ETH zürich

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

edgeGuard: Multimodal On-Device Perception for First Responders in Emergency Scenarios

Integrated Systems Laboratory (ETH Zürich)

Alice Pagnoux[1] apagnoux@iis.ee.ethz.ch
Marco Giordano[1] marco.giordano@ethz.ch
Prof. Michele Magno[1] michele.magno@pbl.ee.ethz.ch
Dr. Victor Kartsch[1] victor.kartsch@iis.ee.ethz.ch
Prof. Luca Benini[1,2] lbenini@iis.ee.ethz.ch

[1]*Integrated Systems Laboratory (IIS), ETH Zürich, Switzerland*
[2]*Department of Electrical, Electronic and Information Engineering (DEI), University of Bologna, Italy*

PULP Platform
Open Source Hardware, the way it should be!

@pulp_platform
pulp-platform.org
youtube.com/pulp_platform

The problem

[1]Firefigher's Raw POV - https://www.youtube.com/watch?v=mphHFk5IXsQ

Indoor mission threats

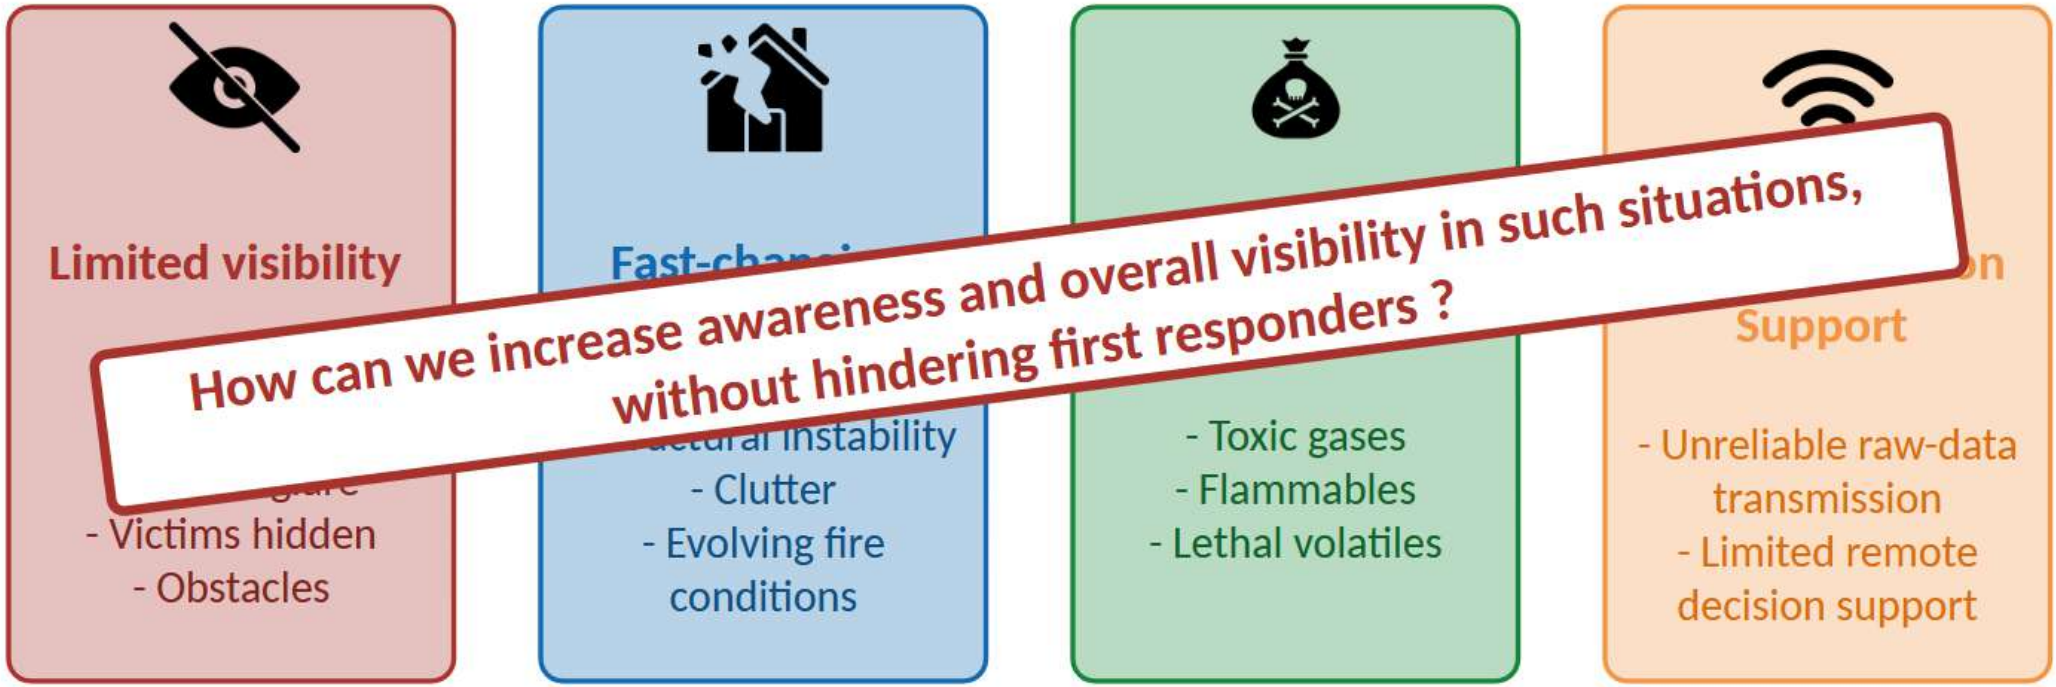

What exists today — and what is missing

Commercial wearables

+ Re… navigation
- No onboard multimodal scene understanding, limited compute

Academic edge vision

…on algorithms demonstrated
- Multi-watts platforms, hard to helmet-mount, heavy

Electronic noses

…lassifiers

+ Used mainly for data loggers
- No tailored for field

Gap: A compact, mission-ready wearable that onboards multimodal intelligence for the extreme fire conditions

[1] Axon Body 4 - https://www.axon.com/products/axon-body-4
[2] Motorola VB400 - https://www.motorolasolutions.com/en_us/video-security-access-control/body-worn-cameras/vb400.html
[3] Qwake C-THRU - https://www.qwake.tech/
[4] Raspberry Pi 5 - https://www.pi-shop.ch/raspberry-pi-5-starter-kit-pi-5-4gb
[5] Orange Pi - http://www.orangepi.org/
[6] Jetson Orin - https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/

edgeGuard - Hardware Architecture

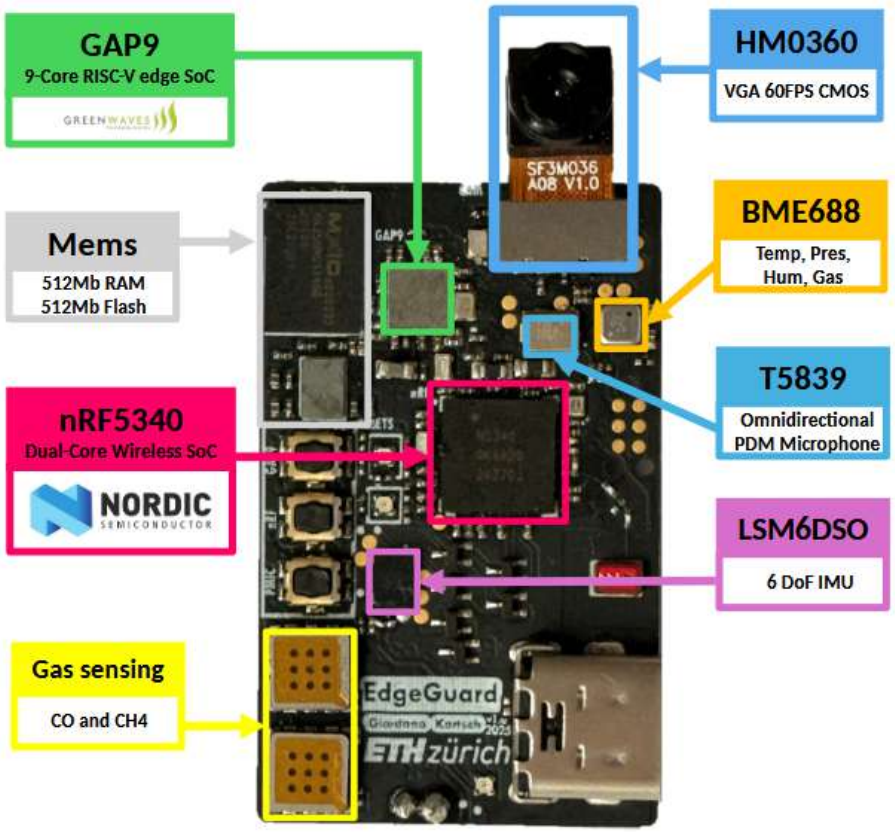

Top View

GAP9, multicore AI-MCU

9-core RISC-V cluster @ 370MHz, **NE16 NN accelerator**, 1.5MB L2 + Octal-SPI Flash & PSRAM

nRF5340, dual-core MCU

Dual Cortex-M33, BLE 5.4; samples sensors, forwards fused packets

Multimodal sensing array

HM0360 VGA camera (MIPI CSI-2), PDM mic, 6-DoF IMU, BME688 + dedicated CO / CH_4 channels

nPM1300 PMIC

Battery charging + system rails from one 600 mAh Li-ion cell

edgeGuard at a glance

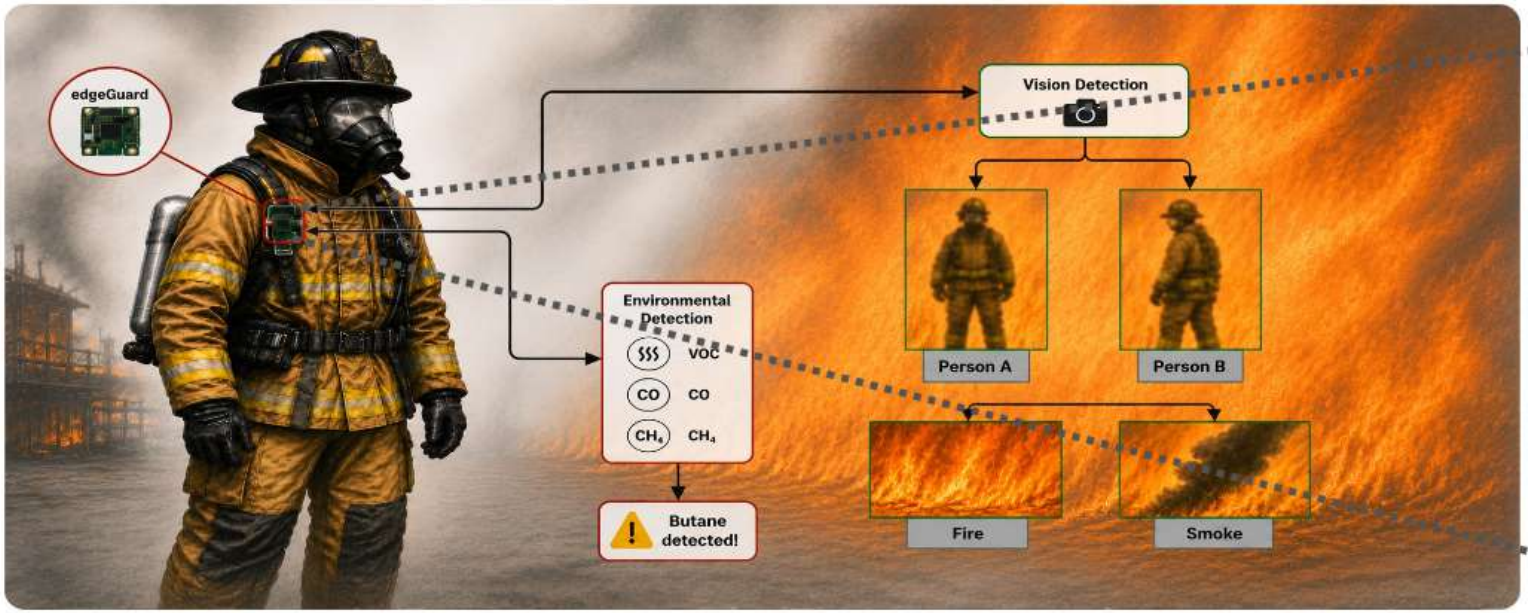

Two perception branches, both fully on-device:

Vision

Lightweight **YOLO** for fire/smoke/person detection on **GAP9 microcontroller**

Environmental

SVM hazard classification of gases & flammables on **nRF5340**

Vision branch — data curation process

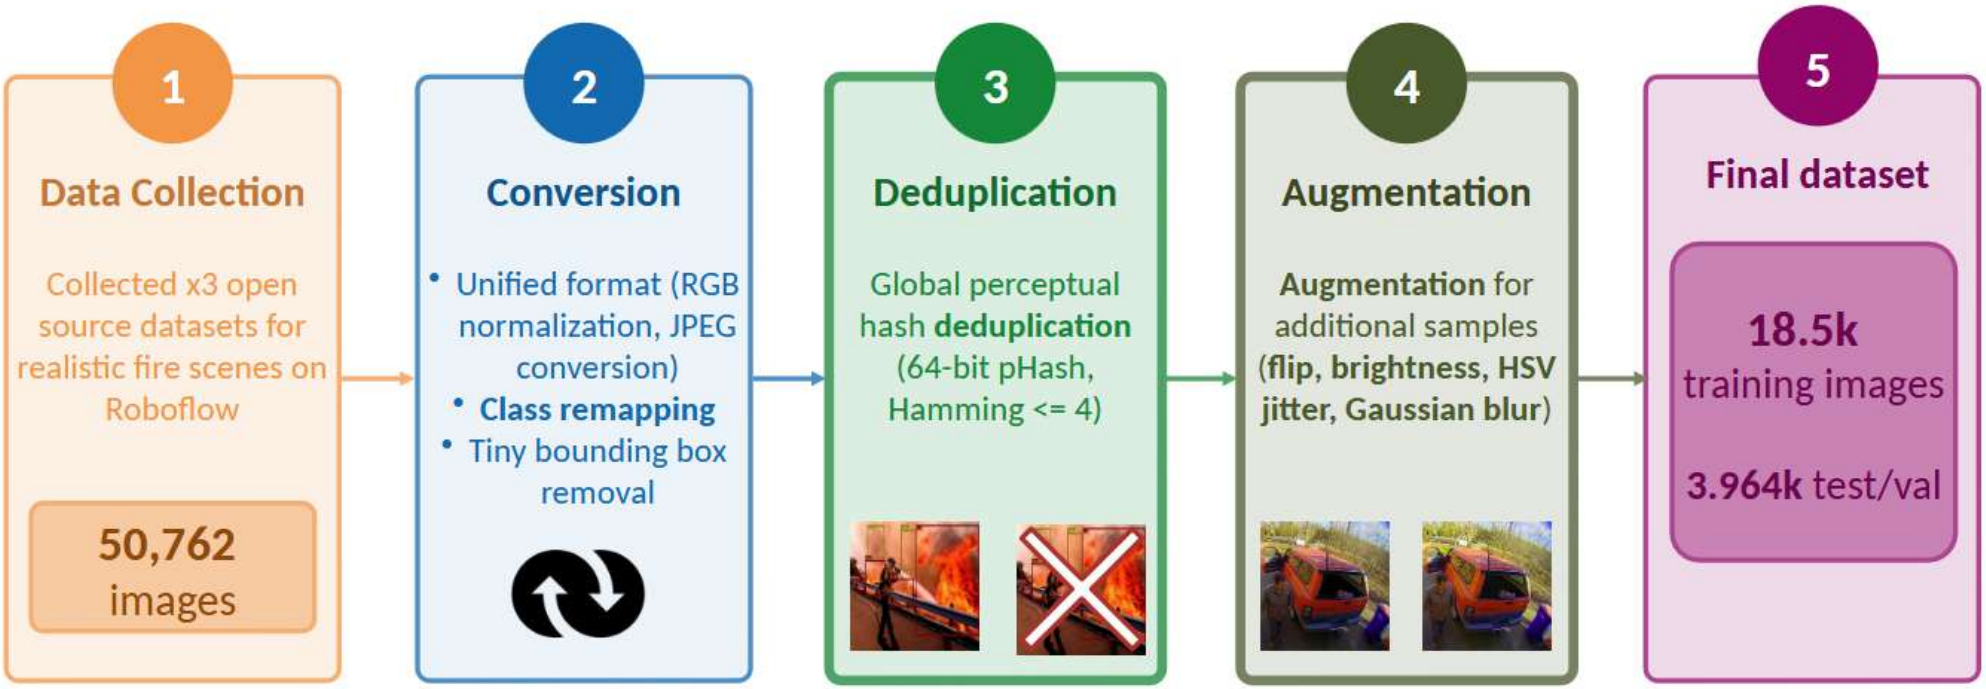

TinyissimoYOLO architecture trade-offs

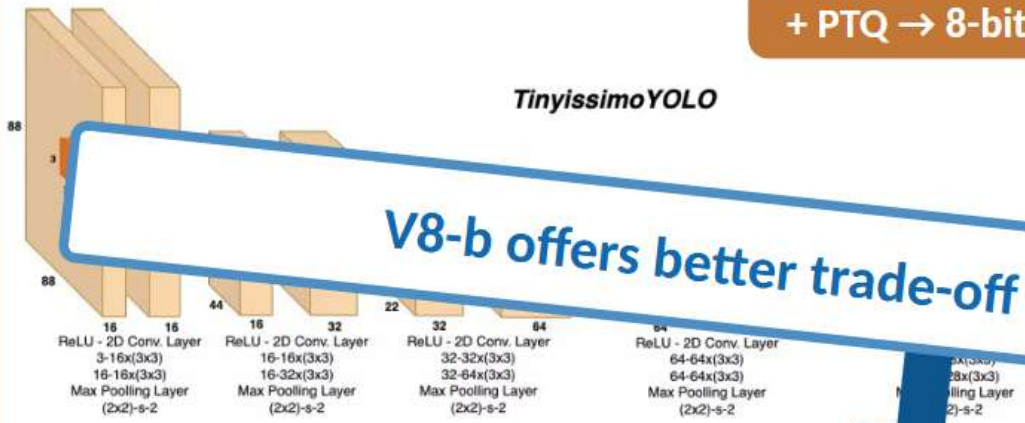

+ PTQ → 8-bit ready model

Designed for **embedded real-time detection**

Ultra-compact: only a **few hundred kB**

MCU

V8-b offers better trade-off for embedd deployment

v8-b Compact option

mAP50	precision
0.787	**0.828**
parameters	FPS
0.824M	**29.96**

v8-n Accuracy option

mAP50	precision
0.808	**0.858**
parameters	FPS
1.441M	**25.56**

[1]Moosmann, J., Giordano, M., Vogt, C., & Magno, M. (2023). Tinyissimoyolo: A quantized, low-memory footprint, tinyml object detection network for low power microcontrollers. *arXiv preprint arXiv:2306.00001*.

Per class detection accuracy

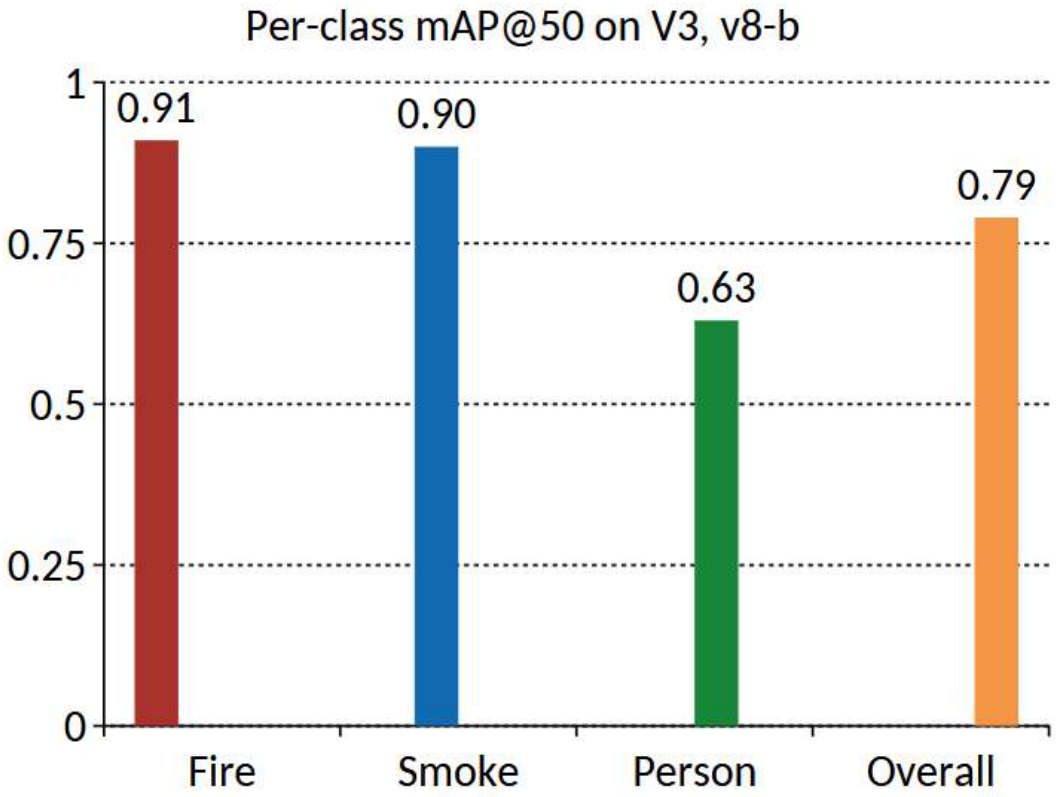

Fire and smoke are the strongest classes

Person detection is a bottleneck

Overall detection performance of 0.79 across all classes

Vision branch - Deployment on GAP9

Offline deploym

FP32-traine

pytorch + A
ONNX

AP9

C / binaries

On-device runtir

Input
HM03

raw Bayer f

rocess
MS
BBoxes
luster)

Deployment Results

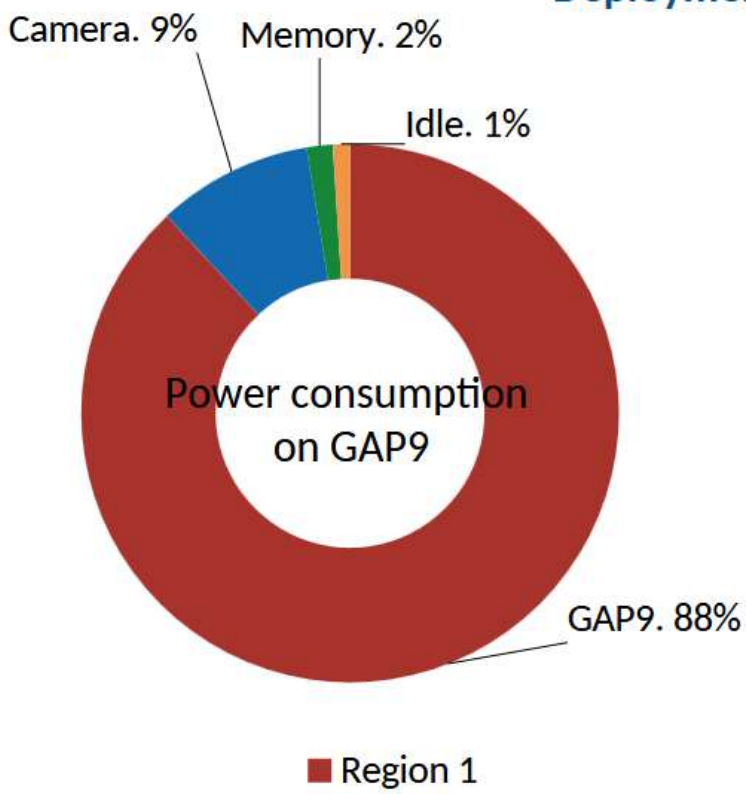

29.96 FPS
inference time

50.39 mW
vision branch

→ **88.1%** of vision branch power is GAP9 compute

Environmental branch - methodology and classification

Sensors: CO | CH^4 | Humidity | Gas resistance | VOC

Classes: Alcohol | Penetrating oil | Cooking oil | Butane | Air

Key idea: classify hazards from patterns across multiple sensor channels

1 — Baseline

Mean over last 120s before exposure

2 — Window

30s stable windows maintained over time

3 — 15-D features

Per channel:
- max deviation
- min deviation
- end-to-start change

4 — Linear SVM

StandardScaler
soft-margin linear SVM
5-fold stratified CV

Environmental results - hazard classification

91.6%

accuracy across five balanced classes

Takeways:

- Perfect separation for Alcohol and Butane
- Penetrating Oil and Cooking Oil can be further improved
- Class are separable, PoC validated

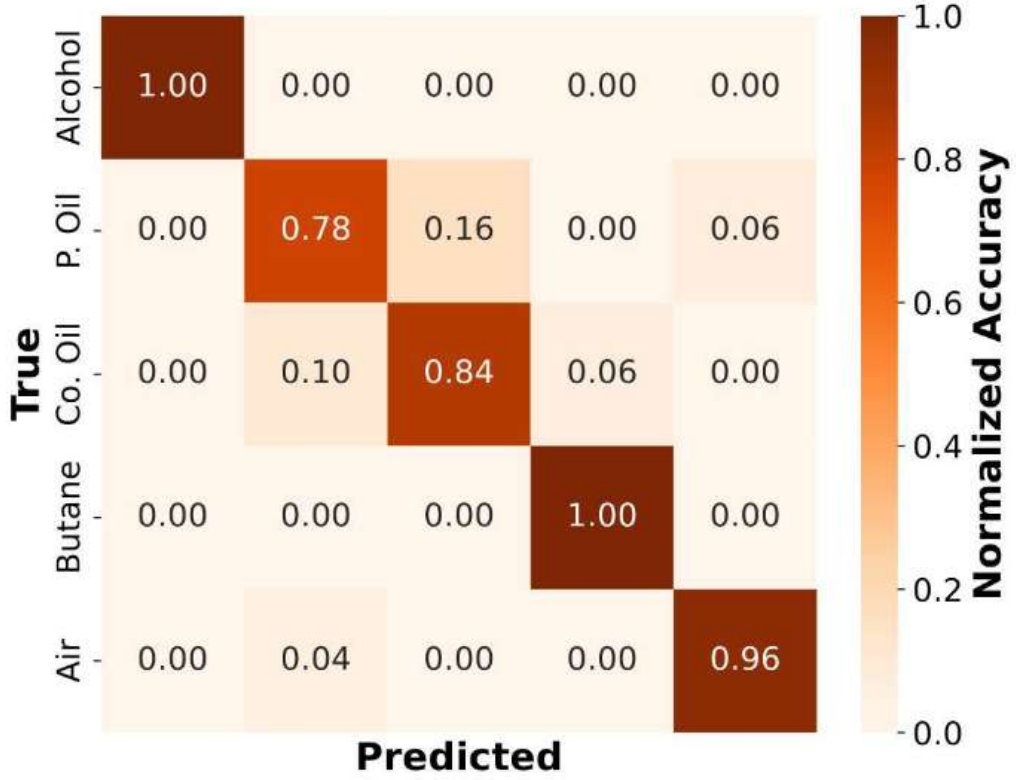

SVM confusion matrix

Environmental branch - Deployment on nRF5340

Host-side classi

Trained Li

standardised

ware

e -> table of

On-device runti

Sensi
5-chan

CO CH4, gas, VC

ification
SVM
-> hazard class

Deployment Results

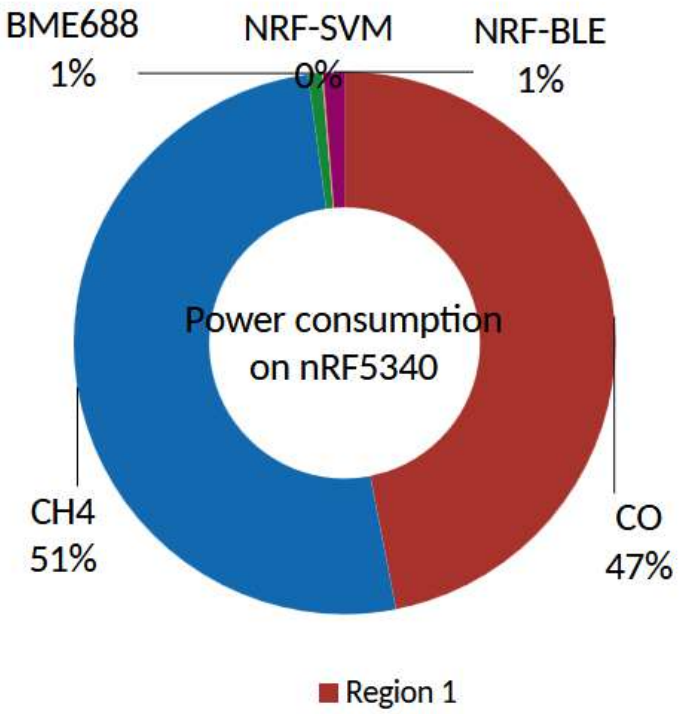

0.69
ms

164.76
mW env branch

→ **97.8%** of env branch power is CO + CH^4 sensing

Real-time, end-to-end operation

Output

Env. class + GAP9 class BBox and accu. + status → **nRF5340 packet** → **BLE / USB**

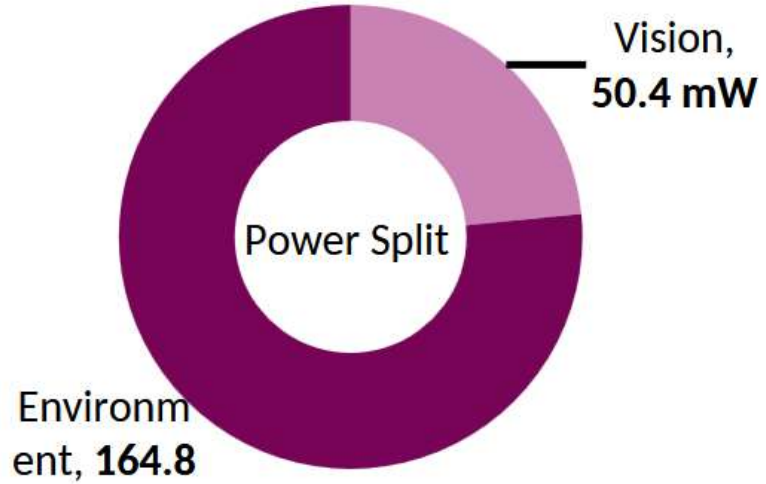

→ **The largest draw comes from chemiresistive gas sensors**

10.3 h
estimated lifetime from 600 mAh battery

215 mW
total power

2 Hz
environmental sampling rate

10 Hz
visual pipeline operating rate

Takeaways & Next steps

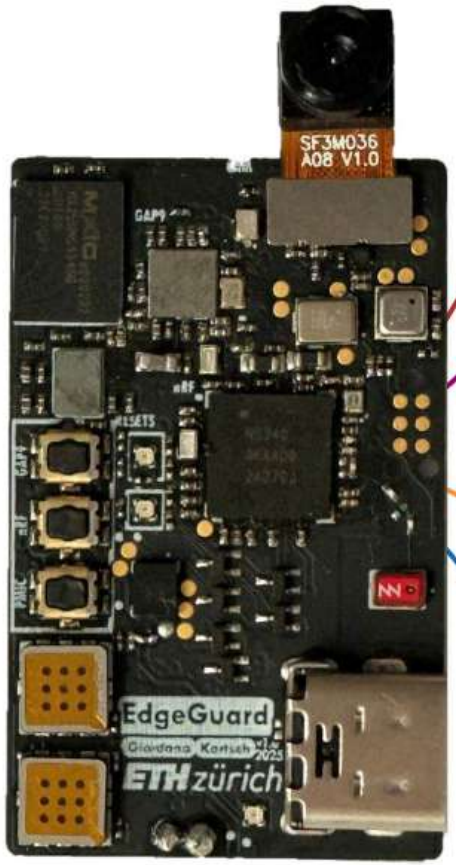

4 x 2.5cm
mountable PCB

Dual-sensing
Vision branch - **29.96 FPS, 0.787 mAP50**
Environment branch - **0.69 ms, 91.6% accuracy**

215.15mW
total power

>10 h
battery lifetime

1

Field validation

Move to **realistic live-fire validation** on target hardware

2

Class Balance

Improve **dataset balance for person detection,** especially in fire scenes conditions

3

Process more modalities

Extract vitals using current sensors (mic) and add **more vision modalities** (depth, thermal, etc)

4

Enable cross-modal fusion

Fuse both vision and environmental modalities for smarter alerts

Marco Giordano
PhD Student at ETH Zurich

Dr. Victor Kartsch Morinigo
Post Doctoral Researcher at ETH Zurich

Prof. Michele Magno
Professor at the ETHZ

Prof. Luca Benini
Professor at the ETHZ and University of Bologna

ETH zürich

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Thank you for your attention and stay tuned for updates!

github.com/pulp-bio/EdgeGuard