Is TinyML Sustainable?

Assessing the Environmental Impacts of Machine Learning on Microcontrollers

Full-length Manuscript Appeared in Communications of the ACM (CACM), November 2023

Shvetank Prakash* Matthew Stewart* Colby Banbury* Mark Mazumder*
Pete Warden§ Brian Plancher† Vijay Janapa Reddi*
*Harvard University §Stanford University †Barnard College, Columbia University

Abstract

The sustained growth of carbon emissions and global waste elicits significant sustainability concerns for our environment's future. The growing Internet of Things (IoT) has the potential to exacerbate this issue. However, an emerging area known as Tiny Machine Learning (TinyML) has the opportunity to help address these environmental challenges through sustainable computing practices. TinyML, the deployment of machine learning (ML) algorithms onto low-cost, low-power microcontroller systems, enables on-device sensor analytics that unlocks numerous always-on ML applications. This article discusses the potential of these TinyML applications to address critical sustainability challenges, as well as the environmental footprint of this emerging technology. Through life cycle analysis (LCA), we find TinyML systems present opportunities to offset their carbon emissions by enabling applications that reduce emissions from other sectors. Nevertheless, when globally scaled, the carbon footprint of TinyML systems is not negligible, necessitating designers factor in environmental impact when formulating new devices.[1]

1 Introduction

The continued growth of carbon emissions and global waste presents a great concern for our environment, increasing calls for a more sustainable future. In response, the United Nations' (UN) 2030 Agenda for Sustainable Development established a shared framework aiming toward peace and prosperity for people and the planet. At its core are 17 Sustainable Development Goals (SDGs) [28], a call to action for all countries to work towards a more environmentally, economically, and socially sustainable future.

Tiny machine learning (TinyML), which enables ML on microcontroller devices, holds potential for addressing numerous SGDs, particularly those related to environmental sustainability (see Figure 1). While TinyML's operational benefits for sustainability are often highlighted, it is crucial to consider the entire life cycle of both applications and hardware to ensure a net carbon reduction. This paper contributes by (1) presenting case studies illustrating TinyML's sustainability benefits, (2) examining the environmental impacts of TinyML at both MCU and system levels through a life cycle analysis (LCA), and (3) identifying future research directions for sustainable TinyML.

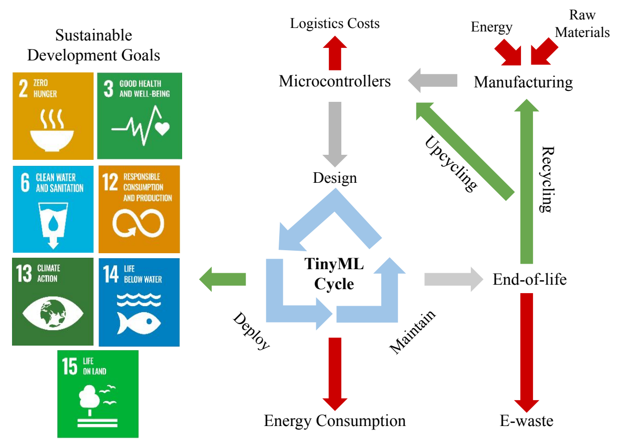

Figure 1. We show the positive (green arrows) and negative (red arrows) environmental footprint of the complete life cycle of TinyML systems as well as how TinyML can contribute to the UN's environmental sustainability goals.

2 Tiny Machine Learning (TinyML)

TinyML is the deployment of machine learning (ML) algorithms onto low-cost, low-power, and resource-constrained MCU systems. TinyML stores neural network models directly within memory (e.g., flash) and runs inference directly on the output of onboard sensors. This approach enables intelligent on-device sensor analytics unavailable with traditional Internet of Things (IoT) approaches, which instead typically rely on external cloud processing. Importantly, TinyML achieves this using a fraction of the compute resources needed for traditional ML systems. Table 1 shows how TinyML requires orders of magnitude fewer resources across compute, memory, storage, power, and cost than traditional BigML (such as cloud and mobile systems). With more than 250 billion MCUs deployed globally today, and the cost of MCUs expected to drop below $0.50 per unit, this number is expected to grow, eclipsing 40 billion MCUs shipped per annum in the next decade [22]. For these reasons, along with bandwidth, latency, energy, reliability, and privacy concerns, running ML directly on these embedded edge devices is growing in popularity. As such, TinyML will become an ever-present technology. But the question we must ask ourselves is *do we*

[1]Please refer to our full-length manuscript for citations: https://dl.acm.org/doi/abs/10.1145/3608473.

Platform	Freq.	Memory	Storage	Power	Price	CO_2-eq Footprint
Cloud	GHz	10+GB	TBs-PBs	~1 kW	$1000+	Hundreds of kgs
Mobile	GHz	Few GB	GBs	~1 W	$100+	Tens of kgs
Tiny	**MHz**	**KBs**	**Few MB**	**~1 mW**	**$10**	**Single kgs**

Table 1. Cloud and mobile ML systems compared with TinyML across frequency, memory, storage, power, price, and footprint. The footprint of TinyML systems is far less.

run the risk of producing an Internet of Trash over the course of TinyML devices' lifetime?

3 Applications of TinyML for Sustainability

To fairly evaluate the environmental impacts of machine learning on microcontrollers, we must consider TinyML's benefits. Typical well-known consumer-facing applications of TinyML include keyword spotting, image classification, and anomaly detection [3]. However, many emerging applications of TinyML can be used to enable a more sustainable future [2] and aid environment-related SDGs (Figure 1).

For example, Nuru, a mobile and cloud-based ML app from the PlantVillage project, helps increase agriculture production by detecting plant diseases and enabled one farmer to increase her revenue by 55% and yields by 146% [5, 17]. TinyML can also be used to aid in our health and well-being. Using Edge Impulse [10], a development platform for TinyML, a system was prototyped to identify the deadliest mosquitoes using wing beats sound classification with 88.3% accuracy [27].

TinyML can also boost conservation and biodiversity efforts by enhancing distributed sensing networks, such as resolving human-elephant conflicts in Asia and Africa. By only transmitting notifications of elephant detection instead of full video streams to the cloud, RESOLVE's WildEyes AI camera can run for more than 1.5 years on a single battery [8].

Combating climate change is another SDG that TinyML is well-suited for through environmental monitoring applications. For example, the SmartForest project utilizes a remote monitoring system to understand tree growth patterns. This replaced the need for 150-160 employees to regularly go into the field with a single sensor install trip [7]. Climate change has contributed to the widespread decline of essential pollinators like bumble bees [23]. TinyML can help provide intelligence to artificial pollinators like the Robobee [30]. TinyML can also further improve upon the 20-40% reduction in building energy usage [1, 16] enabled by smart occupancy systems that control lighting, automated window shading, and HVAC. See the full manuscript for further examples.

4 Quantifying the Sustainability of TinyML

The benefits of ML on microcontrollers for environmental sustainability and beyond will continue to fuel the Internet of Things (IoT) revolution, connecting billions of devices around us. To better understand the environmental costs associated with TinyML, an LCA of the complete TinyML system (i.e., MCU plus peripherals and power supply) is performed. This analysis demonstrates that the footprint of MCUs and TinyML systems *individually* is relatively small. When this analysis is expanded to consider the global scaled impact of TinyML, the impact could be substantial if TinyML is not used for sustainable applications.

4.1 Environmental Impact of MCUs

The TinyML life cycle analysis starts at the MCU level with publicly accessible data from STMicroelectronics [25].[2] The hardware life cycle of an MCU can typically be broken down into five stages: 1) extraction and treatment of raw materials, 2) product manufacturing, 3) transport and distribution, 4) product use, and 5) end of life. Taking these stages into account, there are four different environmental indicators, as shown in Figure 3, that can be used to analyze the footprint of the processing hardware required for TinyML: water demand, freshwater eutrophication, photochemical oxidant formation, and climate change. Across all four indicators, production is the dominant driver of an MCU's environmental footprint, as noted in previous work [14, 31]. However, the exact breakdown varies across indicators (Figure 3).

Overall, we find that the carbon footprint of an MCU is 390 g CO_2-eq. For perspective, this footprint is equivalent to a gasoline-powered car driving 1.6 km. Given that cars typically drive hundreds of thousands of miles over their lifetime, a single MCU alone has minimal impact in the context of everyday human actions. In the following section, CO_2 emissions are used as the primary measure due to their wide acceptance for assessing environmental impact.

4.2 Footprint of TinyML Systems

MCUs are the heart of TinyML systems, but we must consider the additional components that constitute a complete TinyML system to get a more accurate picture of the complete footprint. To do so, we developed an open-source *TinyML Footprint Calculator*.[3] Our calculator leverages the raw data from a study by Pirson and Bol [18] assessing the embodied carbon footprint of IoT devices. Pirson and Bol break down the general architecture and hardware profile of an IoT edge device into a collection of basic functional blocks but only capture the embodied footprint. As such, we additionally model and capture the product use stage (i.e., operational footprint) and end-of-life stage of the hardware life cycle.

Breaking Down TinyML's Footprint. The calculated footprint of TinyML systems is broken down into three scenarios. The "Low-Cost Profile" scenario represents a keyword spotting application that requires only a simple microphone sensor. The "Medium-Cost Profile" scenario represents an image classification application that requires a much larger camera sensor. The "High-Cost Profile" scenario

[2] The general trends hold for other MCU manufacturers.
[3] Available at https://github.com/harvard-edge/TinyML-Footprint

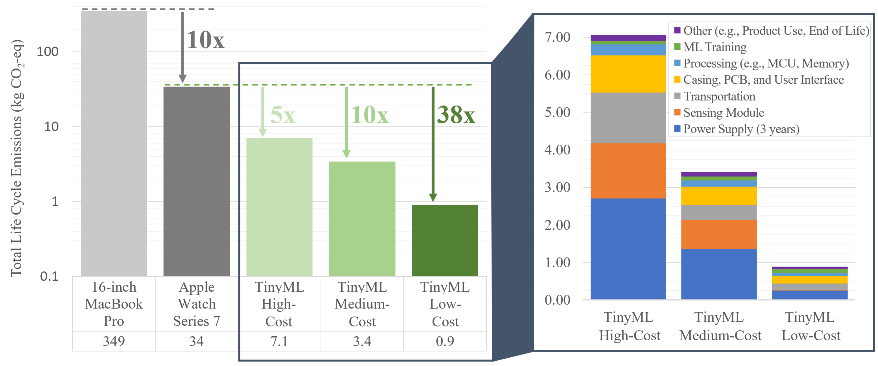

Figure 2. A breakdown of different TinyML system footprints highlights that the footprint is largely attributable to the embodied footprint of the power supply, onboard sensors, and transportation. Note that actuator and connectivity blocks from Pirson and Bol [18] are encapsulated in "Other" and "Processing", respectively, while "Product Use" captures the operational footprint. The carbon footprint of Apple's Series 7 Watch [12] and 16-inch MacBook Pro [11] are also provided for reference. For more details and to compute the footprint of your own TinyML system, see https://github.com/harvard-edge/TinyML-Footprint.

Total Impact	390g CO_2-eq 1.6km by car	23L 23 bottles of water	120mg P-eq 0.2 washing cycles	823mg NMVOC 2.7km by car
	Climate Change	Water Demand	Freshwater Eutrophication	Protochemical Oxidant Formation
End of Life	<1%	<1%	<1%	<1%
Logistics	1%	<1%	<1%	1%
Use	8%	6%	28%	8%
Raw Materials	10%	41%	27%	10%
Production: Other	25%	15%	18%	2%
Production: Energy Consumption	56%	39%	27%	71%

Figure 3. Four different environmental indicators measuring the impact of MCUs on our environment. Each footprint contains both the operational and embodied footprint of the device, including the five-stage life cycle of an MCU. Data courtesy of STMicroelectronics [25].

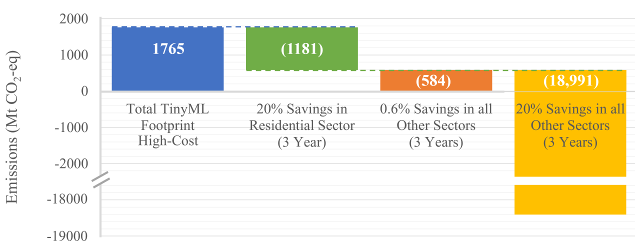

Figure 4. If all 250 billion MCUs were TinyML systems with three-year lifespans, their worst-case footprint would be 1765 million metric tons of CO_2. If these systems enabled a 20% emissions reduction for the residential sector and only a 0.6% reduction for all other sectors, the total footprint would be net-zero. Anything larger (e.g., 20%) results in more carbon savings from TinyML than emissions.

again uses image classification, instead using upper bound carbon emission values for components provided in Pirson and Bol [18].

As the stacked bar graph on the right side of Figure 2 shows, the embodied footprint of all components is much greater than the system's operational footprint (captured in "Product Use"). This result aligns with previous literature suggesting that manufacturing dominates the environmental footprint of small electronics [14, 31]. Moreover, the figure highlights the embodied footprint of the additional components excluding the MCU and the manufacturing and distribution costs. In particular, the embodied footprint of tne battery dominates all other components. To provide a baseline reference, we compared our results with the Apple Watch Series 7 [12] (representative of an "edge" device) and a 16-inch MacBook Pro (representative of traditional computing hardware), as shown on the left of Figure 2. We find that a TinyML system's footprint is 5-38× smaller than an Apple Watch [12] and 49-392× smaller than a Macbook [11].

4.3 TinyML at Scale

To better understand the net effect of TinyML at scale, this section assesses what happens to TinyML's footprint if these systems are scaled to the number of MCUs deployed globally, which currently sits at around 250 billion, using the "High-cost Profile", to provide an upper-bound. This scenario results in a combined, non-trivial global carbon footprint of 1765 million metric tons of CO_2-eq. However, there are existing examples (e.g., [1, 16]) of simple, intelligent IoT devices which can reduce building CO_2 emissions by at least 20%. If such savings were applied to the *entire* residential home sector over three years (*green bar* in Figure 4), 1181 million metric tons of CO_2-eq would be avoided. These avoided emissions alone would offset 67% of the worst-case costs of TinyML. As the residential sector only represents 6% of total global emissions, if remaining TinyML devices were able to reduce emissions from all other sectors by as little as 0.6% on average (*orange bar*), then TinyML would break even from an emissions standpoint. Furthermore, if we were to extrapolate this 20% reduction in the residential sector to all sectors (*yellow bar*) we would see a net reduction in global CO_2 emissions by over 18.4 billion metric tons.

5 Discussion

Prior claims regarding the use of digital technologies for greenhouse gas emissions mitigation do not always address critical aspects which can result in overestimated benefits [20]. Thus, in this section, we recognize limitations of our analysis and discuss factors to be considered in future work.

Limitations of Our Study: One major limitation is the lack of publicly available data on the environmental impact of modern digital electronics which makes it difficult for our analysis to be detailed and precise. Another important consideration is Jevons' paradox (or rebound effect), which suggests that advancements in efficiency can lead to an overall increase in consumption and a negative impact on the environment. Finally, it is also important to note that our analysis approximates carbon savings from TinyML solutions by comparing with a baseline with *no intervention*. However, alternative (non-TinyML) approaches could also be made to save emissions (e.g., behavioral changes or manual efforts by humans to reduce building emissions) that future works should compare with.

Considerations for Future Studies Several additional factors should be considered in future complementary studies. For example, our study assumes all 250 billion existing MCUs are TinyML systems but exponential IoT device growth could limit TinyML's impact. Moreover, our analysis assumes the MCU is fabricated using 90nm CMOS technology [18]. However, the environmental impact of semiconductor manufacturing increases with each successive technology node. For example, Bardon et al. [4] shows a 2.5× increase in greenhouse gas emissions *per wafer* when scaling from 28nm to 3nm. Furthermore, our study assumes a three-year device lifetime in order to compare with LCA's from other vendors. Finally, while our study concludes that TinyML devices can elicit an overall positive impact on the environment (with respect to carbon emission savings and global warming), it is important consider the many other environmental factors (Section 4.1), as well as societal and human costs.

6 Future Sustainable TinyML

In this section, we will discuss the broader implications of our study and suggest ways to make TinyML more sustainable.

Energy Harvesting. Our analysis in Section 4.2 revealed that the batteries used to power TinyML devices dominate their environmental impact. Batteries also present several other environmental issues, such as pollution and the release of carcinogens [15], particularly due to the extraction of lithium [26]. Research in energy harvesting [21] needs to be prioritized to make "batteryless" TinyML the standard practice. Furthermore, advancements in intermittent computing [9] could further reduce power requirements.

Efficient Sensing. Using smaller (e.g., camera vs. inertial measurement unit) or lower-quality sensors (e.g., low- vs. high-resolution camera) combined with more advanced TinyML models, or leveraging sensor fusion, using multiple small sensors, could potentially reduce the overall footprint while achieving the same performance [6].

Datasheets for ML Sensors. Greater transparency regarding the system's data and costs is needed to deploy these TinyML devices safely and ethically. One solution to address privacy concerns is to separate the input sensor data and ML processing from the rest of the system at the hardware level [29]. Also, new supplementary information is needed in the form of a datasheet that builds upon traditional datasheets used for electrical components to enable transparency to end users, including information about the environmental impact and LCA of the device [24].

Datasets for Low-Resource Domains. Many TinyML applications depend on real-world data, which can be challenging to obtain, particularly in public domains. To foster TinyML development, there is a need for extensive, open-access datasets focused on low-resource, high-impact sensor-based problems, akin to ImageNet for TinyML.

Emerging Technologies. New technologies are being developed that could lead to more sustainable TinyML practices. One example includes flexible electronics: PragmatIC Semiconductor has reported less than half of a single gram of CO_2-eq manufacturing such integrated circuits [19].

Recycle and Upcycle. TinyML can potentially exacerbate the problem of electronic waste. However, recycling and reusing TinyML devices is a viable option as many of the algorithms can run on standard MCU hardware, extending the MCU life and reducing the amount of landfill waste.

Accessibility. Finally, for TinyML to have a significant impact on a global scale, there is a need for global access to hardware and educational resources. Recent efforts, led by the TinyML foundation and the TinyML Open Education Initiative[4], among others, have developed open-source materials and provided low- or no-cost hardware to learners [13].

7 Conclusion

ML on microcontrollers can have a significant impact on environmental sustainability, potentially improving efficiency in various sectors and enabling significant reductions in carbon emissions. This assessment shows that TinyML's carbon footprint could be offset by using the technology to reduce emissions from other economic sectors. However, TinyML's footprint is not negligible when scaled globally, and thus designers must be mindful and factor in sustainability when developing new devices. Emerging technologies may further enable more sustainable computing practices and cement the net-positive potential of TinyML.

8 Acknowledgements

We thank Carole-Jean Wu, Udit Gupta, David Brooks, Danilo Pau, and Paul Whatmough for their guidance and support.

[4]https://www.tinyml.org/ and https://www.tinymledu.org/

References

[1] Parth Asopa, Poorva Purohit, Rahul Reddy Nadikattu, and Pawan Whig. 2021. Reducing Carbon Footprint for Sustainable development of Smart Cities using IoT. In *International Conference on Intelligent Communication Technologies and Virtual Mobile Networks*. 361–367. https://doi.org/10.1109/ICICV50876.2021.9388466

[2] Hatim Bamoumen, Anas Temouden, Nabil Benamar, and Yousra Chtouki. 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. https://doi.org/10.1109/3ICT56508.2022.9990661

[3] Colby Banbury, Vijay Janapa Reddi, Peter Torelli, Nat Jeffries, Csaba Kiraly, Jeremy Holleman, Pietro Montino, David Kanter, Pete Warden, et al. 2021. MLPerf Tiny Benchmark. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung (Eds.), Vol. 1.

[4] M Garcia Bardon, P Wuytens, L-Å Ragnarsson, G Mirabelli, D Jang, G Willems, A Mallik, A Spessot, J Ryckaert, and B Parvais. 2020. DTCO including sustainability: Power-performance-area-cost-environmental score (PPACE) analysis for logic technologies. In *2020 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 41–4.

[5] CGAIR: Platform for Big Data in Agriculture. 2018. PlantVillage Nuru. https://bigdata.cgiar.org/digital-intervention/plantvillage-nuru-pest-and-disease-monitoring-using-ai/

[6] Eyup Cinar. 2022. A Sensor Fusion Method Using Deep Transfer Learning for Fault Detection in Equipment Condition Monitoring. In *2022 International Conference on INnovations in Intelligent SysTems and Applications (INISTA)*. IEEE, 1–6.

[7] Rodrigo de Oliveira Andrade. 2019. Sensors take the manual work out of forest monitoring. https://www.scidev.net/global/news/sensors-take-the-manual-work-out-of-forest-monitoring/

[8] Eric Dinerstein and Sanjiv Fernando. 2020. WildEye AI. https://www.resolve.ngo/blog/WildEyes-AI-Helping-to-Save-Wild-Elephants-and-Prevent-Human-Elephant-Conflict.htm

[9] Graham Gobieski, Brandon Lucia, and Nathan Beckmann. 2019. Intelligence beyond the edge: Inference on intermittent embedded systems. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*. 199–213.

[10] Shawn Hymel, Colby Banbury, Daniel Situnayake, Alex Elium, Carl Ward, Mat Kelcey, Mathijs Baaijens, Mateusz Majchrzycki, Jenny Plunkett, David Tischler, et al. 2022. Edge Impulse: An MLOps Platform for Tiny Machine Learning. *arXiv preprint arXiv:2212.03332* (2022).

[11] Apple Inc. 2021. *Product Environmental Report: 16-inch MacBook Pro*. Technical Report. 9 pages. https://www.apple.com/environment/pdf/products/notebooks/16-inch_MacBook_Pro_PER_Oct2021.pdf

[12] Apple Inc. 2021. *Product Environmental Report: Apple Watch Series 7*. Technical Report. 9 pages. www.apple.com/in/environment/pdf/products/watch/Apple_Watch_Series7_PER_Sept2021.pdf

[13] Vijay Janapa Reddi, Brian Plancher, Susan Kennedy, Laurence Moroney, Pete Warden, Lara Suzuki, Anant Agarwal, Colby Banbury, Massimo Banzi, Matthew Bennett, Benjamin Brown, Sharad Chitlangia, Radhika Ghosal, Sarah Grafman, Rupert Jaeger, Srivatsan Krishnan, Maximilian Lam, Daniel Leiker, Cara Mann, Mark Mazumder, Dominic Pajak, Dhilan Ramaprasad, J. Evan Smith, Matthew Stewart, and Dustin Tingley. 2022. Widening Access to Applied Machine Learning With TinyML. *Harvard Data Science Review* (27 1 2022). https://doi.org/10.1162/99608f92.762d171a

[14] Alex K Jones, Yiran Chen, William O Collinge, Haifeng Xu, Laura A Schaefer, Amy E Landis, and Melissa M Bilec. 2013. Considering fabrication in sustainable computing. In *IEEE/ACM International Conference on Computer-Aided Design*. 206–210.

[15] Daniel Hsing Po Kang, Mengjun Chen, and Oladele A Ogunseitan. 2013. Potential environmental and human health impacts of rechargeable lithium batteries in electronic waste. *Environmental science & technology* 47, 10 (2013), 5495–5503.

[16] Jennifer King and Christopher Perry. 2017. *Smart buildings: Using smart technology to save energy in existing buildings*. Amercian Council for an Energy-Efficient Economy.

[17] Latifa M Mrisho, Neema A Mbilinyi, Mathias Ndalahwa, Amanda M Ramcharan, Annalyse K Kehs, Peter C McCloskey, Harun Murithi, David P Hughes, and James P Legg. 2020. Accuracy of a smartphone-based object detection model, PlantVillage Nuru, in identifying the foliar symptoms of the viral diseases of cassava–CMD and CBSD. *Frontiers in plant science* (2020), 1964.

[18] Thibault Pirson and David Bol. 2021. Assessing the embodied carbon footprint of IoT edge devices with a bottom-up life-cycle approach. *Journal of Cleaner Production* 322 (2021), 128966.

[19] PragmatIC. 2023. Pragmatic Life-Cycle Analysis White-Paper. https://www.pragmaticsemi.com/app/uploads/2023/04/Pragmatic-Life-Cycle-Analysis-White-Paper.pdf

[20] Aina Rasoldier, Jacques Combaz, Alain Girault, Kevin Marquet, and Sophie Quinton. 2022. How realistic are claims about the benefits of using digital technologies for GHG emissions mitigation? (2022).

[21] Adnan Sabovic, Michiel Aernouts, Dragan Subotic, Jaron Fontaine, Eli De Poorter, and Jeroen Famaey. 2022. Towards Energy-Aware Tinyml on Battery-Less IoT Devices. *Case Studies in Chemical and Environmental Engineering* (2022).

[22] Nikolaos Schizas, Aristeidis Karras, Christos Karras, and Spyros Sioutas. 2022. TinyML for Ultra-Low Power AI and Large Scale IoT Deployments: A Systematic Review. *Future Internet* 14, 12 (2022), 363.

[23] Peter Soroye, Tim Newbold, and Jeremy Kerr. 2020. Climate change contributes to widespread declines among bumble bees across continents. *Science* 367, 6478 (2020), 685–688.

[24] Matthew Stewart, Pete Warden, Yasmine Omri, Shvetank Prakash, Joao Santos, Shawn Hymel, Benjamin Brown, Jim MacArthur, Nat Jeffries, Brian Plancher, and Vijay Janapa Reddi. 2023. Datasheets for Machine Learning Sensors. arXiv:2306.08848 [cs.LG]

[25] STMicroelectronics. 2013. Footprint of a Microcontroller. https://www.st.com/content/st_com/en/about/st_approach_to_sustainability/sustainability-priorities/sustainable-technology.html

[26] Todd C. Frankel and Peter Whoriskey. 2016. TOSSED ASIDE IN THE 'WHITE GOLD' RUSH. https://www.washingtonpost.com/graphics/business/batteries/tossed-aside-in-the-lithium-rush/

[27] Kirankumar Trivedi and Harsh Shroff. 2021. Identification of deadliest mosquitoes using wing beats sound classification on tiny embedded system using machine learning and Edge Impulse platform. In *ITU Kaleidoscope: Connecting Physical and Virtual Worlds*. IEEE, 1–6.

[28] United Nations. 2021. The Sustainable Development Goals Report 2021. https://unstats.un.org/sdgs/report/2021/The-Sustainable-Development-Goals-Report-2021.pdf

[29] Pete Warden, Matthew Stewart, Brian Plancher, Sachin Katti, and Vijay Janapa Reddi. 2023. Machine Learning Sensors: A Design Paradigm for the Future of Intelligent Sensors. *Commun. ACM* 66, 11 (Nov 2023), 25–28. https://doi.org/10.1145/3586991

[30] Robert Wood, Radhika Nagpal, and Gu-Yeon Wei. 2013. Flight of the robobees. *Scientific American* 308, 3 (2013), 60–65.

[31] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. 2022. Sustainable ai: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems* 4 (2022), 795–813.