
EDGE IMPULSE: AN MLOPS PLATFORM FOR TINY MACHINE LEARNING

Shawn Hymel* Colby Banbury* Daniel Situnayake Alex Elium Carl Ward Mat Kelcey Mathijs Baaijens Mateusz Majchrzycki Jenny Plunkett David Tischler Alessandro Grande Louis Moreau Dmitry Maslov Artie Beavis Jan Jongboom Vijay Janapa Reddi

ABSTRACT

Edge Impulse is a cloud-based machine learning operations (MLOps) platform for developing embedded and edge ML (TinyML) systems that can be deployed to a wide range of hardware targets. Current TinyML workflows are plagued by fragmented software stacks and heterogeneous deployment hardware, making ML model optimizations difficult and unportable. We present Edge Impulse, a practical MLOps platform for developing TinyML systems at scale. Edge Impulse addresses these challenges and streamlines the TinyML design cycle by supporting various software and hardware optimizations to create an extensible and portable software stack for a multitude of embedded systems. As of Oct. 2022, Edge Impulse hosts 118,185 projects from 50,953 developers.

1 INTRODUCTION

Machine learning (ML) has become an increasingly important tool in embedded systems for solving difficult problems, enhancing existing Internet of Things (IoT) infrastructure, and offering unique ways to save power and bandwidth in sensor networks. ML inference on *TinyML* systems has facilitated the development of technologies in low-power devices such as wakeword detection (Gruenstein et al., 2017), predictive maintenance (Susto et al., 2015), anomaly detection (Koizumi et al., 2019), visual object detection (Chowdhery et al., 2019), and human activity recognition (Chavariaga et al., 2013). According to ABI Resarch, a global technology intelligence firm, the “installed base of devices with edge AI chipset will exceed 5 billion by 2025.” Additionally, the embedded ML market is expected to reach US\$44.3 billion by 2027 (abi, 2021).

Despite the promising advances, the embedded ML development ecosystem has lagged behind the demand for applications. The embedded ML development workflow often requires specific expertise. For instance, embedded ML developers often have to learn a new set of tools for training new models and porting them to an embedded framework written in C or C++, while managing conflicting library dependencies. Additionally, hardware vendor specific frameworks often lock a developer into a particular ecosystem,

which limits the flexibility and scalability of an application.

While popular frameworks such as TensorFlow Lite for Microcontrollers (TFLM) (David et al., 2021) help address optimization and compression of neural networks for embedded devices, adoption has been slow due to challenges that are unique to the *embedded* machine learning ecosystem. Broadly, these include the following:

1. *Data collection challenge.* There is no large-scale, curated, public sensor data set for the embedded ecosystem. Currently, it is difficult to efficiently collect, and analyze such datasets from a rich variety of sensors. Additionally, data cleaning and labeling are essential to ML development but are expensive, labor intensive processes without tooling or automation.
2. *Data preprocessing challenge.* Digital signal processing (DSP) is a critical stage of the ML stack and has strong interactions with the ML model, that are sometimes hard to quantify. Yet there are a lack of automated machine learning tools for the embedded ecosystem that include the DSP component, which hinders the development of efficient preprocessing methods for these systems by non-domain experts.
3. *Development challenge.* Matching TensorFlow and TFLM dependency versions across training and deployment infrastructure is challenging. To ensure functional reliability, it is often necessary for developers to have extensive knowledge in multiple domains, such as Python, machine learning, TensorFlow, C/C++, and embedded systems.
4. *Deployment challenge.* Scaling ML deployment is hampered by the rich heterogeneity of embedded architectures and development frameworks, which restricts code portability, particularly as a result of architecture-

*Equal Contribution

Corresponding author: <cbanbury@g.harvard.edu>.

Colby Banbury and Vijay Janapa Reddi are with the John A. Paulson School of Engineering and Applied Sciences, Harvard University. All others are with Edge Impulse.

Edge Impulse website: <http://www.edgeimpulse.com/>
Proceedings of the 6th MLSys Conference, Miami Beach, FL, USA, 2023. Copyright 2023 by the author(s).

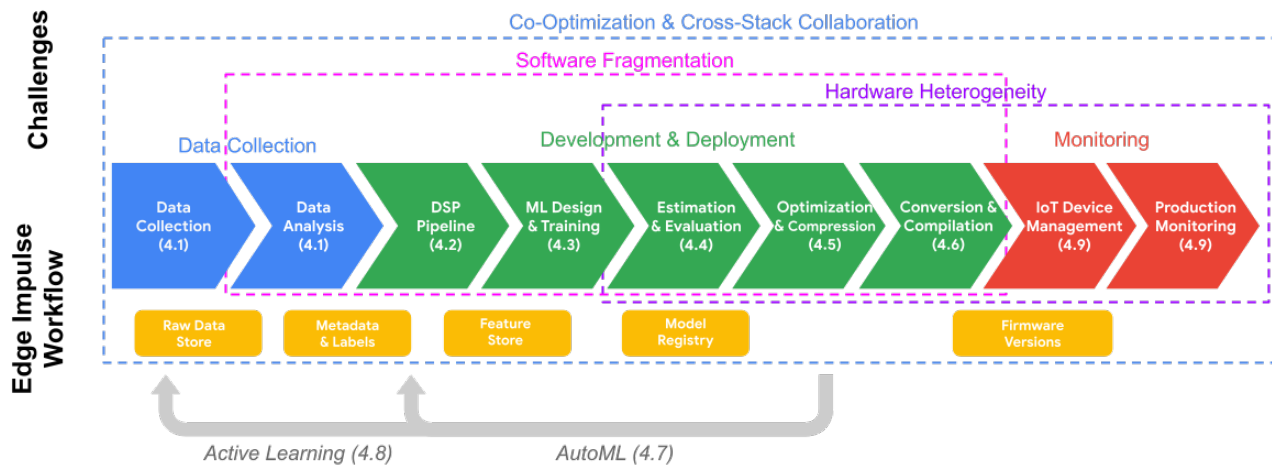


Figure 1. The challenges associated with the ML Workflow and features of Edge Impulse that solve those challenges.

specific optimization strategies. In addition, the absence of automated machine learning (AutoML) tools to assist non-domain experts in developing model architectures for embedded systems restricts accessibility.

5. *Monitoring challenge.* In contrast to traditional cloud-based machine learning systems that rely on mature software and hardware stacks, there is no unified MLOps framework for programmatically updating datasets, training models, and deploying them to embedded devices. Additionally, few benchmarking tools exist to quantify model performance on diverse embedded architectures that are highly heterogeneous.

Edge Impulse, an online platform designed to simplify the process of collecting data, train deep learning models, and deploy them to embedded and edge computing devices, allows us to address these aforementioned issues. Edge Impulse targets customers in the business sector who want to develop edge machine learning (ML) solutions for a variety of problems. However, the Edge Impulse platform also facilitates a research- and classroom-friendly environment.

Figure 1 illustrates the end-to-end ML workflow of Edge Impulse. Edge Impulse simplifies the process of data collection and curation for users and streamlines the training and evaluation of models. Users can interact with the training and deployment process via a combination of a web-based graphical user interface (GUI) and an API. Edge Impulse also provides an extensible and portable C/C++ library that encapsulates the preprocessing code and trained model to make inferencing simple across a wide range of target devices as well as a number of target-specific optimizations to reduce inference time and model memory consumption.

Edge Impulse offers several key technical contributions that are unique. The first contribution is a data collection system that helps users collect and store training and test data

alongside their model and deployment code. Rather than relying on prebuilt datasets or requiring users to construct their own data gathering technology, Edge Impulse offers a variety of methods to gather data in real-world environments. The second contribution is pairing preprocessing feature extraction with deep learning, which allows users to explore a range of possible solutions to their individual problem or task. The Edge Optimized Neural (EON) Tuner assists in this task by automatically exploring a user-defined search space of both preprocessors and ML models. The third contribution is an extensible and portable inferencing library that can be deployed across a wide range of edge and embedded systems. The EON Compiler removes the overhead required by the TFLM interpreter, thereby reducing usage of the limited RAM and flash space.

In this paper, we outline the challenges of developing and deploying ML models to embedded devices from an industry practice perspective. Next, we describe the architecture and use cases for a platform designed specifically to address these obstacles. We provide several examples to illustrate how this platform has been utilized successfully in industry, academia, and research institutions to develop novel machine learning-based solutions. Finally, we provide an evaluation of the performance and portability of the platform-generated inference code.

2 EMBEDDED ECOSYSTEM CHALLENGES

In Section 1, we highlight the challenges faced by an embedded ML developer. In this section, we highlight the challenges posed by the Embedded ecosystem which make platform and framework development difficult.

2.1 Device Resource Constraints

TinyML systems often have very limited computational capabilities, due to their small size, cost, and energy budget. Microcontrollers, which are the most common and general purpose processors in the TinyML space, often have much lower clock speeds and fewer architectural features (Banbury et al., 2021b) than their mobile or server-class counterparts. This becomes a challenge when trying to keep pace with the flow of data from a sensor or hit latency constraints.

ML workloads often require gigabytes of working memory and storage to store activations and model weights, but TinyML systems are often equipped with only a few hundred kilobytes of SRAM and a few megabytes of eFlash. This enforces a strict constrain on the models. TinyML systems often have very flat memory hierarchies, due to small or non-existent caches and often no off-chip memory (Banbury et al., 2021b). This means the typical data access patterns that neural networks have been designed around no longer apply, which has forced the design of new model architectures (Banbury et al., 2021b; Lin et al., 2020b).

Finally, many TinyML applications operate on battery power, and the battery life of the system directly impacts the usefulness of the application. Due to the small size and cost of TinyML systems, these batteries are often small and low capacity (e.g. a coin cell). Due to the limited energy budget, any wireless transmission can quickly deplete the battery (Siekkinen et al., 2012). Since data is often only transmitted once a specific prediction is made (e.g. “OK Google”, “Alexa”, “Hey Siri”, etc.), false positives contribute to battery drain with no benefit. Therefore, the accuracy of a model can directly impact the energy consumption of the system. These device constraints force TinyML application developers to leverage every compression and optimization technique at their disposal, which, as described in the next sections, poses it’s own challenges.

2.2 Hardware Heterogeneity

Despite resource limitations being fairly constant across TinyML hardware, the embedded computing systems themselves are quite diverse. TinyML devices range from microcontrollers (Eggimann et al., 2019) and digital signal processors (Gruenstein et al., 2017), to application specific accelerators (Prakash et al., 2022; ARM, 2022) and neuromorphic processors (Qiao et al., 2015). The STM32 32-bit Arm Cortex MCU family alone, for example, includes 17 series of microcontrollers. Each STM32 microcontroller series is based on an ARM processor core that is either Cortex-M7, Cortex-M4, Cortex-M33, Cortex-M3, Cortex-M0+, or Cortex-M0 (STM). Their capabilities can also vary at the instruction set architecture level. The same is true of other vendors. Each hardware platform supports different deployment processes, model types, numerical formats, and

memory access patterns, which often makes TinyML applications difficult to port across devices. This complexity is exacerbated when a creating an application at scale, which must be deployed to a wide variety of devices, each with their own libraries and deployment method.

2.3 Software Fragmentation

Due to TinyML’s infancy, the software stack has not yet reached a state of stability concerning particular formats and best practices. Occasionally, TinyML applications are deployed with a full operating system (OS) like Linux, a real-time OS like Zephyr (Zep), an inference framework like TFLM (David et al., 2021), or even a bare-metal implementation as a C++ library with no external dependencies. This diversity restricts the interoperability of new optimizations and tools A standard TinyML training pipeline incorporates tools and techniques from multiple sources, resulting in a tangled web of software versions and ports that can hinder collaboration, portability, robustness, and reproducibility.

2.4 Co-Optimization and Cross-Stack Collaboration

Each software and hardware optimization depends on the other layers of the development and deployment stack to be effective. This necessitates a complex optimization problem with many interconnected knobs to tune for optimal performance, as TinyML applications have stringent constraints.

Even without the addition of model hyper-parameters and optimizations such as quantization, applying consistent preprocessing across projects is a complicated jumble of hyperparameters that often requires deep, domain-specific insights into a signal. Due to this inherent complexity, TinyML development is a time consuming process that requires a wide range of specific technical expertise that is often not readily available in the industry. In addition, the cross-product of options and versions at each layer complicates collaboration and the reproducibility of ML applications.

Additionally, data consistency poses it’s own challenges, especially when using an internal or collected dataset. Significant operational challenges are posed by maintaining train, validation, and test splits, adding or removing individual samples, and preserving metadata. To facilitate large-scale collaborative projects and aid in the resolution of the ML reproducibility crisis (Hutson, 2018), one must version control the data, preprocessing, model, and deployment code while tracking a complex web of external dependencies.

3 OVERVIEW AND DESIGN OBJECTIVES

Edge Impulse is a combination of software-as-a-service, developer tooling, embedded software, and documentation to help embedded development teams create software that makes use of embedded machine learning at scale. At the

time of writing, Edge Impulse is being used by 50,953 developers in 118,185 projects, of which, 3,219 have been made public; it is in use at over 5500 enterprise organizations, excluding universities and other educational institutions.

Edge Impulse is designed and engineered according to the following seven guiding principles, based on the developer challenges as well as the embedded ecosystem challenges that we described in Section 1 and Section 2, respectively.

1. Accessible. Edge Impulse’s primary objective is to make embedded machine learning simpler and more accessible while focusing on producing Tiny ML solutions for resource-constrained devices (Section 2.1). This effectively broadens the pool of potential embedded ML developers by helping the embedded engineers with ML and the ML engineers with embedded systems (Section 2.4).

2. End-to-end. Edge Impulse provides users the ability to easily experiment with the end-to-end ML workflow holistically (Figure 1). Using Edge Impulse, one (or a team) could collect a dataset, train a efficient, optimized model, evaluate its performance, and deploy embedded firmware.

3. Data-centric. Edge Impulse prioritizes a data-centric approach because data collection and analysis has been historically slowed down in the ML pipeline. Given the scarcity of sensor datasets in the embedded ecosystem (Challenge #1, Section 1), Edge Impulse enables users to ingest data from various sources. Therefore, Edge Impulse encourages a data-centric approach to ML development, rather than (over) emphasizing a model-centric approach.

4. Iterative. Since cross-stack optimization is critical (Section 2.4), Edge Impulse promotes short developer feedback loops that allows developers to quickly experiment and iterate over different design space optimizations. To this end, Edge Impulse strives to provide a rich set of AutoML tools. Short design cycles and AutoML tools removes some of the burden of expertise (Section 2.4).

5. Easy integration and extensible. Edge Impulse prioritizes integration and extensibility to address the challenge of software fragmentation (Section 2.3) and cross-stack collaboration (Section 2.4). Experts should be able to connect the technology with their preferred downstream stacks, ideally using open standards where possible, and deploy to a wide variety of embedded and edge platforms (Section 2.2).

6. Team Oriented. To scale well (Challenge #5, Section 1), Edge Impulse facilitates the teamwork and communication required for many embedded machine learning projects by supporting multiple users on projects, versioning of projects, and sharing of projects. In addition, it is well-documented with accessible content that serves every user type.

7. Community supported. Finally, the technology should promote a strong commitment to the community and exist

within an ecosystem of community users, tools, and content.

When a user creates a project, they are guided through the process of gathering data, analyzing that data, creating a DSP preprocessing block, training a machine learning model, evaluating that model, and ultimately deploying it to a hardware platform of their choice. These steps are shown in the ML workflow in Figure 1. Figure 2 shows a user’s view inside an Edge Impulse Studio project with the ML workflow steps shown on the left side of the page. Projects in Edge Impulse are divided into a series of *blocks* that represent the dataflow. For this keyword spotting example, data arrives (i.e. from a microphone) in the left block labeled “Time series data” is preprocessed into Mel-frequency cepstral coefficients (MFCCs) in the middle block, and then sent to a NN for inference in the block labeled “Classification (Keras).” In the rest of the project design, users can modify block parameters to adjust functionality or create their own blocks to transform the data. s

With the design objectives established, a few things are beyond the scope of Edge Impulse. Edge Impulse is not intended to eliminate the need for a design process informed by domain expertise, stakeholder consultation, and machine learning workflow insight. For instance, a team of engineers utilizing Edge Impulse must still assess the suitability of machine learning as a solution to the problem they are attempting to solve. The team must have the necessary domain expertise to comprehend the problem and develop a responsible solution. They must still comprehend the general nature of the machine learning (ML) workflow, including iterative development and adopt appropriate evaluation metrics.

4 IMPLEMENTATION

In this section, we describe all the different aspects of the end-to-end flow that Edge Impulse supports, as illustrated in Figure 1. For each stage, we present the rationale for the stage and describe its implementation specifics.

4.1 Data Collection and Analysis

Since every ML project begins with data that is often hard to gather easily, Edge Impulse provides a number of features designed to help users collect data, manage their dataset, and perform feature extraction through digital signal processing (DSP). Edge Impulse projects can accept data stored in a several file formats: CSV, CBOR, JSON, WAV, JPG, or PNG. The platform also offers several methods to help users gather data for their project, including command line interface (CLI) tools that interface with device firmware to ingest data in real time and web-based API to upload data directly or from an existing cloud-based store (e.g. AWS S3 bucket). A GUI allows users to visualize training and test set split as well as class allocation grouped into buckets.

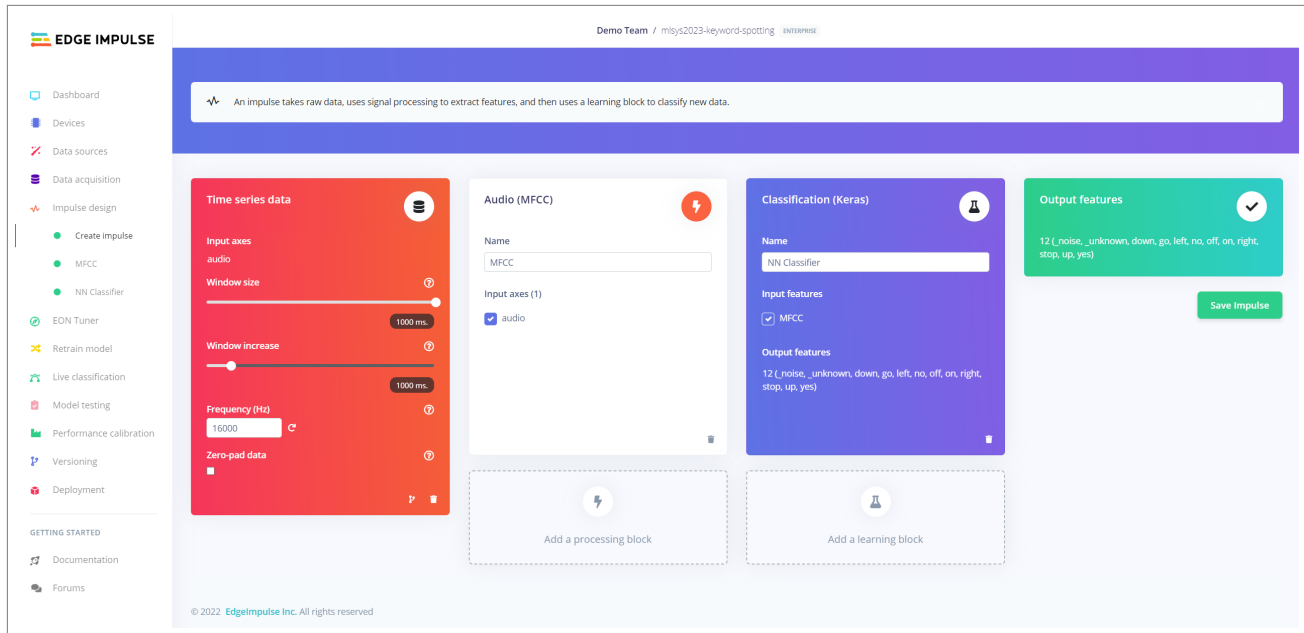


Figure 2. Screenshot showing the user’s view inside an Edge Impulse project where the *blocks* are connected depicting the dataflow.

Users can also examine raw data in each sample through time series plots or images, depending on the data type.

4.2 DSP Pipeline

Many of the embedded machine learning applications or use cases rely on sensor data preprocessing. Edge Impulse offers the ability to perform various preprocessing of raw signal data automatically prior to use in training or inference. This preprocessing step is known as *digital signal processing (DSP)* in the Edge Impulse workflow. By preprocessing raw data, model size can often be reduced, and preprocessing can incorporate more efficient algorithms than can be realized with typical neural net architectures. For example, an FFT is an $O(n \cdot \log(n))$ algorithm for extracting frequency information, whereas using 1D convolutional layers to accomplish the same thing would require $O(n^2)$ operations.

Edge Impulse simplifies the preprocessing workflow by providing a rich array of continuum blocks that trade off model size and complexity (edg), a visual explorer for tuning DSP block hyperparameters (frame length, stride, window size, number of coefficients, etc.), and estimates of memory and latency requirements for a given choice of hyperparameters.

Edge Impulse offers sensible defaults for a variety of tasks to ensure minimal knowledge is required by users, though domain experts can choose preprocessing steps and hyperparameters that reduce the amount of training noise. Additionally, users can automatically select these hyperparameters via the DSP autotune feature, or optimize them via the Eon Tuner (Sec. 4.7)

4.3 ML Design and Training

Traditionally, a user would need to write code to train a neural network on their data. But to make machine learning more accessible to everyone, Edge Impulse employs a visual editor that allows a user to train on their data without entering any code. There are preset neural network architectures that are suggested based on the type of data coming into the machine learning block. However, the layers or the network can be customized by the user. Advanced users can download the containerized code for the block to train locally or use expert mode to customize further using the Keras framework (Chollet, 2015) and Python.

Arbitrary combinations of building blocks (as shown in the Figure 2 screenshot) allows for rich flexibility in model architecture, but it is important that the model is trainable. Edge Impulse provides a number of subtle, but important, optimisation pieces to ensure stable training including, but not limited to, learning rate finding, classifier bias initialisation, best model checkpoint restoration. Such building blocks and optimizations make the training process accessible to machine learning novices while the extensibility of the expert mode allow domain experts to develop more complex ML models on Edge Impulse.

Edge Impulse provides a transfer learning (Pan & Yang, 2009) block for audio keyword detection. This allows the users to quickly develop a robust keyword spotting application, even when working with a relatively small dataset.

Edge Impulse maintains partnerships with silicon vendors who have developed specific neural network accelerator

hardware, such as the Syntiant NDP101. In addition to generating models for general purpose processors, Edge Impulse supports a variety of architecture-specific devices and optimizations, such as CMSIS-NN (Lai et al., 2018) to maximize the performance and minimize the memory footprint of neural networks on Cortex-M processor cores, thus alleviating many issues found with hardware heterogeneity.

Edge Impulse also supports several unsupervised learning algorithms to tackle anomaly detection problems. At the moment, Edge Impulse uses K-means clustering and will support Gaussian mixture models (GMM) in the near future.

4.4 Estimation and Evaluation

Since embedded systems are resource-constrained, developers can benefit from having estimations of model inference latency time, RAM usage, and flash memory usage during the early-stage design space exploration. Edge Impulse uses Renode (Hoľenko, 2017) and device-specific benchmarking to produce estimates of preprocessing and model inference times. Models are also compiled with varying options (non-quantized vs. quantized, TFLM vs. EON Compiler) to produce initial insights into RAM and flash memory usage.

Edge Impulse offers a number of tools to assist in evaluating the effectiveness of model performance. A confusion matrix can be generated from the holdout set to provide overall or per-class accuracy and F1 scores. For supported hardware, new data can be collected to perform live inference. Such evaluation options assist users in identifying trade-offs between model performance and model size and latency.

In addition to model evaluation, Edge Impulse enables post-processing evaluation and tuning using a tool known as *performance calibration* (Situnayake, 2022) for projects that identify events in streaming data. The tool accepts an input of user-supplied raw data or synthetically generated data along with the trained model. Using a genetic algorithm, it suggests a number of optimal post-processing configurations that trade off false acceptance rate (FAR) and false rejection rate (FRR). Suggesting optimal post-processing methods significantly reduces the engineering risk associated with a project and increases the quality of its performance.

4.5 Compression and Optimization

Many different optimization types can be used to improve the performance of ML and DSP algorithms when deployed to edge devices. Several types of optimization are supported by Edge Impulse, either out-of-the-box or via extensibility. The optimization areas are model compression and optimization, code optimization, and device-specific optimization.

Model compression and optimization techniques are applied either during or after training and result in models with a reduced size or computational burden when deployed to edge

devices. Compression techniques available out-of-the-box in Edge Impulse include fully int-8, weight and activation quantization (Jacob et al., 2017) and operator fusion (goo, 2022). Quantization-aware training is supported when converting a model to Brainchip’s Neuromorphic format (bus, 2022).

Code optimization involves optimizations to run an algorithm or set of algorithms on a given target. This includes model-specific code generation (via EON Compiler (Jongboom, 2022)), ML kernels optimized for particular processor architectures (e.g. ARM CMSIS-NN (Lai et al., 2018)), and quantization optimized DSP algorithms. The software development kit (SDK) is designed to make use of available optimizations depending on the compiler flags that are set.

Device-specific optimizations are those that apply to specific targets due to the requirement of hardware support. Examples include the training and conversion of spiking neural networks (supported for specific targets via integration) and sparse neural networks. Additional targets and optimizations can be added using the platform’s extensibility via custom processing, learning, and deployment blocks.

Edge Impulse’s EON compiler (Jongboom, 2022) compiles TFLM neural networks to C++ source code. The EON Compiler eliminates the need for the TFLM interpreter by generating code that directly calls the underlying kernels and enables the linker to eliminate unused instructions. This effort reduces the RAM and ROM usage for neural network implementations, as we show in Section 5.3.

4.6 Conversion and Compilation

Edge Impulse offers several possibilities for DSP and model deployment to target embedded and edge devices, such as standalone C++ library, Arduino library, process runner for Linux, WebAssembly library (Web), and precompiled binaries for a variety of supported boards. A deployed project includes both DSP preprocessing and trained machine learning model that have been optimized for a given architecture.

Edge Impulse provides a firmware SDK for collecting data directly on a device that will be used at inference time. This SDK can be built by a user or is available in binary format for a variety of popular microcontrollers, such as Arduino, Raspberry Pi Pico, etc. As a library, the SDK contains several public-facing functions for performing inference (Hymel, 2022). The precompiled binary presents a simple set of AT commands for usage over a serial port.

The SDK provides a pathway to out-of-the-box operation that uses a combination of code generation, macros, and runtime checks to include more efficient algorithms and optimizations where possible, but it falls back on pure C++ where needed to run on a wide range of processor architectures. Porting to a new processor requires an allocator for

the desired memory pool to service the SDK (typically a wrapper to `malloc`). `printf` and timing functions help with debug and profiling, but these are not strictly necessary.

In addition to microcontrollers, the Edge Impulse inference SDK can be built and run on any Linux board utilizing x86 or ARM architectures (ei2, 2022b). The DSP block and trained model is downloaded from the user’s project in EIM format, which is a compiled, native binary application that exposes the I/O interface for use by any number of programming languages (Python, Go, C++, Node.js, etc.).

4.7 AutoML

The accuracy of any deep learning system depends critically on identifying a proper choice of hyperparameters. The inherent resource constraints on embedded targets also limits the selection of hyperparameters. For example, a trade-off has to be made between allocating resources for the DSP and deep learning algorithms. For a novice user, the relationship between hyperparameters can often be difficult to grasp. Therefore, Edge Impulse provides a suite of automated machine learning (AutoML) techniques to assist non-experts in creating usable models and tune hyperparameters.

To ensure low burden on the user, Edge Impulse’s EON Tuner (Jenny Plunkett, 2021) assists in the hyperparameter selection process while taking into account available RAM, ROM, and CPU clock speed of the target device. The EON Tuner helps select a number of hyperparameter configurations, including DSP preprocessing settings. It then trains the associated models to determine their accuracy. From these results, the user can select a preferred configuration (e.g. based on accuracy/F1 score or resource usage) and update the associated project to this configuration.

Figure 3 shows the user’s perspective after EON Tuner has been successfully run. Note the stacked bar plots showing the estimated latency, RAM, and flash usage (Sec. 4.4) for each combination of preprocessing (DSP) and model blocks based on the selected target (e.g. Arduino Nano 33 BLE Sense). Model details, including the specific DSP and NN configuration, are shown, which allows user to select the best combination of blocks to meet their accuracy requirements within the desired hardware constraints.

To select hyperparameter configurations, the EON Tuner combines a random search algorithm (Bergstra et al., 2011) with a heuristic to quickly estimate the performance of the configurations. Future work includes optimizing search methods using a combination of a Bayesian (Eggenberger et al., 2013) and Hyperband (Li et al., 2017) search algorithms. Users have the option of overriding the default search algorithm with their own search methods.

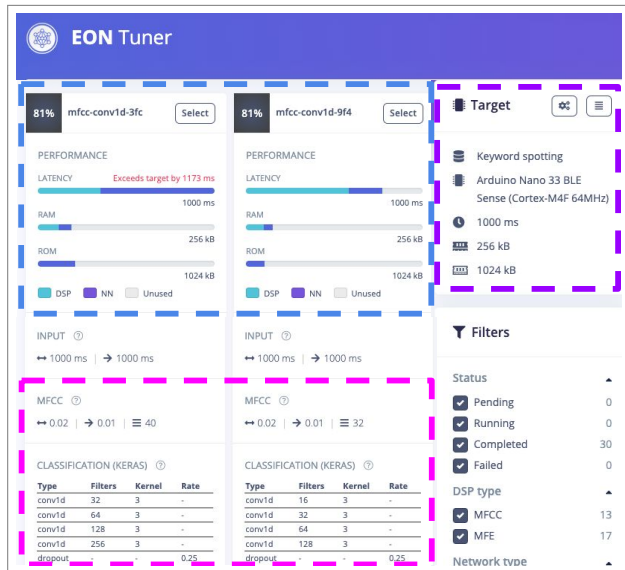


Figure 3. Screenshot of the EON Tuner. Features are annotated with color coded dotted boxes that correspond to the challenges in Figure 1. Purple (top right): The tuner allows users to select the target hardware, which will then inform the constraints set on the search. Blue (top left): The tuner computes the configuration’s accuracy and predicts the resource consumption of the DSP and NN components. Pink (bottom): The tuner searches for optimal DSP and NN combinations and displays their configuration.

4.8 Active Learning

Datasets can be iteratively improved by leveraging a partially trained model to aid in labeling and data cleaning in a process called active learning (Moreau, 2022). Edge Impulse employs an active learning loop for the embedded sensor ecosystem where you can: (1) train a model on a small, labeled subset of your data, (2) generate semantically meaningful embeddings using an intermediate layer of the trained model, (3) visualize the embeddings (non-labeled and labeled samples) in 2D space using a dimensionality reduction algorithm (Umap (McInnes et al., 2018) or t-SNE (Van der Maaten & Hinton, 2008)), and (4) manually or automatically label or remove samples based on their proximity to existing class clusters. This process can drastically speed up the labeling and data cleaning processes, which can lead to major gains in model performance.

4.9 Extensibility

Edge Impulse supports the majority of the workflow shown in Figure 1, with the exception of IoT device management and production monitoring. However, all Edge Impulse functionality is exposed via publicly accessible REST APIs (ei2, 2020), which allows users to automate the data collection, model training, and deployment processes. This API can be integrated into custom workflows and third party so-

| Platform | Processor | Clock | Flash | RAM |
|-----------------------|----------------|---------|-------|--------|
| Nano 33 BLE Sense | Arm Cortex-M4 | 64 MHz | 1 MB | 256 kB |
| ESP-EYE (ESP32) | Tensilica LX6 | 160 MHz | 4 MB | 8 MB |
| Ras. Pi Pico (RP2040) | Arm Cortex-M0+ | 133 MHz | 16 MB | 264 kB |

Table 1. Embedded platforms used for evaluation.

lutions to augment IoT device management and production monitoring, such as Microsoft Azure IoT (Azu).

In addition, users are able to create their own blocks via Docker images to transform raw data taken from an existing store (e.g. AWS S3), perform feature extraction via DSP, train a custom ML model, or deploy a model. Finally, the Edge Impulse inferencing SDK library allows users to develop complete embedded ML solutions that include a variety of model compression and optimization techniques.

Futhermore, Edge Impulse can integrate with existing ML development pipelines via the Edge Impulse Python SDK (Situnayake, 2023). This allows users to use specific features, such as profiling or deployment (Sec. 4.4 & 4.6), without needing to use the graphical interface.

4.10 Scalable Infrastructure

Edge Impulse employs AWS Elastic Kubernetes Service (Man) to dynamically scale compute resources based on workload requirements. All workloads are containerised, which has proven vital for efficient dependency management. Often, ML software infrastructure requires a wide-range of dependencies and versions of dependencies that are not always mutually compatible. The choice of Kubernetes over a vendor-specific tool, such as AWS Elastic Container Service (ama), is to enable migration of the Edge Impulse infrastructure to a different cloud provider or on-premise with a reasonable (1-6 months) amount of effort.

5 PERFORMANCE EVALUATION

ML development often narrowly focuses on the model performance in isolation (Richins et al., 2021) due to the complexity of co-optimization (Sec. 2.4). However, the DSP stage can be a dominant factor in the overall latency and memory consumption of a TinyML application. Edge Impulse is designed to quantify the DSP overhead and allow users to explore the rich DSP and NN co-design space.

In this section, we characterize the latency, SRAM, and flash consumption of TinyML workloads across multiple devices, optimizations, and AutoML defined configurations, thereby showing Edge Impulse’s ability to address the challenges of hardware heterogeneity (Sec. 2.2), software fragmentation (Sec. 2.3), and cross-stack optimization (Sec. 2.4).

5.1 Experimental Setup

We evaluated three representative hardware designs. The details for the platforms are shown in Table 1. We chose

| | Nano 33 BLE Sense | | ESP-EYE | | Ras. Pi Pico | |
|--|-------------------|--------|---------|--------|--------------|---------|
| | Float | Int8 | Float | Int8 | Float | Int8 |
| Keyword Spotting (KWS) inference times | | | | | | |
| Preprocessing | 141.65 | 138.76 | 305.53 | 304.11 | 590.74 | 590.87 |
| Inference | 2866.11 | 322.71 | 648.42 | 314.14 | 5700.03 | 1117.65 |
| Total | 3007.91 | 461.62 | 954.02 | 618.35 | 6290.95 | 1708.71 |
| Visual Wake Words (VWW) inference times | | | | | | |
| Preprocessing | - | 9.98 | 24.25 | 9.07 | - | 56.44 |
| Inference | - | 754.74 | 2309.15 | 662.85 | - | 2205.76 |
| Total | - | 816.56 | 2346.03 | 702.63 | - | 2286.68 |
| Image Classification (IC) inference times | | | | | | |
| Preprocessing | 1.36 | 1.14 | 1.09 | 1.03 | 4.57 | 6.46 |
| Inference | 1518.64 | 229.54 | 340.45 | 191.15 | 3048.05 | 554.04 |
| Total | 1520.25 | 232.56 | 341.62 | 197.36 | 3048.05 | 561.86 |

Table 2. Preprocessing and inference times (in milliseconds). ‘-’ indicates the model did not fit due to flash or RAM constraints.

these platforms for their differences in clock speeds, flash storage and RAM capacity. We chose several models to evaluate the platforms to demonstrate the capabilities of each. These models were created to solve three tasks outlined in the MLPerf Tiny Benchmark (Banbury et al., 2021a): keyword spotting (KWS), visual wake words (VWW), and image classification. KWS is a common task in embedded devices that require wake word detection, such as “Alexa” or “OK Google.” We chose a DS-CNN model (Sørensen et al., 2020) that achieved at least 78% on a test set from the Google Speech Commands dataset (Warden, 2018). For the VWW task, MobileNetV1 was trained using the visual wake words dataset (Chowdhery et al., 2019), which was derived from the Microsoft COCO dataset (Lin et al., 2014). This dataset is a balanced set of “person” and “non-person” images used to train an image classification model. We achieved at least 72% accuracy on a hold-out set. Finally, we trained a simple convolutional neural network (CNN) on CIFAR-10 (Krizhevsky et al., 2009).

5.2 Cross-Hardware Inference Latency Comparison

Table 2 displays three sets of *end-to-end* timing results using provided hardware timers. The table shows the latency of the KWS, VWW, and image classification tasks for both floating point and quantized integer (8-bit) models across the three platforms. The preprocessing and classification tasks are timed from within the Edge Impulse SDK, and the total time is taken by measuring the difference between timestamps taken around the call to run classification, which is a combination of preprocessing and inference plus some overhead not measured in either preprocessing or inference.

On some tasks, such as keyword spotting, the preprocessing time can easily equal or exceed the inference time of the unoptimized model. Therefore, optimizing the network inference via quantization, etc. will not yield the typical magnitude of latency reduction. Edge Impulse allows users to look at the end-to-end performance of a task, rather than focus only on isolated network performance.

Edge Impulse helps users experiment with preprocessing

| Preprocessing | Model | ↓ Acc. | Latency (ms) | | | RAM (kB) | | | Flash (kB) | | |
|------------------------|-----------------------|--------|--------------|--------|-------|----------|--------|-------|------------|--------|-------|
| | | | DSP | Infer. | Total | DSP | Infer. | Total | DSP | Infer. | Total |
| MFE (0.02, 0.01, 40) | MobileNetV2 0.35 | 85% | 332 | 2420 | 2752 | 25 | 468 | 493 | - | 2242 | 2242 |
| MFCC (0.02, 0.01, 40) | 4x conv1d (32 to 256) | 75% | 770 | 437 | 1207 | 30 | 35 | 65 | - | 645 | 645 |
| MFCC (0.02, 0.01, 32) | 4x conv1d (16 to 128) | 73% | 579 | 197 | 776 | 26 | 20 | 46 | - | 221 | 221 |
| MFE (0.02, 0.01, 32) | 3x conv1d (32 to 128) | 72% | 319 | 174 | 493 | 21 | 31 | 52 | - | 231 | 231 |
| MFE (0.02, 0.02, 32) | 2x conv1d (32 to 64) | 70% | 163 | 109 | 272 | 14 | 17 | 31 | - | 125 | 125 |
| MFCC (0.05, 0.025, 40) | 3x conv1d (16 to 64) | 69% | 327 | 48 | 375 | 19 | 10 | 29 | - | 98 | 98 |
| MFE (0.05, 0.025, 32) | 2x conv1d (32 to 64) | 69% | 161 | 67 | 228 | 15 | 14 | 29 | - | 56 | 56 |
| MFE (0.032, 0.016, 32) | 2x conv1d (16 to 32) | 66% | 217 | 91 | 308 | 16 | 19 | 35 | - | 56 | 56 |

Table 3. Preprocessing blocks and models explored with EON Tuner for the keyword spotting task. Latency, RAM, and flash estimates from EON Tuner for the keyword spotting task on the Arduino Nano 33 BLE Sense (float32 inference, using TFLM).

and models to quickly iterate on designs to find acceptable solutions to such problems. Table 3 shows how users can choose to use different preprocessing blocks, Mel-filterbank energy (MFE) or MFCCs, and sweep different model architectures with the EON Tuner for optimizing latency, accuracy, RAM, and Flash storage. There is no ideal solution; the ultimate choice is up to the user as they know their deployment constraints. Edge Impulse simply automates the possibilities and displays suggested configurations.

5.3 Memory Optimization with EON Compiler

Embedded systems are constrained by their memory and storage capacity (Section 2.1), Table 4 details the estimated memory usage, RAM and flash, for all three tasks. The Edge Impulse EON Compiler removes the need for the TFLM interpreter for on-device inference, thus reducing the required RAM and flash in most cases. A consistent decrease in memory utilization is seen when enabling the EON Compiler as well as quantizing to an INT8 model. Quantization can decrease the accuracy of the model due to the lower precision, but in some instances (e.g. the image classification task) it improves the accuracy due to regularization. These optimizations do not impact the preprocessing stage.

| | Keyword Spotting | | | Visual Wake Words | | | Image Classification | | |
|---------------|------------------|-------|------|-------------------|-------|------|----------------------|-------|------|
| | RAM | Flash | Acc. | RAM | Flash | Acc. | RAM | Flash | Acc. |
| Preprocessing | 13.0 | - | - | 4.0 | - | - | 4.0 | - | - |
| FP (TFLM) | 115.8 | 148.0 | 78.5 | 398.4 | 904.4 | 81.1 | 195.8 | 107.5 | 70.9 |
| FP (EON) | 96.8 | 106.7 | | 327.7 | 861.4 | | 162.7 | 78.7 | |
| Int8 (TFLM) | 38.5 | 98.1 | 78.5 | 124.8 | 361.2 | 79.9 | 51.9 | 63.1 | 71.1 |
| Int8 (EON) | 36.4 | 65.3 | | 131.0 | 309.5 | | 44.0 | 42.1 | |

Table 4. Memory estimation (all memory estimates given in kilobytes, accuracy in percentage based on the holdout set). Flash utilization estimation not provided for DSP preprocessing.

5.4 Design Space Exploration with EON Tuner

Table 3 shows a number of configurations that are searched by the EON Tuner in order to find an optimal keyword spotting model. A user can find a model that balances the resources allocated to the DSP and NN stages in order to meet the hardware constraints of the application while maximizing accuracy. For example, the configurations on the 3rd and 4th lines from the bottom balance the DSP and NN

differently, leading to a slightly higher accuracy and lower latency (more NN) compared to less RAM and Flash consumption (more DSP). This process accelerates the initial exploratory phase of ML development and makes cross-stack optimizations accessible for novice ML developers.

6 ECOSYSTEM ENABLEMENT

Since its launch in late 2019, Edge Impulse has seen excitement around embedded ML deployments in a variety of domains. We highlight a few exemplar use cases here.

6.1 Education and Learning

Edge Impulse provides both a graphical user interface through Studio as well as an extensible web-based API. Consequently, it is well suited for classroom activities, as it provides a series of parameters, plots, and visualizations within the Studio to assist newcomers in building end-to-end embedded/edge machine learning systems.

We saw a large interest in embedded machine learning courses upon the delivery of two separate massively open online courses (MOOCs). Between September 2020 and June 2021, over 43,000 students enrolled in three of the Tiny Machine Learning courses on EdX ([Tin](#)). In less than two years, over 75,000 students have benefited from these courses alone ([Janapa Reddi et al., 2022](#)). Additionally, over 30,000 students enrolled in the two Embedded Machine Learning courses between February 2021 and October 2022 that use Edge Impulse ([Int](#)). In addition, throughout 2021 and 2022, the TinyML4D Academic Network ([tin, 2022](#)), which focuses on improving access to edge ML education and technology in developing countries by running a series of workshops for Africa, Asia, and Latin America regions, used Edge Impulse to teach embedded ML to professors, lecturers, and students. These 2022 workshops had 216 attendees from 48 different countries. The high attendance in the MOOCs along with the apparent growing desire for various schools to teach ML demonstrates the need for approachable tools for education.

6.2 Industry Use Cases, Research and Dissemination

Edge Impulse (and indeed, edge ML) is being used by a number of companies to solve unique problems. For example, Oura Ring, which uses ML models to identify sleep patterns (Sec. 8.1), relies on Edge Impulse. Edge Impulse is also being adopted by academic researchers as a tool for enabling scientific research and rapid problem-solving using ML. Since Edge Impulse simplifies the machine learning training and deployment process, it allows researchers to focus on solving the problem within their domain expertise rather than setting up an end-to-end system using another framework and dealing with low-level details. For example, recent work shows how to easily develop an “intelligent chemical sniffer, capable of detecting hazardous volatile organic compounds” (Shamim, 2022b). Another recent work identifies cancerous growths on oral tongue lesions using an embedded device (Shamim, 2022a). Other areas of research include creating low-power solutions to classify mosquito wingbeat sounds in order to identify mosquito species (Al-tayeb et al., 2022). There are several other such examples that are enabling non-ML scientists to easily adopt Edge Impulse to realize edge ML deployments.

6.3 Open Source, Community Development

Developing the open source embedded ML community is a must for the success of the TinyML industry. To this end, Edge Impulse maintains many open source repositories on GitHub (ei2, 2022a), including repositories related to machine learning, device firmware, and sample use-cases. As a result, there is a community of users who contribute code, issues, and bug-fixes beyond the Edge Impulse team.

This emphasis on community and open-source principles carry beyond the GitHub repositories. For example, one feature intended to share knowledge is the concept of *public projects*, where developers can make their work inside the Edge Impulse Studio public for other developers to review and clone, thus helping share applied techniques and best practices in Edge Impulse projects. When a project is set to public, it is also aggregated and searchable via the Projects page (ei2, 2022c) on the Edge Impulse website. This searchable index allows developers to sort, filter, and search for relevant examples and public work. At the time of writing, there are 3,219 public Edge Impulse projects available.

In addition to public projects, preprocessing and model inference libraries (along with the trained model) may be downloaded from Edge Impulse and shared via open source licenses (Apache 2.0, BSD 3-Clause, MIT). The ability to share projects and code allows researchers and developers to disseminate their trained ML systems for reproducibility.

Collaboration can also occur inside of an Edge Impulse project. Through the use of *Organizations*, more than one

| | Data Collection & Analysis | DSP & Model Design | Embedded Deployment | AutoML & Active Learning | IoT Management & Monitoring |
|---------------------|----------------------------|--------------------|---------------------|--------------------------|-----------------------------|
| Edge Impulse | ✓ | ✓ | ✓ | ✓ | ~ |
| Amazon SageMaker | ✓ | ~ | ~ | ✓ | ~ |
| Google VertexAI | ✓ | ~ | ✗ | ✓ | ~ |
| Azure ML & IoT | ✓ | ~ | ~ | ✓ | ✓ |
| Neuton AI | ✗ | ~ | ✓ | ~ | ✗ |
| Latent AI | ✗ | ✓ | ✓ | ✗ | ✗ |
| NanoEdge | ~ | ✓ | ✓ | ~ | ✗ |
| Imagimob | ✓ | ✓ | ✓ | ~ | ✗ |

Table 5. Comparison of supported features of MLOps platforms. ✓:Fully Supported, ~:Partially Supported, ✗:Not Supported.

developer can access or participate in the data upload, analysis, learning block, model creation and testing, or other aspects of the platform or project. Having multiple users can speed up project delivery of course, but the ability to spread knowledge and share best practices can positively impact users who might be starting out with edge ML.

7 RELATED WORK

MLOps platforms like Amazon SageMaker (Rauschmayr et al., 2021) and Google VertexAI (Ver) exist, but they cater to cloud-scale environments that are vastly different from the TinyML ecosystem. Microsoft’s Azure IoT (Klein, 2017) is a platform specifically designed to connect, analyze, and automate data analytics from the edge to the cloud. Edge Impulse uniquely encompasses a large slice of the embedded ecosystem. The platform operates as a dataset analysis and preprocessing tool, a training framework, an optimizer, and an inference engine. Additionally, it readily integrates with other cloud services.

TinyMLOps frameworks focus on the ML training and compression stages of the ML pipeline (Leroux et al., 2022), while Edge Impulse spans data collection to deployment. Neuton AI (Neu) allows users to train a TinyML model from CSV data, but it does not include features for data collection and limits the customization of the preprocessing and model architecture. Latent AI (Lat) provides a platform to compress and compile a model for efficient deployment but does not support data collection and training workflows. NanoEdge AI Studio (Nan) is a TinyML development tool primarily focusing on time series data exclusively for STM devices. Imagimob (ima, 2022) offers end-to-end TinyML development with various preprocessing options, AutoML for model selection, and model training, evaluation, and deployment to optimized C code. However, it does not offer a web-based API that allows for full MLOps automation. Table 5 illustrates the level of support the related MLOps platforms have for each set of relevant features.

TFLM (David et al., 2021) is the standard inference engine for TinyML use cases because it is flexible, open-source, and associated with the TensorFlow ecosystem. Edge Impulse’s EON Compiler uses less memory and storage compared to TFLM (Section 5.3). Arm’s uTen-

processor (uTe) and Microsoft’s ELL (The) are both TinyML inference engines that cross-compile to specific target embedded hardware, but both projects are no longer supported. TinyEngine (Lin et al., 2020a) and uTVM (mic) compile the network graph and perform inter-layer optimizations. TinyEngine is a research project and therefore it is not designed for production scale and uTVM currently supports very few boards. Vendor-specific inference engines, such as STM32Cube.AI (XCU), provides optimized inference but only support vendor-specific hardware, thus limiting application portability. All of these engines, however, can be adopted and integrated into Edge Impulse.

8 INDUSTRY CASE STUDIES

This section focuses on real-world insights from two industry case studies: sleep tracking with the Oura Ring and detecting heat exhaustion with SlateSafety Band V2.

8.1 Oura Ring

The Oura Ring is designed to track the user’s sleep patterns using an ML model that predicts the stages of sleep based on measured physiological signals. In order to improve their model, Oura conducted a large-scale sleep study (~100 participants) and collected a large dataset for model training.

8.1.1 Challenge

Incomplete, noisy, and inconsistent data is unavoidable when collecting real-world datasets and, given the scale of the study, it would be an arduous process to aggregate, scrub, and analyze the data before using it to train the next generation sleep tracking model. Many current platforms provide little assistance for the critical stages before training, and stand-alone data analysis and visualization tools require significant additional effort to set up ad hoc ML pipelines.

8.1.2 Edge Impulse Solution

Edge Impulse is able to ingest and align data from multiple sources and sensors by comparing sensor signatures, which simplifies and speeds up a typically manual and error-prone process. EI integrates analysis tools that enable domain experts to make design decisions on which sensors and features are meaningful for sleep prediction. Through this process, Oura created a model focusing only on heart rate, motion, and body temperature and achieves a best-in-class 79% correlation accuracy when compared to polysomnography and human scorers, which use expensive measurements, such as brain waves (de Zambotti et al., 2019).

8.1.3 Focus Areas for Future ML Systems Research

Current ML Systems researchers too often ignore data-centric techniques and instead focus on stages, like model

design, with easier-to-quantify metrics such as accuracy and latency (Mazumder et al., 2022). Areas for high-impact systems research include: (1) Aggregating and aligning data from unstructured sources and multiple sensors. (2) Visualizing statistically meaningful correlations of data from multiple sensors without hallucinations. (3) Leveraging non-technical domain experts for data cleaning and selection.

8.2 SlateSafety

The SlateSafety BAND V2 is a wearable device that monitors physiological signals of first responders and industrial workers. Due to the lack of reliable wireless connection in the deployment environment, SlateSafety required on-device inference to predict heat exhaustion in users.

8.2.1 Challenge

SlateSafety aimed to leverage existing hardware that was already in the field instead of going through the expensive process of developing an entirely new, ML-capable platform. Therefore, the resulting model had to run in real-time on an existing microcontroller with limited memory capacity.

8.2.2 Edge Impulse Solutions

Edge Impulse’s EON Tuner and Compiler were able to automatically design a custom model and deploy it efficiently to the existing microcontroller via an automatic over-the-air update. This means that SlateSafety could deploy a new feature to its users seamlessly and without a long development cycle around new hardware.

8.2.3 Focus Areas for Future ML Systems Research

Much of existing ML research target state-of-the-art hardware that has been designed with ML deployment in mind. However, in practice, ML-based features are deployed to existing systems. In order to enable broader adoption of AI and prevent existing hardware from being replaced and thrown out (Prakash et al., 2023), ML researchers need to develop optimization techniques that can be backward compatible with past generations of hardware.

9 CONCLUSION

Edge Impulse is a framework focused on building machine learning systems for resource-constrained devices. The framework is built around the principles of accessibility, data-centric co-optimization, and cross-stack collaboration. As a result, Edge Impulse has put AI in the hands of the individuals it impacts by reducing the expertise and computing resources required to participate. Edge Impulse has already been used in various industrial, research, and educational applications, and the lessons learned from these deployments can focus future systems research on high-impact problems.

REFERENCES

- Azure iot – internet of things platform — microsoft azure. <https://azure.microsoft.com/en-us/solutions/iot/#overview>. (Accessed on 10/27/2022).
- Introduction to embedded machine learning — coursera. <https://www.coursera.org/learn/introduction-to-embedded-machine-learning>. (Accessed on 10/25/2022).
- Latent ai – adaptive ai for the intelligent edge. <https://latent.ai/>. (Accessed on 10/27/2022).
- Managed kubernetes service – amazon eks – amazon web services. <https://aws.amazon.com/eks/>. (Accessed on 10/27/2022).
- Nanoedgeaistudio - automated machine learning (ml) tool for stm32 developers. <https://www.st.com/en/development-tools/nanoedgeaistudio.html>. (Accessed on 10/27/2022).
- Neuton.ai - no-code artificial intelligence for all. <https://neuton.ai/>. (Accessed on 10/27/2022).
- Stm32 arm cortex mcus - 32-bit microcontrollers - stmicroelectronics. <https://www.st.com/en/microcontrollers-microprocessors/stm32-32-bit-arm-cortex-mcus.html#overview>. (Accessed on 10/23/2022).
- The embedded learning library - embedded learning library (ell). <https://microsoft.github.io/ELL/>. (Accessed on 10/27/2022).
- Tiny machine learning (tinyml) professional certificate — edx. <https://www.edx.org/professional-certificate/harvardx-tiny-machine-learning>. (Accessed on 10/25/2022).
- Vertex ai — google cloud. <https://cloud.google.com/vertex-ai>. (Accessed on 10/27/2022).
- Webassembly. <https://webassembly.org/>. (Accessed on 10/25/2022).
- X-cube-ai - ai expansion pack for stm32cubemx - stmicroelectronics. <https://www.st.com/en/embedded-software/x-cube-ai.html>. (Accessed on 10/27/2022).
- Zephyr project - zephyr project. <https://zephyrproject.org/>. (Accessed on 10/23/2022).
- Amazon elastic container service (amazon ecs). <https://aws.amazon.com/ecs/>. (Accessed on 10/27/2022).
- edgeimpulse/processing-blocks: Signal processing blocks. <https://github.com/edgeimpulse/processing-blocks>. (Accessed on 10/24/2022).
- microtvm: Tvm on bare-metal — tvm 0.11.dev0 documentation. <https://tvm.apache.org/docs/topic/microtvm/index.html>. (Accessed on 10/27/2022).
- utensor/utensor: Tinyml ai inference library. <https://github.com/uTensor/uTensor>. (Accessed on 10/27/2022).
- Edge impulse api, 2020. URL <https://docs.edgeimpulse.com/reference/edge-impulse-api/edge-impulse-api>.
- The edge ai ecosystem. Technical Report AN-5334, ABI Research, May 2021. URL <https://www.abiresearch.com/market-research/product/7778739-the-edge-ai-ecosystem/>.
- Brainchip awarded us patent for accessing learned functions in an intelligent target device, February 2022. URL <https://www.businesswire.com/news/home/20220201006227/en/BrainChip-Awarded-US-Patent-for-Accessing-Learned-Functions-in-an-Intelligent-Target-Device>.
- Edge impulse github repositories, 2022a. URL <https://github.com/edgeimpulse/>.
- Edge impulse for linux, 2022b. URL <https://docs.edgeimpulse.com/docs/edge-impulse-for-linux/edge-impulse-for-linux>.
- Edge impulse projects, 2022c. URL <https://www.edgeimpulse.com/projects/overview>.
- Tensorflow operation fusion, June 2022. URL https://www.tensorflow.org/lite/models/convert/operation_fusion.
- Imagimob ai, 2022. URL <https://www.imagimob.com/>.
- Tinyml4d academic network, 2022. URL <https://tinyml.seas.harvard.edu/4D/AcademicNetwork>.
- Altayeb, M., Zennaro, M., and Rovai, M. Classifying mosquito wingbeat sound using tinyml. In *Proceedings of the 2022 ACM Conference on Information Technology for Social Good*, GoodIT '22, pp.

- 132–137, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392846. doi: 10.1145/3524458.3547258. URL <https://doi.org/10.1145/3524458.3547258>.
- ARM. Micro npu ethos-u65, 2022. URL <https://www.arm.com/products/silicon-ip-cpu/ethos/ethos-u65>.
- Banbury, C., Reddi, V. J., Torelli, P., Holleman, J., Jeffries, N., Kiraly, C., Montino, P., Kanter, D., Ahmed, S., Pau, D., et al. Mlperf tiny benchmark. *arXiv preprint arXiv:2106.07597*, 2021a.
- Banbury, C., Zhou, C., Fedorov, I., Matas, R., Thakker, U., Gope, D., Janapa Reddi, V., Mattina, M., and Whatmough, P. Micronets: Neural network architectures for deploying tinyml applications on commodity microcontrollers. *Proceedings of Machine Learning and Systems*, 3:517–532, 2021b.
- Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. Algorithms for hyper-parameter optimization. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, pp. 2546–2554, Red Hook, NY, USA, 2011. Curran Associates Inc. ISBN 9781618395993.
- Chavarriaga, R., Sagha, H., Calatroni, A., Digumarti, S. T., Tröster, G., del R. Millán, J., and Roggen, D. The opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters*, 34(15):2033–2042, 2013. ISSN 0167-8655. doi: <https://doi.org/10.1016/j.patrec.2012.12.014>. URL <https://www.sciencedirect.com/science/article/pii/S0167865512004205>. Smart Approaches for Human Action Recognition.
- Chollet, F. keras, 2015. URL <https://github.com/fchollet/keras>.
- Chowdhery, A., Warden, P., Shlens, J., Howard, A., and Rhodes, R. Visual wake words dataset. *CoRR*, abs/1906.05721, 2019. URL <http://arxiv.org/abs/1906.05721>.
- David, R., Duke, J., Jain, A., Janapa Reddi, V., Jeffries, N., Li, J., Kreeger, N., Nappier, I., Natraj, M., Wang, T., et al. Tensorflow lite micro: Embedded machine learning for tinyml systems. *Proceedings of Machine Learning and Systems*, 3:800–811, 2021.
- de Zambotti, M., Rosas, L., Colrain, I. M., and Baker, F. C. The sleep of the ring: Comparison of the Ōura sleep tracker against polysomnography. *Behavioral Sleep Medicine*, 17(2):124–136, 2019. doi: 10.1080/15402002.2017.1300587. URL <https://doi.org/10.1080/15402002.2017.1300587>. PMID: 28323455.
- Eggensperger, K., Feurer, M., Hutter, F., Bergstra, J., Snoek, J., Hoos, H., Leyton-Brown, K., et al. Towards an empirical foundation for assessing bayesian optimization of hyperparameters. In *NIPS workshop on Bayesian Optimization in Theory and Practice*, volume 10, 2013.
- Eggimann, M., Mach, S., Magno, M., and Benini, L. A risc-v based open hardware platform for always-on wearable smart sensing. In *2019 IEEE 8th International Workshop on Advances in Sensors and Interfaces (IWASI)*, pp. 169–174. IEEE, 2019.
- Gruenstein, A., Alvarez, R., Thornton, C., and Ghodrati, M. A cascade architecture for keyword spotting on mobile devices. *arXiv preprint arXiv:1712.03603*, 2017.
- Hofenko, M. renode, 2017. URL <https://github.com/renode/renode>.
- Hutson, M. Artificial intelligence faces reproducibility crisis, 2018.
- Hymel, S. Inferencing sdk, 2022. URL <https://docs.edgeimpulse.com/reference/c+-inference-sdk-library/inferencing-sdk>.
- Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., and Kalenichenko, D. Quantization and training of neural networks for efficient integer-arithmetic-only inference, 2017. URL <https://arxiv.org/abs/1712.05877>.
- Janapa Reddi, V., Plancher, B., Kennedy, S., Moroney, L., Warden, P., Suzuki, L., Agarwal, A., Banbury, C., Banzi, M., Bennett, M., Brown, B., Chitlangia, S., Ghosal, R., Grafman, S., Jaeger, R., Krishnan, S., Lam, M., Leiker, D., Mann, C., Mazumder, M., Pajak, D., Ramaprasad, D., Smith, J. E., Stewart, M., and Tingley, D. Widening access to applied machine learning with tinyml. *Harvard Data Science Review*, 2022. doi: 10.1162/99608f92.762d171a. URL <https://hdsr.mitpress.mit.edu/pub/0gbwdele>. <https://hdsr.mitpress.mit.edu/pub/0gbwdele>.
- Jenny Plunkett, M. B. Introducing the eon tuner: Edge impulse’s new automl tool for embedded machine learning, July 2021. URL <https://www.edgeimpulse.com/blog/introducing-the-eon-tuner-edge-impulses-new-automl-tool-for-embedded-machine-learning>.
- Jongboom, J. Introducing eon: Neural networks in up to 55% less ram and 35% less rom, 2022. URL <https://www.edgeimpulse.com/blog/introducing-eon>.

- Klein, S. *IoT Solutions in Microsoft's Azure IoT Suite*. Springer, 2017.
- Koizumi, Y., Saito, S., Uematsu, H., Harada, N., and Imoto, K. Toyadmos: A dataset of miniature-machine operating sounds for anomalous sound detection. pp. 313–317, 10 2019. doi: 10.1109/WASPAA.2019.8937164.
- Krizhevsky, A., Nair, V., and Hinton, G. Cifar-10 (canadian institute for advanced research). 2009. URL <http://www.cs.toronto.edu/~kriz/cifar.html>.
- Lai, L., Suda, N., and Chandra, V. CMSIS-NN: efficient neural network kernels for arm cortex-m cpus. *CoRR*, abs/1801.06601, 2018. URL <http://arxiv.org/abs/1801.06601>.
- Leroux, S., Simoens, P., Lootus, M., Thakore, K., and Sharma, A. Tinymlops: Operational challenges for widespread edge ai adoption. In *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pp. 1003–1010. IEEE, 2022.
- Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., and Talwalkar, A. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1):6765–6816, 2017.
- Lin, J., Chen, W.-M., Lin, Y., Cohn, J., Gan, C., and Han, S. Mcunet: Tiny deep learning on iot devices, 2020a. URL <https://arxiv.org/abs/2007.10319>.
- Lin, J., Chen, W.-M., Lin, Y., Gan, C., Han, S., et al. Mcunet: Tiny deep learning on iot devices. *Advances in Neural Information Processing Systems*, 33:11711–11722, 2020b.
- Lin, T., Maire, M., Belongie, S. J., Bourdev, L. D., Girshick, R. B., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. URL <http://arxiv.org/abs/1405.0312>.
- Mazumder, M., Banbury, C., Yao, X., Karlaš, B., Rojas, W. G., Diamos, S., Diamos, G., He, L., Kiela, D., Jurado, D., et al. Dataperf: Benchmarks for data-centric ai development. *arXiv preprint arXiv:2207.10062*, 2022.
- McInnes, L., Healy, J., and Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Moreau, L. Create active learning pipelines with data sources and data explorer features, 2022. URL <https://www.edgeimpulse.com/blog/create-active-learning-pipelines-with-data-sources-and-data-explorer-features>.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10): 1345–1359, 2009.
- Prakash, S., Callahan, T., Bushagour, J., Banbury, C., Green, A. V., Warden, P., Ansell, T., and Reddi, V. J. Cfu playground: Full-stack open-source framework for tiny machine learning (tinyml) acceleration on fpgas. *arXiv preprint arXiv:2201.01863*, 2022.
- Prakash, S., Stewart, M., Banbury, C., Mazumder, M., Warden, P., Plancher, B., and Reddi, V. J. Is tinyml sustainable? assessing the environmental impacts of machine learning on microcontrollers. *arXiv preprint arXiv:2301.11899*, 2023.
- Qiao, N., Mostafa, H., Corradi, F., Osswald, M., Stefanini, F., Sumislawska, D., and Indiveri, G. A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128k synapses. *Frontiers in neuroscience*, 9:141, 2015.
- Rauschmayr, N., Kumar, V., Huilgol, R., Olgiati, A., Bhattacharjee, S., Harish, N., Kannan, V., Lele, A., Acharya, A., Nielsen, J., et al. Amazon sagemaker debugger: a system for real-time insights into machine learning model training. *Proceedings of Machine Learning and Systems*, 3:770–782, 2021.
- Richins, D., Doshi, D., Blackmore, M., Nair, A. T., Pathapati, N., Patel, A., Daguman, B., Dobrijalowski, D., Illickal, R., Long, K., et al. Ai tax: The hidden cost of ai data center applications. *ACM Transactions on Computer Systems (TOCS)*, 37(1-4):1–32, 2021.
- Shamim, M. Z. M. Hardware deployable edge-ai solution for pre-screening of oral tongue lesions using tinyml on embedded devices. *IEEE Embedded Systems Letters*, pp. 1–1, 2022a. doi: 10.1109/LES.2022.3160281.
- Shamim, M. Z. M. Tinyml model for classifying hazardous volatile organic compounds using low-power embedded edge sensors: Perfecting factory 5.0 using edge ai. *IEEE Sensors Letters*, 6(9):1–4, 2022b. doi: 10.1109/LESENS.2022.3201398.
- Siekkinen, M., Hienkari, M., Nurminen, J. K., and Nieminen, J. How low energy is bluetooth low energy? comparative measurements with zigbee/802.15. 4. In *2012 IEEE wireless communications and networking conference workshops (WCNCW)*, pp. 232–237. IEEE, 2012.
- Situnayake, D. Announcing performance calibration, October 2022. URL <https://www.edgeimpulse.com/blog/announcing-performance-calibration>.

Situnayake, D. Unveiling byom and the edge impulse python sdk, 2023. URL <https://www.edgeimpulse.com/blog/unveiling-the-new-edge-impulse-python-sdk>.

Sørensen, P. M., Epp, B., and May, T. A depthwise separable convolutional neural network for keyword spotting on an embedded system. *EURASIP Journal on Audio, Speech, and Music Processing*, 2020(1):1–14, 2020.

Susto, G. A., Schirru, A., Pampuri, S., McLoone, S., and Beghi, A. Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3):812–820, 2015. doi: 10.1109/TII.2014.2349359.

Van der Maaten, L. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Warden, P. Speech commands: A dataset for limited-vocabulary speech recognition. *CoRR*, abs/1804.03209, 2018. URL <http://arxiv.org/abs/1804.03209>.