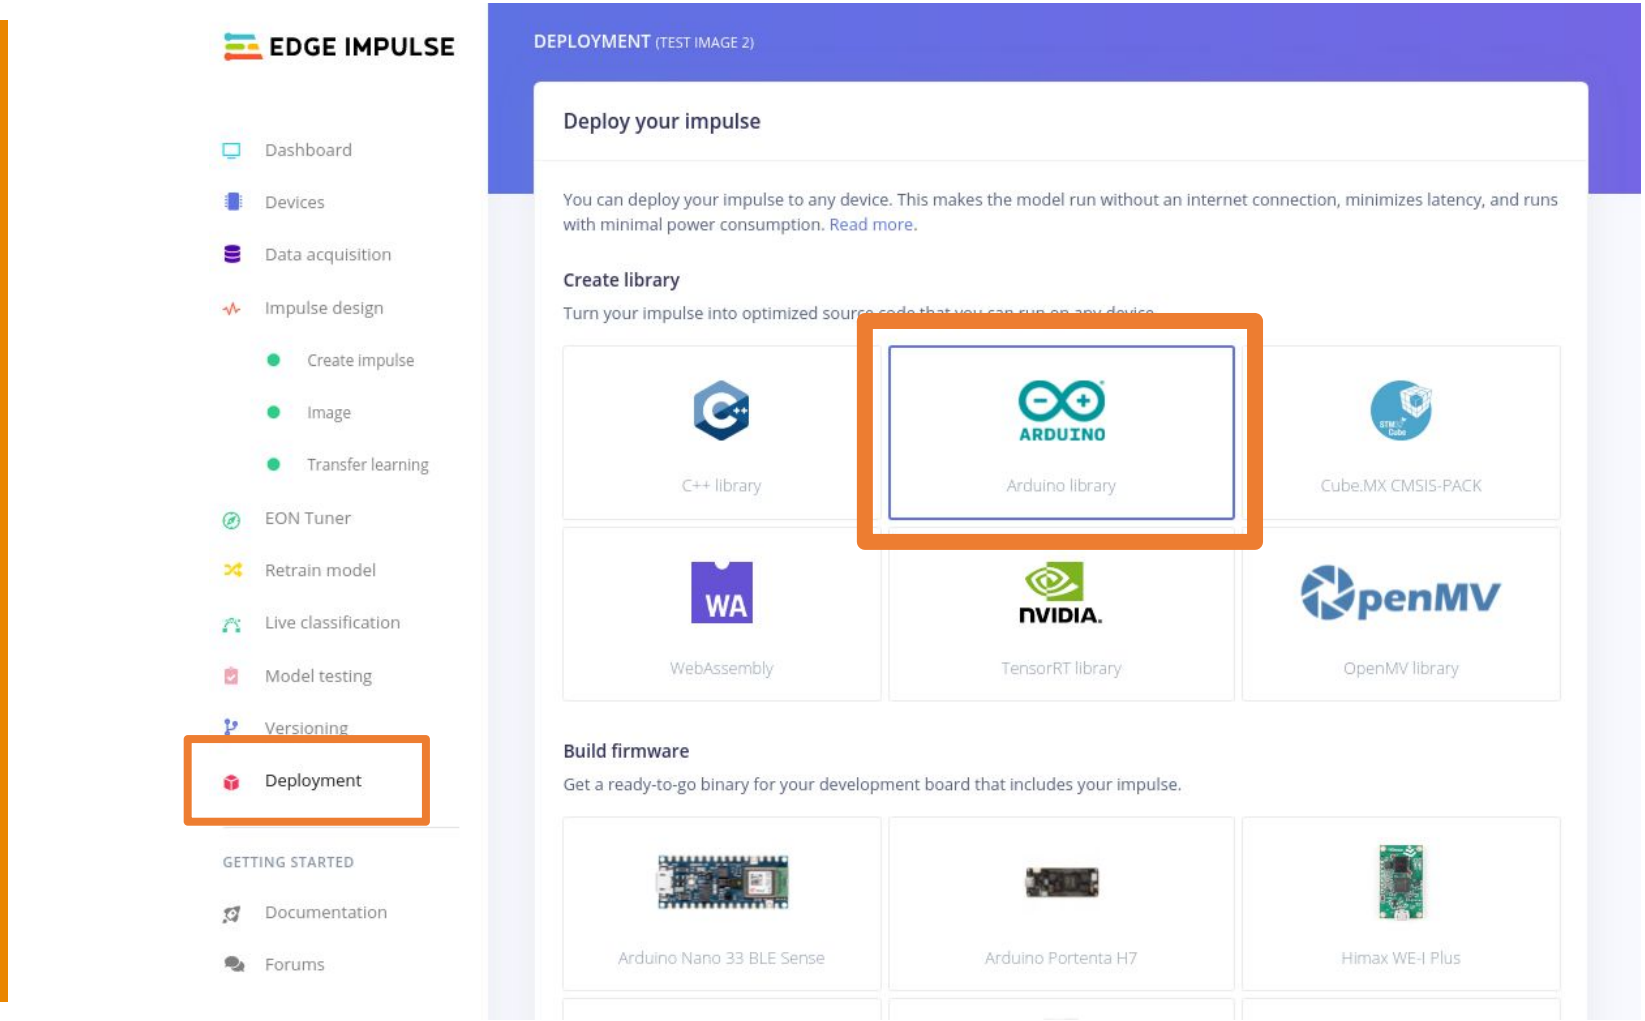

EDGE IMPULSE

- 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	LATENCY
66.1K	58 ms
FLASH USAGE	ACCURACY
108.1K	-

Unoptimized (float32)

Click to select

RAM USAGE	LATENCY
155.6K	43 ms
FLASH USAGE	ACCURACY
193.8K	-

Analyze optimizations

Estimate for Arduino Portenta H7 (Cortex-M7 480MHz)

Build

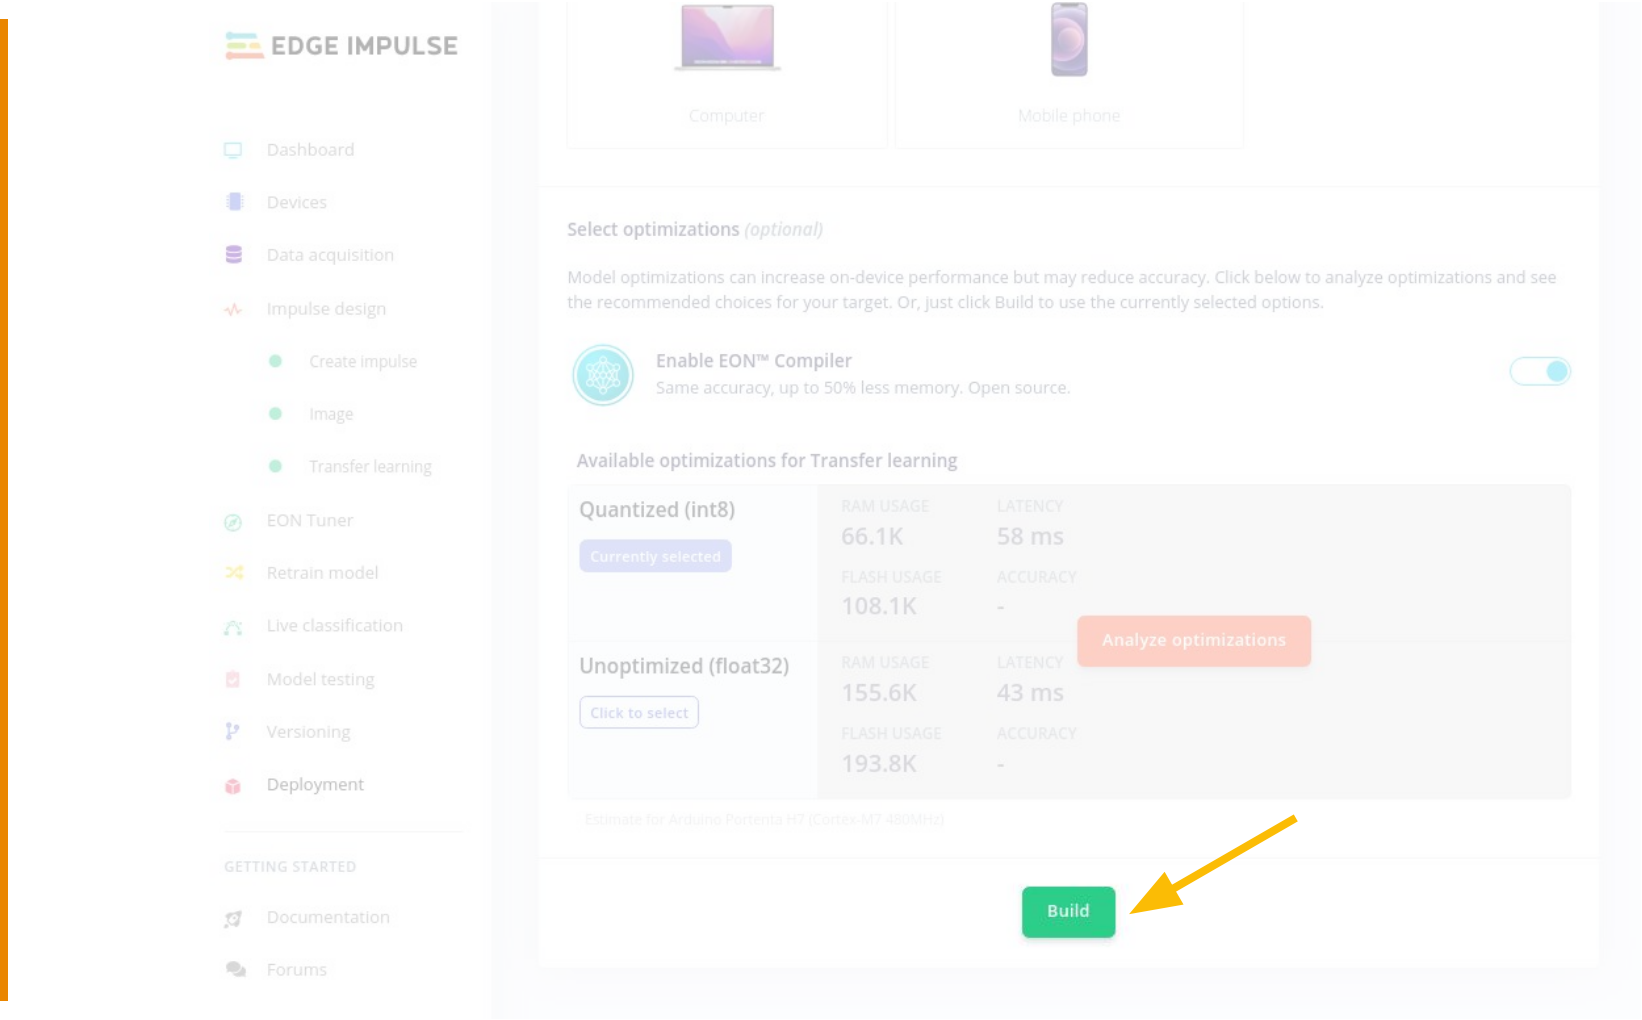

Quantization

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

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

```
max: 3.40282e+38
min: 1.17549e-38
```

Reducing the Precision

4 bytes per model parameter

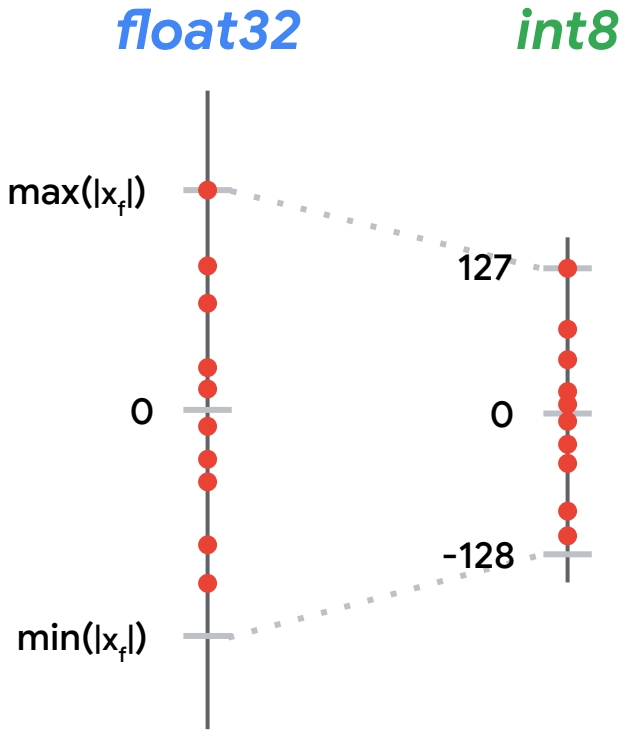

1 byte per model parameter

Tradeoff

	Floating-point Baseline	After Quantization	Accuracy Drop
MobileNet v1 1.0 224	71.03%	69.57%	▼1.46%
MobileNet v2 1.0 224	70.77%	70.20%	▼0.57%
Resnet v1 50	76.30%	75.95%	▼0.35%

float32

Model — Model version: Unoptimized (float32)

Last training performance (validation set)

ACCURACY	LOSS
100.0%	0.25

Confusion matrix (validation set)

	CAR	TRUCK
CAR	100%	0%
TRUCK	0%	100%
F1 SCORE	1.00	1.00

Feature explorer (full training set)

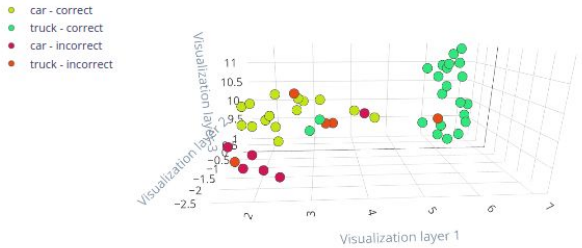

On-device performance

INFERENCING TIME	PEAK RAM USAGE	FLASH USAGE
43 ms.	155.6K	193.8K

int8

Model — Model version: Quantized (int8)

Last training performance (validation set)

ACCURACY	LOSS
70.0%	0.33

Confusion matrix (validation set)

	CAR	TRUCK
CAR	60%	40%
TRUCK	20%	80%
F1 SCORE	0.67	0.73

Feature explorer (full training set)

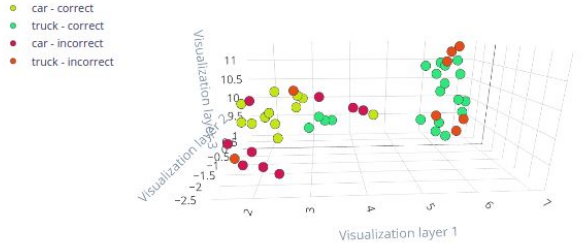

On-device performance

INFERENCING TIME	PEAK RAM USAGE	FLASH USAGE
58 ms.	66.1K	108.1K

float32

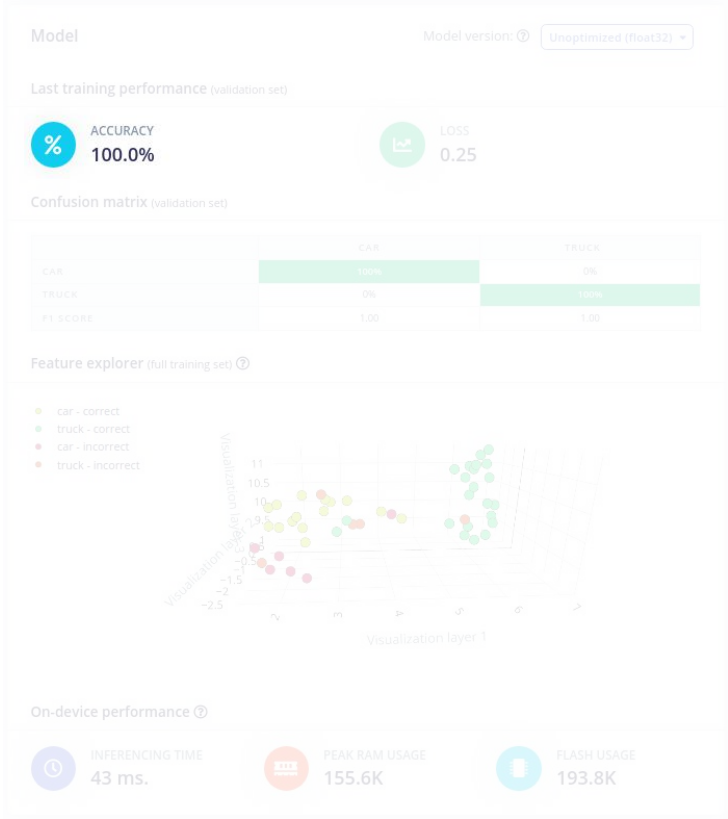

int8

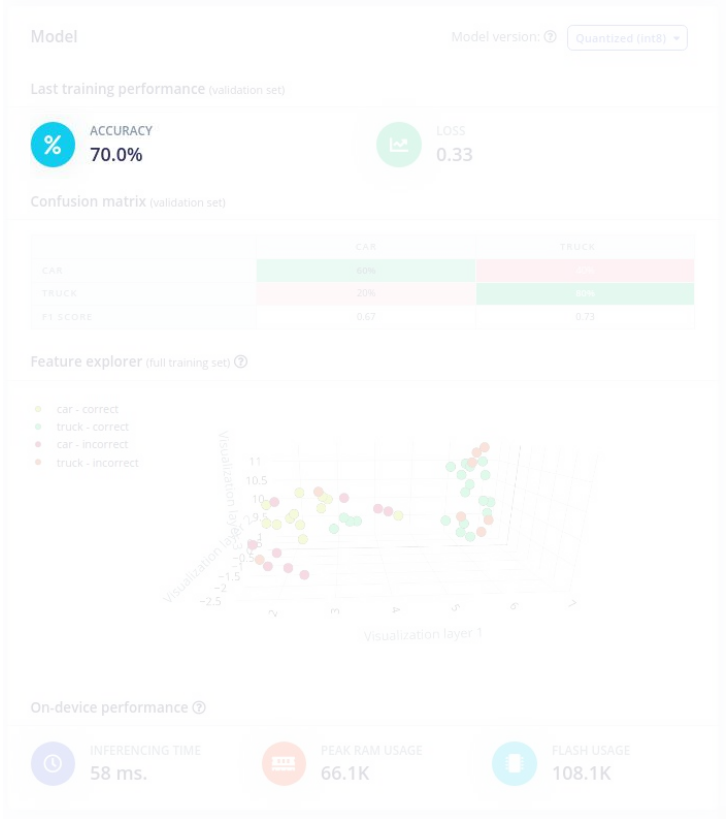

float32

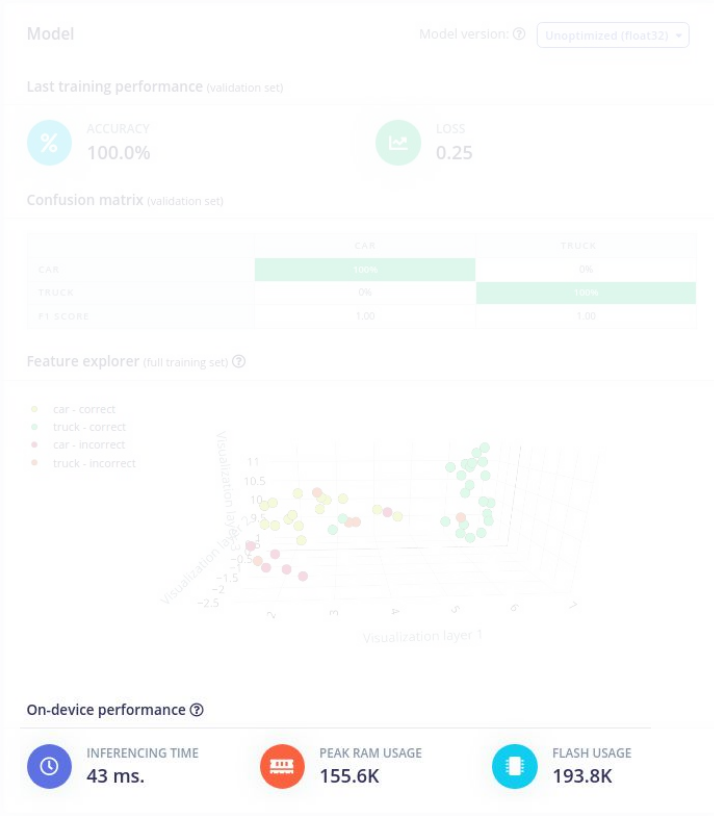

int8

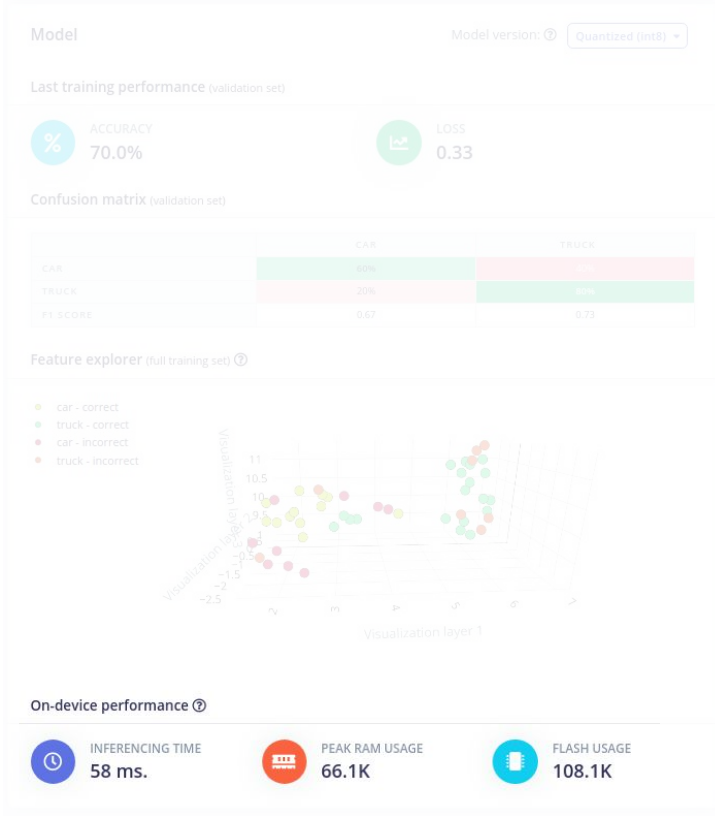

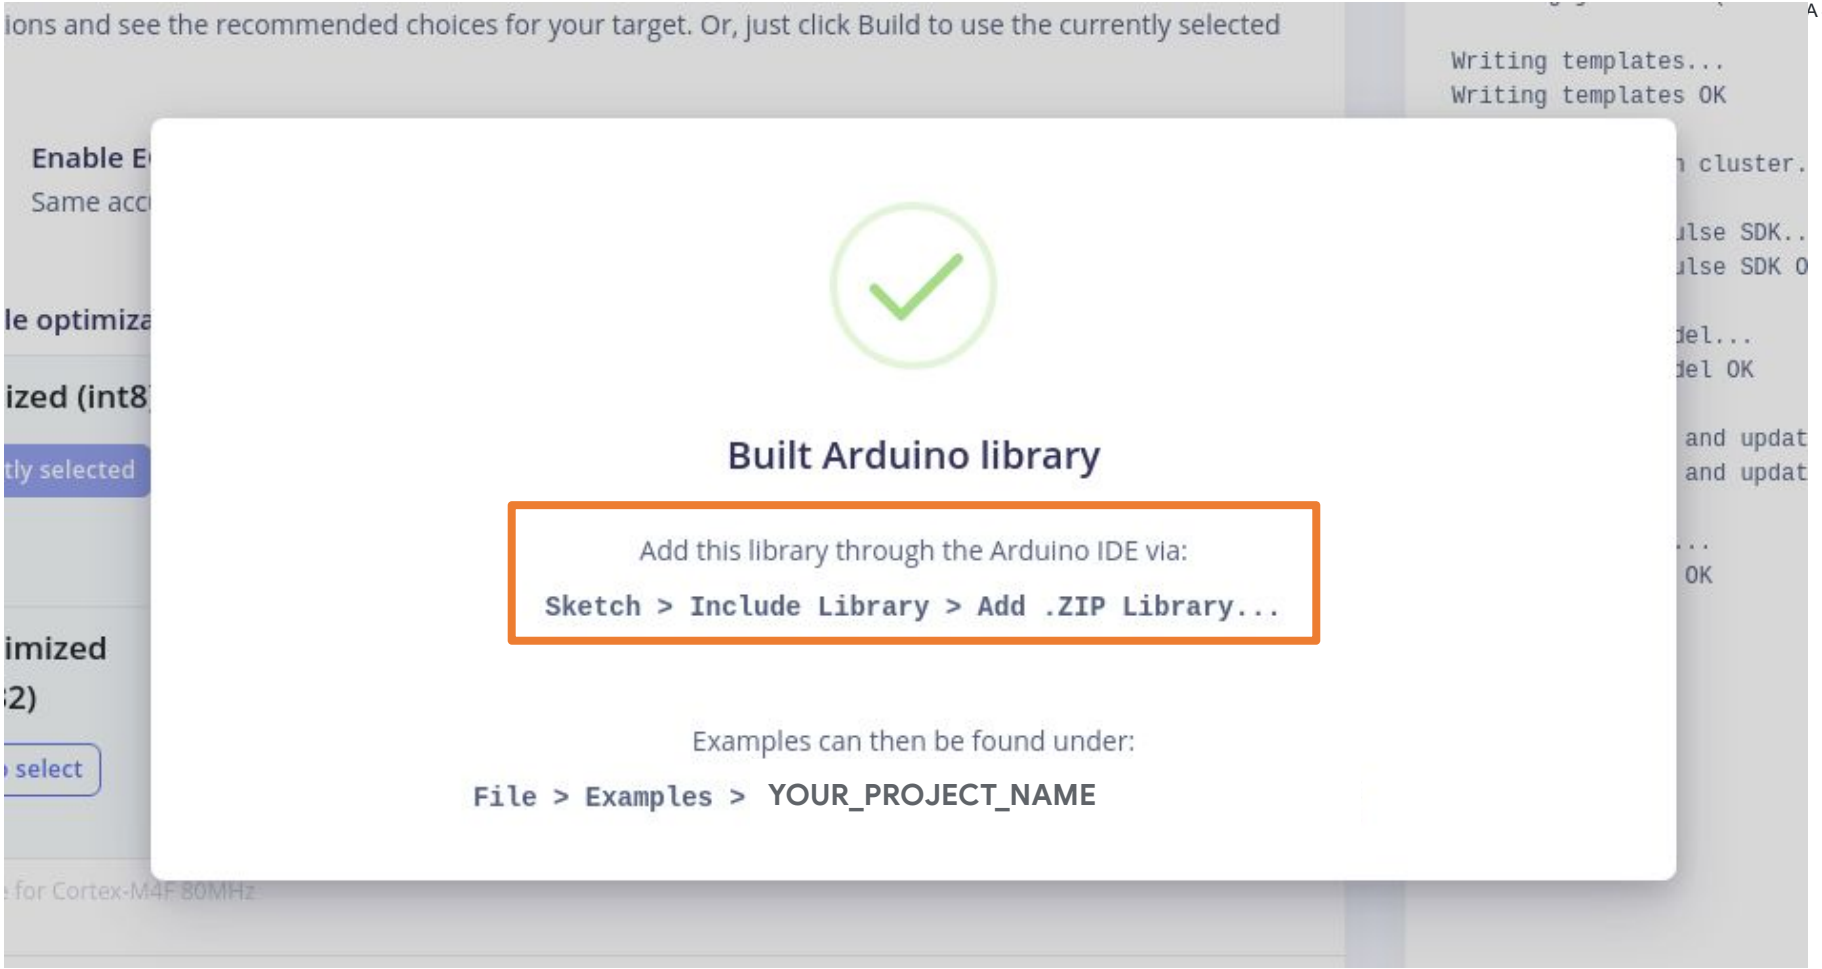

Sketch → **Include Library**
→ **Add .ZIP Library**

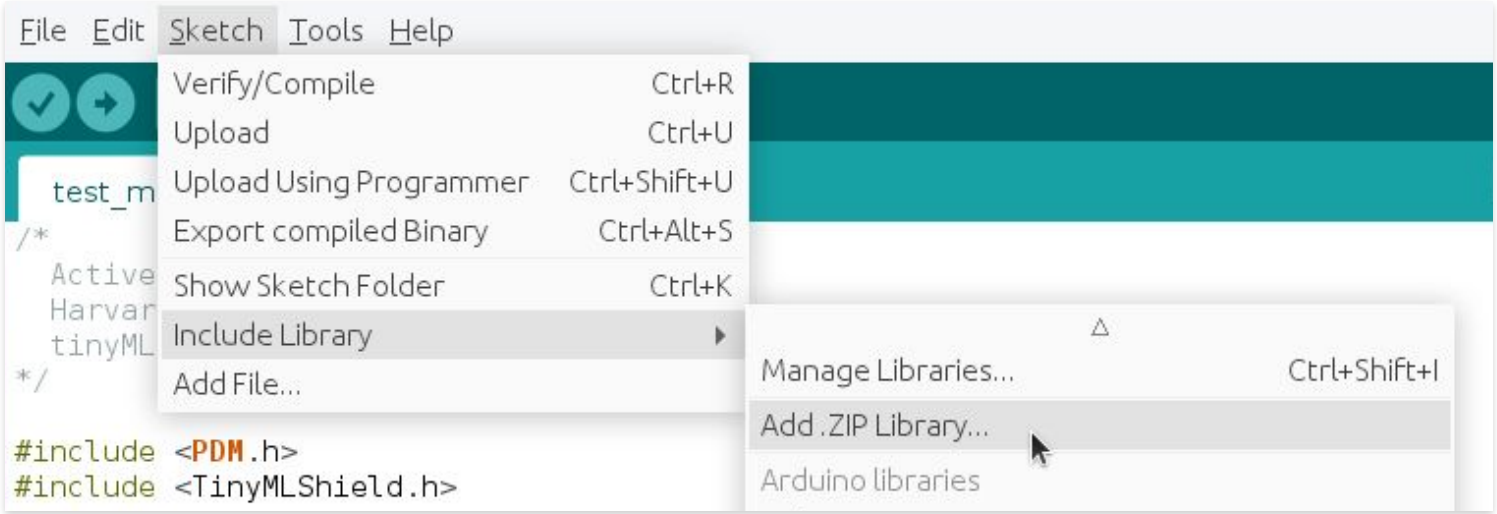

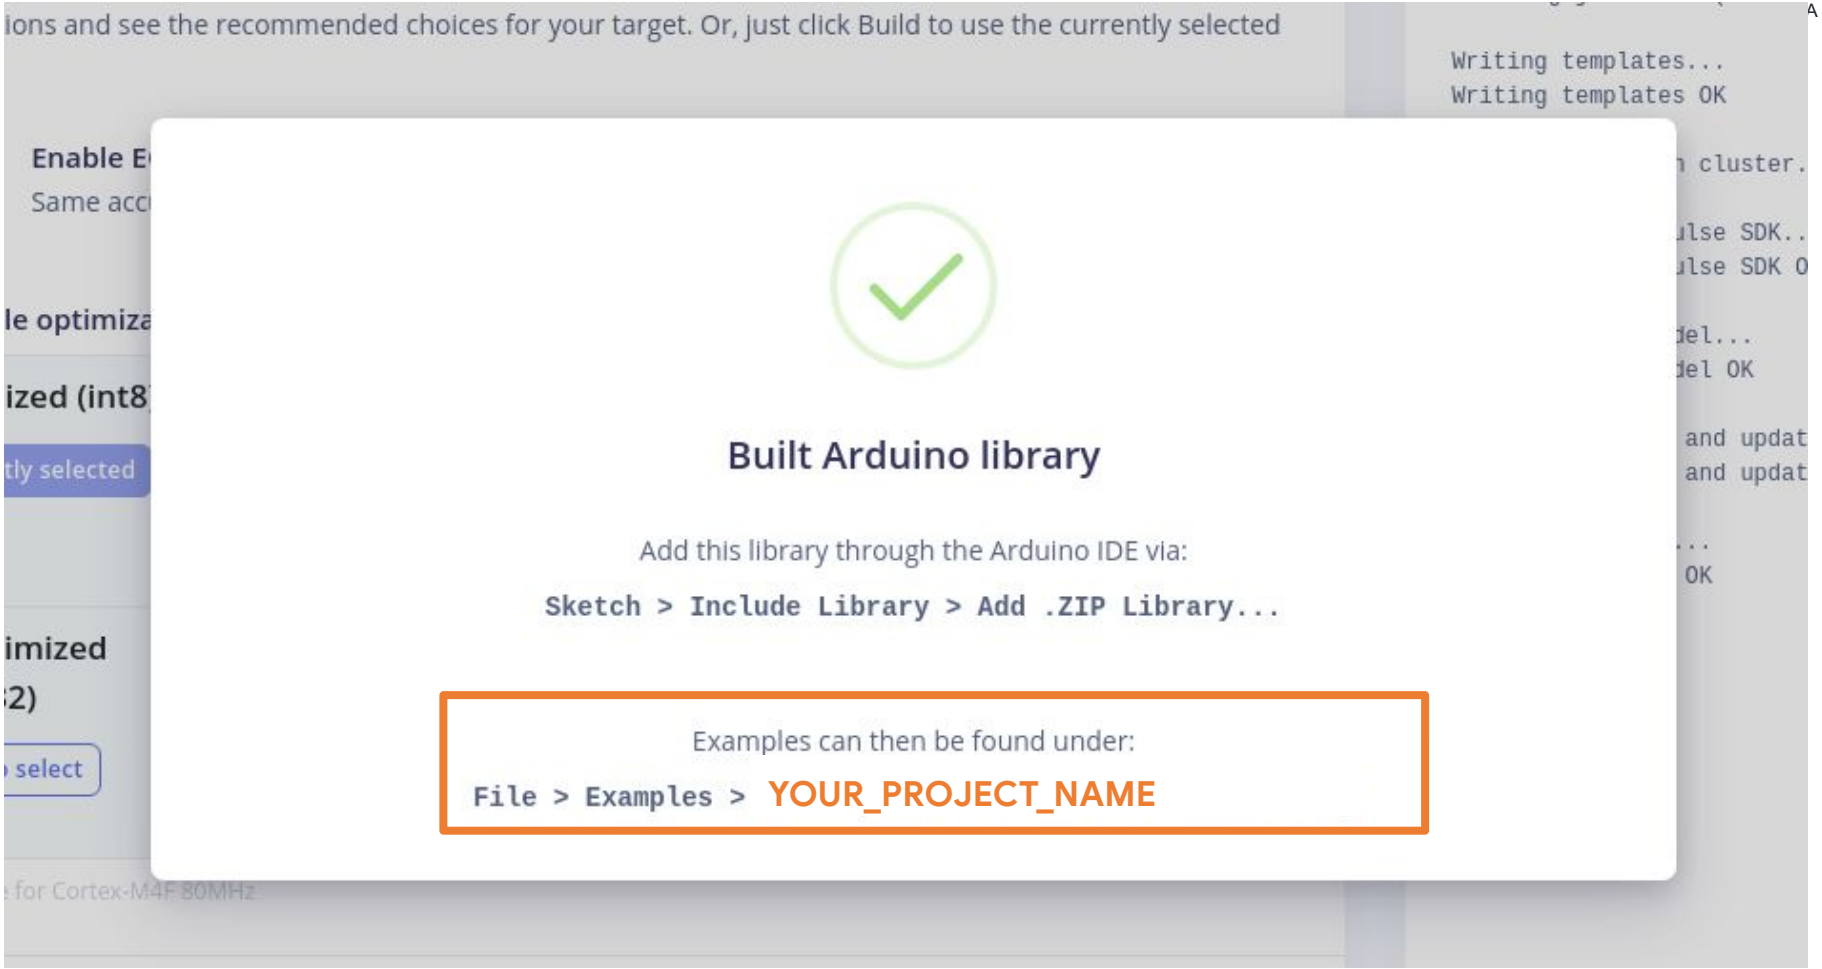

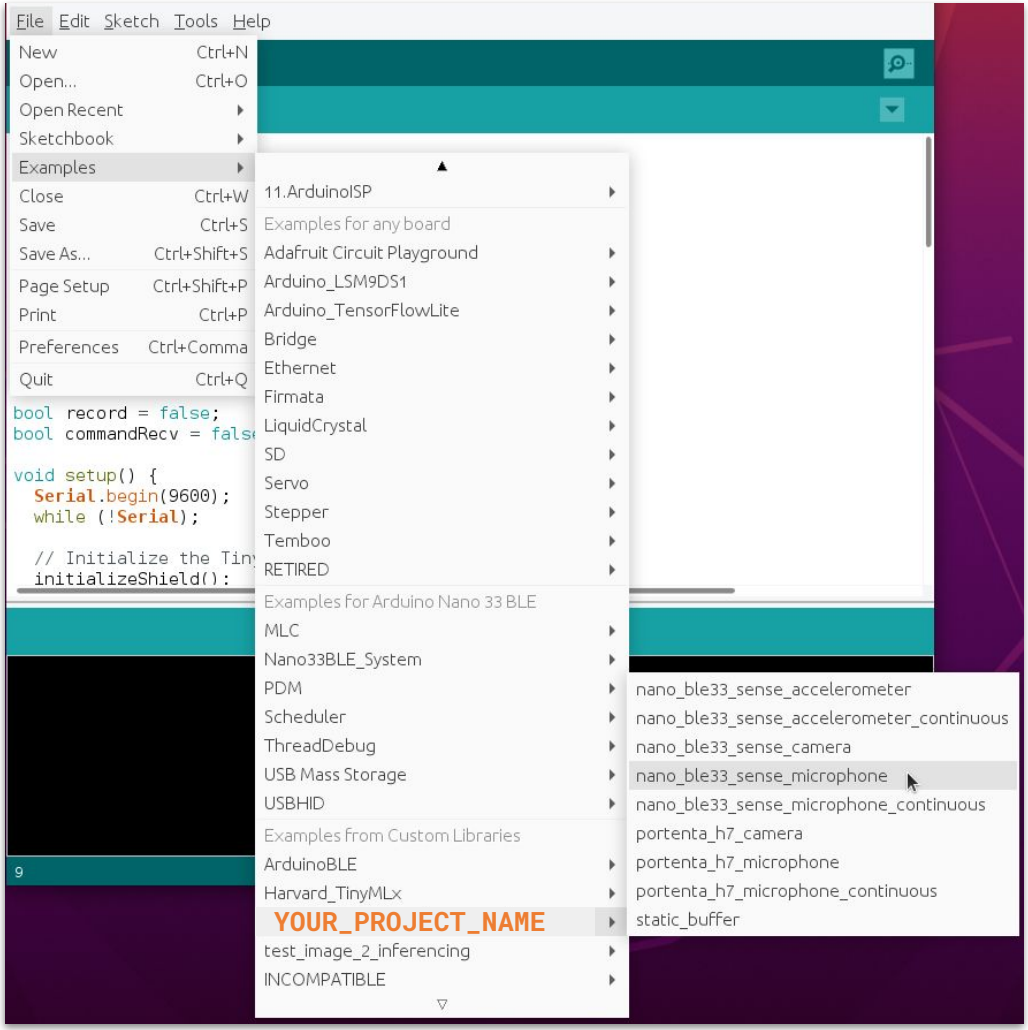

File → Examples →

YOUR_PROJECT_NAME →

nano_ble33_sense_microphone

upload

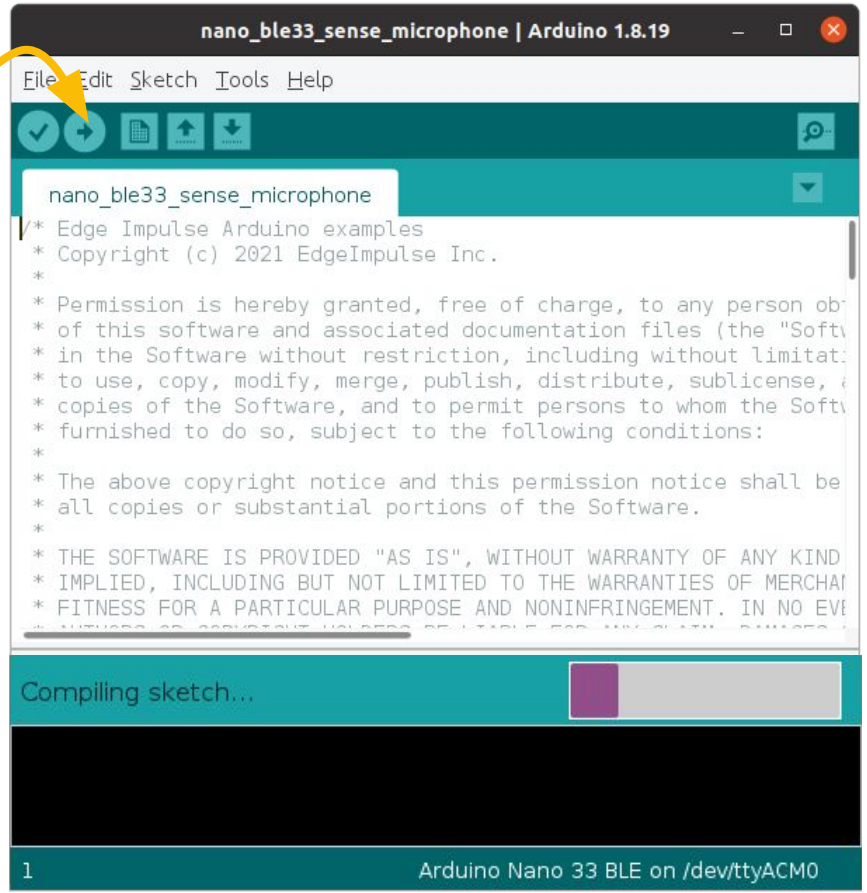

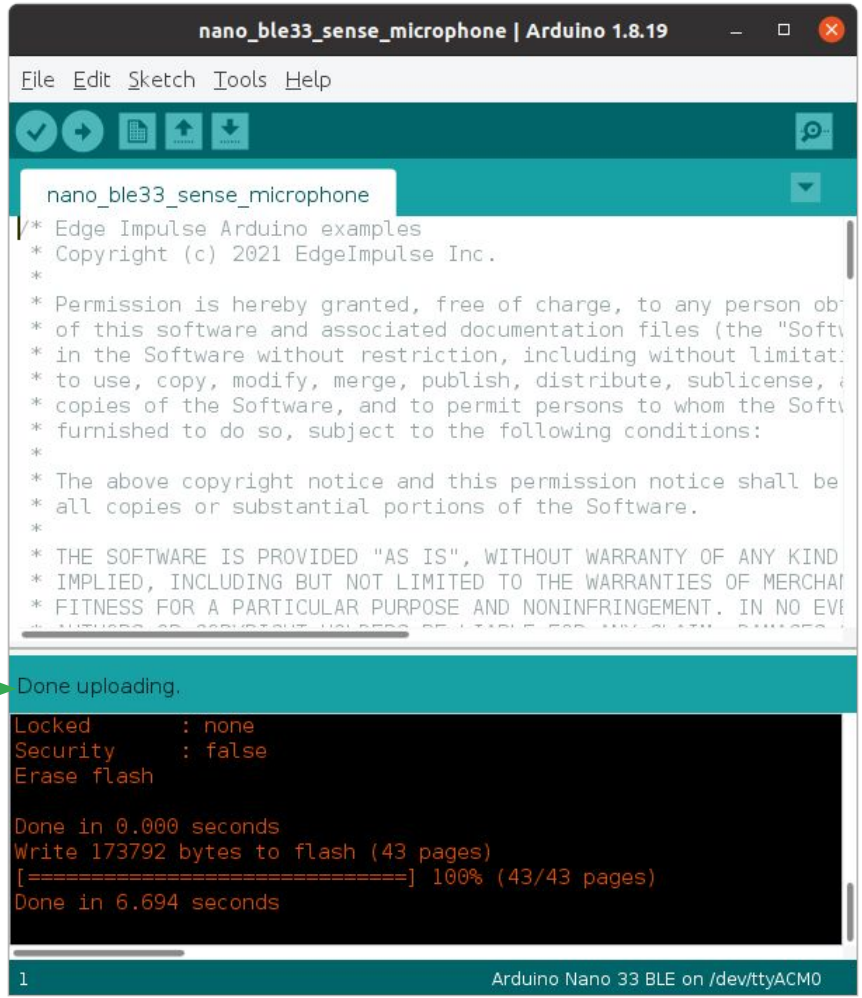

done?

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!

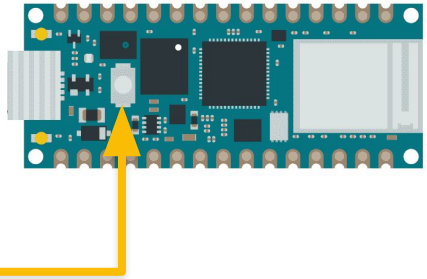

Tools → Serial Monitor

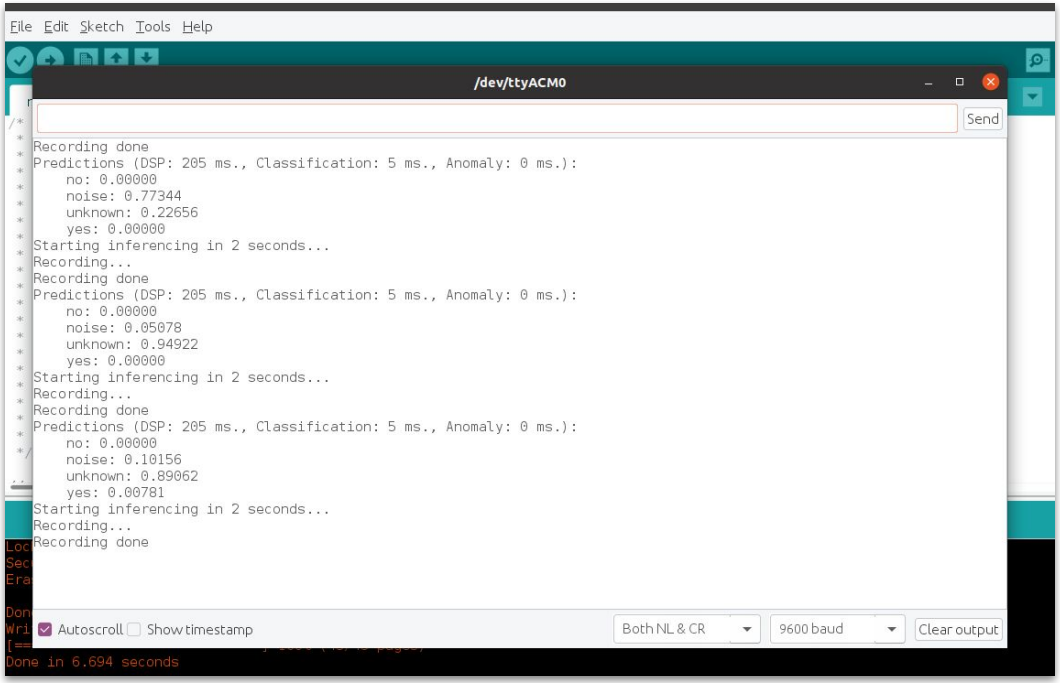

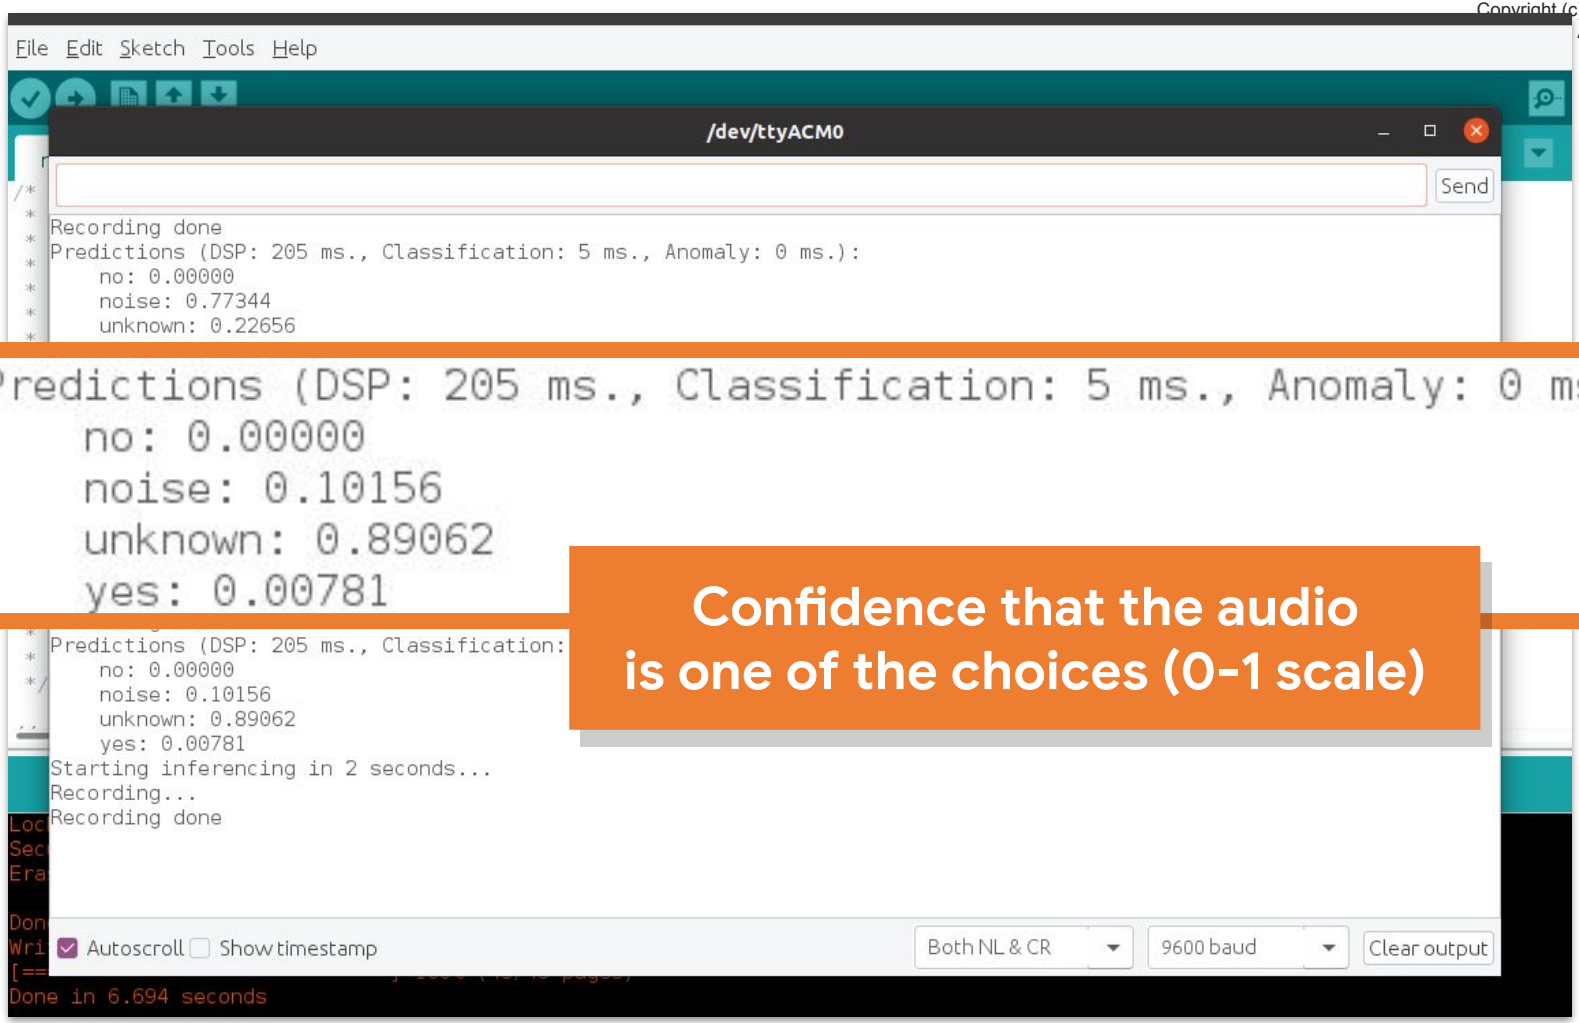

Confidence that the audio is one of the choices (0-1 scale)

Today's **Agenda**

- Review + Why Tiny?
- Hardware Basics
- Installing and Starting the Arduino IDE
- Testing Your TinyML Kit
- Deploying KWS model onto Arduino

Summary & Next Steps

Machine Learning

Deep Learning with Neural Networks

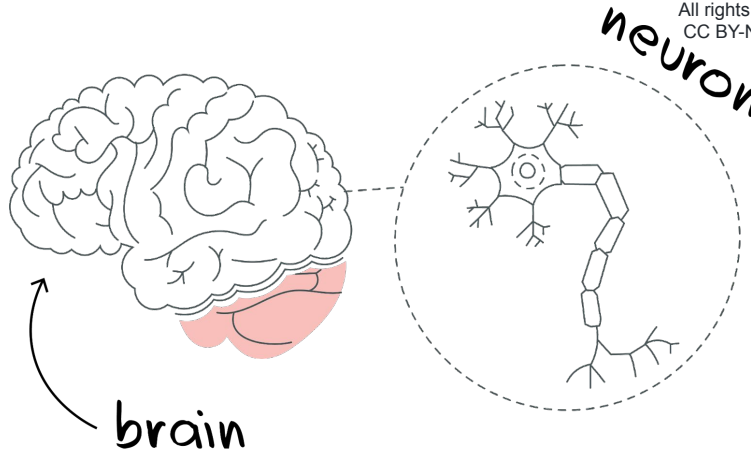

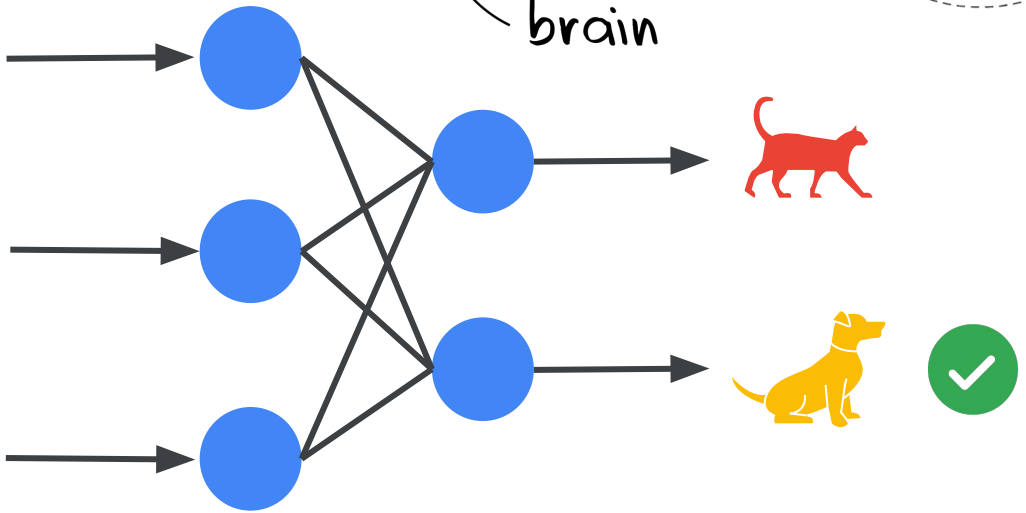

Colab Link

Features can be found with Convolutions

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

The (Tiny) Machine Learning **Workflow**

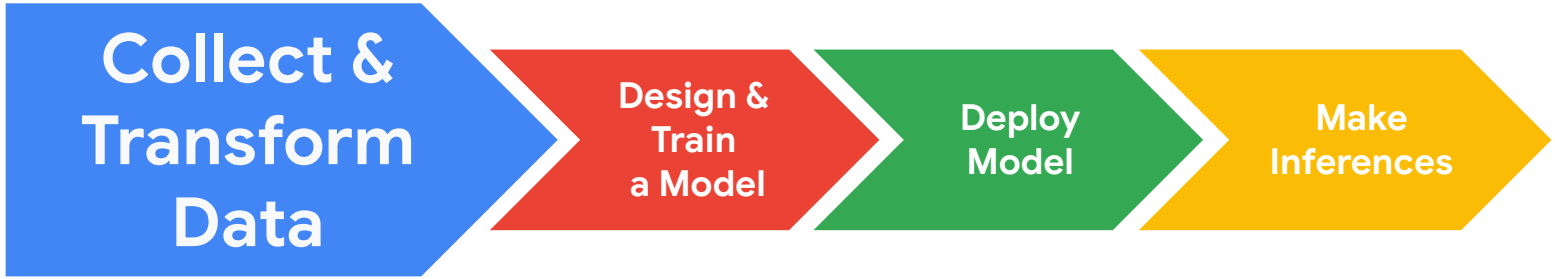

ML is *everywhere* → collect **GOOD** data **RESPONSIBLY**

The (Tiny) Machine Learning **Workflow**

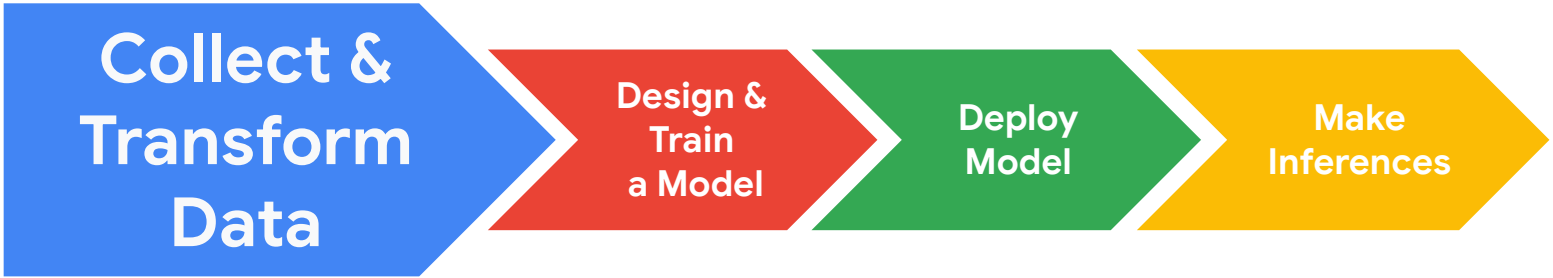

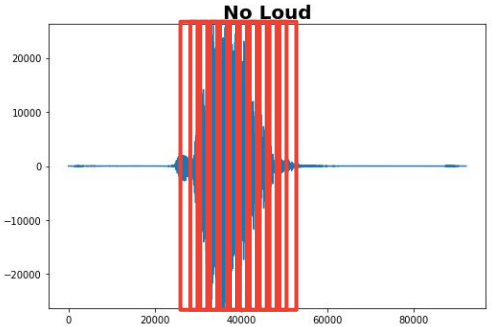

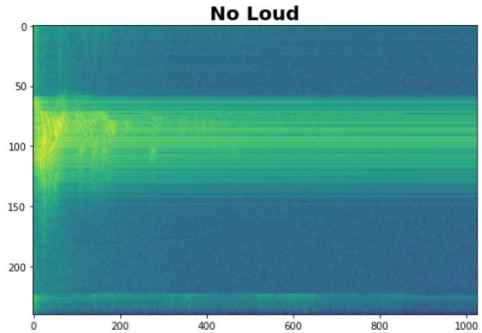

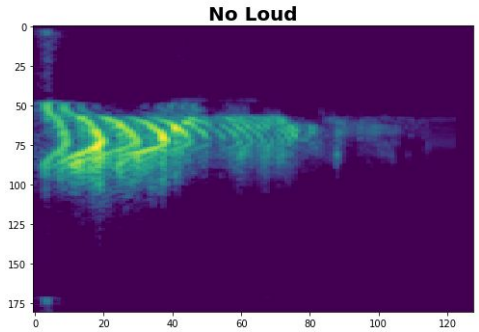

The (Tiny) Machine Learning **Workflow**

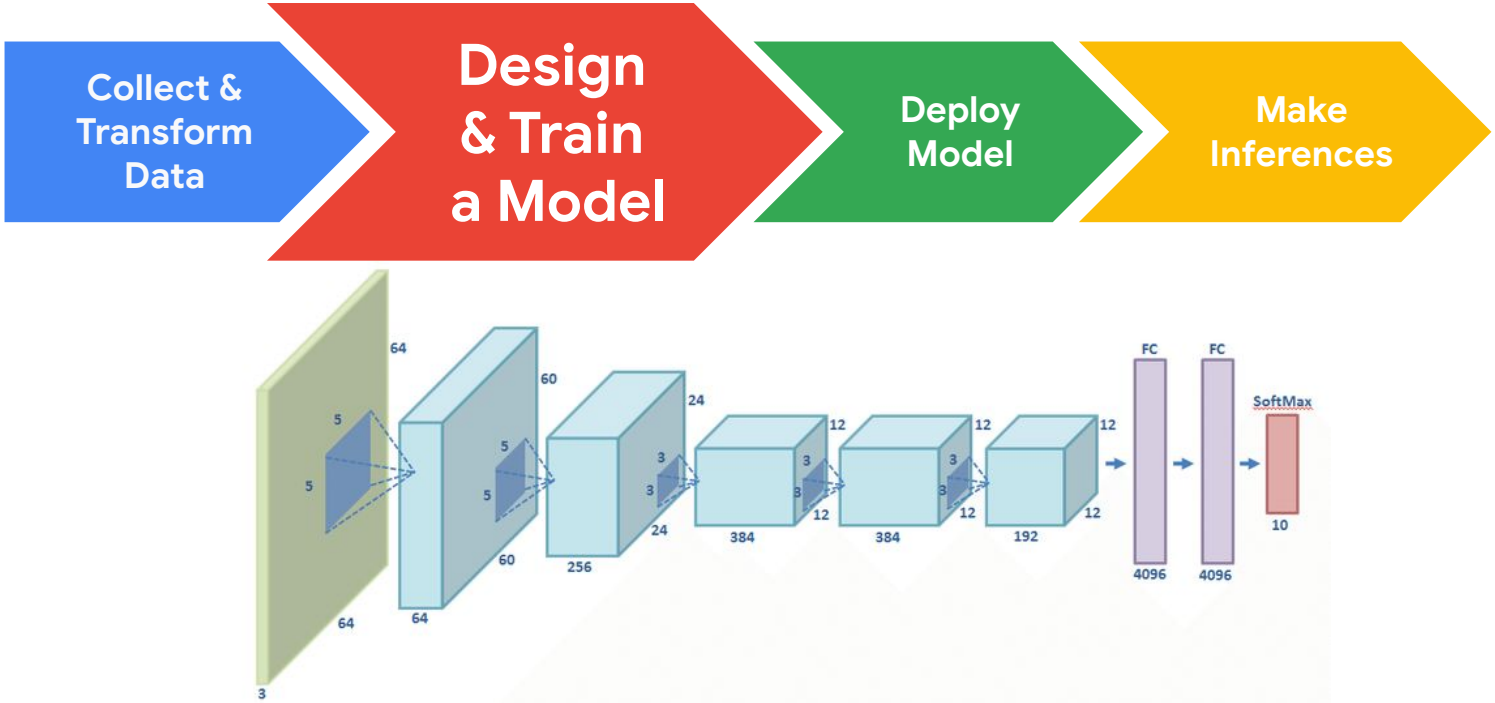

The (Tiny) Machine Learning **Workflow**

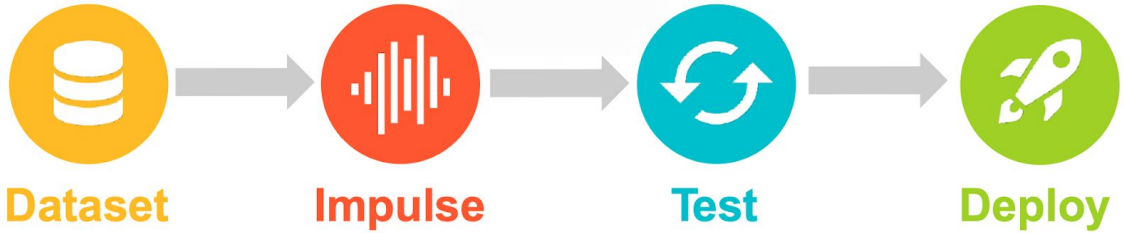

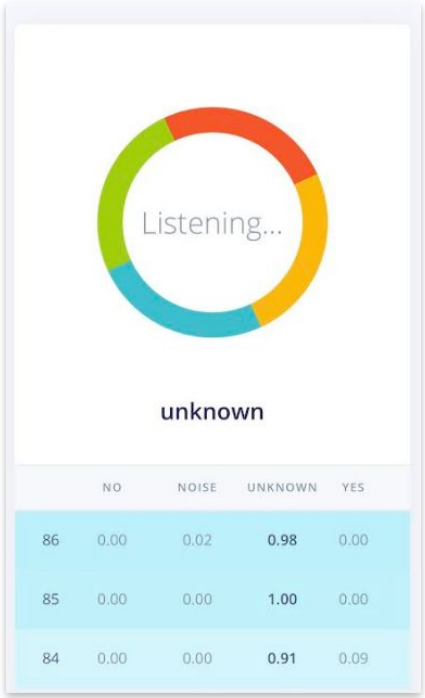

Edge Impulse Simplifies Training and Deployment

The (Tiny) Machine Learning **Workflow**

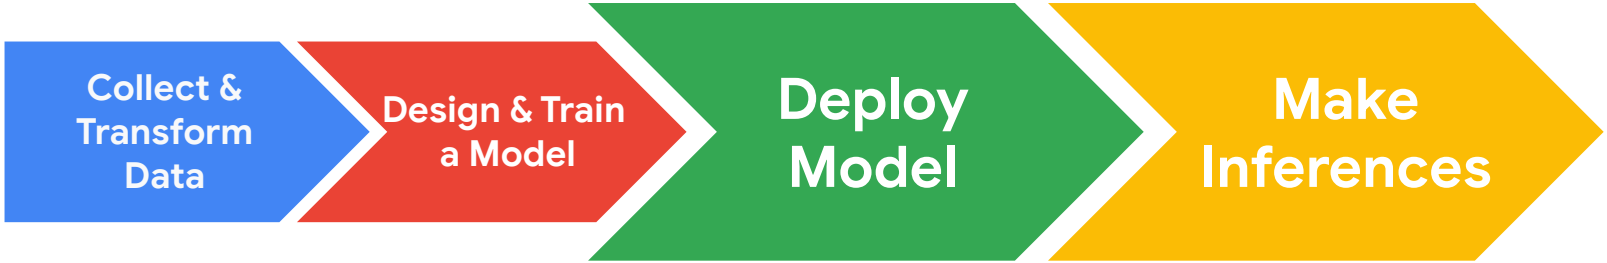

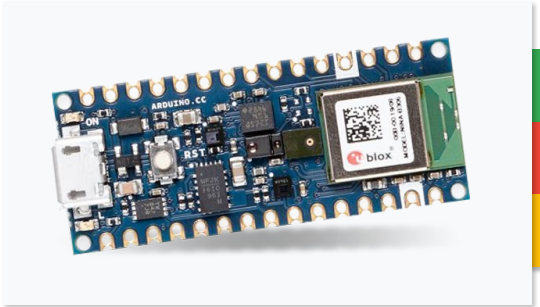

Even less memory

Even less compute power

Also, only focused on ***inference***

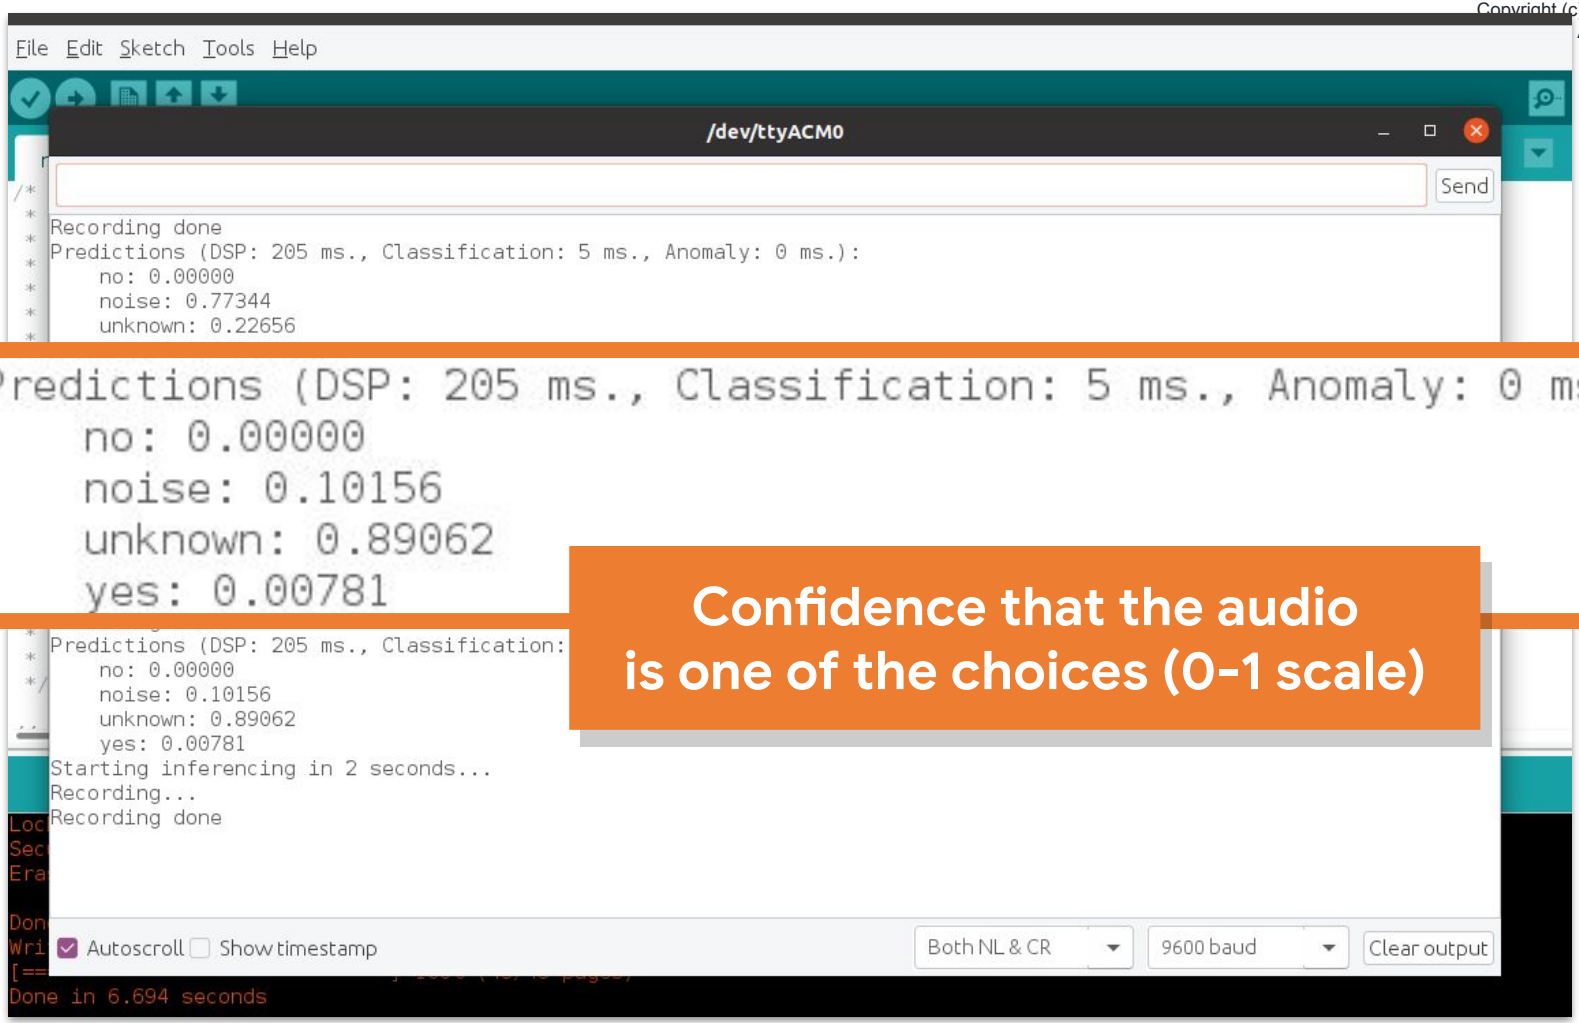

Confidence that the audio is one of the choices (0-1 scale)

Workshop **Agenda**

Day 1	Introduction to AI and (Tiny)ML	Cloud ML
Day 2	Keyword Spotting for the Navajo Language	Mobile ML
Day 3	Bringing AI/ML from the Cloud to the Edge	Embedded ML

Workshop **Agenda**

Day 1	Introduction to AI and (Tiny)ML	Cloud ML
Day 2	Keyword Spotting for the Navajo Language	Mobile ML
Day 3	Bringing AI/ML from the Cloud to the Edge	Embedded ML

Practice with more **machine learning**!

Experiments with Google

Collections ▾ Experiments Search SUBMIT EXPERIMENT

AI + MUSIC

FREDDIEMETER

by Google Research, Google Creative Lab, YouTube Music

An AI-powered singing challenge that rates how closely your singing matches the voice

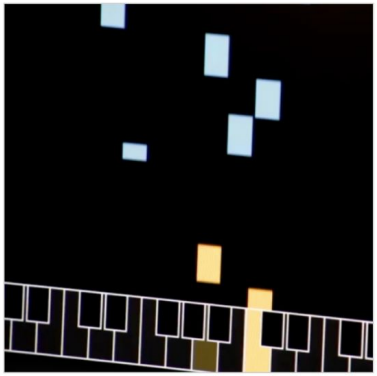

AI DUET

by Yotam Mann

A piano that responds to you.

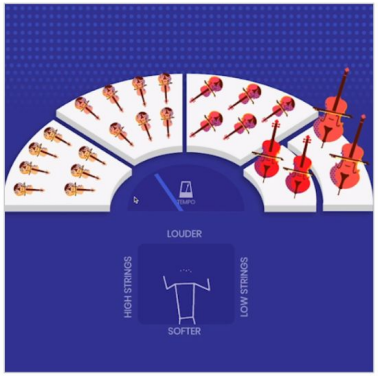

SEMI-CONDUCTOR

by Google Creative Lab

Conduct your own orchestra in the browser by moving your arms

experiments.withgoogle.com

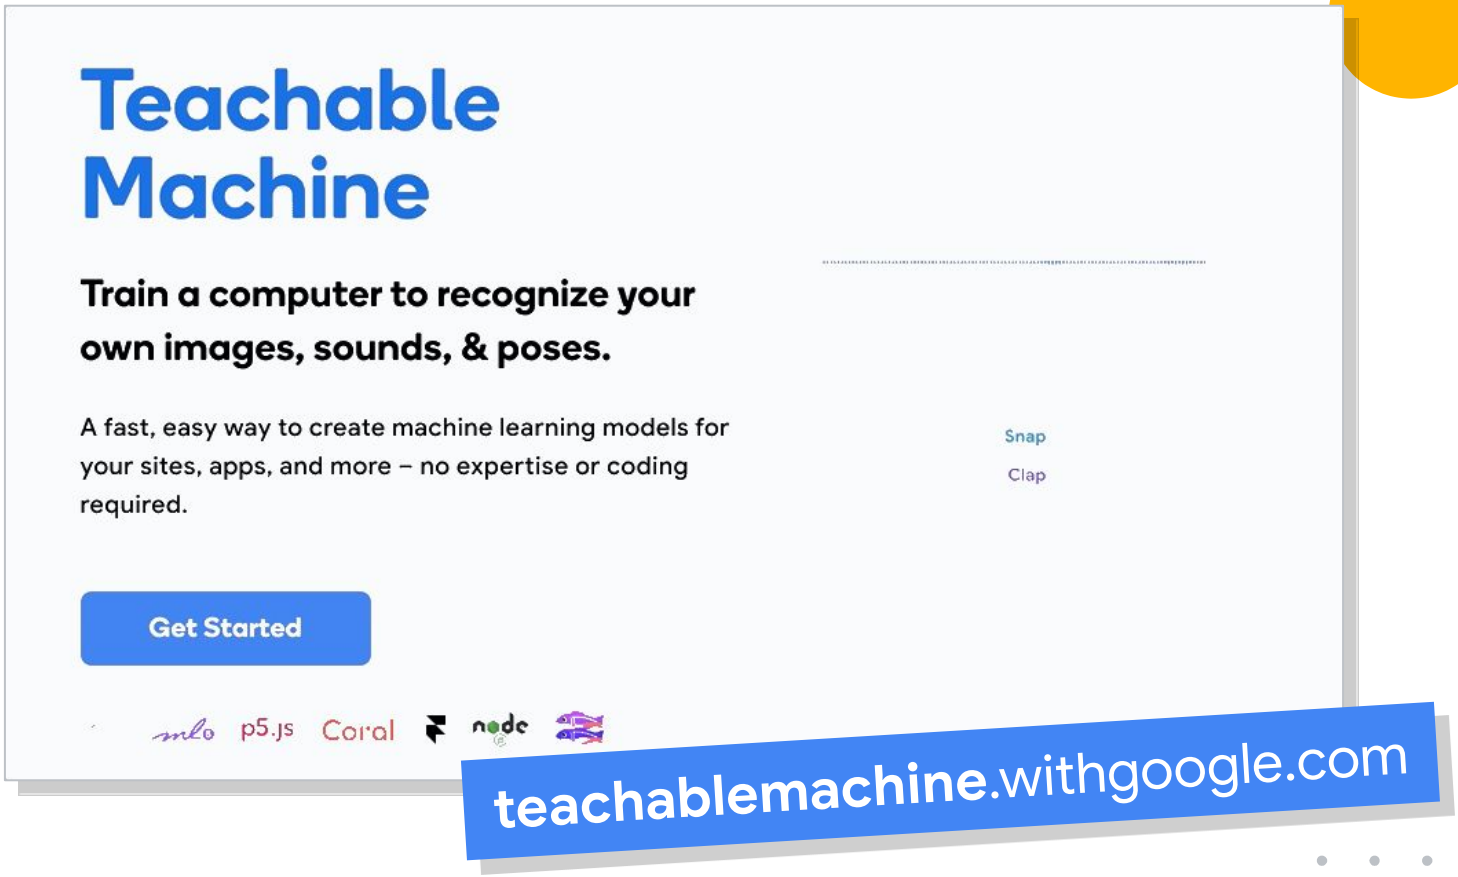

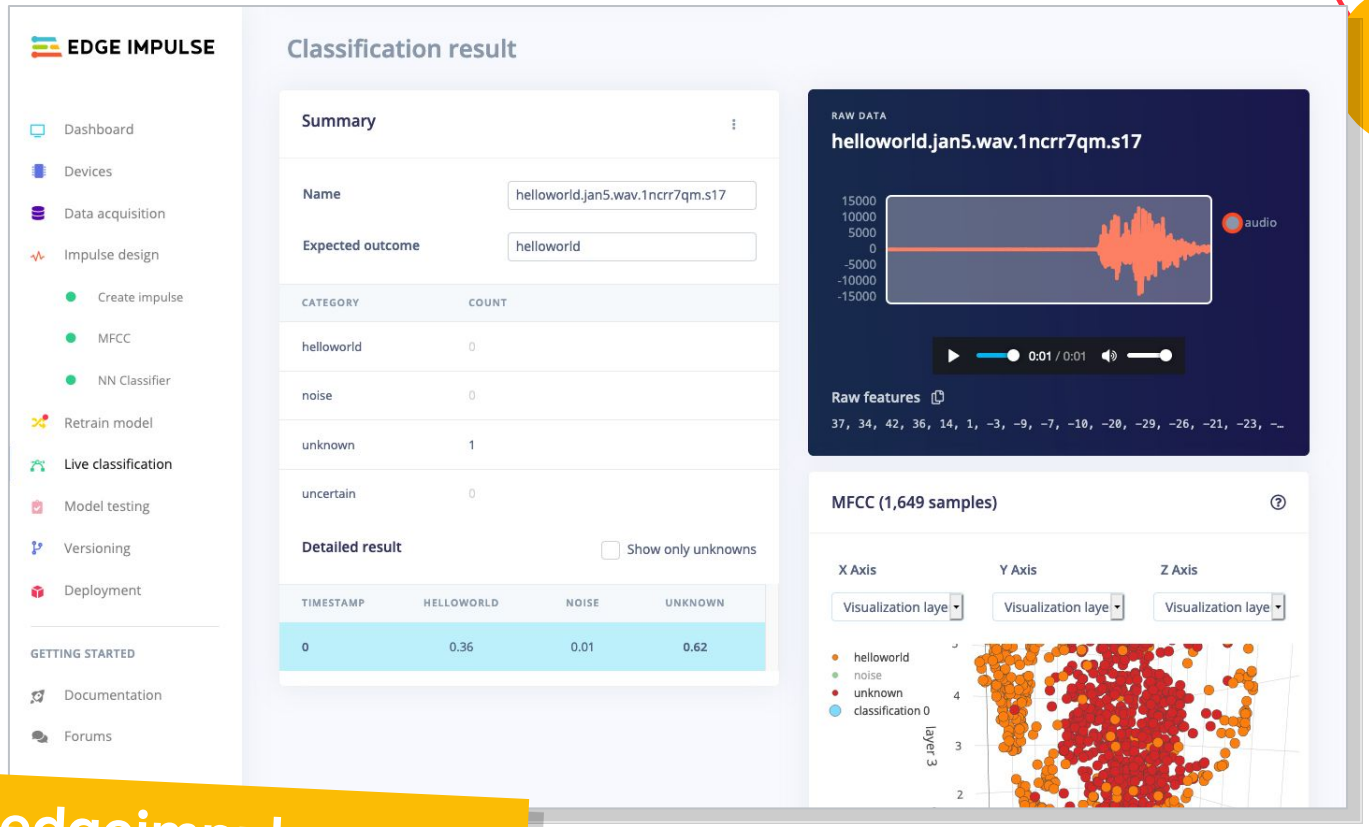

edgeimpulse.com

Neural Networks

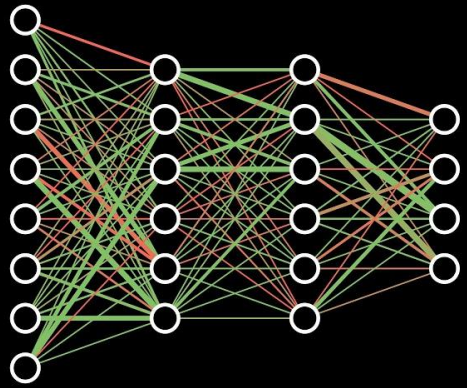

From the ground up

3Blue1Brown

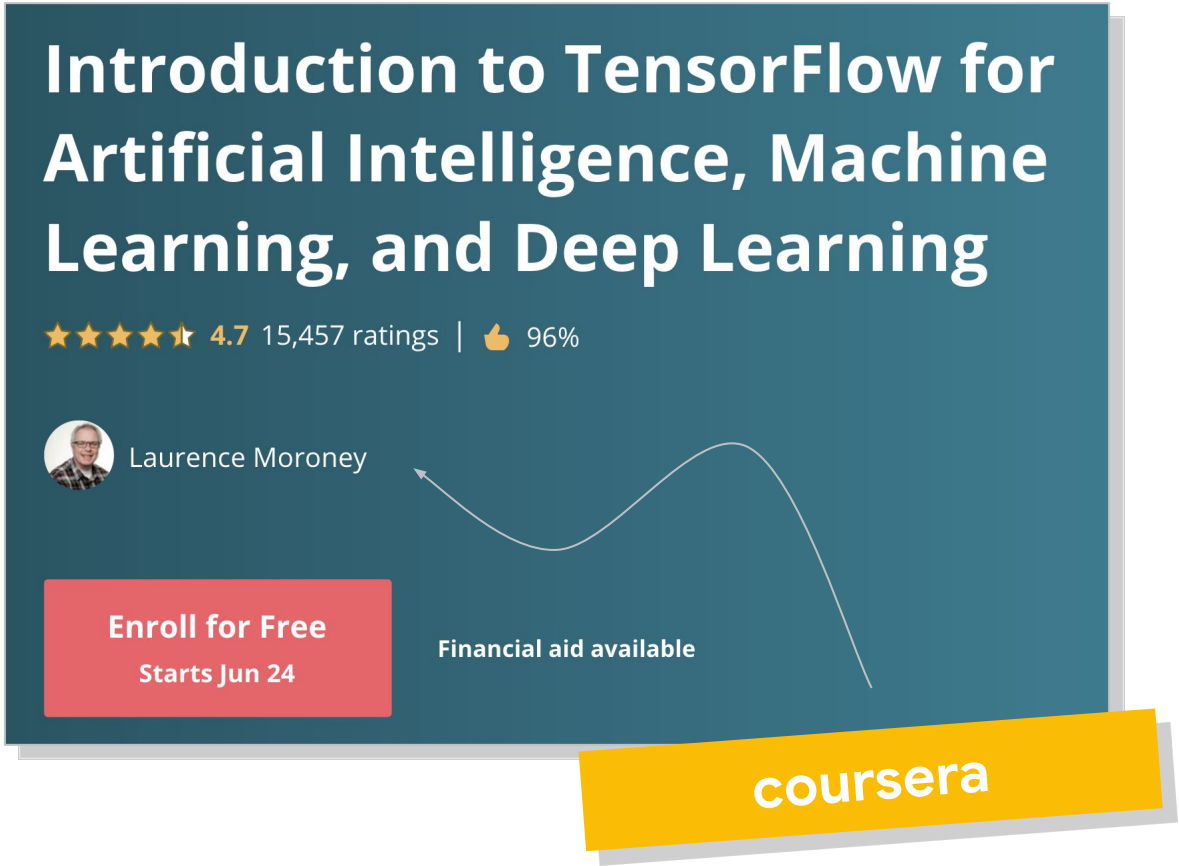

Browse > Data Science > Machine Learning

Offered By

EDGE IMPULSE

Introduction to Embedded Machine Learning

★★★★★ 4.8 335 ratings • 83 reviews

Shawn Hymel +1 more instructor

Enroll for Free
Starts Jun 16

Financial aid available

18,297 already enrolled

coursera

The Future of ML is Tiny and Bright

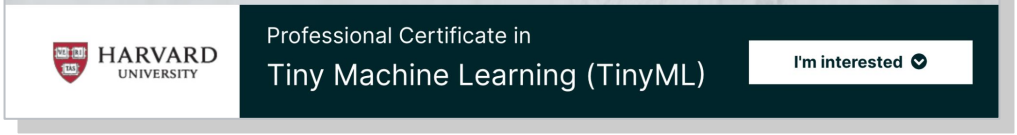

Courses in this program

HarvardX's Tiny Machine Learning (TinyML) Professional Certificate

- **Fundamentals of TinyML**
- **Applications of TinyML**
- **Deploying TinyML**

edx.org

Explore projects

Trending | All difficulties | Featured | Any type

Getting Started with the Raspberry Pi Pico

Arduino “having11” Guy

15 2.5K

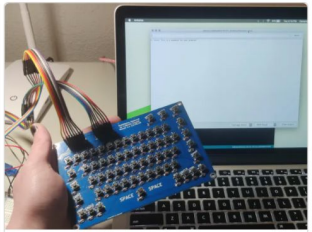

64-Key Prototyping Keyboard Matrix for Arduino

Cameron Coward

19 6.8K

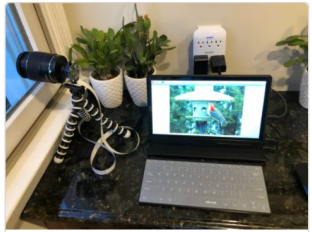

ML-Based Bird and Squirrel Detector (Raspberry Pi and AWS)

Mike Sadowski

31 5.3K

Self-Playing Melodica

touchmysound

30 3.9K

hackster.io

Our website!

tinyMLedu.org/EASI-22

home base for **all information**!

survey!

bit.ly/EASI22-Post

help us make this workshop **better**!

Our team!

with help from **many more**

hágoónee’ 👋

thank you and **keep exploring with ML**