

TinyML4D: Scaling Embedded Machine Learning Education in the Developing World

Brian Plancher^{1*}, Sebastian Buttrich², Jeremey Ellis³, Neena Goveas⁴, Laila Kazmierski^{5,6,7}, Jesus Lopez Sotelo⁸, Milan Lukic⁹, Diego Mendez¹⁰, Rosdiadee Nordin¹¹, Andres Oliva Trevisan^{5,6}, Massimo Pavan¹², Manuel Roveri¹², Marcus Rüb¹³, Jackline Tum¹⁴, Marian Verhelst¹⁵, Salah Abdeljabar¹⁶, Segun Adebayo¹⁷, Thomas Amberg¹⁸, Halleluyah Aworinde¹⁷, José Bagur¹⁹, Gregg Barrett²⁰, Nabil Benamar²¹, Bharat Chaudhari²², Ronald Criollo²³, David Cuartielles²⁴, Jose Alberto Ferreira Filho²⁵, Solomon Gizaw²⁶, Evgeni Gousev²⁷, Alessandro Grande²⁸, Shawn Hymel²⁸, Peter Ing²⁹, Prashant Manandhar³⁰, Pietro Manzoni³¹, Boris Murmann³², Eric Pan³³, Rytis Paskauskas³⁴, Ermanno Pietrosemoli³⁴, Tales Pimenta²⁵, Marcelo Rovai²⁵, Marco Zennaro³⁴, Vijay Janapa Reddi³⁵

¹Barnard College, Columbia University

²IT University of Copenhagen

³School District 75 Mission

⁴BITS Pilani K K Birla - Goa Campus

⁵Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET)

⁶Centro Atómico Bariloche (CAB)

⁷Comisión Nacional de Energía Atómica (CNEA)

⁸Universidad Autónoma de Occidente

⁹University of Novi Sad

¹⁰Pontificia Universidad Javeriana

¹¹Sunway University

¹²Politecnico di Milano

¹³Technische Universität München

¹⁴Dedan Kimathi University of Technology

¹⁵KU Leuven

¹⁶King Abdullah University of Science and Technology (KAUST)

¹⁷Bowen University

¹⁸University of Applied Sciences Northwestern Switzerland FHNW

¹⁹Universidad del Valle de Guatemala

²⁰Cirrus

²¹Moulay Ismail University

²²MIT World Peace University

²³Escuela Superior Politécnica del Litoral

²⁴Malmö University

²⁵Universidade Federal de Itajubá (UNIFEI)

²⁶Addis Ababa University

²⁷Qualcomm

²⁸Edge Impulse

²⁹TFG (The Foschini Group)

³⁰Center for Information and Communication Technology for Development

³¹Universitat Politècnica de València

³²University of Hawaii

³³Seed Studio

³⁴The Abdus Salam International Centre for Theoretical Physics (ICTP)

³⁵Harvard University

*Corresponding author contact: bplancher@barnard.edu
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Abstract

Embedded machine learning (ML) on low-power devices, also known as “TinyML,” enables intelligent applications on accessible hardware and fosters collaboration across disciplines to solve real-world problems. Its interdisciplinary and practical nature makes embedded ML education appealing, but barriers remain that limit its accessibility, especially in developing countries. Challenges include limited open-source software, courseware, models, and datasets that can be used with globally accessible heterogeneous hardware. Our vision is that with concerted effort and partnerships between industry and academia, we can overcome such challenges and enable embedded ML education to empower developers and researchers worldwide to build locally relevant AI solutions on low-cost hardware, increasing diversity and sustainability in the field. Towards this aim, we document efforts made by the TinyML4D community to scale embedded ML education globally through open-source curricula and introductory workshops co-created by international educators. We conclude with calls to action to further develop modular and inclusive resources and transform embedded ML into a truly global gateway to embedded AI skills development.

Introduction

Embedded machine learning (ML) has the potential to enable intelligent on-device interactions and sustainable computing. This involves the ability to run ML locally on low-powered devices, often referred to as “TinyML” (Warden and Situnayake 2019). Potential applications leveraging embedded ML span areas such as healthcare care, climate monitoring, agriculture, and transportation (Prakash et al. 2023; Tsoukas et al. 2021; Janapa Reddi et al. 2023).

However, adopting embedded ML globally, especially in developing countries, faces critical challenges. First, embedded ML requires the integration of multidisciplinary software and hardware concepts, a significant barrier for students and educators from varying backgrounds. The lack of unified software toolchains and hardware platforms also creates a bottleneck for developing canonical examples to teach and test ideas. Finally, interdisciplinary topics like embedded ML face obstacles to widespread faculty adoption despite the availability of curricula. This highlights the importance of well-coordinated global initiatives that incorporate accessibility as a fundamental principle to support the growth of this area, particularly in developing countries.

This paper presents strides that the TinyML for Developing Countries (TinyML4D) community has made toward overcoming these obstacles. The TinyML4D community is an international group of educators, makers, and experts in embedded systems, machine learning, and embedded applications, who have come together to promote the use of TinyML technologies in resource-constrained regions across the globe. Academics contribute their knowledge of cutting-edge research and applications. Industry leaders share their technical skills. Those with on-the-ground experience in developing countries provide critical insights into local challenges and opportunities. Together, through grassroots efforts, open-source projects, workshops, and educational initiatives, TinyML4D aims to empower local communities to leverage AI and ML effectively, ethically, and sustainably.

Over the past two years, the community has developed and released open-source at tinyMLEdu.org, embedded ML courses and examples of canonical embedded ML tasks on various hardware platforms. To support this courseware, we recently developed a foundational textbook, *Machine Learning Systems with TinyML* (Janapa Reddi et al. 2024). In addition, we have developed a university network spread throughout the global south (Plancher and Janapa Reddi 2022) that has taught courses and workshops in multiple languages (e.g., English, Spanish, Portuguese) in more than a dozen countries spanning five continents. These efforts have reached over a thousand students and led to more than 15 peer-reviewed academic publications.

TinyML4D’s goal is to catalyze truly global access to embedded ML education, requiring strategic efforts spanning academia, industry, nonprofits, and governments. To do so, we are developing a modular open-source embedded ML curriculum that includes core software and hardware skills before branching into application domains like health and agriculture. Central to these efforts will be cross-sector collaborations to ensure that our curriculum teaches real-world skills, that hardware resources are accessible globally, and that a repository of portable and maintained models and datasets is developed. We also advocate for increased reputable scholarly venues focused on the topic, such as the TinyML Research Symposium (tinyML Foundation 2024), to strengthen incentives for faculty adoption. Finally, as we aim to broaden access, we must develop outreach efforts and intentionally develop an inclusive community.

In the following sections, we explain these obstacles, outline our initial efforts to address them, and present our recommendations for further action.

Embedded Machine Learning Background

Embedded ML defines the process of deploying ML models on low-power devices at the network edge. Also called “TinyML” (Warden and Situnayake 2019), this approach combines optimizations from machine learning and embedded systems. By running models locally, embedded ML unlocks order-of-magnitude improvements in cost, power, connectivity, reliability, and privacy. As such, this technology is poised to advance the United Nations Sustainable Development Goals (SDGs) in areas such as hunger, clean water, climate action, and more (Prakash et al. 2023; Zennaro, Plancher, and Janapa Reddi 2022, 2023). For example, embedded ML unlocks applications in remote scenarios lacking reliable infrastructure, from wildlife monitoring conservation networks to predictive agriculture for local soil monitoring (Dutta and Bharali 2021b,a; Abadade et al. 2023). And, by processing data internally and never transmitting raw data, privacy and security are inherently upheld, aligning with ethical AI principles (Warden et al. 2023).

Embedded Machine Learning Education

Despite rapid advances, cutting-edge machine learning education remains out of reach for millions of aspiring learners worldwide. For example, healthcare workers in low and middle-income countries lack access to impactful AI

skills training opportunities (Mollura et al. 2020), with similar trends in other fields (Pedro et al. 2019; James 2021). Prohibitive barriers around connectivity, computational resources, language, and cost prevent countless students and institutions from transforming their disciplines. This underscores the urgent need for alternative channels that radically democratize access to applied ML education, targeting learners of all backgrounds. Embedded ML presents immense opportunities to overcome these challenges and increase the accessibility and impact of ML education globally. However, numerous challenges prevent its scaled adoption. In this section we present both in detail.

Global Opportunities

By enabling operation with minimal resources, promoting interdisciplinary collaboration, and facilitating hands-on applied learning, embedded ML can dramatically widen access to applied ML at a global scale.

Low Resource Requirements: Embedded ML courses can leverage low-cost microcontroller hardware in low-connectivity and low-power environments. This allows impactful ML education to finally be within reach for communities lacking reliable infrastructure, from rural villages to underserved urban neighborhoods.

Interdisciplinary Focus: For aspiring ML engineers and scientists from all fields, embedded ML offers a unique interdisciplinary learning experience, spanning both software and hardware concepts, as well as integration with adjacent scientific fields, such as environmental science, conservation, and healthcare. Learners can develop coveted cross-stack skill sets by applying algorithms, designing inference pipelines, optimizing models, and integrating hardware components such as sensors.

Applied Learning: Embedded ML education centers on hands-on experiential learning for community impact. Students exercise the entire process of identifying local problems, designing ML solutions, and deploying and measuring real-world efficacy. For example, a recent project leveraged TinyML to develop a low-cost mosquito detection system to reduce the spread of diseases in low-resource communities (Kimutai and Förster 2023).

Global Challenges

Realizing embedded ML’s full potential requires an honest confrontation with existing barriers and core challenge areas, which we detail below. We underscore these challenges not to discourage progress but to galvanize strategic solutions to unlock embedded ML at scale.

Software and Hardware Fragmentation: The immense diversity of embedded software frameworks and hardware devices poses the most immediate barrier to unlocking embedded ML globally. Each microcontroller and sensor combination requires tailored models and toolchains optimized for specific instruction sets, memory profiles, and data pipelines. The large universe of devices from the many popular manufacturers and technologies (e.g., ARM, RISC-V, Arduino, Raspberry Pi¹) combined with the rapid pace

of updates across software tools existing courses rely upon (e.g., Edge Impulse Studio, TensorFlow Lite Micro²) multiplies course creation and maintenance efforts. This heterogeneity, combined with uneven global availability of tools and devices, has stymied attempts at a global curriculum.

Affordability Barriers: Although embedded ML relies on lower-cost hardware, acquisition expenses still often prove prohibitively high for widespread adoption in economically disadvantaged regions. Microcontrollers that are affordably priced, in theory, face substantial markups from shipping costs, taxes, and import duties in practice, which can often be more expensive than the device itself.

Localization Roadblocks: The immense diversity of cultures and languages across emerging economies poses additional roadblocks for embedded ML globally. Effective localized education requires adapting materials beyond a word-for-word translation and customizing them to regional contexts. For instance, an agriculture use case in Mexico should reference endemic crops such as agave and climate conditions in the arid northern regions. This level of localization is resource-intensive yet quintessential for relevance.

Educator Readiness: Embedded ML’s multidisciplinary potential necessitates cultivating adept educators fluent across electronics, coding, and ML concepts. However, traditional domain silos slow this transition. Electrical engineering curricula often lack software depth, and computer science programs do not often integrate circuit fundamentals. This constrains the supply of qualified educators.

Research Incentives: Without sufficient avenues for publishing scientific advances or curriculum innovations, critical insights are lost, and professors also lack external validation from peers essential for institutional endorsement and career advancement.

The TinyML4D Network

In 2020, assisted by the International Centre for Theoretical Physics (ICTP) and Harvard University, the Tiny Machine Learning Open Education Initiative (TinyMLedu) was established as an international consortium comprising scholars and professionals. The primary objective of this initiative is to enhance the availability of educational resources worldwide in the area of embedded ML, specifically focusing on the lowest-power applications. As such, to overcome the challenge of jump-starting the adoption of TinyML across the Global South, TinyMLedu’s first significant initiative was the launch of the TinyML4D Academic Network in 2021 to catalyze localized course offerings through shared best practices and resources. Since its launch in 2021 with 20 founding institutional partners from across the Global South, the TinyML4D network has doubled to include over 40 official members by 2023 (see Figure 1 left). Additionally, several institutions are waiting to join the network, and a number of supporting institutions from the Global North are active supporters. In the remainder of this section, we describe TinyML4D’s early success in addressing the global challenges outlined above.

¹arm.com, riscv.org, arduino.cc, raspberrypi.com

²edgeimpulse.com, tensorflow.org/lite/microcontrollers



Figure 1: Left: The 2023 TinyML4D Academic Network. Right: The Participants of the Workshop on Widening Access to TinyML Network by Establishing Best Practices in Education held at ICTP in Trieste, Italy, in July 2023.

Software and Hardware Fragmentation & Affordability Barriers: To help address these structural issues, we provided all 40 network members with 10 standardized Arduino hardware kits. These kits pair with the existing TinyML edX course materials, which we released open-source (TinyML EdX Team 2021). Since then, additional industrial collaborations with Arduino, Seeed Studio and Edge Impulse³ have enabled academic pricing schemes and the direct distribution of additional hardware to universities around the globe.

Educator Readiness and Localization Roadblocks: To increase global educator access to high quality training, TinyML4D has launched several hands-on virtual, hybrid, and in-person workshops and courses in multiple languages (e.g., English, Spanish, Portuguese) across more than a dozen countries spanning five continents.

Our flagship workshop, *SciTinyML: Scientific Use of Machine Learning on Low-Power Devices*, a 5-day hands-on, virtual workshop for university students and professors exploring real-world applications of TinyML and their impact on the developing world, was run globally in 2021 with 216 participants from 48 countries, regionally in 2022 for Africa (187 from 29), Asia (100 from 8), and Latin America (200 from 17), globally in 2023 (418 from 76), and will be run globally in 2024. While most workshop participants came from network member institutions to whom we provided hardware resources, we also opened the workshops to the broader community. We developed off-ramps for hands-on exercises that enabled those without the requisite hardware to still complete most activities. Excitingly, we have trained more than 1,000 participants through our workshops, from students seeking broader skills to professionals upskilling to even high schoolers tackling technical challenges.

One of our core learnings is that focused regional workshops provide an excellent platform for regional cooperation. As such, while we aim to continue to teach online global workshops to ensure maximal accessibility of these topics, we are also working to develop more focused regional, in-person workshops. This has led to several one-day and one-week, in-person and hybrid workshops that have been run or are scheduled to run in the next few years spanning the globe, including in Brazil, Ecuador, Ghana, Macao, Morocco, Nepal, Ethiopia, Malaysia, Panama, and Saudi

Arabia. Notably, many current educators in the field were first participants in our workshops and courses, and many of these regional efforts were designed specifically to launch embedded ML activities through a train-the-trainer design.

Research Incentives: We have worked to overcome these challenges in two ways. First, we have helped develop reputable venues, including supporting the aforementioned TinyML Research Symposium (tinyML Foundation 2024), developing a special issue in IEEE Micro (Reddi and Murrmann 2023), and an upcoming special session at the 2024 IEEE World Congress on Computational Intelligence. Second, to catalyze early research efforts, we launched a “Show and Tell” online series, which provides a more informal setting for young researchers and practitioners to present in-progress projects and gain feedback from the community. We have found this format to be more inclusive than traditional academic conferences and received so many contributions (over 50) for our first four planned events, that we have continued the series quarterly. A sample of the projects presented includes: Hand Gesture Recognition for Mute People from Algeria, Monitoring Bees from Kenya, Irrigation Prediction from Colombia, Non-Invasive Anaemia Detection from Peru, Shelf Life Estimation of Date Palms from Saudi Arabia, the development of an Intelligent Personal Trainer System from Brazil and Cashew Nut Disease Detection from India. Building on this, researchers in our network have over 15 peer-reviewed embedded ML publications including: (Ooko et al. 2021; Avellaneda, Mendez, and Fortino 2022; Altayeb, Zennaro, and Rovai 2022; Bamoumen et al. 2022; Mihigo et al. 2022; Avellaneda, Mendez, and Fortino 2023; Bordin Yamashita and Leite 2023; Kazimierski et al. 2023; Meza-Rodriguez, De La Cruz, and Cáceres-DelAguila 2023; Mallick et al. 2023; Silva et al. 2022; Neuman et al. 2022; Atanane et al. 2023).

Foundations of a Modular, Global, Open-Source Curriculum

TinyML4D efforts have been greatly facilitated by the development of open-source educational resources, allowing educators worldwide to easily create new courses without starting from scratch. While this approach has shown early success, to increase our global impact further, we organized an in-person workshop, *Edge ML University Program 2023*:

³arduino.cc, seedstudio.com, edgeimpulse.com

Optional Modules					Core Modules			Canonical Hands-On Examples
Software Focus	Embedded Software Engineering Deep Dive	ML Compilers and Optimizers	Neural Network Architecture and Design		Machine Learning and Deep Learning Fundamentals (E.g., Models, Training, Overfitting, Regression vs. Classification, Neural Networks)			From data collection to model training to deployed inference
Hardware Focus	Electronics and IoT Deep Dive	Sensor Paradigms and Design	Device Design and Deployment		Data Centric AI (E.g., Data Collection, Pre- and Post-Processing)	Embedded Systems (E.g., Microcontrollers, Embedded Programming, Basic Electronics & IoT)		Audio Keyword Spotting
Domain-Specific Focus	Conservation	Predictive Maintenance	Smart Cities		Responsible AI (E.g., Bias, Privacy, Security)			Image Classification
Course Lengths	<ul style="list-style-type: none"> 3-5 day short courses include one hands-on example from each core module and a high level overview of the theory. Time permitting they include optional modules. Micro-credential courses dive deeper into the theory of each core module and select optional modules. Full semester-long courses explore a full track of optional modules and the core modules in detail with multiple hands-on examples and theoretical derivations and explorations. 							

Figure 2: A high-level overview of our curriculum framework revealing our common core and examples of extension modules for the various tracks designed to support courses of varying lengths, contexts and degree levels.

Workshop on Widening Access to TinyML Network by Establishing Best Practices in Education, at ICTP in Trieste, Italy, in July 2023. The workshop was attended by 42 participants from 25 different countries (see Figure 1 right). In particular, with the assistance of ICTP, we offered travel grants and accommodation to participants from the Global South. The resulting open-source modular curriculum framework for embedded ML caters to various skill levels, backgrounds, and areas of interest, with various tracks centered around a common core (see Figure 2). In the remainder of this section, we will elaborate on the framework, its different tracks, and our initial endeavors in creating its open-source materials.

A Common Core with Canonical Hands-on Examples: Regardless of the course’s specific focus, we identified several fundamental modules that should be included in every embedded ML course. This guarantees that all students acquire the essential skills needed to excel in a career in embedded ML and possess the necessary language and foundational knowledge to pursue advanced courses in the various fields that support embedded ML. These encompass:

- **Embedded Systems** (E.g., Microcontrollers, Embedded Programming, Basic Electronics & IoT)
- **Machine Learning and Deep Learning Fundamentals** (E.g., Models, Training, Overfitting, Regression vs. Classification, Neural Networks)
- **Data-Centric AI** (E.g., Data Collection, Pre- and Post-Processing)
- **Responsible AI** (E.g., Bias, Privacy, Security)

To help students navigate the entire applied machine learning journey, we developed three canonical examples of TinyML: audio keyword spotting (e.g., identifying phrases like “OK Google”), multi-class image classification using image datasets, and motion anomaly detection using inertial measurement units (IMUs). Importantly, these three examples leverage different sensors, some of which may be external to the embedded system and require some wiring, require pre- or post-processing of data, and will be easier or

harder to train or generalize the resulting model. As such, we feel that learners can extend the workflow to their future endeavors through these hands-on examples.

Multiple Tracks of Courses: As embedded ML is an interdisciplinary field and can be taught in a variety of different departments, contexts, and environments, our framework includes three main tracks of focus:

- **Software Focus:** Equipping students with a foundation in algorithms, machine learning theory, and software development. Ideal for students in fields such as computer science, data science, and information technology.
- **Hardware Focus:** Deepening the student’s understanding of embedded systems, microcontrollers, and sensor integration. Ideal for students in fields such as computer, electrical, and mechanical engineering.
- **Domain-Specific Focus:** Tailoring embedded ML applications to specific sectors, such as environmental monitoring or assistive technologies. Applicable to a wide range of STEM fields that want to leverage embedded ML as a tool for analysis and application.

Through these pathways, embedded ML can be integrated into various departments and fields, lowering barriers-to-entry for both students and educators. Finally, as the field progresses, this method should streamline the process of maintaining the curriculum by allowing for the easy updating and substitution of individual modules within courses.

Courses of Different Lengths: Our curriculum accommodates courses of varying lengths, categorized into three groups to meet the diverse needs and backgrounds of learners and educators. Although the length of each category varies significantly, each addresses the core topics and then explores specific additional modules, time permitting.

First, there is a comprehensive semester-long course designed for integration into computer science or engineering degree programs, with the goal of establishing a solid understanding of both machine learning and embedded systems. We have already made significant strides by aggregat-

ing the courseware of semester-long courses from Harvard, MIT, UPenn, and UNIFEI in Brazil.

Second, the micro-credential course, modeled after offerings on popular online MOOC platforms, caters to a broad audience, including self-directed adult learners, undergraduates, and graduates. These courses offer flexibility and accessibility, while still providing technical insights and practical experiences. We have already provided course materials for MOOCs from Harvard on edX and Edge Impulse on Coursera⁴. We also developed an open-source textbook to pair with this courseware, *Machine Learning Systems with TinyML* (Janapa Reddi et al. 2024).

Lastly, the short course, lasting 3 to 5 days, is designed to introduce individuals with or without computer science or engineering backgrounds to the embedded ML environment. This format is particularly useful for early exposure and basic comprehension. Importantly, all courses include a hands-on component or live demonstration. Several of these can be found on the *tinyMLedu.org* website.

Finally, to guarantee accessibility and user-friendliness, we are utilizing a decentralized hosting strategy with centralized coordination. In particular, our *tinyMLedu.org* website serves as the centralized front-end linking to materials developed globally and hosted on dependable and well-known cloud platforms, which ultimately store the content.

Calls to Action

Building on our early results and to propel the field of embedded ML forward globally, it is essential to take action on multiple fronts. The following calls to action address key areas that require the attention of researchers, educators, industry professionals, policymakers, and the general public.

Assessing Our Educational Programs: Continuous assessment and improvement of our educational programs is paramount for long term success. As such, to develop an evidence base for what works globally, we must build unified pre- and post-course surveys into our framework. Similarly, we need to establish shared standards for content quality, relevance, and pedagogical effectiveness, while allowing for multiple interpretations and implementations. This should include a transparent peer review system, mechanisms for continuous updates, and a concept guide in the style of the NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing (Prasad et al. 2011).

Maintaining Open-Source Software and Courseware: The challenges of heterogeneity in hardware and software tools within the embedded ML ecosystem mean that even once our modular curriculum is fully developed, it will take significant effort to maintain. We call on industry partners to play a pivotal role in maintaining and supporting open-source software workflows. In particular, for our set of canonical hands-on examples, we ask hardware vendors to ensure that these few examples can work for a five-year long-term support commitment.⁵ To encourage such an effort, we advocate for the development of a certification that notes if a product is “Embedded ML Ready” and comes with

⁴edx.org/learn/tinyml, coursera.org/instructor/shawnhymel

⁵This term follows the Ubuntu long-term support model.

the promise of long-term support. Similarly, we ask that all academics who help develop and share their course materials provide some long-term support. To encourage those efforts, we also want to develop a certification process for “Gold Star Courses,” with a guaranteed maintenance term.

Embedded ML Model and Data Zoo: Building on the foundation of maintained open-source software and courseware, and the continued growth and prominence of data-centric AI (Mazumder et al. 2022; Zha et al. 2023), there is an imperative need to create an “embedded ML model and data zoo.” This will encompass benchmark datasets and models similar to ImageNet for computer vision (Deng et al. 2009). Building on early successes like the Visual Wake Words Dataset (Chowdhery et al. 2019), such a resource will catalyze the rapid prototyping and deployment of embedded ML systems, enabling benchmarks for comparison across the field. Collaborating with initiatives like MLPerf (Mattson et al. 2020), who already produced a TinyML benchmark (Banbury et al. 2021), could accelerate such efforts. To further enhance accessibility and sharing, data sets can be made available through platforms such as Mendeley Data, Huggingface, Zenode (CERN), OSF, OpenML, or GitHub⁶. It is also important to tag datasets and models with key metadata. For example, knowing the mounting location of a sensor and having both raw and cleaned data would enable researchers to develop algorithms that anticipate future challenges with real-world deployments.

Improving Accessibility of Hardware: Efforts should be made to make hardware accessible at a low cost *after* considering shipping, taxes, and import duties. Collaboration between regional leaders and organizations can help overcome regional purchasing challenges. Exploring industry partnerships can help secure hardware for courses while enabling manufacturers to promote their hardware and recruit future talent. Governments can also provide support and enable their populations to gain skills in emerging technologies. Simultaneously, simulators for code verification and execution can help educators develop experiments and courses even *without* hardware resources. Early efforts towards this goal, such as Renode and QEMU⁷, can already support embedded ML applications (Iodice 2022). Hardware testbeds could also be developed based on similar successful efforts for wireless sensor networks (Werner-Allen, Swieskowski, and Welsh 2005; Doddavenkatappa, Chan, and Ananda 2012)).

Growing a Research Community: To sustain and expand our progress, it is imperative to foster a diverse and vibrant research community with opportunities for collaboration and recognition. In particular, we call upon academic leaders to establish a high profile conference or journal specifically tailored to embedded ML’s interdisciplinary nature. The addition of online competitions can also help. We also advocate for industry to provide direct financial support and grant opportunities. These efforts will help educators gain the resources and institutional support needed to teach embedded ML and develop open-source courseware.

⁶data.mendeley.com, huggingface.co, zenodo.org, osf.io, openml.org, github.com

⁷renode.io, qemu.org

Finally, while early research has shown the potential for embedded ML to support efforts toward responsible AI practices (Warden et al. 2023; Stewart et al. 2023), additional work must be done to expand transparency, improve explainability, and reduce the bias of real-world embedded ML systems (e.g., by building on recent efforts in large-scale machine learning (Mitchell et al. 2019; Roscher et al. 2020; Mehrabi et al. 2021; Selbst et al. 2019)).

Outreach and Diversity Efforts: While our global design already reaches a relatively diverse audience geographically and socioeconomically, reaching out to audiences beyond the confines of higher education is also crucial. For example, KU Leuven's InnovationLab⁸ developed a "teach the teacher" partnership with local primary and secondary schools to enable hands-on engineering experiences that integrate directly into existing curricula. Furthermore, initiatives such as webinars, public lectures, podcasts, interactive demonstrations, citizen science projects, and exhibitions can showcase real-world embedded ML applications, making the technology tangible and relatable. Importantly, while conducting these outreach efforts, a particular focus should also be made to foster diversity, inclusion, and belonging within the embedded ML community. As such, we must develop inviting online communities and application-driven content that addresses interdisciplinary global challenges.

Conclusion and Future Work

Embedded ML shows immense promise in enabling global AI education through accessible hands-on learning. Our initiatives establishing open-source courseware, the TinyML4D global university network, and modular curriculum lay the groundwork for an inclusive TinyML community and have early global momentum. However, realizing our full vision requires coordinated next-generation efforts between stakeholders, maintaining software infrastructure, aligning industry-academia content, and providing reputable publication venues as rewards. Through such collaborative efforts, we aim to construct an embedded ML ecosystem where learners of all backgrounds, geographies, and generations can shape our emerging AI landscape.

References

Abadade, Y.; Temouden, A.; Bamoumen, H.; Benamar, N.; Chtouki, Y.; and Hafid, A. S. 2023. A comprehensive survey on tinyml. *IEEE Access*.

Altayeb, M.; Zennaro, M.; and Rovai, M. 2022. Classifying Mosquito Wingbeat Sound Using TinyML. In *Proceedings of the 2022 ACM Conference on Information Technology for Social Good*, GoodIT '22, 132–137. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392846.

Atanane, O.; Mourhir, A.; Benamar, N.; and Zennaro, M. 2023. Smart Buildings: Water Leakage Detection Using TinyML. *Sensors*, 23(22): 9210.

⁸eng.kuleuven.be/innovationlab

Avellaneda, D.; Mendez, D.; and Fortino, G. 2023. A TinyML Deep Learning Approach for Indoor Tracking of Assets. *Sensors*, 23(3).

Avellenada, D.; Mendez, D.; and Fortino, G. 2022. BLE-based Indoor Positioning Platform Utilizing Edge Tiny Machine Learning. In *2022 IEEE Intl Conf on Pervasive Intelligence and Computing (PiCom)*, 1–8.

Bamoumen, H.; Temouden, A.; Benamar, N.; and Chtouki, Y. 2022. How TinyML Can be Leveraged to Solve Environmental Problems: A Survey. In *2022 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT)*, 338–343.

Banbury, C.; Reddi, V. J.; Torelli, P.; Holleman, J.; Jeffries, N.; Kiraly, C.; Montino, P.; Kanter, D.; Ahmed, S.; Pau, D.; et al. 2021. Mlperf tiny benchmark. *arXiv preprint*.

Bordin Yamashita, J. V. Y.; and Leite, J. P. R. 2023. Coffee disease classification at the edge using deep learning. *Smart Agricultural Technology*, 4: 100183.

Chowdhery, A.; Warden, P.; Shlens, J.; Howard, A.; and Rhodes, R. 2019. Visual wake words dataset. *arXiv preprint*.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

Doddavenkatappa, M.; Chan, M. C.; and Ananda, A. L. 2012. Indriya: A low-cost, 3D wireless sensor network testbed. In *Testbeds and Research Infrastructure. Development of Networks and Communities: 7th International ICST Conference, TridentCom 2011, Shanghai, China, April 17–19, 2011, Revised Selected Papers* 7, 302–316. Springer.

Dutta, D. L.; and Bharali, S. 2021a. TinyML Meets IoT: A Comprehensive Survey. *Internet of Things*, 16: 100461.

Dutta, L.; and Bharali, S. 2021b. Tinyml meets iot: A comprehensive survey. *Internet of Things*, 16: 100461.

Iodice, G. M. 2022. *TinyML Cookbook: Combine artificial intelligence and ultra-low-power embedded devices to make the world smarter*. Packt Publishing Ltd.

James, J. 2021. Confronting the scarcity of digital skills among the poor in developing countries. *Development Policy Review*, 39(2): 324–339.

Janapa Reddi, V.; Elijum, A.; Hymel, S.; Tischler, D.; Situnayake, D.; Ward, C.; Moreau, L.; Plunkett, J.; Kelcey, M.; Baaijens, M.; et al. 2023. Edge Impulse: An MLOps Platform for Tiny Machine Learning. *Proceedings of Machine Learning and Systems*, 5.

Janapa Reddi, V.; et al. 2024. Machine Learning Systems with TinyML. https://harvard-edge.github.io/cs249r_book/. Accessed: 2024-03-31.

Kazimierski, L. D.; Oliva Trevisan, A.; Kubisch, E.; Laneri, K.; and Catalano, N. 2023. Design and Development of a Family of Integrated Devices to Monitor Animal Movement in the Wild. *Sensors*, 23(7): 3684.

Kimutai, G.; and Förster, A. 2023. A low-cost TinyML model for Mosquito Detection in Resource-Constrained Environments. In *Proceedings of the 2023 ACM Conference on Information Technology for Social Good*, 23–30.

Mallick, S.; Gawali, S.; Onime, C.; and Goveas, N. 2023. Sleep Apnea Detection System Using Machine Learning on Resource-Constrained Devices. In *2023 IEEE International Systems Conference (SysCon)*, 1–6.

Mattson, P.; Reddi, V. J.; Cheng, C.; Coleman, C.; Diamos, G.; Kanter, D.; Micikevicius, P.; Patterson, D.; Schmuelling, G.; Tang, H.; et al. 2020. MLPerf: An industry standard benchmark suite for machine learning performance. *IEEE Micro*, 40(2): 8–16.

Mazumder, M.; Banbury, C.; Yao, X.; Karlaš, B.; Rojas, W. G.; Diamos, S.; Diamos, G.; He, L.; Parrish, A.; Kirk, H. R.; et al. 2022. Dataperf: Benchmarks for data-centric ai development. *arXiv preprint*.

Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.

Meza-Rodriguez, M.; De La Cruz, L.; and Cáceres-DelAguila, J. A. 2023. Development of an electrocardiographic signal classifier for bundle branch blocks, applying Tiny Machine Learning. In *2023 IEEE XXX International Conference on Electronics, Electrical Engineering and Computing (INTERCON)*, 1–6.

Mihigo, I. N.; Zennaro, M.; Uwitonze, A.; Rwigema, J.; and Rovai, M. 2022. On-Device IoT-Based Predictive Maintenance Analytics Model: Comparing TinyLSTM and Tiny-Model from Edge Impulse. *Sensors*, 22(14).

Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, 220–229.

Mollura, D. J.; Culp, M. P.; Pollack, E.; Battino, G.; Scheel, J. R.; Mango, V. L.; Elahi, A.; Schweitzer, A.; and Dako, F. 2020. Artificial intelligence in low-and middle-income countries: innovating global health radiology. *Radiology*, 297(3): 513–520.

Neuman, S. M.; Plancher, B.; Duisterhof, B. P.; Krishnan, S.; Banbury, C.; Mazumder, M.; Prakash, S.; Jabbour, J.; Faust, A.; de Croon, G. C.; and Janapa Reddi, V. 2022. Tiny Robot Learning: Challenges and Directions for Machine Learning in Resource-Constrained Robots. In *2022 IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. Incheon, Korea.

Ooko, S. O.; Muyonga Ogore, M.; Nsenga, J.; and Zennaro, M. 2021. TinyML in Africa: Opportunities and Challenges. In *2021 IEEE Globecom Workshops (GC Wkshps)*, 1–6.

Pedro, F.; Subosa, M.; Rivas, A.; and Valverde, P. 2019. Artificial intelligence in education: Challenges and opportunities for sustainable development.

Plancher, B.; and Janapa Reddi, V. 2022. TinyMLEdu: The Tiny Machine Learning Open Education Initiative. In *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education*, 1159. New York, NY, USA.

Prakash, S.; Stewart, M.; Banbury, C.; Mazumder, M.; Warden, P.; Plancher, B.; and Reddi, V. J. 2023. Is TinyML Sustainable? Assessing the Environmental Impacts of Machine Learning on Microcontrollers. *Communications of the ACM*, 66(11): 68–77.

Prasad, S. K.; Chtchelkanova, A.; Das, S.; Dehne, F.; Gouda, M.; Gupta, A.; Jaja, J.; Kant, K.; La Salle, A.; LeBlanc, R.; et al. 2011. NSF/IEEE-TCPP curriculum initiative on parallel and distributed computing: core topics for undergraduates. In *Proceedings of the 42nd ACM technical symposium on Computer science education*, 617–618.

Reddi, V. J.; and Murmann, B. 2023. Special Issue on TinyML. *IEEE Micro*, 43(06): 7–10.

Roscher, R.; Bohn, B.; Duarte, M. F.; and Garcke, J. 2020. Explainable machine learning for scientific insights and discoveries. *Ieee Access*, 8: 42200–42216.

Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, 59–68.

Silva, G. V.; Lima, M. D.; Filho, J. A.; and Rovai, M. J. 2022. Atrial Fibrillation and Sinus Rhythm Detection Using TinyML (Embedded Machine Learning). In *Latin American Conference on Biomedical Engineering*, 633–644. Springer.

Stewart, M.; Warden, P.; Omri, Y.; Prakash, S.; Santos, J.; Hymel, S.; Brown, B.; MacArthur, J.; Jeffries, N.; Plancher, B.; and Janapa Reddi, V. 2023. Datasheets for Machine Learning Sensors. *arXiv preprint*.

TinyML EdX Team. 2021. TinyML EdX Courseware. <https://github.com/tinyMLx/courseware/tree/master/edX>.

tinyML Foundation. 2024. Research Symposium 2024. <https://www.tinyml.org/event/research-symposium-2024/>.

Tsoukas, V.; Boumpa, E.; Giannakas, G.; and Kakarountas, A. 2021. A review of machine learning and tinyml in healthcare. In *Proceedings of the 25th Pan-Hellenic Conference on Informatics*, 69–73.

Warden, P.; and Situnayake, D. 2019. *Tinyml: Machine learning with tensorflow lite on arduino and ultra-low-power microcontrollers*. O'Reilly Media.

Warden, P.; Stewart, M.; Plancher, B.; Katti, S.; and Reddi, V. J. 2023. Machine Learning Sensors: A Design Paradigm for the Future of Intelligent Sensors. *Communications of the ACM*, 66(11): 25–28.

Werner-Allen, G.; Swieskowski, P.; and Welsh, M. 2005. Motelab: A wireless sensor network testbed. In *IPSN 2005. Fourth International Symposium on Information Processing in Sensor Networks*, 2005., 483–488. IEEE.

Zennaro, M.; Plancher, B.; and Janapa Reddi, V. 2022. TinyML: Applied AI for Development. In *The UN 7th Multi-stakeholder Forum on Science, Technology and Innovation for the Sustainable Development Goals*.

Zennaro, M.; Plancher, B.; and Janapa Reddi, V. 2023. Bridging the Digital Divide: the Promising Impact of TinyML for Developing Countries. In *The UN 8th Multi-stakeholder Forum on Science, Technology and Innovation for the Sustainable Development Goals*.

Zha, D.; Bhat, Z. P.; Lai, K.-H.; Yang, F.; Jiang, Z.; Zhong, S.; and Hu, X. 2023. Data-centric artificial intelligence: A survey. *arXiv preprint*.