tinyML® Talks

Enabling Ultra-low Power Machine Learning at the Edge

"Standing at the Edge, Looking into the Future"

Alasdair Allan - Raspberry Pi

December 16, 2021

www.tinyML.org

tinyML Talks Strategic Partners

Additional Sponsorships available – contact Olga@tinyML.org for info

Arm: The Software and Hardware Foundation for tinyML

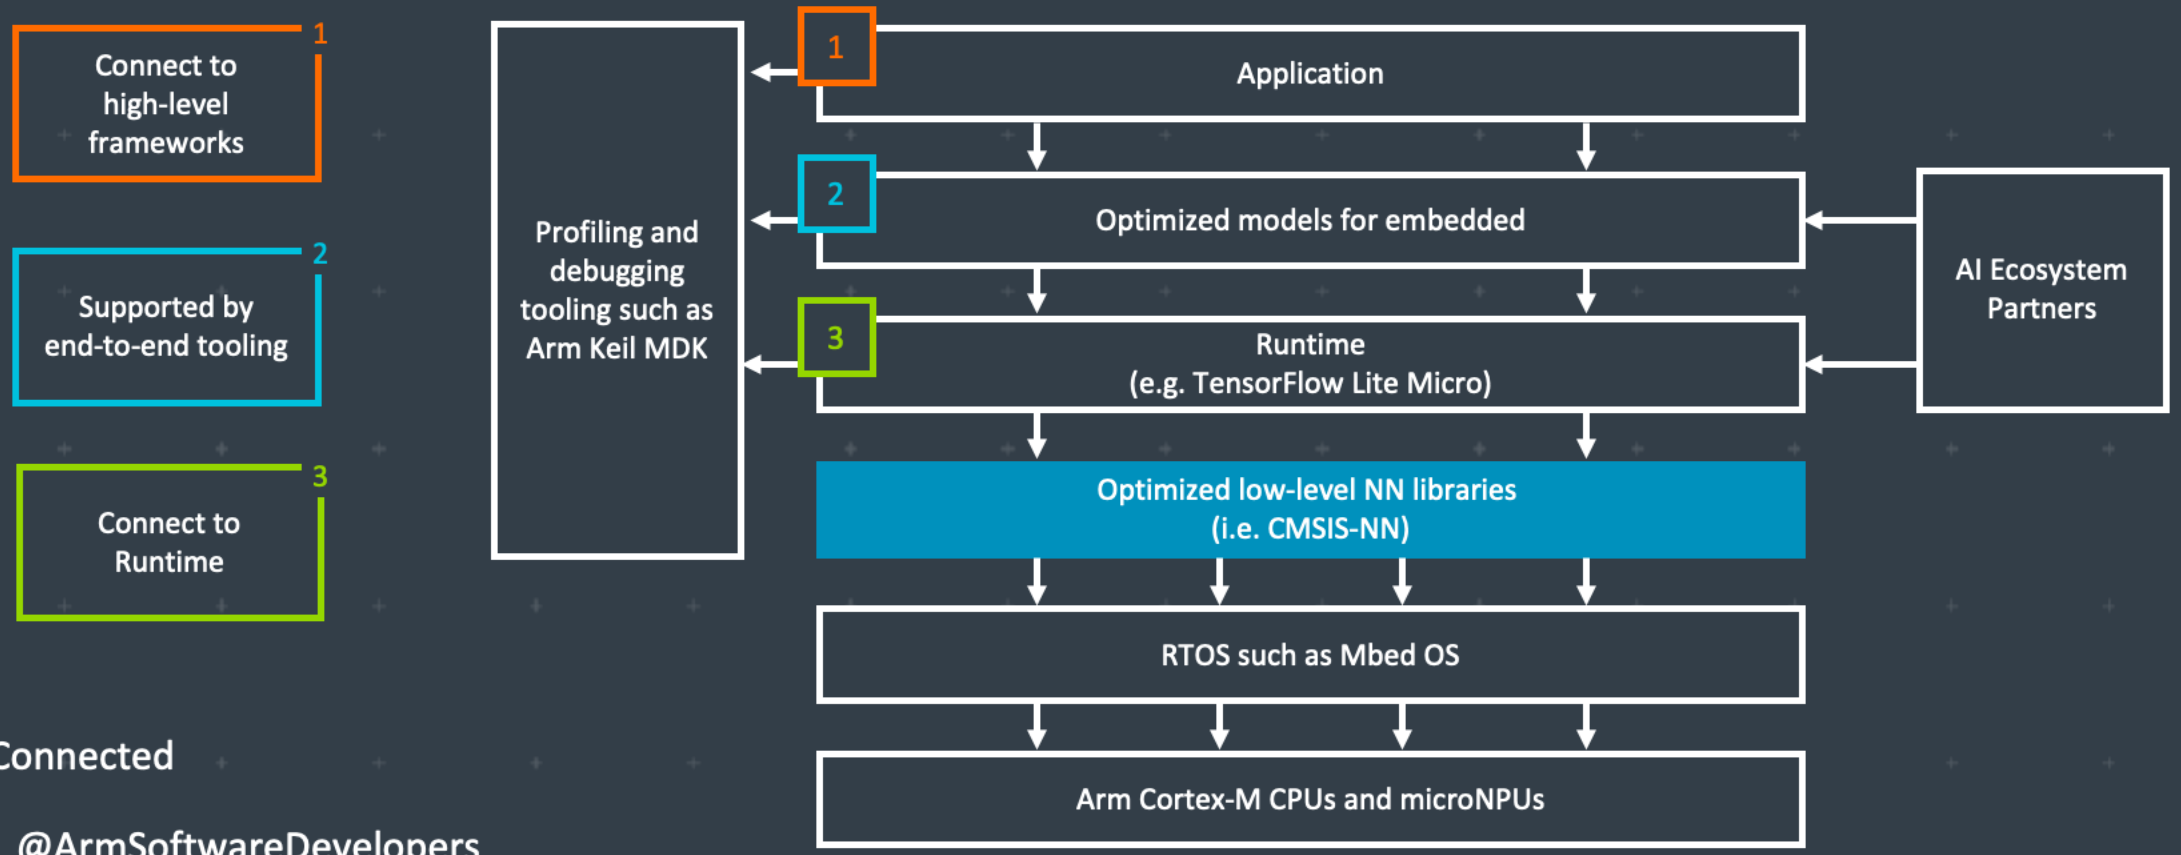

Stay Connected

@ArmSoftwareDevelopers

@ArmSoftwareDev

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

arm

Deeplite

WE USE AI TO MAKE OTHER AI FASTER, SMALLER AND MORE POWER EFFICIENT

Automatically compress SOTA models like MobileNet to <200KB with **little to no drop in accuracy** for inference on resource-limited MCUs

Reduce model optimization trial & error from weeks to days using Deeplite's **design space exploration**

Deploy more models to your device without sacrificing performance or battery life with our **easy-to-use software**

BECOME BETA USER **bit.ly/testdeeplite**

mobilityXlab arm AI 100 CBINSIGHTS

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

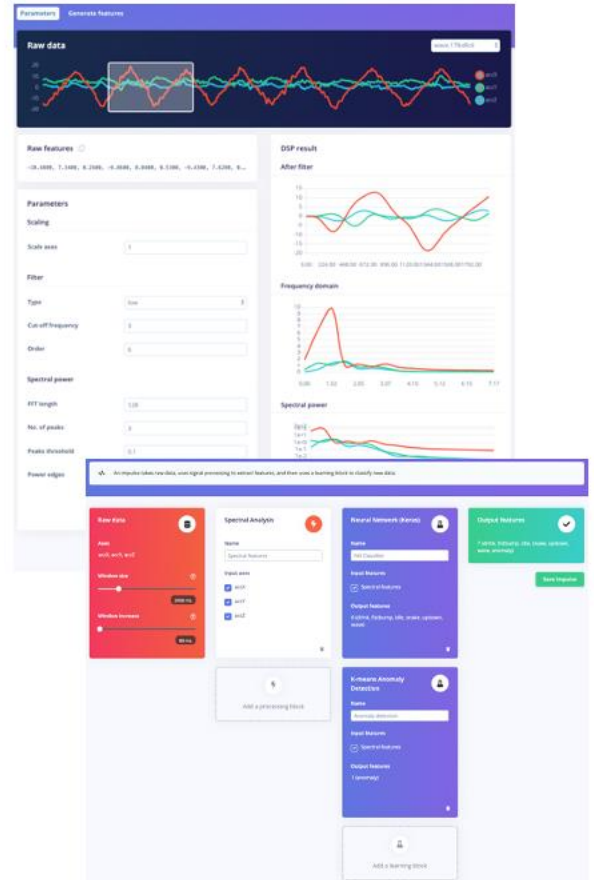

www.edgeimpulse.com

emza
visual sense

The Eye in IoT

Edge AI Visual Sensors

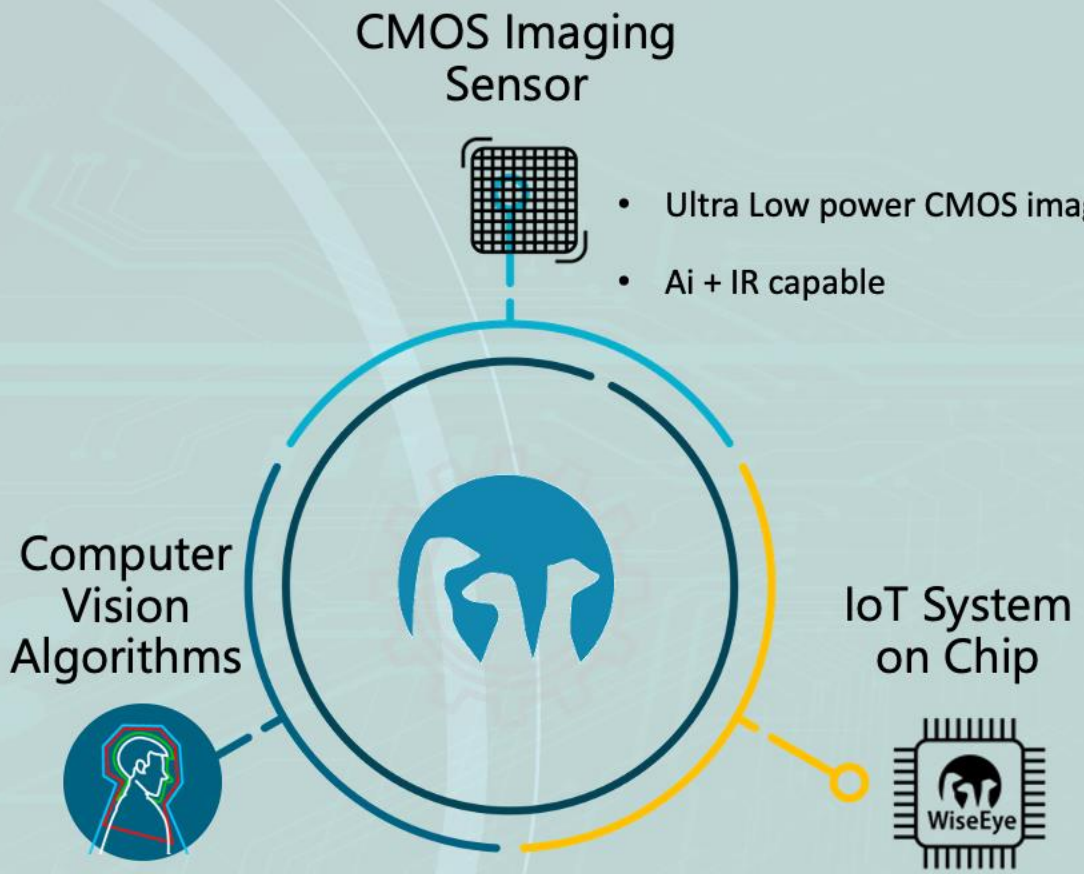

- Ultra Low power CMOS imager
- Ai + IR capable

- Machine Learning algorithm
- <1MB memory footprint
- Microcontrollers computing power
- Trained algorithm
- Processing of low-res images
- Human detection and other classifiers

- Machine Learning edge computing silicon
- <1mW always-on power consumption
- Computer Vision hardware accelerators

info@emza-vs.com

Enabling the next generation of Sensor and Hearable products to process rich data with energy efficiency

Visible Image

Sound

IR Image

Radar

Bio-sensor

Gyro/Accel

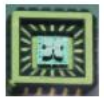

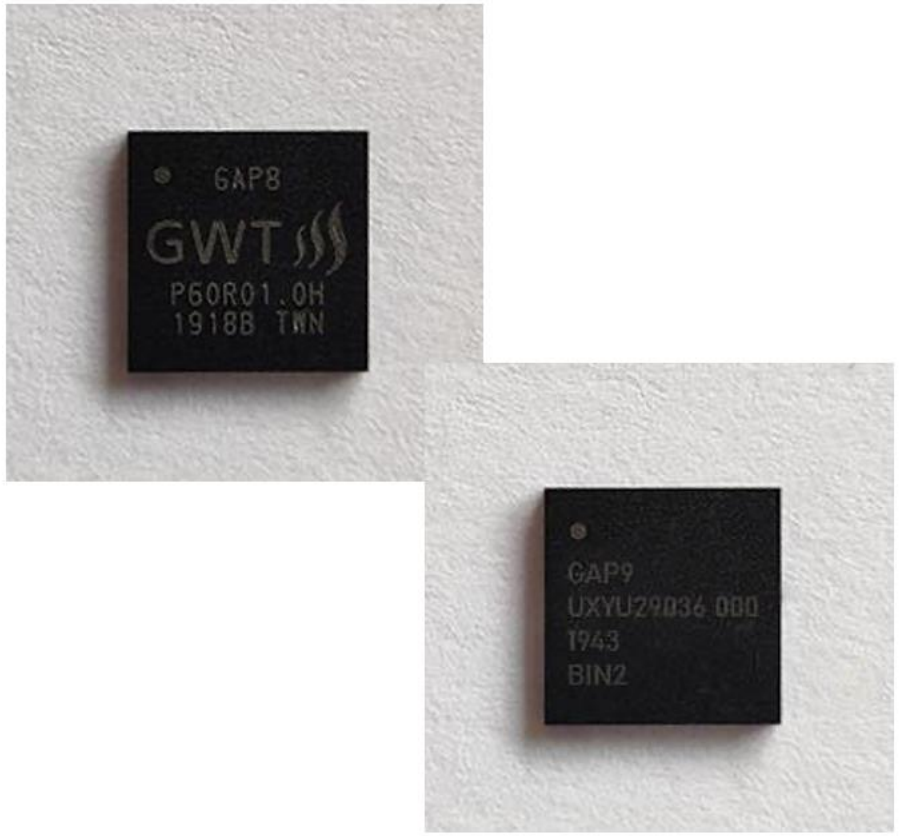

Wearables / Hearables

Battery-powered consumer electronics

IoT Sensors

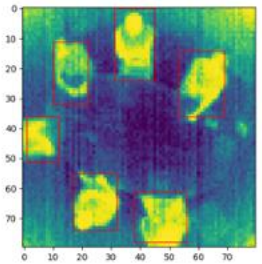

GREENWAVES TECHNOLOGIES

Grovety Inc.

SOFTWARE DEVELOPMENT SERVICES FOR TINYML SOLUTIONS

1 **Development tools**
SDK, IDE, compilers, leveraging on TVM, uTVM & LLVM

2 **Firmware**
Drivers, BSP, protocols, etc.

arm AI PARTNER

Distributed infrastructure for TinyML apps

HOTG
Decoupling intelligence

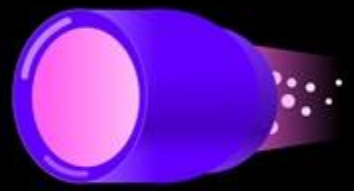

Develop at warp speed

Automate deployments

Device orchestration

HOTG is building the **distributed infrastructure** to pave the way for **AI** enabled **edge applications**

LatentAI

Adaptive AI for the Intelligent Edge

Latentai.com

Maxim Integrated: Enabling Edge Intelligence

Advanced AI Acceleration IC

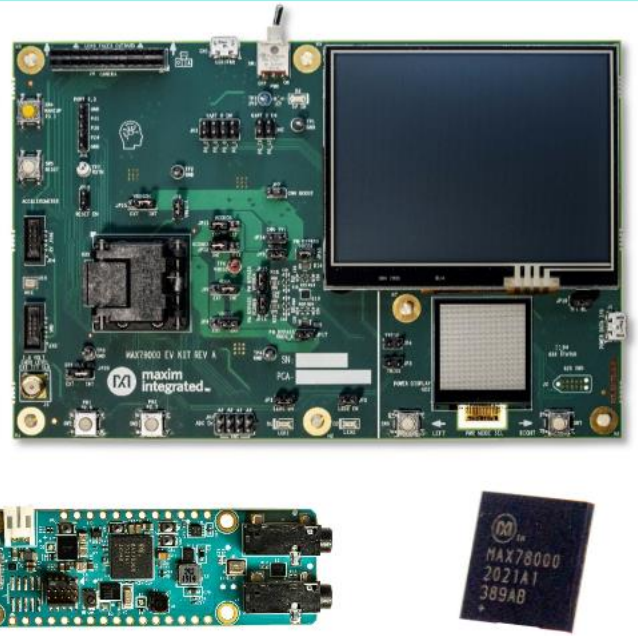

The new MAX78000 implements AI inferences at low energy levels, enabling complex audio and video inferencing to run on small batteries. Now the edge can see and hear like never before.

www.maximintegrated.com/MAX78000

Low Power Cortex M4 Micros

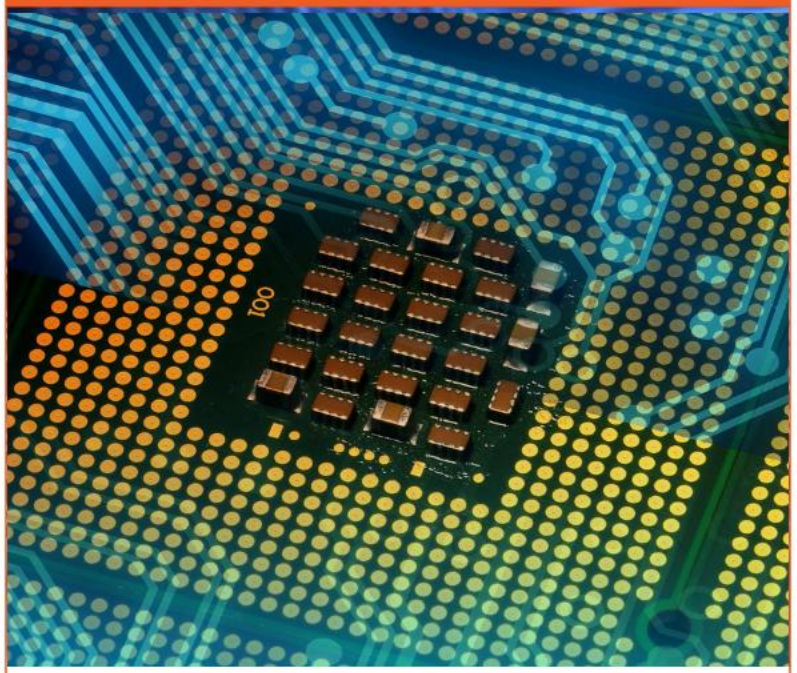

Large (3MB flash + 1MB SRAM) and small (256KB flash + 96KB SRAM, 1.6mm x 1.6mm) Cortex M4 microcontrollers enable algorithms and neural networks to run at wearable power levels.

www.maximintegrated.com/microcontrollers

Sensors and Signal Conditioning

Health sensors measure PPG and ECG signals critical to understanding vital signs. Signal chain products enable measuring even the most sensitive signals.

www.maximintegrated.com/sensors

Qeexo AutoML

Automated Machine Learning Platform that builds tinyML solutions for the Edge using sensor data

Key Features

- Supports 17 ML methods:
  - Multi-class algorithms: GBM, XGBoost, Random Forest, Logistic Regression, Gaussian Naive Bayes, Decision Tree, Polynomial SVM, RBF SVM, SVM, CNN, RNN, CRNN, ANN
  - Single-class algorithms: Local Outlier Factor, One Class SVM, One Class Random Forest, Isolation Forest
- Labels, records, validates, and visualizes time-series sensor data
- On-device inference optimized for low latency, low power consumption, and small memory footprint applications
- Supports Arm® Cortex™- M0 to M4 class MCUs

End-to-End Machine Learning Platform

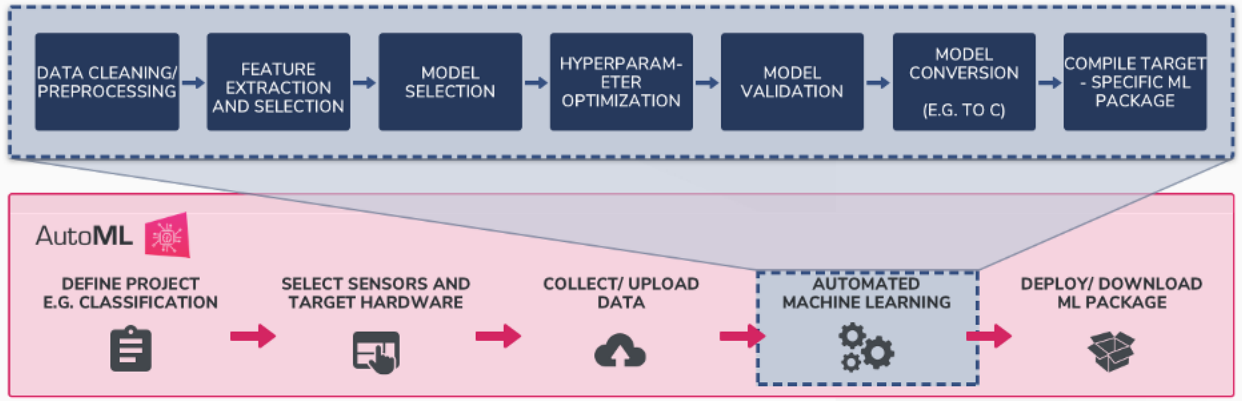

For more information, visit: www.qeexo.com

Target Markets/Applications

- Industrial Predictive Maintenance
- Smart Home
- Wearables
- Automotive
- Mobile
- IoT

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

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

BROAD AND SCALABLE EDGE COMPUTING PORTFOLIO

Microcontrollers & Microprocessors

Arm® Core

RA — Arm® Cortex®-M 32-bit MCUs
Arm ecosystem, Advanced security, Intelligent IoT

RZ — Arm®-based High-end 32 & 64-bit MPUs
High-resolution HMI, Industrial network & real-time control

RE — Arm® Cortex®-M0+ Ultra-low Power 32-bit MCUs
Innovative process tech (SOTB), Energy harvesting

Renesas Synergy™ — Arm®-based 32-bit MCUs for Qualified Platform
Qualified software and tools

Renesas Core

RL78 — Ultra-low Energy 8 & 16-bit MCUs
Bluetooth® Low Energy, SubGHz, LoRa®-based Solutions

RX — High Power Efficiently 32-bit MCUs
Motor control, Capacitive touch, Functional safety, GUI

RH850 — 40nm/28nm process Automotive 32-bit MCUs
Rich functional safety and embedded security features

Core technologies

AI

A broad set of high-power and energy-efficient embedded processors

Security & Safety

Comprehensive technology and support that meet the industry's stringent standards

Digital & Analog & Power Solution

Winning Combinations that combine our complementary product portfolios

Cloud Native

Cross-platforms working with partners in different verticals and organizations

BIG IDEAS FOR EVERY SPACE RENESAS

seeed studio

The IoT Hardware Enabler

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

SynSense

SynSense builds **sensing and inference** hardware for **ultra-low-power** (sub-mW) **embedded, mobile and edge** devices. We design systems for **real-time always-on smart sensing**, for audio, vision, IMUs, bio-signals and more.

https://SynSense.ai

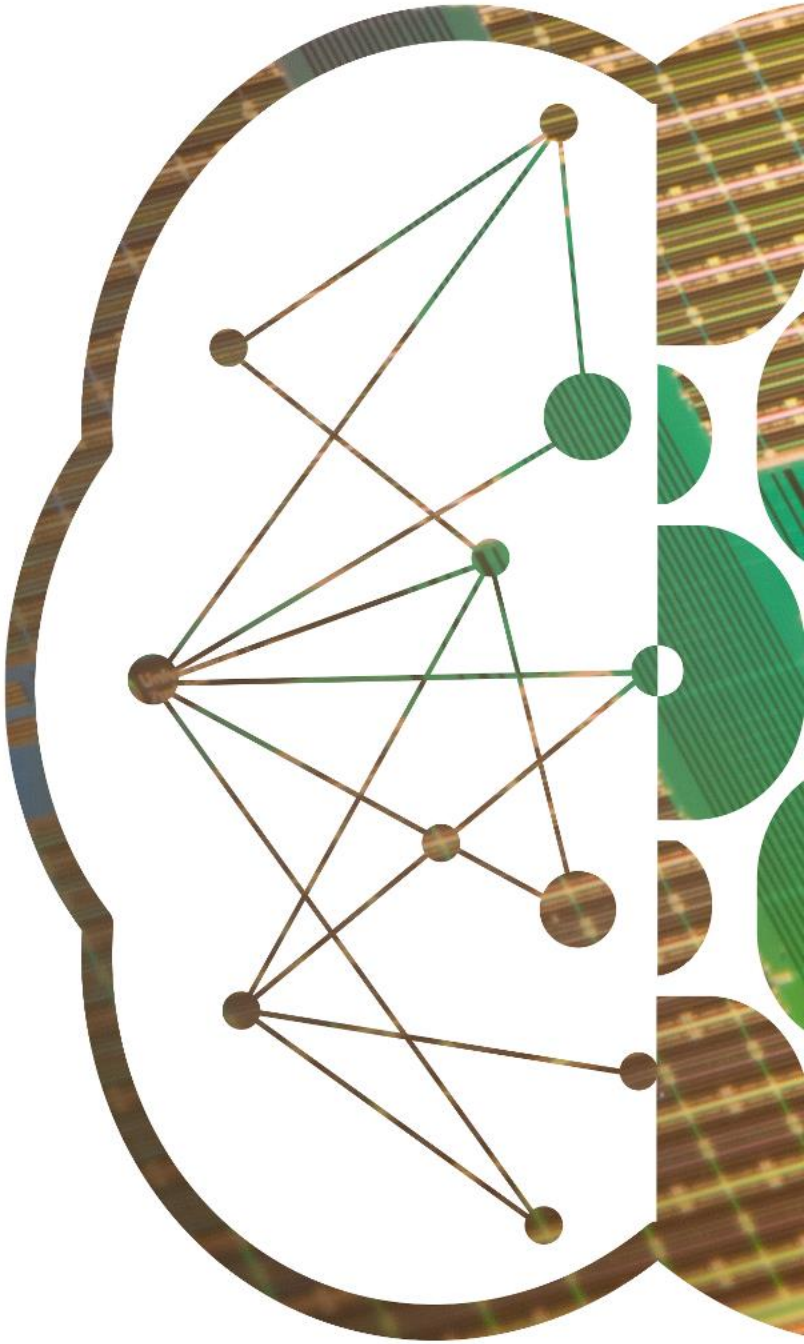

SYNTIANT

End-to-End Deep Learning Solutions for TinyML & Edge AI

Neural Decision Processors

- At-Memory Compute
- Sustained High MAC Utilization
- Native Neural Network Processing

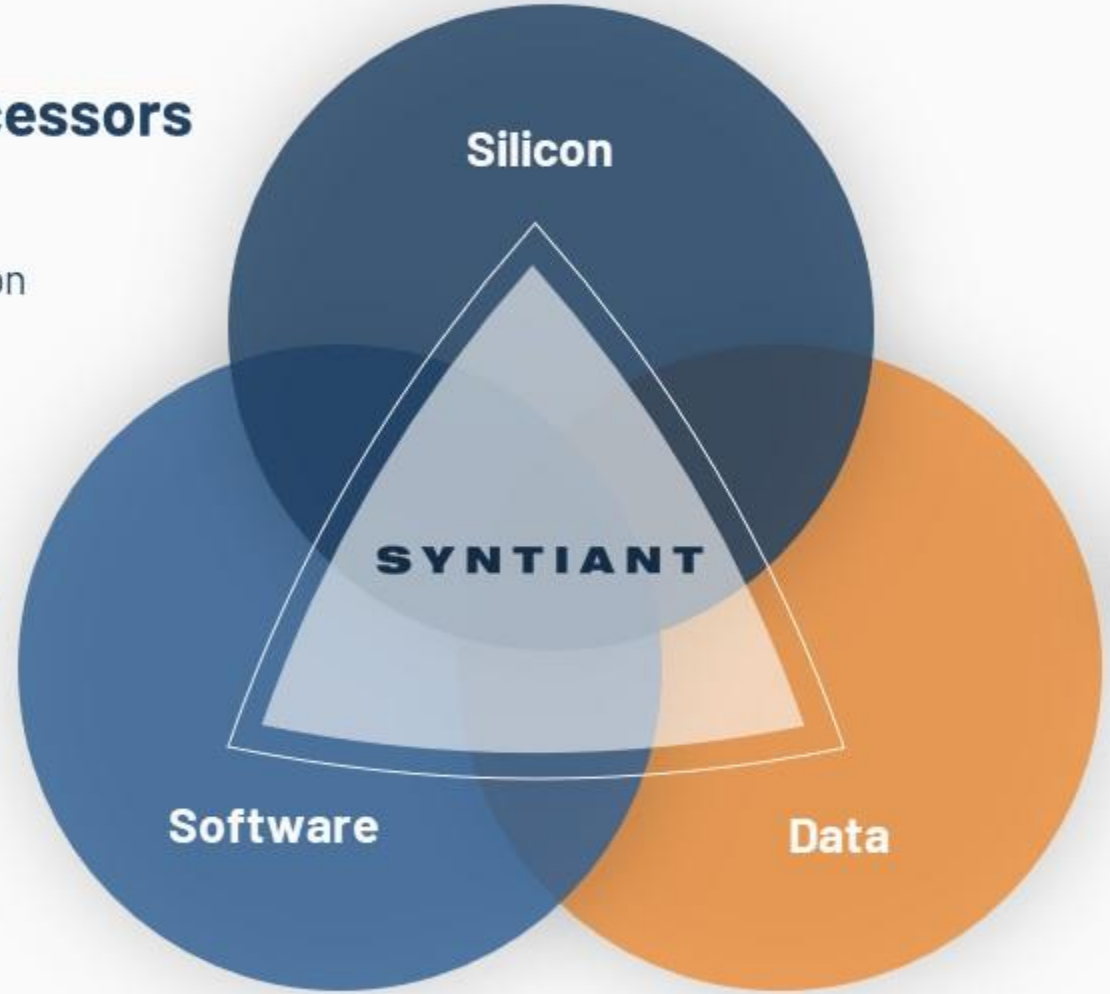

ML Training Pipeline

- Enables Production Quality Deep Learning Deployments

Data Platform

- Reduces Data Collection Time and Cost
- Increases Model Performance

partners@syntiant.com

www.syntiant.com

tinyML Summit 2022

Miniature dreams can come true…

March 28-30, 2022

Hyatt Regency San Francisco Airport

https://www.tinyml.org/event/summit-2022/

Registration will be open on **December 15**, 2021.

Deadline for poster submission is **December 17**.

*The Best Product of the Year and the Best Innovation of the Year awards are open for nominations between **November 15 and February 28.***

tinyML Research Symposium 2022

March 28, 2022

https://www.tinyml.org/event/research-symposium-2022

Call for papers – Submission deadline is **December 17**, 2021.

More sponsorships are available: sponsorships@tinyML.org

Next tinyML Talks

Date	Presenter	Topic / Title
Friday, December 17	Odin Shen, Principal Field Application Engineer, Arm	tinyML Hardware/Software co-design, security

Webcast start time is 3:00 am Pacific time

Please contact talks@tinyml.org if you are interested in presenting

TINY ML TALKS webcast

Reminders

Slides & Videos will be posted tomorrow

PDF

tinyml.org/forums

youtube.com/tinyml

Please use the Q&A window for your questions

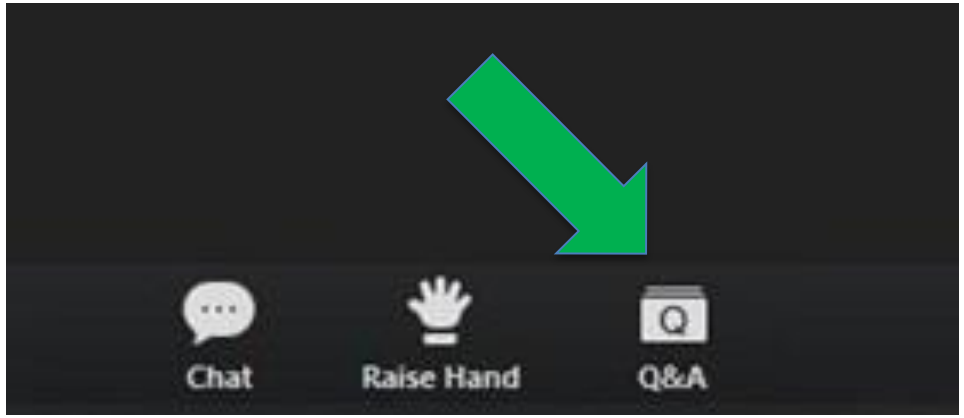

TINY ML TALKS webcast

Alasdair Allan

Alasdair Allan is a scientist, author, hacker, maker, and journalist. He currently works for Raspberry Pi, and is responsible for their documentation. However in the past he has worked as a consultant and journalist, focusing on open hardware, machine learning, big data, and emerging technologies — with expertise in programming, electronics, and especially wireless devices and distributed sensor networks. A former astronomer, he built a peer-to-peer network of telescopes that, acting autonomously, reactively scheduled observations of time-critical events. Notable successes included contributing to the detection of what—at the time—was the most distant object yet discovered.

Standing on the Edge, Looking into the Future

TinyML Virtual Meetup

Dr Alasdair Allan
@aallan

Cliffs of Moher, County Clare, Ireland
Photo Credit: Michal Osmenda

Photo Credit: Max Morse/Techcrunch

A combination of this "carry anywhere" device and a global information utility such as the ARPA network or two-way cable TV, will bring the libraries and schools (not to mention stores and billboards) of the world to the home. One can imagine one of the first programs an owner will write is a filter to eliminate advertising!

Photo Credit: The Guardian


CloudPets
CloudPets
CloudPets
CloudPets

Photo Credit: Amy Osborne/Getty Images

Photo Credit: Jamison Riley

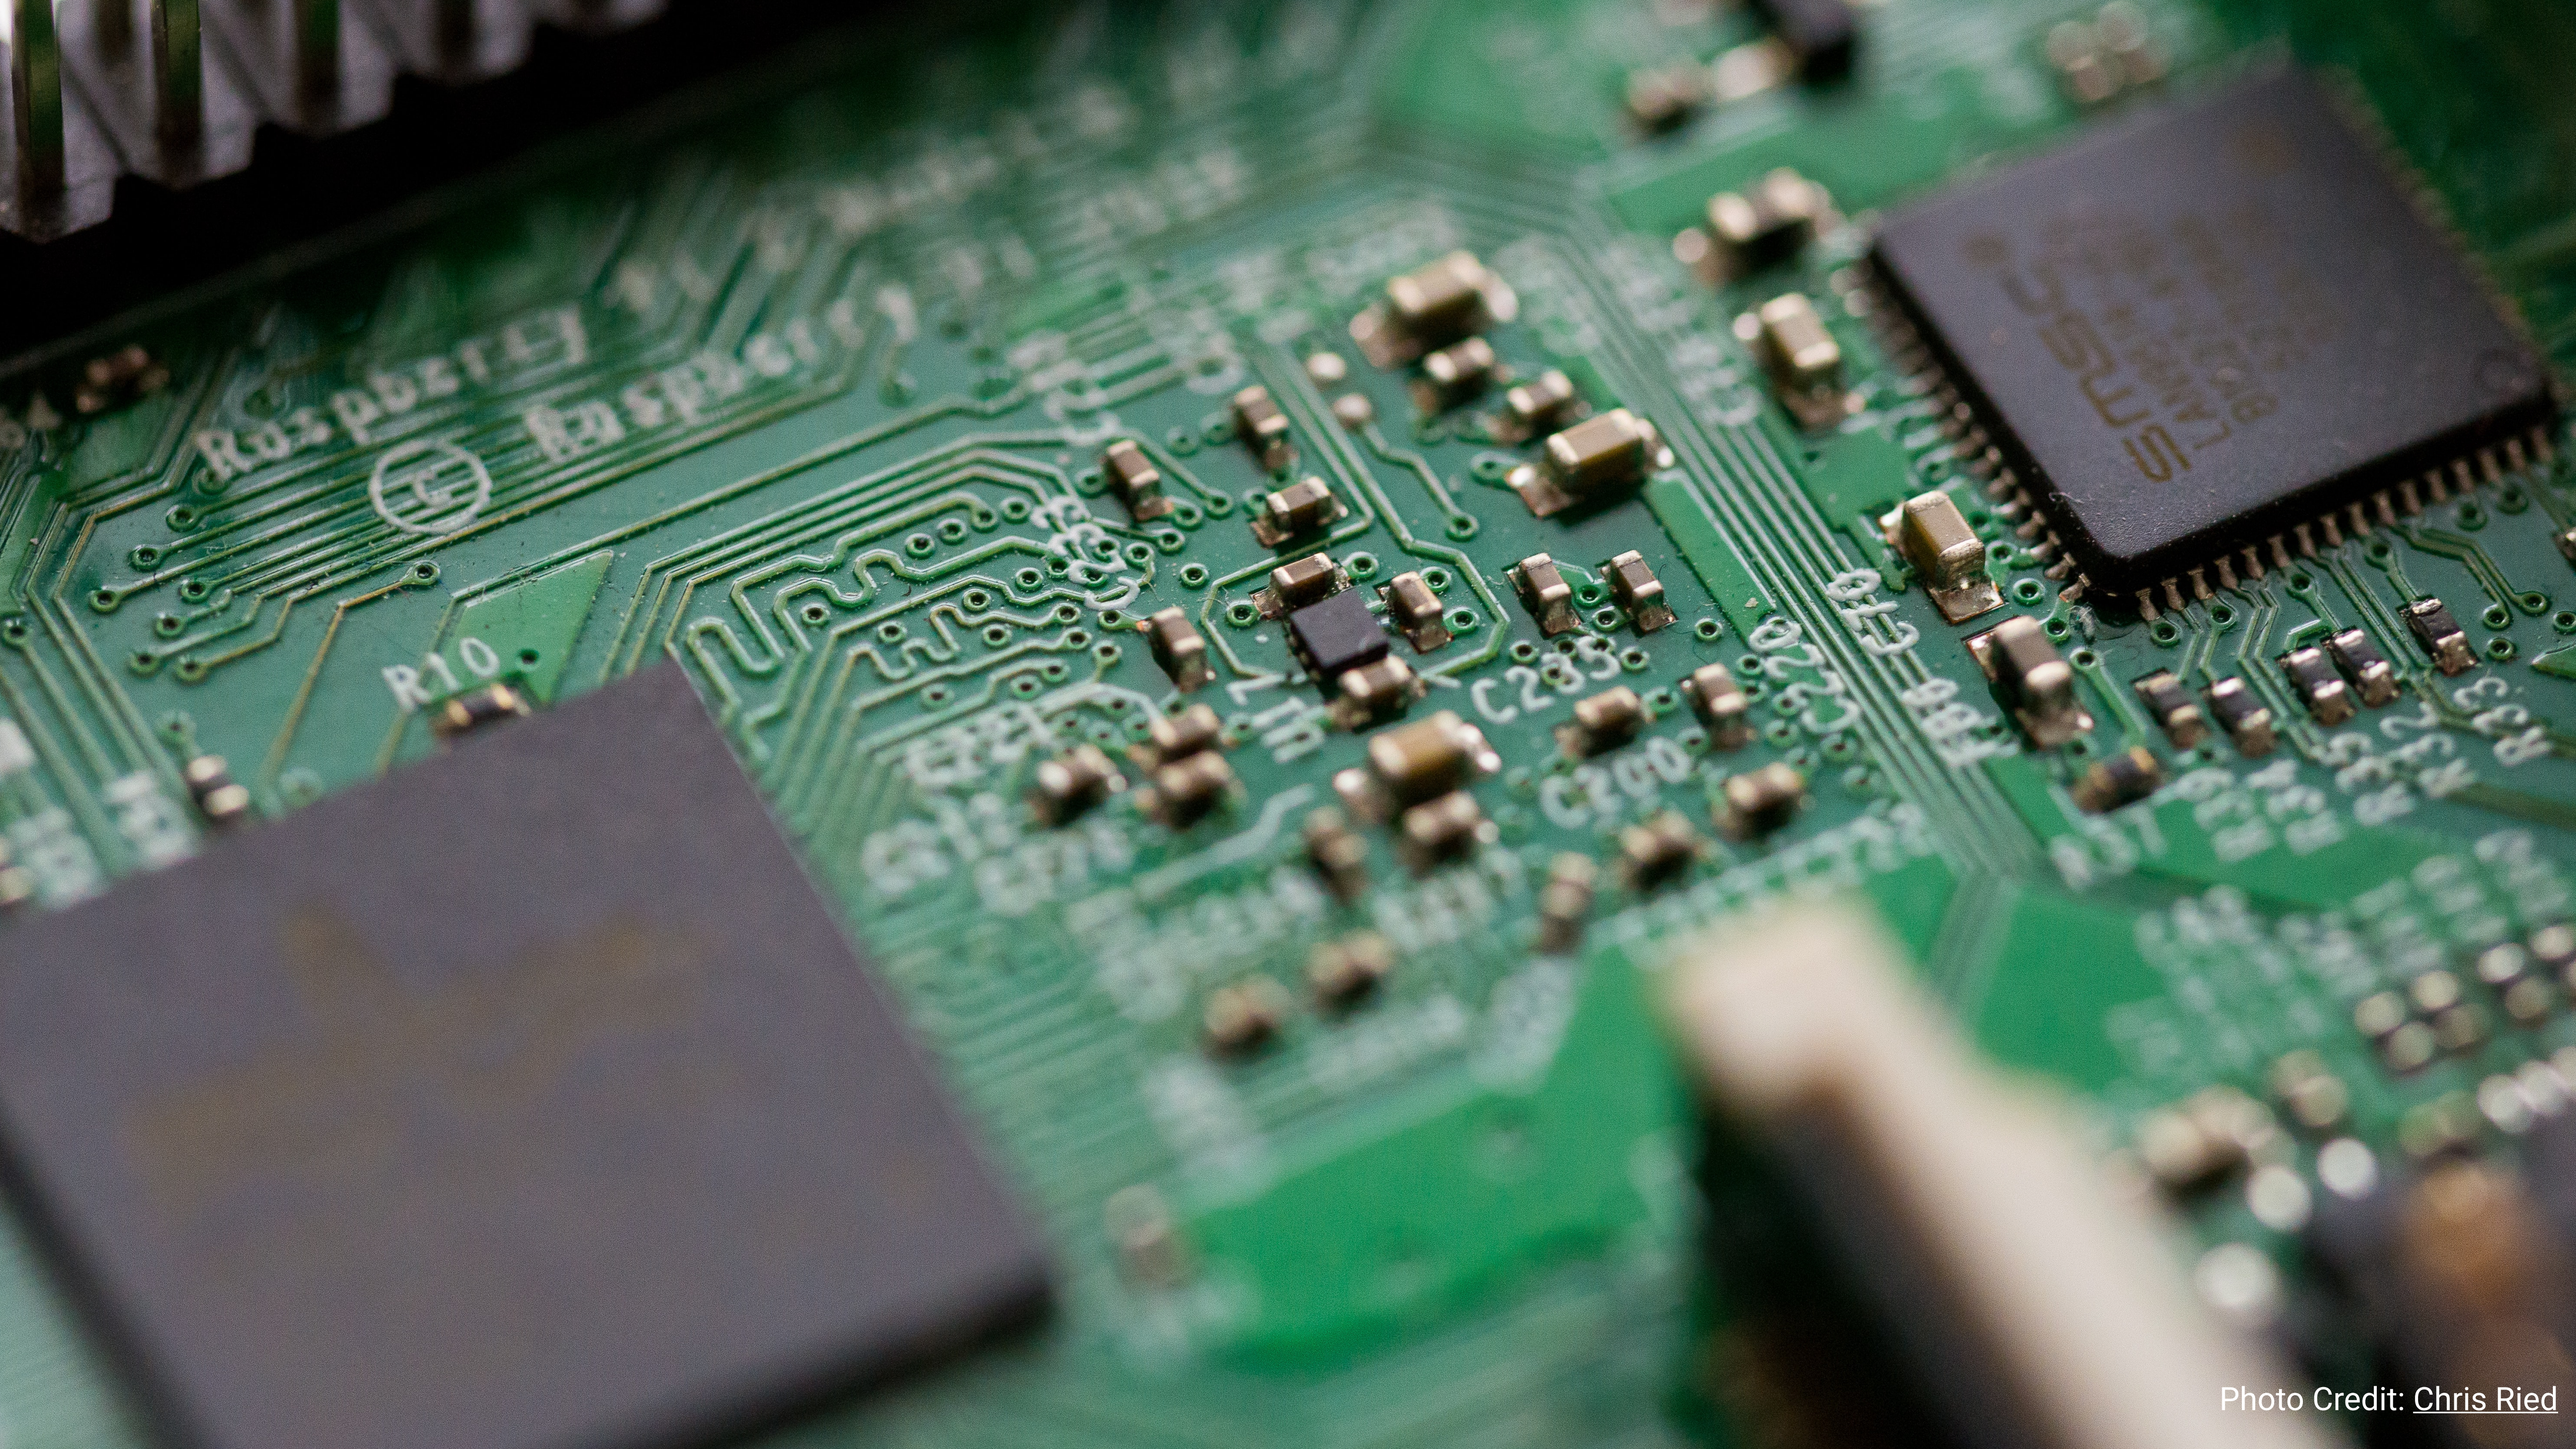

Photo Credit: Chris Ried

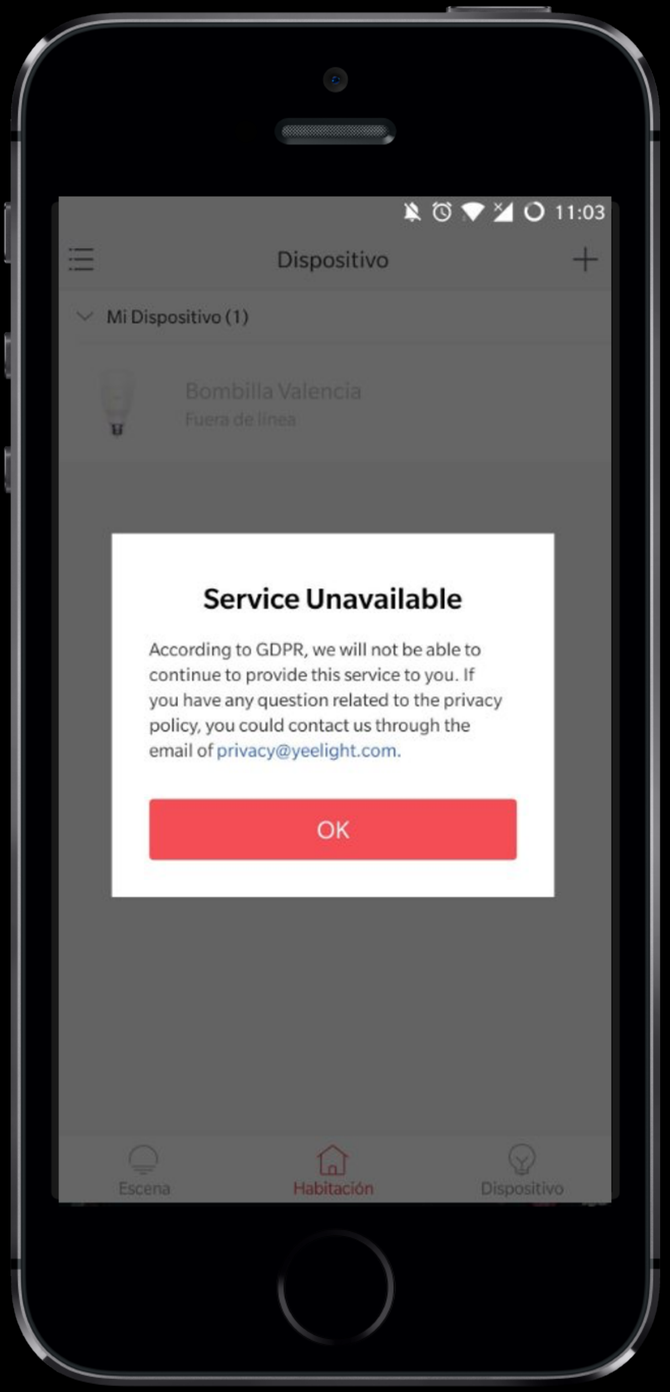
11:03
Dispositivo
Mi Dispositivo (1)
Bombilla Valencia
Fuera de línea
Service Unavailable
According to GDPR, we will not be able to continue to provide this service to you. If you have any question related to the privacy policy, you could contact us through the email of privacy@yeelight.com.
OK
Escena
Habitación
Dispositivo

Photo Credit: Bekky Bekks

Photo Credit: Rodion Kutsaev


Photo Credit: Eberhard Grossgasteiger

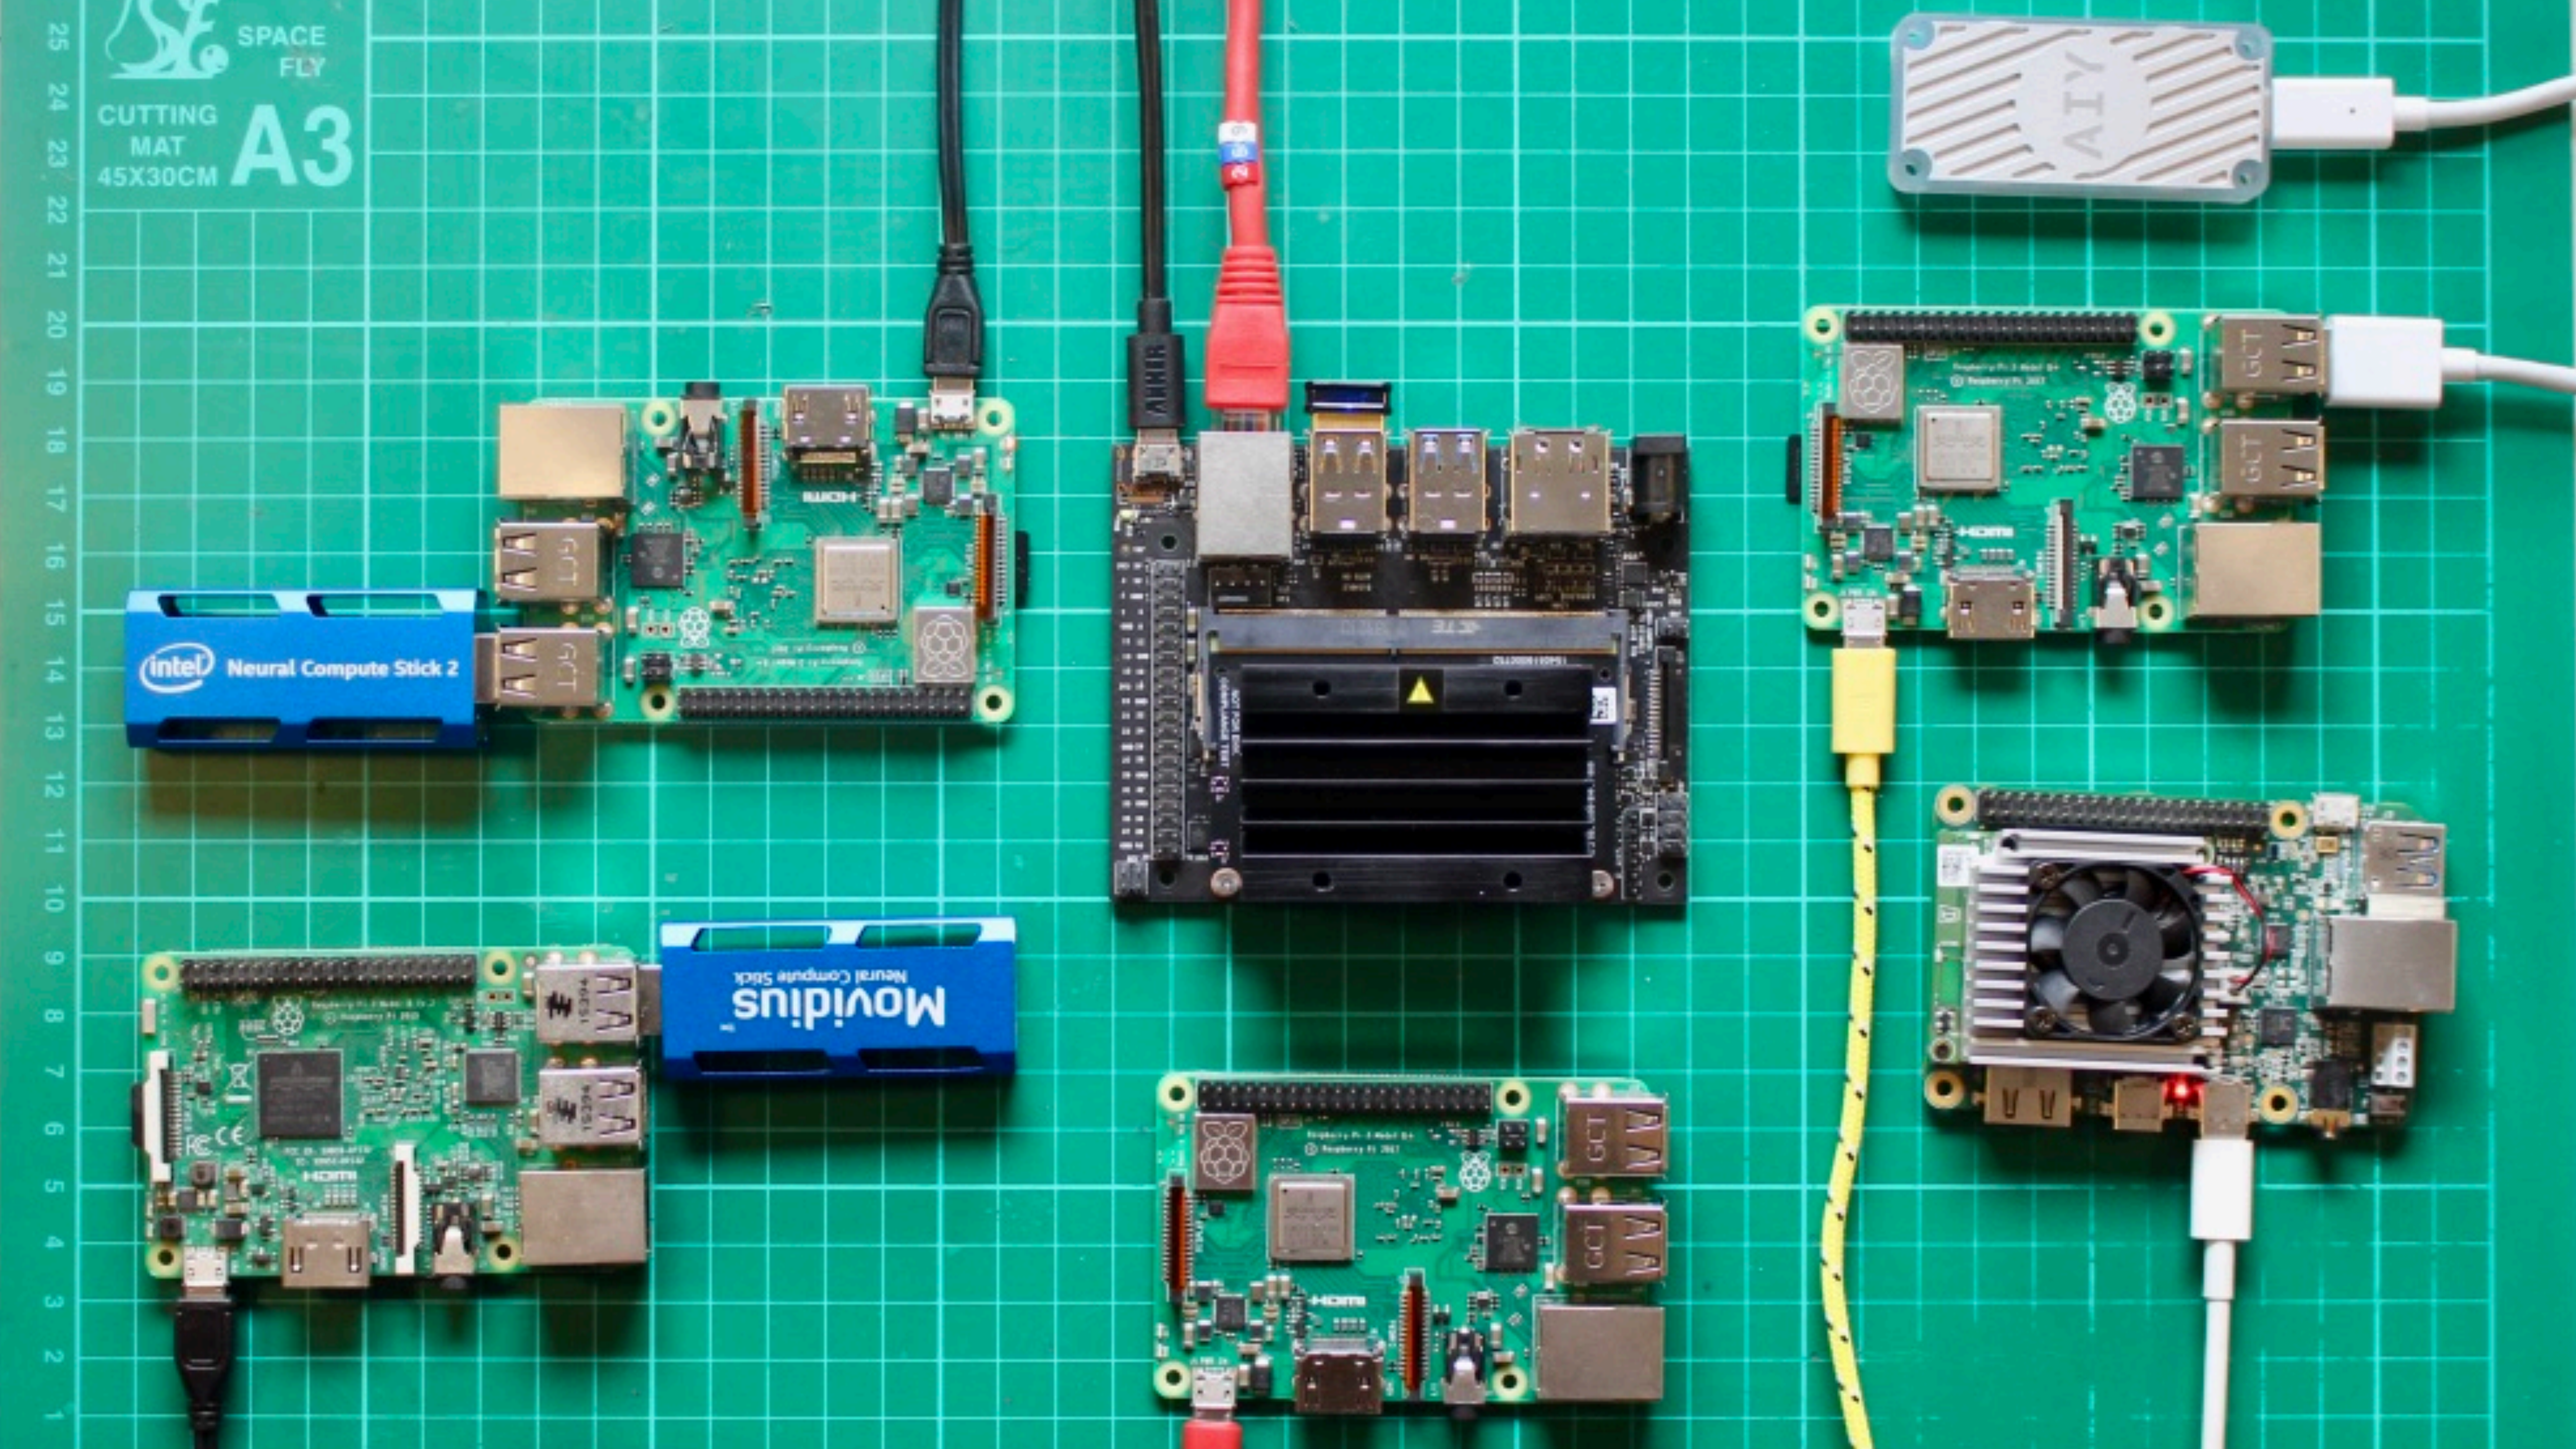
SPACE FLY
CUTTING MAT 45X30CM A3
intel Neural Compute Stick 2
Movidius Neural Compute Stick

Board	MobileNet v1 (ms)	MobileNet v2 (ms)
Coral Dev Board	15.7	20.9
NVIDIA Jetson Nano (TF)	276.0	309.3
NVIDIA Jetson Nano (TF-TRT)	61.6	72.3
MacBook Pro	33.0	71.0

Board	3, Model B+		4, Model B	
	MobileNet v1 (ms)	MobileNet v2 (ms)	MobileNet v1 (ms)	MobileNet v2 (ms)
Coral USB Accelerator (USB2)	49.3	58.1	81.5	102.3
Coral USB Accelerator (USB3)		N/A	14.9	18.2
Movidius NCS (USB2)	115.7	204.5	114.8	202.9
Movidius NCS (USB3)		N/A	88.4	176.4
Intel NCS2 (USB2)	87.2	118.6	85.9	116.7
Intel NCS2 (USB3)		N/A	52.8	80.4
Raspberry Pi (TF Lite)	271.5	379.6	82.7	122.6
Raspberry Pi (TF)	480.3	654.0	263.9	483.5

Raspberry Pi (Xnor)[1]	147.9	79.5

[1] The model used by the Xnor Platform is a proprietary binary convolution network.

http://bit.ly/great-big-roundup

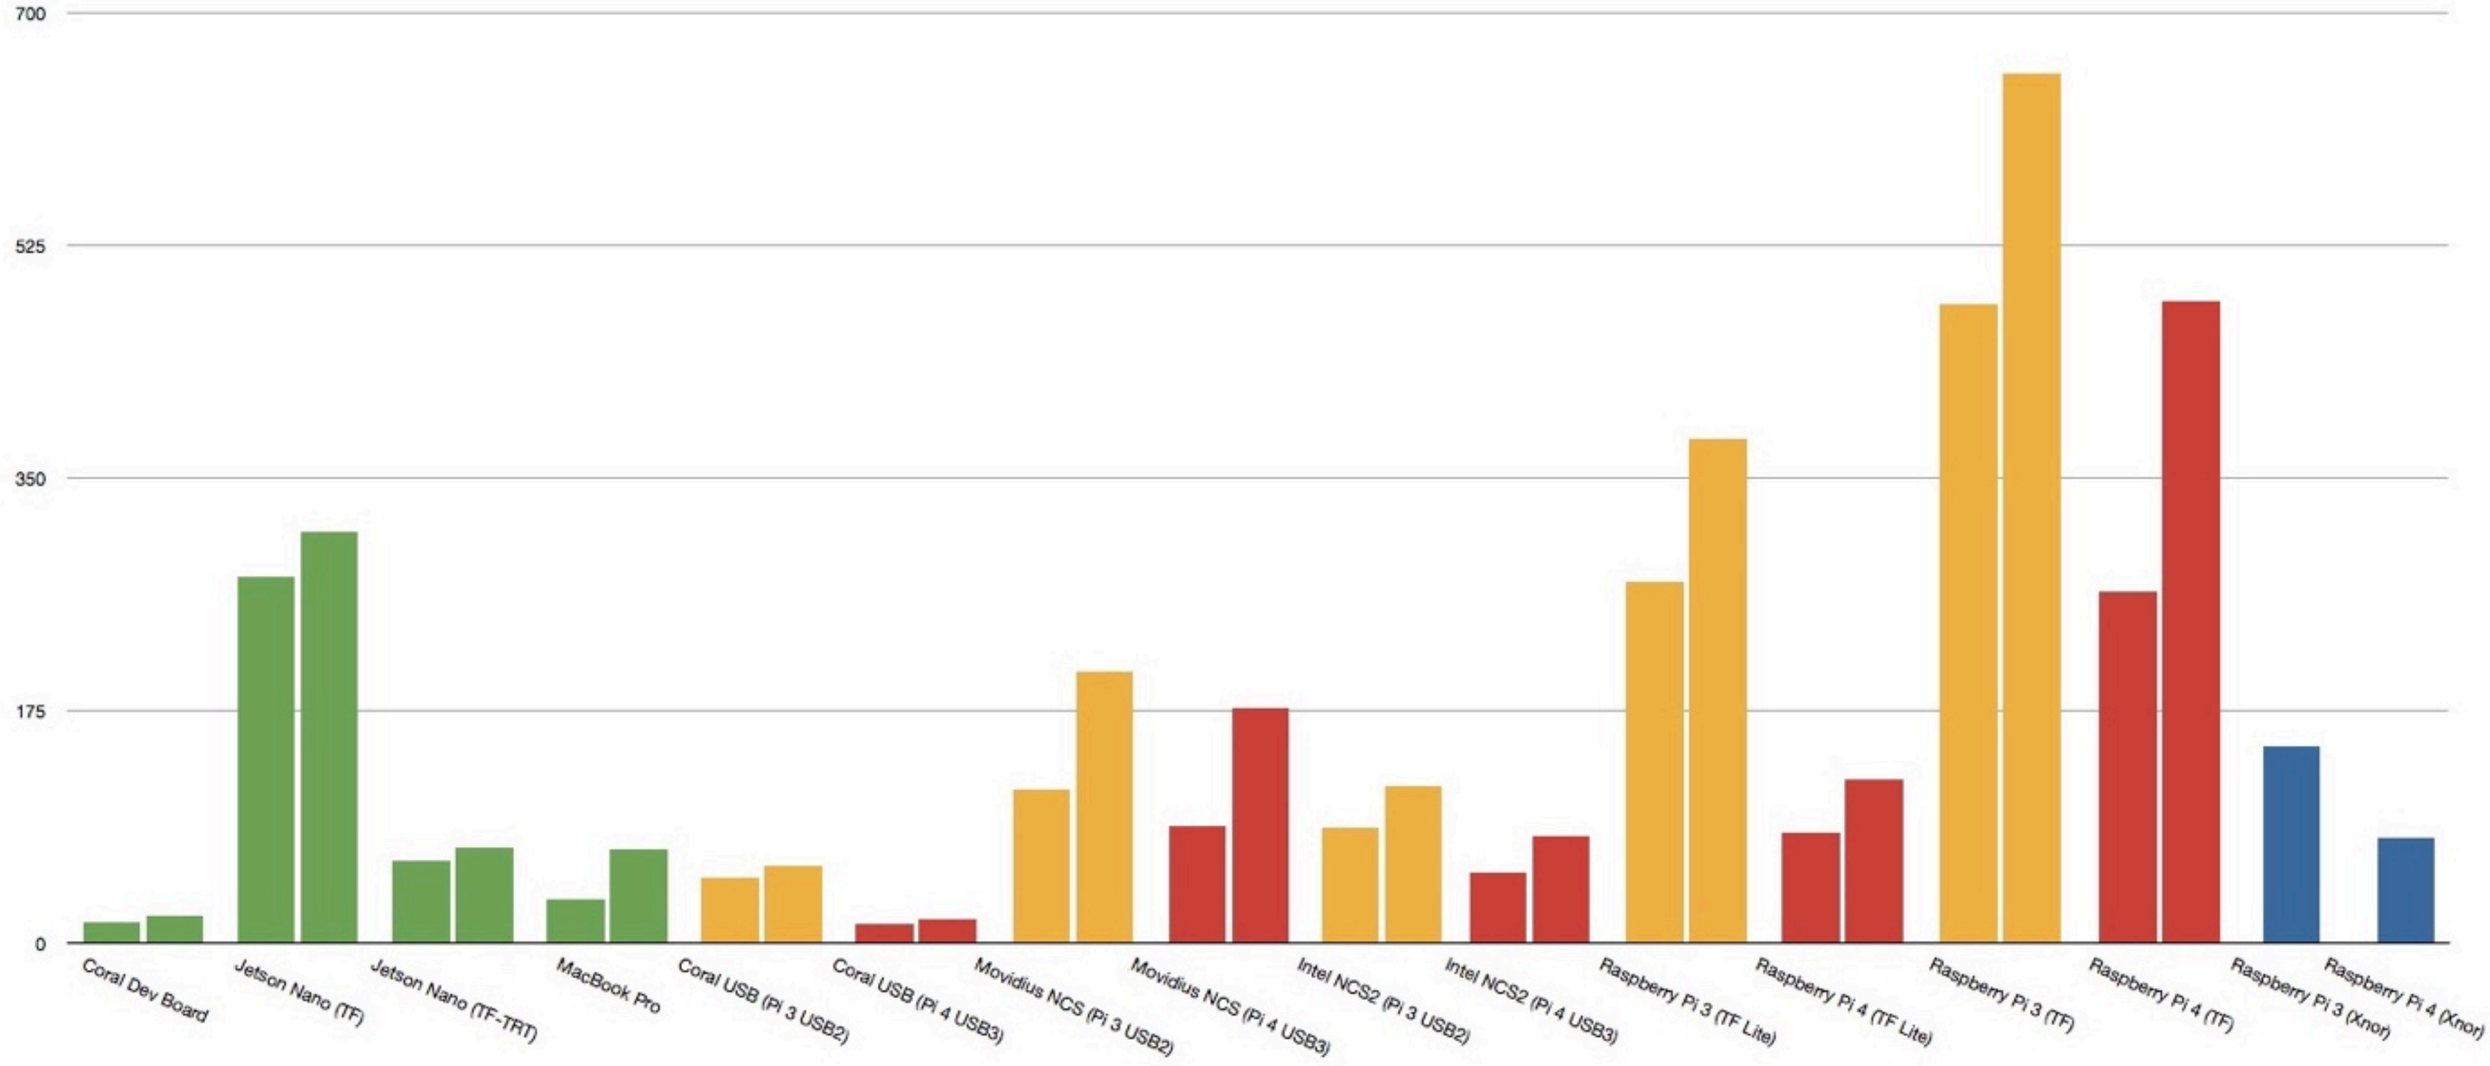

http://bit.ly/great-big-roundup

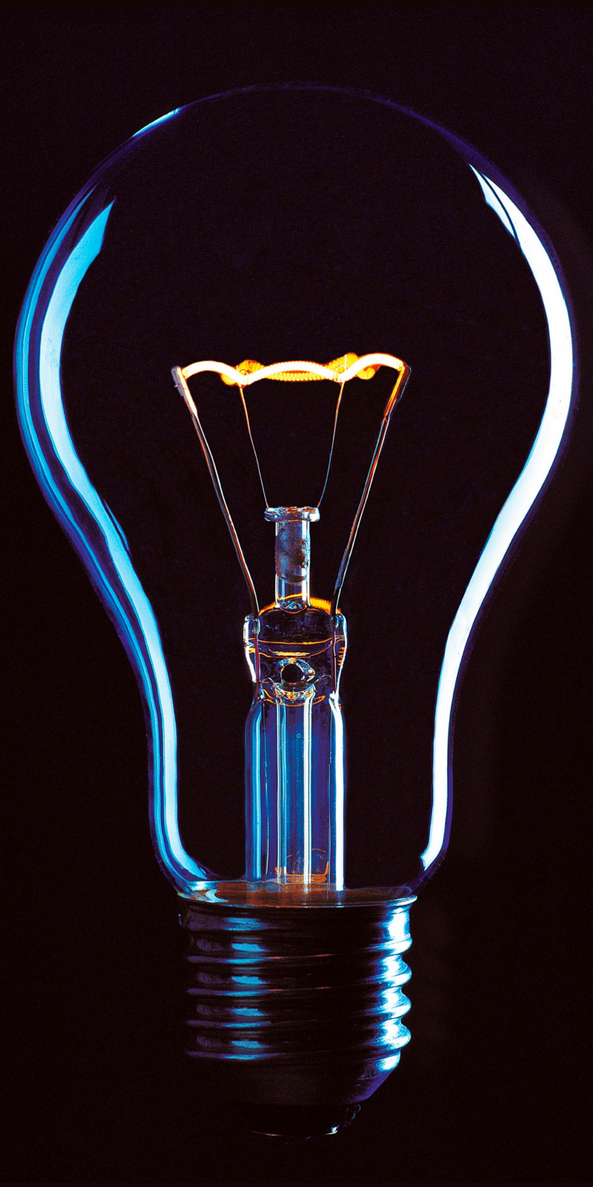

Photo Credit: Alessandro Bianchi

Photo Credit: Ricardo Arce

Photo Credit: Jason Blackeye

Photo Credit: Oziel Gómez

Photo Credit: Daiga Ellaby

Photo Credit: Yulia Khlebnikova

Photo Credit: Bruno Nascimento

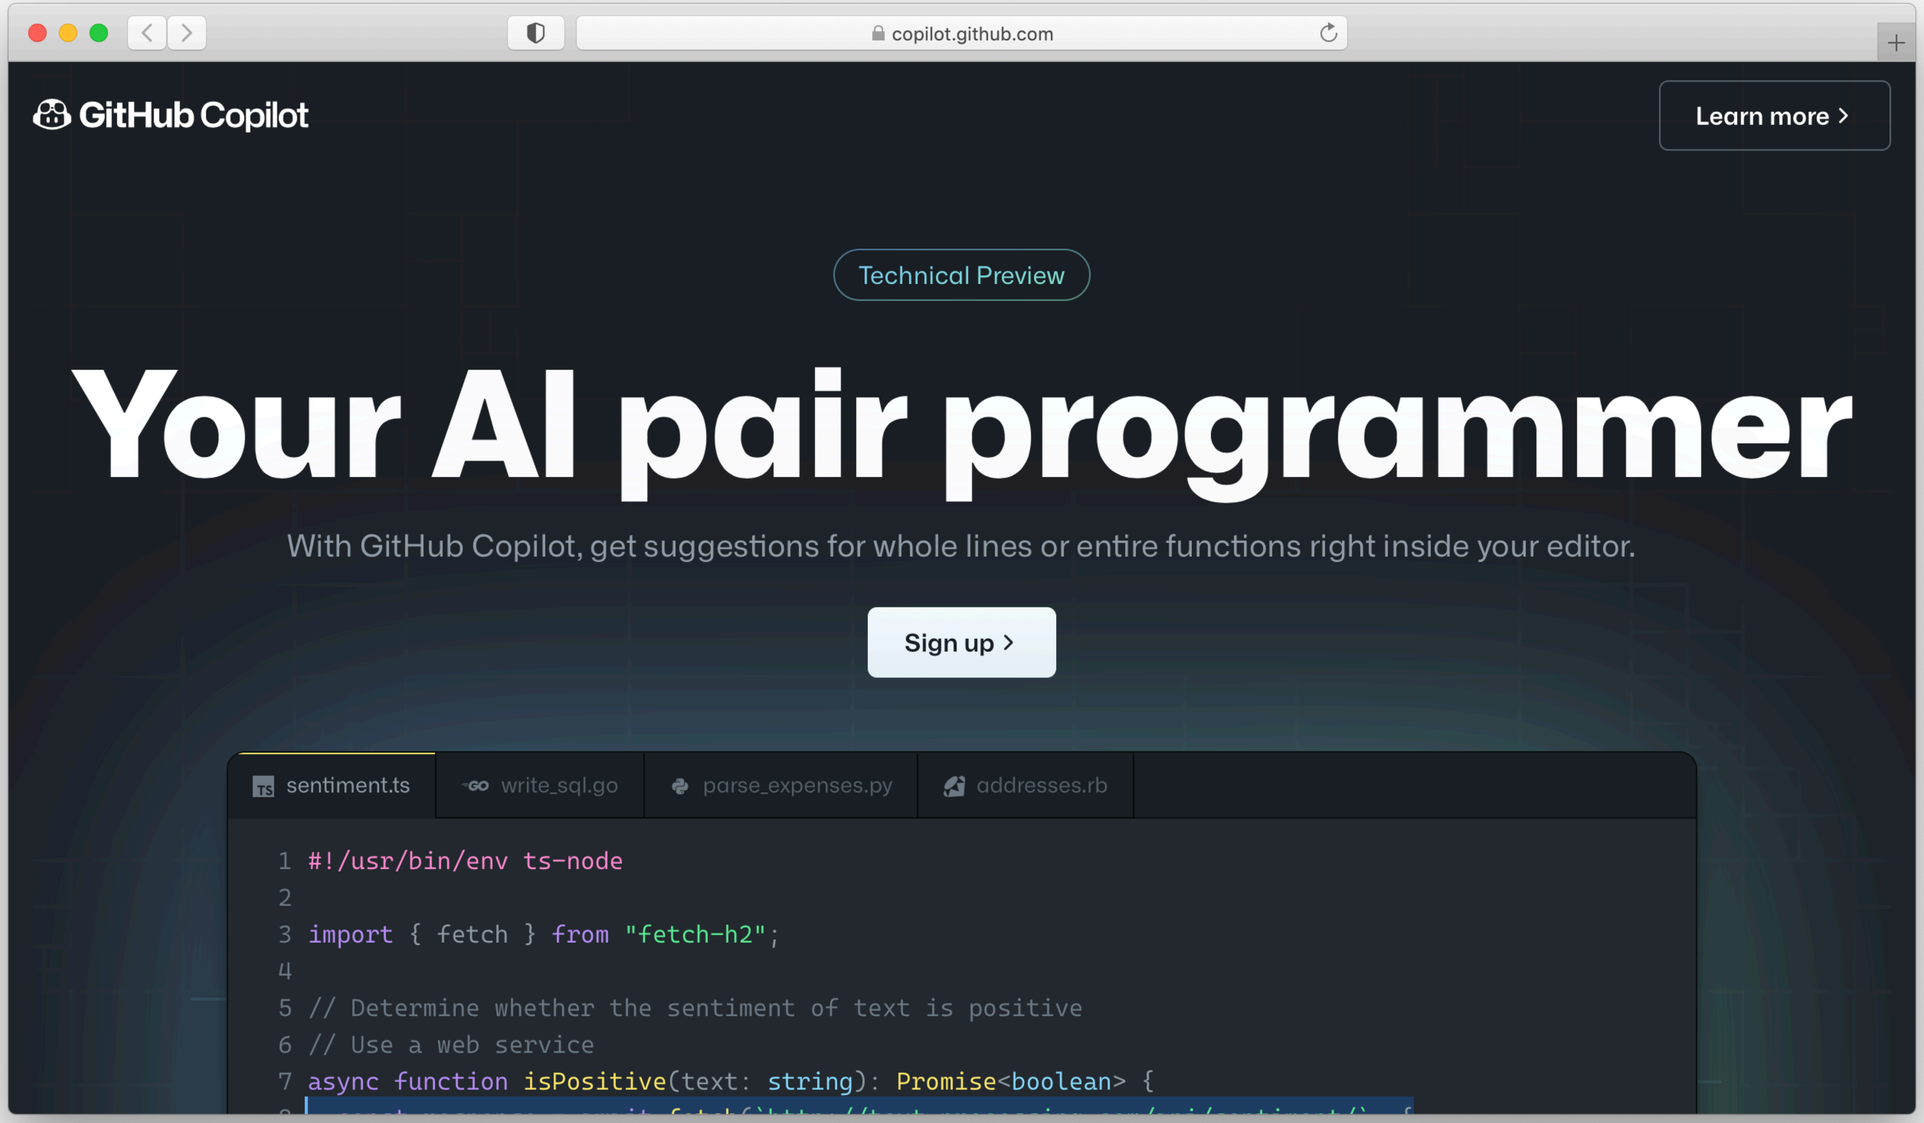
copilot.github.com
GitHub Copilot
Learn more >
Technical Preview
Your AI pair programmer
With GitHub Copilot, get suggestions for whole lines or entire functions right inside your editor.
Sign up >
sentiment.ts
write_sql.go
parse_expenses.py
addresses.rb
#!/usr/bin/env ts-node
import { fetch } from "fetch-h2";
// Determine whether the sentiment of text is positive
// Use a web service
async function isPositive(text: string): Promise<boolean> {


IMAGENET

Photo Credit: Elijah Macleod

Photo Credit: Bence Balla-Schottner

Photo Credit: Erwan Hesry

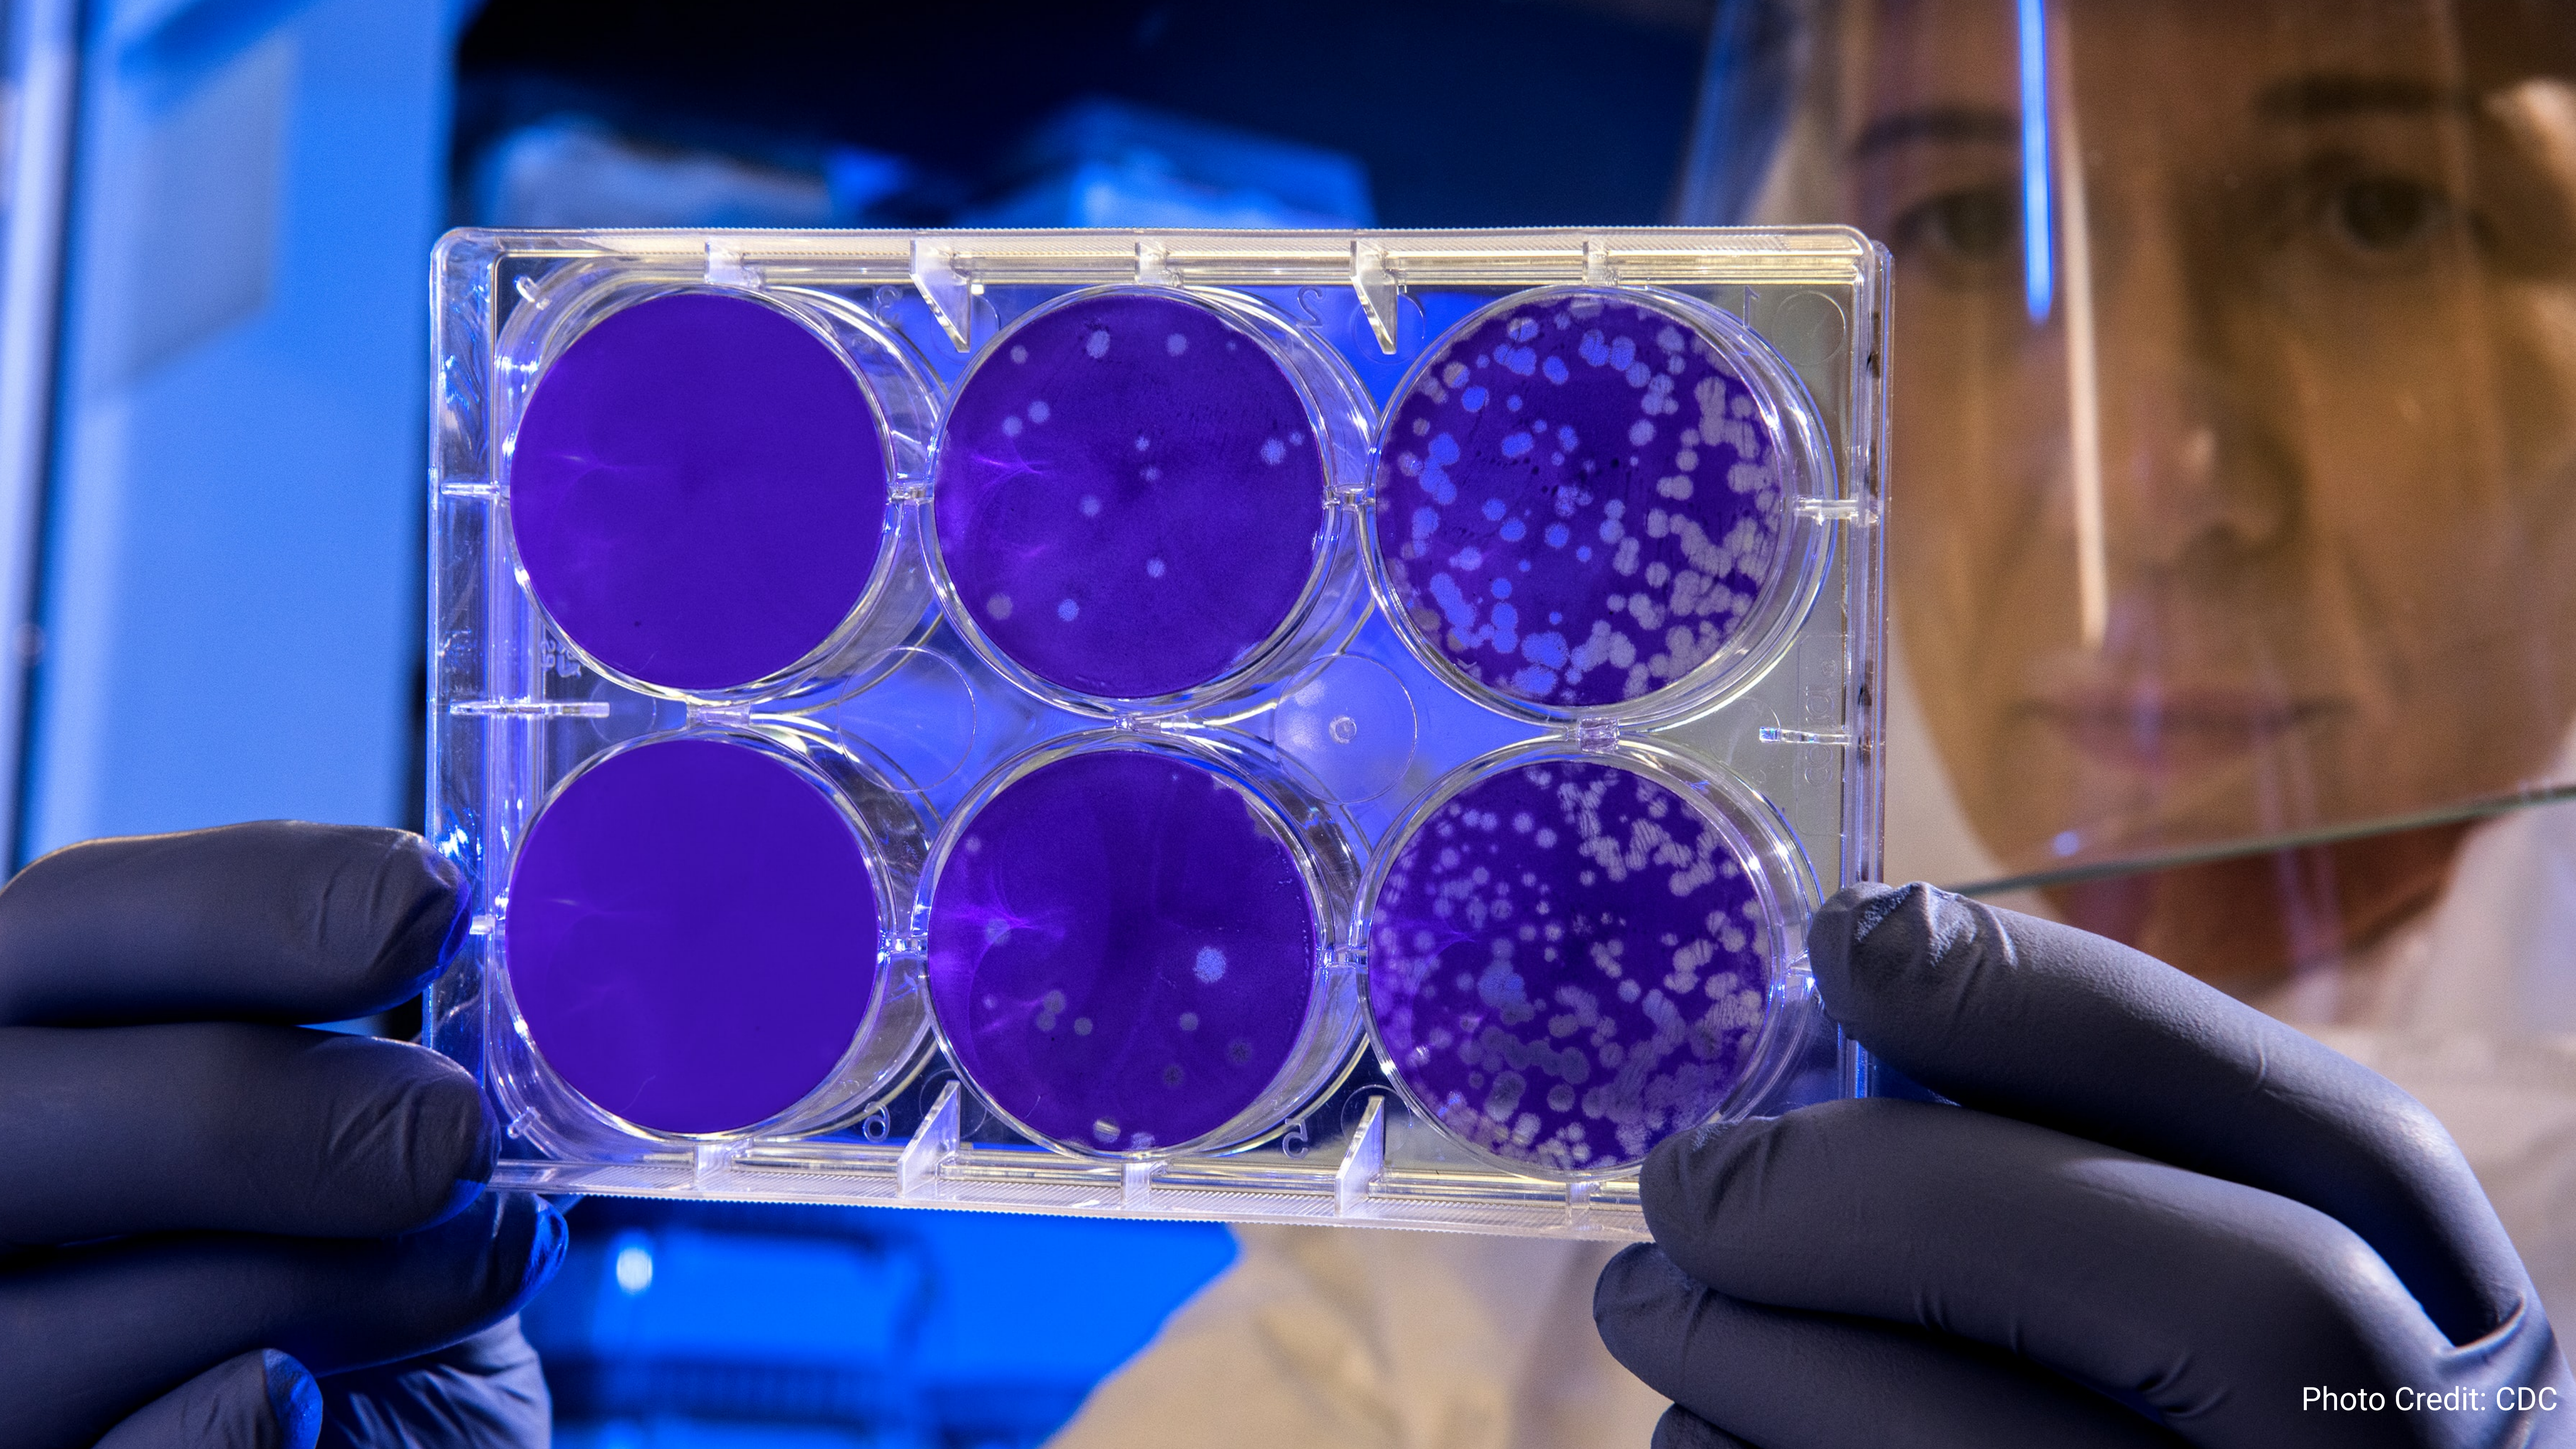

Photo Credit: CDC


Meet the cast:
ABCD
EFGHIJK
LMNOP
QRSTUV
WXYZ


Photo Credit: Patrick Fore

Photo Credit: Ross Findon

Photo Credit: Emiliano Bar

Photo Credit: Micaela Parente

https://arxiv.org/pdf/1707.08945.pdf

Photo Credit: Jesse Martini

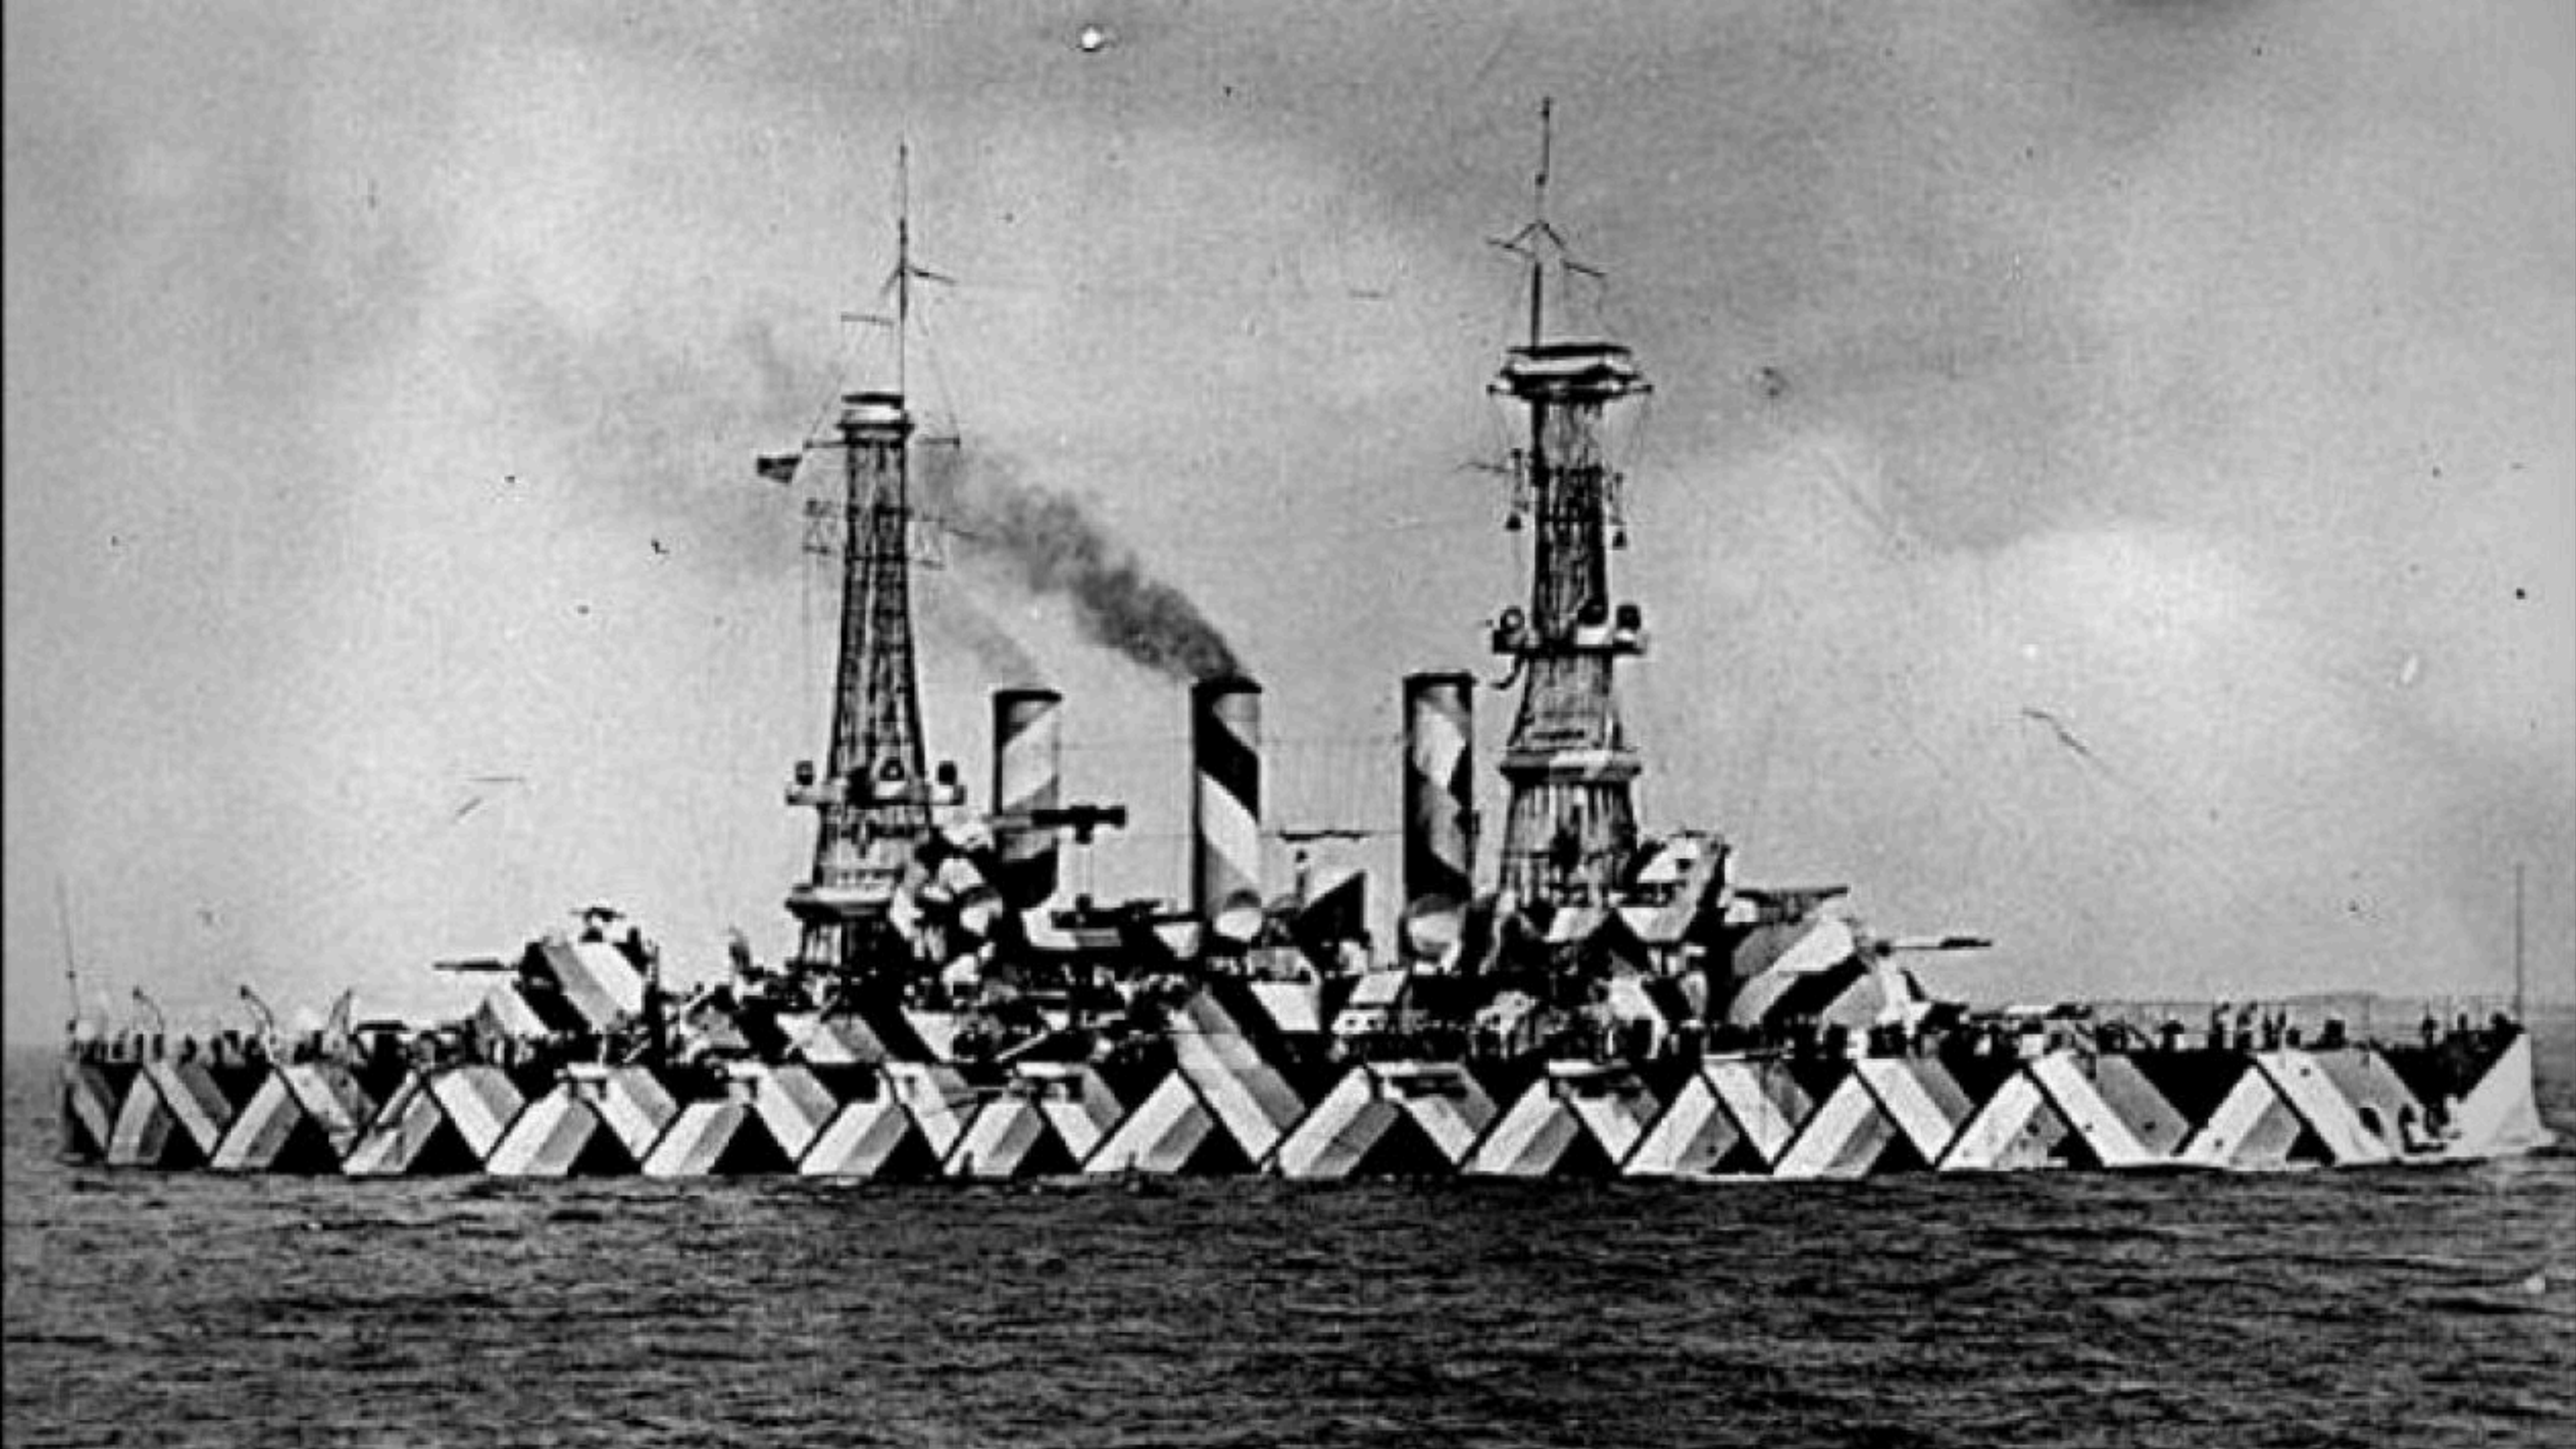

Photo Credit: Fabian Gieske

Photo Credit: Phil Shaw

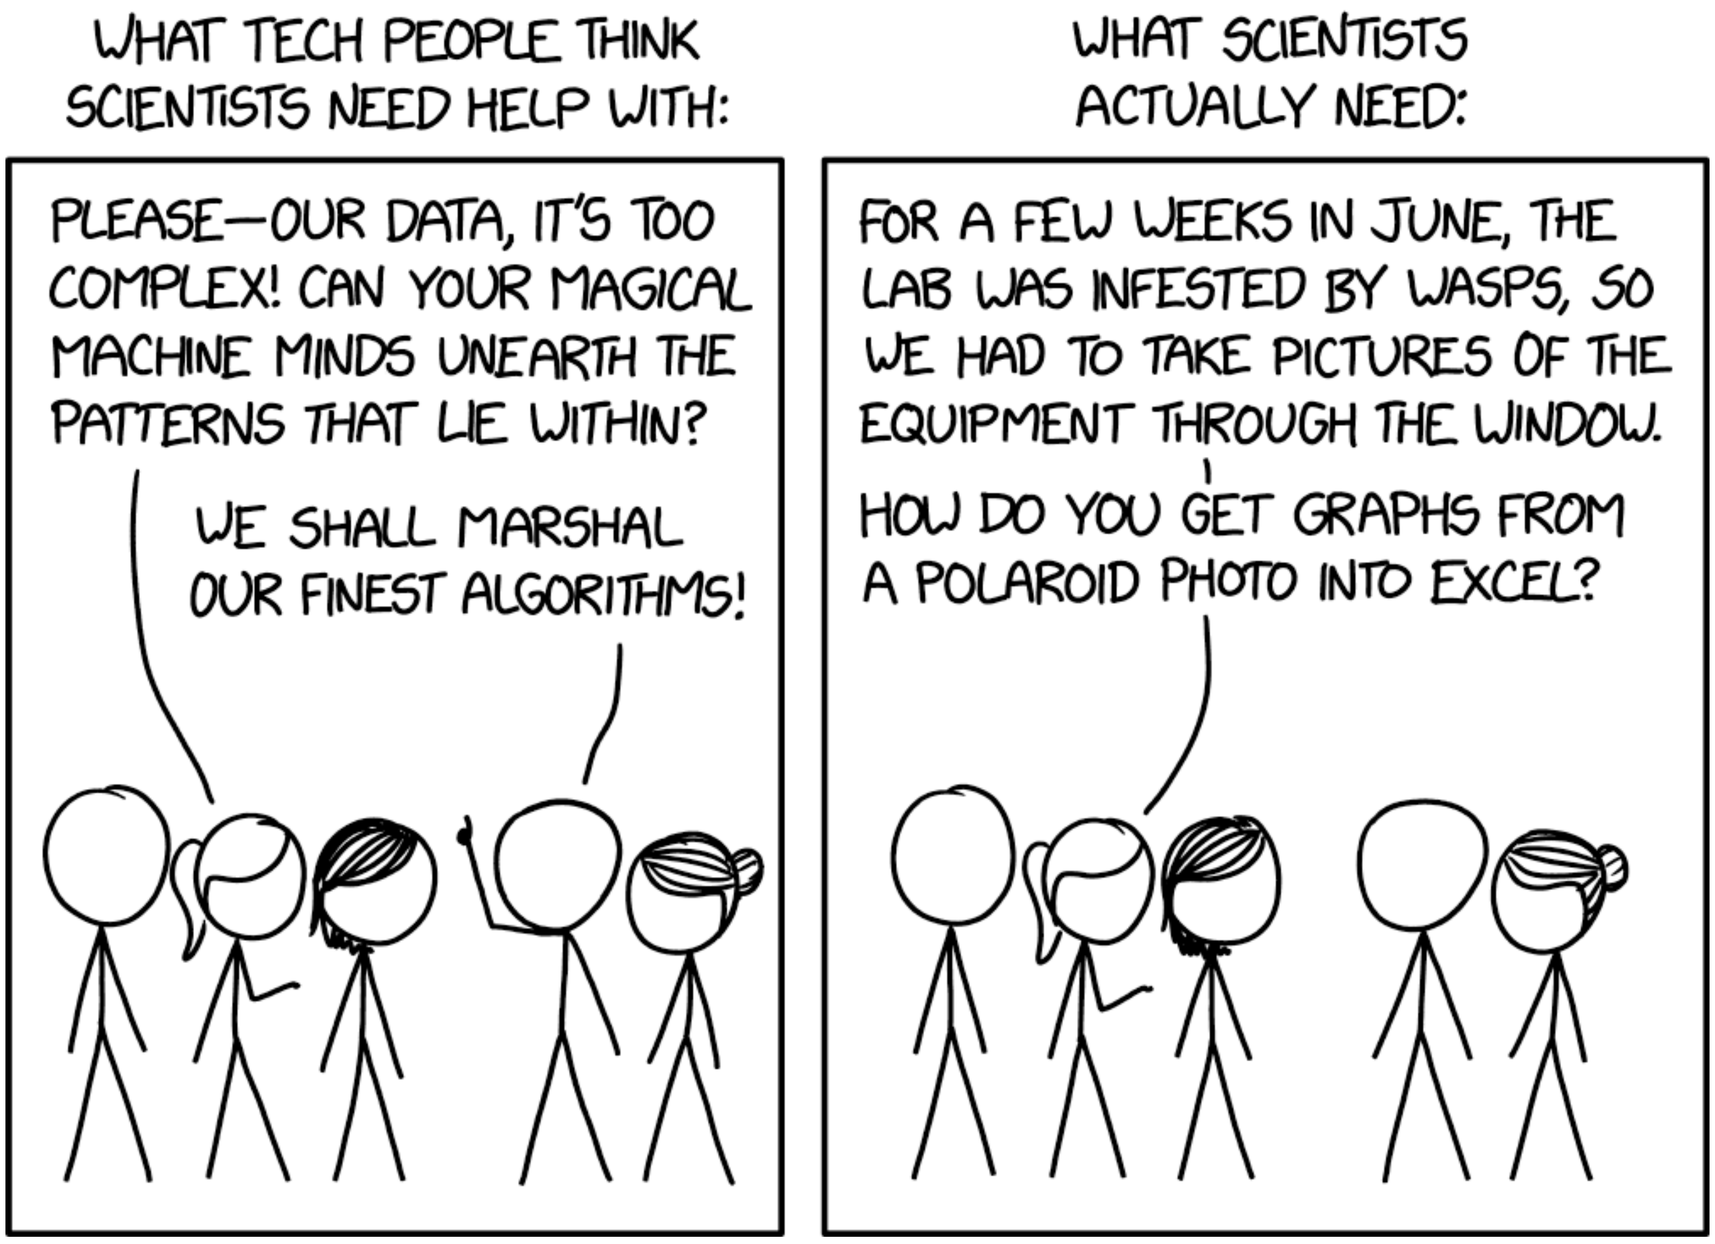


Photo Credit: Oscar Nord

Thank You!

Dr Alasdair Allan
@aallan

Copyright Notice

www.tinyML.org