tinyML® Talks

Enabling Ultra-low Power Machine Learning at the Edge

“tinyML: Enabling Ultra-low Power Always-On Computer Vision at Qualcomm”

Ravishankar Sivalingam - Qualcomm AI Research

October 7, 2021

www.tinyML.org

tinyML Talks Sponsors and Strategic Partners

AONdevIces
tinyML Strategic Partner

arm
tinyML Strategic Partner

Deeplite
tinyML Strategic Partner

EDGE IMPULSE
tinyML Strategic Partner

emza visual sense
tinyML Strategic Partner

GREENWAVES TECHNOLOGIES
tinyML Strategic Partner

LatentAI
Adaptive AI for a Smarter Edge
tinyML Strategic Partner

HOTG
tinyML Strategic Partner

imagimob
tinyML Strategic Partner

maxim integrated
NOW PART OF ANALOG DEVICES

Qeexo
tinyML Strategic Partner

Qualcomm
tinyML Strategic Partner

RealityAI
tinyML Strategic Partner

seeed studio
The IoT Hardware Enabler
tinyML Strategic Partner

SensiML
tinyML Strategic Partner

SynSense
tinyML Strategic Partner

SYNTIANT
tinyML Strategic Partner

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

Arm: The Software and Hardware Foundation for tinyML

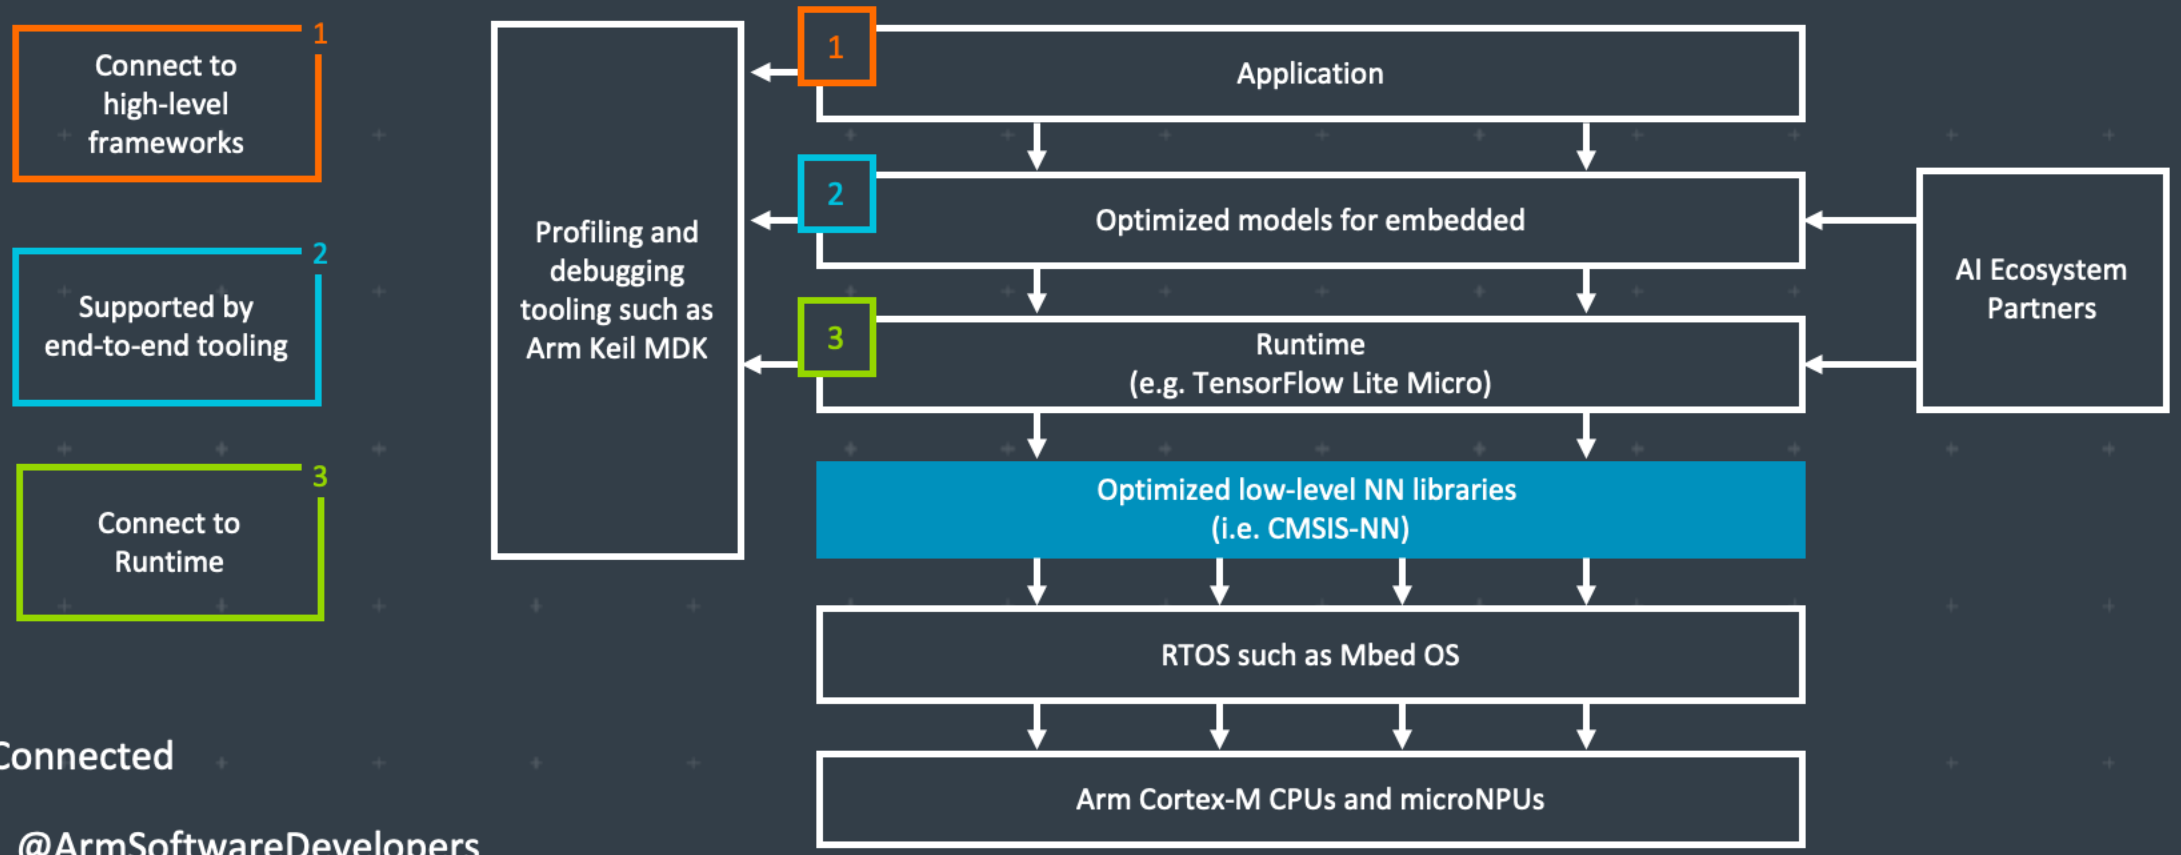

Stay Connected

@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

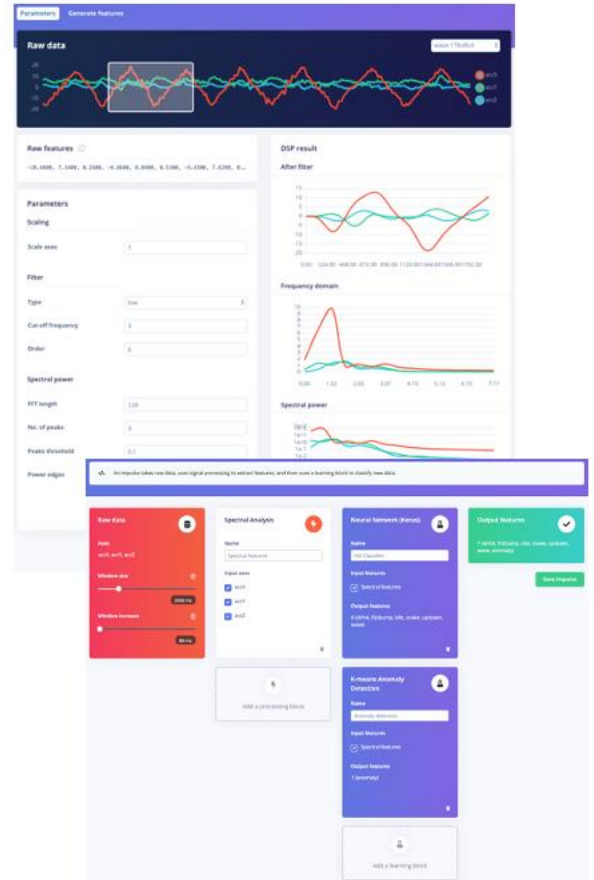

www.edgeimpulse.com

emza
visual sense

The Eye in IoT

Edge AI Visual Sensors

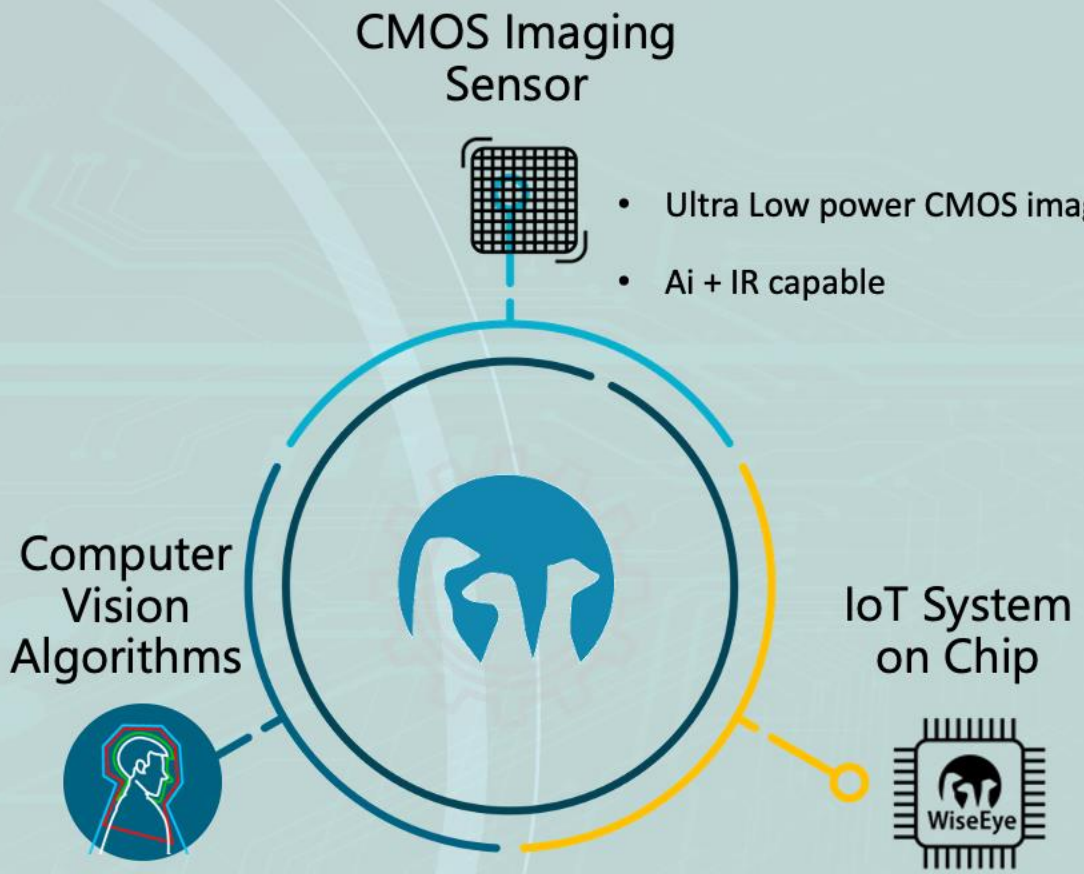

- Ultra Low power CMOS imager
- Ai + IR capable

- 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

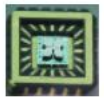

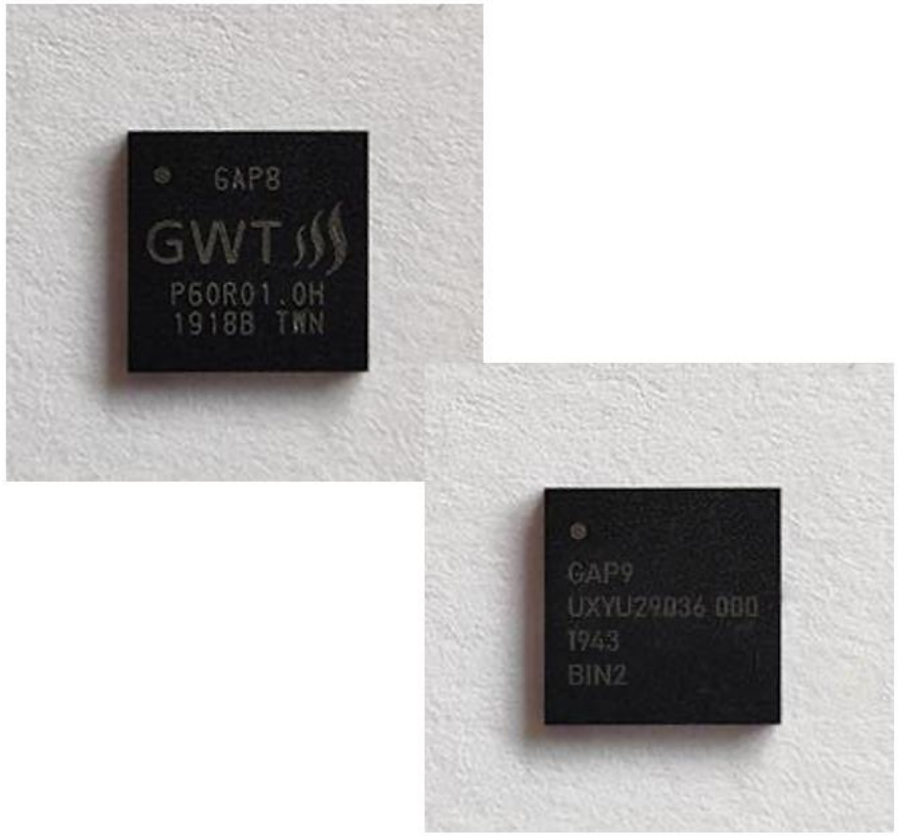

Wearables / Hearables

Battery-powered consumer electronics

IoT Sensors

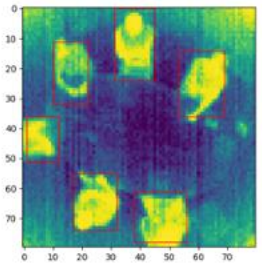

GREENWAVES TECHNOLOGIES

Distributed infrastructure for TinyML apps

HOTG
Decoupling intelligence

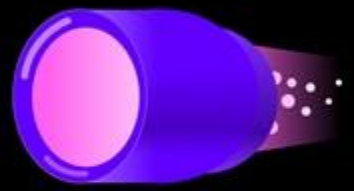

Develop at warp speed

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

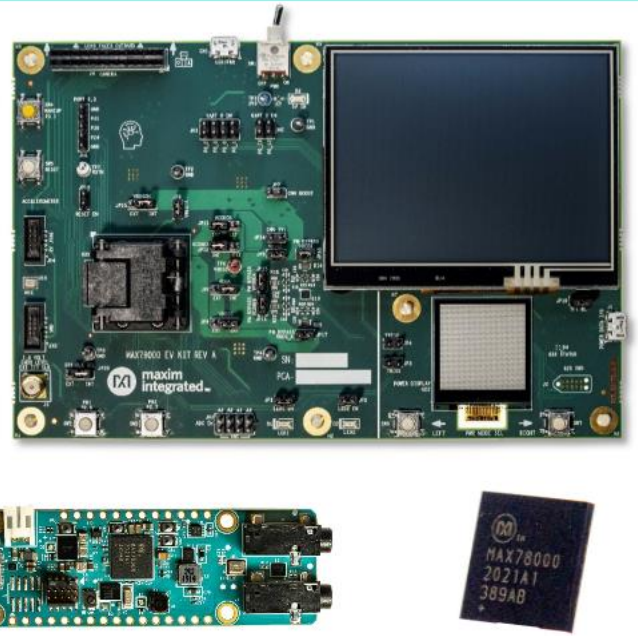

The new MAX78000 implements AI inferences at low energy levels, enabling complex audio and video inferencing to run on small batteries. Now the edge can see and hear like never before.

www.maximintegrated.com/MAX78000

Low Power Cortex M4 Micros

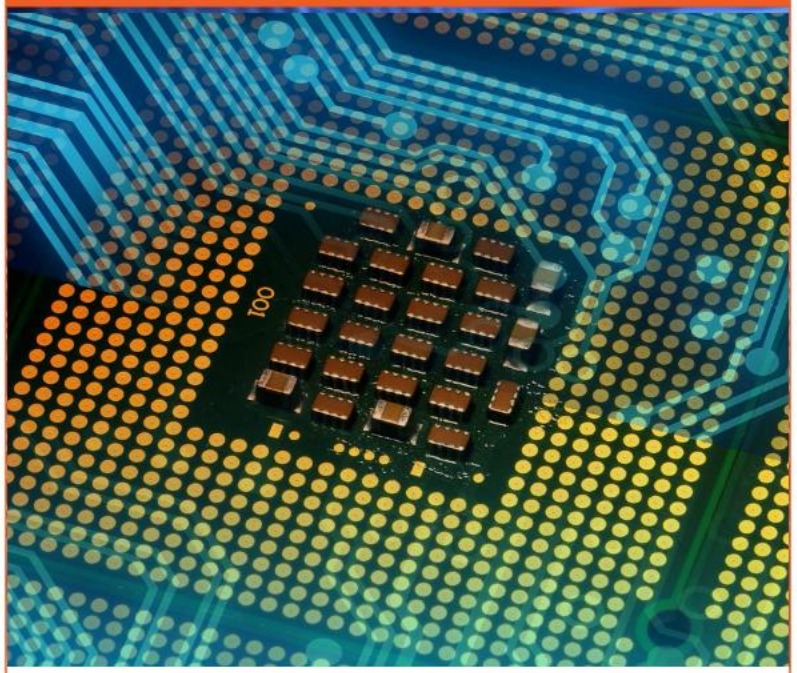

Large (3MB flash + 1MB SRAM) and small (256KB flash + 96KB SRAM, 1.6mm x 1.6mm) Cortex M4 microcontrollers enable algorithms and neural networks to run at wearable power levels.

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

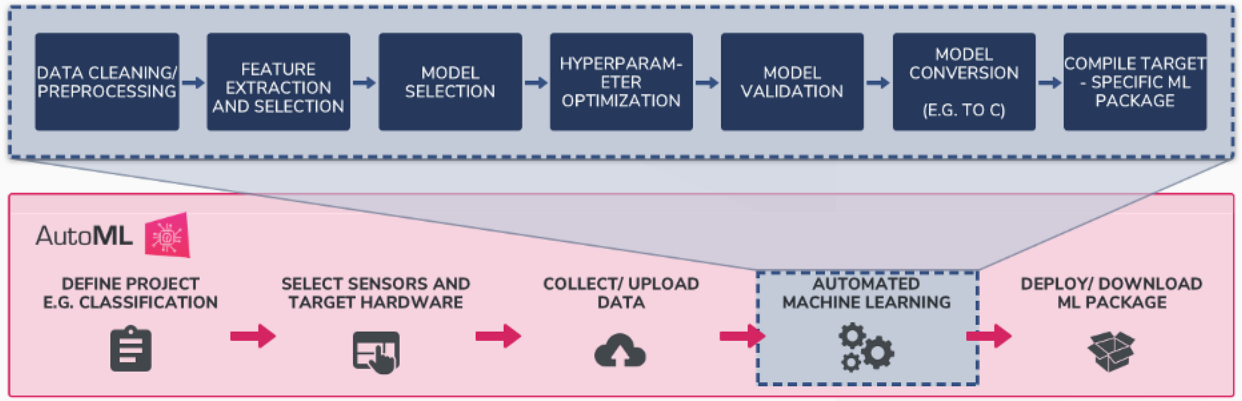

For more information, visit: www.qeexo.com

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

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

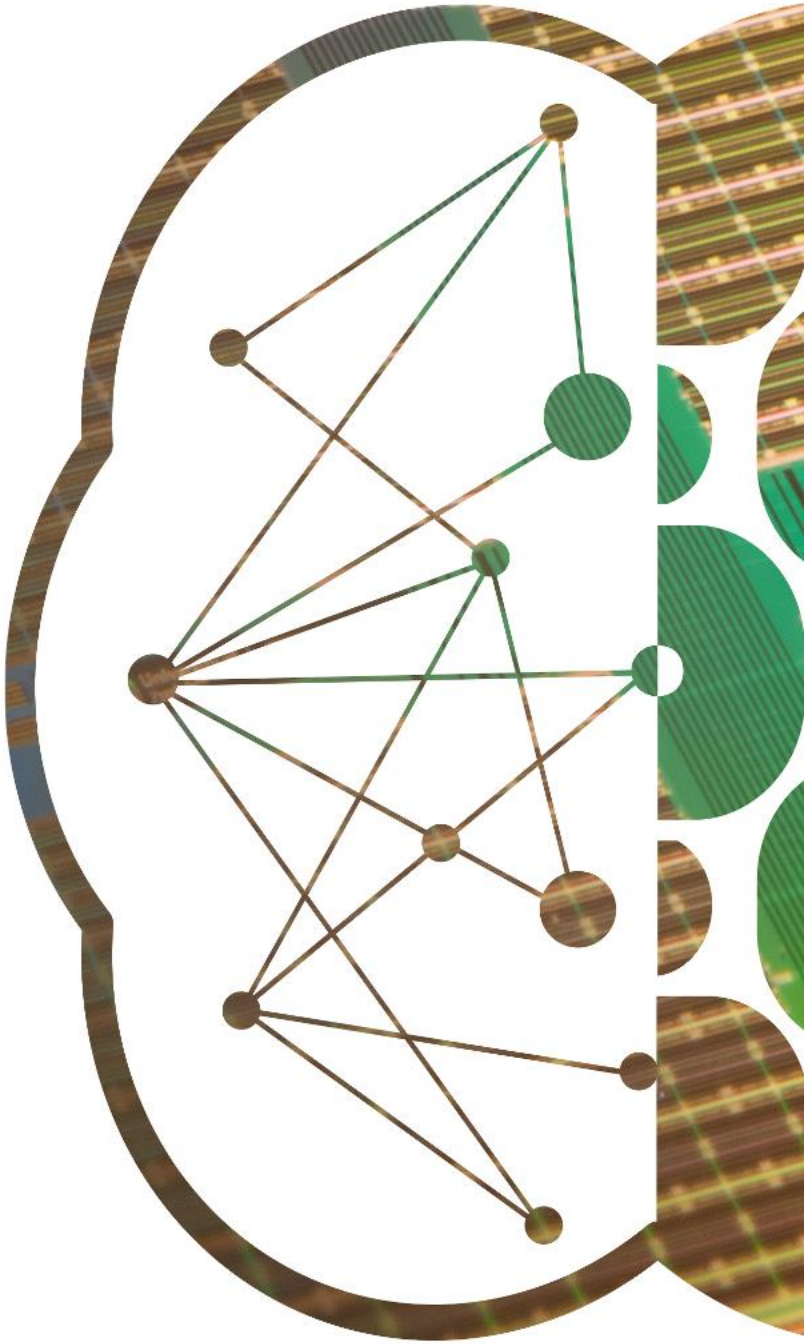

SYNTIANT

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

Neural Decision Processors

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

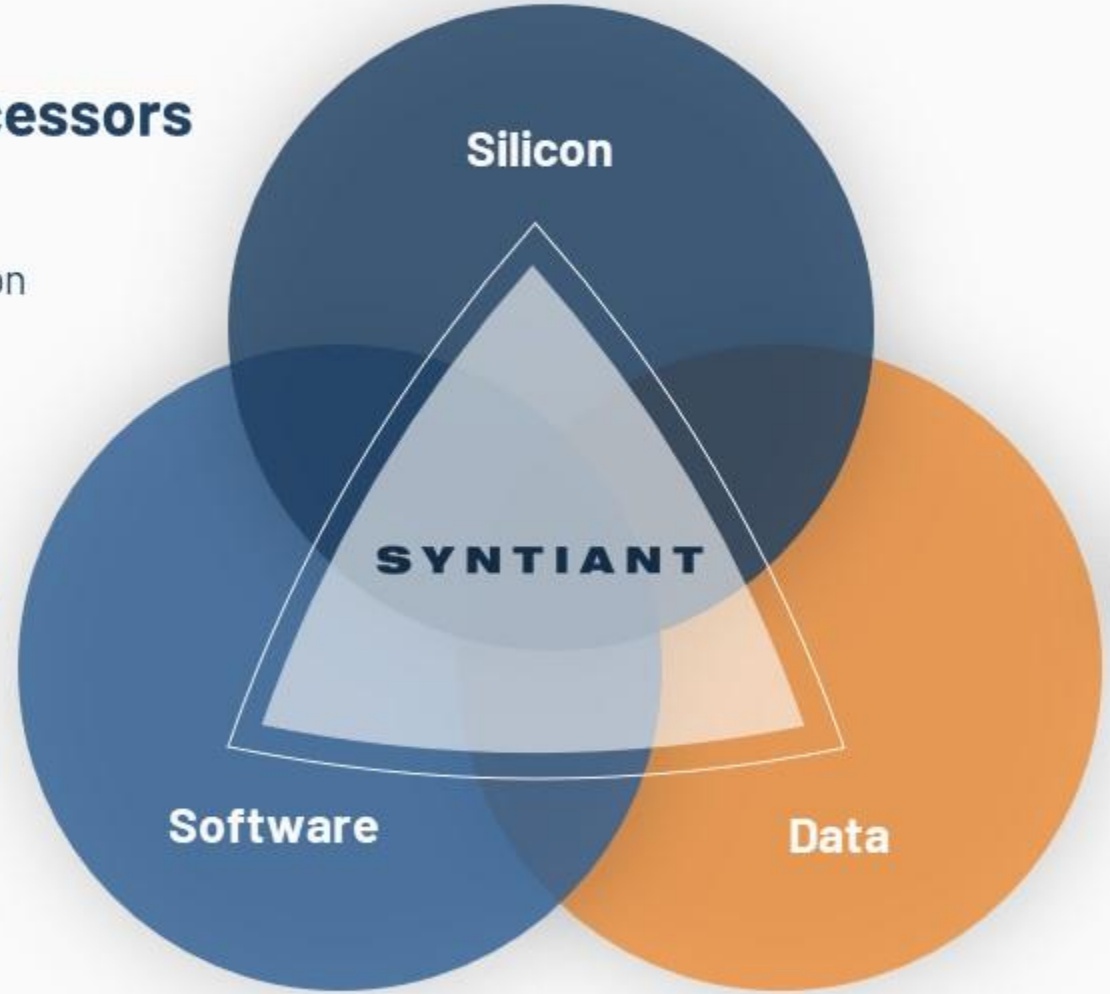

ML Training Pipeline

- Enables Production Quality Deep Learning Deployments

Data Platform

- Reduces Data Collection Time and Cost
- Increases Model Performance

partners@syntiant.com

www.syntiant.com

TINY ML ASIA

LIVE ONLINE **November 2-5, 2021**

(9-11:30 am China Standard time)

https://www.tinyml.org/event/asia-2021/

Technical Program Committee

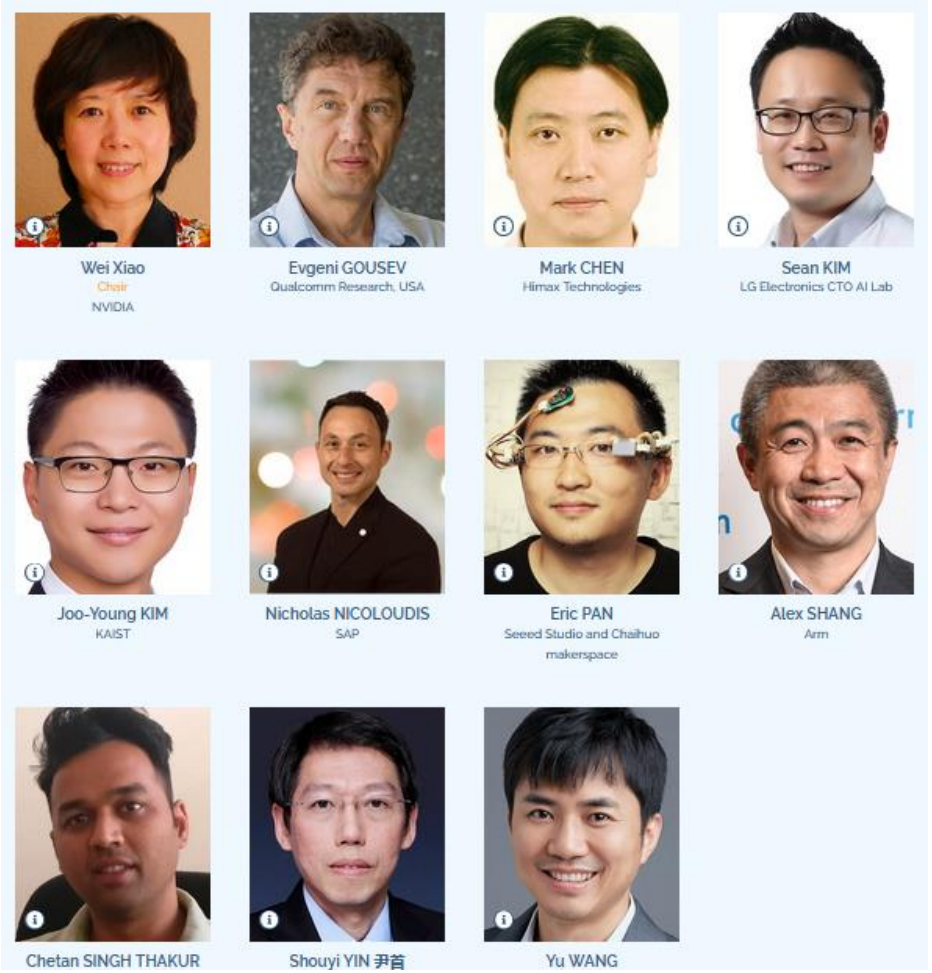

Register today!

Free event courtesy of our sponsors and strategic partners

AONdevices arm EDGE IMPULSE emza visual sense

GREENWAVES TECHNOLOGIES Groovety Inc. HOTG imagimob

LatentAI Adaptive AI for a Smarter Edge Qualcomm Qeexo RealityAI

seeed The IoT Hardware Enabler SensiML SynSense SYNTIANT

More sponsorships are available: sponsorships@tinyML.org

collaboration with hackster.io

AN AVNET COMMUNITY

Focus on:
(i) developing new use cases/apps for tinyML vision; and (ii) promoting tinyML tech & companies in the developer community

Submissions accepted until September 17th, 2021

Winners announced on October 5th, 2021 ($6k value)

Sponsorships available: *sponsorships@tinyML.org*

https://www.hackster.io/contests/tinyml-vision

Next tinyML Talks

Date	Presenter	Topic / Title
Tuesday, October 12	**David Schwarz,** User Success Engineer, Edge Impulse	AutoML + TinyML with Edge Impulse's EON Tuner

Webcast start time is 8 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

tinyml.org/forums

youtube.com/tinyml

Please use the Q&A window for your questions

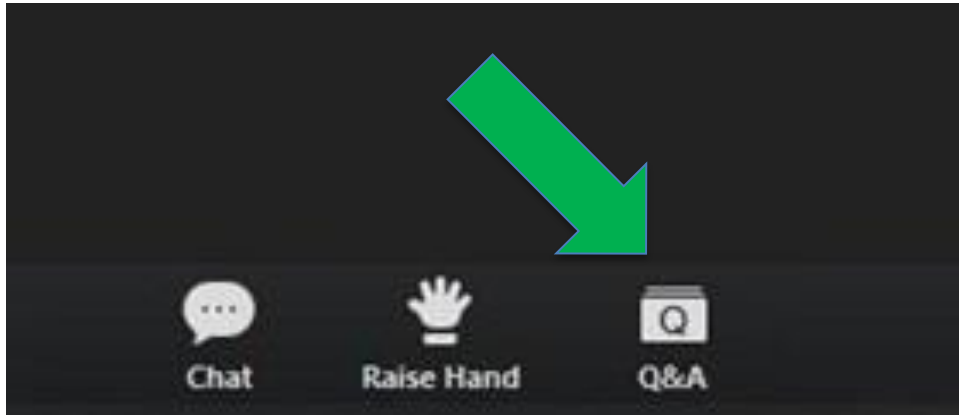

TINY ML TALKS webcast

Ravishankar Sivalingam

Ravishankar Sivalingam obtained his M.S. in Computer Science, M.S. in Electrical Engineering, and Ph.D. in Electrical Engineering from the University of Minnesota, Twin Cities, specializing in sparse modeling for computer vision. He was a founding member of the Computational Intelligence team at 3M Corporate R&D, Minnesota, where he applied machine learning and computer vision to applications in the diverse businesses operated by 3M, ranging from biometrics to dental & orthodontic products. He also worked at June Life, a startup bringing computer vision technology to the smart kitchen of the future. Currently, Ravi develops ML algorithms for ultra-low power computer vision at Qualcomm AI Research, in Santa Clara, CA.

10/07/2021

@qualcomm

Qualcomm

tinyML - Enabling Ultra-low Power Always-On Computer Vision with Qualcomm® QCC112

Ravishankar Sivalingam, PhD
Sr. Staff Engineer/Mgr.
Qualcomm® Artificial Intelligence (AI) Research
Qualcomm Technologies, Inc.

TINY ML

tinyML & “Why the Future of Machine Learning is Tiny*”

tinyML is broadly defined as

- machine learning architectures, techniques, tools and approaches
- capable of performing on-device analytics
- for a variety of sensing modalities (vision, audio, motion, environmental, human health monitoring etc.)
- at “mW” (or below) power range
- targeting predominately battery operated devices (IoT, bioelectronics, …)

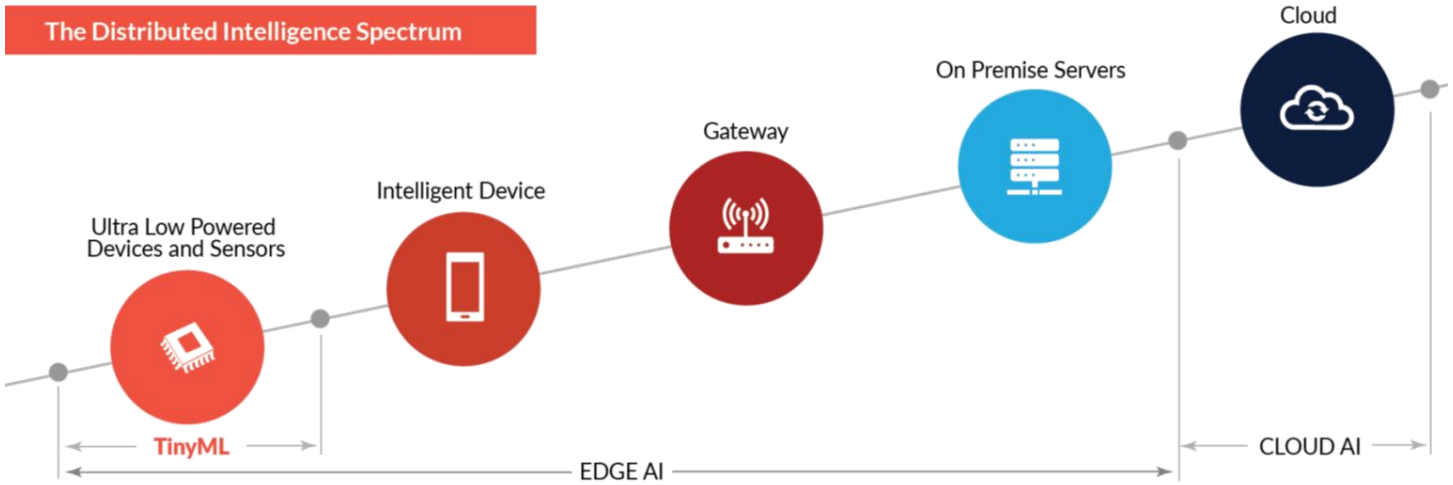

ABIresearch

ABI Research:
TinyML market – est. 1B tinyML devices shipped in 2024, installed based of 5.4B tinyML devices in 2026.

* https://petewarden.com/2018/06/11/why-the-future-of-machine-learning-is-tiny/- Pete Warden

tinyML is growing fast

	2019 Summit *(March 2019)*	2020 Summit *(Feb 2020)*	2021 Summit *(March 2021)*
Attendees	**160**	**400+**	**5000+**
Companies	**90**	**172**	**300+**
LinkedIn members	**0**	**798**	**2258**
Meetups members	**0**	**1140**	**6511**
YouTube subscribers	**0**	**0**	**5120**

Also started in Asia: tinyML WeChat and BiliBili

tinyML Foundation

tinyML.org

300+ videos on
youtube.com/tinyML

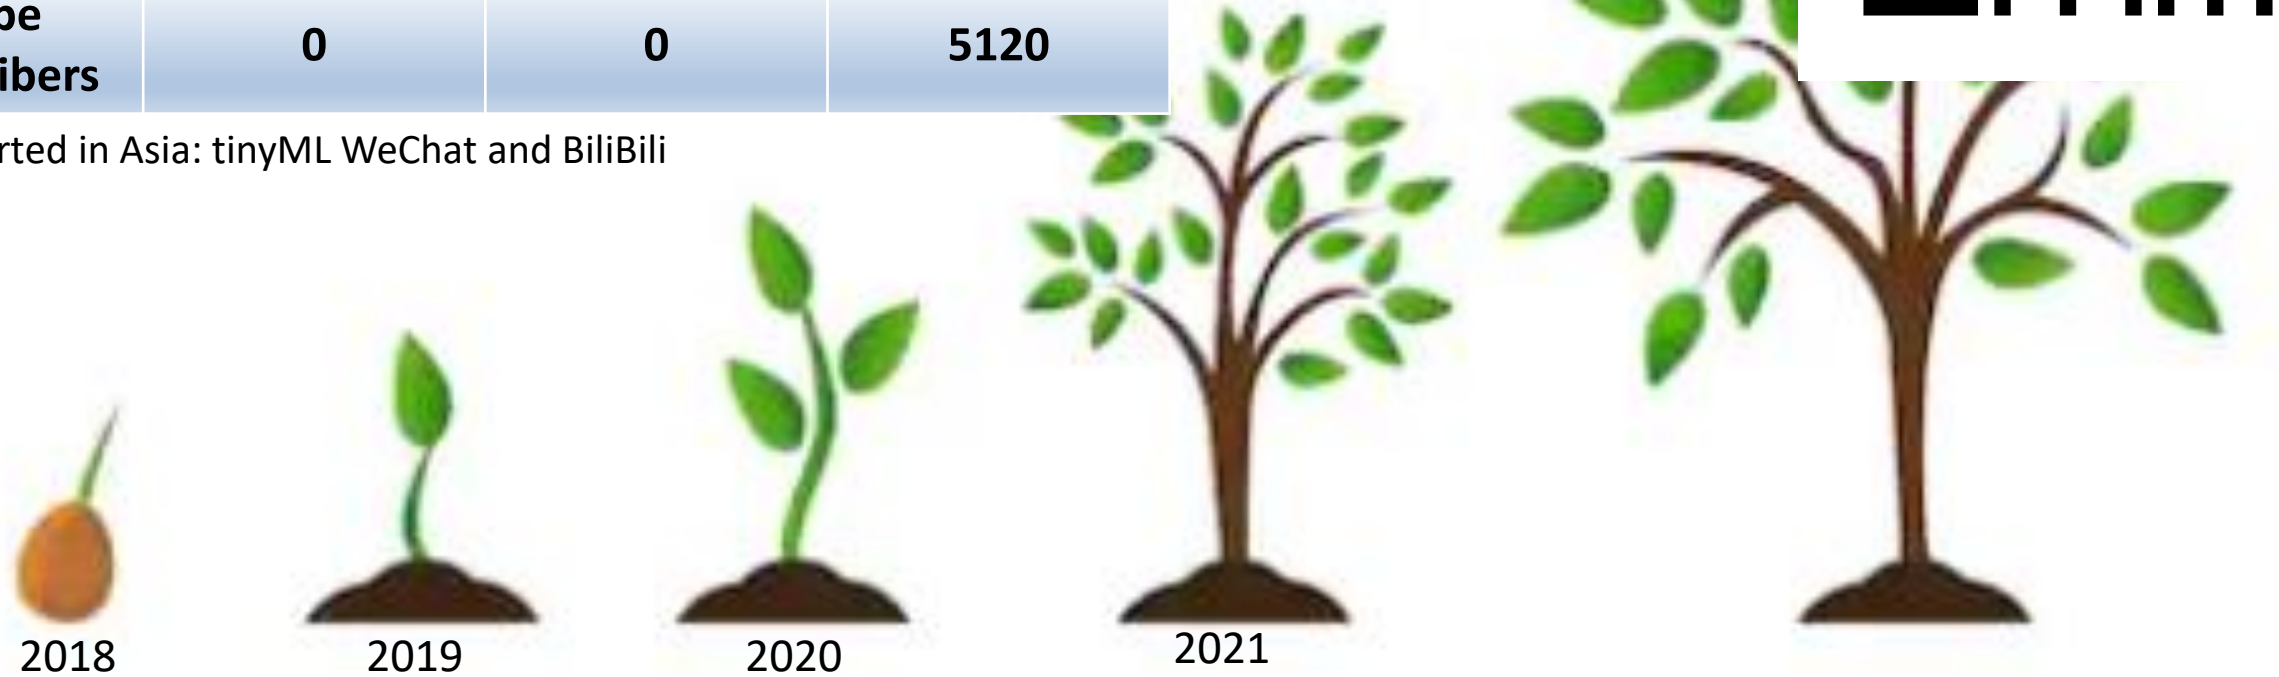

Meetups

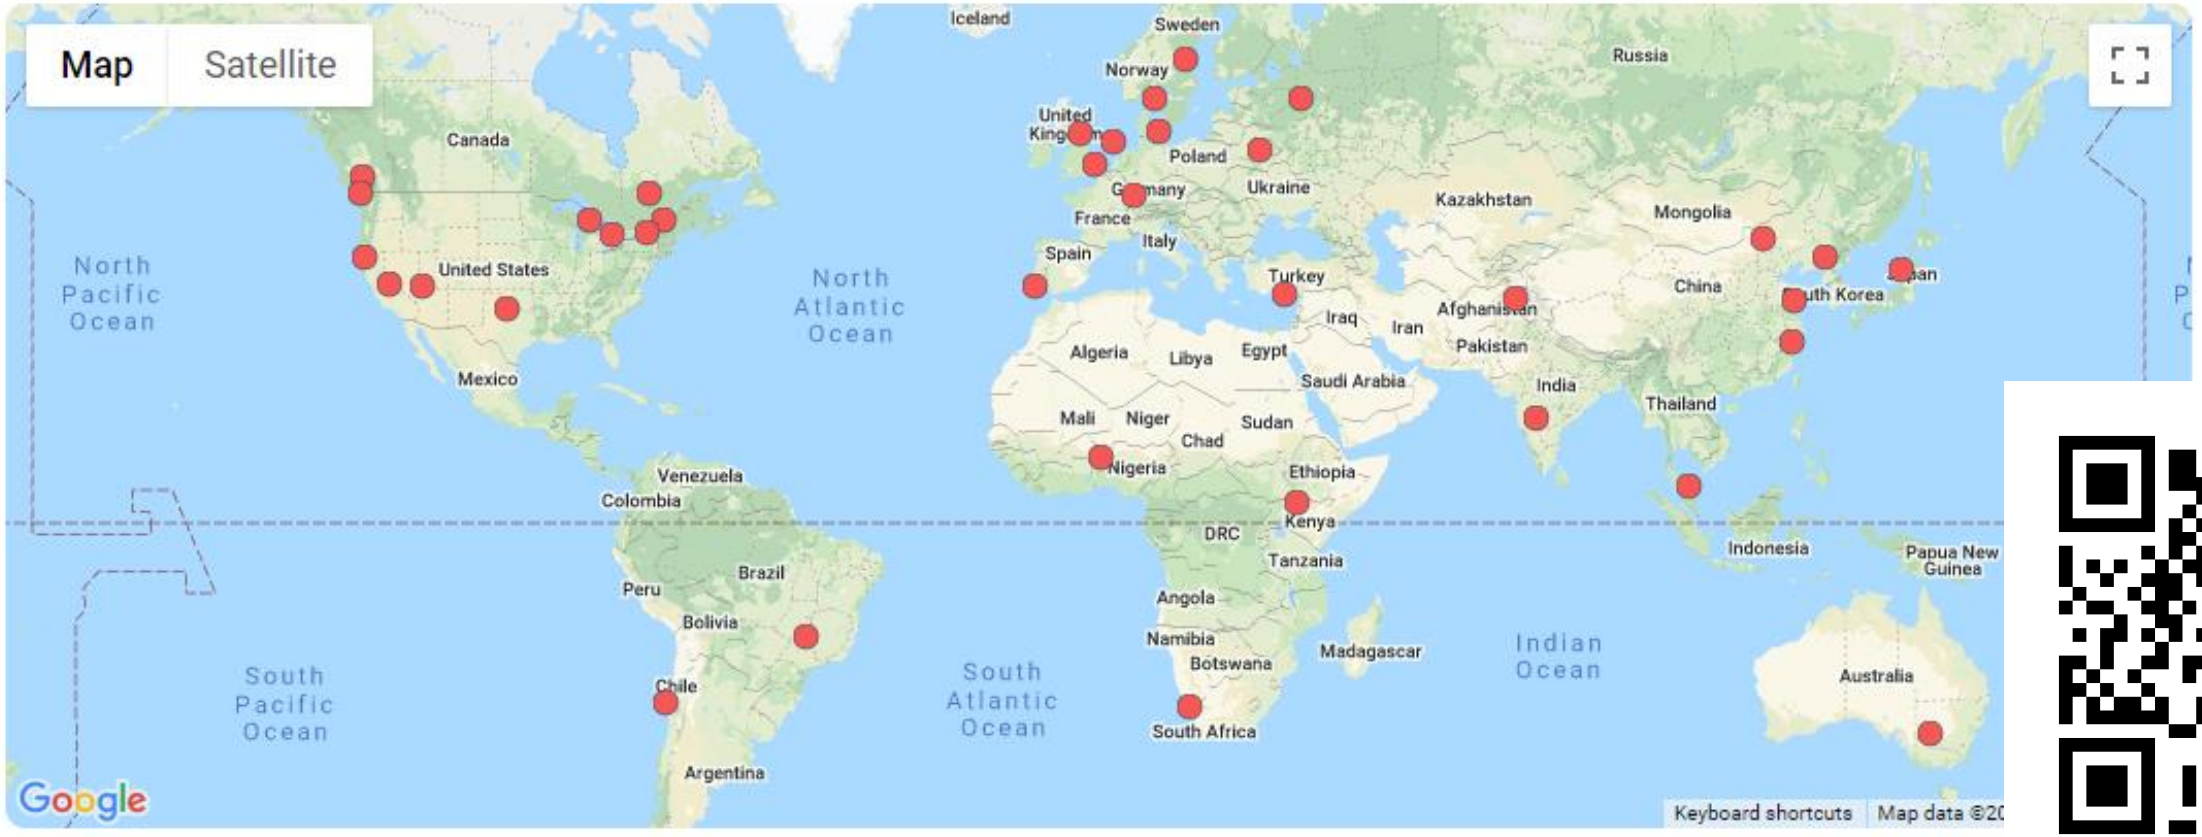

tinyML **meetup.com/pro/tinyML**

Members	Groups	Countries
6,511	36	29

Vision Challenge-2021 Winners

https://www.hackster.io/contests/tinyml-vision

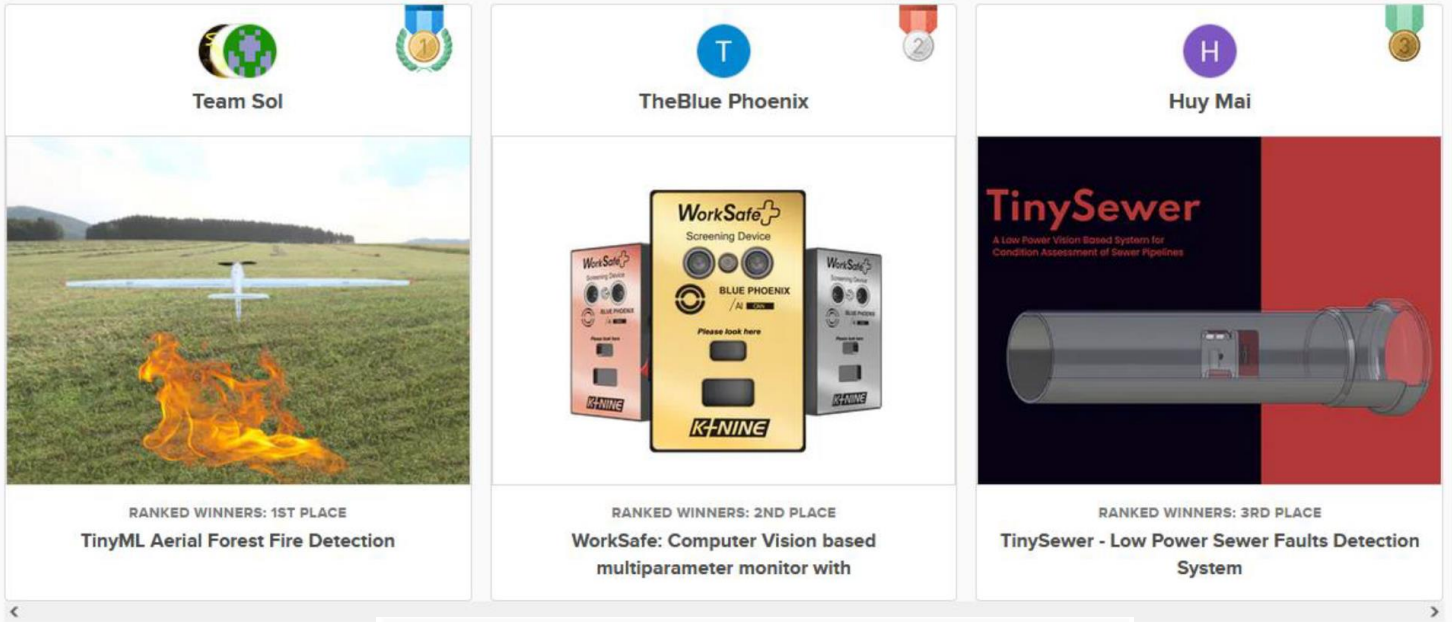

Hackster Impact Prize

The winner was awarded a $250 Gift card + Video interview + More ($530 value)

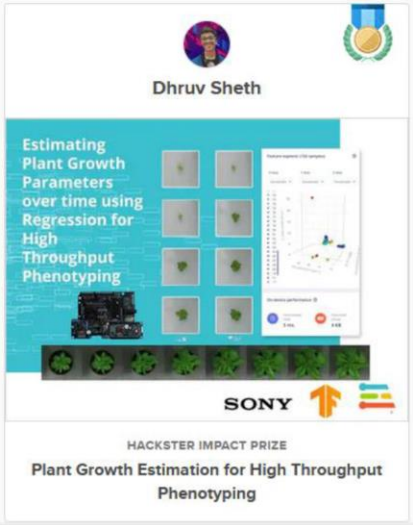

Honorable Mention

Each runner up was awarded a $500 Gift Card ($500 value)

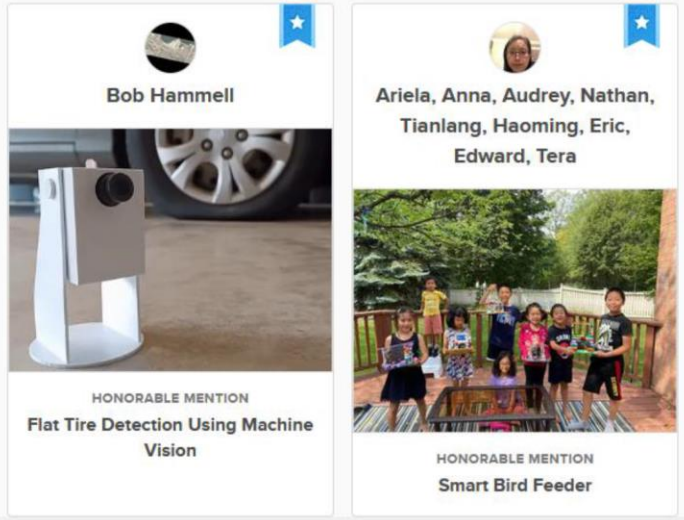

NEW: tinyML Success Stories Series

- Inspiration and educational series
- 1 hour on-line LIVE and interactive interviews with tinyML “movers and shakers”
- Recent M&A stories, VC views, new products, breakthrough research in the academia
- Starting in October 2021
- Hosted by renowned Kurt Keutzer (UC-Berkeley) and Chris Rowen (CISCO)

NEW: tinyML Panel Sessions

- tinyML for Conservation Panel – scheduled for Oct 19th

  Moderator: Dan Situnayake, Edge Impulse

  Panelists:

Cirra Maina

DeKUT

Sam Kelly

Conservation X Labs

Stephanie O'Donnell

WILDLABS

- tinyML Vision Challenge Winners' Panel – late October/early November

TINY ML SUMMIT

- Venue: Hyatt (near SFO), Burlingame, Silicon Valley, CA
- Dates: March 28-30, 2022
  - + Exhibition
  - + 2nd tinyML Research Symposium
- Targeting 500 attendees (to ensure highest quality and networking value)
- Premier Sponsor: Edge Impulse
- Sponsorships@tinyML.org

Chair:
Adam Fuks
NXP

Vice-Chair:
Luca Benini
ETH-Zurich &
Univ. of Bologna

TINY ML

tinyML @ Qualcomm

Always-On UI and Contextual Awareness

Always-On technologies continues to deliver significant value across the whole ecosystem (> 1B units/year)

Always-On Touch

Always-On Voice

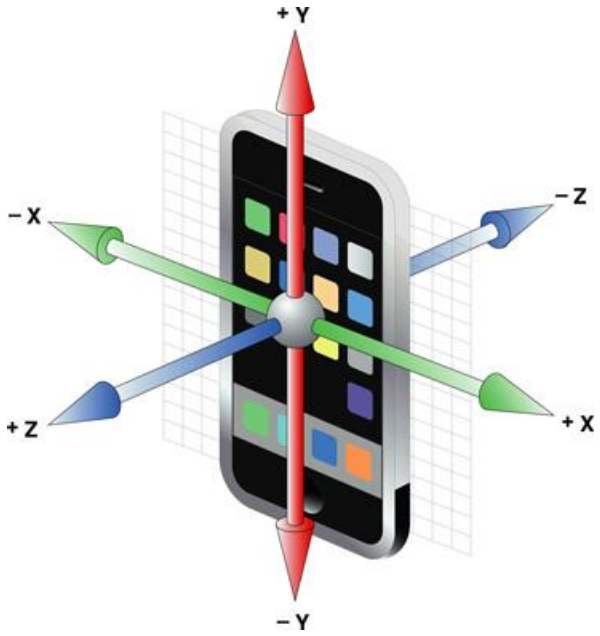

Always-On Motion

Human Insight is Dominated by Vision

83% of our external world perception is through vision

Yet, Always-On vision technologies have been extremely challenging

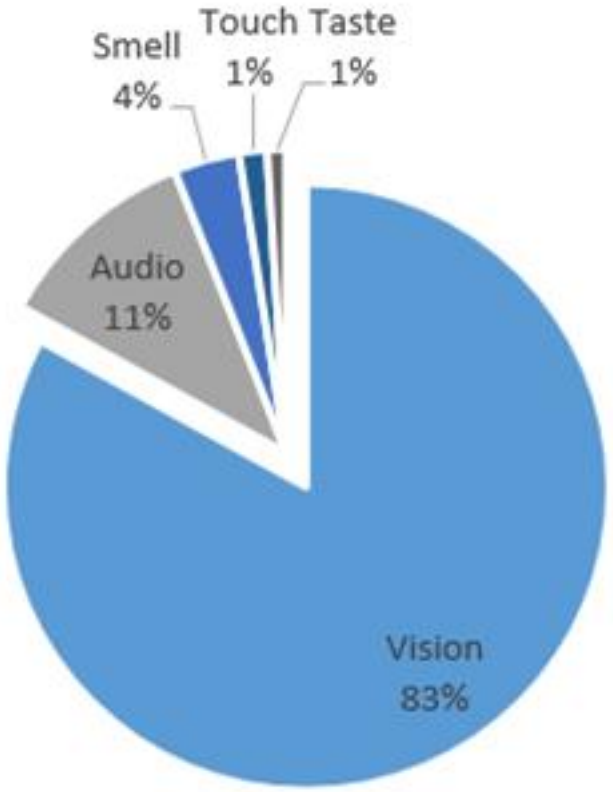

Source: *Hatwell, Y. (1994). Traitéde psychologie expérimentale. Paris, P.U.F.)*

Machine Vision Today

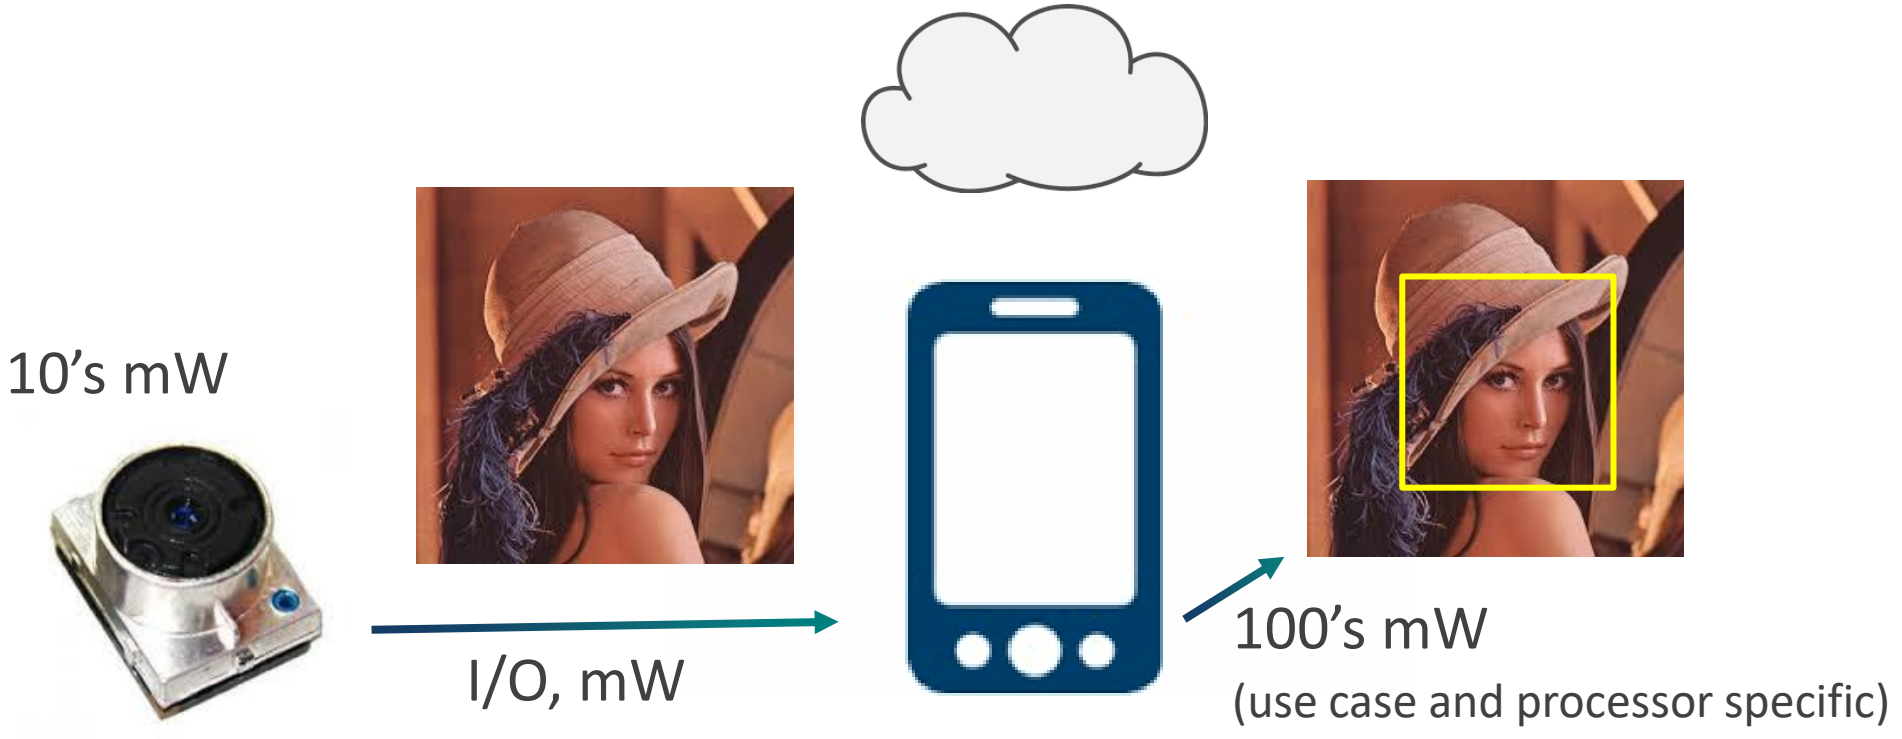

- Conventional approaches are very power hungry:
  - Processor and algorithms
  - Image sensor

Example: Gesture Algorithm Partitioning

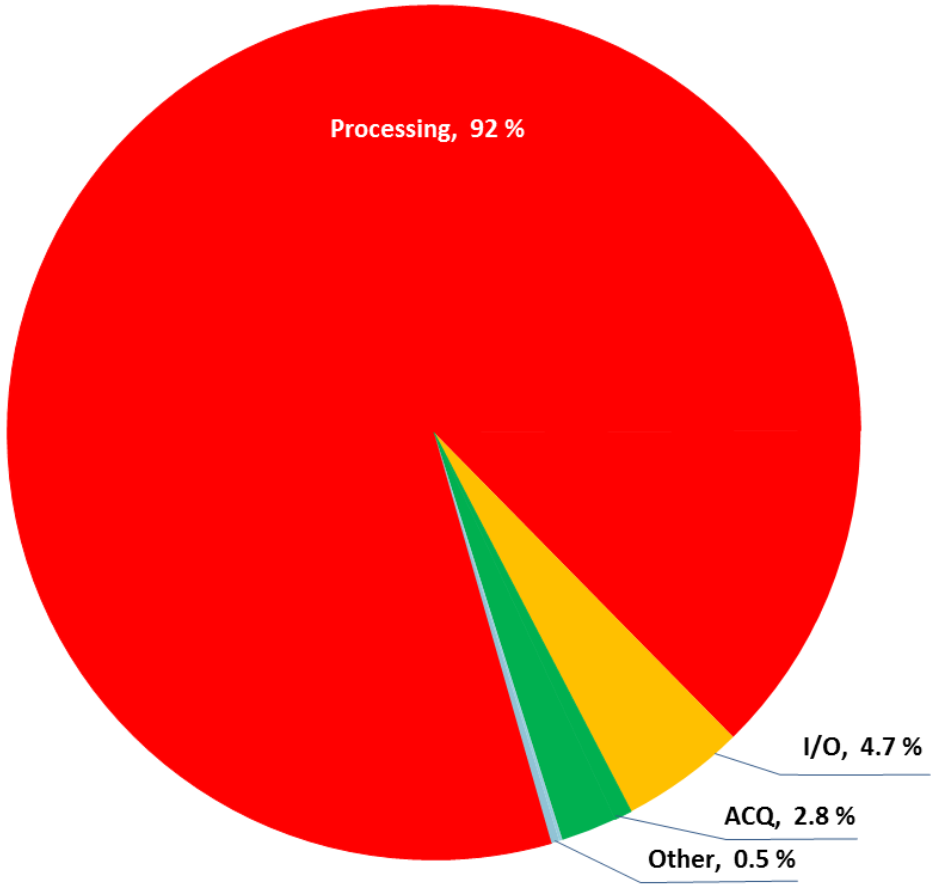

- Vast majority of power is consumed by algorithms running in the processor
- Typically 100's of mW

Vision will enhance many use cases across numerous verticals

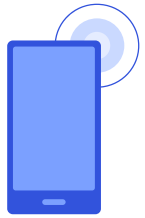

Smartphone

- Face-based auto-wake and auto-sleep
- Always-on trigger for other use cases
- Always-on trigger for iris authentication (removes multiple steps and user initiation)

Smart watch

- Face-based auto-wake and auto-sleep
- Always-on gestures

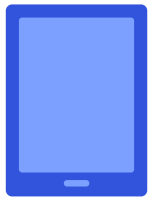

Tablets

- Simple gaze tracking for advertising attribution
- Improved landscape/portrait screen orientation

Virtual reality

- Low power gaze tracking (foveated rendering)
- Low power visual odometry for 6 DoF

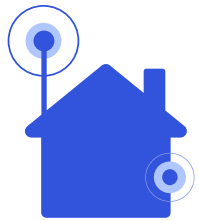

‘Intelligent’ occupancy trigger

- Distinguish humans from other objects
- Add data layer to trigger: How many? Where?
- Trigger on particular events or objects

‘Intelligent’ interactivity trigger

- Face detection as a trigger for interactivity
- Smart appliance can react when a user approaches to engage it

Standalone intelligent data sensor

- Heat maps of how a space is occupied
- Privacy advantages – data only, no images captured

Qualcomm Innovation

- Highly integrated & holistically optimized System
- Ultra-low power designs
- Always-On vision defined and targeted as < 1 mA, **active & end-to-end**

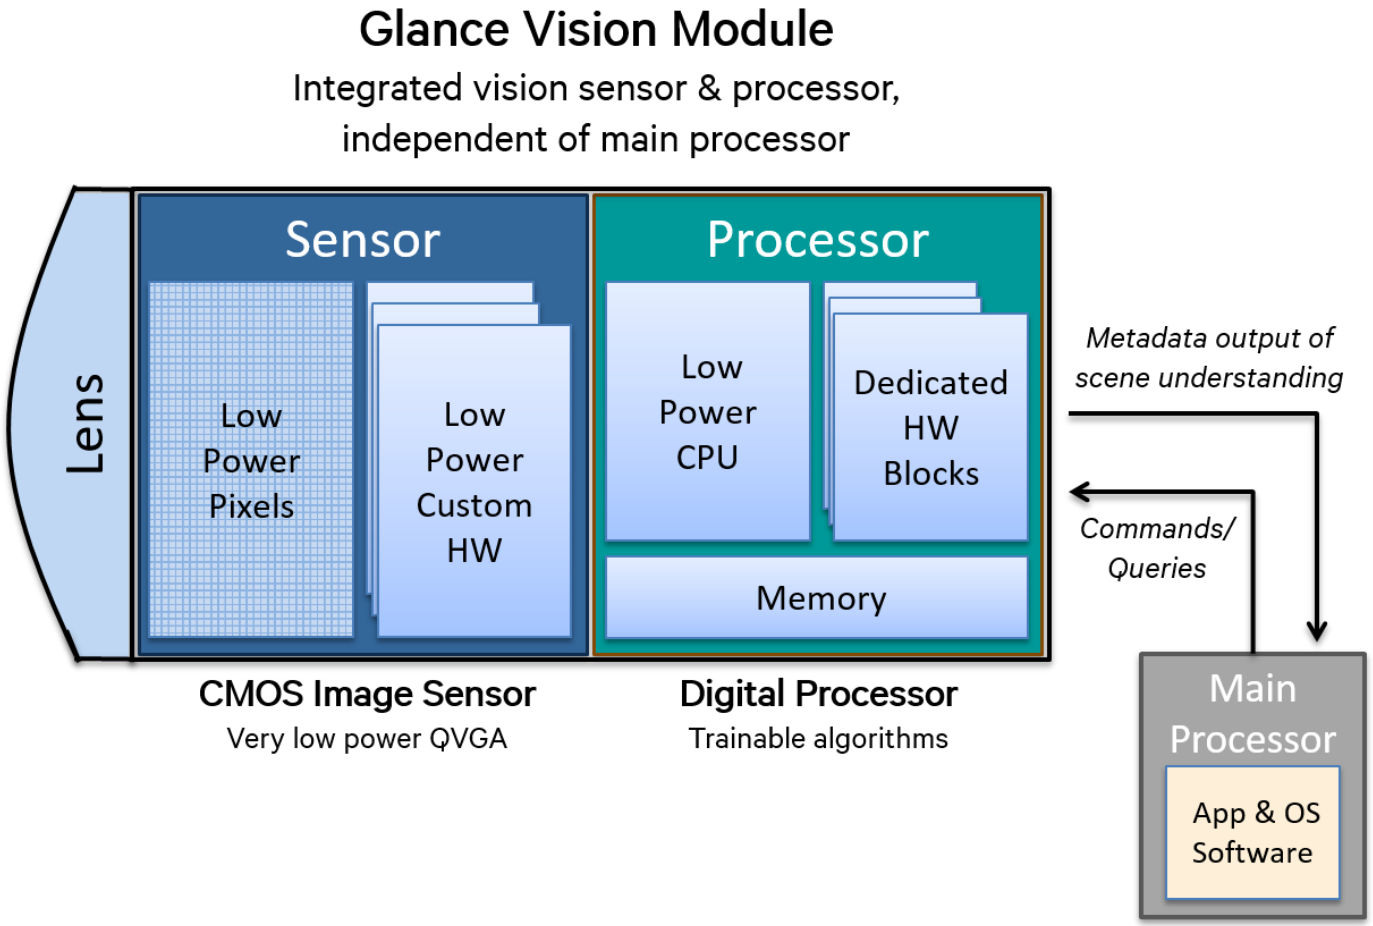

Key Features

- Ultra-low power, < 1 mA (end-to-end, w/ sensor included)
- Small size
- Low cost
- Privacy (output is metadata)
- Configurable for different use cases
- QVGA sensor
- N-IR compatible

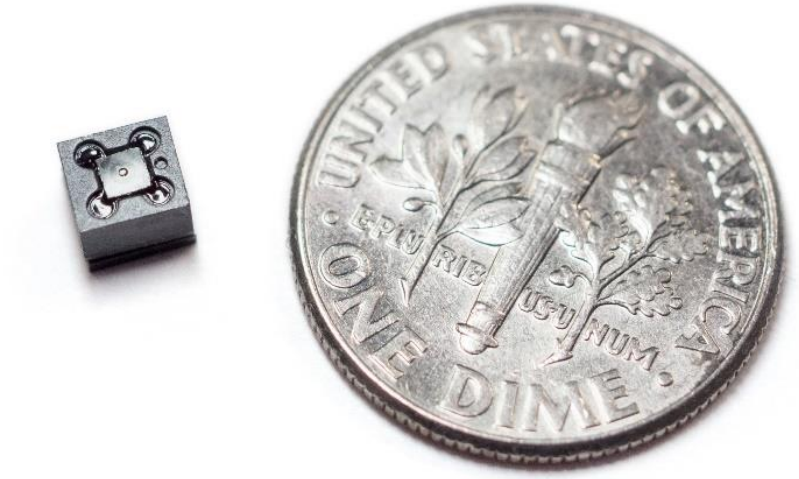

Qualcomm® QCC112

- Commercially available
- Supports many uses at ~1mW, including chip, sensor, and power management
- Features:
  - Ultra-low-power MCU
  - Streaming Array Processor (SuP)
    - Programmable
    - Can be power collapsed
    - Data bursted with DMA into TCM
  - Embedded PMU
  - Vision Accelerator
  - Custom memory
    - 2X lower dynamic power and 3X lower retention power vs. standard memory cell

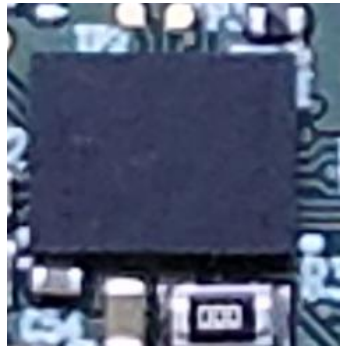

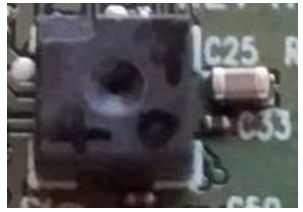

What needs to be done for Always-on Computer Vision

Our Approach	Traditional Approach
Image itself is secondary to information	Image quality paramount
Monochrome works in most cases. ≤ 8-bit sufficient	Color & wide bit-depth preferred
Focus can be good enough in most cases	Focus, autofocus, Bokeh important
Adequate pixel count for applicable distance	Higher pixel count
System power optimized including sensor	Sensor & algorithm/model often split
Images shot in challenging lighting	Camera & subject posed for best image
Inference is heavily weighted	Balance between training & inference time
Algorithms redesigned with memory & power in mind	Built upon available technologies
Metrics may be event-based	Typical metric is frame-based

Our System Approach for Always-on Computer Vision

- Favor algorithms with adaptive compute
  - Only perform computer vision when image/area has changed
  - Run light weight algorithms first
  - Favor algorithms/models with content adaptive capabilities
  - Stop when there is enough information:
    many applications only need to know the presence of ≥1 object vs. count all objects
- Simplify
  - Often easier to run models at different scales than resizing images
  - Optimize brightness to favor detection

Our System Approach for Always-on Computer Vision

- Optimize the entire system end-to-end
  - Use low power sensor
  - Optimize IO
  - Move algorithms to HW when possible
  - Keep memory close to compute engine

Various Use-Cases with Face and Human Detection

Half body

Full body

3/4 body

Change Detection

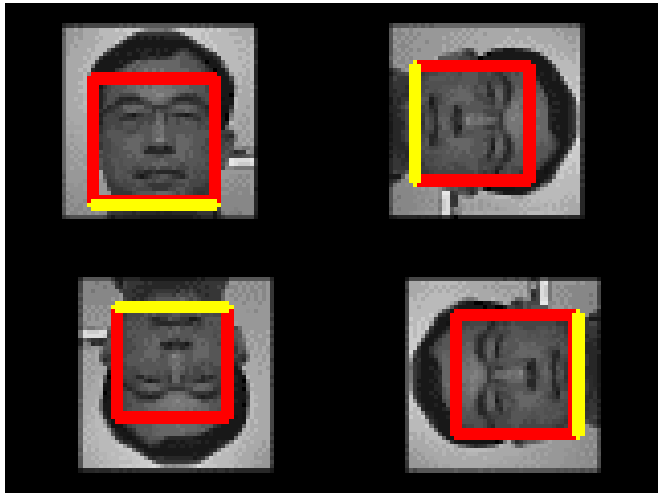

Multiple orientation

Video: Phone/Tablet Use Cases

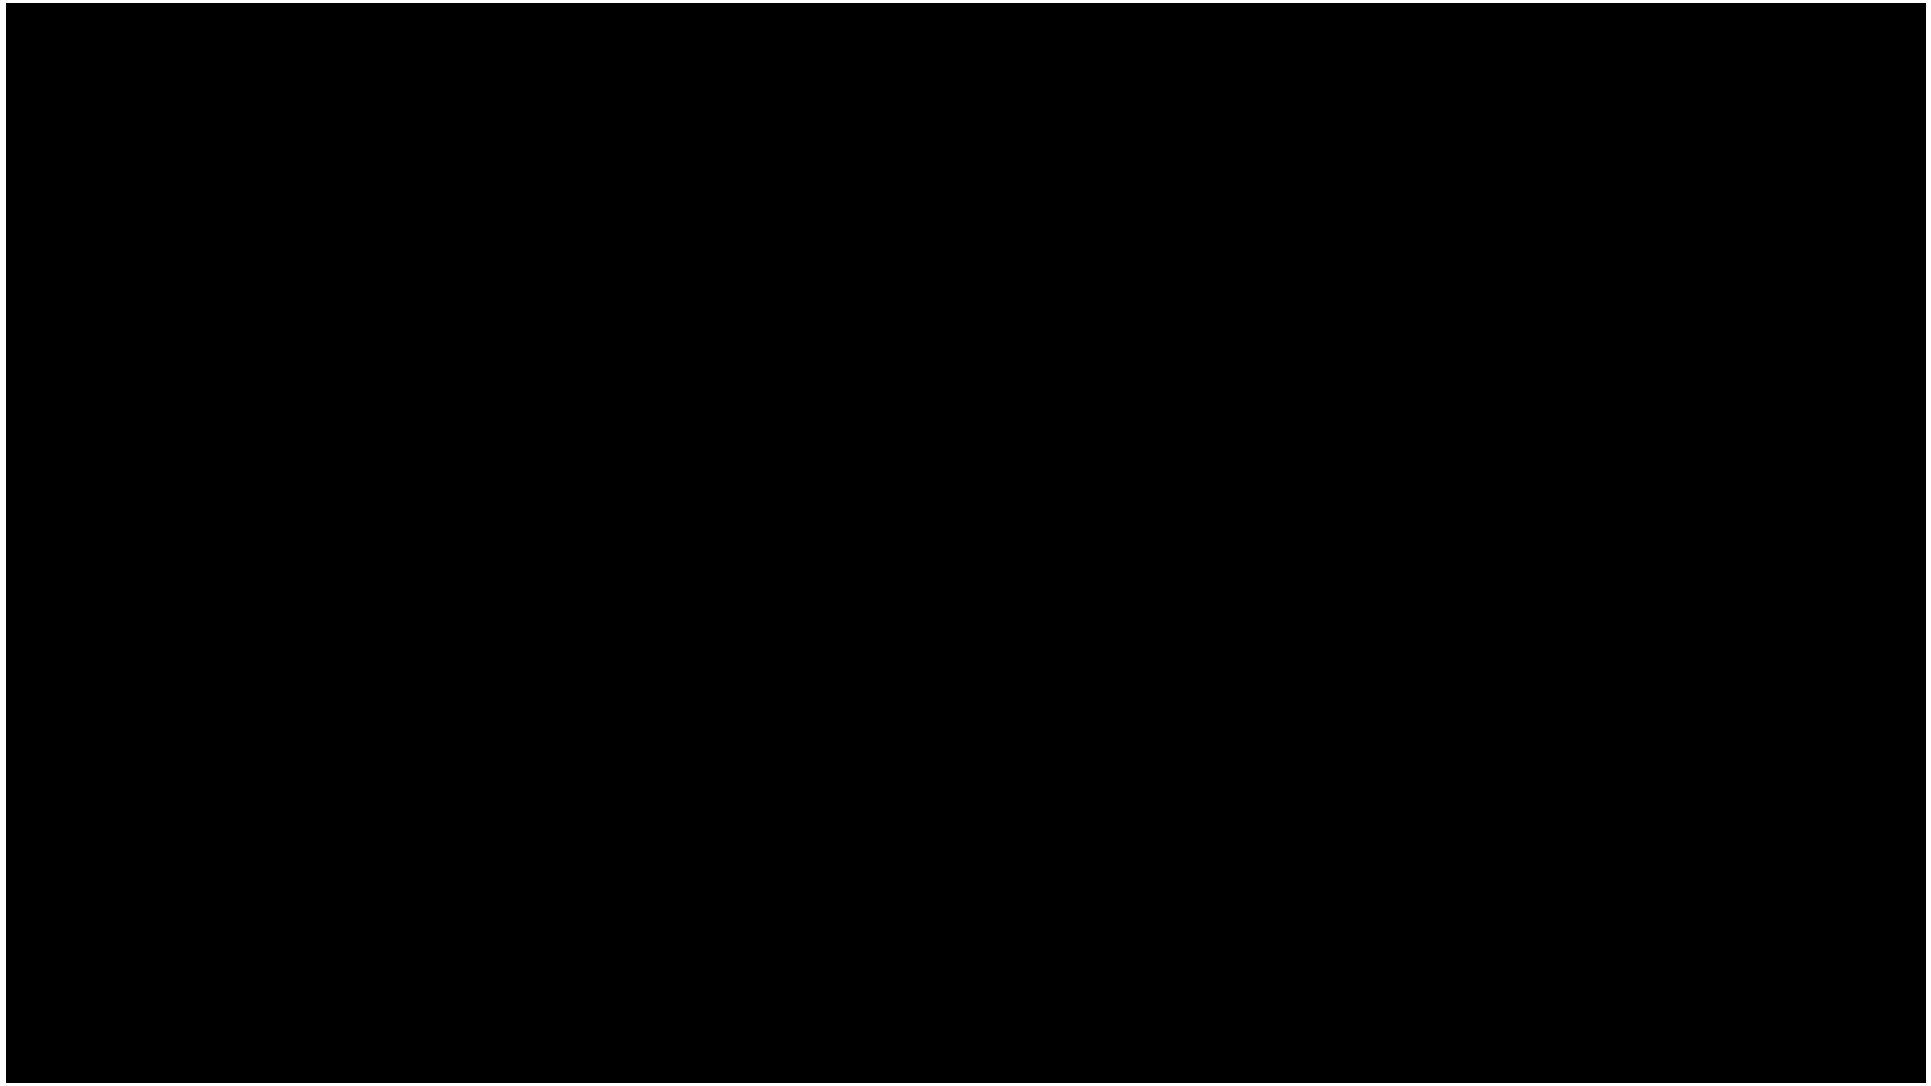

Video: Connected Camera Use Cases

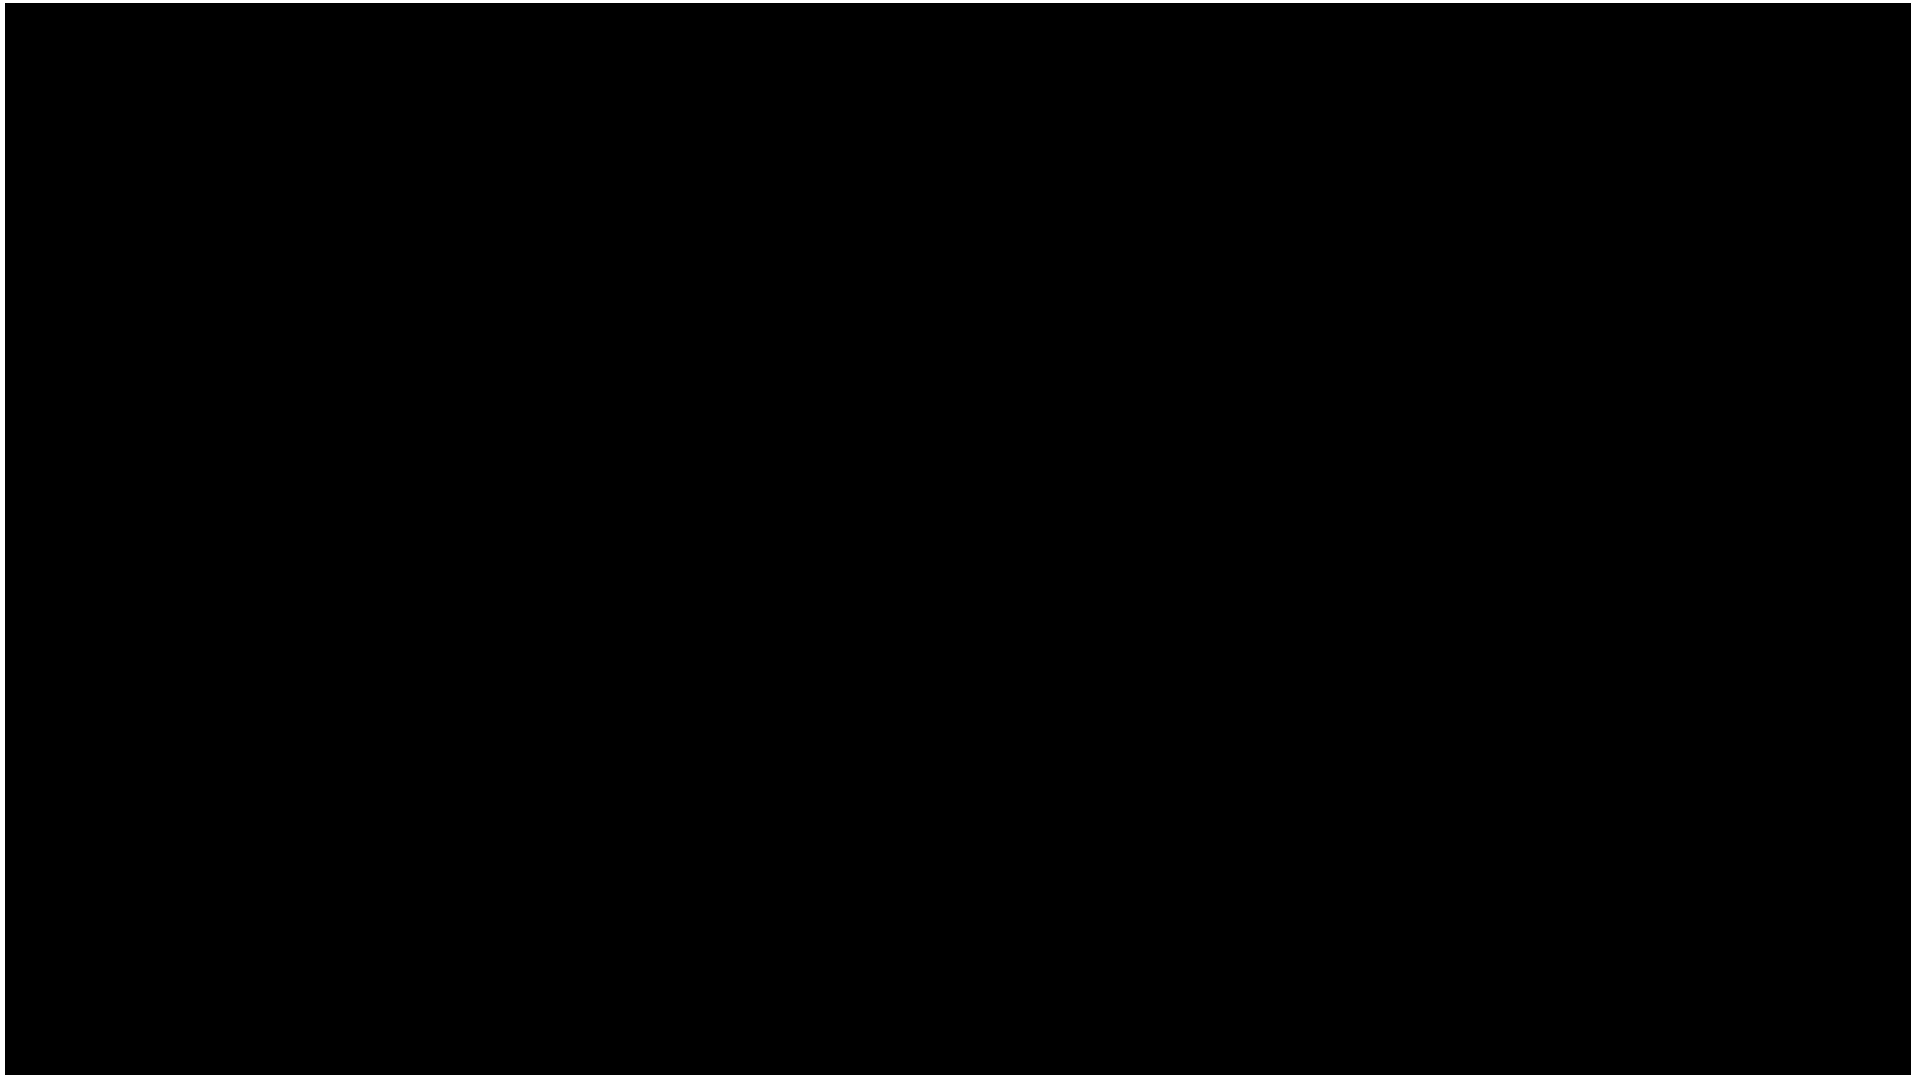

Detect Markers, Logos, Toys, and Many More…

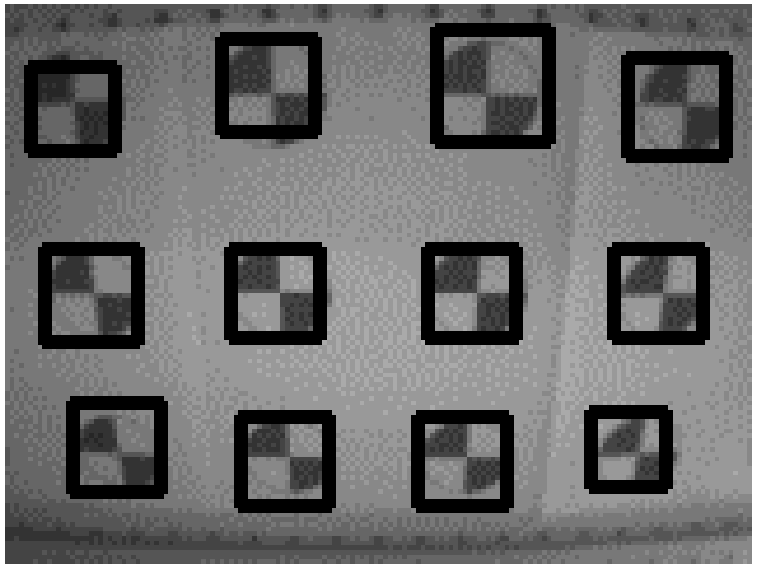

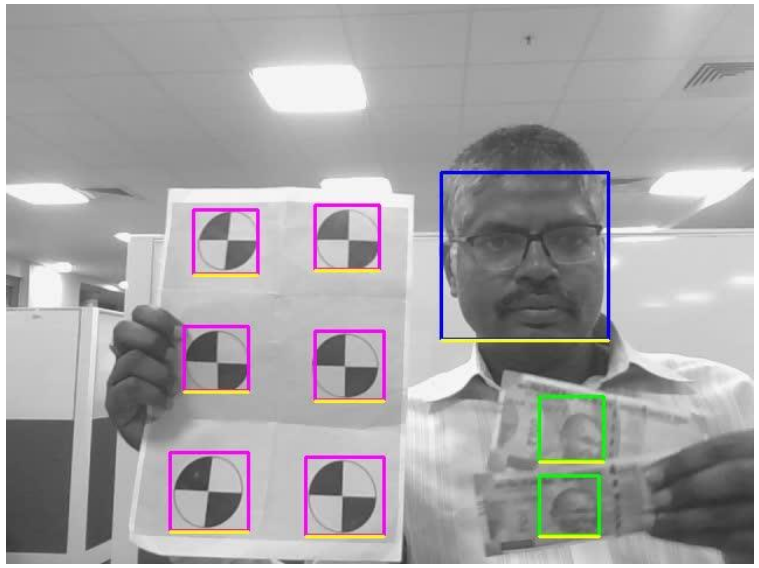

Retail Analytics
(1) Shelf Status
(2) Customer Engagement

Lightweight Gesture Detection

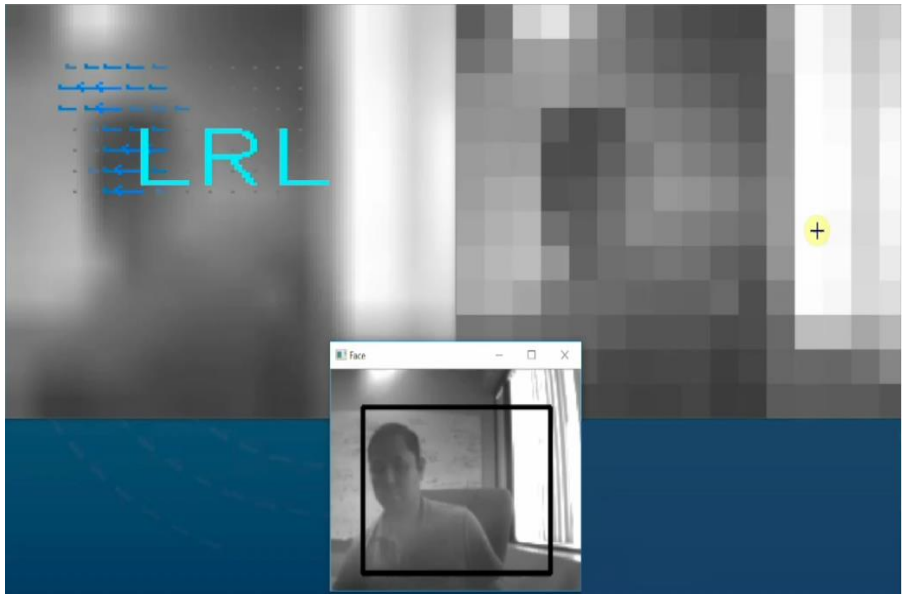

Dynamic Gesture Detection

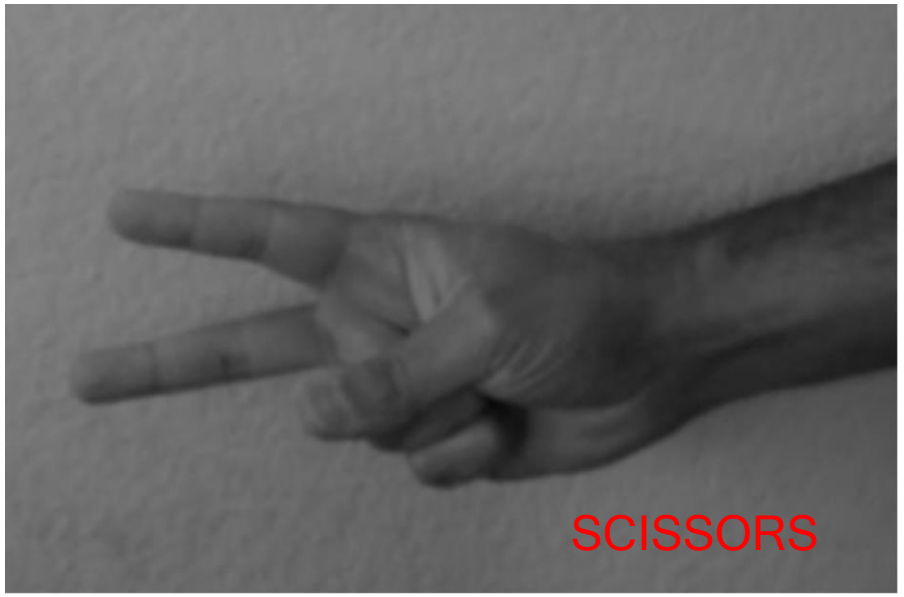

Static Gesture Detection

Gesture Recognition – Rock / Paper / Scissors

Successful detection in challenging light scenarios

- Detection scenario at distance and low light is challenging
- Model and algorithm must be resilient to these scenarios
- Sensor also sensitive to 850 nm IR

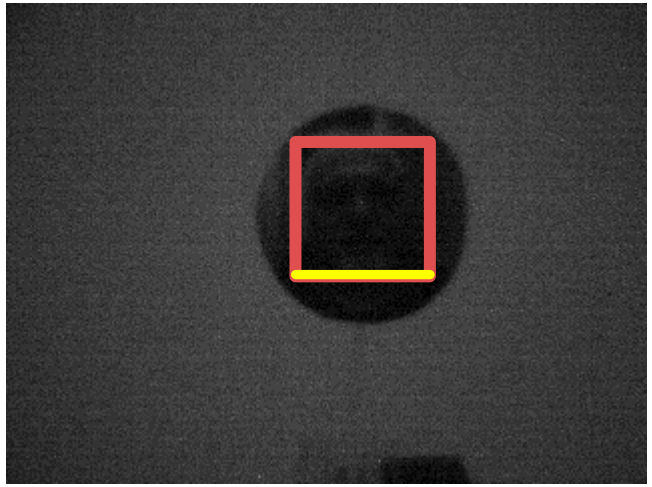

Face at low light (0.1 lux)

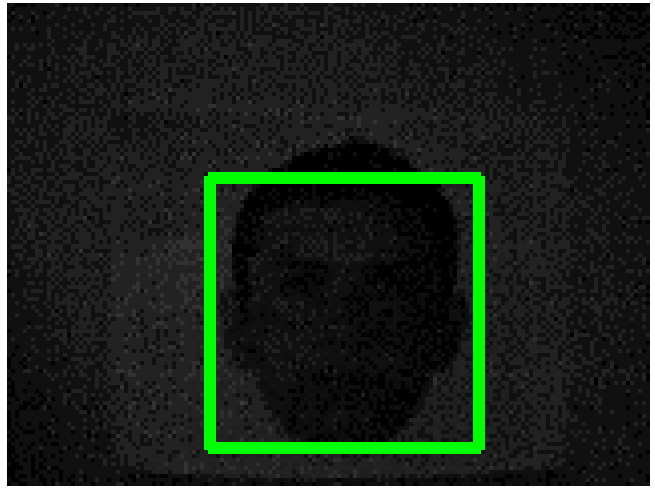

Face detection low light indoor at 768us

Full body at 60 ft

Full Body 20 ft 3 lux

Full Body in direct sunlight

Turbo Mode: 100 FPS Face Detection

End-to-end system optimization allows for enhanced performance at low power

- Tuned exposure control
- Tailored detection models for low light
- Boosted system clock
- Single-digit mW power (1-10 mW)

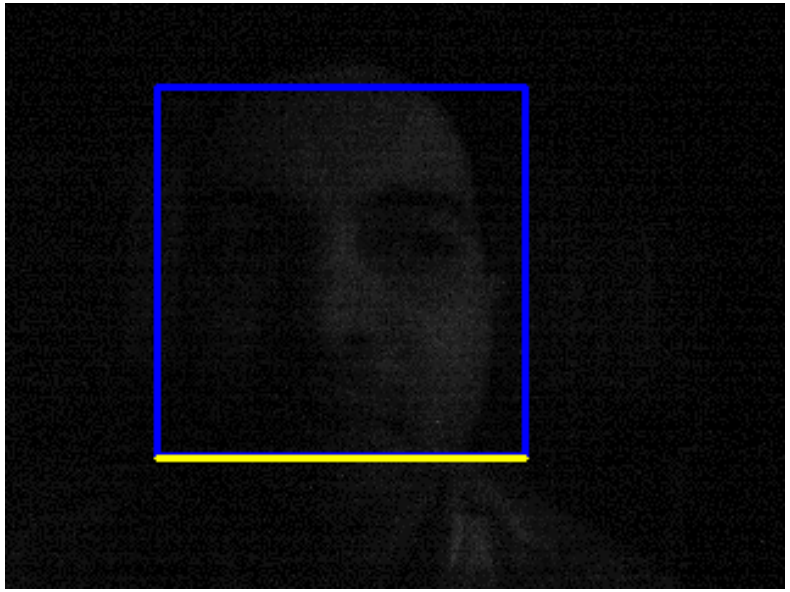

Face Detection @ 10 lux + 4 ms exposure

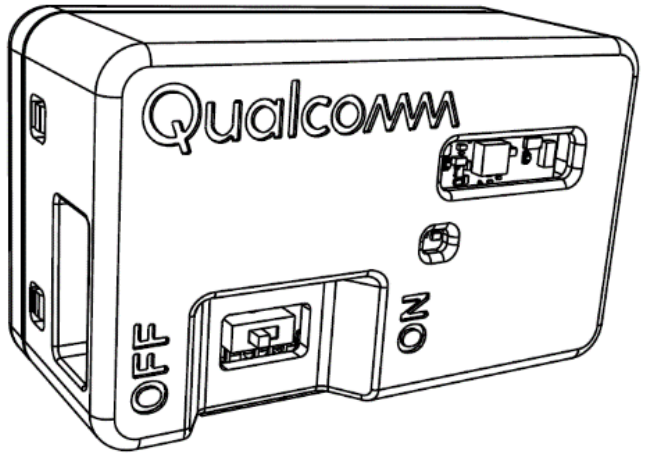

Battery-powered
AOCVS demo

tinyML for AR/VR Applications

Also watch: tinyML Talk by Hans Reyserhove (Facebook Reality Labs)

Embedded Computer Vision Hardware through the Eyes of AR/VR

https://www.youtube.com/watch?v=c4g2zwFR3ps

Building Robust Lightweight Pupil Detection Models

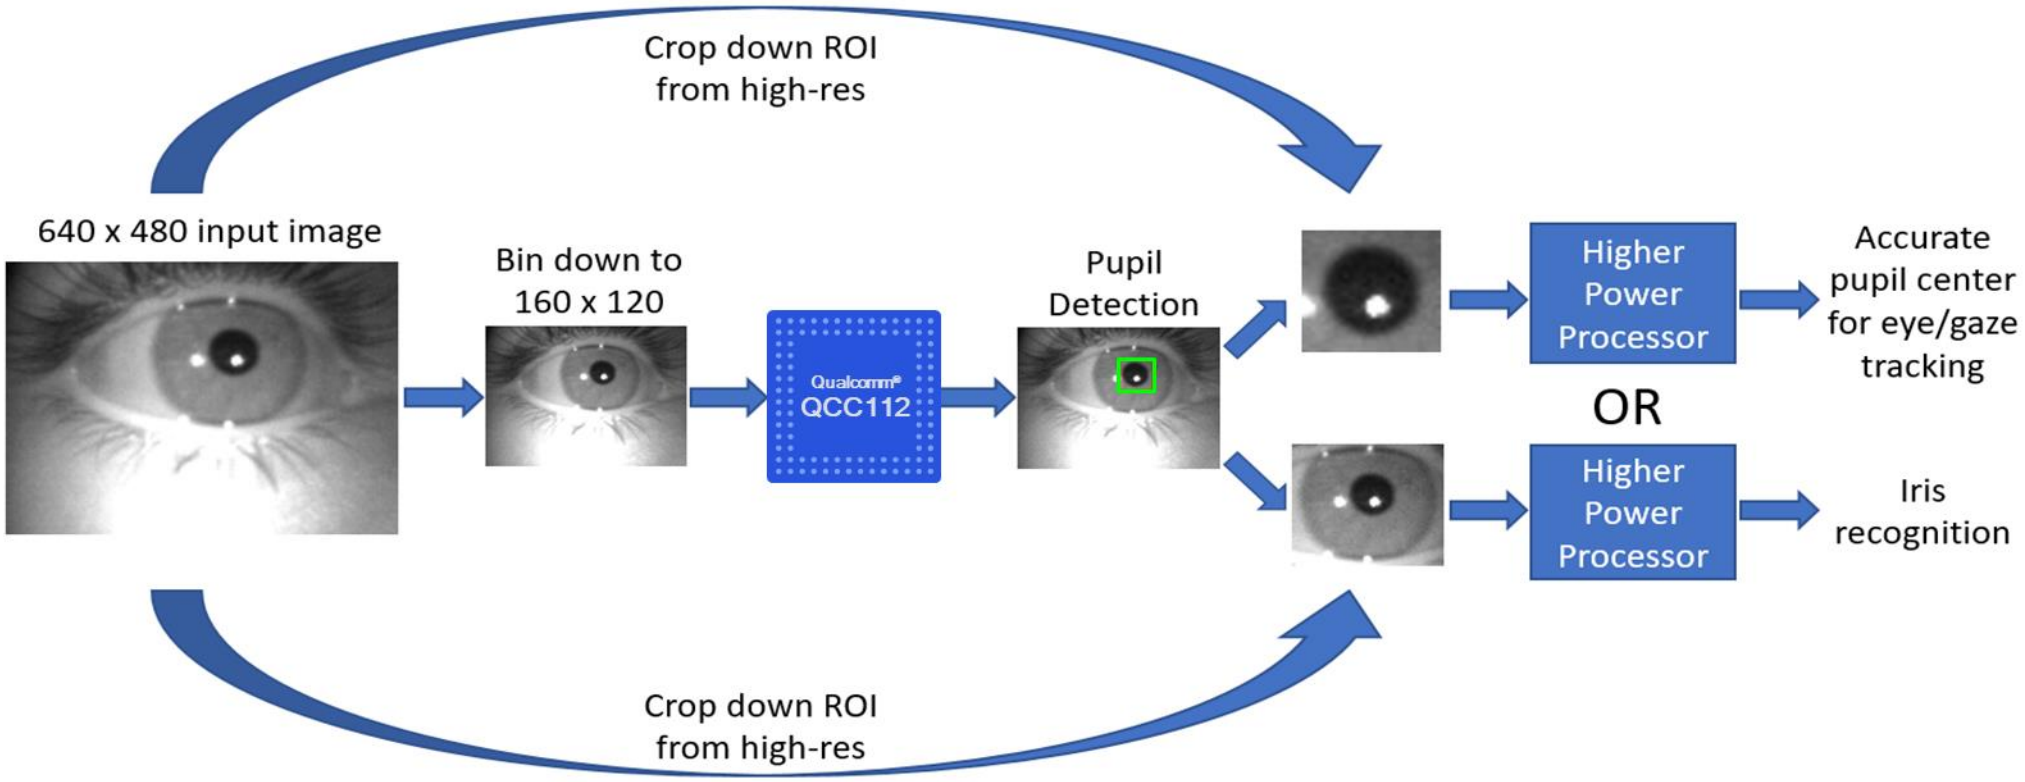

Accurate pupil detection enables downstream AR/VR applications

Detections in Visible/Low/IR Lighting

Visible Light

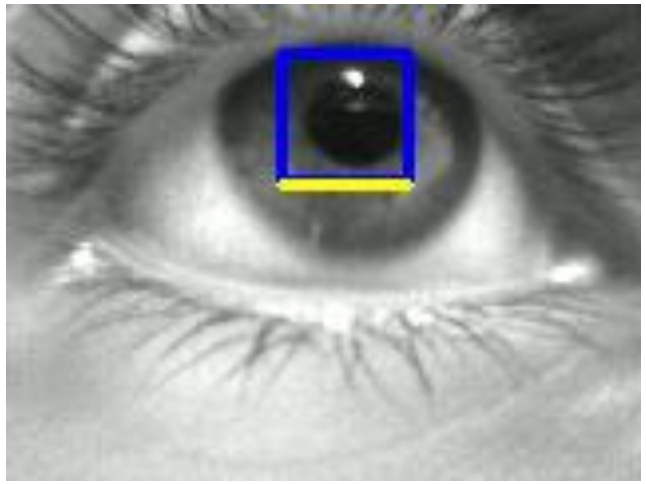

Ambient indoor lighting (~150lux)

Low Light

Raw image taken by sensor at 10 ms exposure time

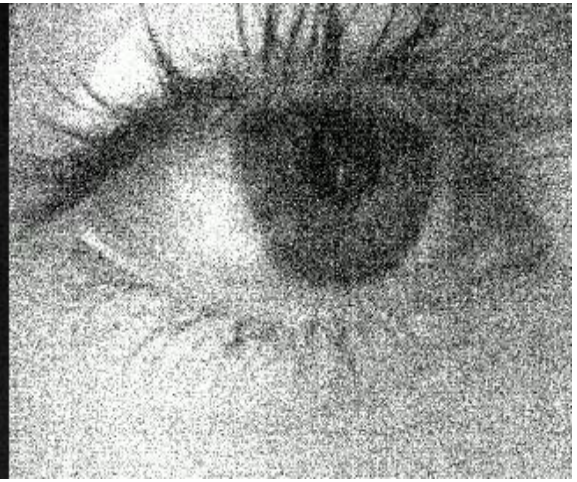

Pupil detections on raw images

Contrast-stretched (for visualization)

Infra Red Light

~1 ms exposure time

How fast can we go?

FPS	Exposure Time (ms)	Detection Time (ms)
60	< 1	6.5
87	< 1	2.7

- IR illumination allowing extremely small exposure time while maintaining image quality
- Image processed at qqVGA res (160x120 pixels)
- Promising results with frame rates hitting as high as 100 FPS!
- Extremely small model size @ 40 kB
- Results at 40MHz clock – can clock higher for higher FPS requirements
- Images not transmitted, only metadata – ROI can be transmitted for downstream processing

Eye tracking demo on computer screen

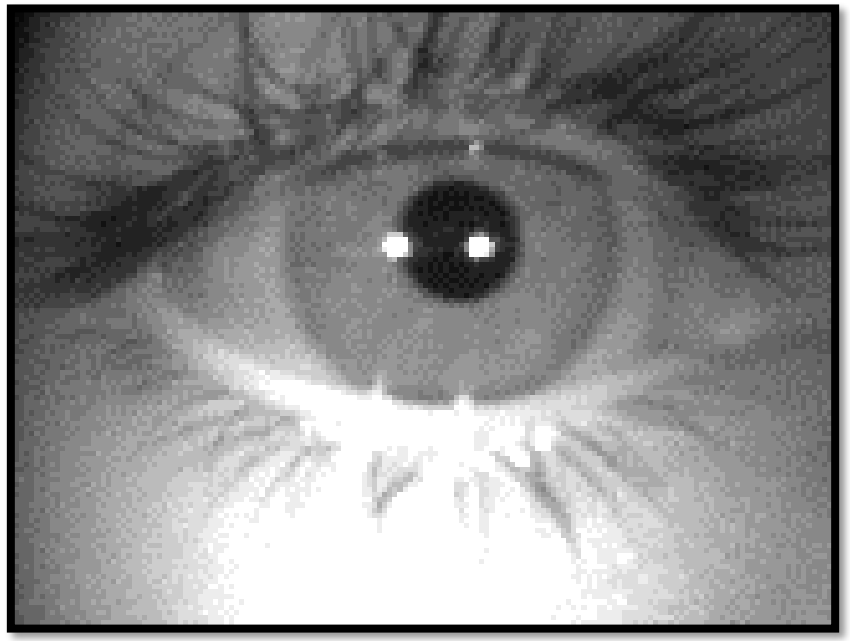

User looking at fixed points on screen

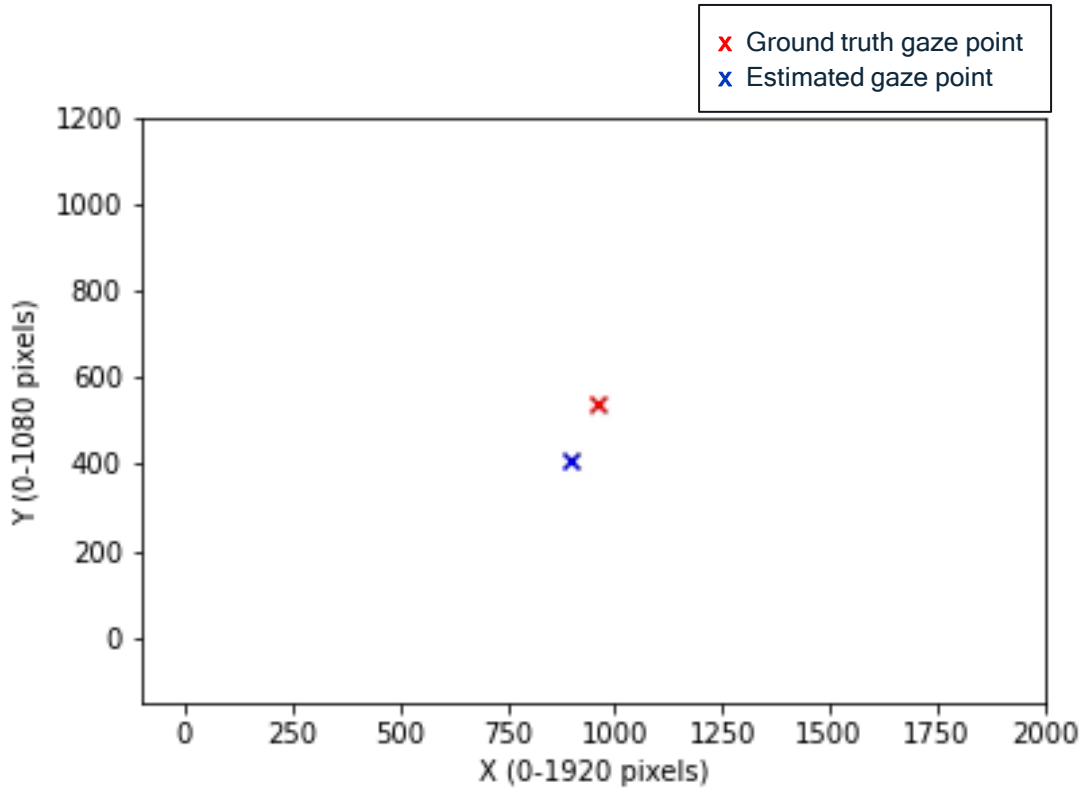

Actual vs estimated gaze points on screen

Screen gaze estimation from regression over pupil detection output

Running on commercial silicon Qualcomm® QCC112

Demonstrated detection at up to 60-100 fps on the hardware

Training Tool: AOCVS (Always-on Computer Vision System) Portal

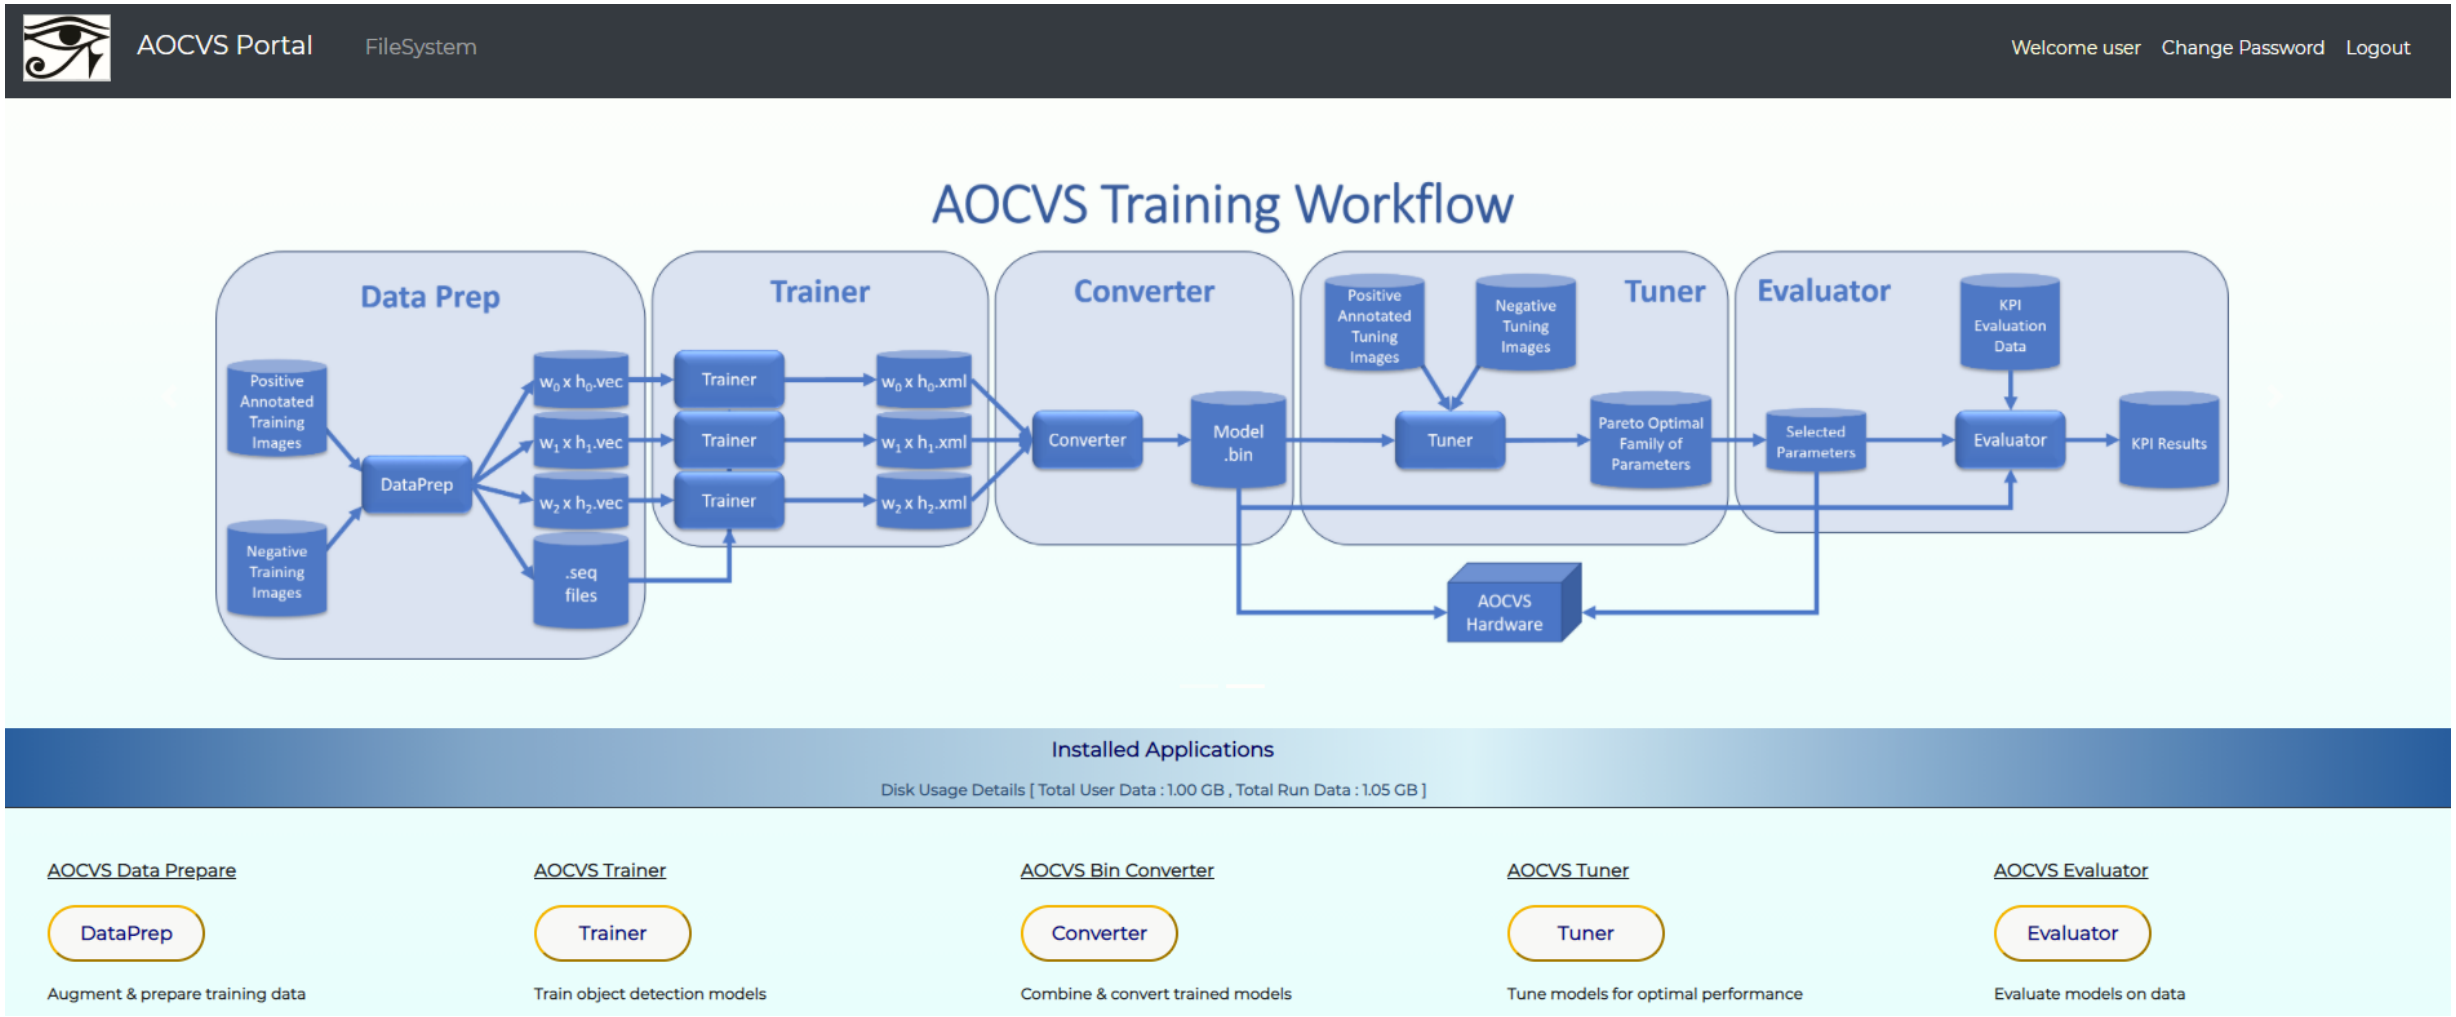

Training Tool: AOCVS (Always-on Computer Vision System) Portal

- Browser-based application front-end
- Backend accessible via command line
- Installation supported on Windows & Linux (Docker)
- Install on remote/cloud machine and launch/monitor jobs from any browser
- Compatible with most modern browsers
- Experiment tracking aids **reproducibility**

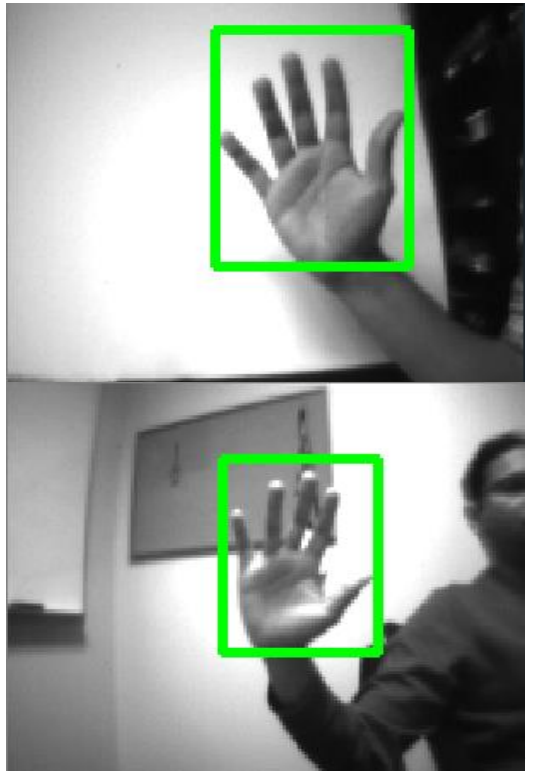

- Synthetic data augmentation enables training with only a few positive images
- Can train models with images from any source, data augmentation matches sensor image quality

Visualization of Results

Show Recent Jobs 25

Select all | Deselect all | Delete Jobs | Stop Jobs

Search:

RunID.	Date	Status	Algo	Model(s)	Input Type	Config	Results	Edit
14	2019-09-16 12:50:29.057	Completed	OD	circle_10_13_16_model.bin	File with list of images	Config	Results	Edit
13	2019-09-16 12:49:16.328	Completed	OD	circle_10_13_16_model.bin	File with list of images	Config	Results	Edit
12	2019-09-05 10:29:54.975	Completed	OD	fullbody_model.bin	List of images	Config	Results	Edit
11	2019-09-05 10:29:21.068	Completed	OD	fullbody_model.bin	List of images	Config	Results	Edit
10	2019-09-05 10:28:41.532	Completed	OD	face_model.bin	List of images	Config	Results	Edit
9	2019-09-05 10:27:50.489	Completed	OD	face_model.bin	List of images	Config	Results	Edit
8	2019-09-05 10:20:05.618	Completed	OD	face_model.bin	Video file	Config	Results	Edit
7	2019-09-05 08:55:10.470	Completed	OD	circle_pattern_10_13_16.bin	File with list of images	Config	Results	Edit
6	2019-09-05 08:54:35.959	Completed	OD	circle_pattern_10_13_16.bin	File with list of images	Config	Results	Edit
5	2019-09-05 08:53:05.583	Stopped	OD	circle_pattern_10_13_16.bin	File with list of images	Config	Results	Edit

Showing 1 to 10 of 14 entries 1 row selected

Previous 1 2 Next

Run ID :14 (Running) Evaluator Log Result CSV Output tarball

Show 10 entries Search:

No.	Processed Frame	Details
1	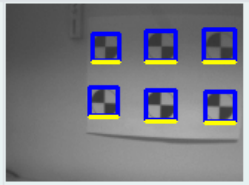	image105_out_14_0.bmp 6 detections LBP Count: 18089 Detect #1 X0 : 58 Y0 : 20 W0 : 19 H0 : 19 Orientation : 0 Score : 11 Detect #2 X1 : 94 Y1 : 18 W1 : 21 H1 : 21 Orientation : 0 Score : 6 Detect #3 X2 : 133 Y2 : 17 W2 : 22 H2 : 22 Orientation : 0 Score : 3 Detect #4 X3 : 56 Y3 : 56 W3 : 20 H3 : 20 Orientation : 0 Score : 9 Detect #5 X4 : 94 Y4 : 58 W4 : 21 H4 : 21 Orientation : 0 Score : 10 Detect #6 X5 : 134 Y5 : 59 W5 : 21 H5 : 21 Orientation : 0 Score : 8
2		image11_out_14_1.bmp

Showing 1 to 10 of 51 entries

Previous 1 2 3 4 5 6 Next

Visualization of Key Performance Metric and Compute Tradeoffs

Show Recent Jobs 25

Select all | Deselect all | Delete Records | Stop Jobs

Search:

RunID.	Date	Status	Model	GT File	MaxFPPI	Config		
6	2019-09-16 12:28:12.552	Completed	circle_10_13_16_model.bin	tuning_data.txt	0.05	Config	Results	Edit
5	2019-09-05 10:23:51.799	Completed	circle_pattern_10_13_16.bin	tuning_data.txt	0.05	Config	Results	Edit
4	2019-09-05 06:50:12.091	Completed	circle_pattern_10_13_16.bin	tuning_data.txt	0.05	Config	Results	Edit
3	2019-09-05 06:49:21.500	Completed	circle_pattern_10_13_16.bin	tuning_data.txt	0.05	Config	Results	Edit
2	2019-09-05 06:46:30.806	Stopped	circle_pattern_10_13_16.bin	tuning_data.txt	0.01	Config	Results	Edit
1	2019-09-05 06:33:44.624	Completed	circle_16_20_25.bin	tuning_data.txt	0.01	Config	Results	Edit

Showing 1 to 6 of 6 entries 1 row selected

Previous 1 Next

Run ID :5 (Completed) Tuner Log GT File Output Pareto CSV Output tarball

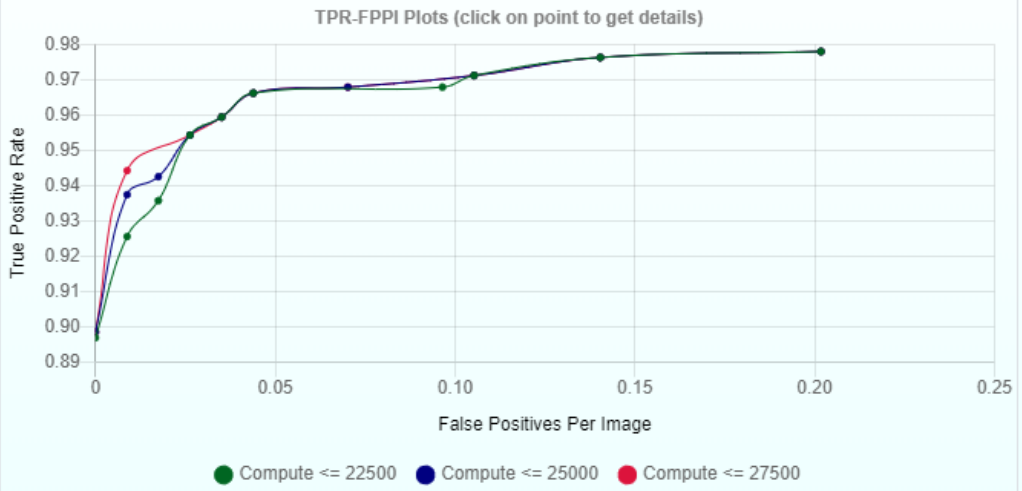

Tuned Parameter	Optimized Value
Row Strides	3,3,1,1
Column Strides	4,3,2,2
Merge Rect Threshold	3
Dist Rect Threshold	5
Min Object Size	1
Max Object Size	160
Num Stages	8

Tuner Result	Computed Values
True Positive Rate	0.9662
False Positive Per Image	0.0439
Precision	0.9913
F1 Score	0.9786
Compute Complexity	22044
Tuner Run Time	62 seconds.

Evaluate Model Object Detect

Next Generation Ultra-Low Power AI Engine

- Highly efficient neural network inference targeted at vision use cases
- Mobile-friendly backbone models such as MobileNets
- Eye gaze tracking with ~ 70 KB CNN model
- Competitive accuracy: < 5° gaze estimation error
- For power/performance reference: face recognition model (0.91 MB) runs at 25 fps at 5 mW power (250 MHz, 7nm)

Qualcomm

Thank you

Follow us on:

For more information, visit us at:

www.qualcomm.com & www.qualcomm.com/blog

Copyright Notice

www.tinyml.org

Copyright Notice

www.tinyML.org