SciTinyML – ICTP Workshop

Scientific Use of Machine Learning on Low Power Devices

Machine Learning Sensors

Acknowledgements: *Z. Asgar, C. Banbury, B. Brown, E. Chen, J. MacArthur, B. Plancher, S. Prakash, S. Katti, 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

Applications of Machine Learning

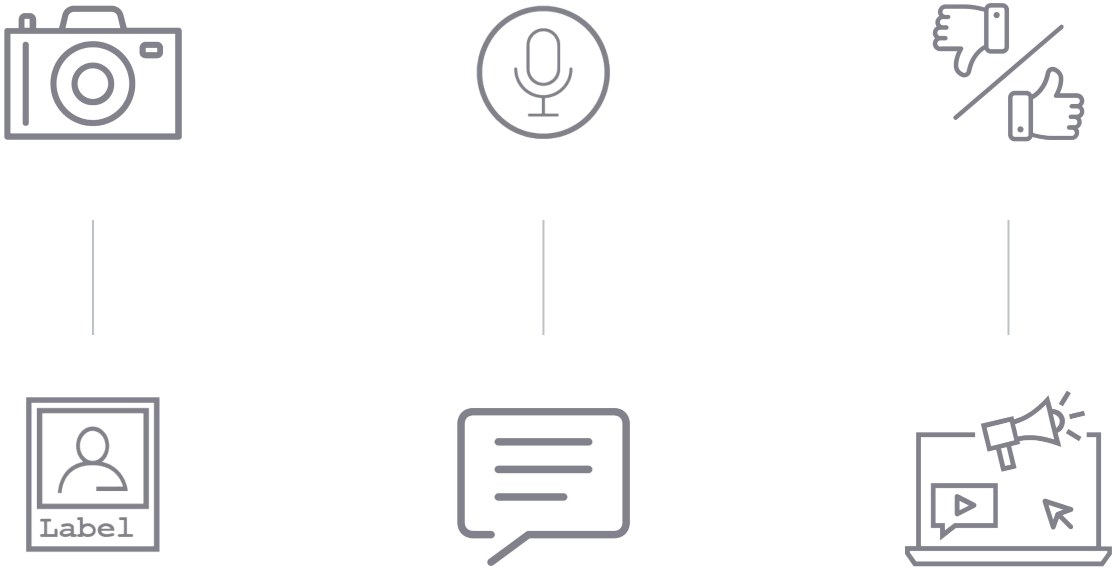

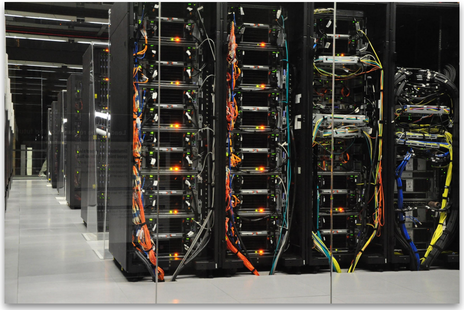


Dinner with Layla in 30 min

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

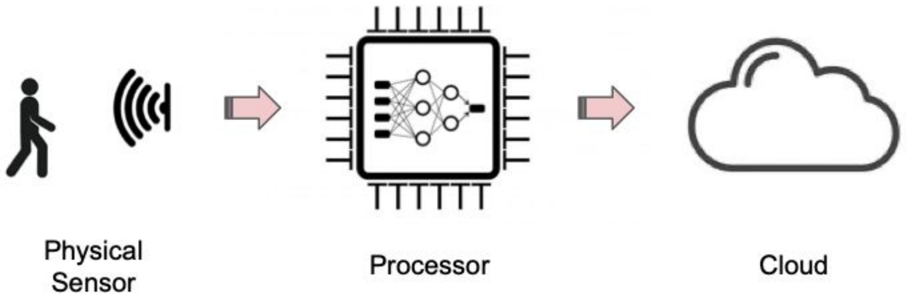

Sensor 1.0

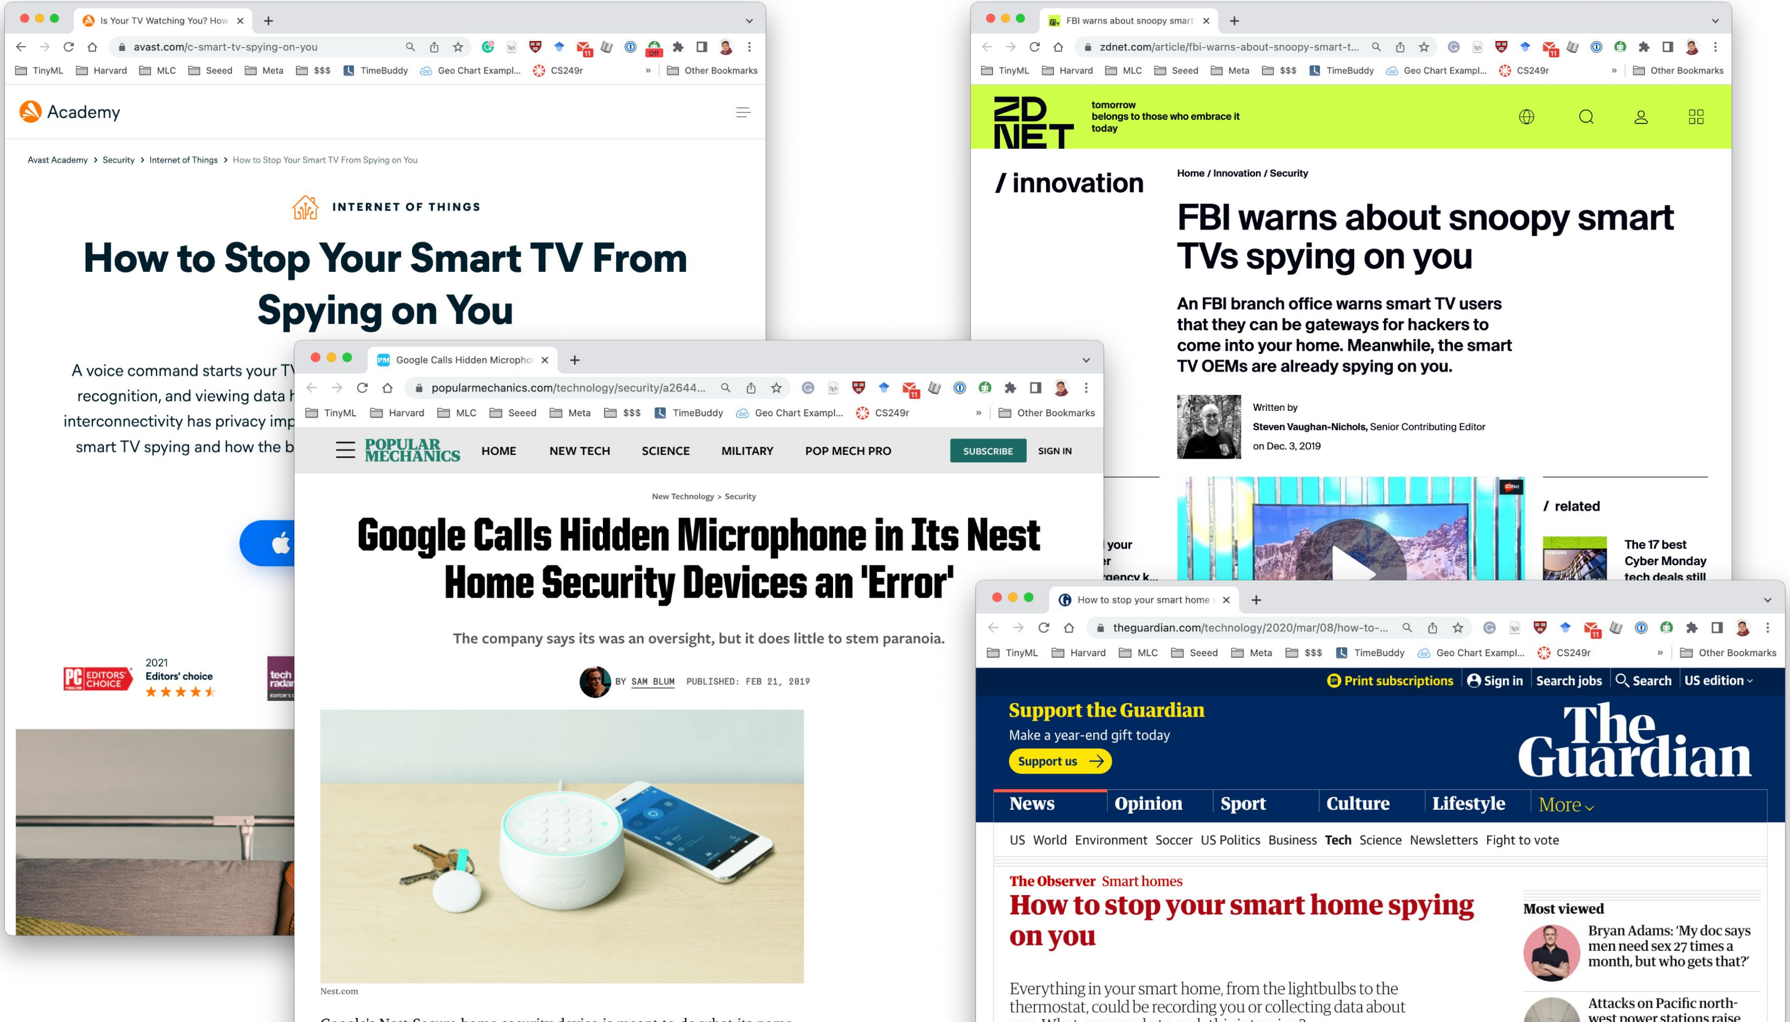
avast.com/c-smart-tv-spying-on-you
Academy
Avast Academy > Security > Internet of Things > How to Stop Your Smart TV From Spying on You
INTERNET OF THINGS
How to Stop Your Smart TV From Spying on You
A voice command starts your T
recognition, and viewing data h
interconnectivity has privacy imp
smart TV spying and how the b
2021
Editors' choice
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
The 17 best Cyber Monday tech deals still
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
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
Everything in your smart home, from the lightbulbs to the thermostat, could be recording you or collecting data about
Most viewed
Bryan Adams: 'My doc says men need sex 27 times a month, but who gets that?'
Attacks on Pacific north-west power stations raise

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

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

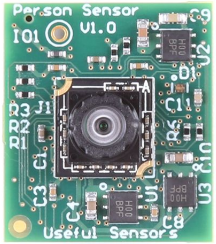

by Useful Sensors

Machine Learning Sensors

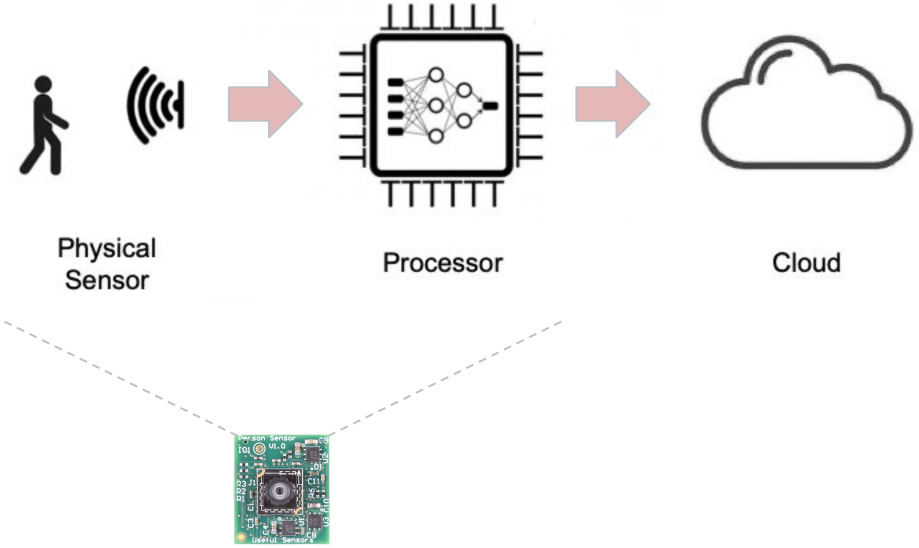

Machine Learning Sensors

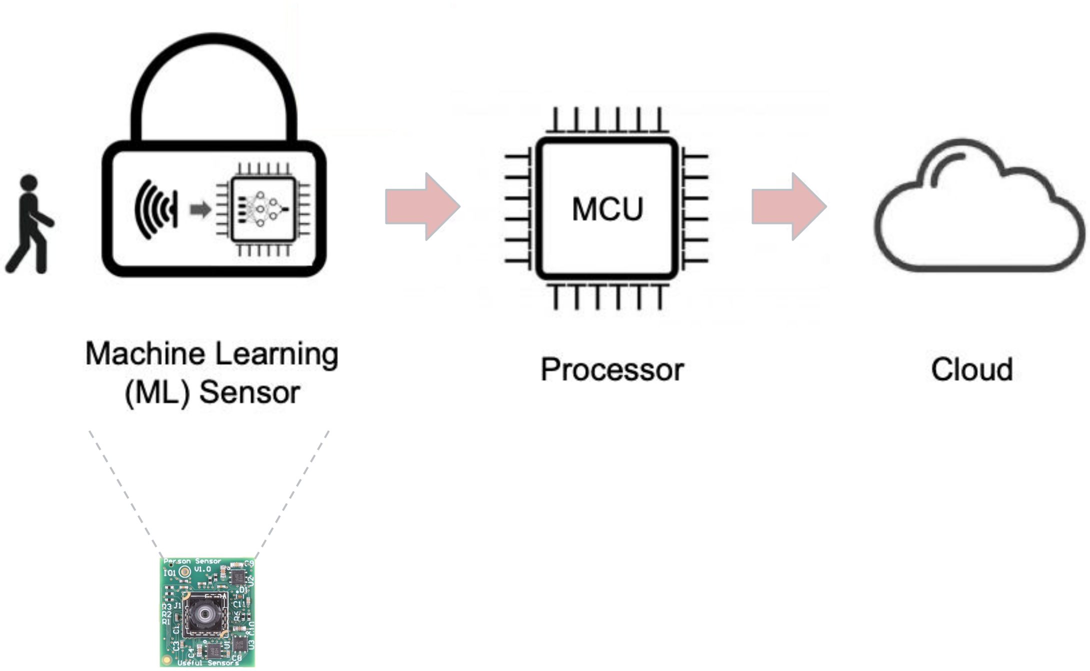

Sensor 2.0

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.

3. 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.

4. 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.

5. 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 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.

3. 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.

4. 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.

5. 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 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.

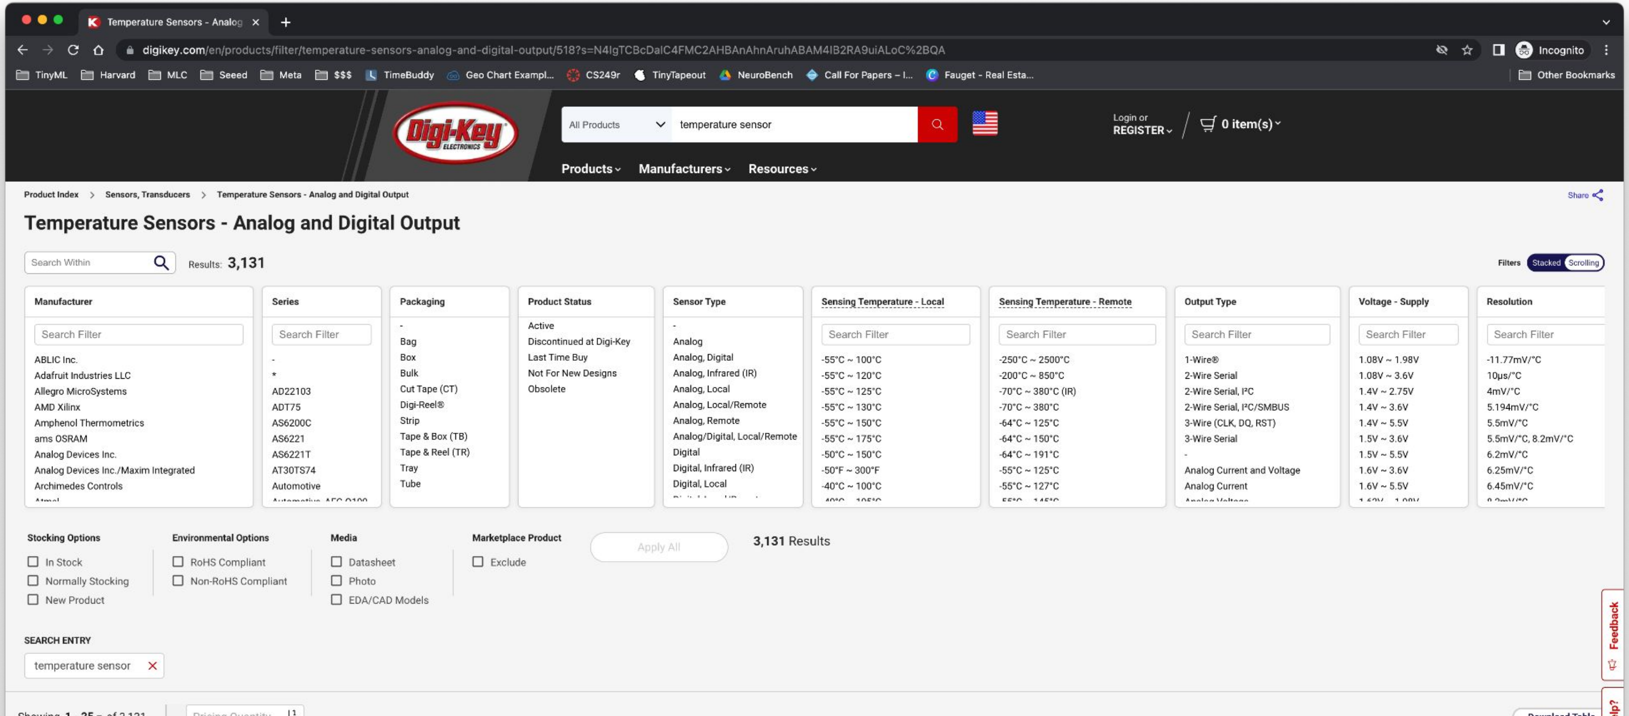

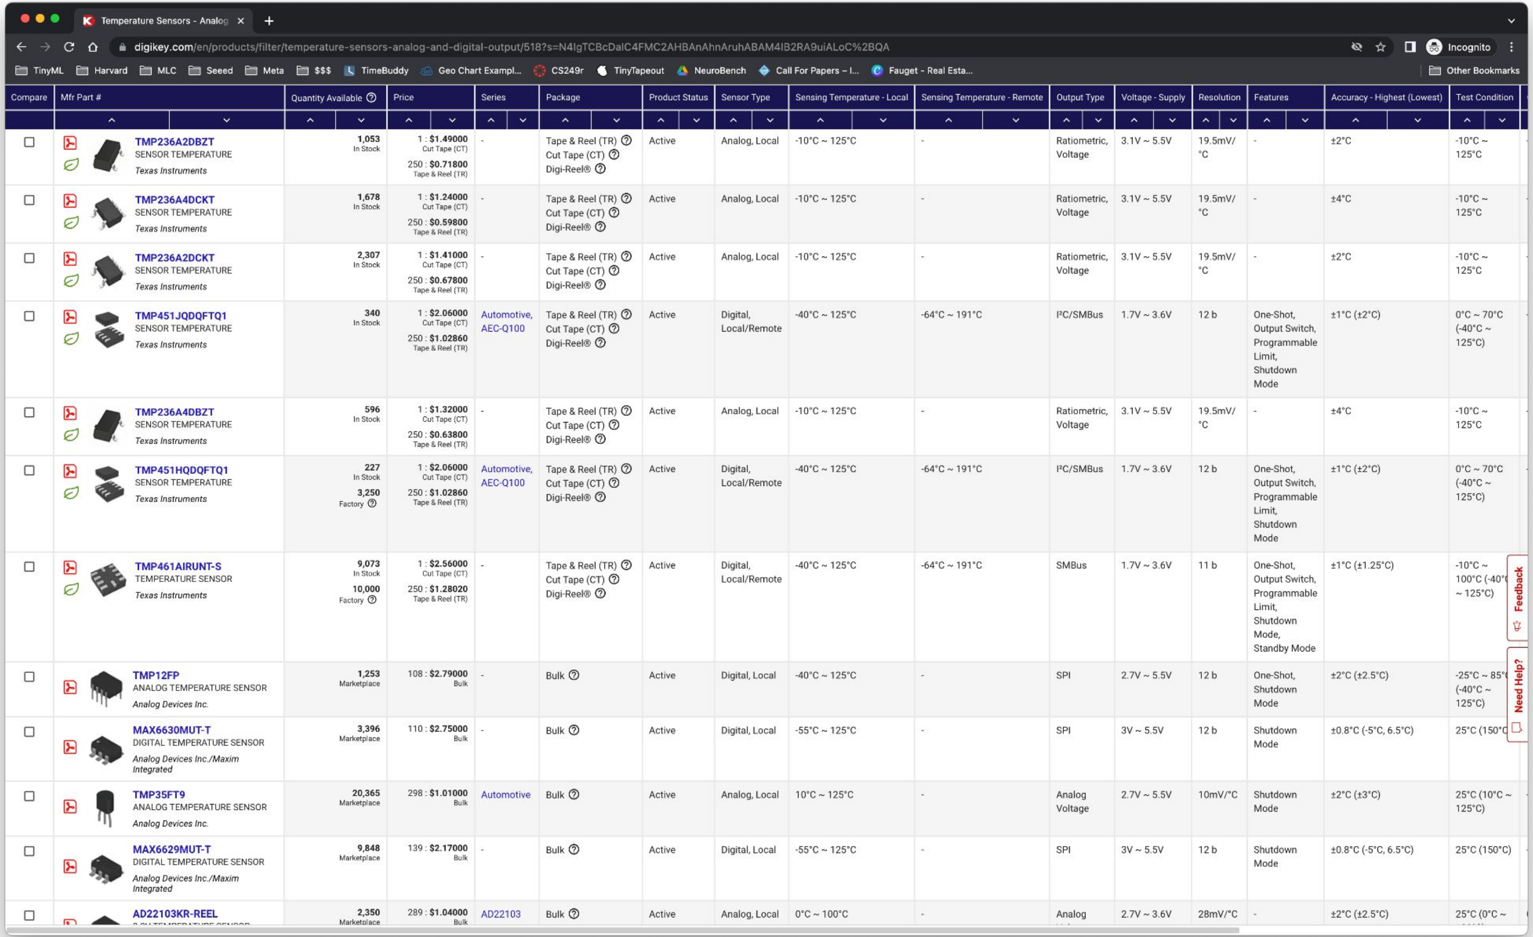

Compare	Mfr Part #	Quantity Available	Price	Series	Package	Product Status	Sensor Type
	TMP236A2DBZT SENSOR TEMPERATURE *Texas Instruments*	1,053 In Stock	1 : $1.49000 Cut Tape (CT) 250 : $0.71800 Tape & Reel (TR)	-	Tape & Reel (TR) Cut Tape (CT) Digi-Reel®	Active	Analog, Local
	TMP236A4DCKT SENSOR TEMPERATURE *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
	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
	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
	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
	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
	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
	TMP12FP ANALOG TEMPERATURE SENSOR *Analog Devices Inc.*	1,253 Marketplace	108 : $2.79000 Bulk	-	Bulk	Active	Digital, Local
	MAX6630MUT-T DIGITAL TEMPERATURE SENSOR *Analog Devices Inc./Maxim Integrated*	3,396 Marketplace	110 : $2.75000 Bulk	-	Bulk	Active	Digital, Local
	TMP35FT9 ANALOG TEMPERATURE SENSOR *Analog Devices Inc.*	20,365 Marketplace	298 : $1.01000 Bulk	Automotive	Bulk	Active	Analog, Local
	MAX6629MUT-T DIGITAL TEMPERATURE SENSOR *Analog Devices Inc./Maxim Integrated*	9,848 Marketplace	139 : $2.17000 Bulk	-	Bulk	Active	Digital, Local
	AD22103KR-REEL	2,350 Marketplace	289 : $1.04000 Bulk	AD22103	Bulk	Active	Analog, Local

Mfr Part #	Sensing Temperature - Local	Sensing Temperature - Remote	Output Type	Voltage - Supply	Resolution	Features	Accuracy - Highest (Lowest)	Test Condition
TMP236A2DBZT	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±2°C	-10°C ~ 125°C
TMP236A4DCKT	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±4°C	-10°C ~ 125°C
TMP236A2DCKT	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±2°C	-10°C ~ 125°C
TMP451JQDQFTQ1	-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	-10°C ~ 125°C	-	Ratiometric, Voltage	3.1V ~ 5.5V	19.5mV/°C	-	±4°C	-10°C ~ 125°C
TMP451HQDQFTQ1	-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	-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	-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	-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	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	-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	0°C ~ 100°C	-	Analog	2.7V ~ 3.6V	28mV/°C	-	±2°C (±2.5°C)	25°C (0°C ~

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.

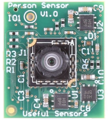

Person detector

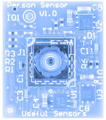

Gaze sensor

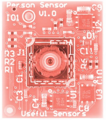

Voice command

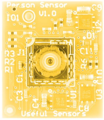

Text recognizer

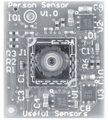

...

Digi-Key Electronics amazon.com Alibaba ...

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.

3. 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.

4. 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.

5. 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 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;
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.

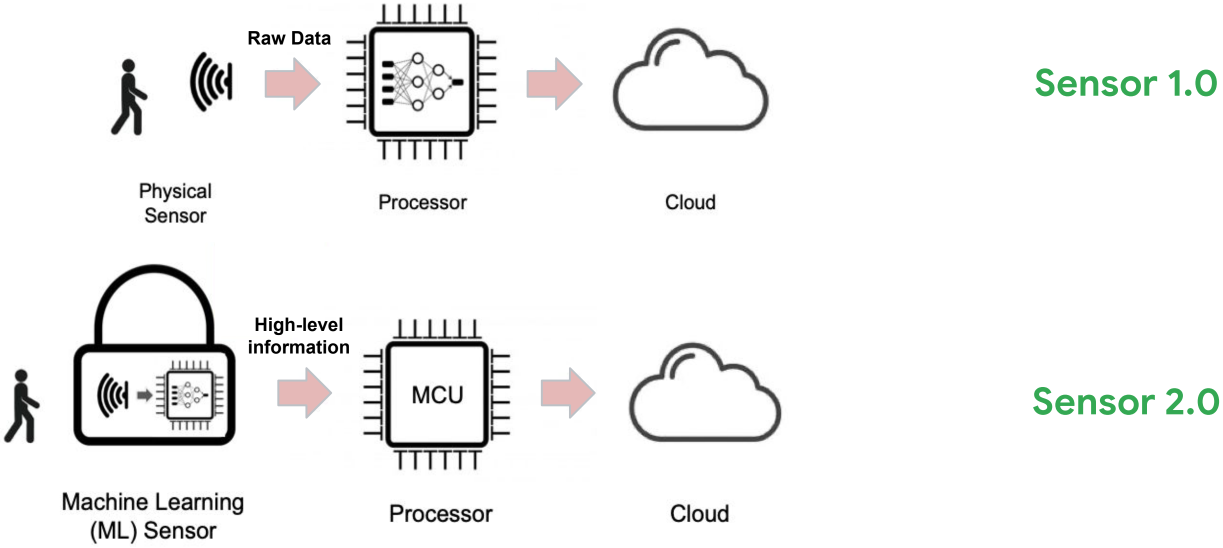

ML Sensors - Guiding Set of Principles

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.

3. 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.

4. 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.

5. 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 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.

3. 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.

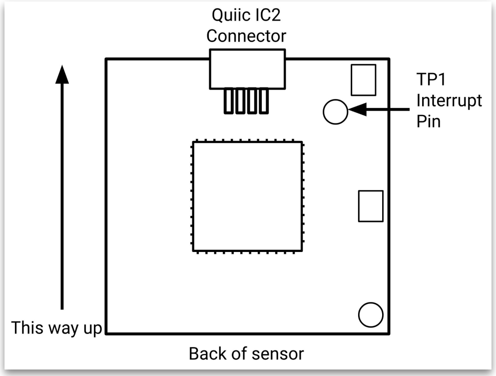

Source: https://github.com/usefulsensors/person_sensor_docs

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

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.

3. 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.

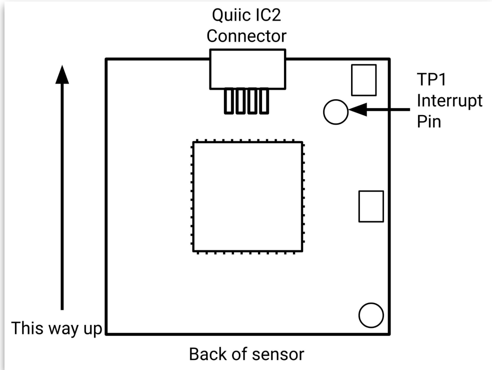

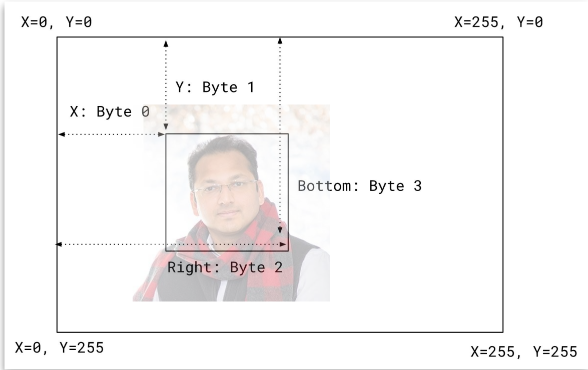

We need to define data formats to enable interoperability and exchange of ML sensors across manufacturers

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

ML 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.

3. 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.

4. 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.

5. 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 Sensors - Guiding Set of Principles

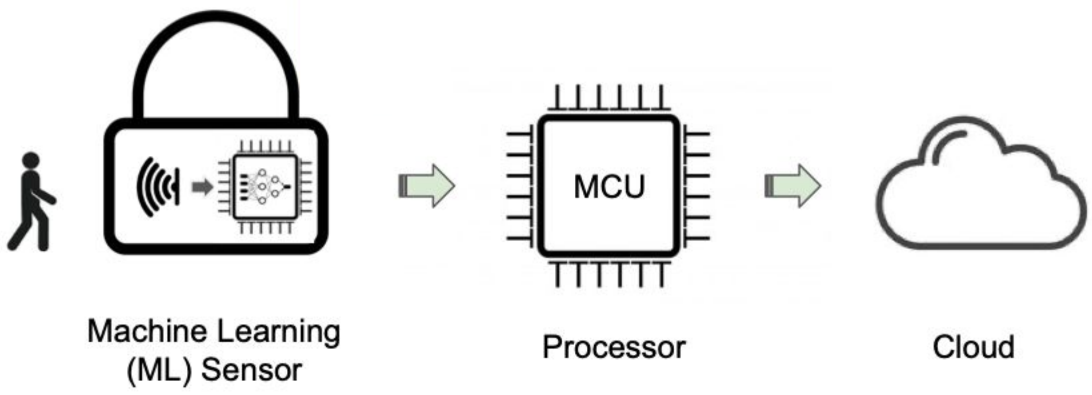

4. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

5. 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 Sensors - Guiding Set of Principles

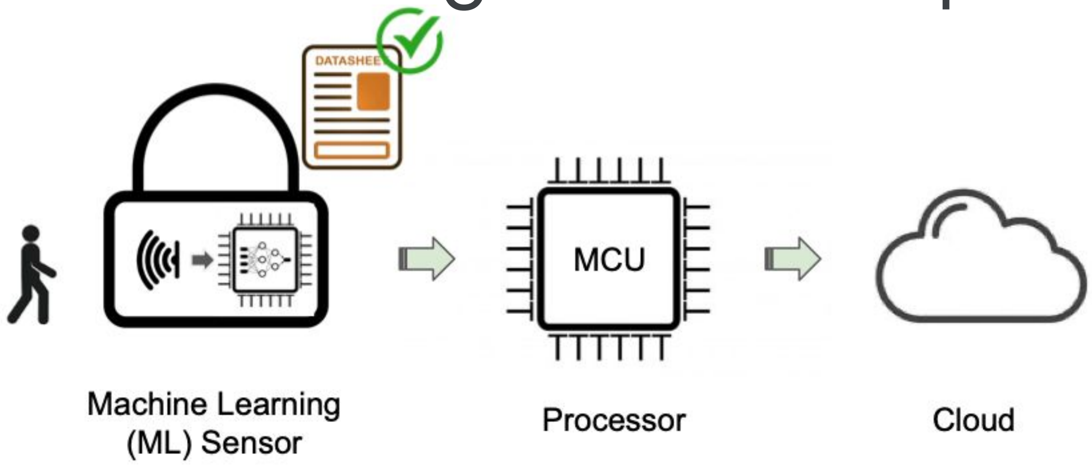

4. ML sensors **must be transparent, indicating in a publicly and freely accessible ML sensor datasheet** all the relevant information to supplement the traditional information available for hardware sensors.

5. 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.

E.g. ML Sensors Datasheets

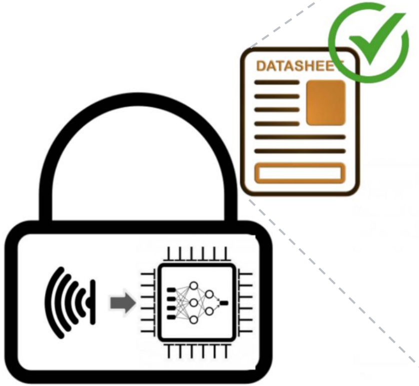

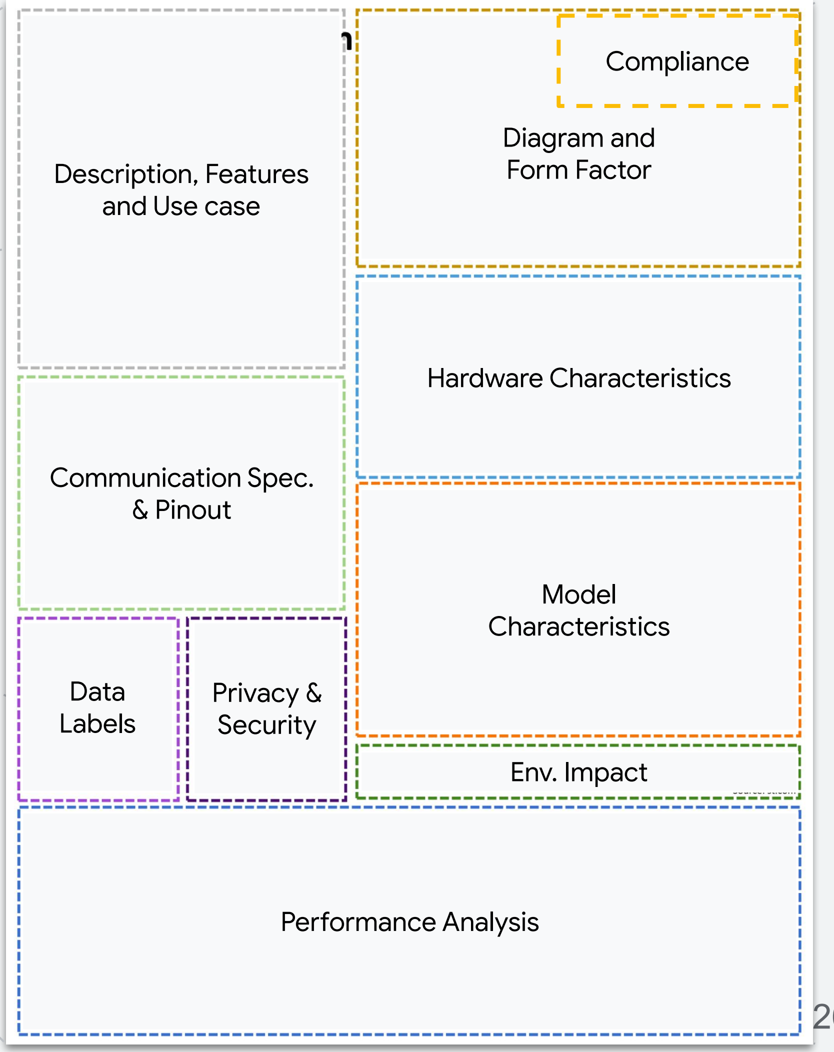

E.g. ML Sensors Datasheets

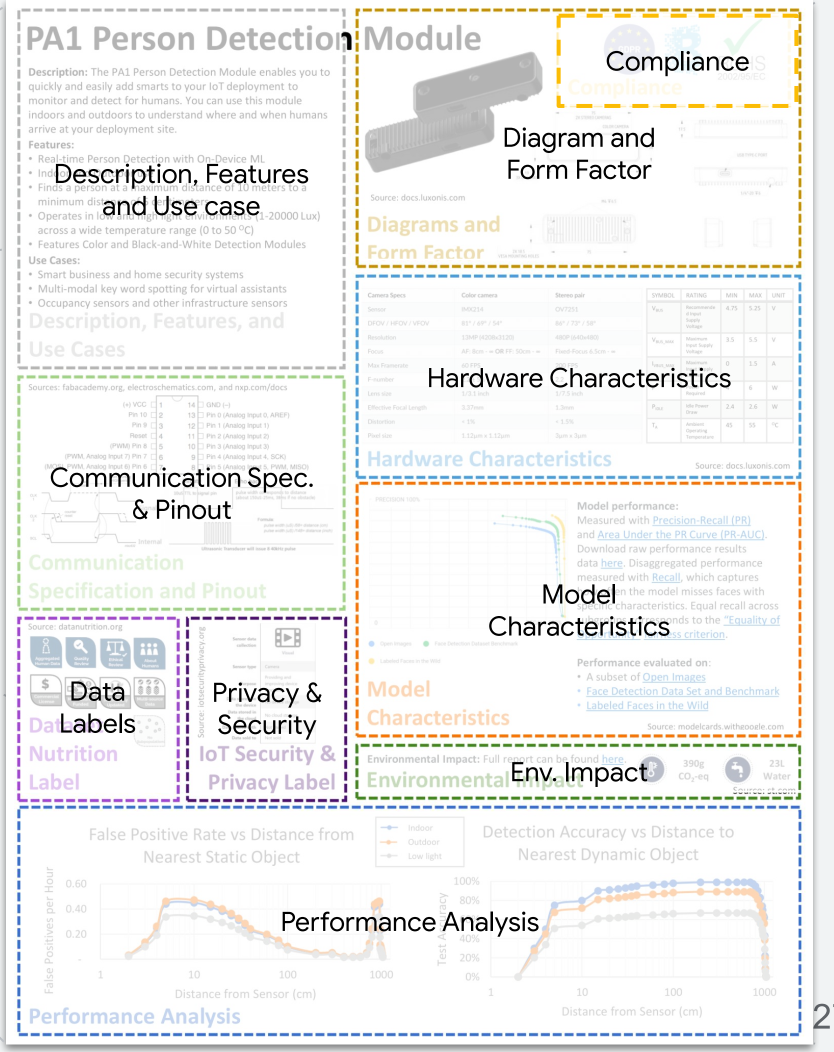

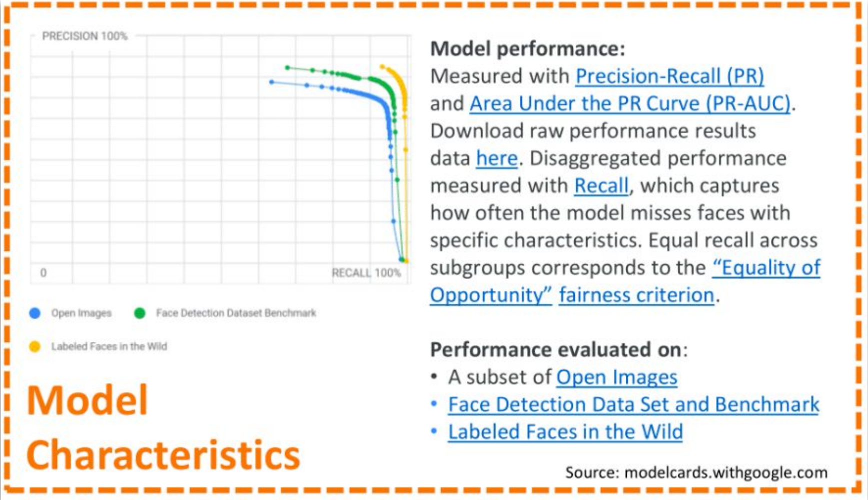

"*Model cards aim to provide a **concise, holistic picture of a machine learning model.** To start, a model card explains **what a model does**, its **intended audience**, and who **maintains it**. A model card also provides insight into the construction of the model, including **its architecture** and the **training data used**.*" – Google Cloud

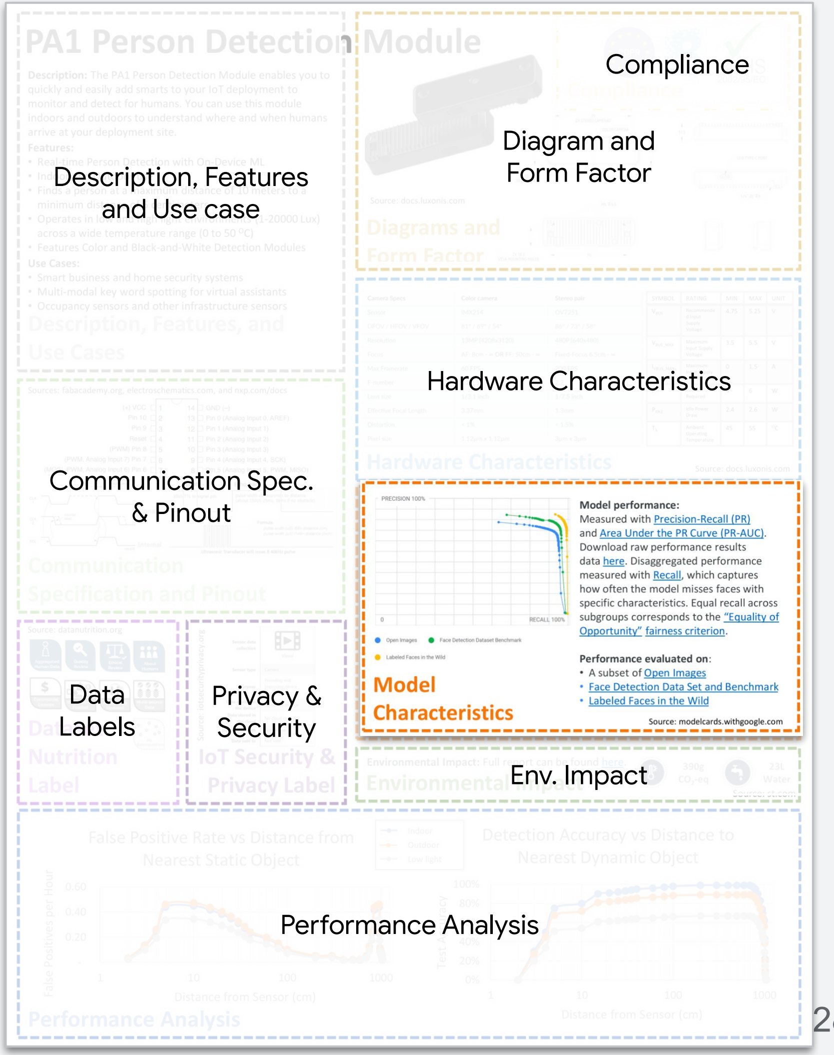

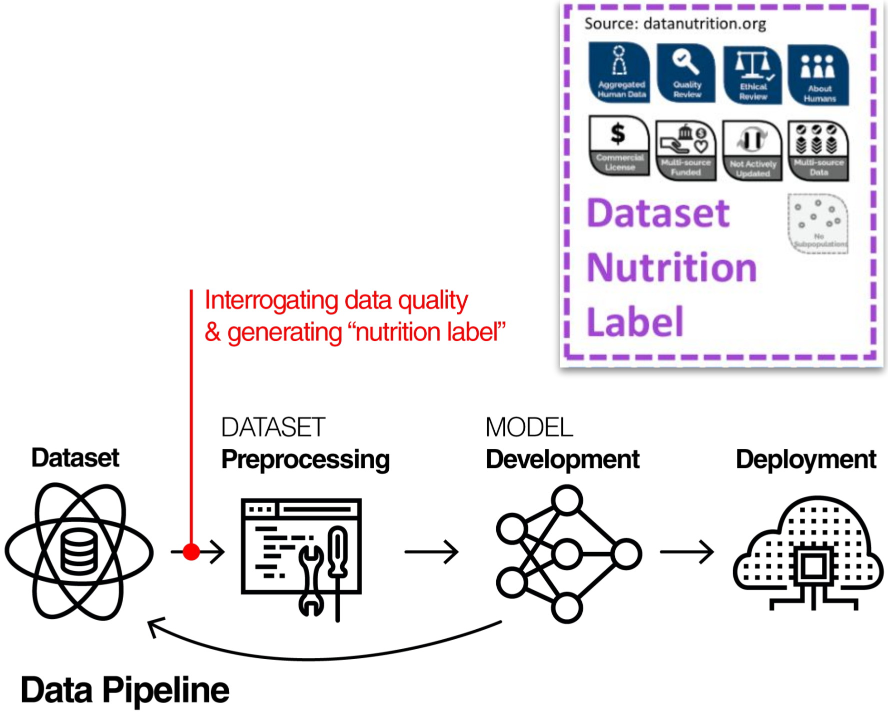

“*There’s a missing step in the AI development pipeline: **assessing datasets** based on **standard quality measures** that are both qualitative and quantitative. We are working on **packaging up these measures** into an easy to use Dataset Nutrition Label.*” - Dataset Nutrition Project

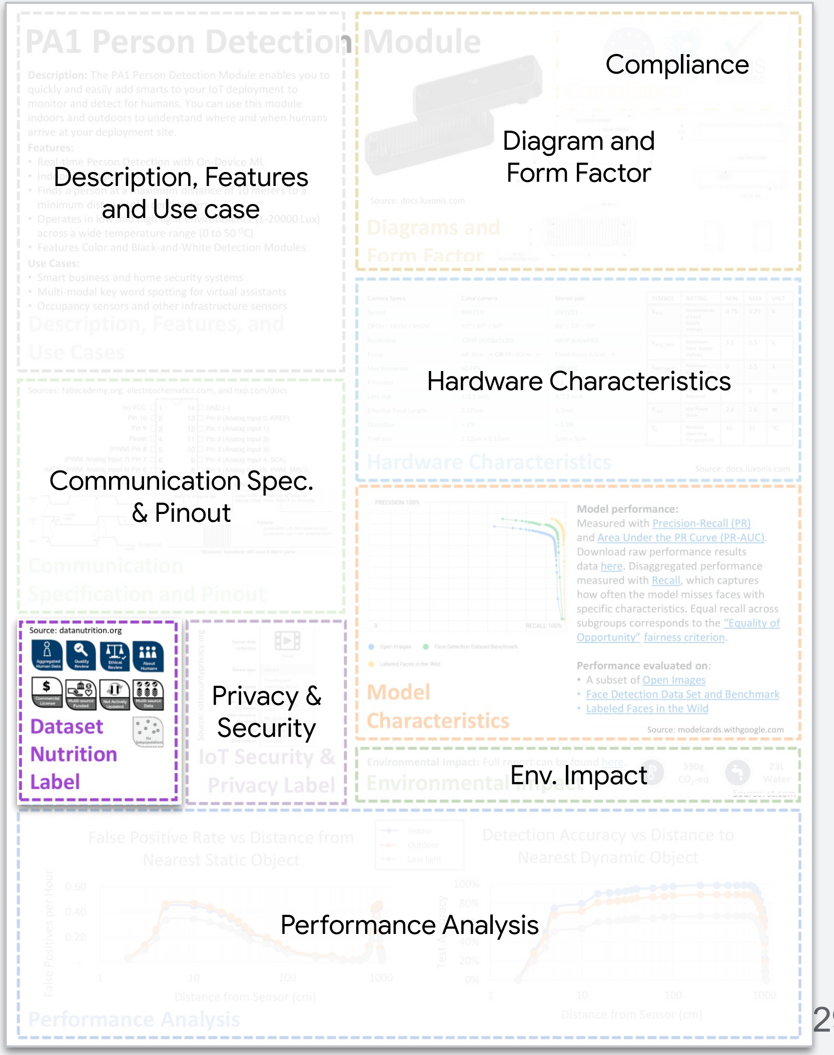

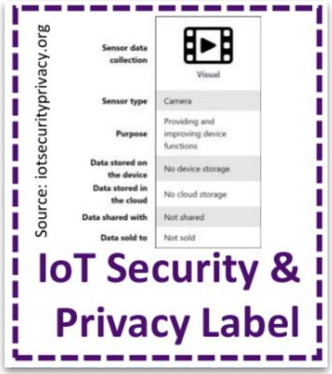

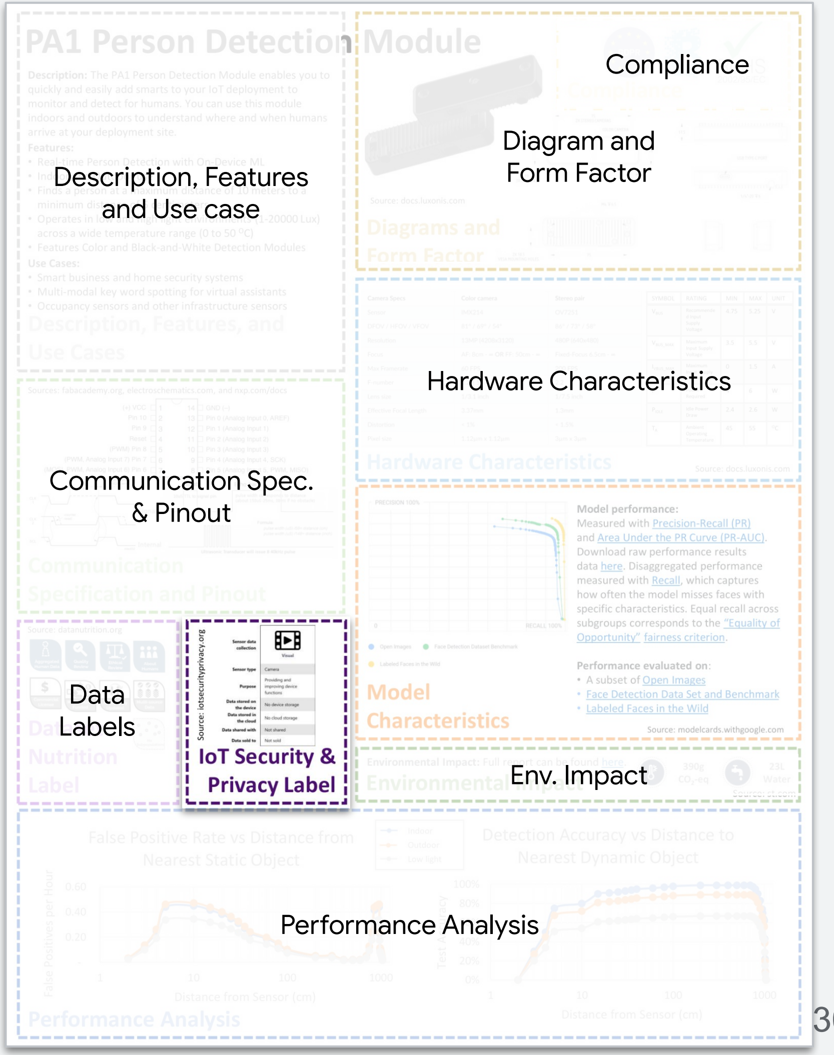

"… *designing a* ***usable security and privacy label*** *for smart devices to* ***help consumers*** *make informed choices about Internet of Things device purchases and encourage manufacturers to* ***disclose their privacy and security practices.***" – IoT Security & Privacy

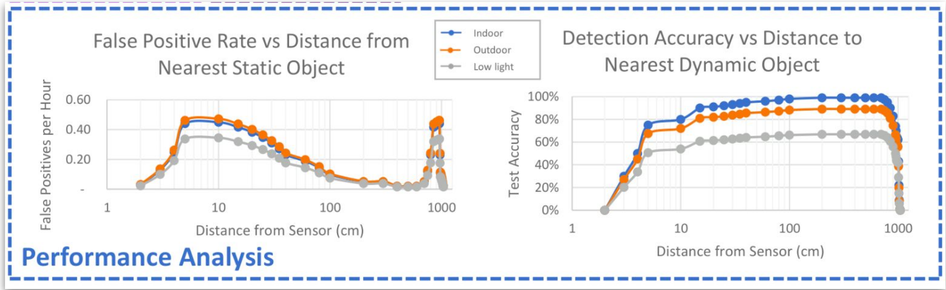

We require ***systematic methodologies*** *to evaluate how an* ***end-to-end system*** *performs under* ***real-world conditions***

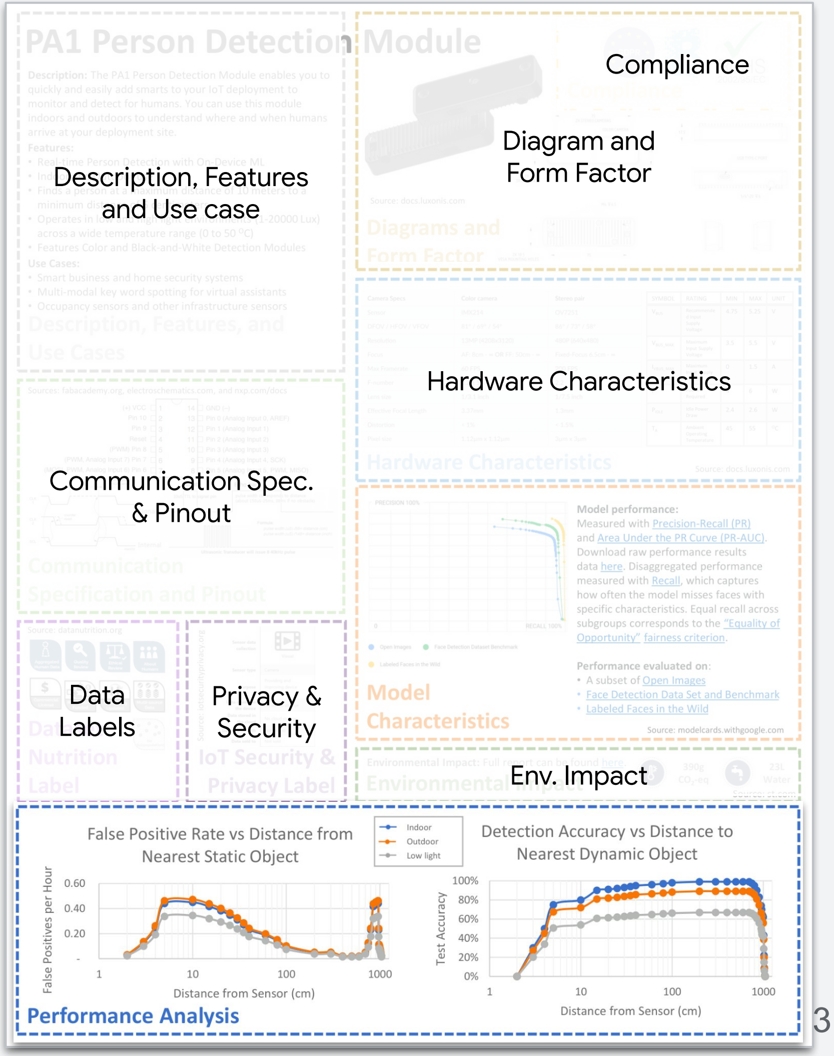

ML sensors ought to be ***tested by 3rd party certification*** *agencies or bodies that* ***specialize in AI/ML technologies.***

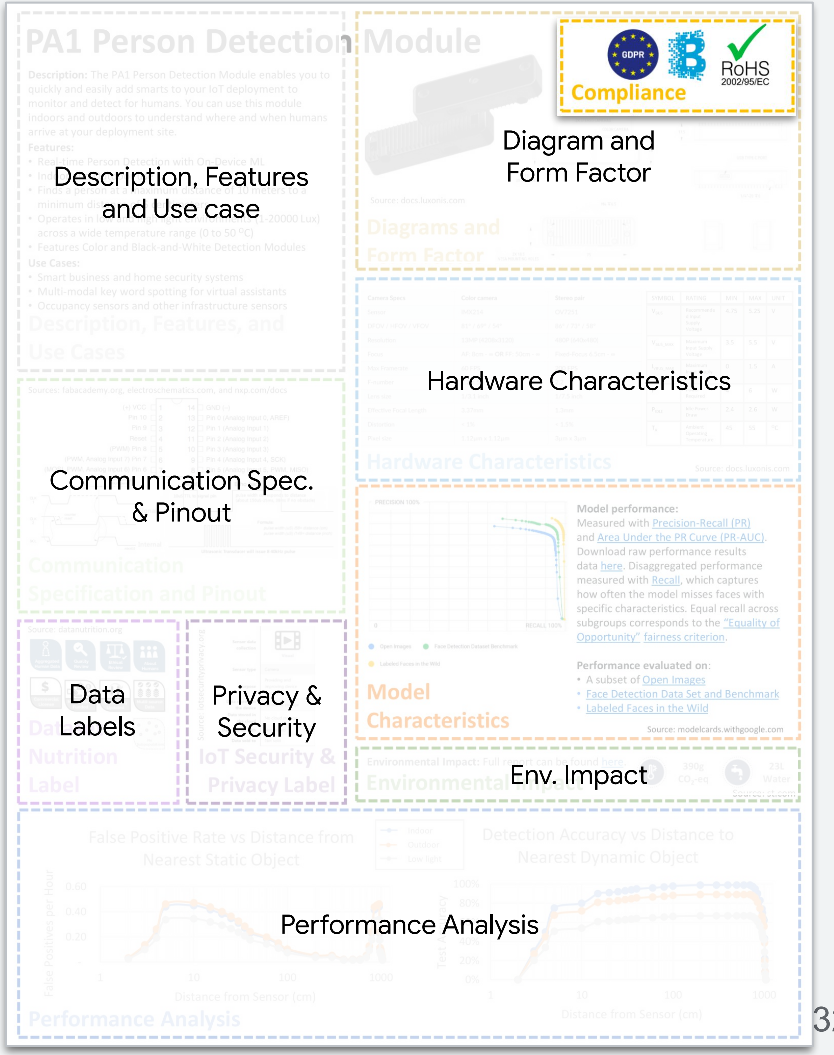

IoT ≈ The Internet of Trash?

We must quantify the effects of ML sensors in terms of carbon emissions. Carbon emissions have two sources: (1) ***operational energy consumption****, and (2)* ***hardware manufacturing and infrastructure****. The former has been decreasing thanks to software and hardware innovations but the* ***total footprint is growing****.*

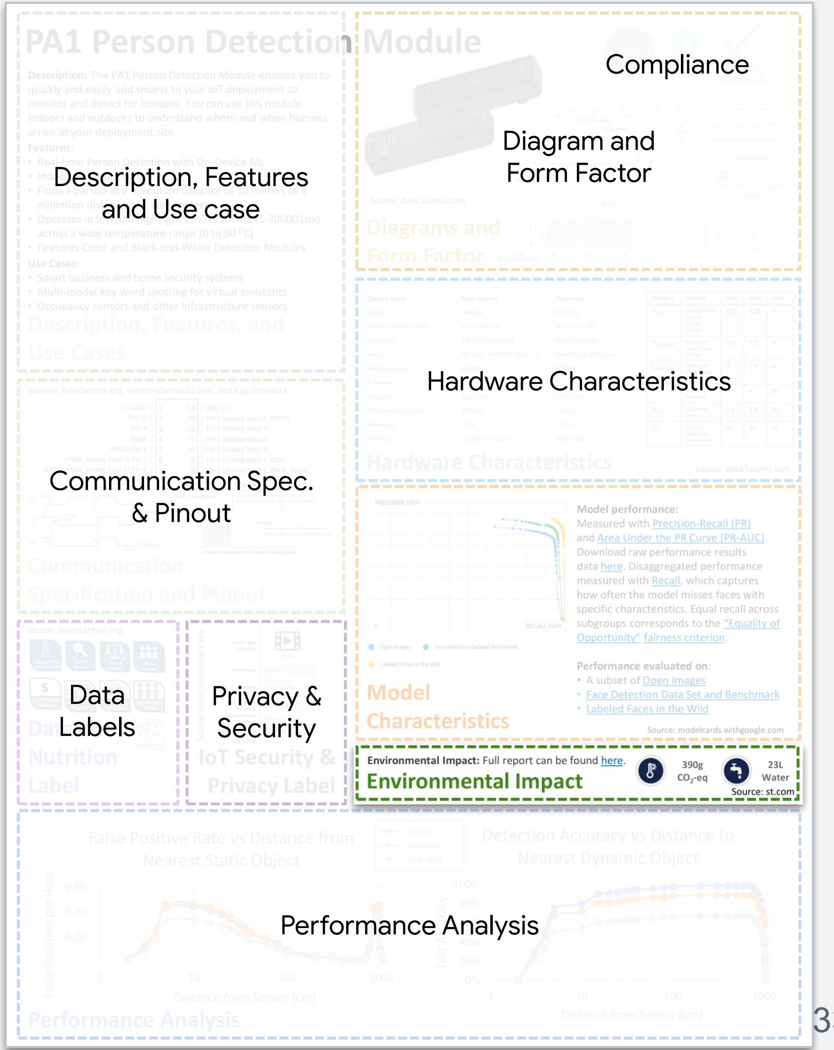

Assessing the Environmental Impact of an **MCU**

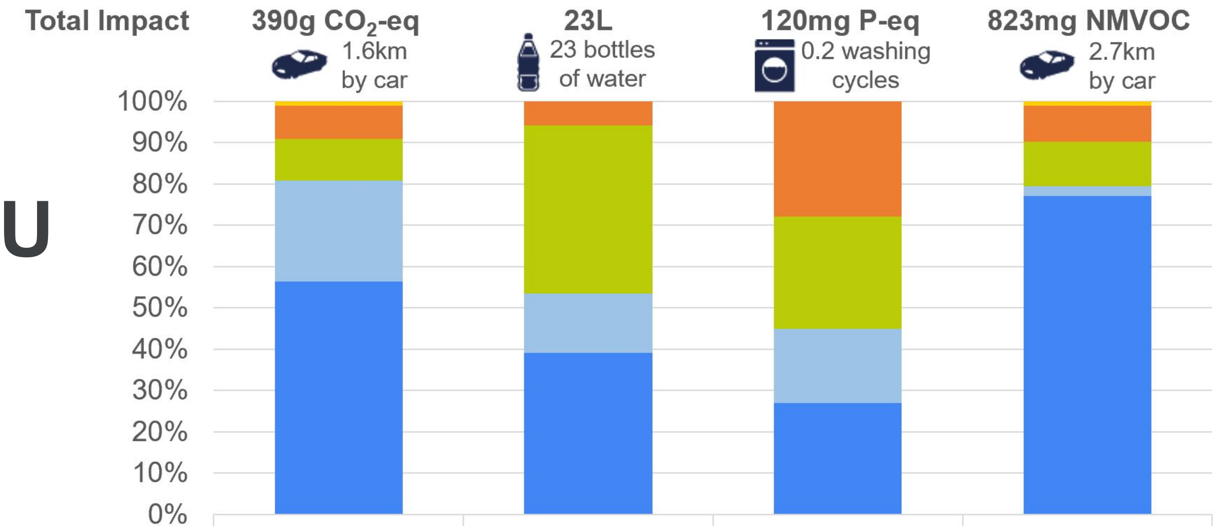

	Climate Change	Water Demand	Freshwater Eutrophication	Protochemical Oxidant Formation
End of Life	<1%	<1%	<1%	<1%
Logistics	1%	<1%	<1%	1%
Use	8%	6%	28%	8%
Raw Materials	10%	41%	27%	10%
Production: Other	24%	15%	18%	2%
Production: Energy Consumption	56%	39%	27%	71%

Source: https://www.st.com/content/st_com/en/about/st_approach_to_sustainability/sustainability-priorities/sustainable-technology/eco-design/footprint-of-a-microcontroller.html

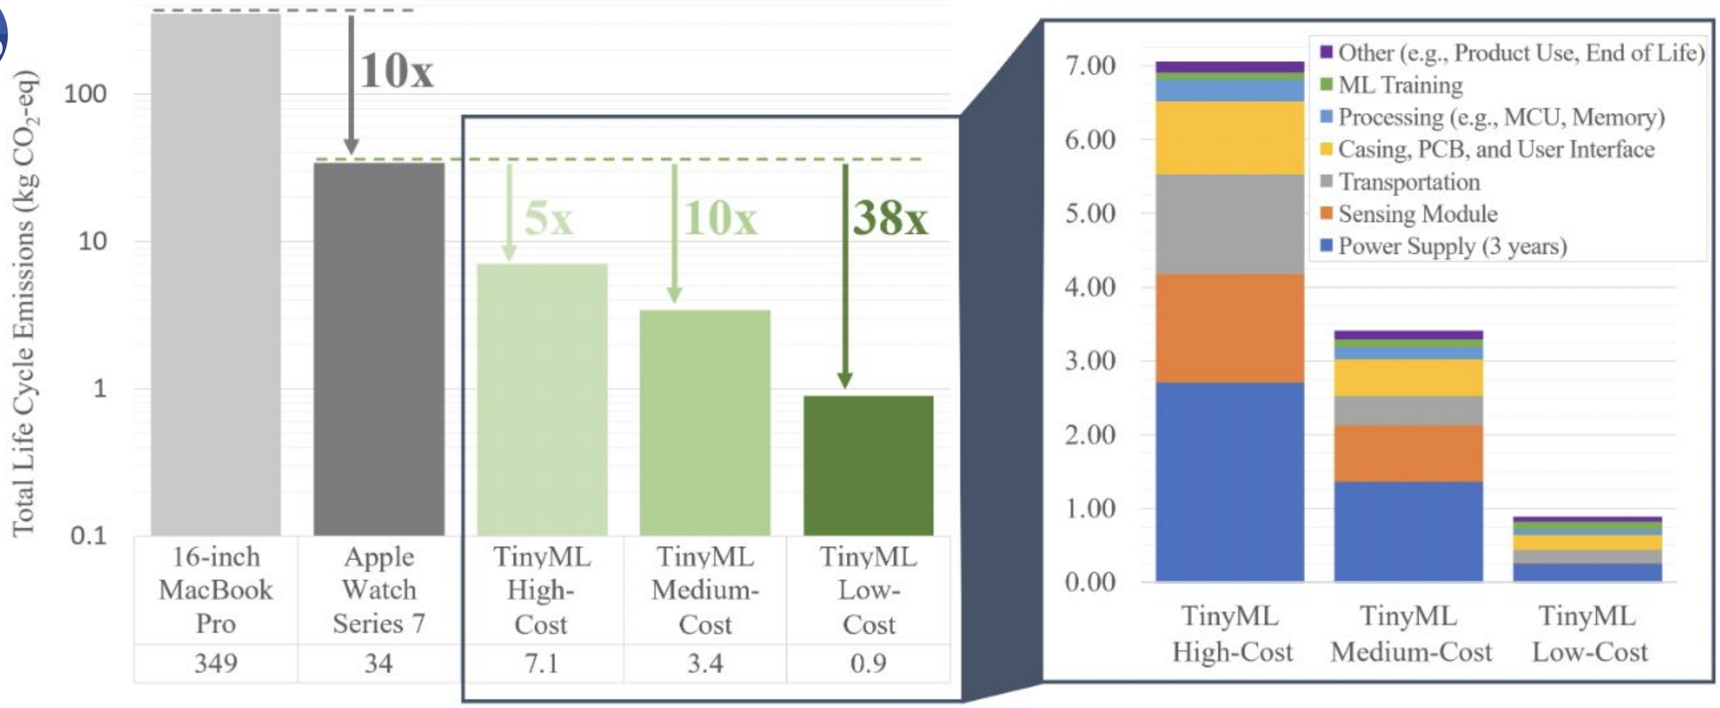

Figure 4. A breakdown of different TinyML system footprints highlights that the footprint is largely attributable to the embodied footprint of the power supply, onboard sensors, and transportation. Note that actuator and connectivity blocks from Pirson and Bol [21] are encapsulated in "Other" and "Processing", respectively, while "Product Use" captures the operational footprint. The carbon footprint of TinyML Systems was also compared with Apple's Series 7 Watch [12] and 16-inch MacBook Pro [11] as baseline references. For more details and to compute the footprint of your own TinyML system see github.com/harvard-edge/TinyML-Footprint.

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.

3. 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.

4. 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.

5. 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.

Machine Learning Sensors - ML

mlsensors.org

TinyML | Harvard | MLC | Research | Seeed | CS141 | TimeBuddy | VJs Funding | Enterprise - Suppl... | Geo Chart Exampl... | Other Bookmarks

MLSensors
Machine Learning Sensors

Home | Whitepaper | 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

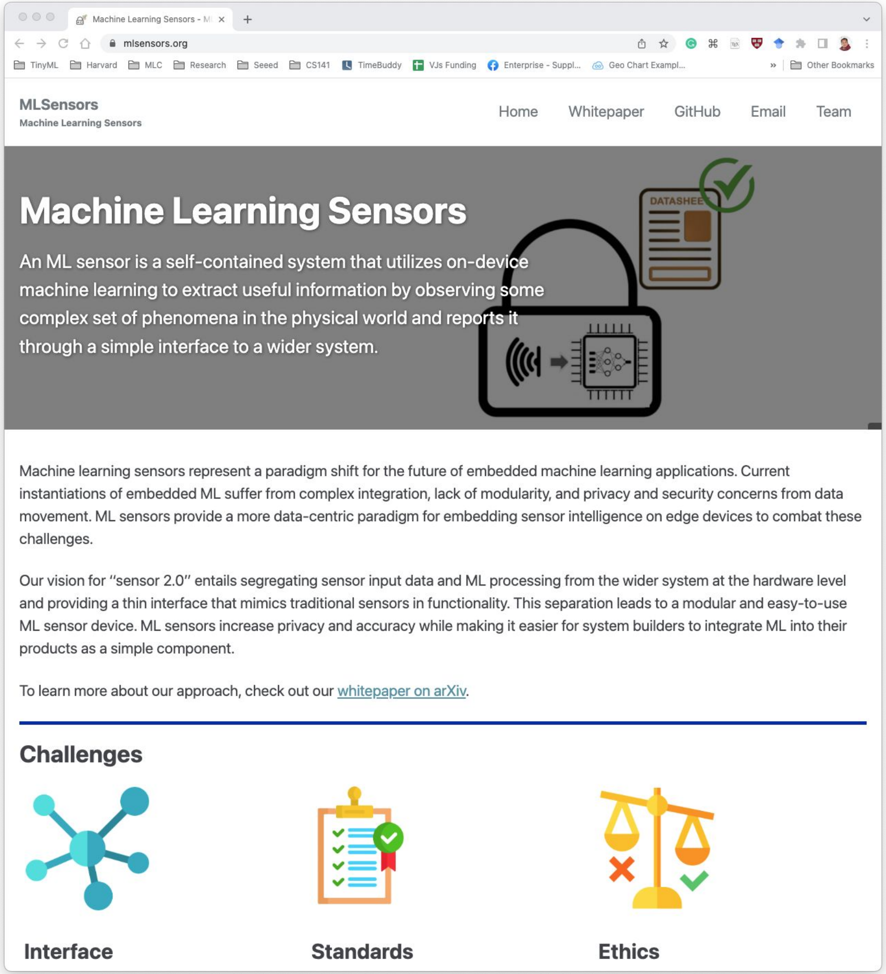

MLSensors
Machine Learning Sensors

Home Whitepaper 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 whitepaper on arXiv.

Challenges

Interface Standards Ethics

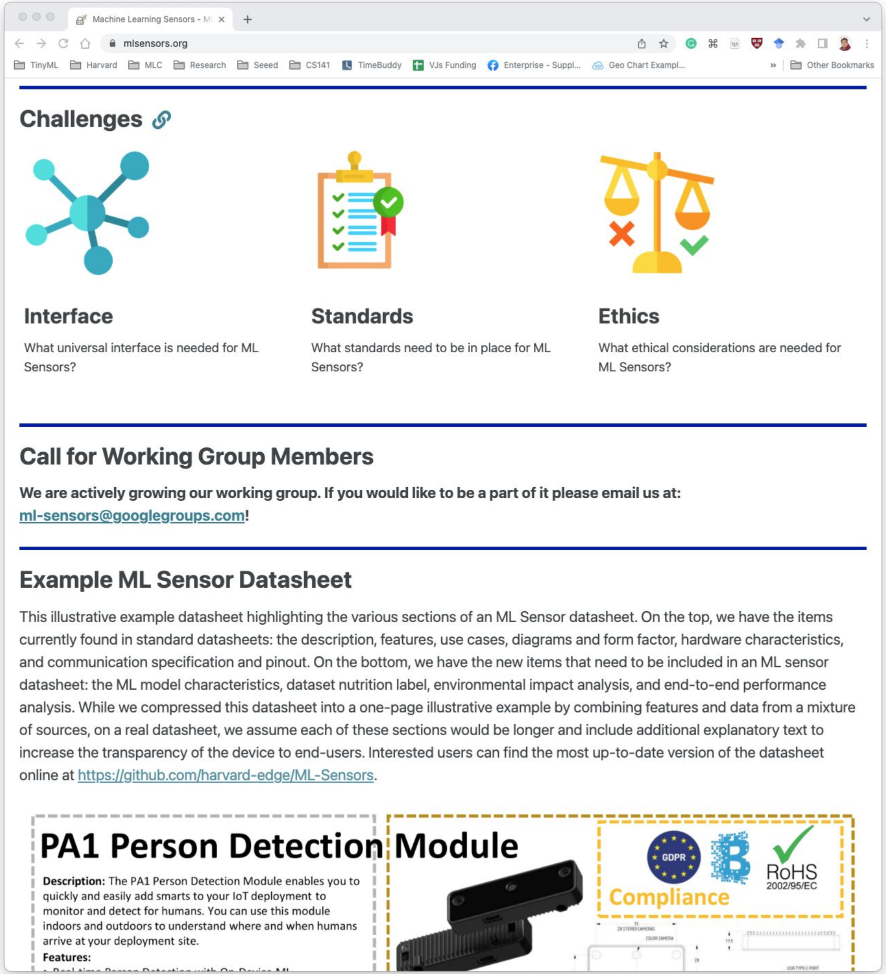

Challenges

Interface
What universal interface is needed for ML Sensors?

Standards
What standards need to be in place for ML Sensors?

Ethics
What ethical considerations are needed for ML Sensors?

Call for Working Group Members

We are actively growing our working group. If you would like to be a part of it please email us at: ml-sensors@googlegroups.com!

Example ML Sensor Datasheet

This illustrative example datasheet highlighting the various sections of an ML Sensor datasheet. On the top, we have the items currently found in standard datasheets: the description, features, use cases, diagrams and form factor, hardware characteristics, and communication specification and pinout. On the bottom, we have the new items that need to be included in an ML sensor datasheet: the ML model characteristics, dataset nutrition label, environmental impact analysis, and end-to-end performance analysis. While we compressed this datasheet into a one-page illustrative example by combining features and data from a mixture of sources, on a real datasheet, we assume each of these sections would be longer and include additional explanatory text to increase the transparency of the device to end-users. Interested users can find the most up-to-date version of the datasheet online at https://github.com/harvard-edge/ML-Sensors.

PA1 Person Detection Module

Description: The PA1 Person Detection Module enables you to quickly and easily add smarts to your IoT deployment to monitor and detect for humans. You can use this module indoors and outdoors to understand where and when humans arrive at your deployment site.
Features:

Compliance GDPR RoHS 2002/95/EC

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.

3. 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.

4. 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.

5. 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.

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?

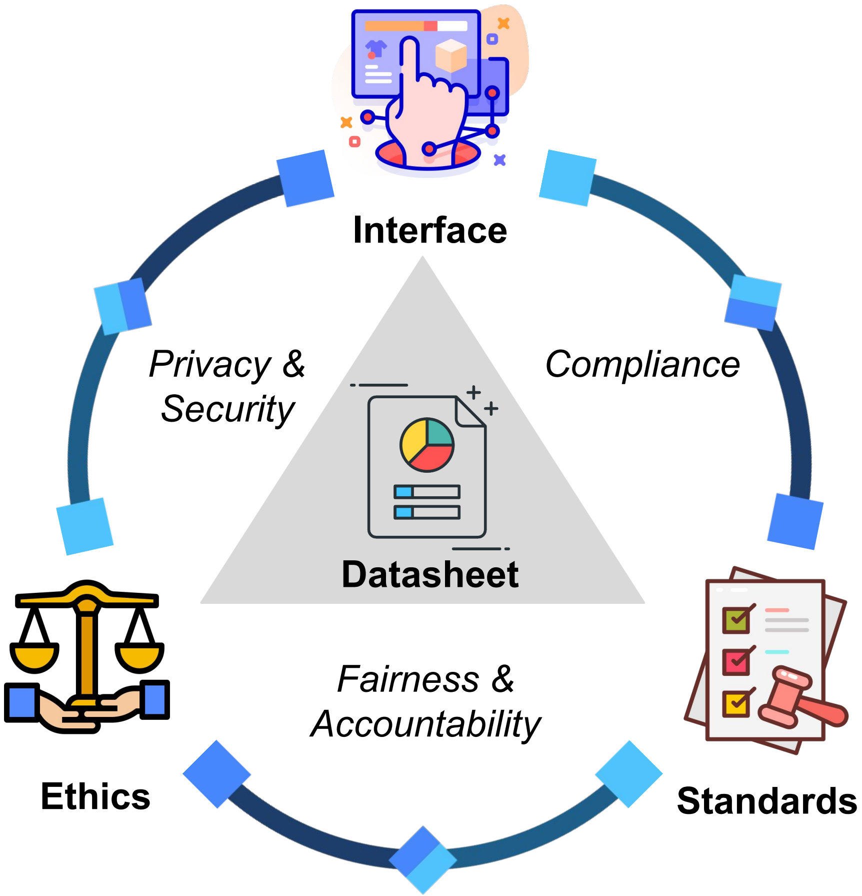

mlsensors.org

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

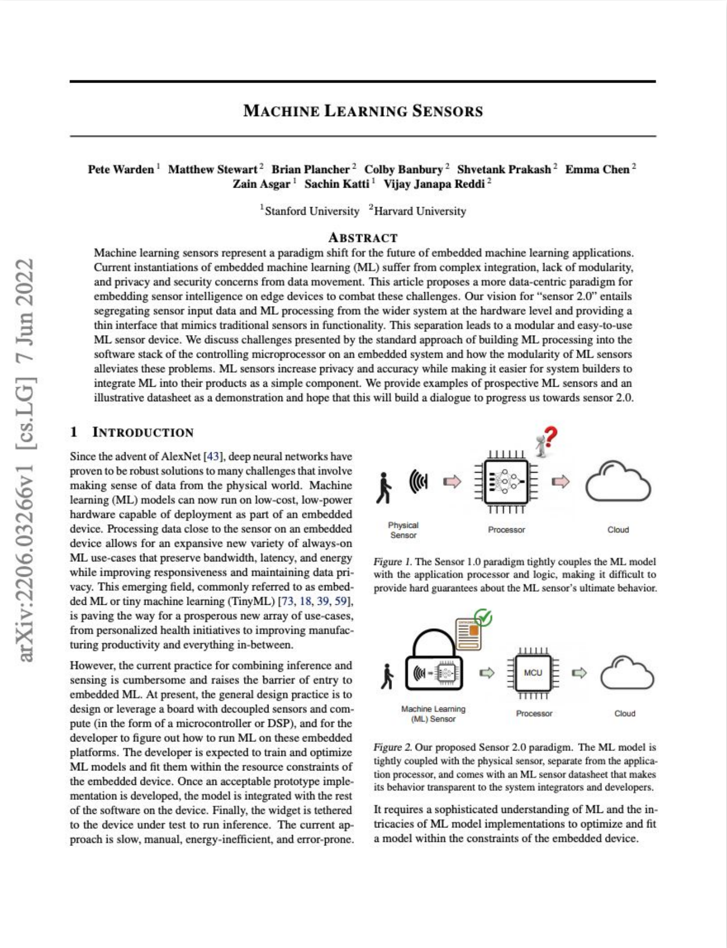

arXiv:2206.03266v1 [cs.LG] 7 Jun 2022

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.