Datasheets for Machine Learning Sensors

Matthew Stewart[1,*], Yuke Zhang[2,*,†], Pete Warden[3,4], Yasmine Omri[3], Shvetank Prakash[1], Jacob Huckelberry[1], Joao Henrique Santos[1], Shawn Hymel[5], Benjamin Yeager Brown[1], Jim MacArthur[1], Nat Jeffries[4], Emanuel Moss[6], Mona Sloane[7], Brian Plancher[8,9], and Vijay Janapa Reddi[1]

***Abstract*—Machine learning (ML) is becoming prevalent in embedded AI sensing systems. These "ML sensors" enable context-sensitive, real-time data collection and decision-making across diverse applications ranging from anomaly detection in industrial settings to wildlife tracking for conservation efforts. As such, there is a need to provide transparency in the operation of such ML-enabled sensing systems through comprehensive documentation. This is needed to enable their reproducibility, to address new compliance and auditing regimes mandated in regulation and industry-specific policy, and to verify and validate the responsible nature of their operation. To address this gap, we introduce the datasheet for ML sensors framework. We provide a comprehensive template, collaboratively developed in academia–industry partnerships, that captures the distinct attributes of ML sensors, including hardware specifications, ML model and dataset characteristics, end-to-end performance metrics, and environmental impacts. Our framework addresses the continuous streaming nature of sensor data, real-time processing requirements, and embeds benchmarking methodologies that reflect real-world deployment conditions, ensuring practical viability. Aligned with the FAIR principles (Findability, Accessibility, Interoperability, and Reusability), our approach enhances the transparency and reusability of ML sensor documentation across academic, industrial, and regulatory domains. To show the application of our approach, we present two datasheets: the first for an open-source ML sensor designed in-house and the second for a commercial ML sensor developed by industry collaborators, both performing computer vision-based person detection.**

I. Introduction

The merging of machine learning (ML), sophisticated data processing methods, and sensor technology is leading to rapid growth in the use of smart sensors [1], and the development of a new generation of intelligent sensing devices—ML sensors [2]. An ML sensor is a standalone unit that employs ML algorithms directly on the device to handle intricate sensor data, providing real-time analysis and data-driven decision making at the point of data collection for complex tasks such as activity recognition [3] and motion tracking [4]. Figure 1 outlines its key components: sensing elements, pre-processing capabilities, ML model execution, and communication. Importantly, unlike traditional sensors which require remote cloud processing, ML sensors deploy the ML model directly on the sensor, enabling both energy-efficient, and privacy-preserving, edge intelligence.

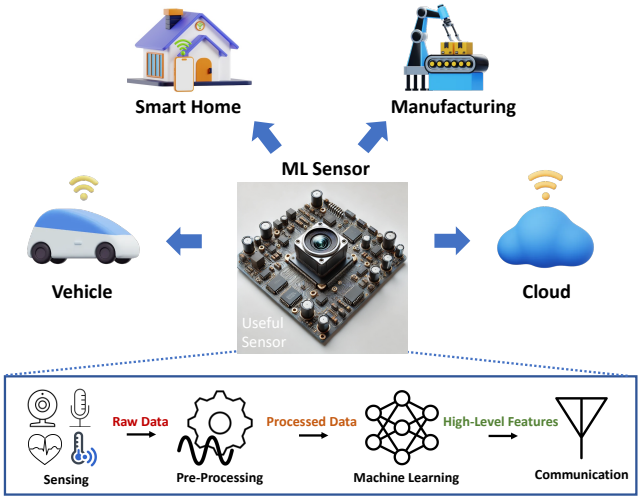

Fig. 1: The ML sensor paradigm: deploying machine learning models directly on the sensor for privacy-promoting and energy-efficient edge intelligence.

Although they represent an important advancement in gathering, analyzing, and responding to environmental data, there is no comprehensive framework for creating ML sensor-specific transparency to support reproducibility, analysis of social impact, and accountability, e.g., ML audits [5]–[7]. However, such transparency can be achieved by building on existing documentation tools and approaches, serving as a bridge between ML sensor producers, end users, and academic researchers. When ML sensor transparency takes the form of standardized documentation, academic researchers can better understand the practical needs and priorities of industry, aligning research with real-world applications, and supporting efforts to make ML sensors auditable and compliant with new and existing regulations (e.g., ML system audits increasingly required under the EU AI Act [6], [8]). Academic researchers can also better understand specifications and performance metrics to develop new applications, comparative studies, and further advancements in sensor technology. At the same time, ML sensor producers can receive valuable feedback from users and researchers, informing future improvements and fostering innovation, while end-users can transparently understand the capabilities of the devices they purchase.

There are existing documentation approaches that have the potential to provide high levels of transparency for ML sensors. However, these existing frameworks often do not fully address the distinct complexities associated with ML sensors. For example, model cards [9] tend to emphasize model design

[1]Harvard University, Cambridge, MA, USA
[2]University of Toronto, Toronto, ON, Canada
[3]Stanford University, Stanford, CA, USA
[4]Useful Sensors, USA
[5]Independent Researcher, USA
[6]Intel Labs, Hillsboro, OR, USA
[7]University of Virginia, Charlottesville, VA, USA
[8]Barnard College, Columbia University, New York, NY, USA
[9]Dartmouth College, Hanover, NH, USA
*These authors contributed equally to this work.
† Corresponding author contact: yuke.zhang@utoronto.ca

and performance metrics, but omit detailed discussions on data management and sensor privacy compliance. Similarly, traditional sensor datasheets for sensor hardware focus on technical details, such as power consumption and latency, without considering the ethical aspects of data collection and processing, or how sensors address bias and privacy issues.

To bridge this gap, we introduce a ML sensor-specific datasheet framework that focuses on creating transparency for reproducibility and auditability of ML sensors. Importantly, our framework was developed by deploying a rigorous engagement approach focused on industry, academia, and government stakeholder engagement through seminars, workshops, and in-depth discussions, that was combined with an iterative refinement process through multiple rounds of testing. The resulting ML sensor datasheet framework offers standardized sections that capture vital information about the sensor's capabilities, data processing flows, and social impact dimensions. To establish transparency for auditability, our datasheet offers comprehensive information about the ML sensor, covering every aspect from the sensing device and hardware specifications to the embedded ML model and system-level performance metrics. Additionally, our datasheet introduces components on environmental impact and end-to-end performance, providing a holistic view of the sensor's sustainability and operational resilience in diverse conditions. To enable public access, for transparency and feedback, we *open-source* our datasheets at: github.com/harvard-edge/ML-Sensors.

Our work supports standardizing the reporting of ML sensor features to facilitate transparency for auditability and to help developers and organizations follow best practices, support ongoing enhancements as standards progress, and ultimately encourage more responsible usage of ML sensors in sensitive and emergent contexts. The contributions of our work are:

- A novel **ML sensor datasheet framework**, developed via industry, academia, and government stakeholder engagement and iterative refinement, that addresses the unique **transparency challenges of ML sensors** that are not fully captured by existing documentation approaches.
- Our **datasheet enables transparency for auditability** of ML sensors, bridging the gap between ML algorithms and embedded systems. It provides a **standardized approach** for documenting important aspects such as data management, privacy compliance, and social impacts, in addition to traditional technical specifications.
- We demonstrate the **practical application of our framework** through two case studies: an **open-source ML sensor** and a **proprietary sensor**. These examples show how our datasheet can be effectively implemented across different types of ML sensors and establish **best practices** for future development and deployment of ML sensors.

II. Background and Related Work

A. Machine Learning Sensors

Embedded sensing systems have a long history and serve as fundamental building blocks for many applications [2], [15], converting physical phenomena into signals for processing. Embedded AI sensing has marked a significant evolution,

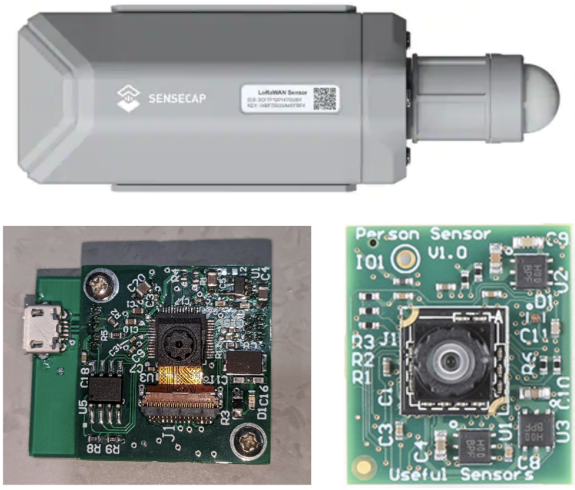

Fig. 2: Examples of existing ML sensors; *(top)* Seeed Studio's SenseCAP LoRaWAN sensor [17] for long-range data collection IoT scenarios like smart farming, *(bottom left)* our own person detection sensor whose design is open-source (link redacted for double blind), and *(bottom right)*, Useful Sensor's person sensor [18]. We use these person detection sensors for our case study applying our ML sensor datasheet to real ML sensors in Section V.

introducing sophisticated on-device processing capabilities that can extract meaningful insights directly from sensor data. However, this evolution also introduced new intertwined challenges of hardware optimization, ML deployment, and system design that developers must balance [2], creating significant barriers to entry.

The "ML sensors" paradigm reimagines this evolution by elevating the abstraction level of embedded AI development [2], [16]. As illustrated in Figure 2, existing commercial ML sensors encapsulate both sensing and computation capabilities within a unified package. These specialized devices focus exclusively on performing ML inference on physical phenomena, communicating only processed outcomes through a streamlined interface to the broader system.

This change transforms embedded AI development through several innovations. The unified design improves the locality of the data and the processing efficiency while significantly improving the modularity and maintainability of the system. The encapsulated approach strengthens privacy and security by limiting the exposure of raw sensor data, while the streamlined interface simplifies reasoning about data flows and system behavior. This modular and self-contained approach to embedded AI significantly simplifies the documentation of technical specifications, model characteristics, and operational constraints through our proposed datasheets as it enables standalone evaluation of the ML sensor product absent from its interaction with large-scale cloud-based data processing.

As such, just as engineers today can readily integrate standard temperature or pressure sensors into their designs, with ML sensors, *that have transparent and comprehensive datasheets*, we envision a future where developers can similarly incorporate pre-packaged "person sensors" or "gesture sensor."

TABLE I: Comparison of ML sensor datasheets with other datasheet types.

Datasheets ↓	Hardware	Privacy & Security	Dataset	Model	Env. Impact	End-to-End
ML Sensors (our work)	✓	✓	✓	✓	✓	✓
Traditional Sensor Datasheet	✓	✗	✗	✗	✗	✗
IoT Security/Privacy Label [10], [11]	✗	✓	✗	✗	✗	✗
Data Nutrition Label [12], [13]	✗	✗	✓	✗	✗	✗
Datasheets for Datasets [14]	✗	✗	✓	✗	✗	✗
Model Cards [9]	✗	✗	✗	✓	✗	✗

B. Datasheets and Documentation in AI

Traditional ML benchmarks and datasheets have predominantly focused on software models, covering data collection, cleaning, labeling, and intended use, both for specific datasets [19]–[21], and for more general frameworks for documenting dataset characteristics [12]–[14]. While efforts have been made to document various features of ML services [22], provide benchmarked evaluations for ML models [9], and include relevant privacy and security information for IoT devices [10], [11], these approaches fall short in capturing the unique challenges posed by the integration of hardware and software in ML sensors. Moreover, recent efforts to highlight potential social impacts of ML and sensor devices [16], [23], [24] as well as efforts to characterize operational and embodied emissions of hardware devices [15], [25], further underscores the need for a new comprehensive datasheet format. As shown in Table I, ML sensors uniquely encompass integrated hardware, software, and machine learning elements, necessitating a novel approach that amalgamates diverse concepts. Therefore, build on prior work and augment traditional sensor datasheets with vital information relevant for ML and IoT, as well as other new elements, resulting in a comprehensive and extendable datasheet that can effectively address the challenges of transparency and auditability in this emerging paradigm.

III. Datasheet Design

A. Design Motivation and Methodology

ML-enabled sensors and devices often provide datasheets with hardware specifications and instructions for using the sensor capabilities. However, these existing datasheets often lack substantial details regarding: (1) information on data characteristics, distributions, and potential biases which would build trust in model performance; (2) model architecture and benchmarks, with details on design choices and performance evaluations on standard datasets allowing better assessment of expected capabilities; (3) environmental impact over the device lifetime, including materials, energy use, and recyclability allowing more sustainable adoption; (4) robustness to anticipated changes in operating conditions once deployed, with testing results on transformations in factors like lighting, backgrounds, weather, and wear guiding appropriate scenarios; and (5) aspects of privacy and security, such as disclosure and protections around data sharing, vulnerabilities, and compliance.

Creating transparency by providing documentation can increase user trust, facilitate accountability, and encourage responsible development and adoption of ML innovations across industries [9]. It also enables those implementing ML systems to ensure that outcomes align with their principles, policies, compliance responsibilities. The guiding motivation is to clearly communicate the inner workings of the ML sensor, including details about its embedded ML models, training data, and performance metrics to ensure that users, developers, stakeholders, and auditors can easily understand how the sensor processes data, makes decisions, and performs under various conditions. To support this, the datasheet includes descriptions of AI models and end-to-end performance analysis summaries that outline how the sensor operates in different environments, as well as sections that document bias mitigation strategies, privacy safeguards, and compliance standards and policies.

In line with the FAIR principles, we ensure that ML sensor datasheets are structured to promote discoverability, integration, and reuse across diverse contexts. These principles are important as ML sensors become part of broader ecosystems that require standardization, transparency, and traceability. Therefore, the development of our datasheet framework for ML sensors was based on a rigorous and multifaceted *engagement process with industry, academia, and government stakeholders* and *iterative refinement*. This process ensured a holistic approach the design process in order to address the complex challenges of ML sensor documentation.

Stakeholder Engagement. We engaged with a broad spectrum of stakeholders, ranging from multinational corporations (e.g., Google, Dyson, LG, Samsung), to emerging startups, to government entities, to academia. The diverse industry and government participation was crucial in capturing real-world implementation needs on various operational scales and use cases. Our academic engagement centered on a two-day seminar at the Harvard Radcliffe Institute, titled "Safeguarding User Privacy in the era of Sensor Intelligence". This event brought together experts from multiple disciplines to critically examine risks and power structures associated with ML sensors, informing the essential components of our data sheet. Finally, to ensure compliance with international legal standards and regulations, we incorporated information on data protection and privacy regulations into our framework (e.g., GDPR [26]).

Iterative Refinement. The overall development of our ML sensor datasheet framework was not a linear process, but rather an iterative process. Our initial template, while comprehensive, served as a starting point for a series of refinements that would shape it into a more robust and universally applicable tool. Throughout this refinement process, we leveraged multiple channels of feedback and expertise.

Our ML sensor-focused workshops held at the TinyML

Summit [27] have proved to be useful forums for engaging with a diverse community of practitioners and researchers. These workshops provided a platform for lively discussions and critical evaluations of our framework, offering insights from various perspectives and use cases that we might not have initially considered. Parallel to these workshops, we entered into a series of in-depth discussions with the National Institute of Standards and Technology (NIST) [28]. These conversations helped us align our framework with emerging standards and best practices in the field of ML and sensor technologies. The expertise provided by NIST was instrumental in ensuring that our datasheet not only met current needs but was also positioned to adapt to future developments in the rapidly evolving landscape.

The most crucial aspect of our refinement process was the practical application of the datasheet to real-world ML sensor products. We tested our framework on a range of devices, from open-source projects to commercial products, spanning various applications and complexities. This hands-on approach, discussed in greater detail in the case studies (Section V-C), allowed us to identify gaps in our documentation, uncover unforeseen challenges, and validate the practical utility of our framework in different contexts.

Each iteration of testing and feedback led to adjustments in our datasheet structure, content, and guidance. We fine-tuned categories, clarified definitions, and added new sections to address emerging concerns. This iterative process of application, feedback, and refinement was crucial to transform our initial concept into a comprehensive, flexible, and practical tool to document ML sensors.

Key Learnings. Our learnings from academia include: (a) the importance of addressing bias in sensor data collection; (b) the need for transparency in algorithmic decision-making processes; (c) emphasis on the sustainability and ethical responsibility. Our insights from industry and government include: (a) the importance of real-time capabilities in ML sensors; (b) the need for clear and transparent metrics; (c) emphasis on integration with existing sensor ecosystems, regulations, and frameworks.

Accessibility and Maintenance. We have open-sourced our datasheet to enhance maintenance, distribution, and accessibility. All stakeholders have easy access to the most current version of our datasheet. We gather feedback from all relevant stakeholders and update the datasheet to reflect any changes.

Outcome. The result of this process is a datasheet framework that is stakeholder-driven and practically validated across a spectrum of real-world applications. However, we recognize that the ML sensors is dynamic and ever-evolving. As such, we view our current framework not as a final product, but as a robust foundation ready for further adaptation and improvement as the technology and its applications continue to advance.

B. Overview of the Proposed Datasheet

Based on the principles and design methodology, our datasheet, shown in Figure 3, is segmented into these areas:

- **Standard Sensor Datasheet Components.** This section mirrors conventional electronics datasheets, encapsulating core sensor information such as a detailed description, features, use cases, communication specifications, pinout details (the function of each connector or pin), compliance with industry standards, as well as physical attributes like diagrams, form factors, and hardware characteristics.
- **IoT Datasheet Components.** Dedicated to the integration of the sensor within Internet of Things ecosystems, this part outlines privacy and security protocols specific to IoT, ensuring that the sensor's deployment aligns with modern cybersecurity practices.
- **AI Datasheet Components.** This segment provides in-depth information about the AI models embedded in the sensor, including the type of algorithms used, the training data, and any pertinent model characteristics that users and developers should be aware of.
- **ML Sensor Datasheet Components.** To address the needs of ML sensors, this section adds more layers of analysis, emphasizing documentation of the ML model as deployed on the specific sensor, which makes the environmental and end-to-end performance analysis possible. Our datasheet includes the environmental impact, showing the device's footprint and sustainability, and an end-to-end performance analysis which reviews the sensor's functionality in real-world scenarios, considering factors like varying environmental conditions and fairness parameters.

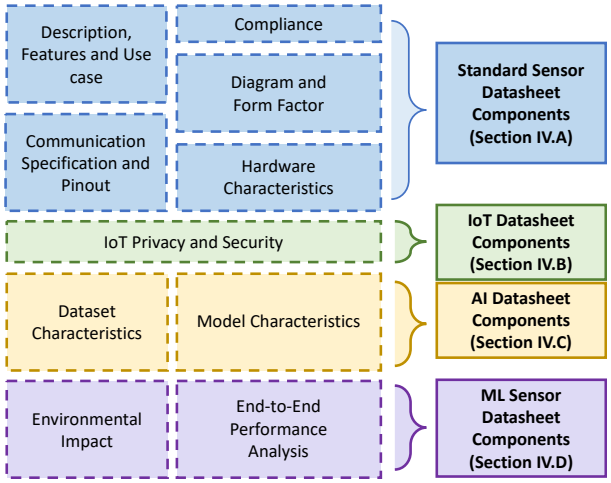

Fig. 3: Schematic of the proposed datasheet template for ML sensors.

Different from prior work [9]–[14], our datasheet presents hardware characteristics, details about the embedded ML model, IoT privacy and security protocols, and, notably, the system-level ML sensor features of environmental impact and end-to-end performance analysis. These system-level features were included based on feedback from seminars and workshops, emphasizing the increasingly urgent need for sustainable design and reliable performance metrics in real-world applications. Our datasheet format also aligns with the FAIR principles by making sensor metadata findable and well-documented, enabling accessible and machine-readable formats, promoting interoperability with existing documentation standards, and supporting the reuse of sensor components, datasets, and models in diverse deployments. In Sections IV-A-IV-D, we provide a detailed breakdown of the four key components.

IV. Datasheet Implementation

In this section, we provide a detailed breakdown of the four key components of the proposed datasheet and clarify the distinct elements relevant to the nine targeted tasks.

A. Standard Datasheet Components

Standard datasheets provide general electronic and physical specifications to help developers, users, and auditors understand the device's basic capabilities, requirements, and form factor (i.e., its size, shape and layout).

Description, Features, and Use Case. *What are high-level characteristics of the sensor?* The description section of the ML sensor datasheet provides an introduction to the device for both technical and non-technical audiences. On the technical side, it includes details about the device's specifications, architecture, and operational principles. For non-technical readers, it offers a more accessible description, explaining the sensor's purpose and function in plain language. This section also highlights key features of the device and a list of common applications of the sensor.

Diagrams and Communication Specification. *What does the device size, shape, and layout look like?* This section provides visual depictions and physical dimensions of the device. It includes detailed diagrams that illustrate the sensor's internal components and their interconnections, offering insights into the design and operation of the sensor. For non-technical audiences, these diagrams can provide a more intuitive understanding of the device, beyond what text descriptions can offer. These diagrams include form factor information which describes the physical shape, size, and layout of the sensor. This data is crucial in planning the sensor's integration into various systems and devices.

Hardware Characteristics. *What components are a part of the device? How do they operate?* This section provides an overview of the physical and functional attributes of the device. It contains specifics about the sensor's integral hardware components, including the processor type, memory capacity, power requirements, and durability under different environmental conditions. In addition, it includes detailed information about the communication protocols supported by the sensor, such as Wi-Fi, Bluetooth, or cellular connectivity, along with data transfer rates. This data is crucial in determining the sensor's compatibility with existing hardware infrastructure and intended deployment environments.

Compliance and Certification. *Which international regulations and industry standards does the device conform to?* This section lists the certifications the sensor has achieved, signifying thorough testing and validation by recognized certification bodies. Compliance with these standards vouches for the sensor's reliability, safety, and overall quality, instilling confidence in developers and end-users about its dependable operation. These may encompass international data privacy regulations like GDPR [29], radio frequency usage guidelines like FCC regulations, or industry-specific requirements like HIPAA [30] or FDA standards in healthcare. This section could also document the ML sensor's adherence to voluntary industry-specific best practices, such as ISO 26262 standard [31] for autonomous vehicles. Additionally, given the rapidly evolving regulatory landscape for AI-driven technologies, frameworks like the EU AI Act are increasingly relevant. Anticipating such developments and providing transparency via the datasheet ensures the sensor and any systems the sensor is embedded in can be audited for compliance with forthcoming standards and facilitates comprehensive, future-proof documentation.

B. IoT Datasheet Components

This section summarizes the ML sensor's safeguards and risks from an IoT perspective, facilitating awareness and informed decision-making when purchasing smart devices.

Security and Privacy. *What IoT security and privacy features does the ML sensor have?* This section promotes transparency and empower consumers, allowing them to make well-informed choices in an increasingly connected world. This section is structured in two distinct layers: a primary layer, which conveys essential privacy and security information in a concise and easily digestible manner, and a secondary layer, which delves into further detail for experts and more technically inclined users [11]. Combined, the layers cover privacy-related aspects such as data collection, retention and transmission practices, security mechanisms (e.g., encryption, automatic security updates), as well as the types of sensors present on the device and their associated data modalities. Information regarding the ability to update the device's ML models could also be included in this section, including details about the frequency, method, and security measures of these updates can further inform users about the longevity and reliability of the device's performance.

In addition to the primary privacy and security elements, other crucial issues may include the risk of unauthorized data access [32], vulnerabilities related to device authentication [33] and cyberattacks [34], [35]. The datasheet can incorporate information regarding measures taken to mitigate these threats, such as multi-factor authentication protocols and tamper detection mechanisms. Furthermore, in privacy-sensitive domains like healthcare or surveillance, datasheets should include results from privacy impact assessments and audits, which demonstrate the sensor's compliance with stringent data protection frameworks [36]. This expanded view ensures that consumers and technical stakeholders have a comprehensive understanding of how the ML sensor aligns with best practices for safeguarding sensitive data in specific use cases.

C. AI Datasheet Components

This section details the data used for training and the resulting model, enabling stakeholders to evaluate aspects like sampling, measurement, and bias.

Dataset Characteristics. *What data is the model trained on?* Outlining dataset characteristics is fundamental for ensuring transparency in ML systems. This transparency is crucial because it allows users, developers, and regulators to understand how and why a ML model makes certain decisions, and to assess its fairness, accuracy, and potential biases. By clearly detailing the nature of the training data, stakeholders can better evaluate the model's applicability to real-world scenarios and

audit the dataset for how well it corresponds with its context of use. Several approaches have been proposed to achieve this, including datasheets for datasets [14] and the data nutrition label [12].

Taking the data nutrition label as an example, this label communicates high-level dataset information to end-users, including (1) the sources of the dataset (i.e., governmental, commercial, academic), (2) licensing details of the dataset, (3) data modality, and (4) context-specific information (e.g., human-labeled, contains information about human individuals), amongst other information. This provides information about the context, content, and quality of dataset(s) used in training the ML model. As such, it makes it easier for developers, researchers, and stakeholders to assess data quality and potential biases such as sampling, measurement, and label bias [37].

Model Characteristics. *What are the characteristics of the trained model? What are the training method and hyperparameters?* This section of the datasheet provides insights into the specific ML model operating within the sensor. This includes important details such as the type of the ML model used, the size of the model in terms of parameters, the type and size of input data it can process, and the nature of output it generates. This section also discusses the model's performance metrics, such as accuracy, precision, recall, F1 score, and performance across demographic categories, measured on a relevant validation dataset. It may also address the model's robustness to variations in input data, its sensitivity to noise, and its generalization capabilities. When applicable and not confidential, details regarding the training parameters used—such as learning rates, batch sizes, and training duration—can be included to provide a more complete understanding of the model's development. This section is vital for users to understand the underlying ML model and its suitability for their specific use cases.

D. ML Sensor Datasheet Components

As IoT proliferation increases e-waste, consumers must be informed of devices' environmental sustainability. Similarly, as these devices are deployed in increasingly diverse environments, users must be able to understand how their device's end-to-end performance will vary. As such, this section outlines the envrionmental impact and end-to-end performance of deploying these devices across real-world settings, enabling the full embodied device to be evaluated by end-users.

Environmental Impact. *How does the device affect the environment during its lifetime?* The Environmental impact component in the ML sensor datasheet provides insights into the sustainability and ecological footprint of ML sensors across different tasks. As ML sensors become widely deployed in various industries, transparency into their environmental impact is crucial for responsible design and usage.

The key element for environmental impact is *Emissions and Carbon Footprint*. To calculate the emissions and carbon footprint of an ML sensor, it's essential to assess its environmental impact throughout its lifecycle—from production and usage to disposal. This begins with gathering the data on material sourcing and manufacturing processes, identifying the materials used (e.g., metals, plastics, semiconductors) and obtaining emissions data associated with their extraction, processing, and assembly. Life cycle assessment (LCA) simulators like the OpenLCA [38] and the TinyML footprint calculator [15] can be used to compute comprehensive modeling of a product's lifecycle emissions.

Other elements within the environmental impact section of the datasheet can include *materials and resource use* and *end-of-life disposal*. *Materials and resource use* list the materials used in each sensor's construction, highlighting any eco-friendly or sustainably sourced components, which is especially important for biometric and environmental sensors in healthcare and monitoring applications that operate in sensitive environments. *End-of-life disposal* instructions specify disposal instructions, recyclability of components, and any options for refurbishment or recycling, helping minimize environmental impact after the sensor reaches end-of-life.

End-to-End Performance Analysis. *How does the device perform as a whole with changing environmental parameters?* This section provides an encompassing evaluation of the sensor's performance from data acquisition to data processing and output generation. This holistic performance analysis may include metrics such as data collection rate, latency, power consumption, and accuracy of the sensor's outputs under a range of conditions. Additionally, it highlights the performance of the ML model when deployed on the sensor hardware, taking into account aspects such as data preprocessing, inference speed, and model accuracy. The analysis could also encompass how the sensor's performance scales with changes in workload or environmental conditions.

This section is important as it helps potential users understand not only the isolated performance of the sensor's components but also how they work together to provide a coherent service. This understanding is vital when integrating the sensor into larger systems or evaluating its fit for particular use-cases. Common factors in end-to-end performance analysis for various ML sensor tasks are demographic bias, distance, class diversity, and environmental conditions.

Performance metrics under these conditions include ML performance metrics such as classification accuracy, detection precision, and recognition accuracy, as well as hardware-specific measurements like power consumption, latency, and throughput. Depending on the intended application, ML sensor manufacturers may adapt these metrics to provide a customized end-to-end performance analysis. This structured approach, combined with adherence to established standards, ensures that the end-to-end performance analysis provides meaningful insights into the sensor's real-world reliability and efficiency.

V. A Case Study in Person Detection

In this section, we present a detailed case study of two person detection ML sensors (shown previously in Figure 2), which are used in applications such as smart farming and occupancy monitoring. These sensors detect the presence of a person without identifying individuals or performing facial recognition. The first is an open-source ML sensor, developed as part of this research initiative. The second is an ML sensor

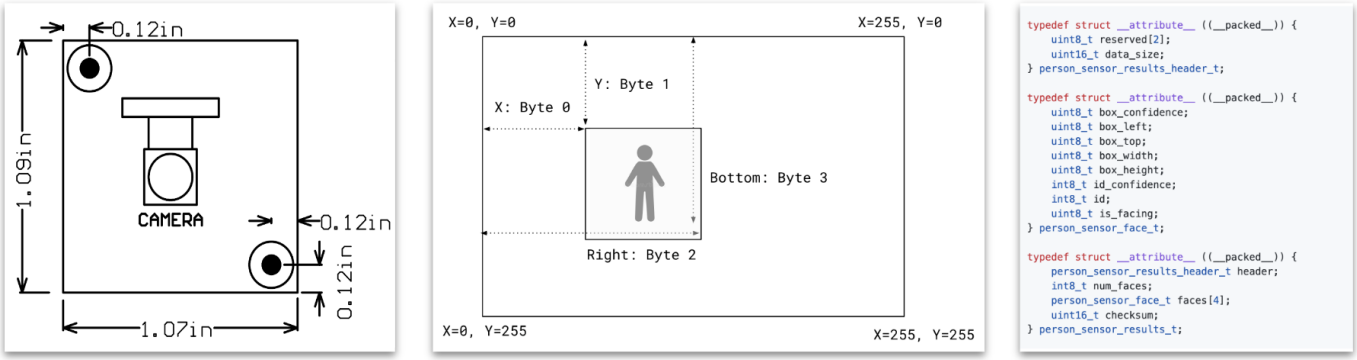

Fig. 4: *(left)* Device diagram of person detection ML sensor, *(middle)* standard for data communication, and *(right)* schema for communication of data off-sensor.

developed by a commercial partner. This partnership allowed us to integrate academic research insights with practical, industry-standard design approaches. The result is a sensor that not only has commercial viability but also achieves increased auditability and transparency. The juxtaposition of these two sensors serves both provides a comprehensive overview of the current state of person detection technology, as well as offers insights into the different design and development methodologies employed in academic versus commercial settings. Finally, this case study aims to highlight the potential synergies between academic research and industry practices, suggesting a model for future collaborative efforts in the fields of ML and sensor technology. We *open-source* the resulting datasheets at: github.com/harvard-edge/ML-Sensors

A. Standard Datasheet Components

In the context of our person detection sensor (Figure 2), the description would be "a device that predicts whether an individual is present in the view of the camera and outputs a corresponding signal response." Examples of diagrams and hardware specifications for our open-source person detector are shown in Figure 4. Figure 4 *(left)* shows our ML sensor with a square form factor and dimensions 27.2mm $\times$ 27.7 mm. It employs the industry-standard Inter-Integrated Circuit (I^2C) interface via a Qwiic connector [39], allowing a data transfer rate of up to 100 kB/s. Figure 4 *(middle)* and *(right)* show the data standard and open-source schema we developed for communication [40]. The sensor communicates through a single byte with values from 0 to 255. The device can accept power at 3.5-5.5 Vs with a 40 mA operating current.

Additionally, while our own ML sensor has not obtained specific certifications or verification, the commercial sensors complies with RoHS [41] for environmental safety and GDPR [42] for data privacy and has been audited by Kodelski Security [43] for security and privacy implications.

B. IoT Datasheet Components

For our ML sensor, the IoT security and privacy label (see Figure 5 left) shows that there is only a camera on the device collecting data continuously, but this data is not being stored or transmitted off-device. The self-contained nature of the ML sensors means that they have limited networking capabilities, and thus privacy concerns from the transmission of raw data are minimal. We note that the commercial sensor has an equivalent IoT privacy and security label.

C. AI Datasheet Components

To evaluate the dataset used in training the on-device model, we utilize the second-generation Dataset Nutrition Label [12], [13]. Summary statistics for the data nutrition label for the open-source person detection sensor are shown on the right side of Figure 5, and the full label is available in our open-sourced datasheet. This label highlights that the dataset, Visual Wake Words [44], is from an upstream source (MS-COCO [45]), contains information about humans obtained without consent, and that the dataset is not currently managed or updated. Figure 5 also shows an example model characteristic of the ML sensor running a MobileNetV1 architecture [46] trained for person detection. In particular, the ROC curve shows that the optimal threshold value lies around 0.52 to balance false positives and negatives. The commercial sensor has a similar, but more complex, software architecture, resulting in a 91.8% accuracy with a threshold of 0.7.

D. ML Sensor Datasheet Components

Environmental Footprint. We captured the carbon footprint of our ML sensor using the methodology in [15]. The calculator includes fields for processing, sensing, power supply, memory, and more, enabling us to input specifications from our bill of materials. Furthermore, we also capture the carbon footprint for the ML sensor's model training, transport, and three-year use. While training costs can be amortized over multiple sensor deployments, we consider them separately to provide a conservative carbon footprint estimate. The total footprint of our ML sensor, including embodied and operational carbon, is approximately 2.34 kg CO_2-eq. Figure 6 (left) shows that the majority of the footprint is attributable to the power supply and camera sensor. We note that other environmental impact indicators beyond carbon footprint should also be included in future datasheets.

End-to-End Performance Analysis: We present an example of end-to-end performance analysis on our open-source person

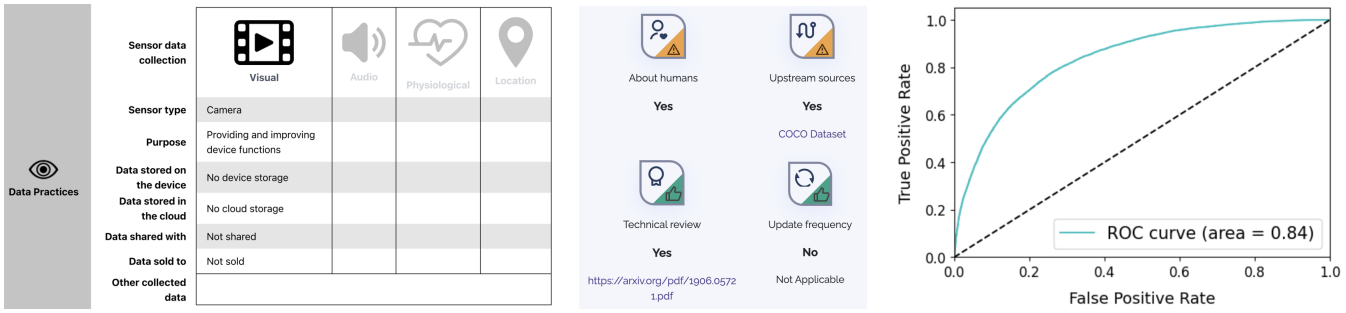

Fig. 5: Primary IoT security and privacy label for the open-source person detection ML sensor *(left)*, as well as its data nutrition label summary statistics *(center)*, and the ROC curve of the person detection model evaluated on a test set *(right)*.

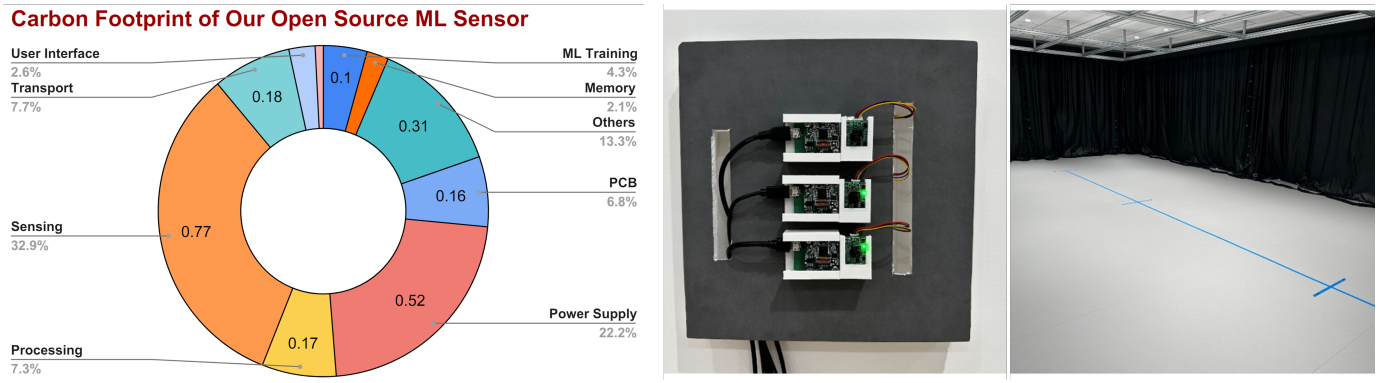

Fig. 6: *(Left)* Breakdown by component of the 2.34 kg CO_2-eq carbon footprint of our ML sensor, using the TinyML Footprint Calculator [15]. *(Center)* The wall-mounted sensor assembly, consisting of sensors developed in-house on the left and those provided by a commercial partner on the right. *(Right)* The experimental environment with 1m, 3m, and 5m distances marked.

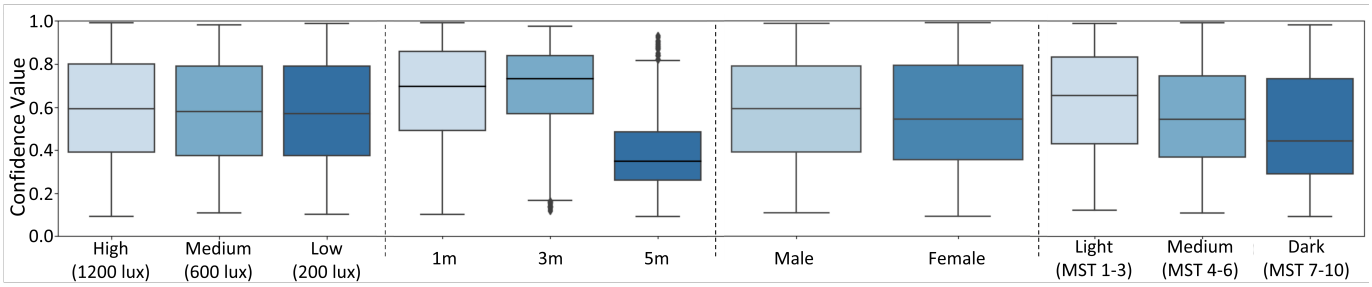

Fig. 7: End-to-end performance analysis of the ML sensor tested on 38 volunteers under controlled laboratory conditions. Confidence across lighting conditions *(far left)*, distances *(center left)*, gender *(center right)*, and skin tone *(far right)* estimated using the Monk Skin Tone (MST) Scale [47].

detection sensor in Figure 7. Such an example study was deemed necessary to assess sensor performance in a deployment environment to determine the extent of dataset shift resulting from the use of different hardware (i.e., the onboard camera), embedded demographic biases, as well as biases from environmental changes (e.g., lighting and distance from the camera).

The study room measured 25 x 31 x 10 ft and contained 32 ceiling lights that were uniformly distributed in a 4 x 8 grid. The lighting conditions were captured quantitatively for each participant using a Lux LCD Illuminance Meter (Precision Vision, Inc.) and a C-800-U Spectrometer (Sekonic Corporation). Sensors were mounted on a wooden board affixed to the wall at a height of 1.5 m above the ground. 39 participants were evaluated at three different distances (1 m, 3 m, and 5 m marked with colored tape) under three lighting conditions (208±31, 584±51, and 1149±59 lux controlled by a dimmer switch). The ambient lighting in the room was provided by artificial lights, and blackout curtains were used to block the ambient lighting from outside (Figure 6 center and right). When all the lights were turned on at full strength, the sensor gave an average reading of 1149 lux. The color temperature of the lighting was measured to be 5600 K.

Participants were asked to provide their gender identity and evaluate their skin tone according to the Monk Skin Tone (MST) Scale. The study evaluated algorithmic bias by bucketing skin tone into three categories: light (MST 0-4), medium (MST 5-7), and dark (MST 8-10). Ten readings from each sensor were averaged at each location and lighting condition. This anonymous study was approved by the Institutional Review Board of Harvard University on 6 April 2023 (Project Code: IRB23-0136).

The accuracy of the model (see Figure 7) is provided as a function of lighting condition, distance, gender identity, and skin tone. We note that overall, 63.2% of the participants were male, and 36.8% were female; the percentage of participants corresponding to each skin tone group was: 47.4% light, 39.4% medium, and 13.2% dark. These analysis provide examples of both device efficacy under changing environmental conditions, a common type of analysis on standard sensor datasheets, as well as possible demographic biases within the model (see Figure 7).

Our performance analysis revealed that while lighting conditions had minimal impact, the model's accuracy degraded sharply when the distance between the subject and the sensor increased from 3 to 5 meters. The model also exhibited potential demographic biases, performing slightly better on men than women and favoring lighter skin tones over darker ones. The diversity of clothing worn by study participants was not fully captured in our testing data, potentially affecting the results. In contrast, the commercial sensor showed increased overall performance and decreased bias, likely due to its more robust custom dataset and advanced software architecture. These findings highlight the importance of considering the sensor's effective range, ensuring diversity and inclusion in training data, actively monitoring and mitigating biases, and investing in high-quality datasets and sophisticated model architectures to enhance accuracy and fairness in ML sensor systems.

VI. Discussion and Limitations

Our datasheet template, as proposed, consolidates multiple areas of critical sensor information and finds relevance in numerous practical applications, including predictive maintenance in industrial settings [48], environmental monitoring [49], healthcare diagnostics [50], autonomous vehicles [51], and smart homes [52]. By detailing the hardware characteristics and conformity with industry and regulatory standards, our datasheet provides developers and users with a dependable tool to assess sensor suitability for their specific use-cases. In this section, we discuss the generalizability and limitations of our approach. Overall, we find that our high-level template can be easily adapted for a wide range of current and future applications but additional development is needed to specify detailed metrics and domain-specific requirements.

Open-Source vs. Commercial Comparison. At a high-level, the datasheet template was found to be applicable for both our open-source sensor, as well as the commercial sensor, with changes only necessary in a limited number of sections. Sections where changes were necessary were mainly in the data nutrition label and the model characteristics in order to obfuscate aspects of the commercial partner's intellectual property, such as proprietary datasets, models, and training procedures. This obfuscation was critical to enable industrial collaboration and care will need to be taken in the future to ensure that the level of obfuscation balances transparency and intellectual property.

Extendability to Varied Data Modalities. The structure of the datasheet remains consistent across different modalities, ensuring a familiar framework for assessing the diverse aspects of ML sensors. For instance, when applied to event cameras used in VR/AR [53], [54], while event-based cameras have different properties than CMOS cameras and utilize alternative approaches such as spiking networks over convolutional networks, datasheets for sensors using either camera type will retain similar sections such as optical properties of the camera and the network training process. Similarly, when applied to audio data, instead of detailing the optical properties of the camera, the datasheet would detail the acoustic properties relevant to the microphone, such as sensitivity ranges, signal-to-noise ratios, and the types of audio processing algorithms used. This adaptability is also true for various model architectures ranging from the CNNs described in our case study to those implementing more basic neural network operations [55]. In all cases, we will still need to document both the dataset, the resulting metrics for the trained model, and end-to-end metrics. As such, while incremental refinement might be needed over time, we believe that our datasheet template, regardless of sensor or model configuration can retain the same structure and sections, and simply adjust the details to support the relevant metrics for the specific device.

Reliance on Self-Reporting. A key limitation of our approach is that the datasheet relies on the accuracy and honesty of the information provided by the manufacturers or developers, with the potential risk of misinformation, misinterpretation, or lack of updates to the datasheet after product updates. Transparency for audit is most useful within relationships of accountability [56], and, as mentioned in Section IV-A, oversight mechanisms such as certification from a trusted third-party entity could resolve this concern, and the use of blockchain technologies could aid in auditability [57].

Cross-Compatibility of Metrics. Our case study provides an important example of another key limitation of our datasheet which is cross-compatibility and standardization of metrics across a wide possible breadth of ML Sensor device designs and implementations. In fact, we found it challenging to directly compare end-to-end results from the commercial and open-source devices due to the differing approaches utilized by the sensors. The commercial sensor utilized a face detection bounding-box model with a detection threshold set at $\sim$0.6, whereas our open-source sensor focused on person detection within the full image. This along with differences in camera specification meant that the open-source device was better at detecting individuals over longer distances, while the commercial sensor had a wider angle of detection. This suggests that future research is needed to design and build methods to fairly compare and evaluate ML sensors as their diversity grows. However, we believe that ensuring that devices are bundled with relevant, even if not cross-compatible, end-to-end metrics is a critical first step to ensuring transparency and auditability for these devices.

VII. Conclusion and Future Work

This paper introduces a novel datasheet template designed specifically for ML sensors to provide meaningful ML sensor transparency. We focus on embedded ML deployments, where models are tightly integrated and fused with devices, warranting

closer introspection in this new space. We demonstrate the practical application of our datasheets by developing them for two real-world ML sensors, enhancing transparency, auditability, and user-friendliness across open-source and commercial devices. By providing a standardized format for documenting and evaluating ML sensors, our work contributes to the advancement of data-centric ML research and the creation of reliable benchmarks for next-generation ML that is deeply fused in with embedded and edge AI systems.

In future work we hope to expand upon our initial case study and provide formal documentation to help others build their own datasheets for their own ML Sensors. We also intend to perform large-scale validation of our datasheet template design, approach, layout, and content to ensure its effectiveness and generality. In particular, large-scale user studies across academia and industry can help analyze the effectiveness of the current template. To that end, future work will also study the effectiveness of the datasheet as a component of audit, to better understand design changes that could promote transparency for auditability across stakeholder groups. Similarly, large-scale design studies to understand what forms of visualizations, and organization of the data contained in these datasheets would make them most effective in conveying their information would further improve the impact of this work. Finally, we hope to find ways to build on recent work (e.g., Croissant-RAI [58]) to ensure that future datasheets are both human and machine readable to increase their impact and usability for a wider range of users and applications.

REFERENCES

[1] P. Yadav, "Advancements in machine learning in sensor systems: Insights from sensys-ml and tinyml communities," in *2024 IEEE 3rd Workshop on Machine Learning on Edge in Sensor Systems (SenSys-ML)*, pp. 21–26, IEEE, 2024.

[2] P. Warden, M. Stewart, B. Plancher, S. Katti, and V. J. Reddi, "Machine learning sensors: A design paradigm for the future of intelligent sensors," *Communications of the ACM*, vol. 66, p. 25–28, Nov 2023.

[3] M. Z. Uddin and A. Soylu, "Human activity recognition using wearable sensors, discriminant analysis, and long short-term memory-based neural structured learning," *Scientific Reports*, vol. 11, no. 1, p. 16455, 2021.

[4] H. Yang, J. Li, X. Xiao, J. Wang, Y. Li, K. Li, Z. Li, H. Yang, Q. Wang, J. Yang, *et al.*, "Topographic design in wearable mxene sensors with in-sensor machine learning for full-body avatar reconstruction," *Nature communications*, vol. 13, no. 1, p. 5311, 2022.

[5] D. Hartmann, J. R. L. de Pereira, C. Streitbörger, and B. Berendt, "Addressing the regulatory gap: moving towards an eu ai audit ecosystem beyond the ai act by including civil society," *AI and Ethics*, pp. 1–22, 2024.

[6] J. Mökander, M. Axente, F. Casolari, and L. Floridi, "Conformity assessments and post-market monitoring: a guide to the role of auditing in the proposed european ai regulation," *Minds and Machines*, vol. 32, no. 2, pp. 241–268, 2022.

[7] G. Falco, B. Shneiderman, J. Badger, R. Carrier, A. Dahbura, D. Danks, M. Eling, A. Goodloe, J. Gupta, C. Hart, *et al.*, "Governing ai safety through independent audits," *Nature Machine Intelligence*, vol. 3, no. 7, pp. 566–571, 2021.

[8] E. Commission, "AI Act," Mar. 2024.

[9] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, "Model cards for model reporting," in *Proceedings of the Conference on Fairness, Accountability, and Transparency*, ACM, jan 2019.

[10] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi, "Ask the experts: What should be on an iot privacy and security label?," in *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 447–464, IEEE, 2020.

[11] P. Emami-Naeini, J. Dheenadhayalan, Y. Agarwal, and L. F. Cranor, "An informative security and privacy "nutrition" label for internet of things devices," *IEEE Security & Privacy*, vol. 20, no. 2, pp. 31–39, 2021.

[12] S. Holland, A. Hosny, S. Newman, J. Joseph, and K. Chmielinski, "The dataset nutrition label: A framework to drive higher data quality standards," *arXiv preprint arXiv:1805.03677*, 2018.

[13] K. S. Chmielinski, S. Newman, M. Taylor, J. Joseph, K. Thomas, J. Yurkofsky, and Y. C. Qiu, "The dataset nutrition label (2nd gen): Leveraging context to mitigate harms in artificial intelligence," 2022.

[14] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. I. au2, and K. Crawford, "Datasheets for datasets," 2021.

[15] S. Prakash, M. Stewart, C. Banbury, M. Mazumder, P. Warden, B. Plancher, and V. J. Reddi, "Is tinyml sustainable? assessing the environmental impacts of machine learning on microcontrollers," 2023.

[16] M. Sloane, E. Moss, S. Kennedy, M. Stewart, P. Warden, B. Plancher, and V. J. Reddi, "Materiality and risk in the age of pervasive ai sensors," *Nature Machine Intelligence*, pp. 1–12, 2025.

[17] Seeed Studio, *SenseCAPA1101 LoRaWAN Vision AI Sensor User Guide*, 2022. Version v1.0.5.

[18] U. Sensors, "Person sensor developer guide." usfl.ink/ps_dev, 2023.

[19] J. Bandy and N. Vincent, "Addressing" documentation debt" in machine learning: A retrospective datasheet for bookcorpus," 2021.

[20] M. Zilka, B. Butcher, and A. Weller, "A survey and datasheet repository of publicly available us criminal justice datasets," *Advances in Neural Information Processing Systems*, vol. 35, pp. 28008–28022, 2022.

[21] R. Srinivasan, E. Denton, J. Famularo, N. Rostamzadeh, F. Diaz, and B. Coleman, "Artsheets for art datasets," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.

[22] M. Arnold, R. K. Bellamy, M. Hind, S. Houde, S. Mehta, A. Mojsilović, R. Nair, K. N. Ramamurthy, A. Olteanu, D. Piorkowski, *et al.*, "Factsheets: Increasing trust in ai services through supplier's declarations of conformity," *IBM Journal of Research and Development*, vol. 63, no. 4/5, pp. 6–1, 2019.

[23] K. L. Boyd, "Datasheets for datasets help ml engineers notice and understand ethical issues in training data," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW2, pp. 1–27, 2021.

[24] R. Owen, J. Stilgoe, P. Macnaghten, M. Gorman, E. Fisher, and D. Guston, "A framework for responsible innovation," *Responsible innovation: managing the responsible emergence of science and innovation in society*, pp. 27–50, 2013.

[25] U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, and C.-J. Wu, "Act: Designing sustainable computer systems with an architectural carbon modeling tool," in *Proceedings of the 49th Annual International Symposium on Computer Architecture*, pp. 784–799, 2022.

[26] "General data protection regulation (gdpr)." Regulation (EU) 2016/679, 2016.

[27] "TinyML Summit 2023." https://www.edgeaifoundation.org/events/summit-2023.

[28] National Institute of Standards and Technology, "Nist - national institute of standards and technology," 2023. Accessed: 2023-11-05.

[29] E. Parliament *et al.*, "Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation)," *Official Journal of the European Union L*, vol. 119, no. 1, 2016.

[30] A. Act, "Health insurance portability and accountability act of 1996," *Public law*, vol. 104, p. 191, 1996.

[31] ISO, "26262: 2018:"road vehicles—functional safety"," *British Standards Institute*, vol. 12, 2018.

[32] M. Aljabri, A. A. Alahmadi, R. M. A. Mohammad, F. Alhaidari, M. Aboulnour, D. M. Alomari, and S. Mirza, "Machine learning-based detection for unauthorized access to iot devices," *Journal of Sensor and Actuator Networks*, vol. 12, no. 2, p. 27, 2023.

[33] A. De Luca, E. Von Zezschwitz, N. D. H. Nguyen, M.-E. Maurer, E. Rubegni, M. P. Scipioni, and M. Langheinrich, "Back-of-device authentication on smartphones," in *Proceedings of the sigchi conference on human factors in computing systems*, pp. 2389–2398, 2013.

[34] T. Talaei Khoei, S. Ismail, and N. Kaabouch, "Dynamic selection techniques for detecting gps spoofing attacks on uavs," *Sensors*, vol. 22, no. 2, p. 662, 2022.

[35] J. Huckelberry, Y. Zhang, A. Sansone, J. Mickens, P. A. Beerel, and V. J. Reddi, "Tinyml security: Exploring vulnerabilities in resource-constrained machine learning systems," *arXiv preprint arXiv:2411.07114*, 2024.

[36] Y. Zhang, D. Chen, S. Kundu, C. Li, and P. A. Beerel, "Sal-vit: Towards latency efficient private inference on vit using selective attention search with a learnable softmax approximation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5116–5125, 2023.

[37] H. Jiang and O. Nachum, "Identifying and correcting label bias in machine learning," in *International Conference on Artificial Intelligence and Statistics*, pp. 702–712, PMLR, 2020.

[38] "Openlca." GreenDelta. Available at: https://www.openlca.org/.

[39] "Qwiic connect system - sparkfun electronics." https://www.sparkfun.com/qwiic. (Accessed on 06/06/2023).

[40] "Github - usefulsensors/person_sensor_docs: Documentation for the person sensor." https://github.com/usefulsensors/person_sensor_docs. (Accessed on 06/06/2023).

[41] "Restriction of hazardous substances (rohs)." European Union Directive 2011/65/EU, 2011.

[42] EU, "Gdpr," 2024.

[43] Kudelski Security, 2024.

[44] A. Chowdhery, P. Warden, J. Shlens, A. Howard, and R. Rhodes, "Visual wake words dataset," *arXiv preprint arXiv:1906.05721*, 2019.

[45] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár, "Microsoft coco: Common objects in context," 2015.

[46] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.

[47] T. Doshi, "Improving skin tone representation across google." https://blog.google/products/search/monk-skin-tone-scale/, 5 2022. (Accessed on 06/07/2023).

[48] E. Njor, J. Madsen, and X. Fafoutis, "A primer for tinyml predictive maintenance: Input and model optimisation," in *Artificial Intelligence Applications and Innovations: 18th IFIP WG 12.5 International Conference, AIAI 2022, Hersonissos, Crete, Greece, June 17–20, 2022, Proceedings, Part II*, pp. 67–78, Springer, 2022.

[49] A. Gkogkidis, V. Tsoukas, S. Papafotikas, E. Boumpa, and A. Kakarountas, "A tinyml-based system for gas leakage detection," in *2022 11th International Conference on Modern Circuits and Systems Technologies (MOCAST)*, pp. 1–5, IEEE, 2022.

[50] V. Tsoukas, E. Boumpa, G. Giannakas, and A. Kakarountas, "A review of machine learning and tinyml in healthcare," in *25th Pan-Hellenic Conference on Informatics*, pp. 69–73, 2021.

[51] M. de Prado, M. Rusci, A. Capotondi, R. Donze, L. Benini, and N. Pazos, "Robustifying the deployment of tinyml models for autonomous mini-vehicles," *Sensors*, vol. 21, no. 4, p. 1339, 2021.

[52] A. Zacharia, D. Zacharia, A. Karras, C. Karras, I. Giannoukou, K. C. Giotopoulos, and S. Sioutas, "An intelligent microprocessor integrating tinyml in smart hotels for rapid accident prevention," in *2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*, pp. 1–7, IEEE, 2022.

[53] A. N. Angelopoulos, J. N. Martel, A. P. Kohli, J. Conradt, and G. Wetzstein, "Event based, near eye gaze tracking beyond 10,000 hz," *arXiv preprint arXiv:2004.03577*, 2020.

[54] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, *et al.*, "Event-based vision: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 1, pp. 154–180, 2020.

[55] T. Ma, Y. Feng, X. Zhang, and Y. Zhu, "Camj: Enabling system-level energy modeling and architectural exploration for in-sensor visual computing," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, pp. 1–14, 2023.

[56] A. F. Cooper, E. Moss, B. Laufer, and H. Nissenbaum, "Accountability in an algorithmic society: relationality, responsibility, and robustness in machine learning," in *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, pp. 864–876, 2022.

[57] O. Ali, A. Jaradat, A. Kulakli, and A. Abuhalimeh, "A comparative study: Blockchain technology utilization benefits, challenges and functionalities," *Ieee Access*, vol. 9, pp. 12730–12749, 2021.

[58] N. Jain, M. Akhtar, J. Giner-Miguelez, R. Shinde, J. Vanschoren, S. Vogler, S. Goswami, Y. Rao, T. Santos, L. Oala, *et al.*, "A standardized machine-readable dataset documentation format for responsible ai," *arXiv preprint arXiv:2407.16883*, 2024.

Appendix A. Example Data Sheet - Open-Source Sensor

OVERVIEW

PA1 Person Detection Module

Compliance and Certifications

The person detection sensor complies with essential industry standards and regulations, including RoHS for environmental safety and GDPR for protecting individual privacy. As of the time of writing, the sensor does not have any certifications from third-party organizations.

Description

The PA1 Person Detection Module is a cost-effective device that uses a machine learning (ML) algorithm to detect the presence of a person within its range. The sensor is equipped with cameras and sensors that capture images and data from the surrounding environment. These images and data are then processed by the on-device ML algorithm to identify people. When a person is detected, the sensor sends an alert or trigger to connected devices or systems, allowing them to perform specific actions such as activating security cameras, turning on lights, or opening doors. The person detection sensor is ideal for use in security, home automation, and other applications that require quick and accurate detection of people.

The sensor has a small form factor and utilizes a monochrome camera with a field of view of 320 x 320 (QVGA). The sensor is equipped with an onboard 3.3V regulator, which enables it to operate with an input voltage range of 3.5V - 5.5V when enabled, or 3.0V - 3.6V when disabled. The typical operating current for the sensor is 40 mA. The sensor communicates via I2C/Qwiic mode, conforming to SparkFun Qwiic electrical/mechanical specifications, and has a maximum cable length of 1 m. The sensor has a maximum data rate of 100 kb/s and a wide sensitivity coverage of 0.1 - 10 klux.

Features

- Real-time person detection with on-device ML
- Indoor and outdoor use
- Low power consumption
- Onboard camera
- Small form factor: 10 x 10 x 2 mm
- I2C serial communication
- Wide sensitivity coverage: 0.1 - 10 klux

Use Cases

- Security
- Home automation
- Consumer appliances

MODEL CHARACTERISTICS

Software Flow Diagram

Grayscale images (320x320) are collected and resized to 96x96 via bilinear interpolation. Images are fed into a MobileNetV1 architecture trained and optimized through Edge Impulse. The output probability is communicated via Qwiic interface to the application processor.

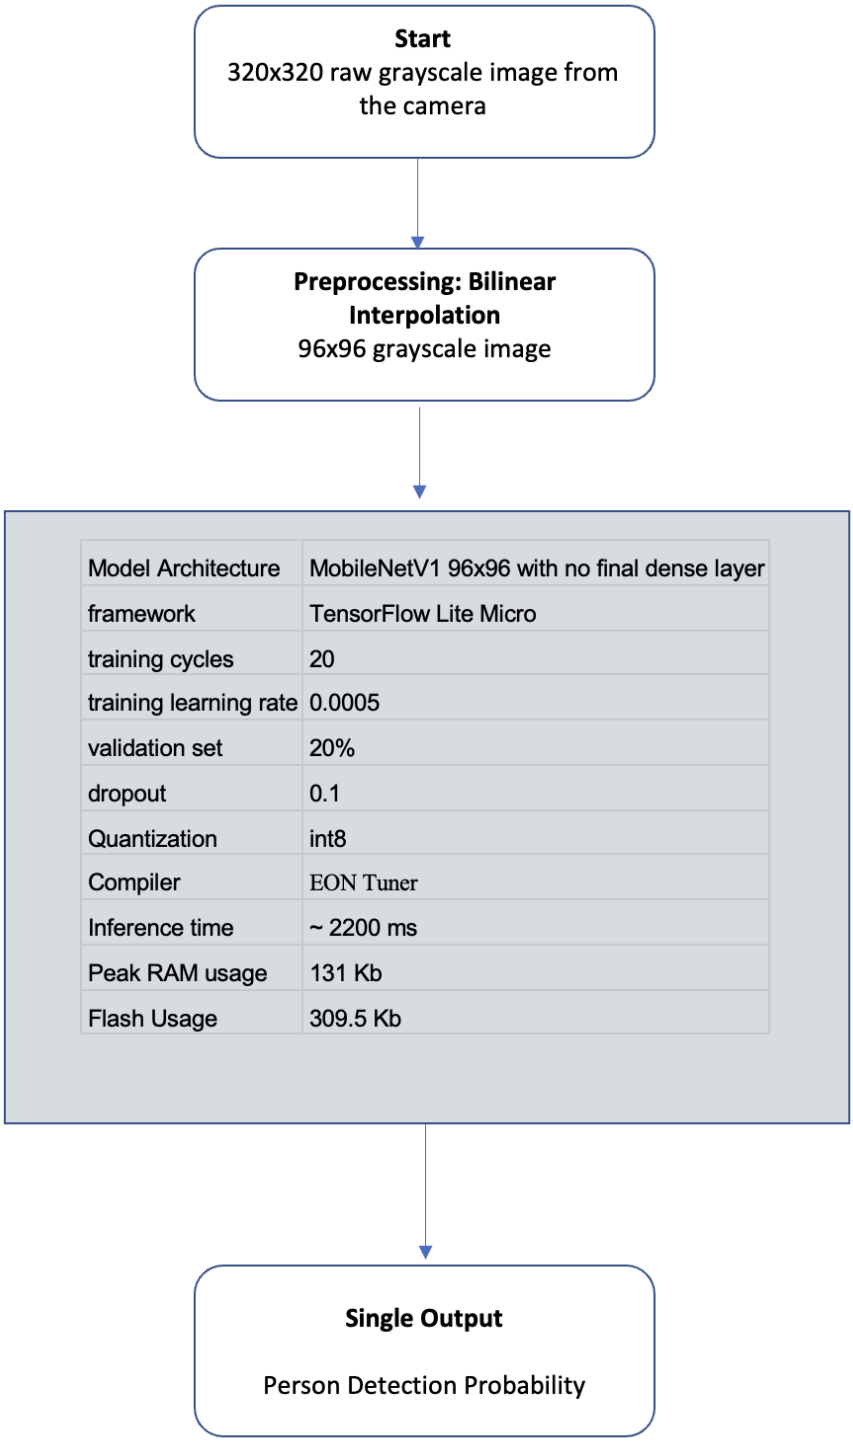

Dataset Nutrition Label

The data nutrition label is publicly available here, with some important features outlined below.

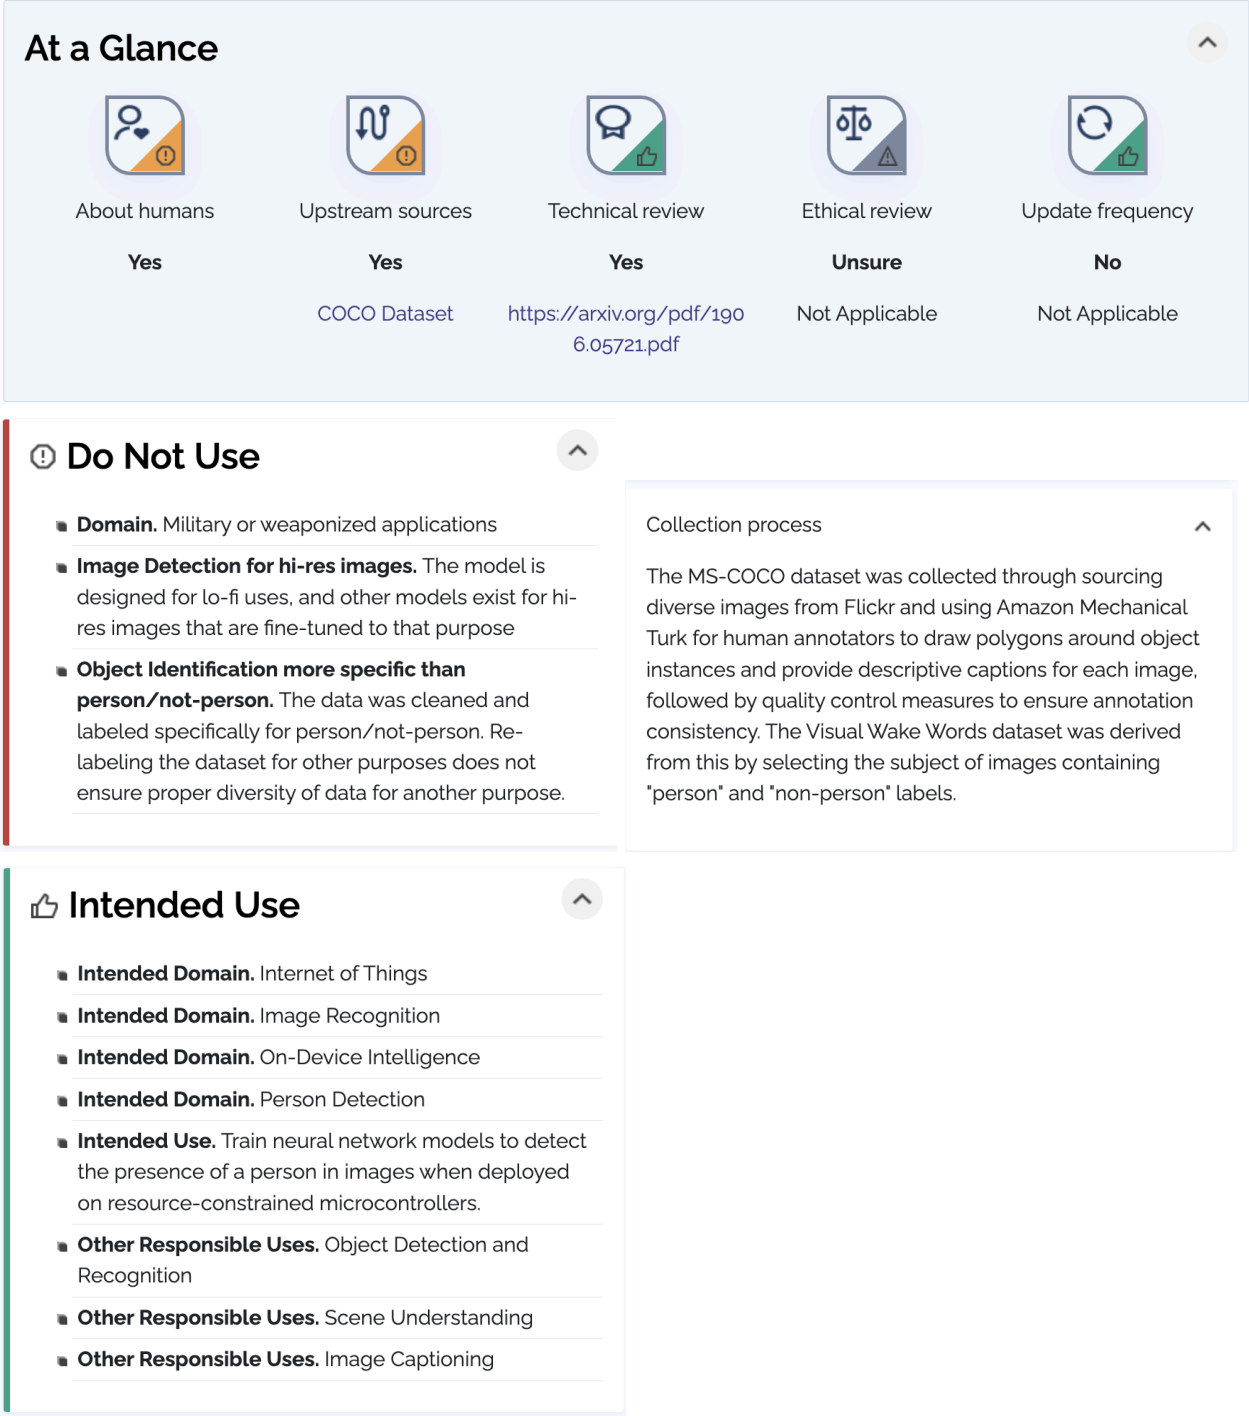

General risks

Any additional risks?

Individual Information

yes

Consent

Consent was not given.

Generalized Inferences

The original source material, from COCO, is mainly made up of photographs from Flickr, and it's not clear to what extent the users of Flickr are representative of the population at large outside the U.S., for instance.

Generalized Inferences - Mitigation

Identifying a specific use case for models made using this dataset, creating a list of situations in which people would be found for that use case, and then reviewing the base dataset to ensure it has a diversity of images related to the situations you identify (this may be a somewhat manual process).

Sensitive Content

Not Applicable

Documented Known Issues

https://medium.com/@jamie_34747/how-i-found-nearly-300-000-errors-in-ms-coco-79d382edf22b

Other Known Issues

Some items in both the person and non-person categories are known to be mislabeled.

Number of issues

Risky	2
Safe	1
Unknown	4

Feature selection

Which columns were chosen and why?

Cultural or Domain Assumptions

Proxy Characteristics

Planning Representation

Domain Knowledge

Some familiarity with the style of how images are labeled in the COCO datasets would be helpful

Number of issues

Risky	1
Safe	1
Unknown	2

Representation

Which rows were included and why?

Subpopulation Information

Not Applicable

Representation

Unknown

Individual Inferences

Decisions or predictions based on the dataset may not accurately account for individual variations, such as clothing and accessories worn by an individual, and could result in overgeneralized outcomes that don't consider unique circumstances or factors. Additionally, the data may include bias due to its data collection practices which may lead to unfair or discriminatory decisions.

Individual Inferences - Mitigation

Collection Representation

Other Representation Issues

Number of issues

Risky	1
Safe	0
Unknown	5

Data values

What values are in each column?

Collection and Labeling Protocols

The data was generalized from its original description to be that of a "person" or "not person", which required scraping of the original dataset based on search parameters entered by the authors of the dataset. The upstream dataset used Amazon Mechanical Turk workers to label pictures as well, on a custom interface created by the upstream dataset authors.

Data Imputation Protocols

Data Manipulation Protocols

Missing Data

The dataset is derived from MS-COCO and thus contains all items within that dataset that include person and non-person tags.

Raw Data

Number of issues

Risky	1
Safe	0
Unknown	4

IoT Security and Privacy Label

This device contains a camera that takes pictures at 1 s intervals. No other sensory data is collected. Raw data is contained solely within the ML module, with only high-level features transmitted to the main processor (i.e., no image data is accessible by the main processor). This module has no internet connectivity or data storage capacity outside the model and software.

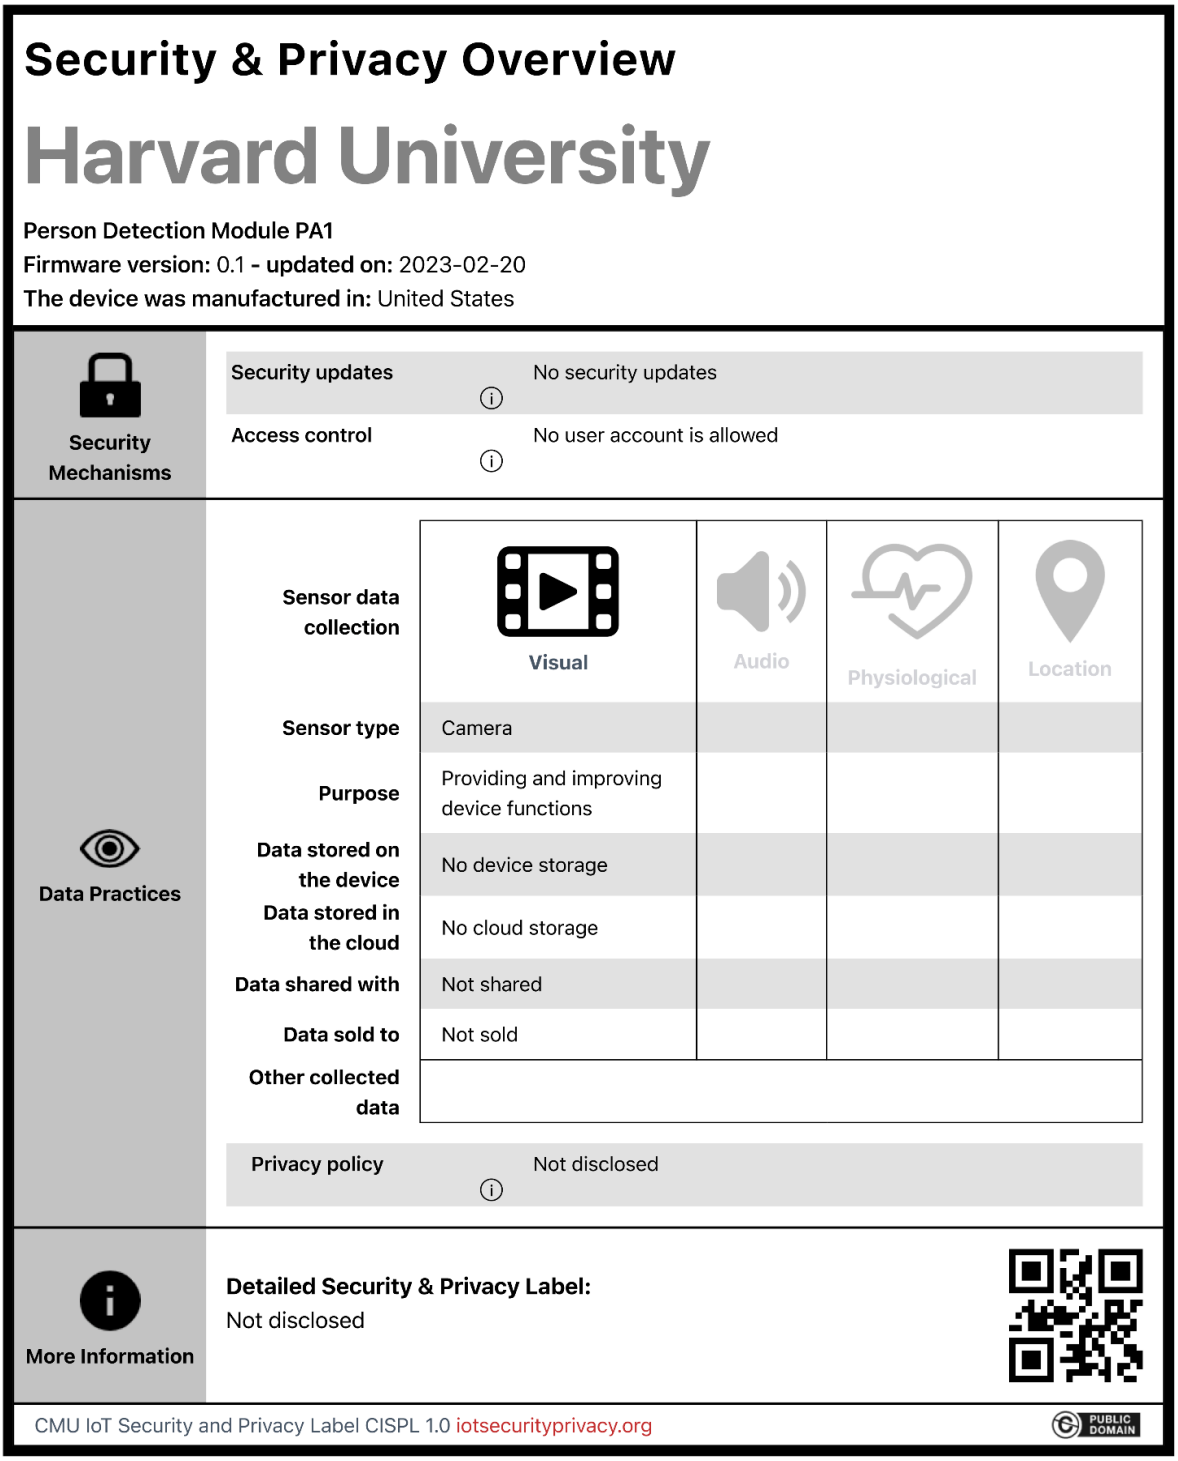

Security & Privacy Details

Harvard University

Person Detection Module PA1
Firmware version: 0.1 - **updated on:** 2023-02-20
The device was manufactured in: United States

Security Mechanisms

Security updates	No security updates
Access control	No user account is allowed.
Security oversight	No security audits
Ports and protocols	Not disclosed
Hardware safety	Not disclosed
Software safety	Not disclosed
Personal safety	Not disclosed
Vulnerability disclosure and management	Not disclosed
Software and hardware composition list	Not disclosed
Encryption and key management	Not disclosed

Data Practices

Sensor data collection	Visual
Sensor type	Camera
Data collection frequency	Continuous
Purpose	Providing and improving device functions
Data stored on the device	No device storage
Local data retention time	No retention
Data stored in the cloud	No cloud storage
Cloud data retention time	No retention
Data shared with	Not shared
Data sharing frequency	Not shared
Data sold to	Not sold
Other collected data	None

Data linkage	Data will not be linked with other data sources
What will be Inferred from User's Data	Presence of a human
Special data handling practices for children	No
In Compliance with	GDPR
Privacy policy	Not disclosed

More Information

Call Harvard University with your questions at	Not disclosed
Email Harvard University with your questions at	ml-sensors@googlegroups.com
Functionality when offline	Full functionality on offline mode
Functionality with no data processing	Not disclosed
Physical actuations and triggers	Device performs customized actions when person is detected.
Compatible platforms	Not disclosed

CMU IoT Security and Privacy Label CISPL 1.0 iotsecurityprivacy.org

PUBLIC DOMAIN

Machine Learning Model Specification

The person detection model was created using transfer learning with the MobileNetV1 neural network (see architecture here) on Edge Impulse. The training and testing of the model were done using a subset of images from the MS-COCO 2017 dataset, which is widely used for image recognition. Only images containing humans were selected from the dataset, totaling 109,604 images. The derived dataset is equivalent to the Visual Wake Words dataset. A train/validation split ratio of 0.8 was used.

The input to the model is a 96x96 raw image in 8-bit grayscale format, equivalent to 9,216 features. The training process was carried out over 20 cycles with a learning rate of 0.0005 and a test set of 20% on MobileNetV1 with a dropout of 0.1 and no final dense layer. The output layer of the model produces a two-class vector of results, indicating the probability of a person being present in the image. The unoptimized (float32) model has an accuracy of 76.3%, with a false positive (FP) and false negative (FN) rate of 20.7% and 26.8%, respectively. The model was quantized to int8 and deployed on Edge Impulse using the integrated EON-Compiler to produce a C++ library. The quantized model has an accuracy of 75.5%, with an FP and FN rate of 23.9% and 25.1%.

To enable live person detection, a set of image provision scripts was added to the software pipeline. The scripts continuously capture data from the onboard camera and pass it to the model in the appropriate scale and format. Using the Arm GNU Toolchain, the Pico-SDK, and the resulting C++ library, the model was built and compiled into a binary file that can be flashed to the ML board [See README/GitHub Repo]. The output of the model is an output vector consisting of a non-person score and a person score, which is communicated through a serial connection and can be viewed on a serial monitor.

Model workflow and characteristics can be viewed through the public Edge Impulse project version here.

(a) Receiver Operating Characteristic Curve

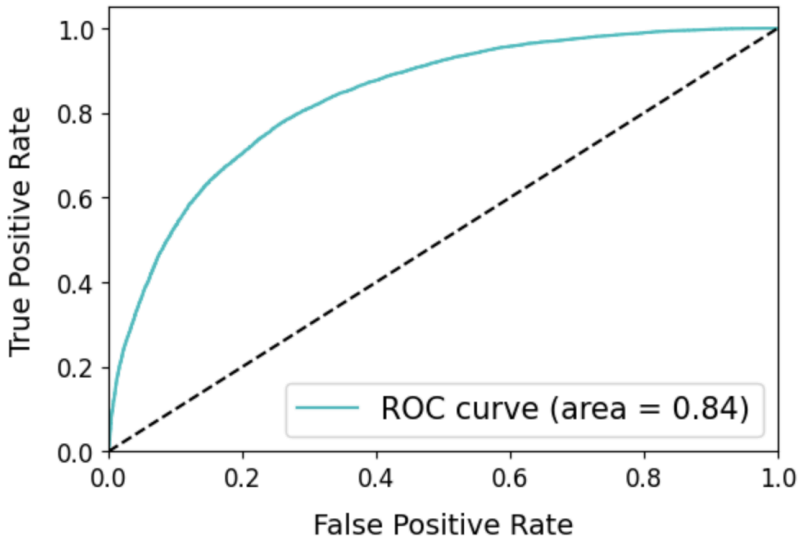

(b) Confusion Matrix

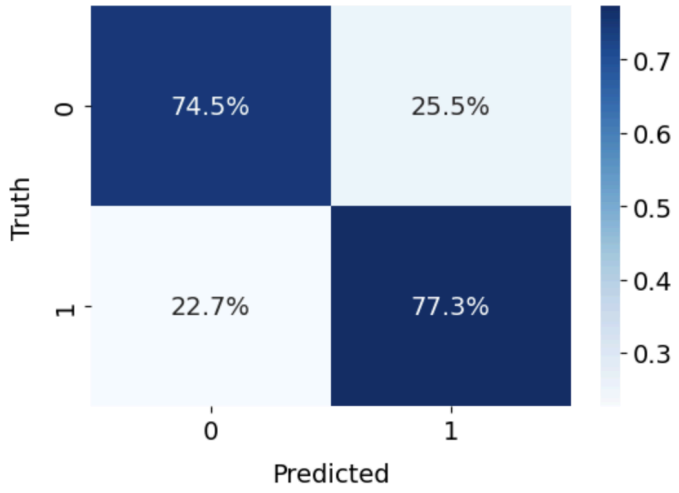

(c) Precision-Recall Curve

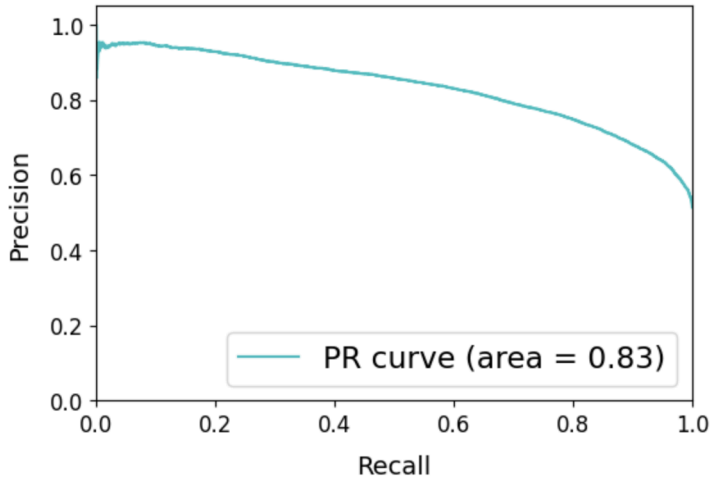

(d) Histogram of Predicted Probabilities

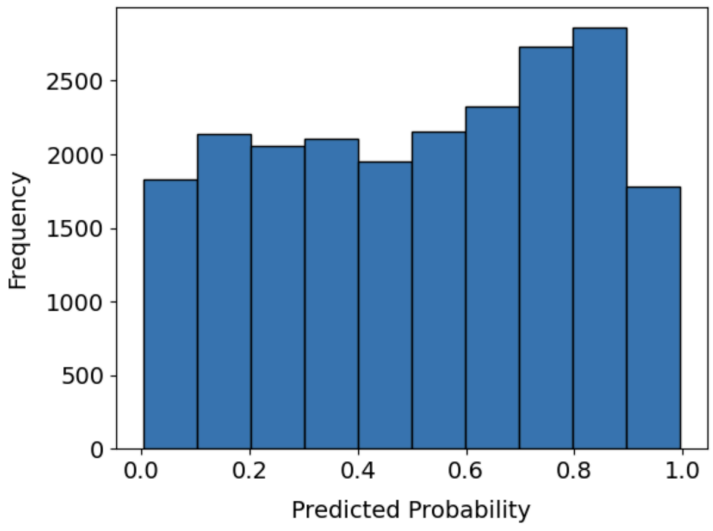

Performance Analysis

The end-to-end performance of the person detection sensor model was tested through an experimental study. The study involved 40 participants and evaluated the accuracy of the model under different lighting conditions using three identical sensors.

The study room measured 25 x 31 x 10 ft and contained 32 ceiling lights that were uniformly distributed in a 4 x 8 grid. The lighting conditions were captured quantitatively for each participant using a Lux LCD Illuminance Meter (Precision Vision, Inc.) and a C-800-U Spectrometer (Sekonic Corporation).

The sensors were mounted on a wooden board affixed to the wall at a height of 1.5 m above the ground. The participants were evaluated at three different distances (1.5 m, 4.5 m, and 7.5 m) from the sensors under each lighting condition. The ambient lighting in the room was provided by artificial lights, and blackout curtains were used to block the ambient lighting from outside.

The lighting levels were controlled using a dimmer switch that had three levels of operation, corresponding to 208±31, 584±51, and 1149±59 lux, respectively. When the lights were turned off, the illuminance meter gave a reading of zero lux. When all the lights were turned on at full strength, the sensor gave an average reading of 1149 lux. The color temperature of the lighting was measured to be 5600 K, corresponding to white light. Colored tape was placed on the ground to demarcate the locations where participants should stand during the experiment (i.e., 1.5, 4.5, and 7.5 m from the sensor array).

Before entering the study environment, the participants were asked to provide their gender identity and evaluate their skin tone according to the Monk Skin Tone (MST) Scale to evaluate algorithmic bias. The study evaluated algorithmic bias by bucketing skin tone into three categories: light (MST 0-4), medium (MST 5-7), and dark (MST 8-10). At each location and lighting condition, ten readings were taken from each sensor and averaged.

Participants were recruited using flyers, and those interested filled out a Study Interest Form. Upon arrival, participants signed a Consent Form indicating their willingness to participate in the study. The accuracy of the model is provided in the following graphs as a function of lighting condition, distance, gender identity, and skin tone. Overall, 63.2% of the participants were male, and 36.8% were female; the percentage of participants corresponding to each skin tone group was: 47.4% light, 39.4% medium, and 13.2% dark.

This anonymous study was approved by the Institutional Review Board of Harvard University on 6 April 2023 (Project Code: IRB23-0136).

Environmental Sensitivity

The device shows a marginal decrease in performance under decreased lighting conditions. A marked drop off in performance is observed at distances 3-5 meters from the sensor.

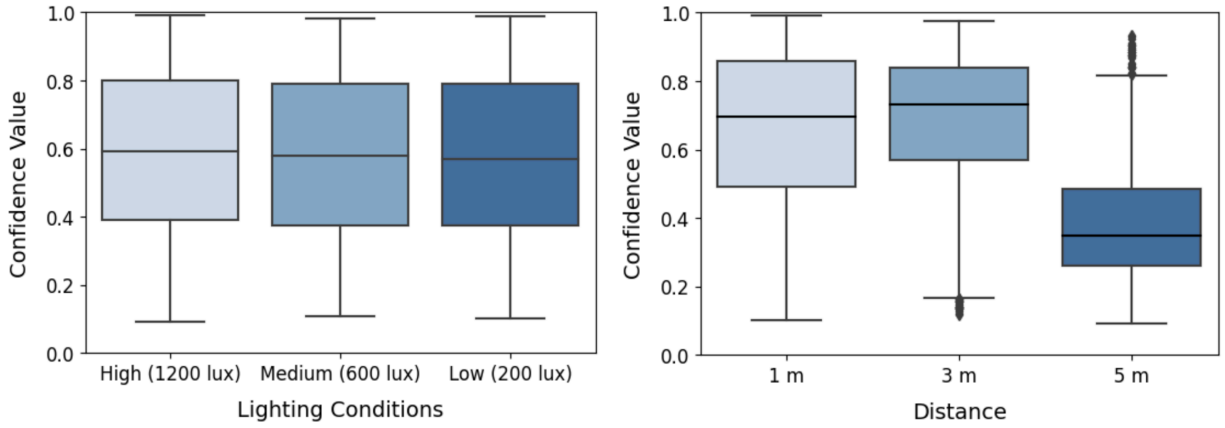

Demographic biases

A small gender bias is observed in model performance. A large skin tone bias was observed, showing approximately a 20% decrease in the confidence value for individuals with a darker skin tone.

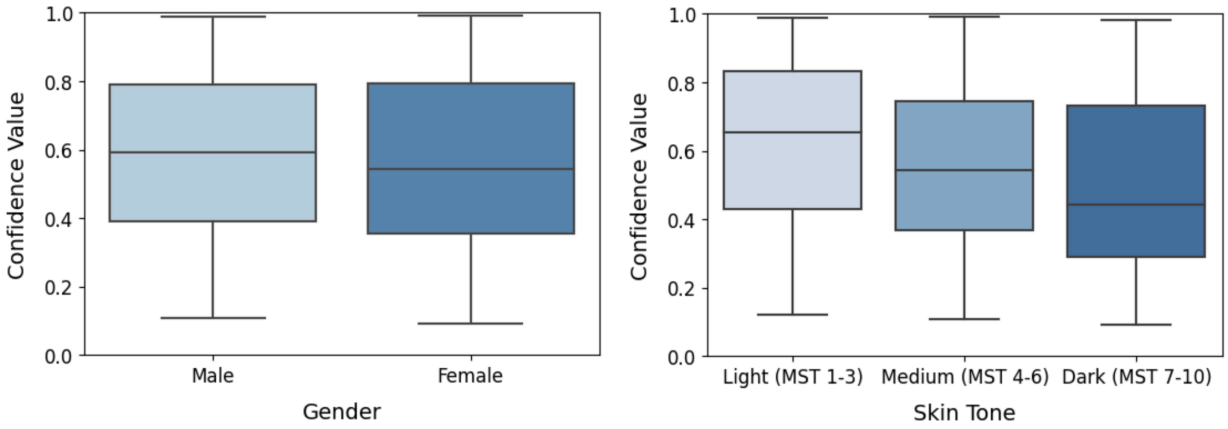

HARDWARE CHARACTERISTICS

Hardware Details

Camera Specifications (see here)	
Field of view (horizontal)	87º
Color Filter Array	Bayer, Monochrome
Frame Rate	45FPS @ 6MHz
Pixel Array (Active/ Effective)	324 x 324 / 320 x 320
Electrical Specifications	
Operating Voltage Range (regulator enabled)	3.5V to 5.5V
Operating Voltage Range (regulator disabled)	3.0V to 3.6V
Operating Current	40 mA
Operating Temperature	-20 ºC to 85 ºC
Communication Specifications	
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

Device Diagrams

Front and side view of sensor.

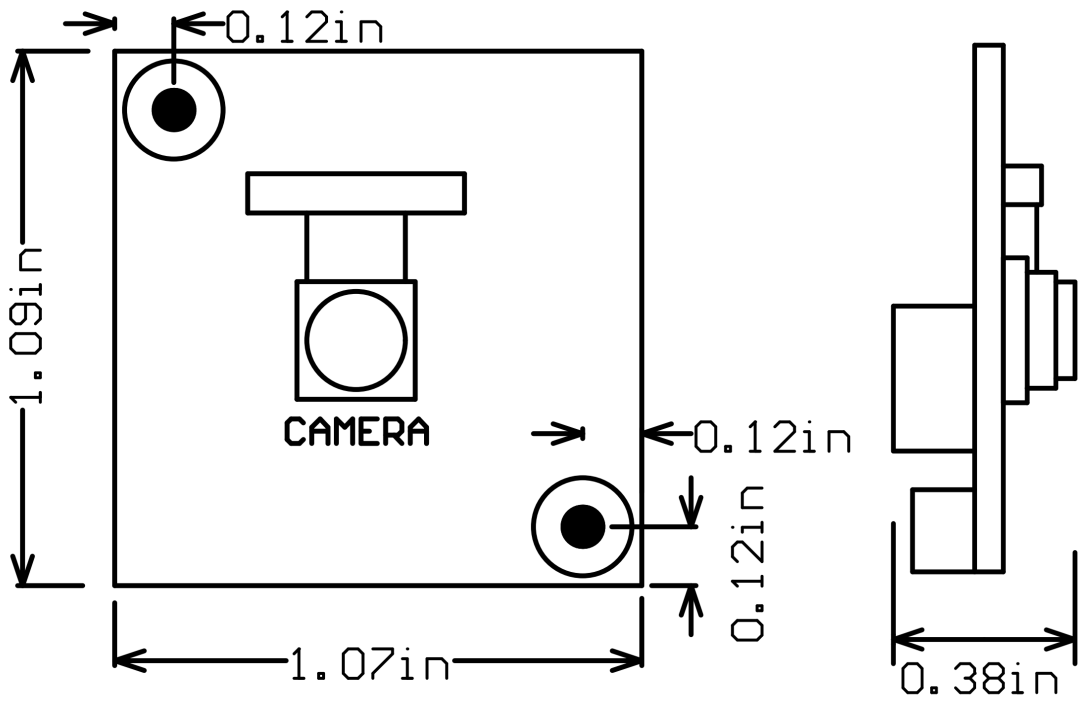

Back view of sensor.

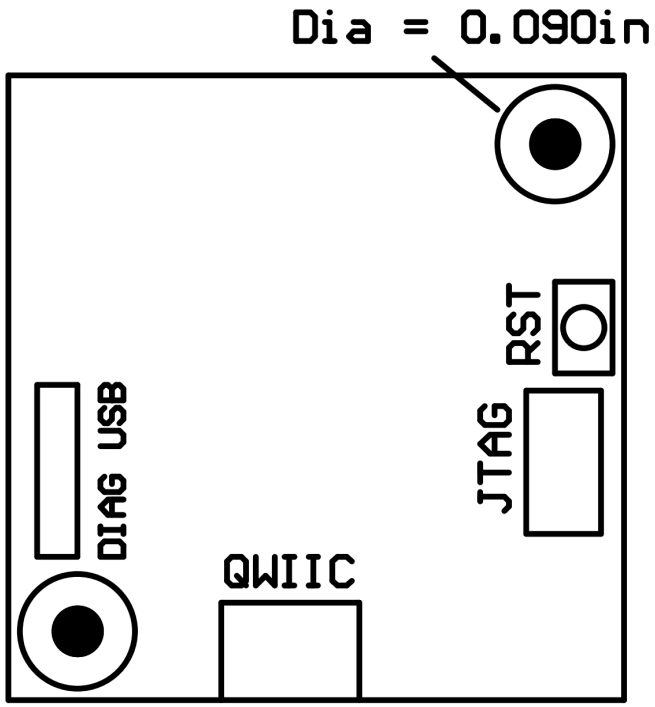

Bill of Materials

The following is a comprehensive list of materials required to assemble the PA1 person detection module, commonly referred to as the bill of materials. All unit cost values quoted in minimum order quantity of one.

Category In TinyML Calculator	Component	Unit Cost ($)	Quantity	Manufacturer	Link to Datasheet (if available)
Functional Components					
✔	RP2040 Microcontroller	1.00	1	Raspberry Pi	https://datasheets.raspberrypi.com/rp2040/rp2040-datasheet.pdf
✔	QVGA Camera Module HM01B0	8.90	1	HiMax	https://cdn.sparkfun.com/assets/7/f/c/8/3/HM01B0-MNA-Datasheet.pdf
✔	Flash Memory W25Q16JVSNIQ	0.36	1	Winbond Electronics	https://www.winbond.com/resource-files/w25q16jv%20spi%20revg%2003222018%20plus.pdf
✔	12 MHz Crystal Oscillator 445C25D12M00000	0.42	1	CTS-Frequency Controls	https://www.mouser.com/datasheet/2/96/008-0360-0-786290.pdf
Power Circuitry					
	Voltage Regulator TLV70228 2.8V	0.69	1	Texas Instruments	https://www.digchip.com/datasheets/download_datasheet.php?id=3747267&part-number=TLV70228
Indication					
✔	LTST-C190KGKT LED	0.05	1	Lite-On Inc.	https://www.digikey.com/htmldatasheets/production/37809/0/0/1/ltst-c190kgkt.pdf
Connectors					
	FFC connector FH26W-31S-0	1.28	1	Hirose Electric Co Ltd	https://www.hirose.com/product/download/?distributor=digikey&type=specSheet&lang=en&num=FH26W-31S-0.3SHW(60)
	Qwiic connector PRT-14417	0.57	1	SparkFun Electronics	https://www.mouser.com/datasheet/2/813/Qwiic_Connector_Datasheet-1223982.pdf
Passive Components					
✔	Resistors	0.01	10	-	N/A
✔	Capacitors (low value)	0.01	15	-	N/A
✔	Capacitors (high value)	0.05	7	-	N/A
✔	Ferrite bead 600Ω	0.07	2	-	N/A
✔	Printed circuit board	0.50	1	-	N/A
	Total	**14.51**			

Environmental Impact

With the widespread deployment of smart sensors, it is essential to consider and be conscious of the environmental impact such ubiquitous computing may have. Thus another component we advocate to be included in the datasheet is an “environmental impact" section that outlines the device footprint. Using the methodology of [9], we generated a sample of what this section might look like as part of the datasheet for our sensor specifically. We capture the carbon footprint (CO2-eq.) of our ML sensor in the chart below. Due to the limited amount of data available on electronic device footprint we were not able to capture every single component. We were able to account for 10 out of 13 components from our bill of materials, though, which we feel captures the concept sufficiently for the sake of demonstration. We were unable to find data for the connectors and voltage regulator. However, in addition to the bill of materials, we capture the carbon footprint for the ML sensor’s model training, transport, and three-year use.

The total carbon footprint, including embodied and operational footprint, of our ML Sensor is approximately **2.34 kg CO2-eq**. The chart below shows how the footprint is broken down. The majority of the footprint can be attributed to the power supply and camera sensor.

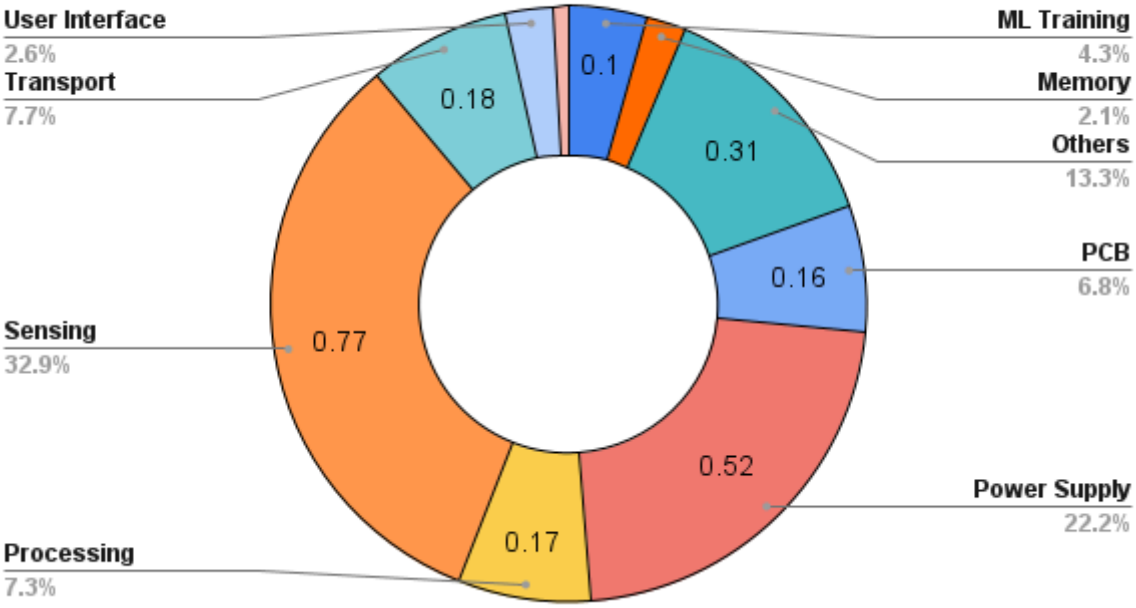

We note that we do not claim that this is 100% accurate but rather a representative approximation of the sensor’s environmental impact and what other future datasheet should aim to include.

Acronyms

Acronym	Description
SNR	Signal-to-noise ratio
COCO	Common Objects in Context
FFC	Flexible Flat Cable
GDPR	General Data Protection Regulation
ML	Machine Learning
I2C	Inter-Integrated Circuit
LED	Light-Emitting Diode
MCU	Microcontroller Unit
SCL	Serial Clock
SDA	Serial Data
GPIO	General Purpose Input Output
SDK	Software Development Kit
QVGA	Quarter Video Graphics Array

Glossary

Lux	Photometric unit of luminance (at 550 nm, 1 lux = 1 lumen/m^2 = 1/683 W/m^2)
Sensitivity	A measure of pixel performance that characterizes the rise of the photodiode or sense node signal in Volts upon illumination with light. Units are typically V/(W/m^2)/sec and are dependent on the incident light wavelength. Sensitivity measurements are often taken with 550 nm incident light. At this wavelength, 683 lux is equal to 1 W/m^2; the units of sensitivity are quoted in V/lux/sec. Note that responsivity and sensitivity are used interchangeably in image sensor characterization literature so it is best to check the units.
SNR	Signal-to-noise ratio. This number characterizes the ratio of the fundamental signal to the noise spectrum up to half the Nyquist frequency.
Inference	The process of applying a trained machine learning model to unseen data for making predictions or classifications. In the context of person detection, it involves analyzing images or video frames to determine if a person is present.
False Positive	A situation in person detection where the system incorrectly identifies an object or pattern as a person when it is not.
False Negative	A situation in person detection where the system fails to identify a person when one is present.
Accuracy	A performance metric that measures the overall correctness of a person detection system, indicating the percentage of correctly identified persons in the total number of instances.
Monk Skin Tone Scale	A 10-shade system, developed by Google, designed to provide a more inclusive representation of diverse skin tones in image-based technologies to address the challenges of representation in image-based technologies, especially for people of color.
Precision	A performance metric that measures the proportion of correctly identified persons among all the instances identified as persons by the system. It quantifies the system's ability to avoid false positives.
Recall (Sensitivity)	A performance metric that measures the proportion of correctly identified persons among all the actual persons present in the data. It quantifies the system's ability to avoid false negatives.
Threshold	A predefined value used to determine whether the output of a person detection system indicates the presence or absence of a person. Adjusting the threshold affects the balance between false positives and false negatives.
Training Set	Labeled examples or samples used to teach a machine learning model to recognize and classify objects accurately. In the case of person detection, it comprises images or videos with annotated information about the presence or absence of people.
Test Set	A subset of the dataset that is strictly used to evaluate the performance of a model after it has been trained. The test set provides an unbiased evaluation of a model's generalization to new, unseen data. It should never be used during training or hyperparameter tuning.
Validation Set	A subset of the dataset, separate from the training set, used to evaluate a model during training. It provides an intermittent check on the model's performance, allowing for hyperparameter tuning and model selection. By evaluating model performance on a validation set, one can detect issues like overfitting (where the model performs exceptionally well on the training set but poorly on new, unseen data). Once the model is optimized using the validation set, its final performance is then assessed on the test set.
Person Detection	The process of identifying the presence and location of a person within an image or video stream.
Sensor	A device that detects and measures physical or environmental properties, such as the presence of a person, and converts them into electrical signals.

USER STUDY FORMS

Study Flyer


Harvard Edge Computing Group
MACHINE LEARNING SENSORS
EXPERIMENTAL STUDY
INFO & SIGN-UP
PARTICIPANTS NEEDED!
Come help us test the first machine learning sensor!
CONTACT :
matthew_stewart@g.harvard.edu
yasmineomri@college.harvard.edu
tinyurl.com/mlsensors

Interest Form

Machine Learning Sensors Experimental Study Interest Form

The Edge Computing Group is seeking participants for an experimental study evaluating a new paradigm of machine learning sensors that we are designing.

We are looking into the next generation of sensors, ML sensors, which use on-device machine learning to extract useful information from the raw data before reporting it to the outer system. The ML sensor paradigm comes with significant benefits when it comes to modularity, composability, power efficiency, privacy and security, and more! Part of establishing the ML sensor paradigm is reimagining what the conventional sensor data sheet. More particularly, we want to test the end-to-end performance of the sensors through a representative study to see how it performs in the real world on a set of people that it was not trained on. In order to investigate potential algorithmic bias, we will be collecting data on participants' sex and skin tone throughout the study.

The study will take place in the SEC, where a room will be set up with a series of sensors that will output the probability of seeing a person. The study will involve varying distances and lighting settings and is estimated to take approximately 2 minutes per person. Participants will be provided with snacks and required to fill out a brief consent form prior to participation. The sensors will take instantaneous images for processing, but none of the information will be saved to ensure privacy preservation.

If you have any questions or are interested in helping out with this study, please contact matthew_stewart@g.harvard.edu or yasmineomri@college.harvard.edu. For more information about our project, please see our high-level blog or this paper.

Name:

_____

Email:

_____

Which of these weeks would you be available to take part in this experiment? We will be reaching out by email closer to the date to organize concrete times.

- ☐ April 3-7
- ☐ April 10-14
- ☐ April 17-21
- ☐ April 24-28

Machine Learning Sensors Experimental Study Consent Form

The Edge Computing Group is looking for participants to take part in an experimental study we are devising to evaluate a new paradigm of intelligent sensor that we are developing.

We are looking into the next generation of sensors, ML sensors, which use on-device machine learning to extract useful information from the raw data before reporting it to the outer system. The ML sensor paradigm comes with significant benefits when it comes to modularity, composability, power efficiency, privacy and security, and more. Part of establishing the ML sensor paradigm is reimagining the conventional sensor data sheet. More particularly, we want to test the end-to-end performance of the sensors through a representative study to see how it performs in the real world on a set of people that it was not trained on.

Please review and confirm your agreement to the following:

What You Will Do in this Study: Participants will be requested to select their gender identity and skin tone (using the Monk Skin Tone Scale). You will be asked to input this information at the end of this consent form and prior to entering the study room. This information will remain anonymous, as no name or additional information about the participant will be recorded.

The participant will then enter the study room and stand in front of a group of six sensors at three varying distances. At each distance, three different light settings will be tested. One individual in the room will be altering the lighting conditions and another taking note of the sensor outputs. The sensor predictions (detection of a person) are recorded and coupled with the skin tone and gender identity to investigate potential algorithmic bias. We expect the experiment to take about 2-3 minutes per person.

Participation in this study is voluntary. You do not have to be in this study if you do not want to and you can quit the study at any time.

Risks and Benefits: There are no known risks to participants. Participants will be compensated with snacks for their time. The results of this study intend to be published and could lead to the development of more accurate and effective person detection technology, which could have a positive impact on various fields, such as security, healthcare, and transportation.

Confidentiality: No privacy-sensitive information will be recorded, such as images or personal details. Skin tone and gender identity will be deidentified prior to data analysis, ensuring participants remain entirely anonymous.

We will only report a summarized analysis of the variability of the sensors' prediction accuracies with skin tone and gender.

If you have questions about the survey or study, contact Matthew Stewart (matthew_stewart@g.harvard.edu), Yasmine Omri (yasmineomri@college.harvard.edu) or Professor Vijay Janapa Reddi (vj@eecs.harvard.edu).

I confirm that I am at least eighteen years of age, I understand and have read the points above, and consent to the collection, use, and sharing of your anonymous responses.

- ☐ I consent

Please select your gender identity:

- ☐ Male
- ☐ Female
- ☐ Transgender Male
- ☐ Transgender Female
- ☐ Non-binary
- ☐ Other (Please Specify)

Which Monk Scale skin tone most closely matches your skin color?

Please refer to the Monk Scale skin tone chart below or on the table. Ask a team member if you are having trouble identifying the appropriate value.

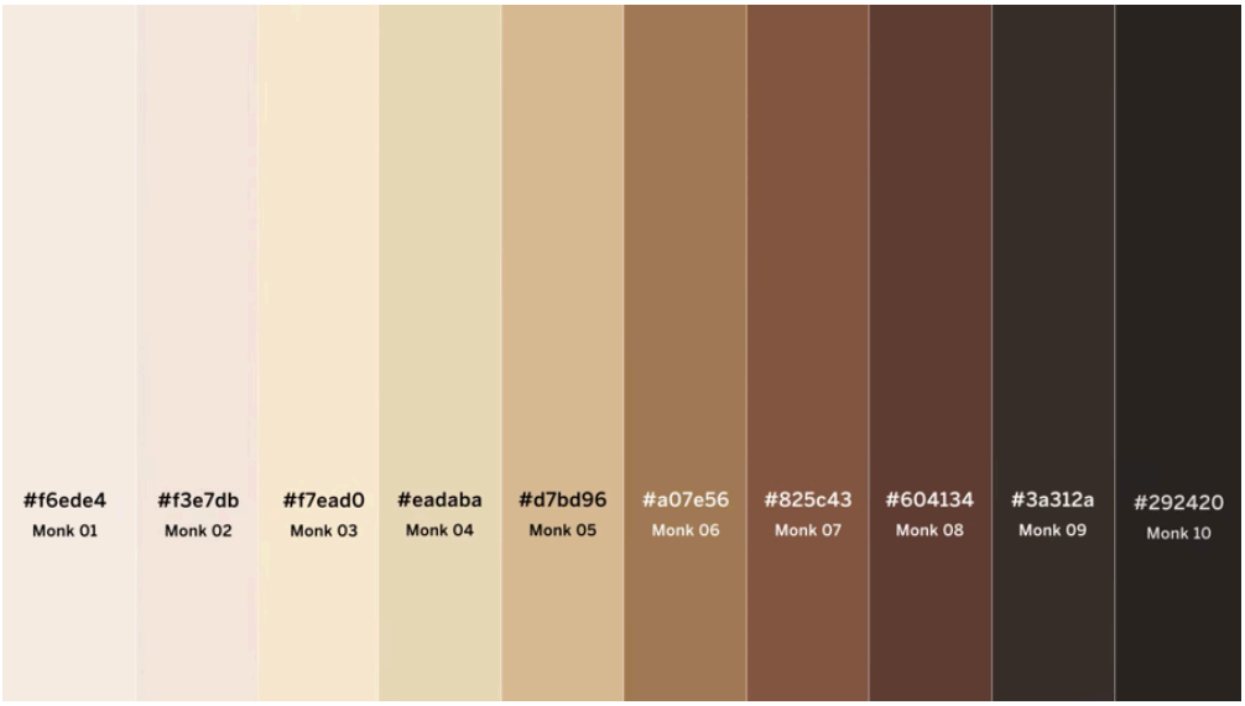

- ☐ 1
- ☐ 2
- ☐ 3
- ☐ 4
- ☐ 5
- ☐ 6
- ☐ 7
- ☐ 8
- ☐ 9
- ☐ 10

Appendix B. Example Data Sheet - Commercial Sensor

OVERVIEW

Person Sensor V1.0
SEN-21231

Compliance and Certifications

The person detection sensor complies with essential industry standards and regulations, including RoHS for environmental safety and GDPR for protecting individual privacy. The sensor has been audited by Kodelski Security for security and privacy implications.

Description

The Person Sensor is a small, low-cost hardware module that detects nearby peoples' faces, and returns information about how many there are, where they are relative to the device, and performs facial recognition. It is designed to be used as an input to a larger system, for example to wake up a kiosk display from sleep mode when somebody approaches, mute a microphone when nobody is present, or orient a fan so it's always pointing at the nearest person.

The sensor has a small form factor and utilizes a monochrome camera with a field of view of 640 x 480 (VGA). The input voltage for the sensor is 3.3V and the typical operating current for the sensor is 40 mA. The sensor communicates via I2C/Qwiic mode, conforming to SparkFun Qwiic electrical/mechanical specifications, and has a maximum cable length of 1 m at 400 kb/s. Longer cables can be used at lower data rates. The sensor has a maximum data rate of 400 kb/s.

Features

- Real-time person + head pose tracking with on-device ML
- Real-time person identification with on-devicce ML
- Low power consumption
- Onboard camera
- Small form factor: 22 x 20 x 10 mm
- I2C serial communication
- Lead-free

Use Cases

- Security
- Home automation
- Consumer appliances

MODEL CHARACTERISTICS

Software Flow Diagram

8-bit grayscale images (640x480) are resized to 192x144 and passed into a RetinaFace model trained to detect human faces. This model outputs a list of faces with coordinates for a bounding box around each face as well as five key facial landmarks. If identity is enabled, bounded faces are cropped out of the original image and rescaled to 47x55 and passed into a DeepID model to generate an embedding. This embedding is compared with saved facial IDs, and the nearest ID is returned along with the bounding box and information about whether the person is facing the sensor. Output information is communicated via Qwiic interface to the application processor.

	Person Detection Model	**Person ID Model**
Architecture	RetinaFace	DeepID
Framework	TFLite Micro	TFLite Micro
Validation Set	20%	-
Quantization	int8	int8
Inference Time	140 ms	125 ms
Peak RAM Usage	442.6 kB	189 kB
Flash Usage	449 kB	397 kB

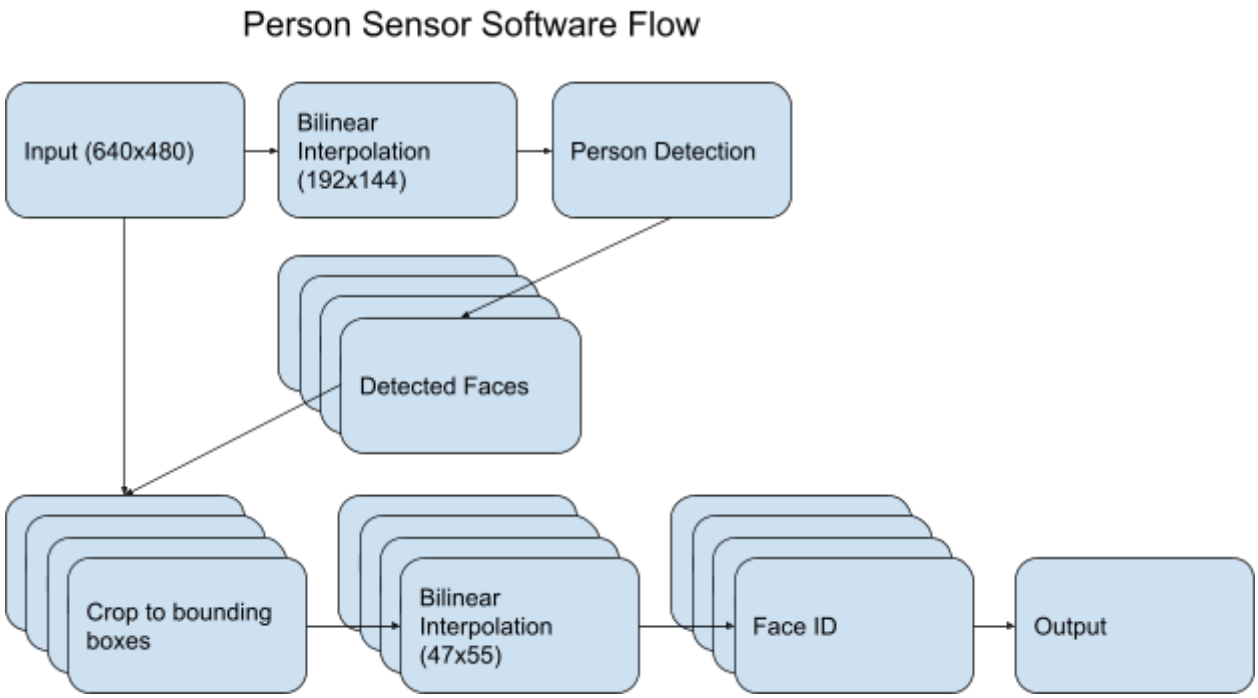

Dataset Nutrition Label

The data nutrition label is publicly available here, with some important features outlined below.

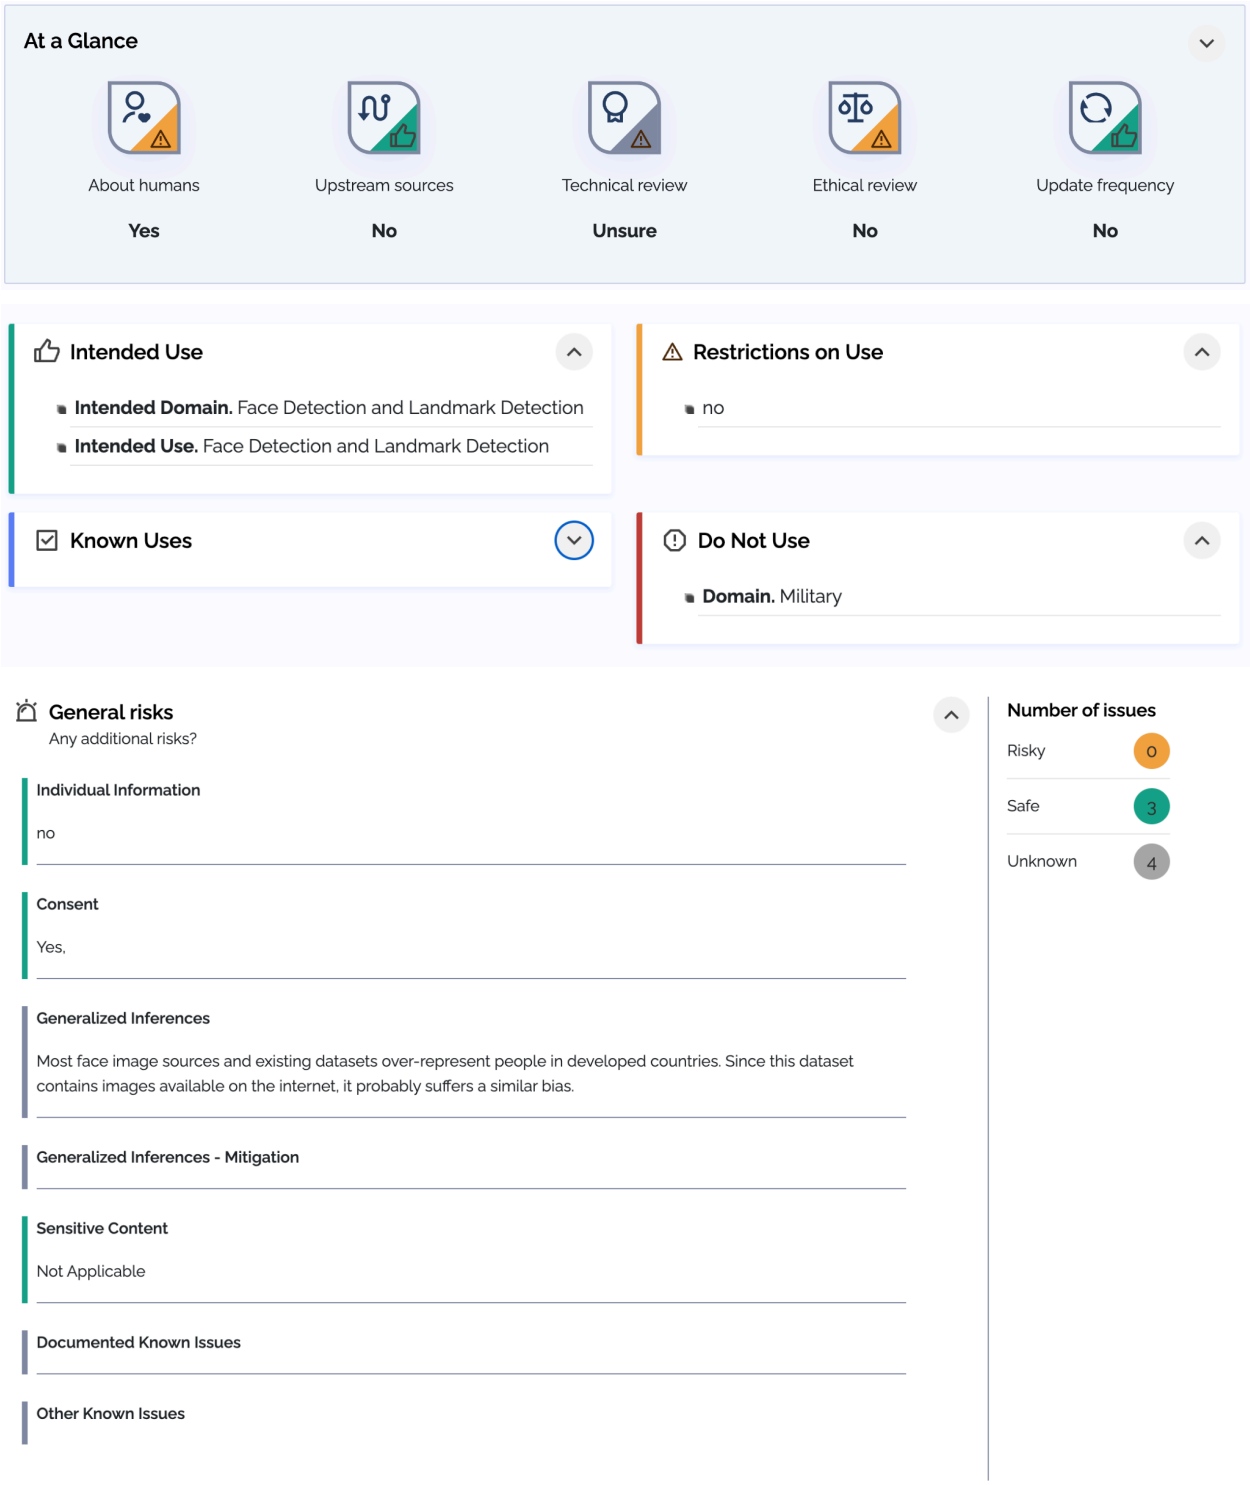

IoT Security and Privacy Label

Security & Privacy Overview

Useful Sensors

Person Sensor V1.0
Firmware version: Not disclosed - **updated on:** 2023-05-03
The device was manufactured in: China

Security Mechanisms

Security updates	No security updates
Access control	Not disclosed

Data Practices

	Visual	Audio	Physiological	Location
Sensor type	Camera			
Purpose	Providing and improving device functions			
Data stored on the device	No device storage			
Data stored in the cloud	No cloud storage			
Data shared with	Not shared			
Data sold to	Not sold			
Other collected data				

Privacy policy	Not disclosed

More Information

Detailed Security & Privacy Label:
Not disclosed

CMU IoT Security and Privacy Label CISPL 1.0 iotsecurityprivacy.org

Security & Privacy Details

Useful Sensors

Person Sensor V1.0
Firmware version: Not disclosed **- updated on:** 2023-05-03
The device was manufactured in: China

Security Mechanisms

Security updates	No security updates *Sensor is a standalone unit. For security and privacy, the firmware cannot be changed and only sensor outputs are available.*
Access control	Not disclosed
Security oversight	Audits performed by third-party security auditors *Third party security audit performed by Kudelski Security*
Ports and protocols	Not disclosed
Hardware safety	Not disclosed
Software safety	Not disclosed
Personal safety	Not disclosed
Vulnerability disclosure and management	Not disclosed
Software and hardware composition list	Not disclosed
Encryption and key management	Not disclosed

Data Practices

Sensor data collection	Visual
Sensor type	Camera
Data collection frequency	Continuous
Purpose	Providing and improving device functions
Data stored on the device	No device storage
Local data retention time	No retention
Data stored in the cloud	No cloud storage
Cloud data retention time	No retention
Data shared with	Not shared
Data sharing frequency	Not shared
Data sold to	Not sold
Other collected data	None

Data linkage	Not disclosed
What will be Inferred from User's Data	Not disclosed
Special data handling practices for children	Not disclosed
In Compliance with	Not disclosed
Privacy policy	Not disclosed

More Information

Call Useful Sensors with your questions at	1 805 813 7571
Email Useful Sensors with your questions at	contact@usefulsensors.com
Functionality when offline	Full functionality on offline mode
Functionality with no data processing	Not applicable
Physical actuations and triggers	Not disclosed
Compatible platforms	Not disclosed

CMU IoT Security and Privacy Label CISPL 1.0 iotsecurityprivacy.org

Machine Learning Model Specification

Person Detection Model

The person detection model was trained on a proprietary dataset of ~30,000 images with 300k labeled faces and five facial landmarks per face. The model input is a 192x144 raw image in 8-bit grayscale format, equivalent to 27,648 features. The training process was performed until the model accuracy ceased to improve. Final model performance achieved a precision of 91.8% on the test set, using a threshold of 0.7. The precision-recall curve of the model on the test set is shown below. The model was quantized to 8-bit integer using post training quantization using the Tensorflow Lite converter and is deployed using the Tensorflow Lite Micro runtime.

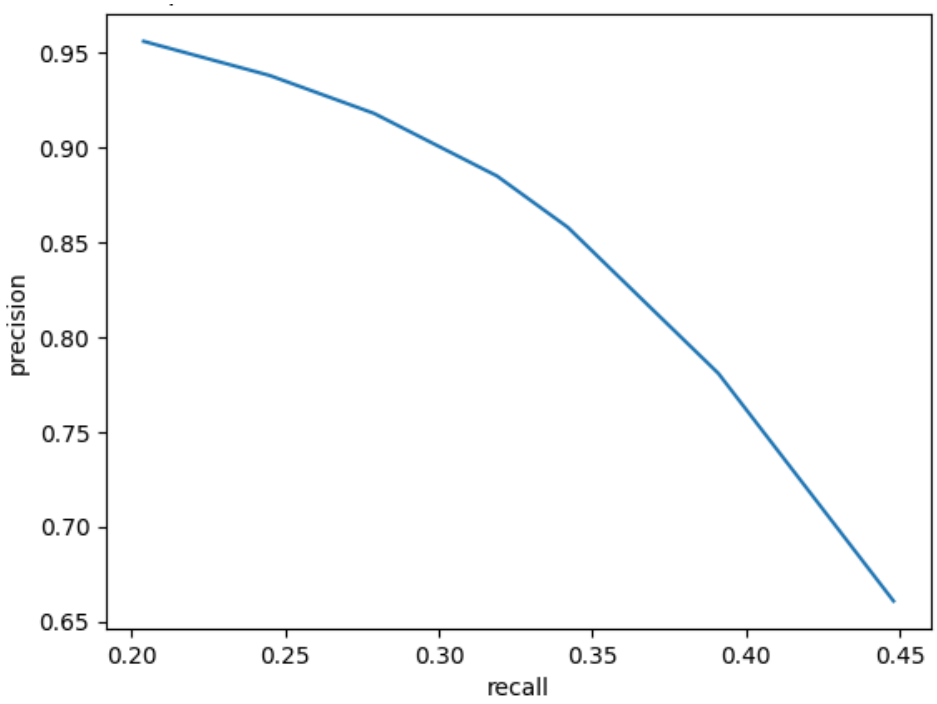

Face Identification Model

The face identification model was fine-tuned on a proprietary dataset encompassing ~4000 images across 5 identities captured using the sensor camera module. The input to the model is a 47x55 raw image in 8-big grayscale format, equivalent to 2,585 features. To produce the best separation between faces a dense classification layer was added to the model, and several iterations of freezing either the classification layer or the model was used to achieve a higher accuracy on the fine-tuning dataset. Finally, the classification layer was removed and embedding separation was evaluated using Principal Component Analysis (PCA) in three dimensions.

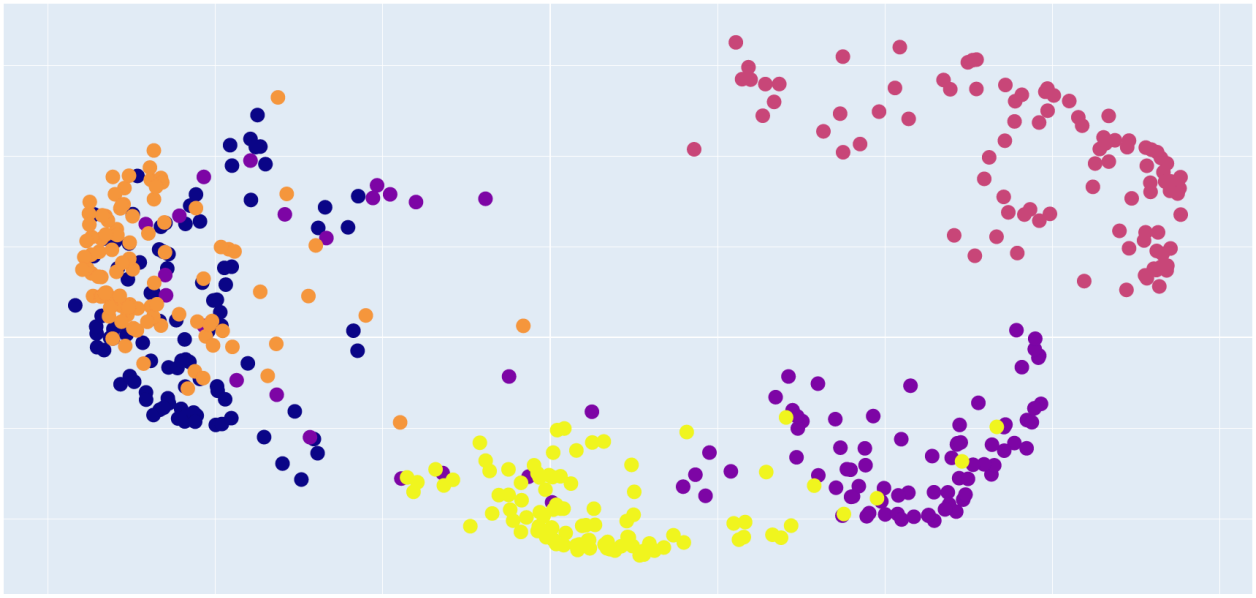

Each color represents one of five unique identities in the validation dataset. Distances between points indicate approximate distances between embeddings simplified to 2-D space.

The model was quantized to 8-bit integer using post training quantization using the Tensorflow Lite converter and is deployed using the Tensorflow Lite Micro runtime. On the sensor, the embedding generated by the Face ID model is compared against registered faces, and if a face with similar enough features is found, that identity is used. Otherwise, an identity of -1 is returned to indicate that no registered identity was found.

Performance Analysis

The end-to-end performance of the person detection sensor model was tested through an experimental study conducted in the Science and Engineering Complex (SEC) at Harvard University. The study involved 40 participants and evaluated the accuracy of the model under different lighting conditions using three identical sensors.

The study room measured 25 x 31 x 10 ft and contained 32 ceiling lights that were uniformly distributed in a 4 x 8 grid. The lighting conditions were captured quantitatively for each participant using a Lux LCD Illuminance Meter (Precision Vision, Inc.) and a C-800-U Spectrometer (Sekonic Corporation).

The sensors were mounted on a wooden board affixed to the wall at a height of 1.5 m above the ground. The participants were evaluated at three different distances (1.5 m, 4.5 m, and 7.5 m) from the sensors under each lighting condition. The ambient lighting in the room was provided by artificial lights, and blackout curtains were used to block the ambient lighting from outside.

The lighting levels were controlled using a dimmer switch that had three levels of operation, corresponding to 208±31, 584±51, and 1149±59 lux, respectively. When the lights were turned off, the illuminance meter gave a reading of zero lux. When all the lights were turned on at full strength, the sensor gave an average reading of 1149 lux. The color temperature of the lighting was measured to be 5600 K, corresponding to white light. Colored tape was placed on the ground to demarcate the locations where participants should stand during the experiment (i.e., 1.5, 4.5, and 7.5 m from the sensor array).

Before entering the study environment, the participants were asked to provide their gender identity and evaluate their skin tone according to the Monk Skin Tone (MST) Scale to evaluate algorithmic bias. The study evaluated algorithmic bias by bucketing skin tone into three categories: light (MST 0-4), medium (MST 5-7), and dark (MST 8-10). At each location and lighting condition, ten readings were taken from each sensor and averaged.

Participants were recruited using flyers, and those interested filled out a Study Interest Form. Upon arrival, participants signed a Consent Form indicating their willingness to participate in the study. The accuracy of the model is provided in the following graphs as a function of lighting condition, distance, gender identity, and skin tone. Overall, 63.2% of the participants were male, and 36.8% were female; the percentage of participants corresponding to each skin tone group was: 47.4% light, 39.4% medium, and 13.2% dark.

This anonymous study was approved by the Institutional Review Board of Harvard University on 6 April 2023 (Project Code: IRB23-0136).

Environmental Sensitivity

The device showed no decrease in performance under decreased lighting conditions. A moderate drop off in performance of around 10% is observed at distances 3-5 meters from the sensor.

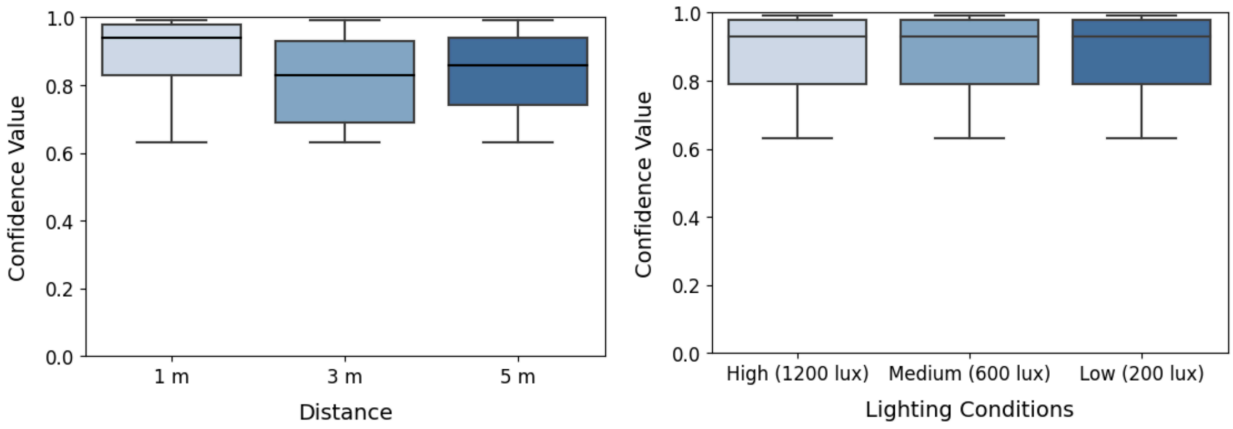

Demographic biases

A small gender bias is observed in model performance. A moderate skin tone bias was observed, showing approximately a 10% decrease in the confidence value for individuals with a darker skin tone.

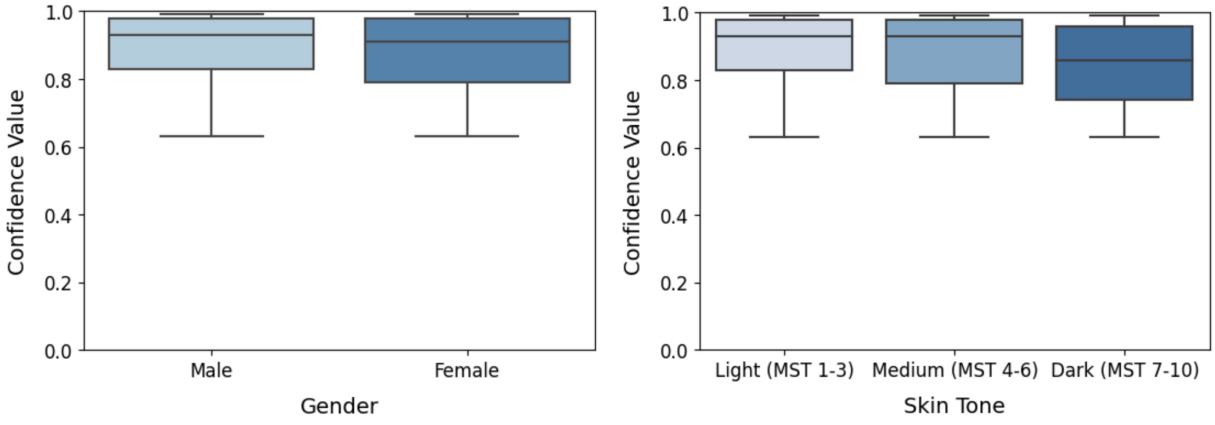

HARDWARE

Hardware Details

Camera Specifications (see here)	
Field of view (horizontal)	110º
Color Filter Array	Bayer, Monochrome
Frame Rate	60FPS @ 48MHz
Pixel Array (Active/ Effective)	644 x 484 / 640x480
Electrical Specifications	
Operating Voltage Range (regulator enabled)	3.1V to 3.5V
Operating Current	40 mA
Operating Temperature	-20 ºC to 85 ºC
Communication Specifications	
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	Text on sensor is upright, up arrow points upwards
GPIO mode	INT pin is high when person is detected
Diagnostic LED	Default behavior of green LED on board: illuminates when person detected.
Data Transfer and Format	See I2C Protocol Table
I2C Address	0x63

I2C Protocol

Address	Name	Default	Description
0x01	Mode	0x01 (continuous)	Mode. See mode table below.
0x02	Enable ID	0x00 (False)	Enable / Disable the ID model. With this flag set to False, only capture bounding boxes.
0x03	Single shot	0x00	Trigger a single-shot inference. Only works if the sensor is in standby mode.
0x04	Label next ID	0x00	Calibrate the next identified frame as person N, from 0 to 7. If two frames pass with no person, this label is discarded.
0x05	Persist IDs	0x01 (True)	Store any recognized IDs even when unpowered.
0x06	Erase IDs	0x0	Wipe any recognized IDs from storage.
0x07	Debug Mode	0x01 (True)	Whether to enable the LED indicator on the sensor.

Mode	Name	Description
0x00	Standby	Lowest power mode, sensor is in standby and not capturing.
0x01	Continuous	Capture continuously, setting the GPIO trigger pin to high if a face is detected.

Device Diagrams

Front and backside of the sensor module.

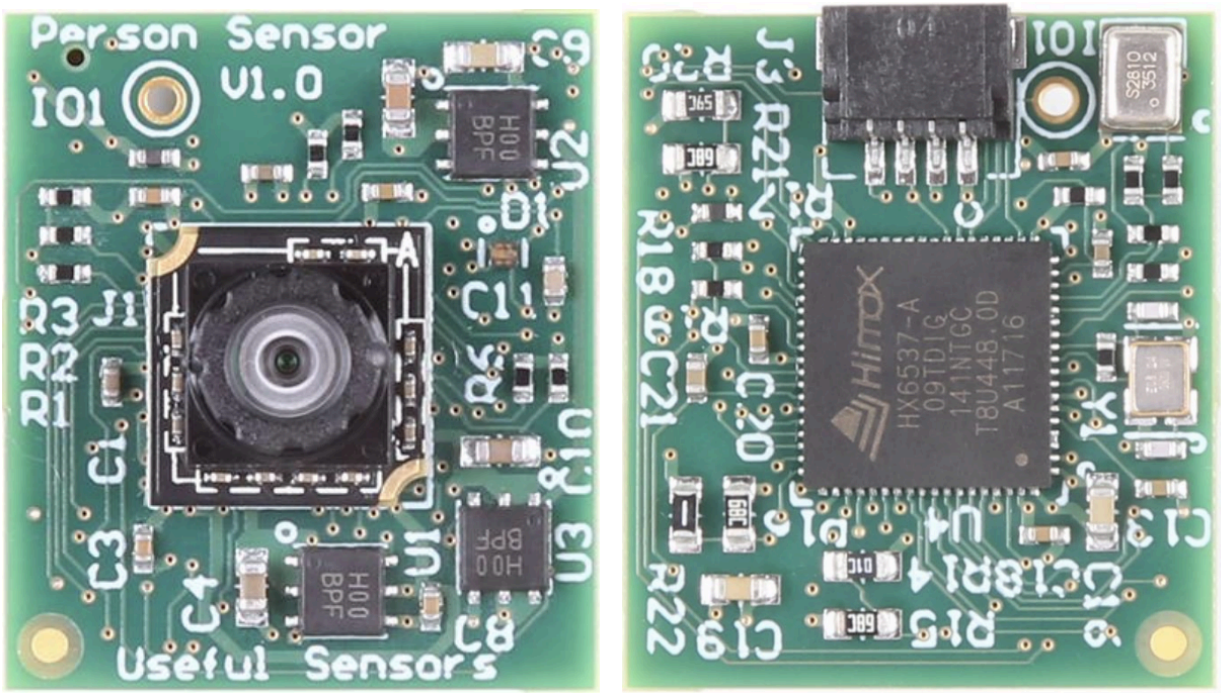

Qwiic connector interface.

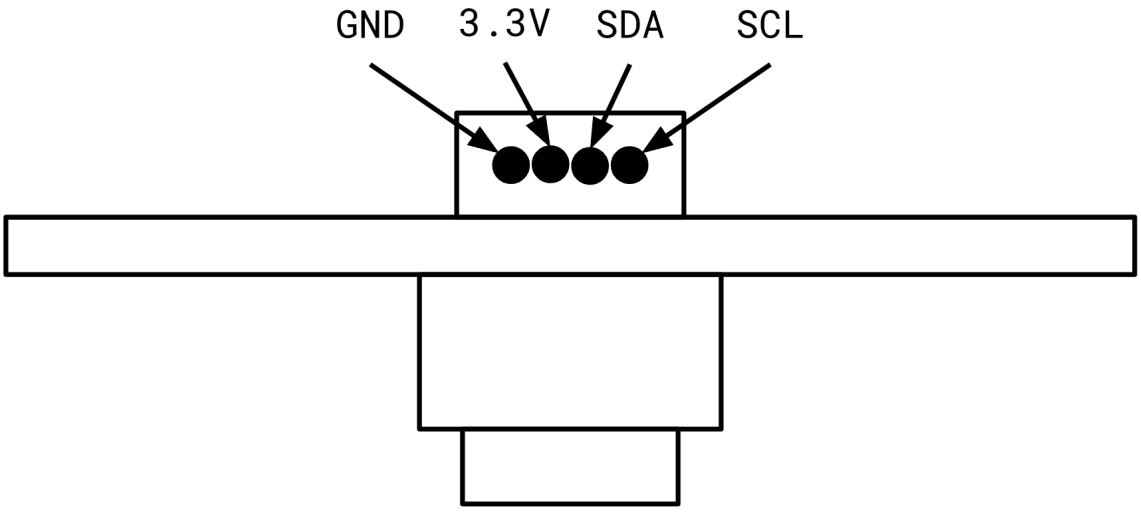

Frontside schematic of sensor.

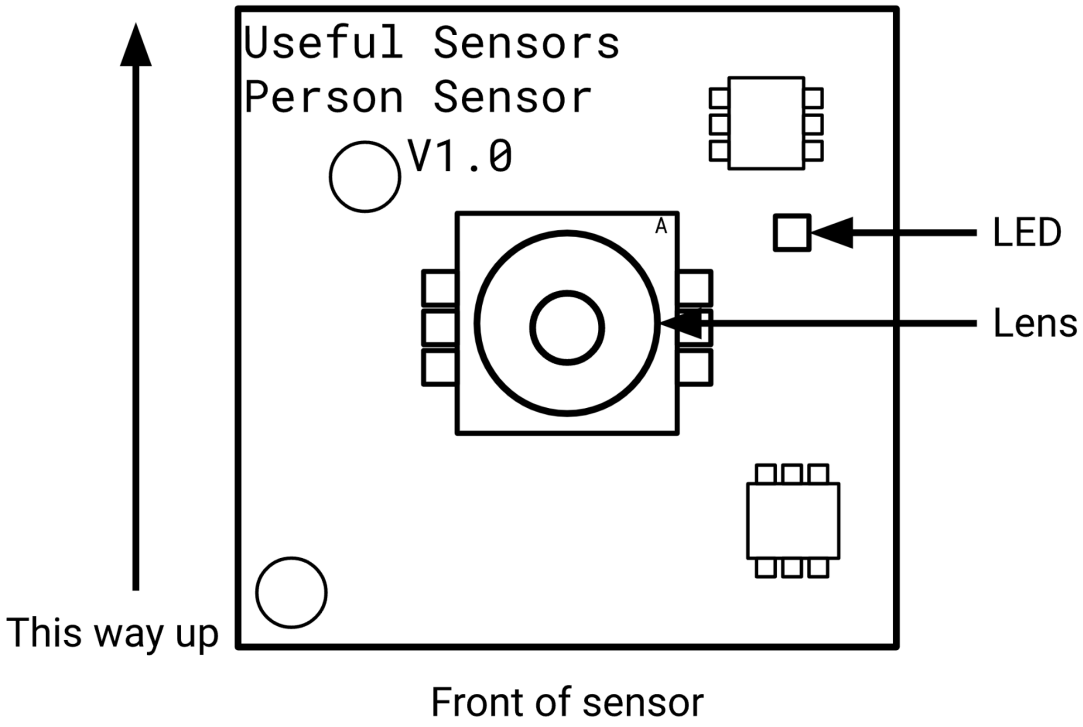

Backside schematic of sensor.

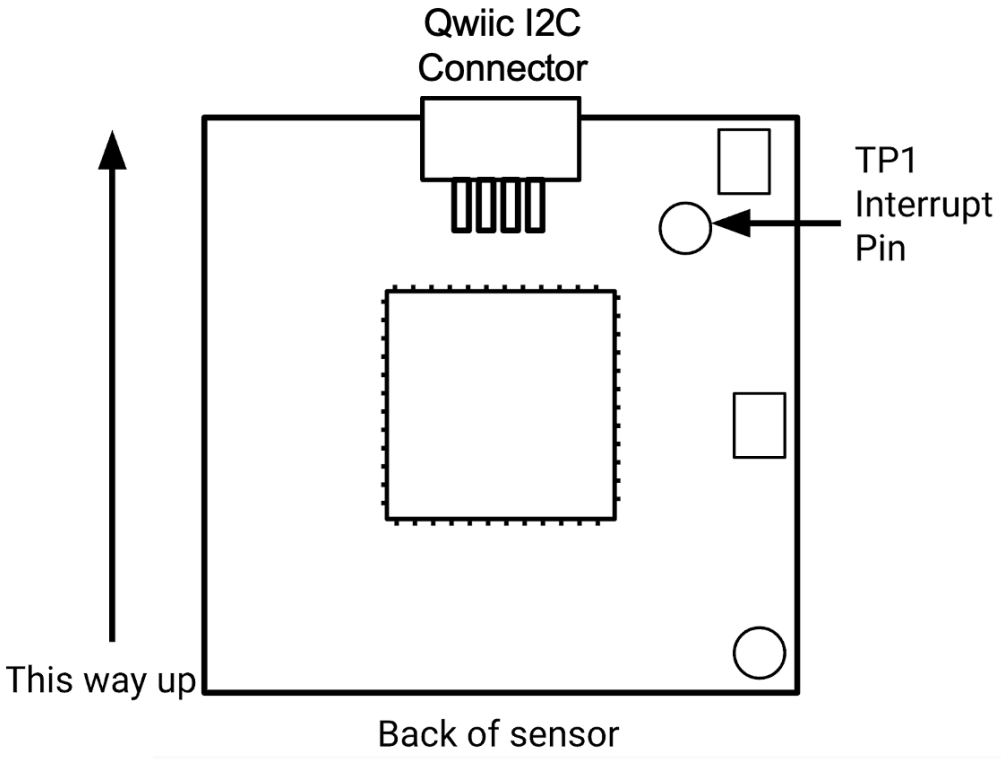

Bill of Materials

The following is a comprehensive list of materials required to assemble the Person Sensor V1.0. All unit cost values quoted in minimum order quantity of one.

Category In TinyML Calculator	Component	Unit Cost ($)	Quantity	Manufacturer	Link to Datasheet (if available)
Functional Components					
✔	Himax MCU HX6537/39/40-A	14.50	1	HiMax	https://cdn.sparkfun.com/assets/6/6/7/e/8/HX6537-A_DS_public_v01_1_.pdf
✔	Camera Module HM0360-MWA	8.81	1	HiMax	https://cdn.sparkfun.com/assets/d/2/9/9/7/Pre-HM0360_DS_preliminary_v04__Ltd._-1.pdf
✔	MEMS Microphone SPH0641LM4H-1	1.05	1	Knowles	https://media.digikey.com/pdf/Data%20Sheets/Knowles%20Acoustics%20PDFs/SPH0641LM4H-1.pdf
✔	Crystal Oscillator ECS-240-10-36-CKM-TR	0.57	1	ECS Inc.	https://ecsxtal.com/store/pdf/ECX-2236.pdf
Power Circuitry					
	Adjustable Linear Voltage Regulator R1173D001B-TR-FE	1.33	3	Nisshinbo Micro Devices Inc.	https://www.nisshinbo-microdevices.co.jp/en/pdf/datasheet/r1173-ea.pdf
Indication					
✔	RGB LED	0.1	1	Harvatek Corporation	https://media.digikey.com/pdf/Data%20Sheets/Harvatek%20PDFs/B39D3RGB-F6C0001HOU1930.pdf
Connectors					
	Board to Camera OK-10F030-04	1.22	1	AliExpress	N/A
	Qwiic JST SH 4-pin Right Angle Connector	0.40	1	Adafruit	N/A
Passive Components					
✔	Misc resistors	0.01	15	-	N/A
✔	Misc capacitors	0.01	17	-	N/A
✔	Misc inductors	0.01	1	-	N/A
	Total	**30.97**			

Environmental Impact

With the widespread deployment of smart sensors, it is essential to consider and be conscious of the environmental impact such ubiquitous computing may have. Thus another component we advocate to be included in the datasheet is an “environmental impact" section that outlines the device footprint. Using the methodology of [9], we generated a sample of what this section might look like as part of the datasheet for our sensor specifically. We capture the carbon footprint (CO2-eq.) of our ML sensor in the chart below. Due to the limited amount of data available on electronic device footprint we were not able to capture every single component. We were able to account for 8 out of 11 components from our bill of materials, though, which we feel captures the concept sufficiently for the sake of demonstration. We were unable to find data for the connectors and voltage regulator. However, in addition to the bill of materials, we capture the carbon footprint for the ML sensor’s model training, transport, and three-year use.

The total carbon footprint, including embodied and operational footprint, of our ML Sensor is approximately **2.82 kg CO2-eq**. The chart below shows how the footprint is broken down. The majority of the footprint can be attributed to the power supply and camera sensor.

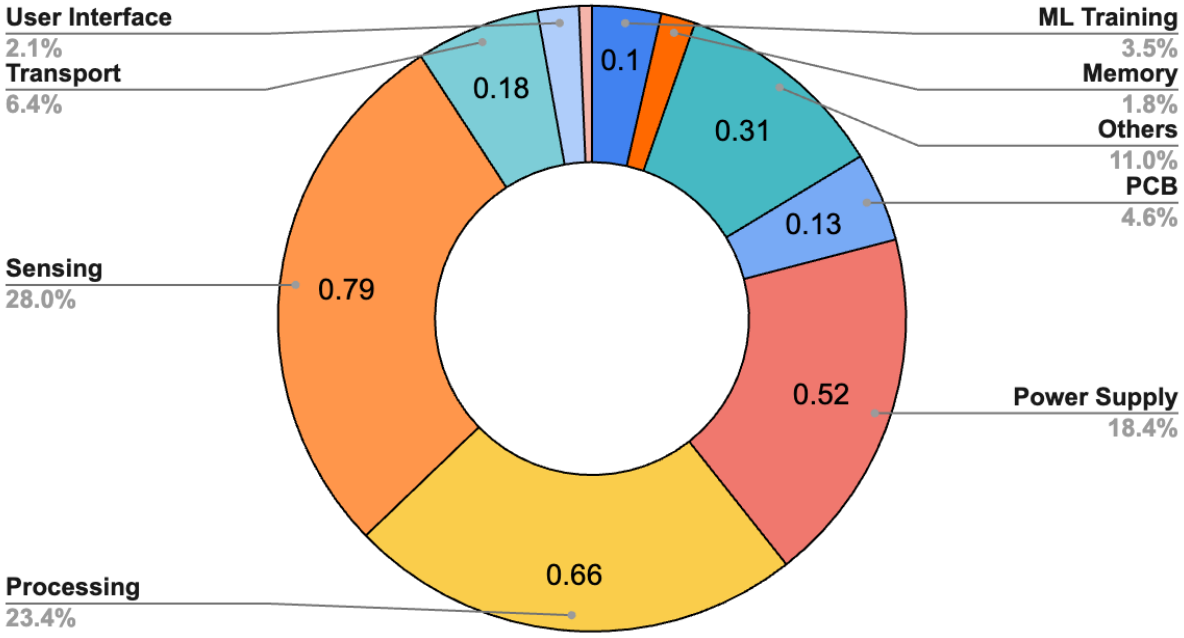

We note that we do not claim that this is 100% accurate but rather a representative approximation of the sensor’s environmental impact and what other future datasheet should aim to include.

Acronyms

Acronym	Description
GDPR	General Data Protection Regulation
GPIO	General Purpose Input Output
ML	Machine Learning
I2C	Inter-Integrated Circuit
ID	Identifier
IoU	Intersection Over Union
LED	Light-Emitting Diode
MCU	Microcontroller Unit
MEMS	Microelectromechanical System
MST	Monk Skin Tone Scale
PCA	Principal Component Analysis
RGB	Red Blue Green

Glossary

Lux	Photometric unit of luminance (at 550 nm, 1 lux = 1 lumen/m^2 = 1/683 W/m^2)
Sensitivity	A measure of pixel performance that characterizes the rise of the photodiode or sense node signal in Volts upon illumination with light. Units are typically V/(W/m^2)/sec and are dependent on the incident light wavelength. Sensitivity measurements are often taken with 550 nm incident light. At this wavelength, 683 lux is equal to 1 W/m^2; the units of sensitivity are quoted in V/lux/sec. Note that responsivity and sensitivity are used interchangeably in image sensor characterization literature so it is best to check the units.
IoU	Intersection Over Union (IoU) is a metric used to evaluate the accuracy of an object detector on a specific dataset. It measures the overlap between the predicted bounding box (from the detector) and the ground truth bounding box. Values range between 0 and 1, where a higher value indicates better prediction accuracy. A value of 0 indicates no overlap, while a value of 1 indicates perfect overlap (the predicted box matches the ground truth exactly).
Inference	The process of applying a trained machine learning model to unseen data for making predictions or classifications. In the context of person detection, it involves analyzing images or video frames to determine if a person is present.
Accuracy	A performance metric that measures the overall correctness of a person detection system, indicating the percentage of correctly identified persons in the total number of instances.
Principal Component Analysis	A statistical procedure that transforms a set of observations of possibly correlated variables into a set of values of linearly uncorrelated variables called principal components. These components are orthogonal to each other and capture the variance in the data in decreasing order.
Monk Skin Tone Scale	A 10-shade system, developed by Google, designed to provide a more inclusive representation of diverse skin tones in image-based technologies to address the challenges of representation in image-based technologies, especially for people of color.
Training Set	Labeled examples or samples used to teach a machine learning model to recognize and classify objects accurately. In the case of person detection, it comprises images or videos with annotated information about the presence or absence of people.
Test Set	A subset of the dataset that is strictly used to evaluate the performance of a model after it has been trained. The test set provides an unbiased evaluation of a model's generalization to new, unseen data. It should never be used during training or hyperparameter tuning.
Validation Set	A subset of the dataset, separate from the training set, used to evaluate a model during training. It provides an intermittent check on the model's performance, allowing for hyperparameter tuning and model selection. By evaluating model performance on a validation set, one can detect issues like overfitting (where the model performs exceptionally well on the training set but poorly on new, unseen data). Once the model is optimized using the validation set, its final performance is then assessed on the test set.
Person Detection	The process of identifying the presence and location of a person within an image or video stream.
Sensor	A device that detects and measures physical or environmental properties, such as the presence of a person, and converts them into electrical signals.

USER STUDY FORMS

Study Flyer


Harvard Edge Computing Group
MACHINE LEARNING SENSORS
EXPERIMENTAL STUDY
INFO & SIGN-UP
PARTICIPANTS NEEDED!
Come help us test the first machine learning sensor!
CONTACT :
matthew_stewart@g.harvard.edu
yasmineomri@college.harvard.edu
tinyurl.com/mlsensors

Interest Form

Machine Learning Sensors Experimental Study Interest Form

The Edge Computing Group is seeking participants for an experimental study evaluating a new paradigm of machine learning sensors that we are designing.

We are looking into the next generation of sensors, ML sensors, which use on-device machine learning to extract useful information from the raw data before reporting it to the outer system. The ML sensor paradigm comes with significant benefits when it comes to modularity, composability, power efficiency, privacy and security, and more! Part of establishing the ML sensor paradigm is reimagining what the conventional sensor data sheet. More particularly, we want to test the end-to-end performance of the sensors through a representative study to see how it performs in the real world on a set of people that it was not trained on. In order to investigate potential algorithmic bias, we will be collecting data on participants' sex and skin tone throughout the study.

The study will take place in the SEC, where a room will be set up with a series of sensors that will output the probability of seeing a person. The study will involve varying distances and lighting settings and is estimated to take approximately 2 minutes per person. Participants will be provided with snacks and required to fill out a brief consent form prior to participation. The sensors will take instantaneous images for processing, but none of the information will be saved to ensure privacy preservation.

If you have any questions or are interested in helping out with this study, please contact matthew_stewart@g.harvard.edu or yasmineomri@college.harvard.edu. For more information about our project, please see our high-level blog or this paper.

Name:

_____

Email:

_____

Which of these weeks would you be available to take part in this experiment? We will be reaching out by email closer to the date to organize concrete times.

- ☐ April 3-7
- ☐ April 10-14
- ☐ April 17-21
- ☐ April 24-28

Machine Learning Sensors Experimental Study Consent Form

The Edge Computing Group is looking for participants to take part in an experimental study we are devising to evaluate a new paradigm of intelligent sensor that we are developing.

We are looking into the next generation of sensors, ML sensors, which use on-device machine learning to extract useful information from the raw data before reporting it to the outer system. The ML sensor paradigm comes with significant benefits when it comes to modularity, composability, power efficiency, privacy and security, and more. Part of establishing the ML sensor paradigm is reimagining the conventional sensor data sheet. More particularly, we want to test the end-to-end performance of the sensors through a representative study to see how it performs in the real world on a set of people that it was not trained on.

Please review and confirm your agreement to the following:

What You Will Do in this Study: Participants will be requested to select their gender identity and skin tone (using the Monk Skin Tone Scale). You will be asked to input this information at the end of this consent form and prior to entering the study room. This information will remain anonymous, as no name or additional information about the participant will be recorded.

The participant will then enter the study room and stand in front of a group of six sensors at three varying distances. At each distance, three different light settings will be tested. One individual in the room will be altering the lighting conditions and another taking note of the sensor outputs. The sensor predictions (detection of a person) are recorded and coupled with the skin tone and gender identity to investigate potential algorithmic bias. We expect the experiment to take about 2-3 minutes per person.

Participation in this study is voluntary. You do not have to be in this study if you do not want to and you can quit the study at any time.

Risks and Benefits: There are no known risks to participants. Participants will be compensated with snacks for their time. The results of this study intend to be published and could lead to the development of more accurate and effective person detection technology, which could have a positive impact on various fields, such as security, healthcare, and transportation.

Confidentiality: No privacy-sensitive information will be recorded, such as images or personal details. Skin tone and gender identity will be deidentified prior to data analysis, ensuring participants remain entirely anonymous.

We will only report a summarized analysis of the variability of the sensors' prediction accuracies with skin tone and gender.

If you have questions about the survey or study, contact Matthew Stewart (matthew_stewart@g.harvard.edu), Yasmine Omri (yasmineomri@college.harvard.edu) or Professor Vijay Janapa Reddi (vj@eecs.harvard.edu).

I confirm that I am at least eighteen years of age, I understand and have read the points above, and consent to the collection, use, and sharing of your anonymous responses.

- ☐ I consent

Please select your gender identity:

- ☐ Male
- ☐ Female
- ☐ Transgender Male
- ☐ Transgender Female
- ☐ Non-binary
- ☐ Other (Please Specify)

Which Monk Scale skin tone most closely matches your skin color?

Please refer to the Monk Scale skin tone chart below or on the table. Ask a team member if you are having trouble identifying the appropriate value.

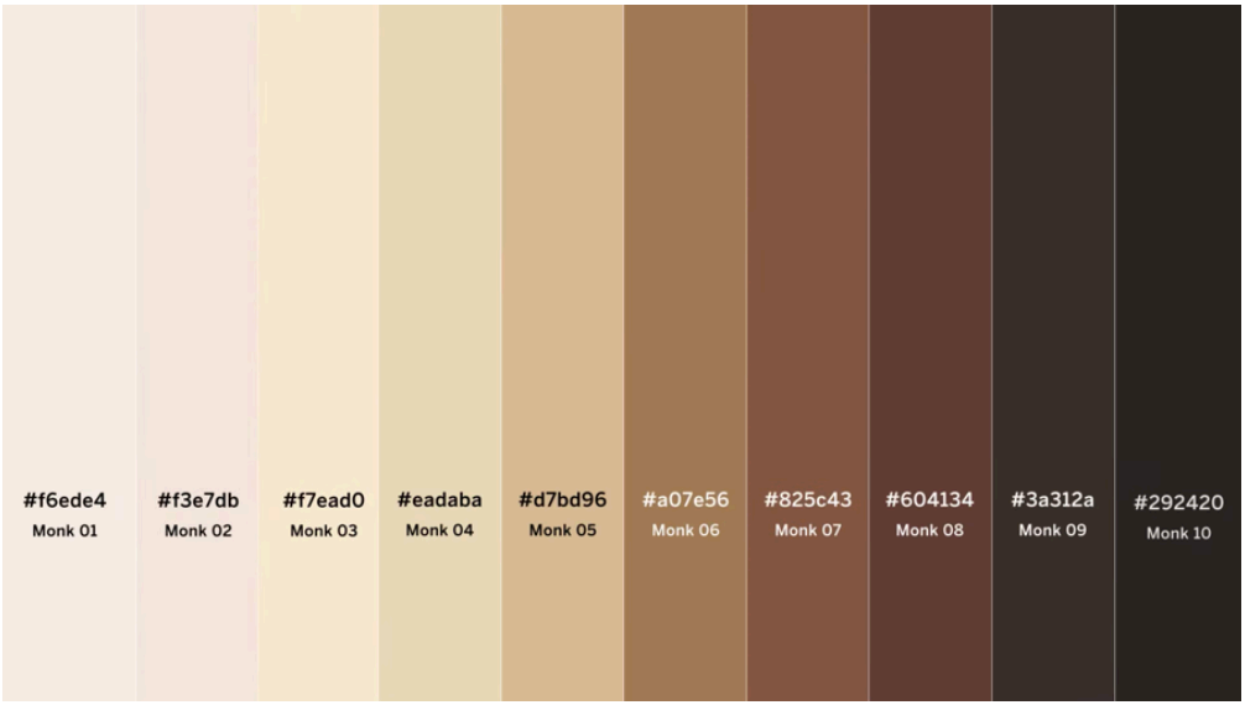

- [ ] 1
- [ ] 2
- [ ] 3
- [ ] 4
- [ ] 5
- [ ] 6
- [ ] 7
- [ ] 8
- [ ] 9
- [ ] 10