tinyML® Talks

Enabling Ultra-low Power Machine Learning at the Edge

"Datasheets for Machine Learning Sensors"

Matthew Stewart – Harvard University

July 11, 2023

TINY ML

www.tinyML.org

Thank you, **tinyML Strategic Partners**, for committing to take tinyML to the next Level, together

Executive Strategic Partners

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

Accelerate Your Edge Compute

SYNTIANT

Making Edge AI A Reality

www.syntiant.com

Platinum Strategic Partners

Renesas is enabling the next generation of AI-powered solutions that will revolutionize every industry sector.

renesas.com

DEPLOY VISION AI AT THE EDGE AT SCALE

Gold Strategic Partners

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Where what if becomes what is.

Witness potential made possible at analog.com.

ARDUINO PRO

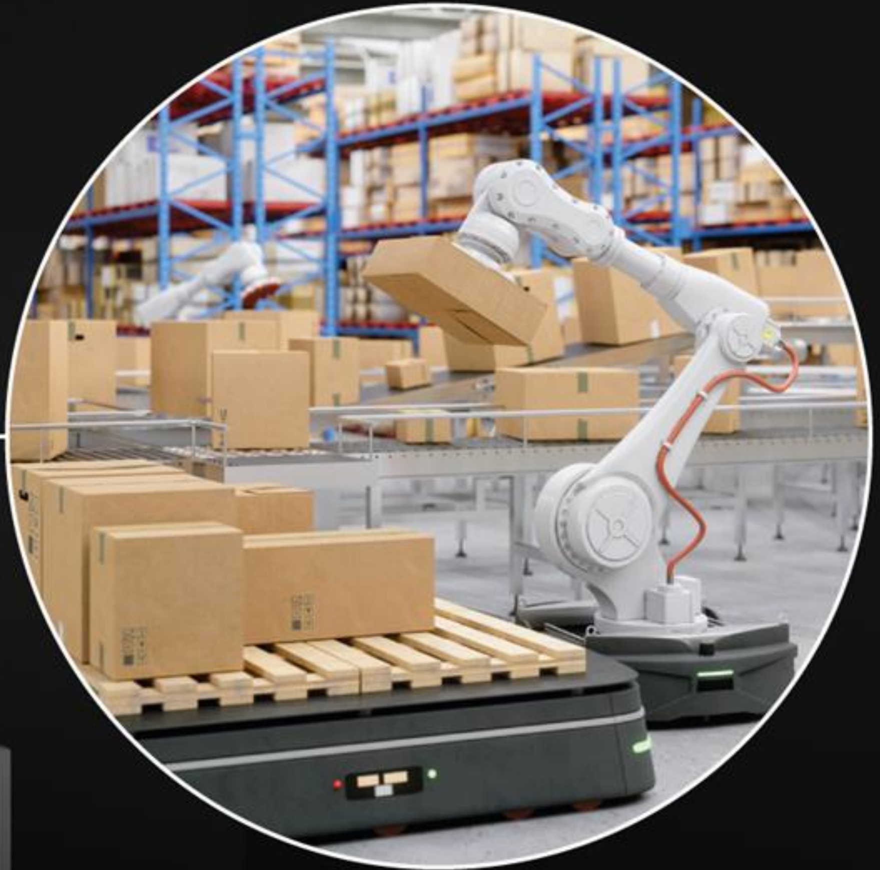

Easily deploy your tinyML solutions with Arduino Pro

arduino.cc/pro

Made In Italy

Build the Future of tinyML

on arm

TINY ML TALKS webcast

EDGE IMPULSE

The Leading Development Platform for Edge ML

edgeimpulse.com

Driving decarbonization and digitalization. Together.

Infineon serving all target markets as
Leader in Power Systems and IoT

www.infineon.com

innatera

NEUROMORPHIC INTELLIGENCE FOR THE SENSOR-EDGE

www.innatera.com

Microsoft

The Right Edge AI Tools Can Make or Break Your Next Smart IoT Product
SensiML
Analytics Toolkit Suite
AutoML
Data Collection
Test & Validation
Data Labeling
Model Building
Code Generation
Team Collaboration
Version Control and Model Management
sensiml.com/tinyML

STMicroelectronics provides extensive solutions to make tiny Machine Learning easy

synaptics

ENGINEERING EXCEPTIONAL EXPERIENCES

We engineer exceptional experiences for consumers in the home, at work, in the car, or on the go.

www.synaptics.com

Silver Strategic Partners

TINY ML TALKS webcast

AIZIP

brainchip

GREENWAVES TECHNOLOGIES

Grovety Inc.

IBM

imagimob

Nota AI

NXP

POLYN TECHNOLOGY

Qeexo

Schneider Electric

SILICON LABS

Join Growing tinyML Communities:

15.6k members in
49 Groups in 41 Countries

tinyML - Enabling ultra-low Power ML at the Edge

https://www.meetup.com/tinyML-Enabling-ultra-low-Power-ML-at-the-Edge/

4k members
&
12.4k followers

The tinyML Community

https://www.linkedin.com/groups/13694488/

TINY ML TALKS webcast

YouTube

9.9k subscribers, 590 videos with 349k views

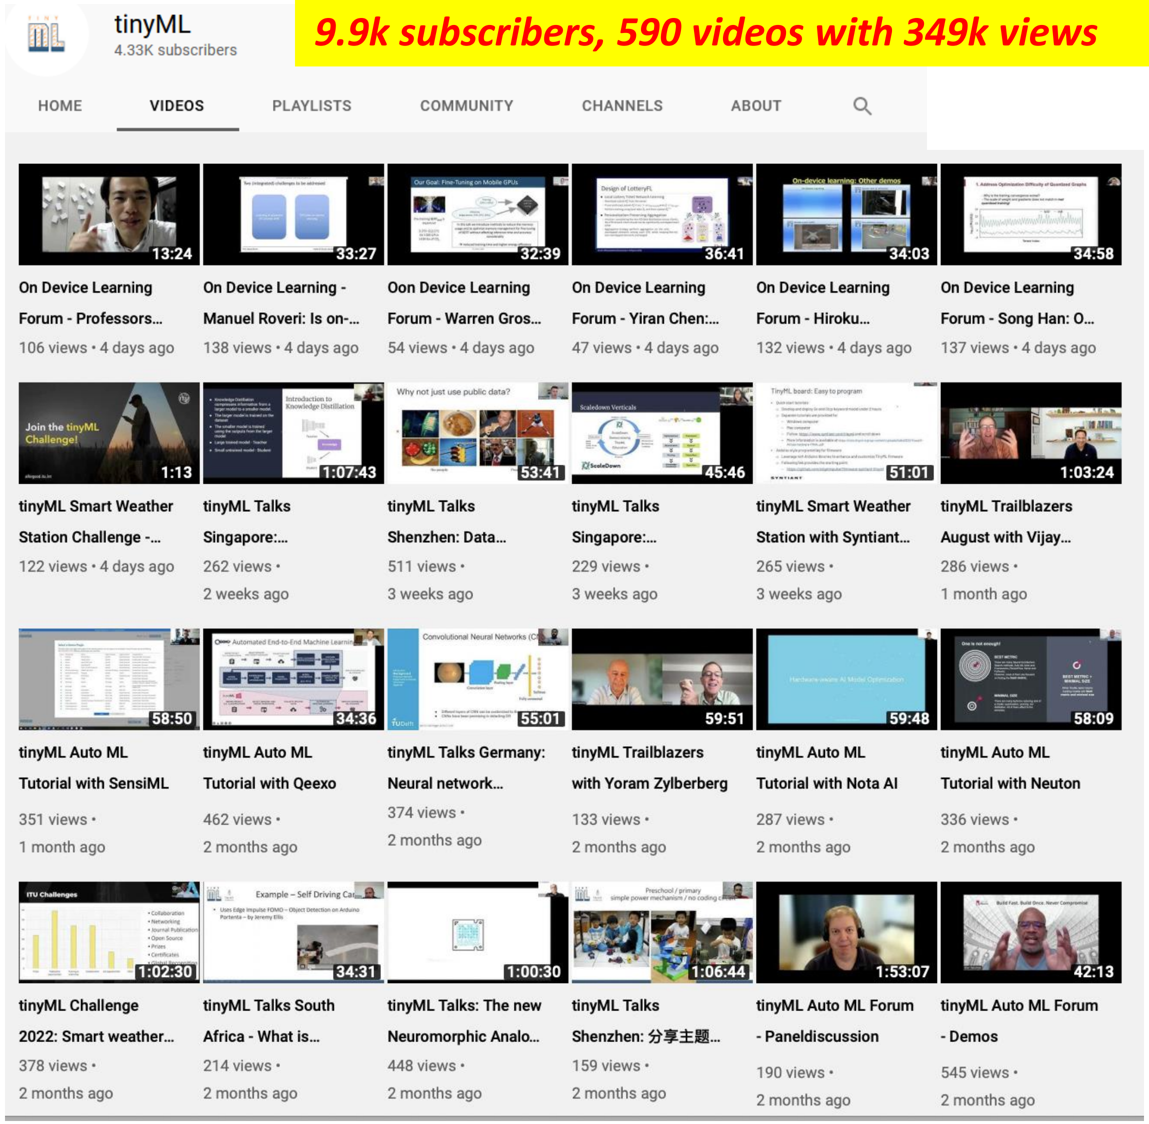

Subscribe to
tinyML YouTube Channel
for updates and notifications
(including this video)

www.youtube.com/tinyML

Reminders

Slides & Videos will be posted tomorrow

tinyml.org/forums   youtube.com/tinyml

Please use the Q&A window for your questions

Matthew Stewart

Matthew Stewart is a postdoctoral researcher in the Edge Computing Lab at Harvard University. He holds a Ph.D. and MSc in Engineering Sciences and Data Science from Harvard University, and an integrated BEng/MEng in Mechanical Engineering from Imperial College London and the National University of Singapore. Matthew's research work is highly interdisciplinary, encompassing embedded machine learning, autonomous vehicles, benchmarking tools for reinforcement learning and robotics, sustainable computing, and machine learning sensors. Matthew is also a part-time blogger for Towards Data Science, a co-creator of the HarvardX tinyML courses, and a research coordinator at MLCommons.

Datasheets for Machine Learning Sensors

Acknowledgements: *B. Brown, Y. Omri, J. Santos, J. MacArthur, B. Plancher, S. Prakash, N. Jeffries, V. J. Reddi, P. Warden & the Useful Sensors Team*

Matthew Stewart, Ph. D. | Postdoctoral Researcher |
John A. Paulson School of Engineering and Applied Sciences | Harvard University |
Web: https://mpstewart.io

Harvard Radcliffe Institute

useful sensors

Miniaturization of Compute

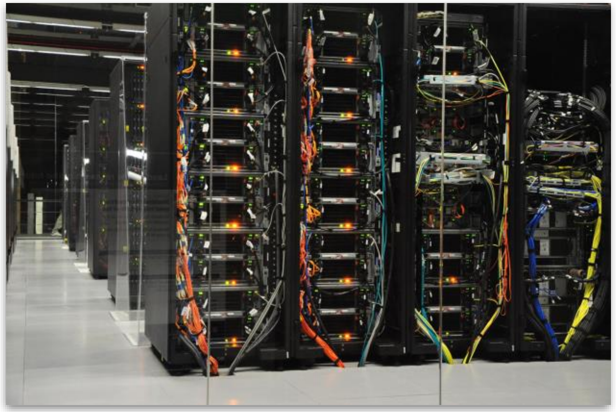

Data Centers

Edge Devices

Internet-of-Things

TinyML

No Good Data Left Behind

5 Quintillion

bytes of data produced every day by IoT

<1%

of unstructured data is analyzed or used at all

Source: Harvard Business Review, [What's Your Data Strategy?](), April 18, 2017
Cisco, Internet of Things (IoT) Data Continues to Explode Exponentially. Who Is Using That Data and How?, Feb 5, 2018

The “Classic” TinyML Paradigm

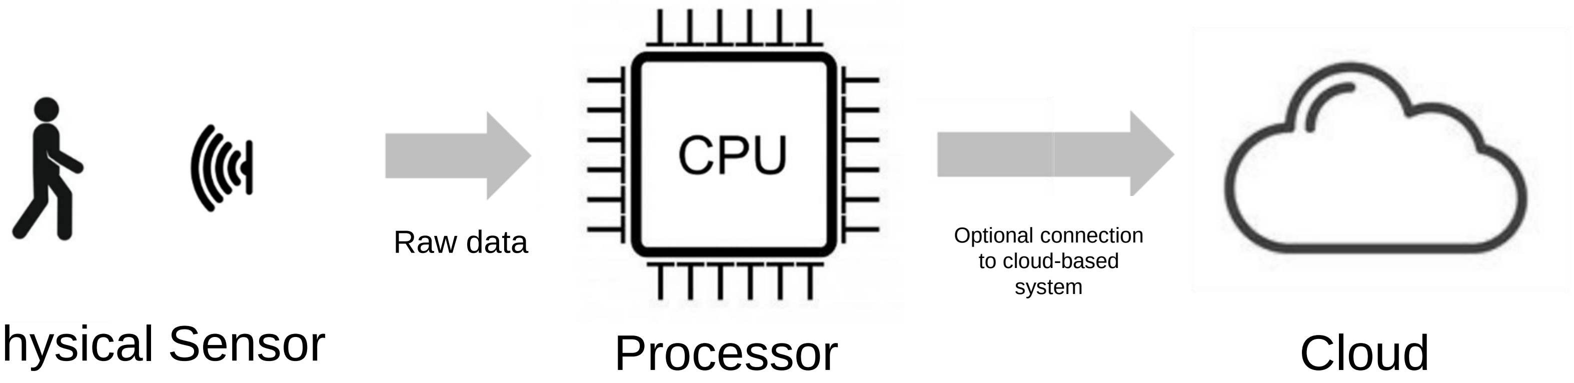

Sensor 1.0

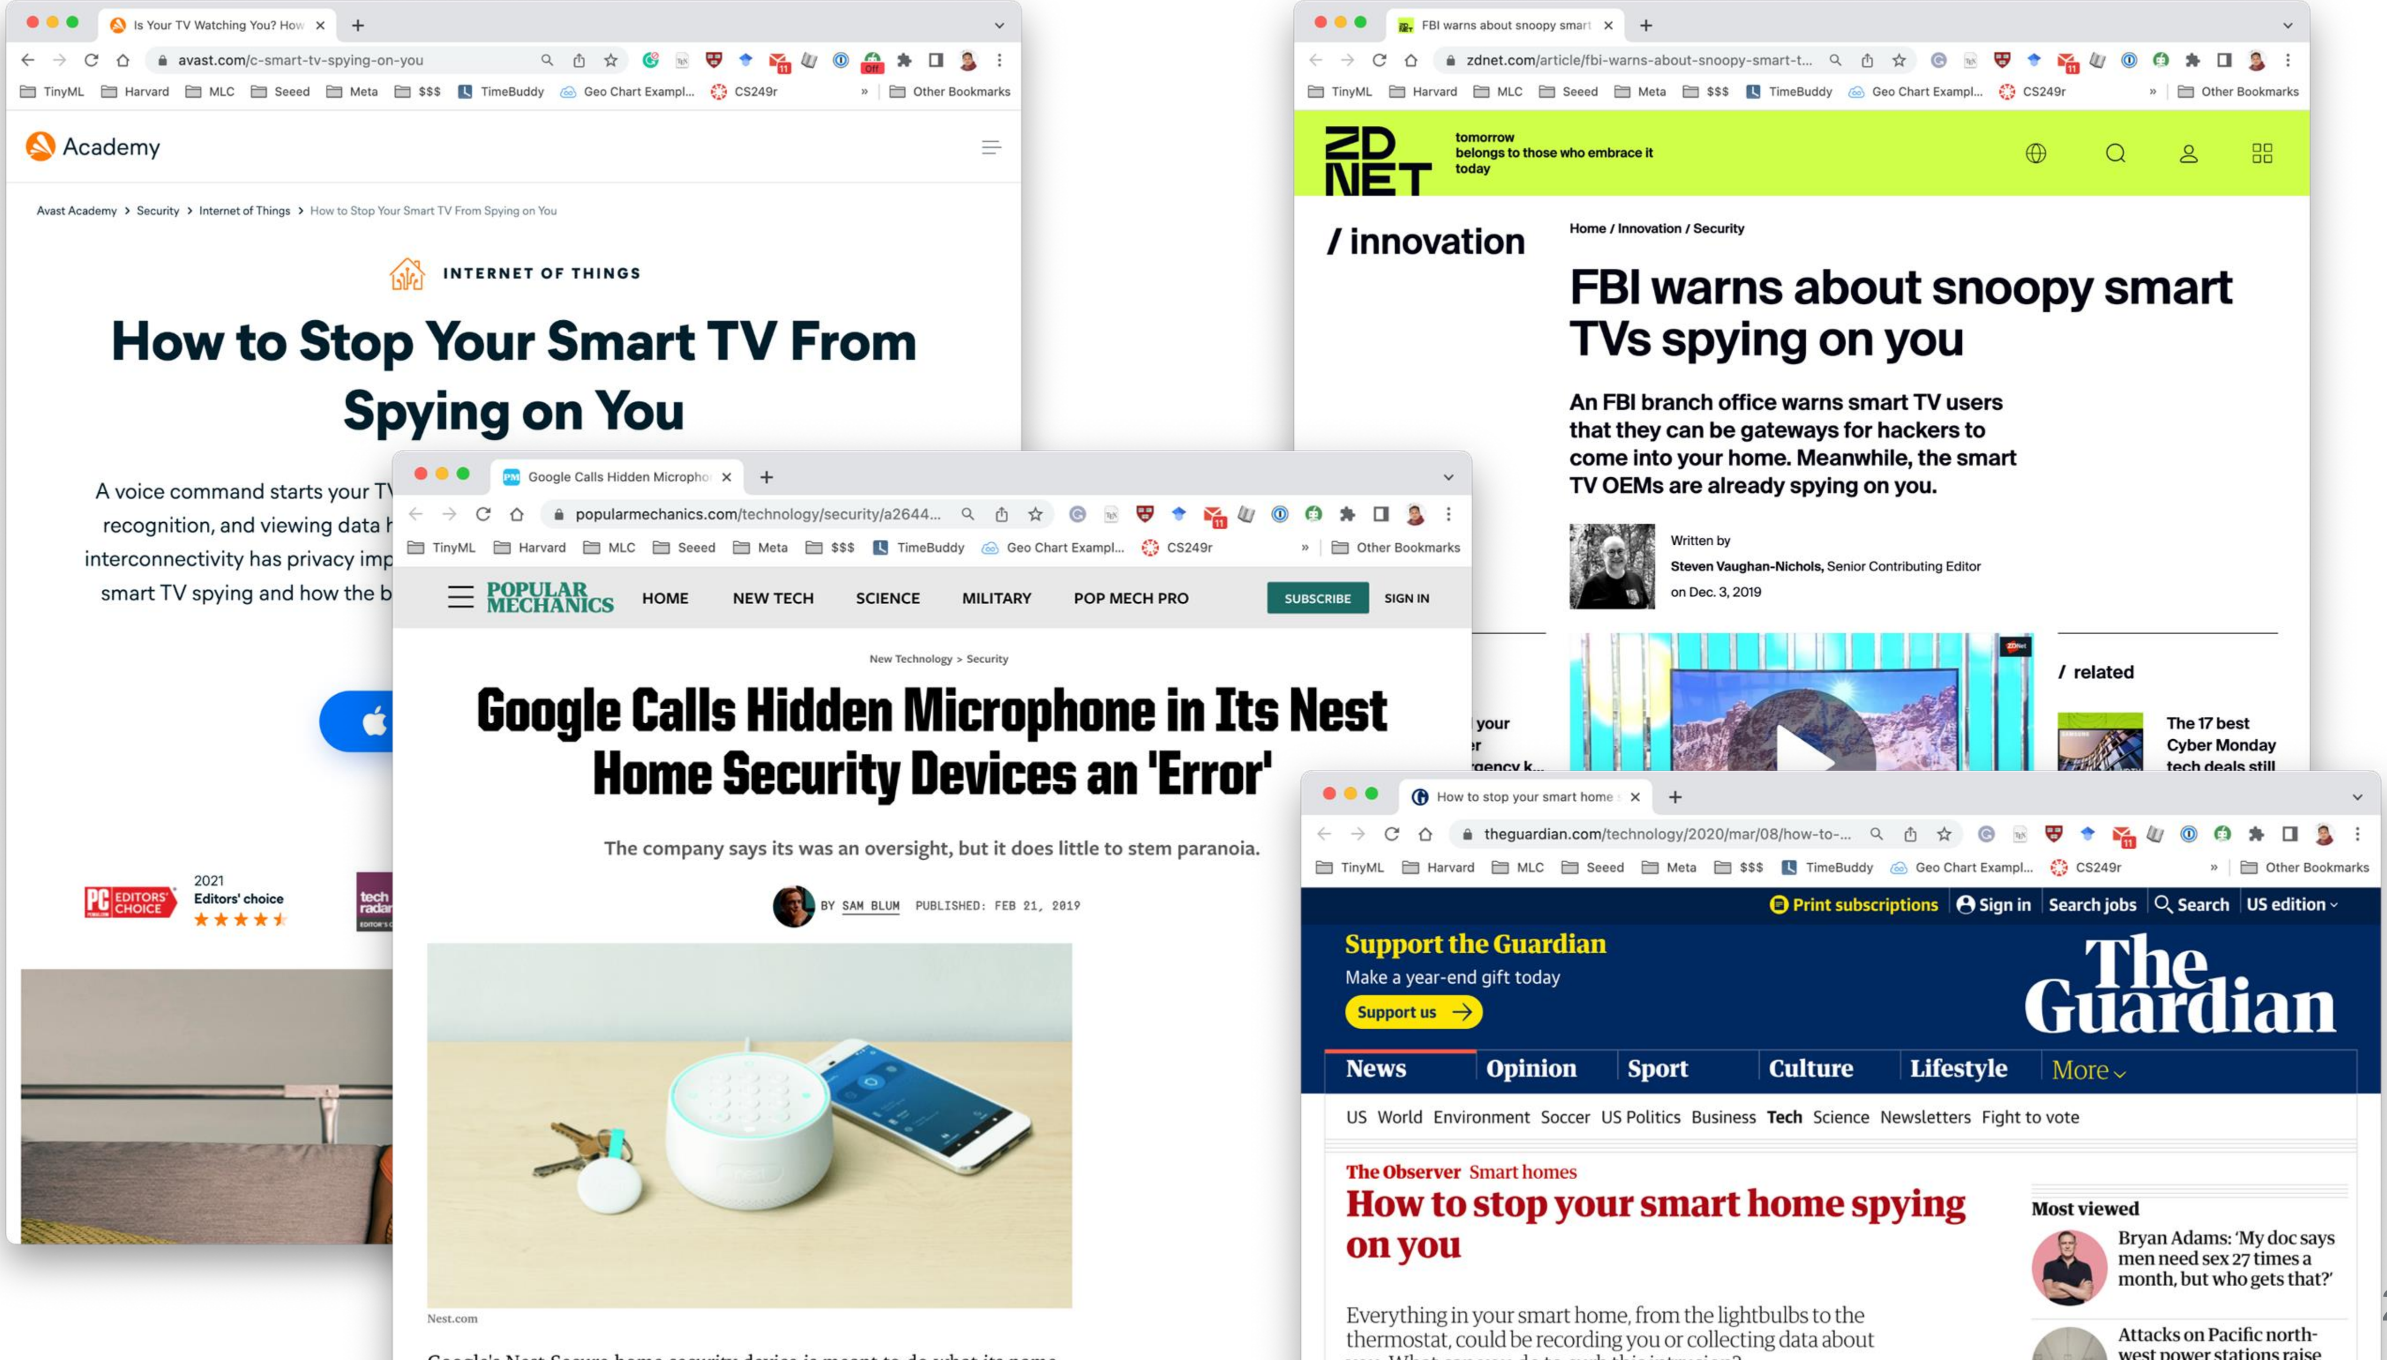
Is Your TV Watching You? How
avast.com/c-smart-tv-spying-on-you
Academy
Avast Academy > Security > Internet of Things > How to Stop Your Smart TV From Spying on You
INTERNET OF THINGS
How to Stop Your Smart TV From Spying on You
2021 Editors' choice
FBI warns about snoopy smart
zdnet.com/article/fbi-warns-about-snoopy-smart-t...
ZD NET
tomorrow belongs to those who embrace it today
/ innovation
Home / Innovation / Security
FBI warns about snoopy smart TVs spying on you
An FBI branch office warns smart TV users that they can be gateways for hackers to come into your home. Meanwhile, the smart TV OEMs are already spying on you.
Written by
Steven Vaughan-Nichols, Senior Contributing Editor
on Dec. 3, 2019
/ related
Google Calls Hidden Micropho
popularmechanics.com/technology/security/a2644...
POPULAR MECHANICS
HOME
NEW TECH
SCIENCE
MILITARY
POP MECH PRO
SUBSCRIBE
SIGN IN
New Technology > Security
Google Calls Hidden Microphone in Its Nest Home Security Devices an 'Error'
The company says its was an oversight, but it does little to stem paranoia.
BY SAM BLUM PUBLISHED: FEB 21, 2019
Nest.com
How to stop your smart home
theguardian.com/technology/2020/mar/08/how-to-...
Print subscriptions
Sign in
Search jobs
Search
US edition
Support the Guardian
Make a year-end gift today
Support us
The Guardian
News
Opinion
Sport
Culture
Lifestyle
More
US World Environment Soccer US Politics Business Tech Science Newsletters Fight to vote
The Observer Smart homes
How to stop your smart home spying on you
Most viewed
Bryan Adams: 'My doc says men need sex 27 times a month, but who gets that?'

How do we architect future Tiny Machine Learning (tinyML) sensors efficiently, effectively and robustly into the embedded ecosystem?

Machine Learning Sensors

Machine Learning Sensors

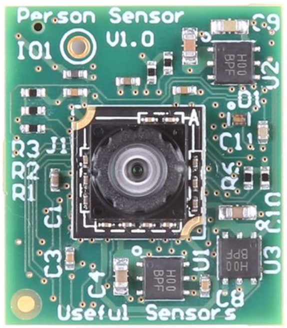

by Useful Sensors

> An ML sensor is a **self-contained system** that utilizes **on-device machine learning** to extract **useful information** by observing some complex set of phenomena in the **physical world** and reports it through a **simple interface** to a wider system.

Machine Learning Sensors

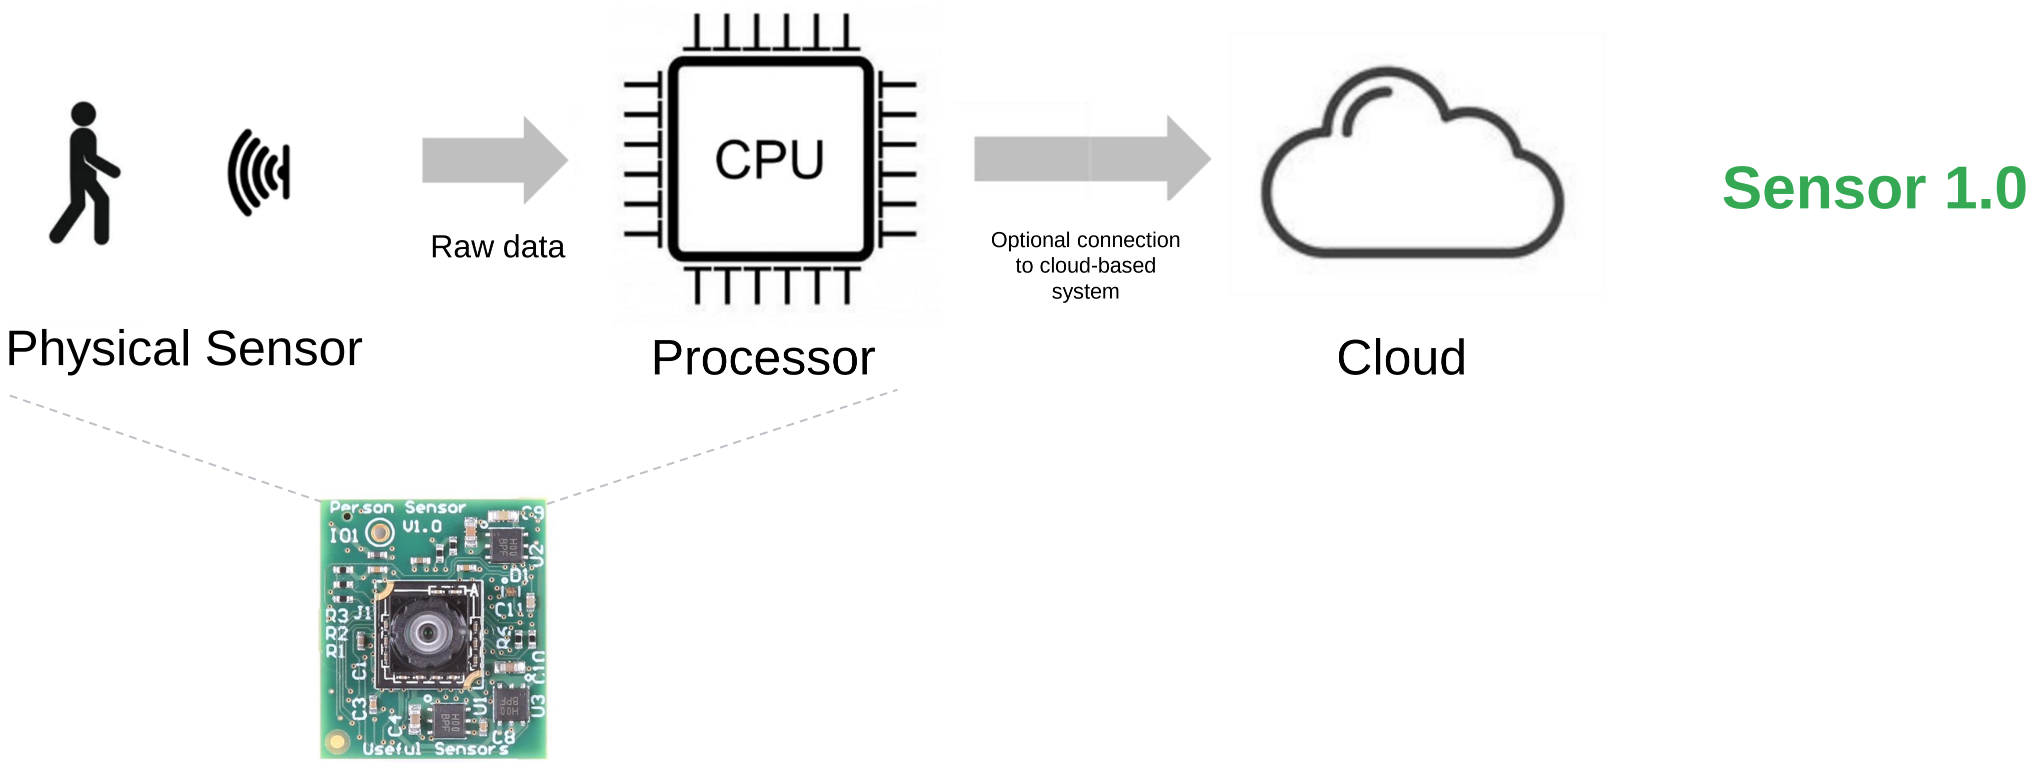

Machine Learning Sensors

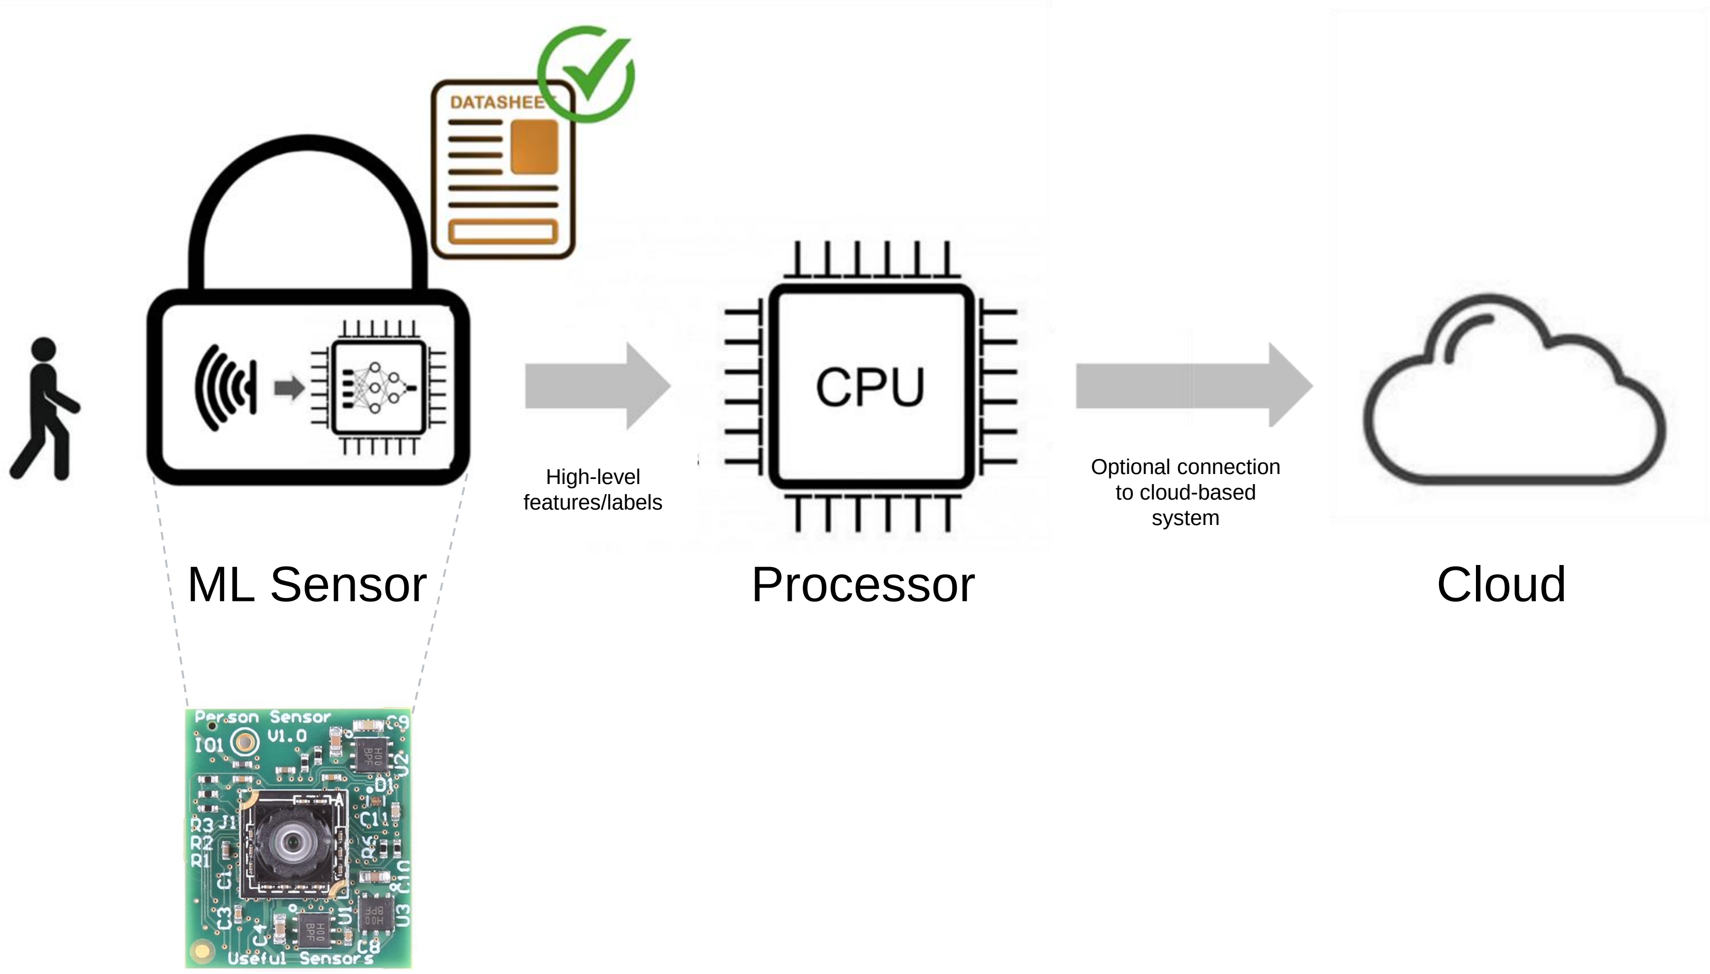

Sensor 2.0

ML Sensors - Guiding Set of Principles

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

ML Sensor Principles — Abstraction

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

ML Sensor Principles — Abstraction

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

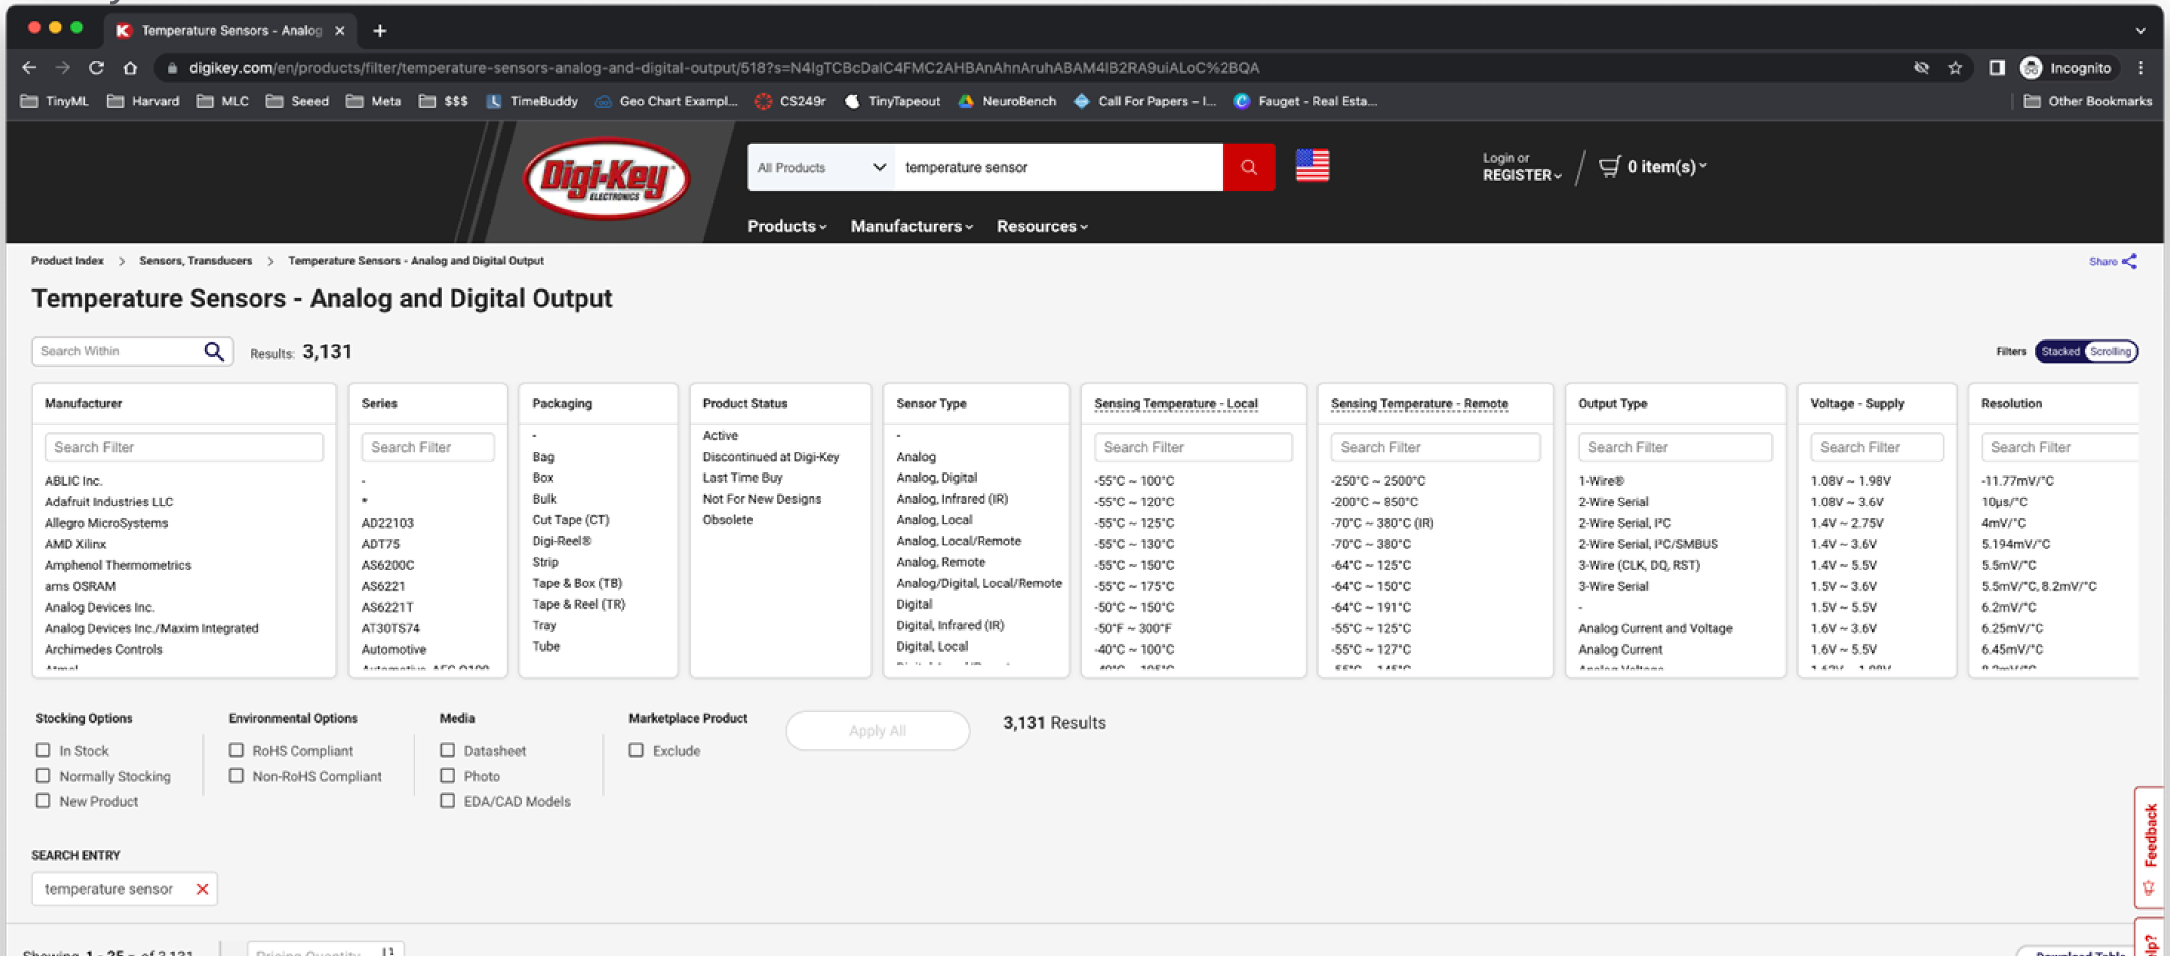

Temperature Sensors - Analog

digikey.com/en/products/filter/temperature-sensors-analog-and-digital-output/518?s=N4IgTCBcDaIC4FMC2AHBAnAhnAruhABAM4IB2RA9uiALoC%2BQA

TinyML Harvard MLC Seeed Meta $$$ TimeBuddy Geo Chart Exampl... CS249r TinyTapeout NeuroBench Call For Papers – I... Fauget - Real Esta... Other Bookmarks

Incognito

Compare	Mfr Part #	Quantity Available	Price	Series	Package	Product Status	Sensor Type	Sensing Temperature - Local	Sensing Temperature - Remote	Output Type	Voltage - Supply	Resolution	Features	Accuracy - Highest (Lowest)	Test Condition
	TMP236A2DBZT SENSOR TEMPERATURE *Texas Instruments*	1,053 In Stock	1 : $1.49000 Cut Tape (CT) 250 : $0.71800 Tape & Reel (TR)	-	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Analog, Local	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±2°C	-10°C ~ 125°C
	TMP236A4DCKT SENSOR TEMPERATURE *Texas Instruments*	1,678 In Stock	1 : $1.24000 Cut Tape (CT) 250 : $0.59800 Tape & Reel (TR)	-	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Analog, Local	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±4°C	-10°C ~ 125°C
	TMP236A2DCKT SENSOR TEMPERATURE *Texas Instruments*	2,307 In Stock	1 : $1.41000 Cut Tape (CT) 250 : $0.67800 Tape & Reel (TR)	-	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Analog, Local	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±2°C	-10°C ~ 125°C
	TMP451JQDQFTQ1 SENSOR TEMPERATURE *Texas Instruments*	340 In Stock	1 : $2.06000 Cut Tape (CT) 250 : $1.02860 Tape & Reel (TR)	Automotive, AEC-Q100	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Digital, Local/Remote	-40°C ~ 125°C	-64°C ~ 191°C	I²C/SMBus	1.7V ~ 3.6V	12 b	One-Shot, Output Switch, Programmable Limit, Shutdown Mode	±1°C (±2°C)	0°C ~ 70°C (-40°C ~ 125°C)
	TMP236A4DBZT SENSOR TEMPERATURE *Texas Instruments*	596 In Stock	1 : $1.32000 Cut Tape (CT) 250 : $0.63800 Tape & Reel (TR)	-	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Analog, Local	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±4°C	-10°C ~ 125°C
	TMP451HQDQFTQ1 SENSOR TEMPERATURE *Texas Instruments*	227 In Stock 3,250 Factory	1 : $2.06000 Cut Tape (CT) 250 : $1.02860 Tape & Reel (TR)	Automotive, AEC-Q100	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Digital, Local/Remote	-40°C ~ 125°C	-64°C ~ 191°C	I²C/SMBus	1.7V ~ 3.6V	12 b	One-Shot, Output Switch, Programmable Limit, Shutdown Mode	±1°C (±2°C)	0°C ~ 70°C (-40°C ~ 125°C)
	TMP461AIRUNT-S TEMPERATURE SENSOR *Texas Instruments*	9,073 In Stock 10,000 Factory	1 : $2.56000 Cut Tape (CT) 250 : $1.28020 Tape & Reel (TR)	-	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Digital, Local/Remote	-40°C ~ 125°C	-64°C ~ 191°C	SMBus	1.7V ~ 3.6V	11 b	One-Shot, Output Switch, Programmable Limit, Shutdown Mode, Standby Mode	±1°C (±1.25°C)	-10°C ~ 100°C (-40°C ~ 125°C)
	TMP12FP ANALOG TEMPERATURE SENSOR *Analog Devices Inc.*	1,253 Marketplace	108 : $2.79000 Bulk	-	Bulk	Active	Digital, Local	-40°C ~ 125°C	-	SPI	2.7V ~ 5.5V	12 b	One-Shot, Shutdown Mode	±2°C (±2.5°C)	-25°C ~ 85°C (-40°C ~ 125°C)
	MAX6630MUT-T DIGITAL TEMPERATURE SENSOR *Analog Devices Inc./Maxim Integrated*	3,396 Marketplace	110 : $2.75000 Bulk	-	Bulk	Active	Digital, Local	-55°C ~ 125°C	-	SPI	3V ~ 5.5V	12 b	Shutdown Mode	±0.8°C (-5°C, 6.5°C)	25°C (150°C)
	TMP35FT9 ANALOG TEMPERATURE SENSOR *Analog Devices Inc.*	20,365 Marketplace	298 : $1.01000 Bulk	Automotive	Bulk	Active	Analog, Local	10°C ~ 125°C	-	Analog Voltage	2.7V ~ 5.5V	10mV/°C	Shutdown Mode	±2°C (±3°C)	25°C (10°C ~ 125°C)
	MAX6629MUT-T DIGITAL TEMPERATURE SENSOR *Analog Devices Inc./Maxim Integrated*	9,848 Marketplace	139 : $2.17000 Bulk	-	Bulk	Active	Digital, Local	-55°C ~ 125°C	-	SPI	3V ~ 5.5V	12 b	Shutdown Mode	±0.8°C (-5°C, 6.5°C)	25°C (150°C)
	AD22103KR-REEL	2,350 Marketplace	289 : $1.04000 Bulk	AD22103	Bulk	Active	Analog, Local	0°C ~ 100°C	-	Analog	2.7V ~ 3.6V	28mV/°C	-	±2°C (±2.5°C)	25°C (0°C ~

Feedback

Need Help?

ML Sensor Principles — Abstraction

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

Person detector

Gaze sensor

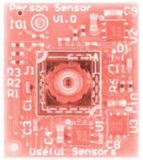

Voice command

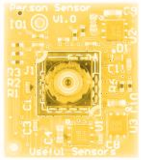

Text recognizer

…

Digi-Key ELECTRONICS amazon.com Alibaba . . .

property.

ML Sensor Principles — Data-centric

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

ML Sensor Principles — Data-centric

1. We need to raise the level of abstraction to enable ease of use for scalable deployment of ML sensors;
2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

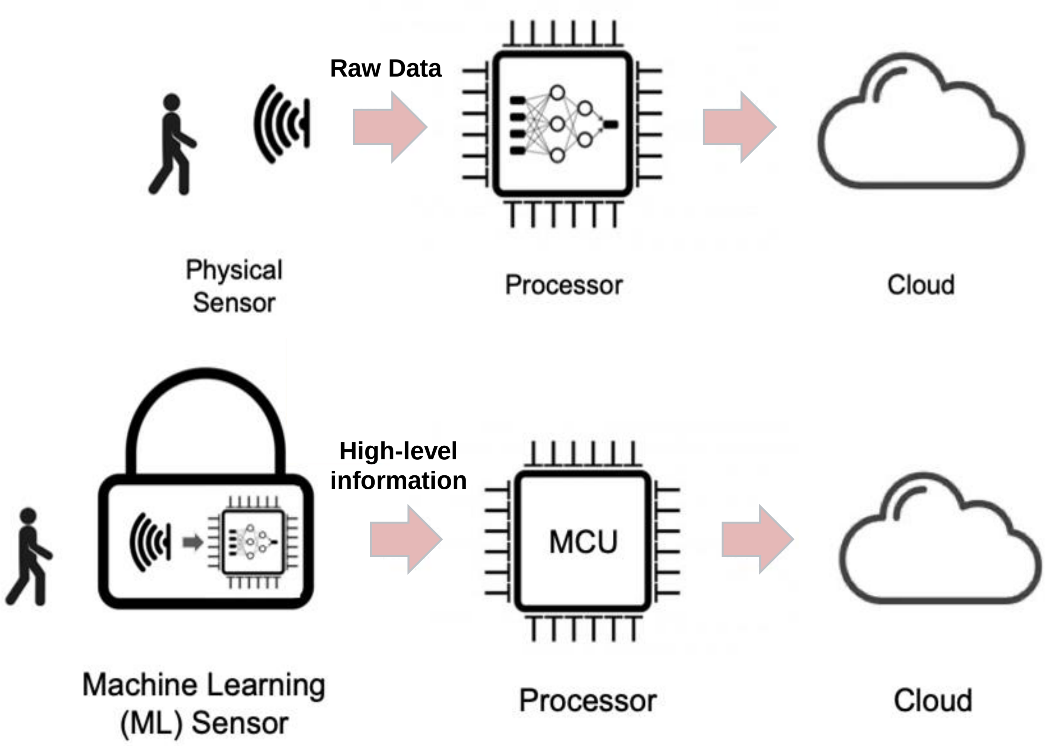

Sensor 1.0

Sensor 2.0

ML Sensor Principles — Simplicity

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

ML Sensor Principles — Simplicity

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

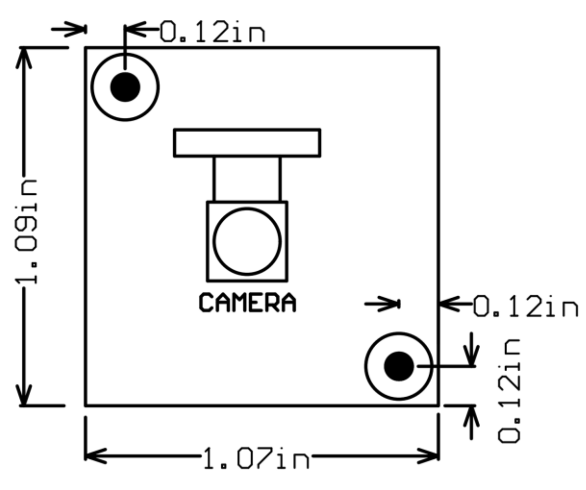

We need to define or rely on standard interfaces and mechanisms for communication with sensors.

Source: https://github.com/usefulsensors/person_sensor_docs

ML Sensor Principles — Simplicity

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

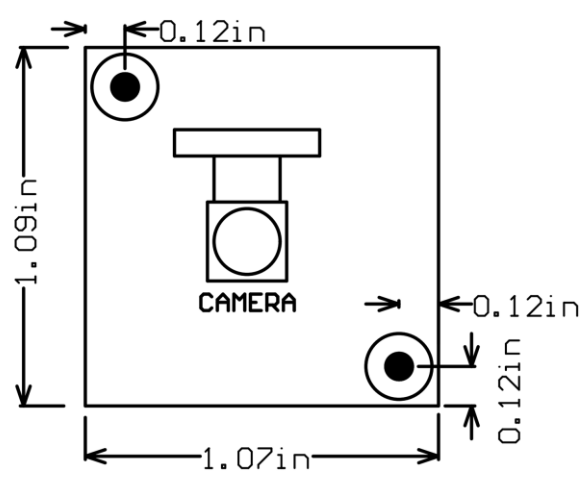

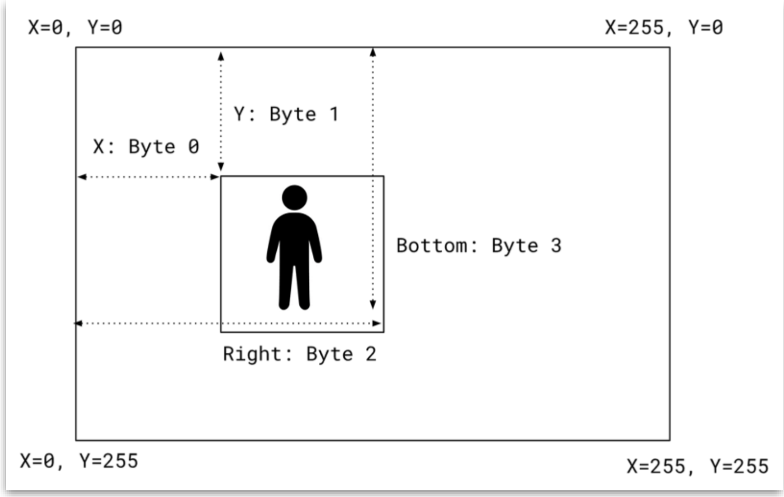

We need to define data formats to enable interoperability and exchange of ML sensors across manufacturers

Source: https://github.com/usefulsensors/person_sensor_docs

ML Sensor Principles — Simplicity

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

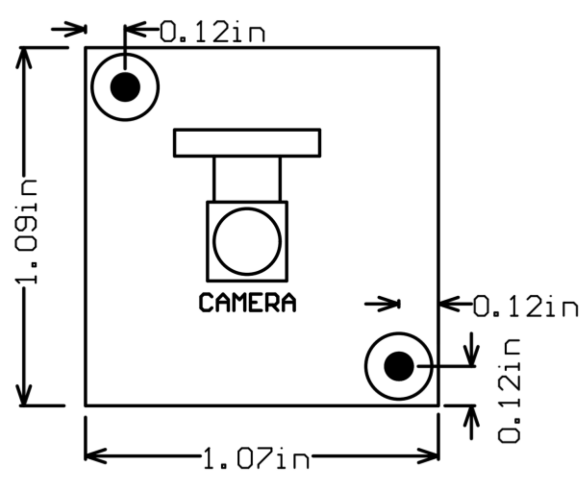

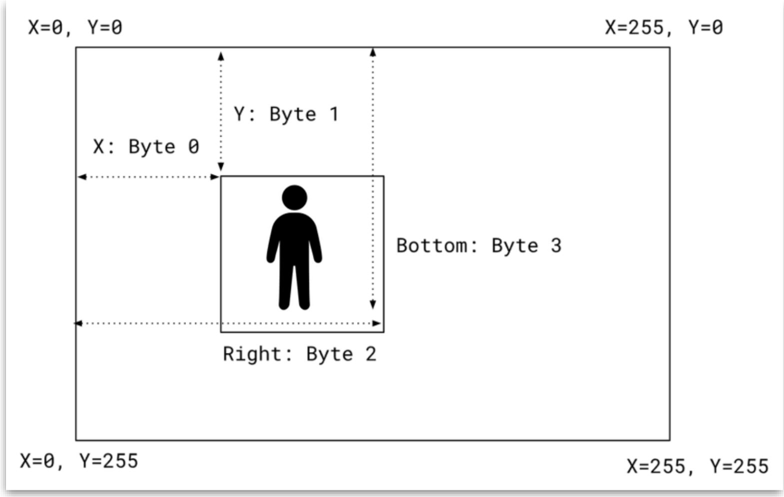


```
typedef struct __attribute__ ((__packed__)) {
    uint8_t reserved[2];
    uint16_t data_size;
} person_sensor_results_header_t;

typedef struct __attribute__ ((__packed__)) {
    uint8_t box_confidence;
    uint8_t box_left;
    uint8_t box_top;
    uint8_t box_width;
    uint8_t box_height;
    int8_t id_confidence;
    int8_t id;
    uint8_t is_facing;
} person_sensor_face_t;

typedef struct __attribute__ ((__packed__)) {
    person_sensor_results_header_t header;
    int8_t num_faces;
    person_sensor_face_t faces[4];
    uint16_t checksum;
} person_sensor_results_t;
```

Source: https://github.com/usefulsensors/person_sensor_docs

ML Sensor Principles — Documentation

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

ML Sensor Principles — Documentation

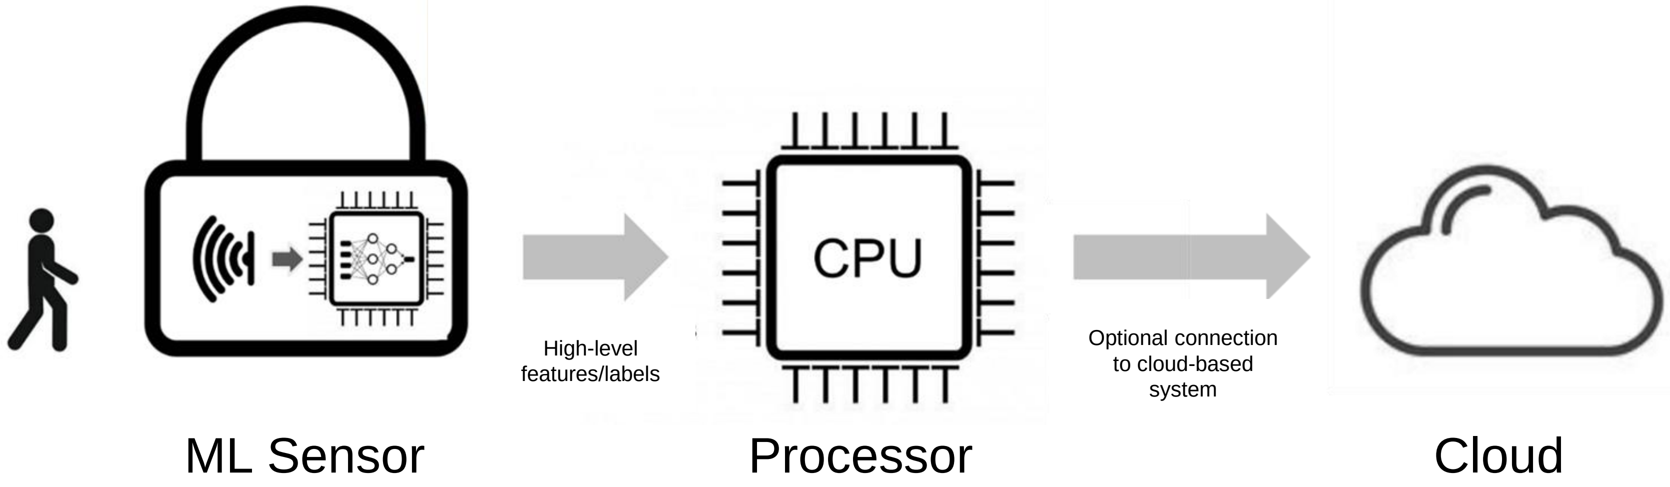

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

ML Sensor Principles — Documentation

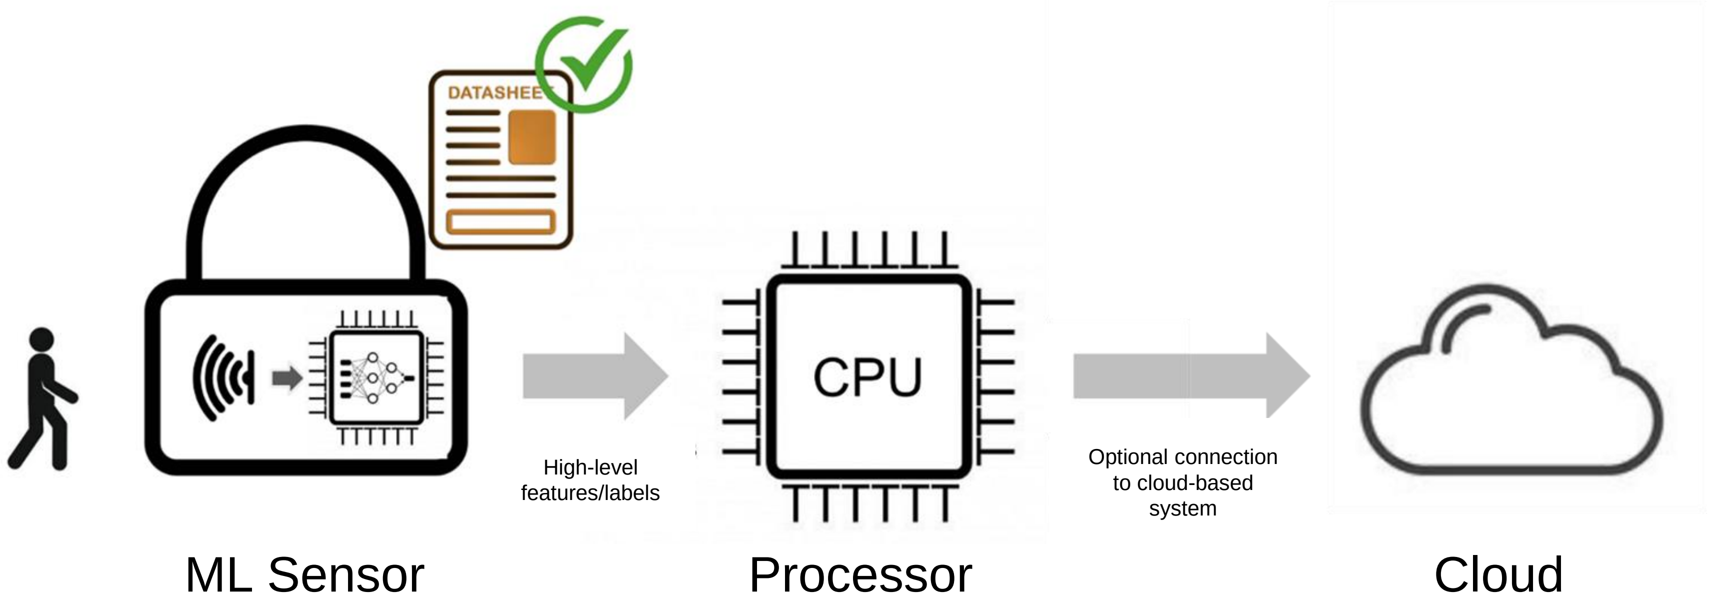

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

Datasheet Template

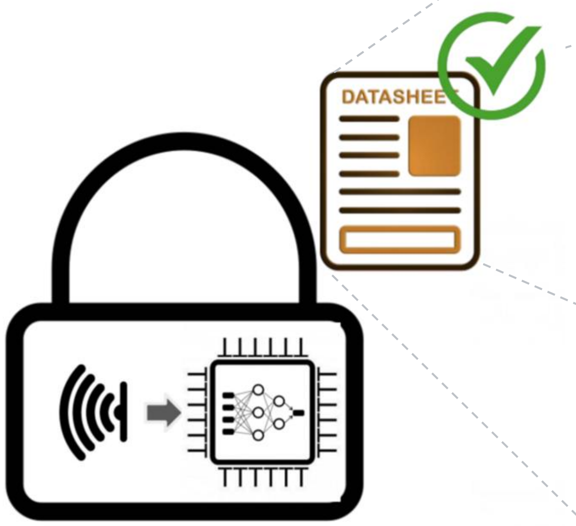

Description, Features and Use case

Compliance

Environmental Impact

Privacy and Security

Diagram and Form Factor

Data Nutrition Label

Hardware Characteristics

Communication Specification and Pinout

Model Characteristics

End-to-End Performance Analysis

Model Characteristics

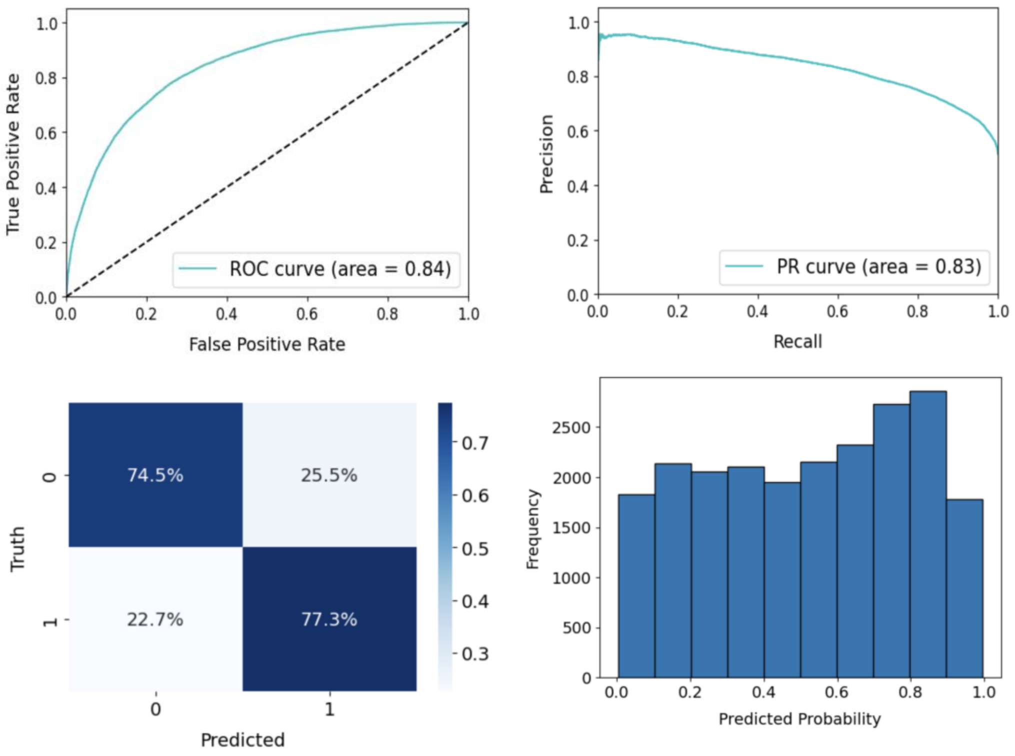

- Provides a **concise, holistic picture of the performance characteristics** of a machine learning model
- For our sensor, this is a **binary classification person detection model** which processes **raw image data**

(Mitchell et al., 2019)

Description, Features and Use case

Compliance

Environmental Impact

Privacy and Security

Diagram and Form Factor

Data Nutrition Label

Hardware Characteristics

Communication Specification and Pinout

Model Characteristics

End-to-End Performance Analysis

Data Nutrition Label

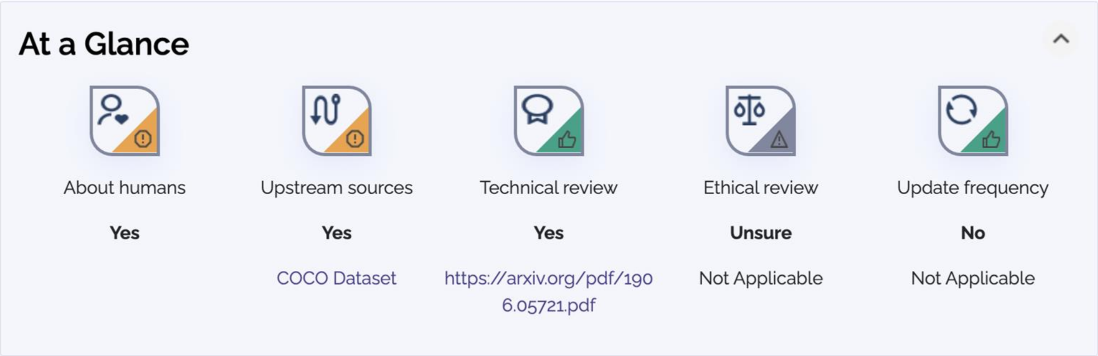

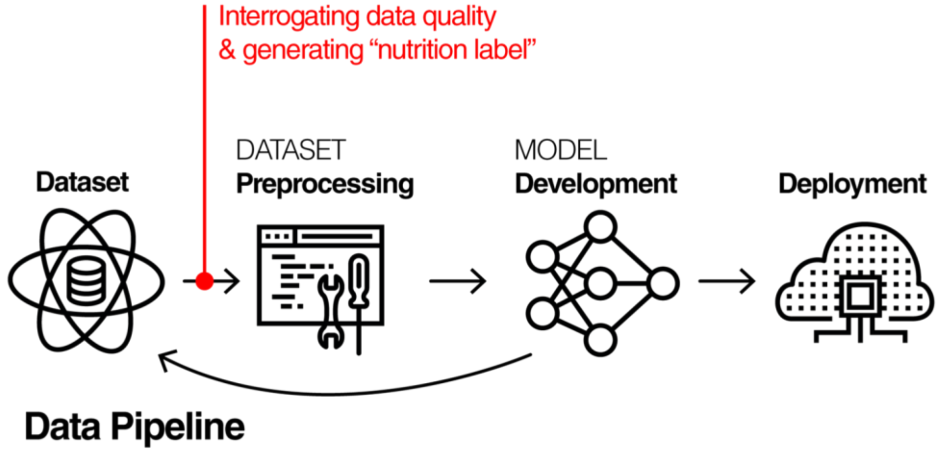

- Offers a **succinct, comprehensive snapshot of dataset attributes** used for model training
- For our sensor, this focuses on attributes of the **visual wake words dataset** derived from MS-COCO.

(Holland et al., 2020; Chmielinski et al., 2022)

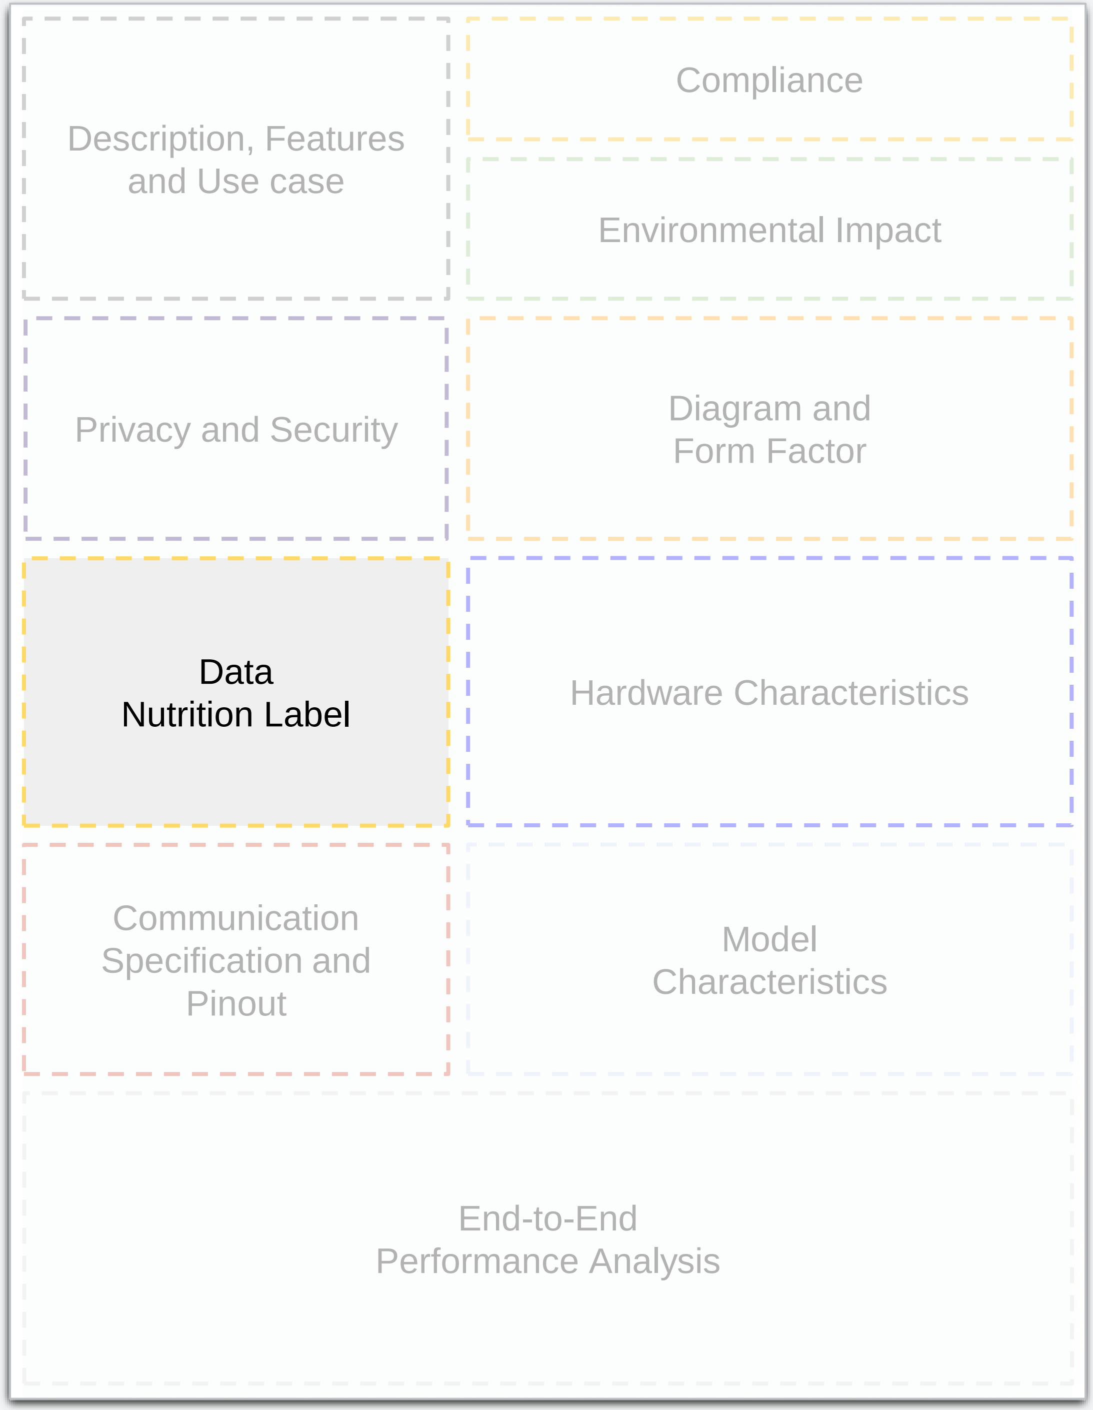

Privacy and Security Label

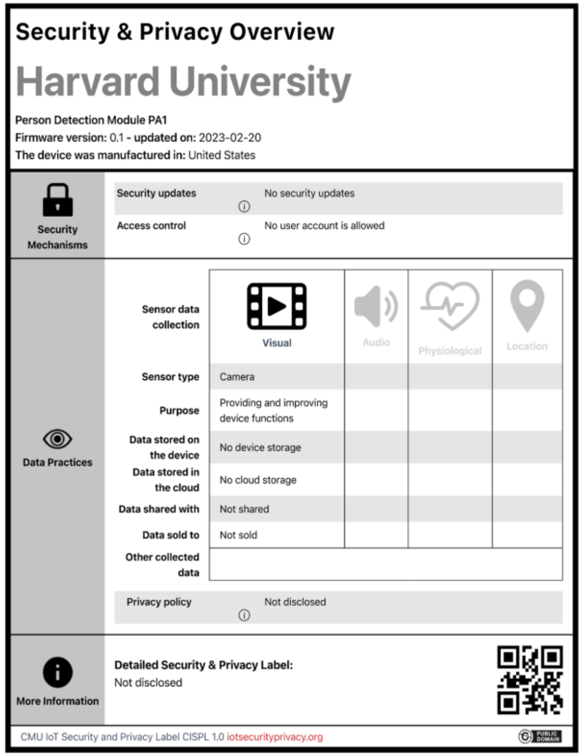

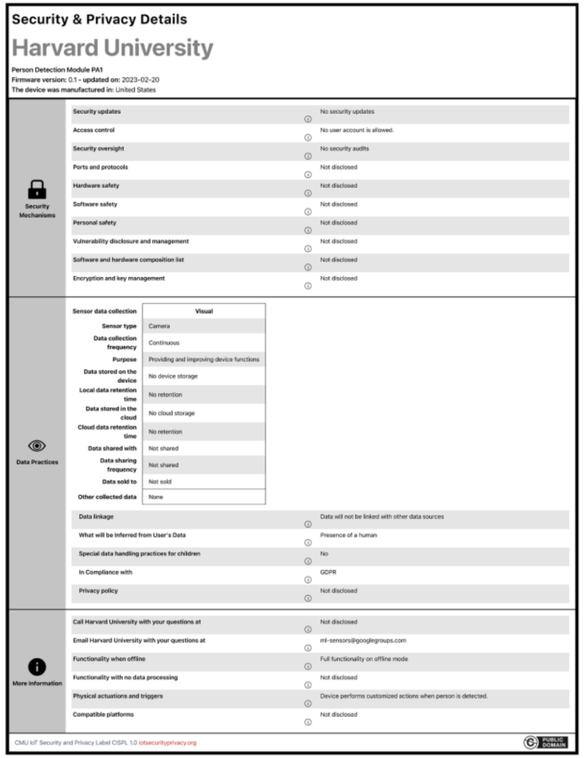

- **Privacy.** Provides **clear and transparent information regarding data capture, usage, and storage** for **each data modality**.
- **Security.** Safety protocols and security mechanisms associated with the device are outlined.

(Emami-Naeini et al., 2021)

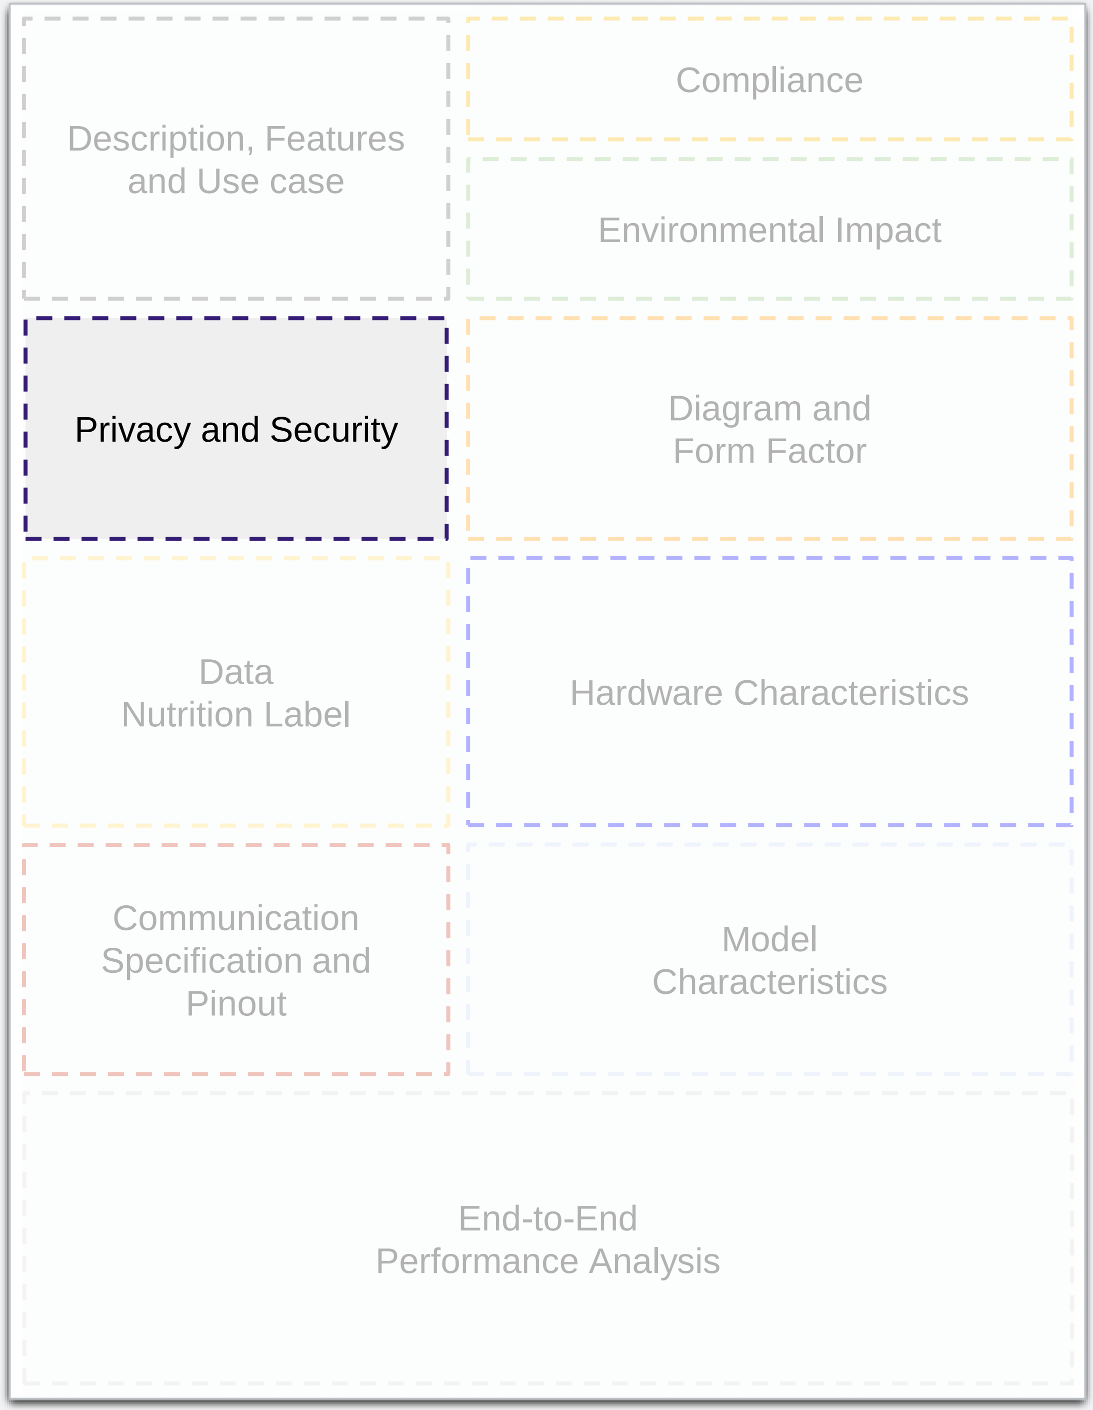

Performance Analysis

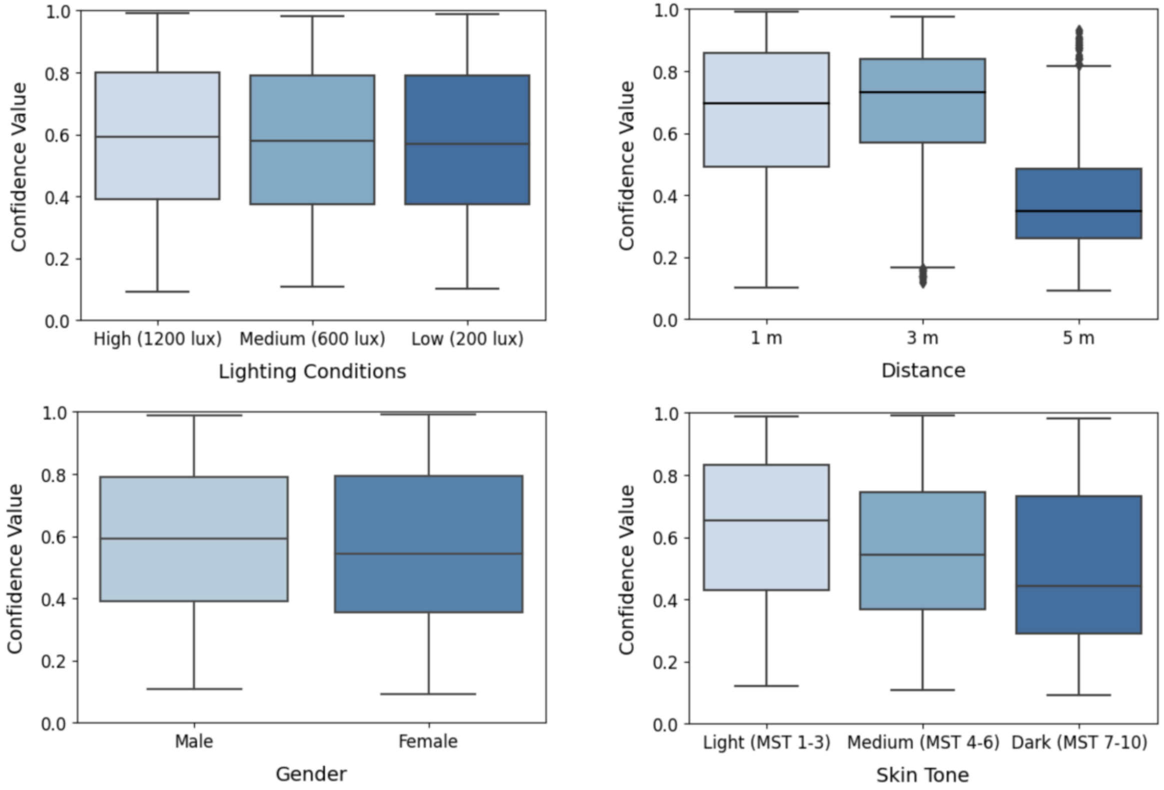

- Provides an indication of **demographic biases** as well as performance changes under **varying environmental conditions**.
- **Experimental study conducted** under different lighting and distances to assess sensor performance in **real-world conditions**.

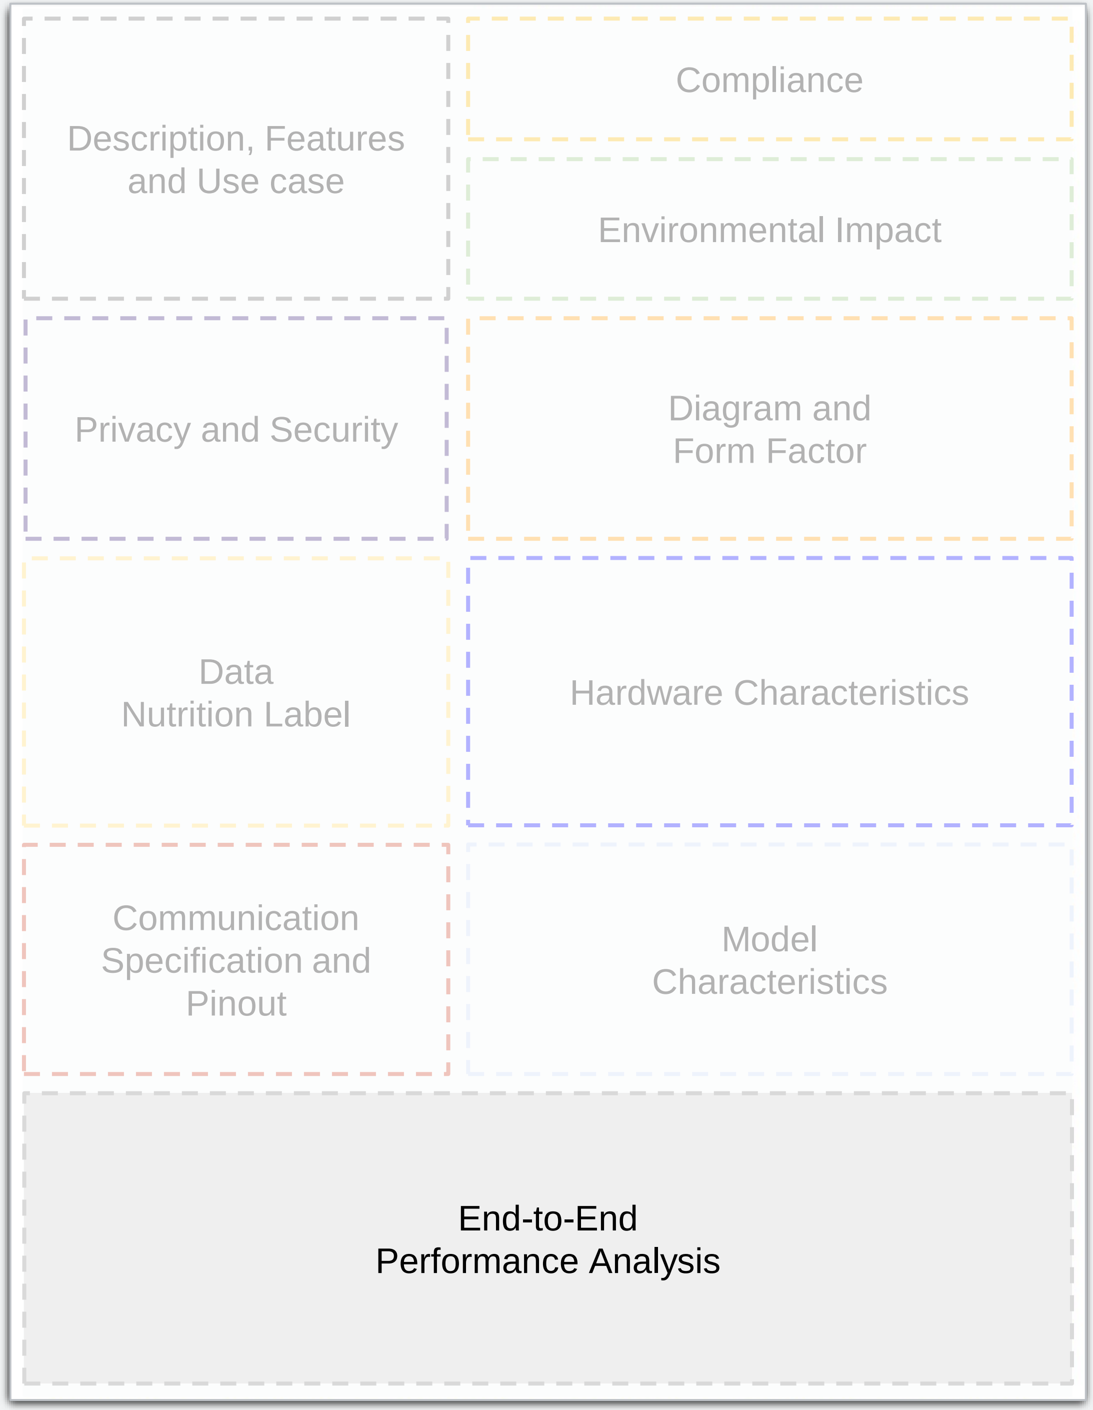

Experimental Study Details

- Data collected from 39 participants
- Each participant tested at three different distances and three lighting levels (nine total scenarios)
- Ten measurements (~10 s) taken at each location and averaged
- Six sensors utilized: three open-source ML sensors and three commercially available (Useful Sensors)
- Neutral background environment with no ambient light exposure; 32 homogeneously distributed overhead lights

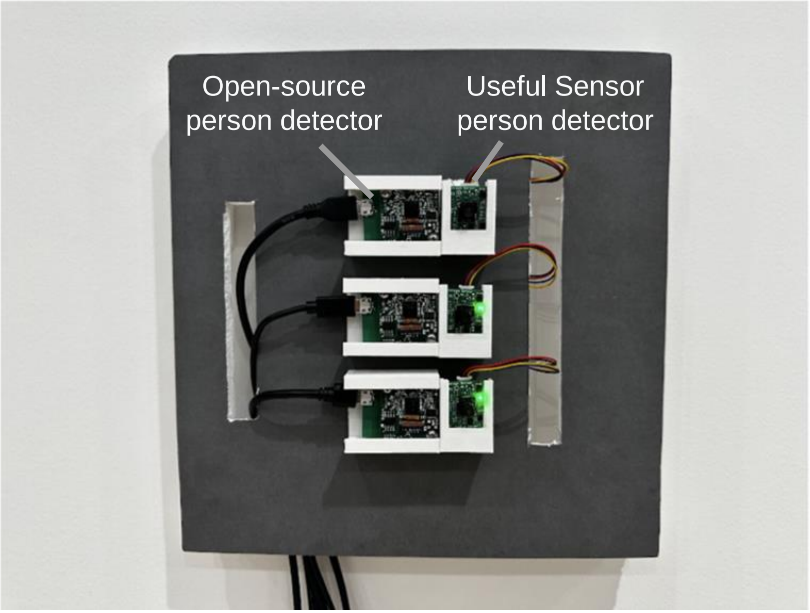

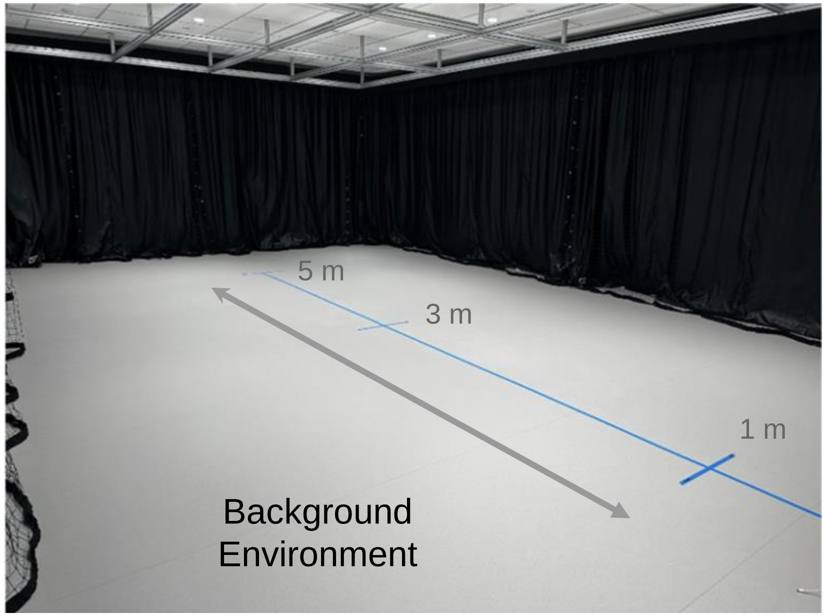

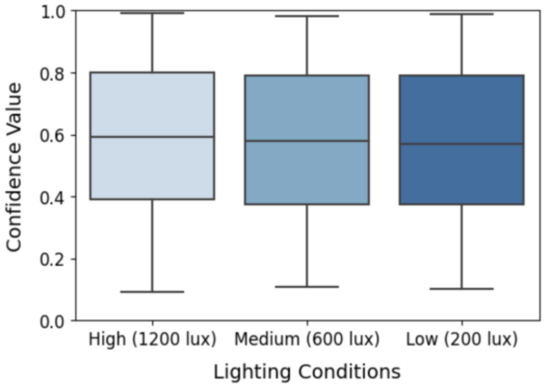

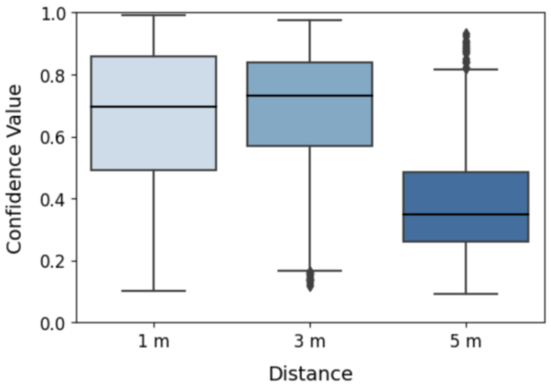

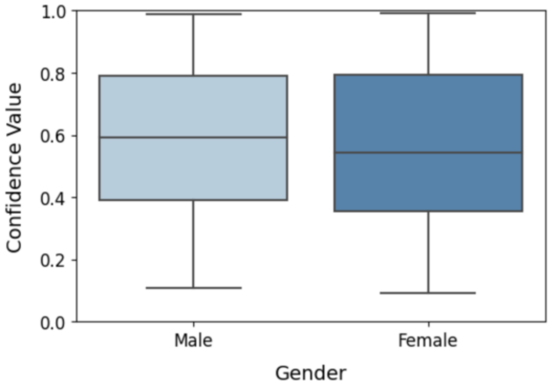

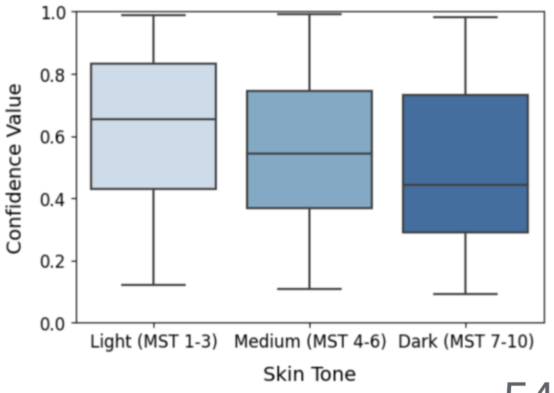

Compliance

ML sensors should be **tested by 3rd party certification** agencies or bodies that **specialize in AI/ML technologies to:**

1. Ensure adherence to **industry standards and regulations**
1. Support **data integrity and accuracy**
1. **Foster trust** in the product's **performance and reliability**

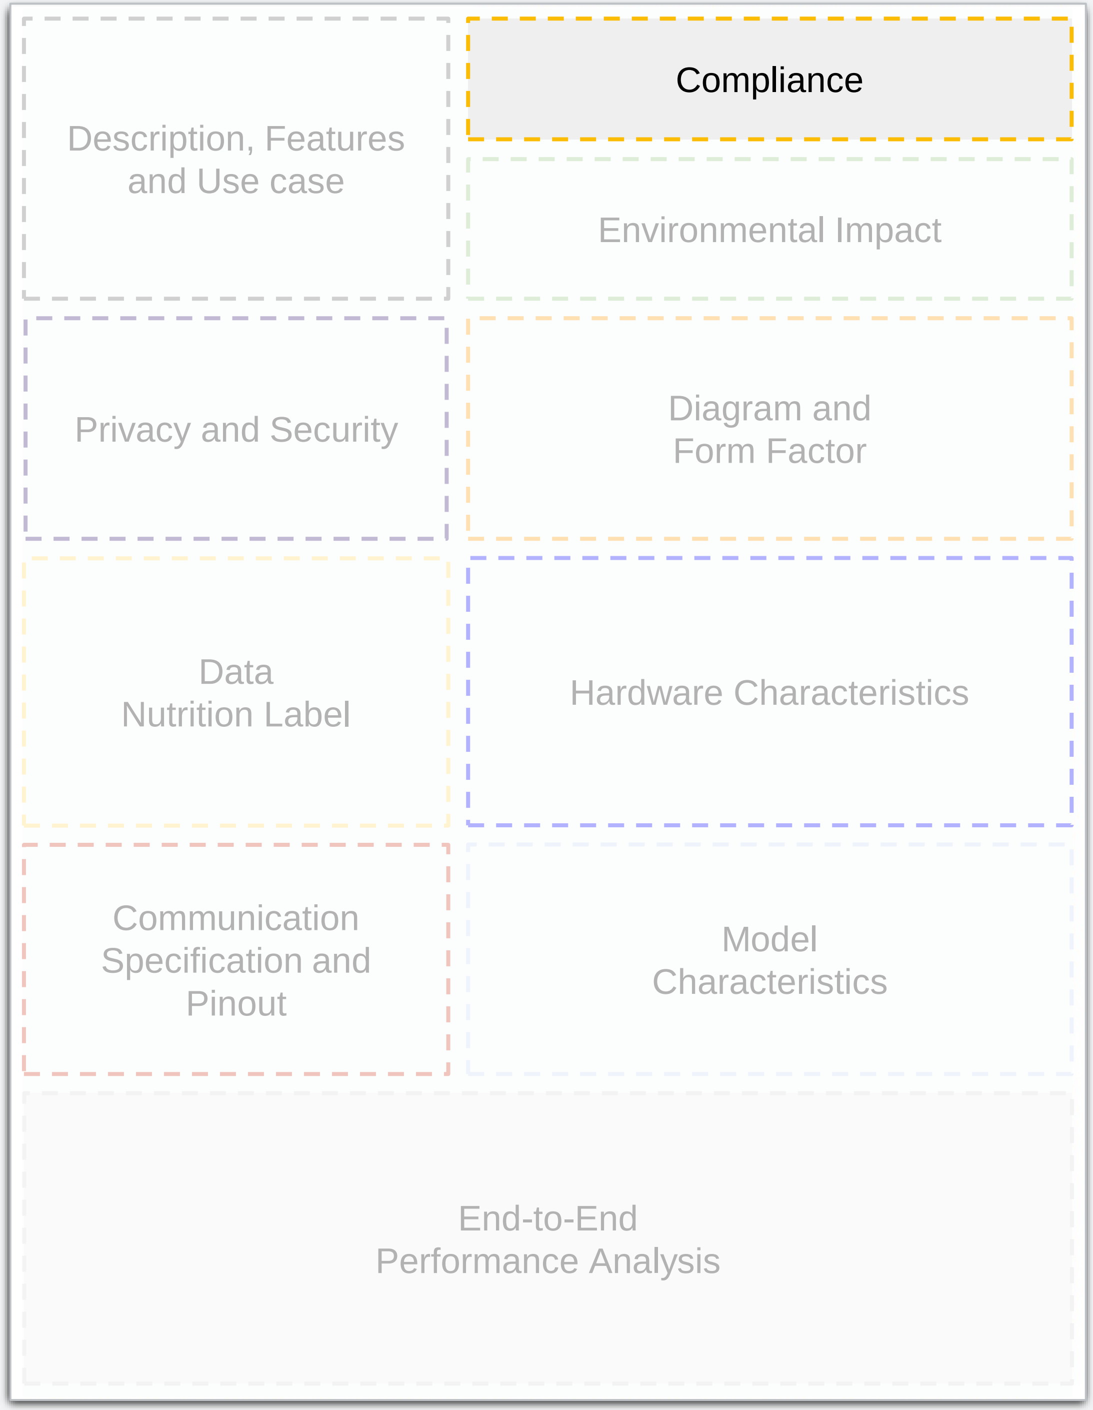

Hardware Details

Diagram and Form Factor

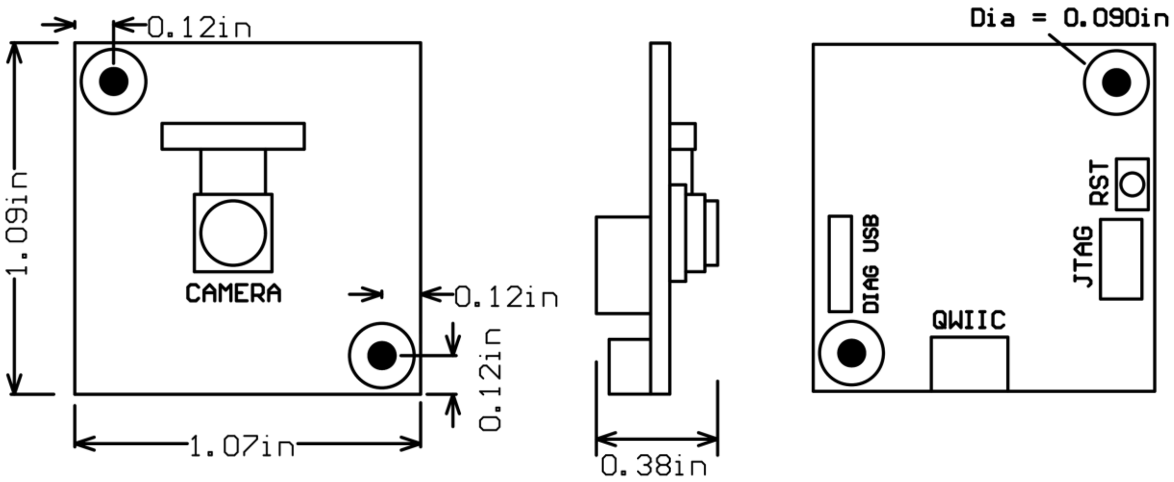

Communication Specification

I2C/Qwiic mode	Conforms with SparkFun Qwiic electrical/mechanical specifications. https://www.sparkfun.com/qwiic
Max cable length	1 m
Max data rate	100 kb/s
Module Orientation	Red arrow on sticker points up.
GPIO mode	SCL/SDA lines can be customized to make programmable flag lines (I_{out} max = 12 mA)
Diagnostic LED	Default behavior of green LED on board: illuminates for one second on power-up, then illuminates when person detected.
Data Transfer and Format	Single byte: number from 0-255 representing confidence score
I2C Address	0x22

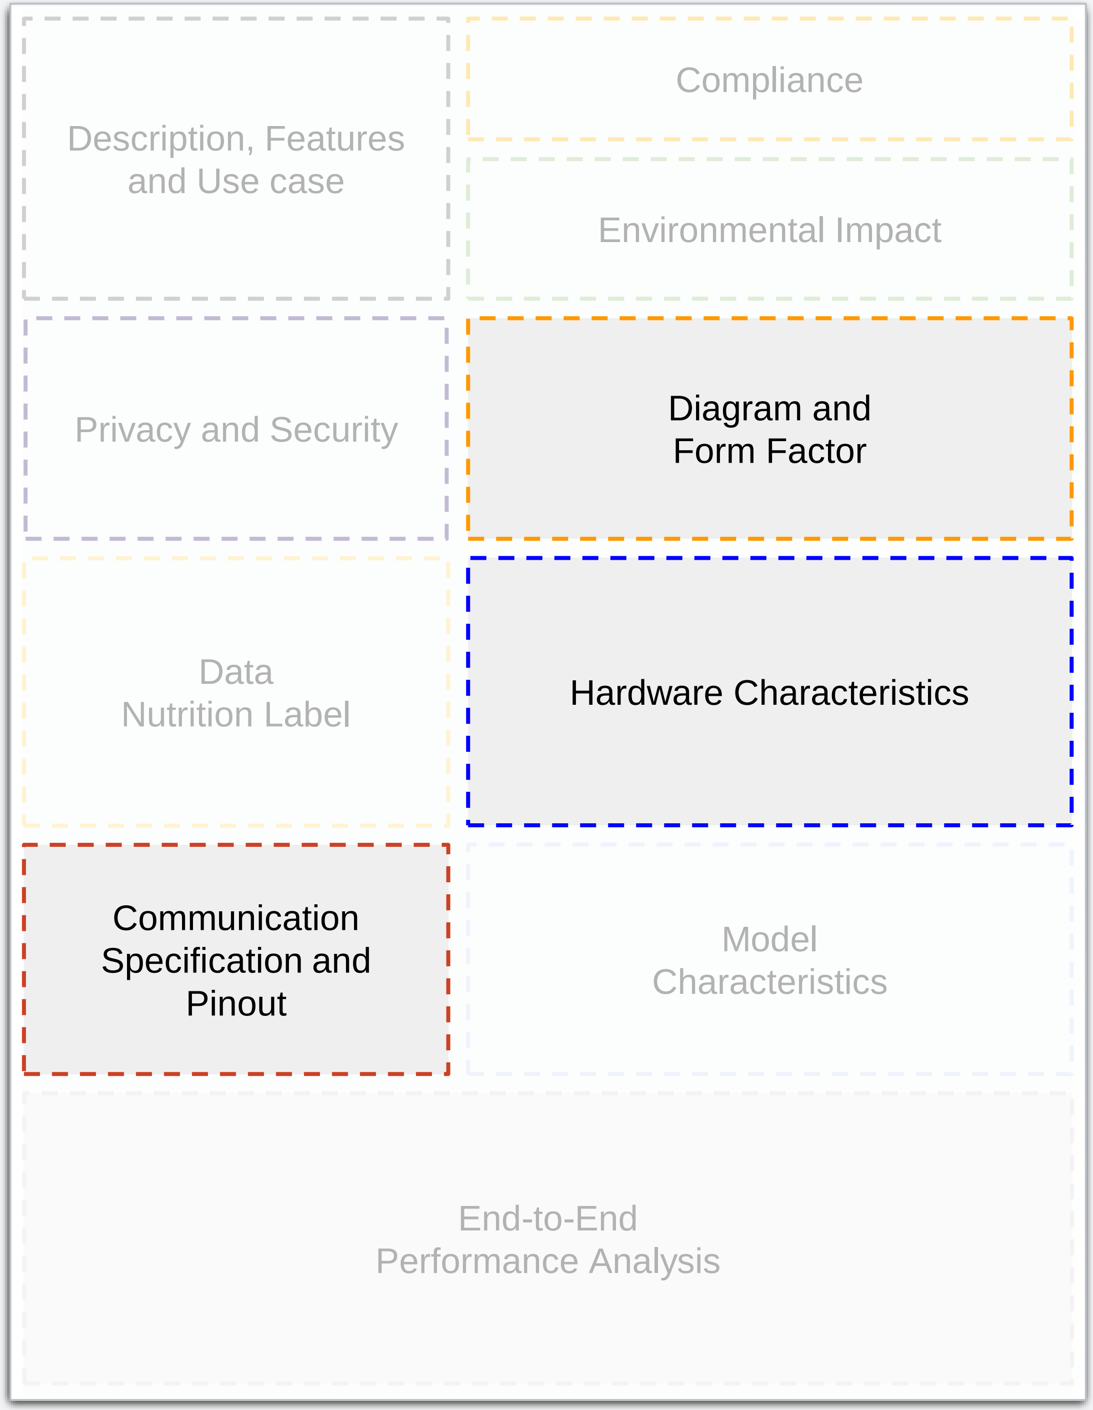

Environmental Impact

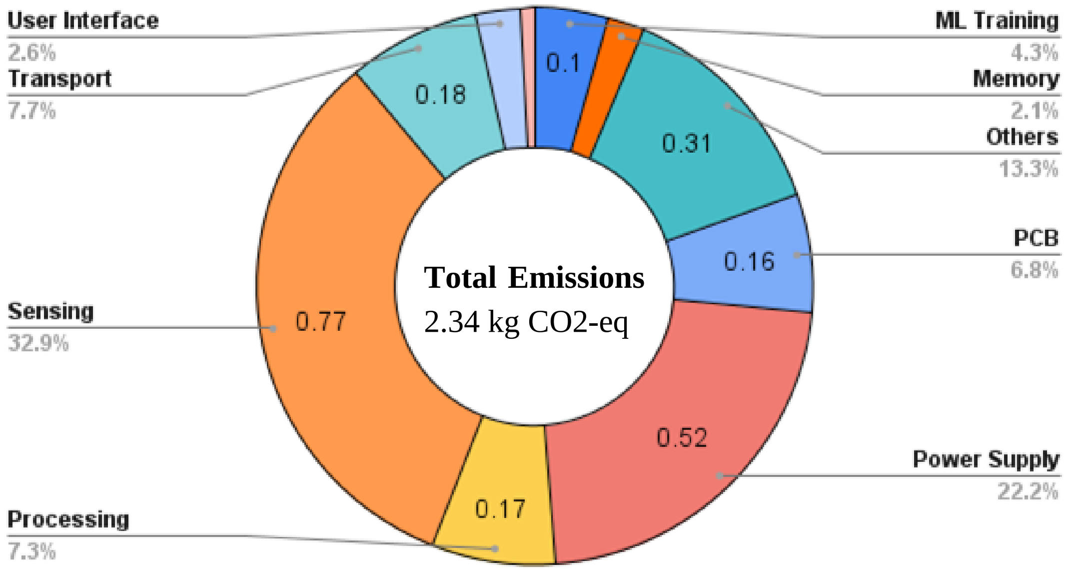

- Need to consider **environmentally-relevant metrics** such as carbon emissions, water usage, and eutrophication potential.
- Carbon emissions have two relevant sources: **operational energy consumption** and **hardware manufacturing and infrastructure**.
- Majority of emissions from our person detection sensor are associated with the **embodied footprint**. Data estimated using our TinyML sustainability calculator.

(Gupta et al., 2022; Prakash et al., 2022)

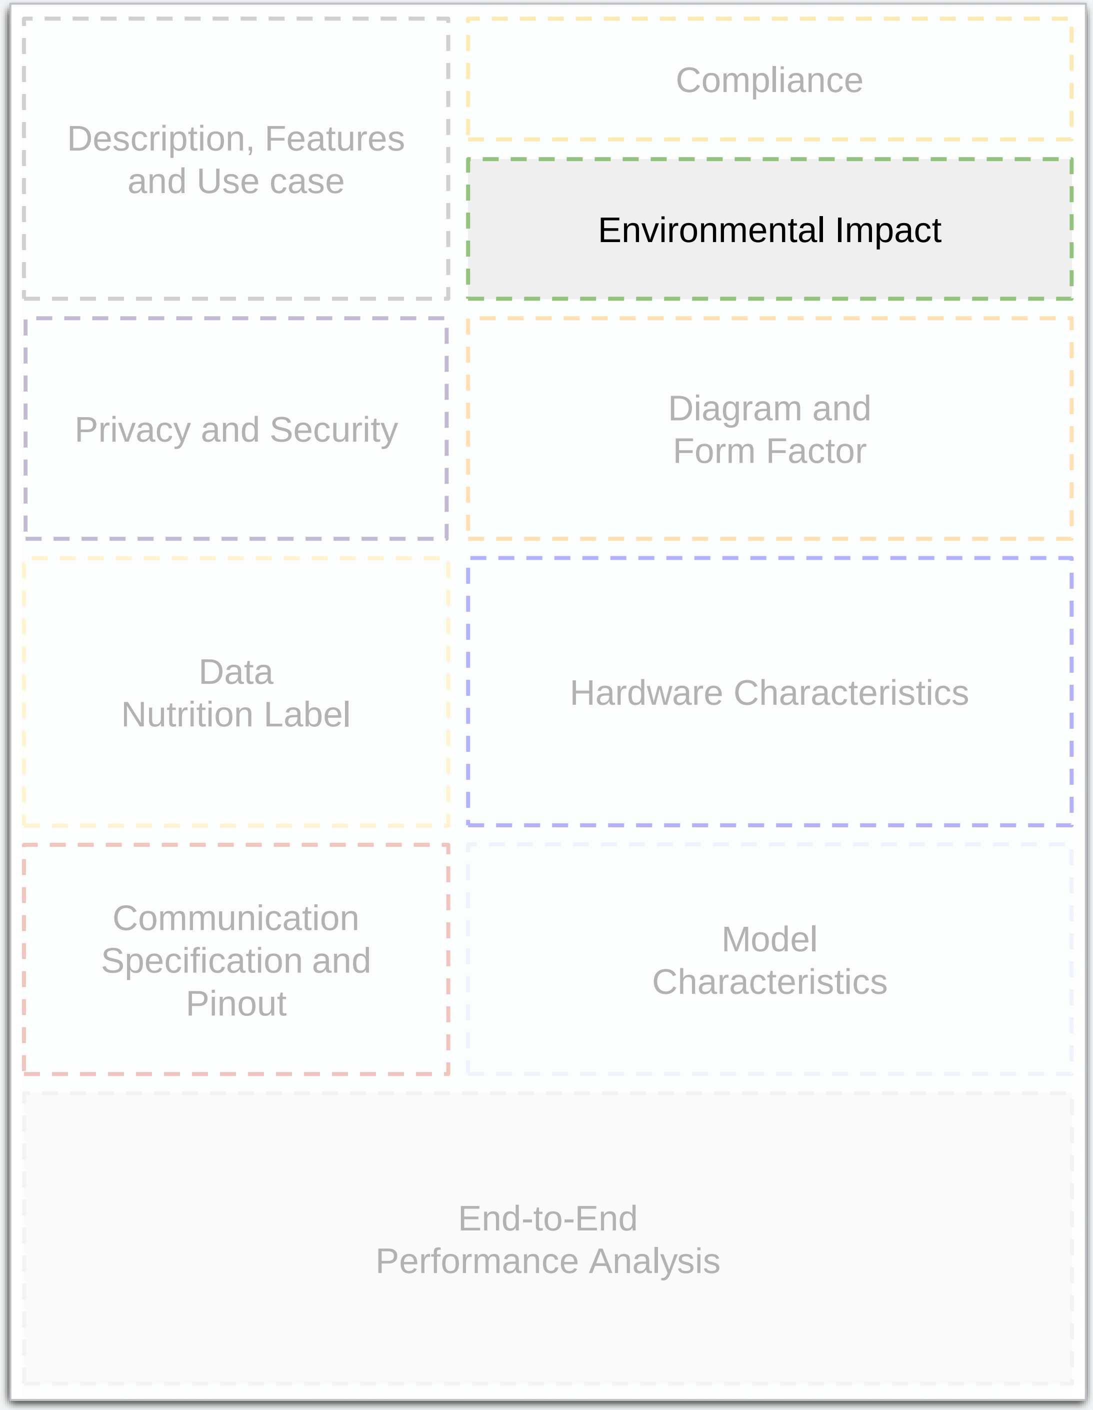

Datasheet Overview

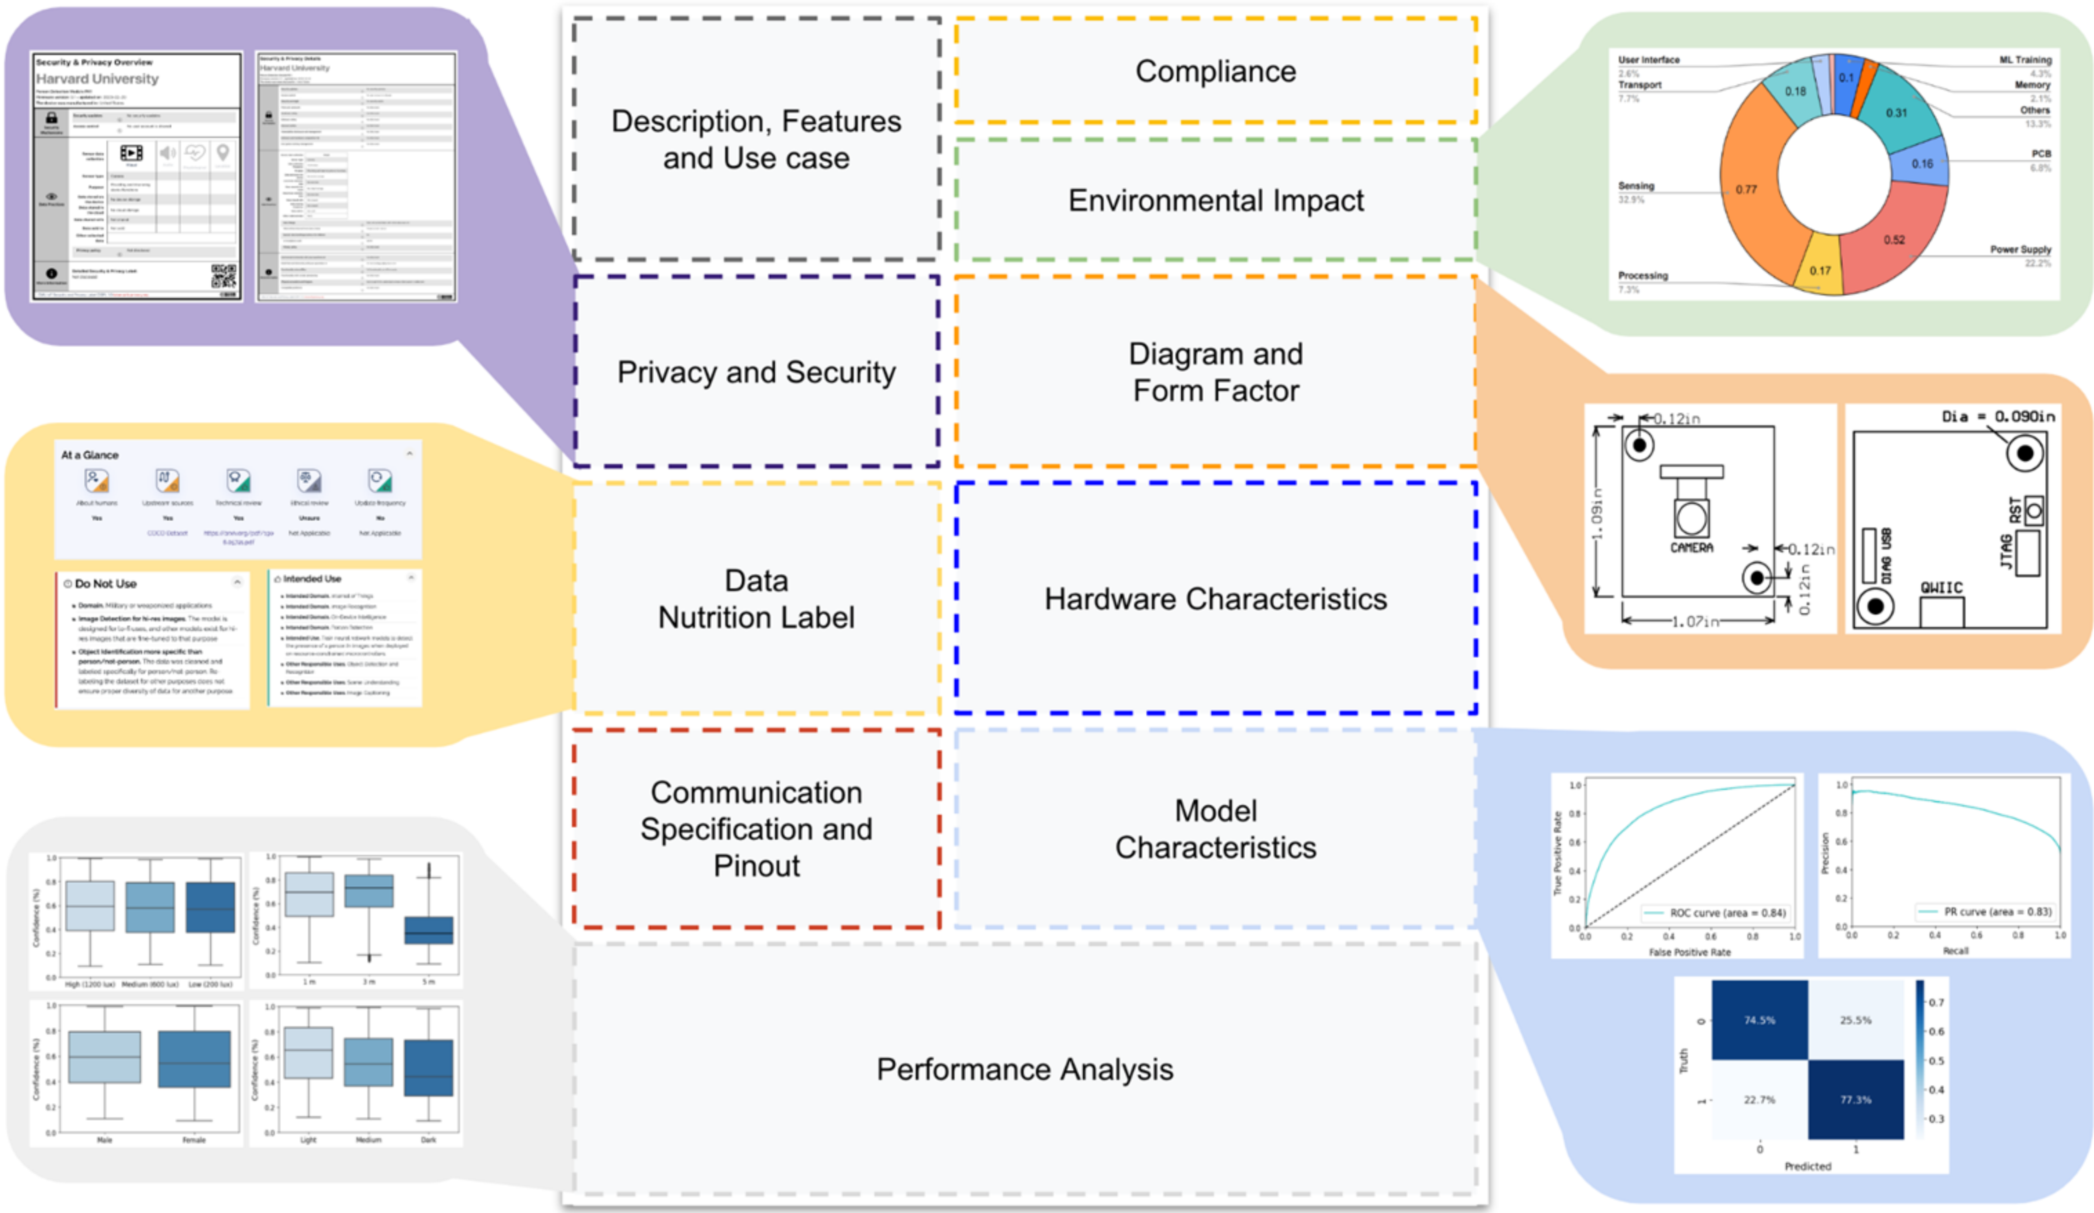

ML Sensors - Guiding Set of Principles

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

mlsensors.org

MLSensors
Machine Learning Sensors

Home | Whitepaper | Datasheets | GitHub | Email | Team

Machine Learning Sensors

An ML sensor is a self-contained system that utilizes on-device machine learning to extract useful information by observing some complex set of phenomena in the physical world and reports it through a simple interface to a wider system.

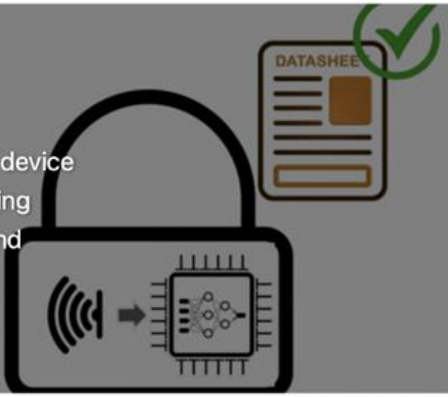

Machine learning sensors represent a paradigm shift for the future of embedded machine learning applications. Current instantiations of embedded ML suffer from complex integration, lack of modularity, and privacy and security concerns from data movement. ML sensors provide a more data-centric paradigm for embedding sensor intelligence on edge devices to combat these challenges.

Our vision for "sensor 2.0" entails segregating sensor input data and ML processing from the wider system at the hardware level and providing a thin interface that mimics traditional sensors in functionality. This separation leads to a modular and easy-to-use ML sensor device. ML sensors increase privacy and accuracy while making it easier for system builders to integrate ML into their products as a simple component.

To learn more about our approach, check out our ML sensor whitepaper, as well as our recent work on datasheets for machine learning sensors.

Challenges

Interface

What universal interface is needed for ML

Standards

What standards need to be in place for ML

Ethics

What ethical considerations are needed for

Recap of ML Sensors

1. We need to **raise the level of abstraction** to enable ease of use for scalable deployment of ML sensors; not everyone should be required to be a developer or an engineer to leverage ML sensors into their ecosystem.

2. The ML sensor's **design should be inherently data-centric** and defined by its input-output behavior instead of exposing the underlying hardware and software mechanisms that support ML model execution.

1. An ML sensor's **implementation must be clean and complexity-free**. Features such as reusability, software updates, and networking must be thought through to ensure data privacy and secure execution.

2. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

3. We as a community should aim to **foster an open ML sensors ecosystem by maximizing data, model, and hardware transparency** where possible, without necessarily relinquishing any claim to intellectual property.

Harvard Radcliffe Institute

Call to Action

Radcliffe exploratory seminar to determine:

- What ethical considerations are necessary when developing ML sensors?
- What compliance standards must be met by ML sensor developer and manufacturers?
- How should ML sensors interface with existing systems?

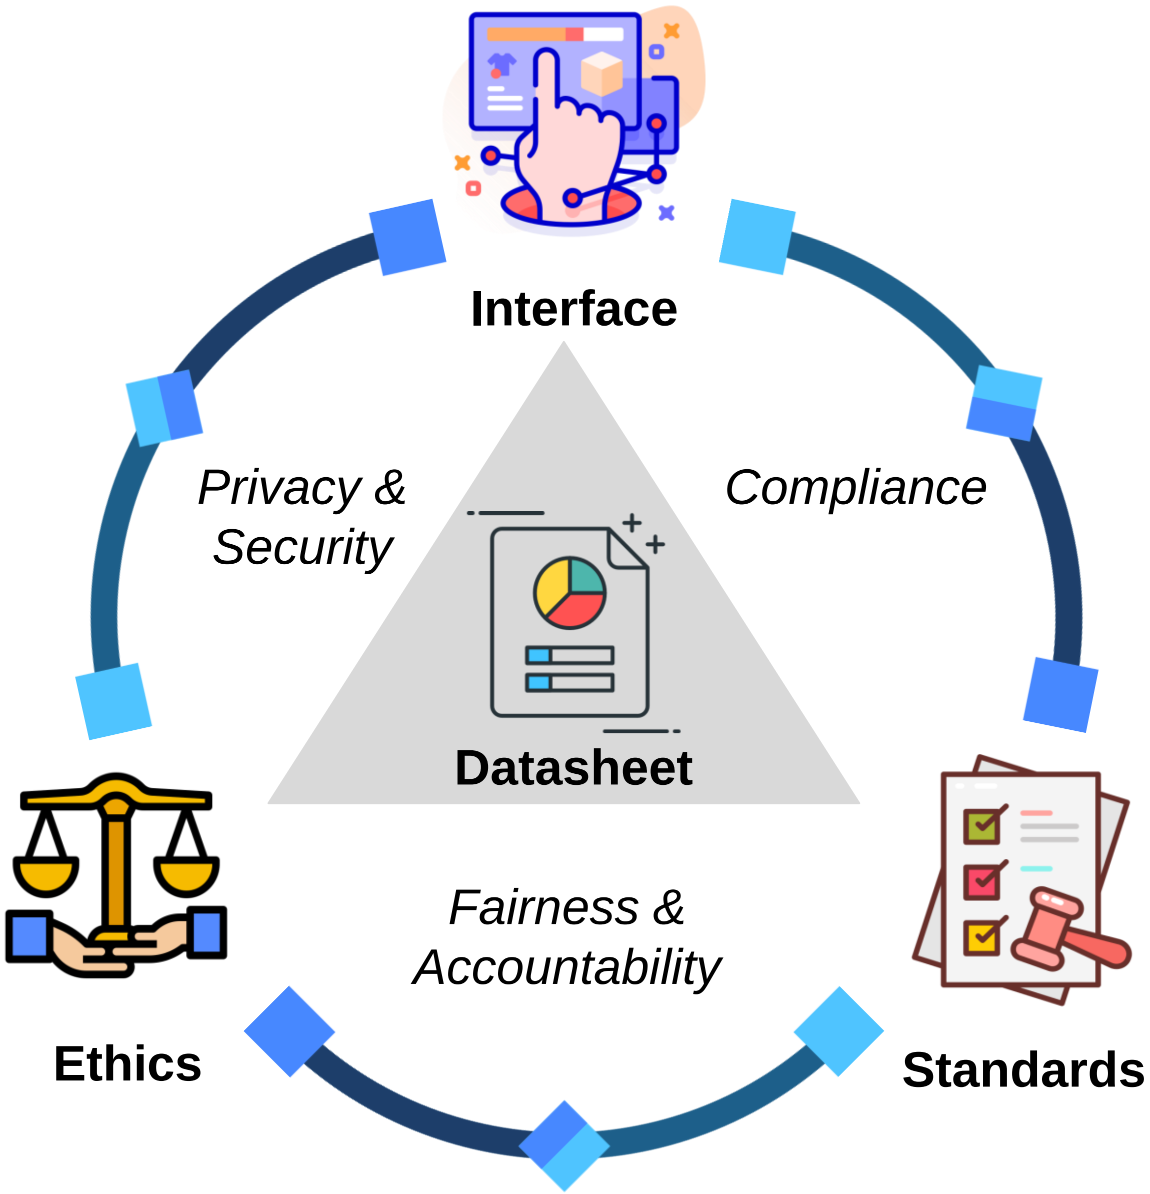

mlsensors.org

https://github.com/harvard-edge/ML-Sensors

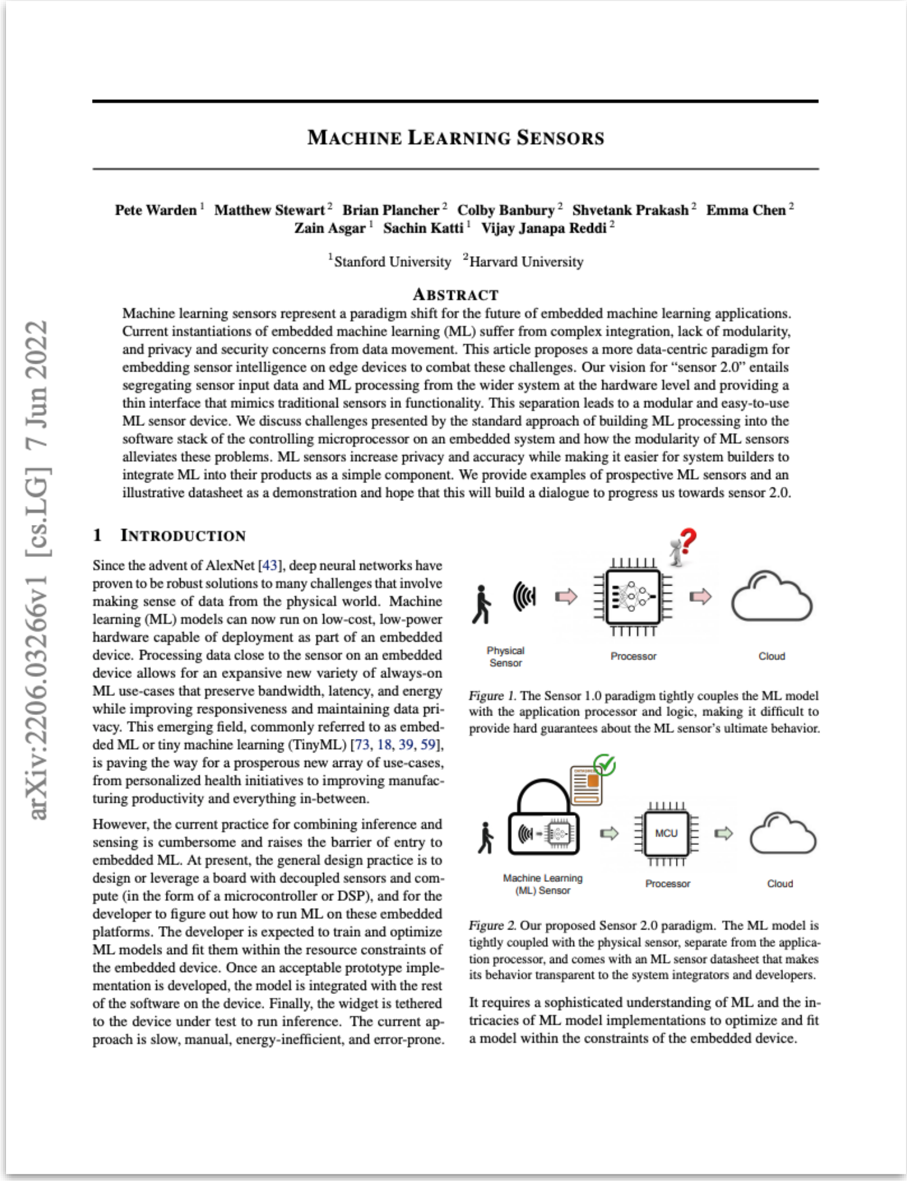

arXiv:2206.03266v1 [cs.LG] 7 Jun 2022

MACHINE LEARNING SENSORS

Pete Warden [1] Matthew Stewart [2] Brian Plancher [2] Colby Banbury [2] Shvetank Prakash [2] Emma Chen [2]
Zain Asgar [1] Sachin Katti [1] Vijay Janapa Reddi [2]

[1]Stanford University [2]Harvard University

ABSTRACT

Machine learning sensors represent a paradigm shift for the future of embedded machine learning applications. Current instantiations of embedded machine learning (ML) suffer from complex integration, lack of modularity, and privacy and security concerns from data movement. This article proposes a more data-centric paradigm for embedding sensor intelligence on edge devices to combat these challenges. Our vision for "sensor 2.0" entails segregating sensor input data and ML processing from the wider system at the hardware level and providing a thin interface that mimics traditional sensors in functionality. This separation leads to a modular and easy-to-use ML sensor device. We discuss challenges presented by the standard approach of building ML processing into the software stack of the controlling microprocessor on an embedded system and how the modularity of ML sensors alleviates these problems. ML sensors increase privacy and accuracy while making it easier for system builders to integrate ML into their products as a simple component. We provide examples of prospective ML sensors and an illustrative datasheet as a demonstration and hope that this will build a dialogue to progress us towards sensor 2.0.

1 INTRODUCTION

Since the advent of AlexNet [43], deep neural networks have proven to be robust solutions to many challenges that involve making sense of data from the physical world. Machine learning (ML) models can now run on low-cost, low-power hardware capable of deployment as part of an embedded device. Processing data close to the sensor on an embedded device allows for an expansive new variety of always-on ML use-cases that preserve bandwidth, latency, and energy while improving responsiveness and maintaining data privacy. This emerging field, commonly referred to as embedded ML or tiny machine learning (TinyML) [73, 18, 39, 59], is paving the way for a prosperous new array of use-cases, from personalized health initiatives to improving manufacturing productivity and everything in-between.

However, the current practice for combining inference and sensing is cumbersome and raises the barrier of entry to embedded ML. At present, the general design practice is to design or leverage a board with decoupled sensors and compute (in the form of a microcontroller or DSP), and for the developer to figure out how to run ML on these embedded platforms. The developer is expected to train and optimize ML models and fit them within the resource constraints of the embedded device. Once an acceptable prototype implementation is developed, the model is integrated with the rest of the software on the device. Finally, the widget is tethered to the device under test to run inference. The current approach is slow, manual, energy-inefficient, and error-prone.

Physical Sensor
Processor
Cloud

Figure 1. The Sensor 1.0 paradigm tightly couples the ML model with the application processor and logic, making it difficult to provide hard guarantees about the ML sensor's ultimate behavior.

Machine Learning (ML) Sensor
MCU
Processor
Cloud

Figure 2. Our proposed Sensor 2.0 paradigm. The ML model is tightly coupled with the physical sensor, separate from the application processor, and comes with an ML sensor datasheet that makes its behavior transparent to the system integrators and developers.

It requires a sophisticated understanding of ML and the intricacies of ML model implementations to optimize and fit a model within the constraints of the embedded device.

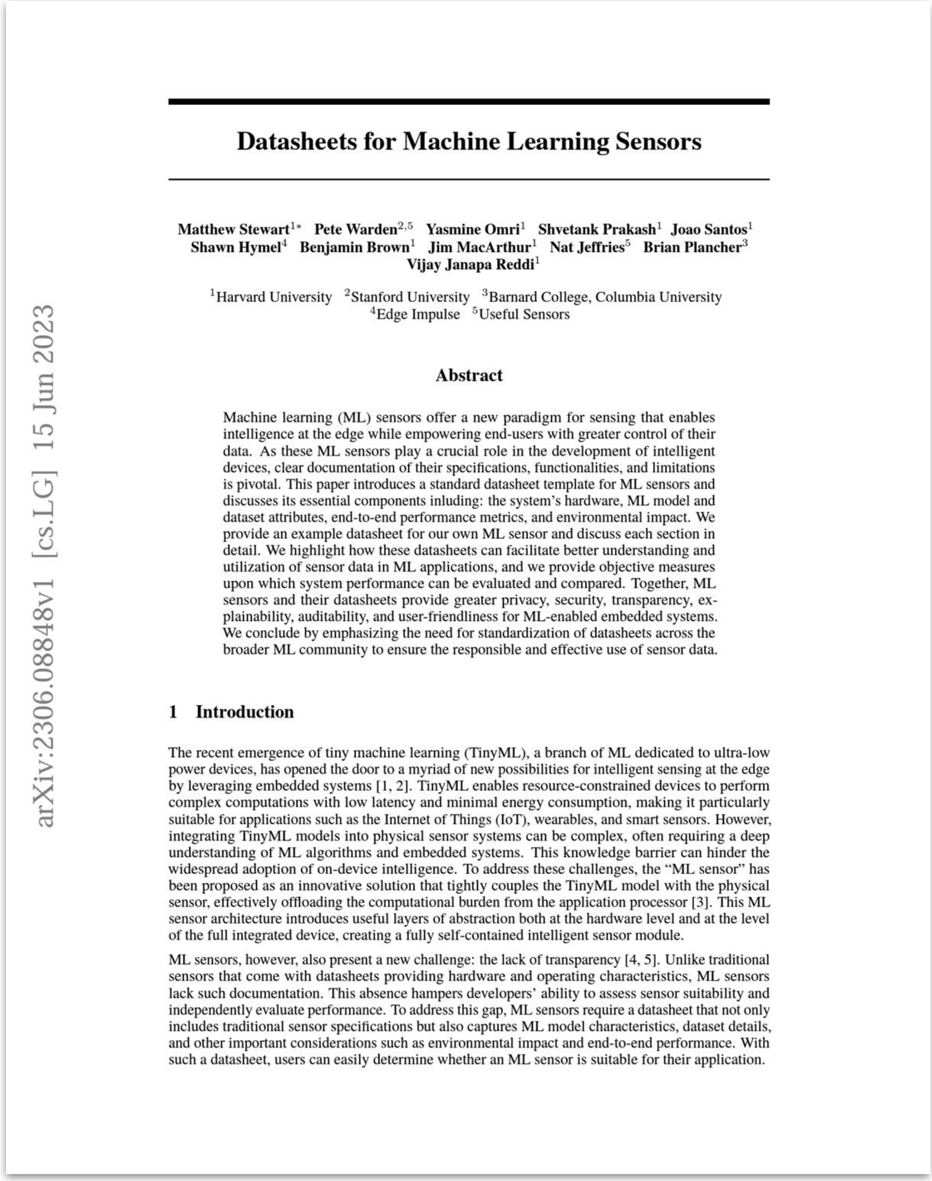

arXiv:2306.08848v1 [cs.LG] 15 Jun 2023

Datasheets for Machine Learning Sensors

Matthew Stewart[1*] Pete Warden[2,5] Yasmine Omri[1] Shvetank Prakash[1] Joao Santos[1]
Shawn Hymel[4] Benjamin Brown[1] Jim MacArthur[1] Nat Jeffries[5] Brian Plancher[3]
Vijay Janapa Reddi[1]

[1]Harvard University [2]Stanford University [3]Barnard College, Columbia University
[4]Edge Impulse [5]Useful Sensors

Abstract

Machine learning (ML) sensors offer a new paradigm for sensing that enables intelligence at the edge while empowering end-users with greater control of their data. As these ML sensors play a crucial role in the development of intelligent devices, clear documentation of their specifications, functionalities, and limitations is pivotal. This paper introduces a standard datasheet template for ML sensors and discusses its essential components inluding: the system's hardware, ML model and dataset attributes, end-to-end performance metrics, and environmental impact. We provide an example datasheet for our own ML sensor and discuss each section in detail. We highlight how these datasheets can facilitate better understanding and utilization of sensor data in ML applications, and we provide objective measures upon which system performance can be evaluated and compared. Together, ML sensors and their datasheets provide greater privacy, security, transparency, explainability, auditability, and user-friendliness for ML-enabled embedded systems. We conclude by emphasizing the need for standardization of datasheets across the broader ML community to ensure the responsible and effective use of sensor data.

1 Introduction

The recent emergence of tiny machine learning (TinyML), a branch of ML dedicated to ultra-low power devices, has opened the door to a myriad of new possibilities for intelligent sensing at the edge by leveraging embedded systems [1, 2]. TinyML enables resource-constrained devices to perform complex computations with low latency and minimal energy consumption, making it particularly suitable for applications such as the Internet of Things (IoT), wearables, and smart sensors. However, integrating TinyML models into physical sensor systems can be complex, often requiring a deep understanding of ML algorithms and embedded systems. This knowledge barrier can hinder the widespread adoption of on-device intelligence. To address these challenges, the "ML sensor" has been proposed as an innovative solution that tightly couples the TinyML model with the physical sensor, effectively offloading the computational burden from the application processor [3]. This ML sensor architecture introduces useful layers of abstraction both at the hardware level and at the level of the full integrated device, creating a fully self-contained intelligent sensor module.

ML sensors, however, also present a new challenge: the lack of transparency [4, 5]. Unlike traditional sensors that come with datasheets providing hardware and operating characteristics, ML sensors lack such documentation. This absence hampers developers' ability to assess sensor suitability and independently evaluate performance. To address this gap, ML sensors require a datasheet that not only includes traditional sensor specifications but also captures ML model characteristics, dataset details, and other important considerations such as environmental impact and end-to-end performance. With such a datasheet, users can easily determine whether an ML sensor is suitable for their application.

Copyright Notice

www.tinyml.org

Copyright Notice

www.tinyml.org