

This is a repository copy of *Towards Battery-Free Machine Learning and Inference in Underwater Environments*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/200234/>

Preprint:

Zhao, Yuchen orcid.org/0000-0003-4780-093X, Afzal, Sayed Saad, Akbar, Waleed et al. (5 more authors) (2022) *Towards Battery-Free Machine Learning and Inference in Underwater Environments*. [Preprint]

<https://doi.org/10.1145/3508396.3512877>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Towards Battery-Free Machine Learning and Inference in Underwater Environments

Yuchen Zhao*
Imperial College London
yuchen.zhao19@imperial.ac.uk

Sayed Saad Afzal*
MIT
afzals@mit.edu

Waleed Akbar
MIT
wakbar@mit.edu

Osvy Rodriguez
MIT
osvyrd@mit.edu

Fan Mo
Imperial College London
f.mo18@imperial.ac.uk

David Boyle
Imperial College London
david.boyle@imperial.ac.uk

Fadel Adib
MIT
fadel@mit.edu

Hamed Haddadi
Imperial College London
h.haddadi@imperial.ac.uk

Abstract

This paper is motivated by a simple question: Can we design and build battery-free devices capable of machine learning and inference in underwater environments? An affirmative answer to this question would have significant implications for a new generation of underwater sensing and monitoring applications for environmental monitoring, scientific exploration, and climate/weather prediction.

To answer this question, we explore the feasibility of bridging advances from the past decade in two fields: battery-free networking and low-power machine learning. Our exploration demonstrates that it is indeed possible to enable battery-free inference in underwater environments. We designed a device that can harvest energy from underwater sound, power up an ultra-low-power microcontroller and on-board sensor, perform local inference on sensed measurements using a lightweight Deep Neural Network, and communicate the inference result via backscatter to a receiver. We tested our prototype in an emulated marine bioacoustics application, demonstrating the potential to recognize underwater animal sounds without batteries. Through this exploration, we highlight the challenges and opportunities for making underwater battery-free inference and machine learning ubiquitous.

CCS Concepts

• **Computing methodologies** → **Supervised learning**; • **Hardware** → **Sensor devices and platforms**; **Renewable energy**; **Wireless integrated network sensors**; • **Applied computing** → **Environmental sciences**.

Keywords

battery-free networking, machine learning, underwater sensing

ACM Reference Format:

Yuchen Zhao, Sayed Saad Afzal, Waleed Akbar, Osvy Rodriguez, Fan Mo, David Boyle, Fadel Adib, and Hamed Haddadi. 2022. Towards Battery-Free

*Co-primary authors

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HotMobile '22, March 9–10, 2022, Tempe, AZ, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9218-1/22/03.

<https://doi.org/10.1145/3508396.3512877>

Machine Learning and Inference in Underwater Environments. In *The 23rd International Workshop on Mobile Computing Systems and Applications (HotMobile '22), March 9–10, 2022, Tempe, AZ, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3508396.3512877>

1 Introduction

In recent years, advances in battery-free sensing have enabled a range of novel applications including localization and environmental sensing. These advances have enabled long-term use of wireless sensors without the need for power supplies or batteries, hence enabling a transition towards a pollution-free sensing infrastructure for a variety of environmental sensing and monitoring applications such as detecting pollution and monitoring biodiversity. These applications are not limited to smart cities and urban environments, but also extend to climate change monitoring and environmental sensing in uninhabited areas such as forests, mines, remote areas, and space [1, 41].

The vast majority of existing battery-free technologies have been designed for land applications and very few have been engineered for the ocean. Only recently have researchers started looking into battery-free sensors for underwater environments. However, existing underwater battery-free technologies stop at energy harvesting for powering up sensors and backscatter for communicating raw sensory data [9, 11, 16]. Meanwhile, research into energy-efficient edge-based machine-learning (ML) models have enabled lightweight models to be used for on-device ML inference for a variety of applications. There have been a number of advances in model reduction, compression, and layer-wise quantization for various objectives such as privacy, energy, and efficiency [3, 21, 30]. A fundamental challenge that remains is to perform time-series data analysis and ML on underwater battery-free sensors. One benefit of such capabilities is that they would enable us to limit data transmission, hence increasing the energy budget on underwater devices and their operation longevity. This is particularly important for usages of underwater battery-free sensors in ambient and remote monitoring applications for the ocean, where communication bandwidth is narrow and the sources for energy harvesting are limited [31, 37].

In this paper, we propose and present a new challenge to the community, investigating the feasibility of battery-free ML in underwater environments, where extremely lightweight and task-specific




Figure 1: Design of battery-free inference on underwater edge nodes

deep neural network (DNN) models are executed on dedicated and highly-efficient underwater sensor nodes. We further investigate the tradeoffs between lightweight on-device analytics requirements for an exemplar case, which is marine mammal recognition in the ocean. This use case has many applications including monitoring biodiversity, understanding marine animal migration patterns, and even supporting the discovery of new species [27].

Our investigation focuses on two critical aspects of underwater battery-free ML, which are (1) the capability (i.e., model accuracy) of lightweight DNN models on the tasks to be conducted in underwater environments, and (2) the feasibility of hosting the DNN models on low-power underwater computational devices such as microcontrollers. Specifically on the task of marine mammal recognition, we design an end-to-end pipeline (i.e., from recording sounds to classification results on an underwater battery-free device) and evaluate our prototype’s accuracy and power consumption. Our preliminary results indicate that, for a task of classifying four marine mammals, lightweight DNN models can achieve decent accuracy. With the help of existing energy harvesting technologies and customized circuit design, it is possible to run lightweight ML on underwater battery-free devices. These results have important implications for both the mobile and ML communities, and pose exciting new opportunities for ubiquitous underwater battery-free ML.

2 Underwater Battery-free Architecture

We design a wireless, battery-free underwater system to perform sensing and inference on edge nodes and communicate with a receiver. To this end, we prototype a system that can perform marine mammal recognition in the ocean. Fig. 1 shows the system design for our prototype. The system harvests energy from underwater sound to power up its processor and on-board sensor that captures animal sounds (§2.1), performs on-board inference (§2.2), and transmits the result on the uplink via backscatter (§2.3).

2.1 Underwater Energy Harvesting

To enable battery-free operation, our sensor needs to harvest energy from ambient underwater sources such as sound, waves, or temperature gradients. Recent past research has demonstrated that sound provides sufficient energy to power up low-power microcontrollers [16]; hence, we design our node to harvest energy from underwater sound. It is worth noting that our node architecture is general and can be adopted to other underwater ambient energy sources.

To harvest energy from sound, our node employs piezoelectric materials, which convert mechanical energy (i.e. sound) to electrical energy. Since the harvested signal is typically an alternating current (sound is a wave), our sensor node rectifies it to a DC voltage using a multi-stage rectifier. Once the node harvests sufficient energy, it

can power up an on-board microcontroller. The microcontroller has an integrated ADC (analog to digital converter) which can be interfaced to a hydrophone (as shown in Fig. 1). The ADC samples received sound and stores it in memory.

2.2 Battery-free Underwater Inference

The next stage in our design is to recognize animals from the recordings. To do this, one option is to program the node to transmit its recordings to a remote receiver that has a dedicated power source (e.g., an underwater drone or a coastal base station) and perform inference there. However, such an approach is undesirable for multiple reasons. First, since the throughput of underwater acoustic channels is limited, transmitting the full recording would require our node to remain powered up for an extended period of time, which will drain its harvested energy. Second, the limited throughput would incur a large delay in data transmission, which is detrimental in time-sensitive scenarios (e.g., detection of endangered species). Furthermore, if a desired receiver is not within the communication range of our node, the node would need to store the recording in its constrained memory, which limits its ability to process new recordings.

To overcome these challenges, instead of transmitting the recordings for inference at a powered base station, our node performs on-board inference, as we describe in this section. This results in a lower-power, time-and-resource efficient, and more scalable system design (as we show in §3).

2.2.1 Offline Training: In order to have an accurate DNN model that can recognize marine mammals from their sounds, we need to use pre-collected audio clips (obtained from publicly available databases) to train the model until it achieves acceptable accuracy. Given the limited memory and computational resources of battery-free devices, it is difficult to perform the training process on them. Thus, we first train the high-level DNN models in a centralized manner (e.g., in the cloud) with offline training and validating data.

To ensure that our model would perform well on the target underwater battery-free device, we introduce a number of pre-processing steps (e.g., re-sampling, normalization, and reshaping) as shown in Fig. 2. These steps are necessary because the audio files (from databases) were collected using different equipment than our battery-free nodes (e.g., different ADC sampling rates and resolution). To mitigate the impact of these differences, the raw data from pre-collected clips is re-sampled with the sample rate of the nodes. The re-sampled data is then normalized and converted into a spectrogram through Short-Time Fourier transform (STFT).

To make the models work on a battery-free device, the high-level models need to be converted into a low-level model that satisfies the constraints of the architecture and memory of the target device. Our model conversion process uses techniques such as model compression and quantization to make the resulting low-level model work on the target device independently without support from run-time libraries.

2.2.2 On-board Edge Inference: After the model is trained and converted, we can load the model parameters to the memory of the target device and then use it for on-board inference. We use the same pre-processing pipeline (normalization, STFT etc.) as for offline training which was discussed in §2.2.1. One challenge with ultra-low-power microcontrollers is that, unlike standard audio

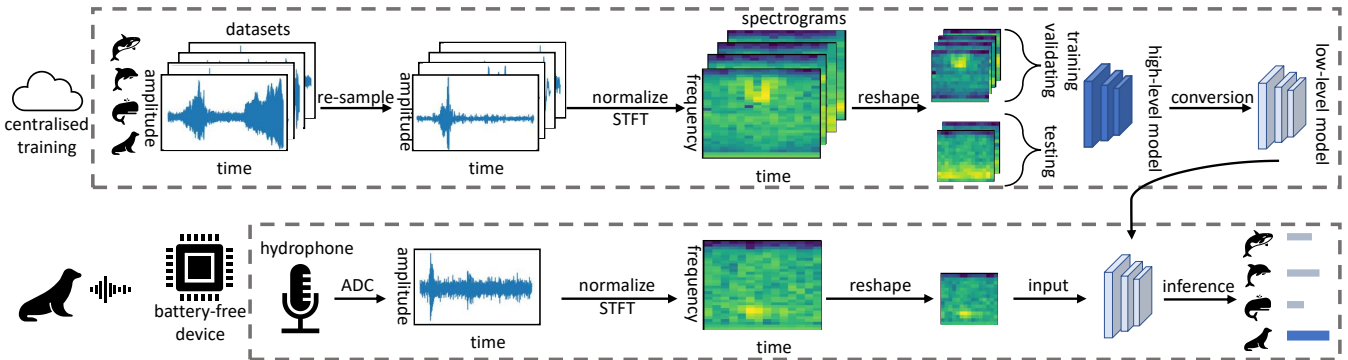


Figure 2: Overview of offline training and on-board inference. Initially, a high-level model is trained, validated, and evaluated offline on pre-collected audio clips from different marine mammals. A model with acceptable offline accuracy is converted to a low-level model that can be supported by standard C libraries on a target battery-free device. The converted model is then deployed in an end-to-end (i.e., from sounds received by a hydrophone to final classifications) marine mammal recognition pipeline on the target device.

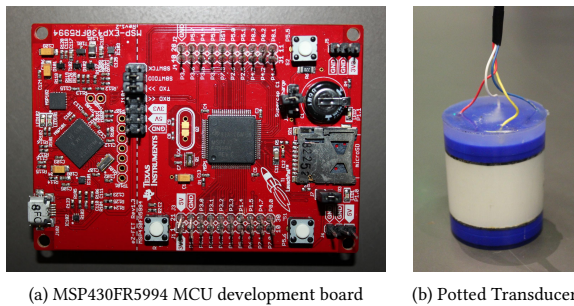


Figure 3: Prototype Hardware (a) shows the processor that our design uses for on-board inference. (b) shows one of our potted transducers which can be used for energy harvesting and communication.

cards, their ADCs cannot sample audio with both positive and negative amplitude. The lack of negative amplitude introduces another discrepancy between the pre-collected audio files and sound recordings collected using the microcontroller. To overcome this challenge, our design adds a clamping circuit to hardware design in order to passively add a DC offset to our input signal before we feed it to the on-board ADC. This DC offset is subsequently filtered in software by applying a high-pass filter; this ensures that the DC offset does not bias the STFT output or negatively impact the model performance.

2.3 From Inference to Communication

The last step in our design is to communicate the inference result to a remote receiver. To do this, we leverage underwater backscatter – a net-zero power communication technology – to communicate on the uplink, which allows us to reuse the same piezoelectric material for both harvesting (§2.1) and communication. In underwater backscatter, the piezoelectric material is connected to an impedance switch that modulates its reflection coefficient. At a high level, by alternating between two impedance states, the node communicates in binary. A remote receiver can sense changes in reflected power and use them to decode the transmitted packet by applying standard underwater decoding mechanisms.

3 Feasibility Study

We evaluated our prototype in terms of both its model accuracy and its power requirements.

3.1 Prototype Hardware

Backscatter node. Each backscatter node consists of a piezoelectric transducer and a hardware controller as shown in Fig. 3. To build our transducers, we used piezo-ceramic cylinders with a nominal resonance frequency of 17 KHz [38]. Similar to our previous work, the nodes were potted, encapsulated for insulation and matching to water, and housed in 3D-printed mold and end-caps [32].

Hardware controller. The node hardware is used for energy harvesting, processing/inference, and interfacing with a sensor (e.g., a hydrophone). Our harvesting architecture consists of a standard multi-stage rectifier followed by a super-capacitor, and low dropout voltage regulator which drives the digital logic unit. We implemented the logic on a MSP430-FR5994 microcontroller [39]. The microcontroller has a 12bit-ADC, which samples the (sound) signal at the rate of 330 samples/sec and stores a window of 512 samples in its SRAM. The microcontroller logic, including the code for downlink and uplink decoding as well as the inference model, is all stored in the FRAM. The classifier output is transmitted on the uplink by controlling the backscatter logic which modulates the load impedance of the backscatter node to change between reflective and non-reflective state.

Receiver. In addition to a backscatter node for energy harvesting and communication, our design uses a hydrophone (Omnidirectional Reson TC4014) that receives and decodes the FM0 encoded backscatter packets [4].

3.2 Offline Pipeline

For the evaluation, we used the Watkins Marine Mammal Sounds Database [33] and chose the sounds from eight mammals including Atlantic spotted dolphins, bearded seals, Beluga white whales, bottlenose dolphins, bowhead whales, harp seals, narwhals, and walrus. Overall we used 364 sound files and in each trial of our offline evaluation, we shuffled and split them as 80%, 10%, and 10% for training, validating, and testing. We re-sampled each file with the sample rate (330) of our prototype’s input device. The decoded sound signals were first normalized to $[-1.0, 1.0]$ and then the first 512 signals were transformed into a spectrogram through STFT (window size = 64, window step = 32).

We used TensorFlow to train high-level Keras models and then used Keras2C [5] to convert them into low-level models that are supported by standard C libraries. This library generates smaller




Figure 4: CNN structure for the classification task. The convolutional layer has 8 filters (size 3×3 , strides 2×2) activated by ReLU. The dense layer is activated by softmax.




Figure 5: Distribution (boxplots and kernel density) of offline test accuracy CNN models with different input layer size (N) and number of mammal classes (C). The converted model of (N-16, C-4) can be deployed on our prototype.

DNN models than TensorFlow Lite does and requires fewer third-party libraries when generating models. We used a lightweight convolutional neural network (CNN) model as shown in Fig. 4. The spectrogram of normalized input signals is first reshaped as a 2-D input layer with size of $N \times N$. The input layer is connected to a convolutional layer that has 8 filters, each of which has the kernel size of 3×3 and strides of 2×2 . The activated output of the convolutional layer is flattened and connected to a dense layer, which outputs the probability distribution of a classification.

3.3 Model Accuracy

We evaluated the offline accuracy of the model with different size N of the input layer and number of mammal classes, as these two hyperparameters affected the size of the model. For 4-mammal classification, we used the data from Atlantic spotted dolphins, bearded seals, beluga white whales, and narwhals. Fig. 5 shows the offline accuracy distribution of each model hyperparameter configuration from 64 trials. For 4-mammal classification, the converted models when $N = 8$ and $N = 16$ can fit in our prototype, while $N = 32$ demands a larger model which didn't fit in the memory of our microcontroller. The average accuracy for $N = 16$ and $N = 32$ is 84% and 88% respectively. For 8-mammal classification, with $N = 32$, the average accuracy is 76%, which requires larger memory than that on our prototype.

We further evaluated the model on the online data generated from the ADC on our prototype. To this end, we deployed the converted 4-mammal classification model into the pipeline on our prototype. We then randomly selected 16 raw sound files (4 samples

for each mammal class) from the dataset. We used SIGLENT SDG 1032X function generator [35] to convert raw sound files (animal sounds [33] of whales, seals, dolphins, and narwhals) to analog signals that were fed to the ADC. For each input signal, the microcontroller sampled the input analogue signal at 330 samples/s, normalized it, and computed the spectrogram of the processed data before running inference on it. The average accuracy of the online model was 63% (while random guess is 25%) which is lower than the offline accuracy. The lower online accuracy can be attributed to low number of testing samples and loss of data resolution during sampling. Nonetheless, this result demonstrates that our prototype was successful in on-board classification of four marine animals with a reasonable accuracy at extremely low power, and this accuracy may be improved as the research evolves and with more comprehensive online evaluation.

3.4 Power Performance

Next, we used the MSP-EXP430FR5994 development kit board to evaluate the power consumption of the microcontroller while it was running the end-to-end pipeline. We computed the power as the product of the current and the voltage when the board was connected to a power supply. We used SIGLENT SDG 1032X function generator [35] to power on the board at 1.9 V. To measure the current, we connected Nordic Semiconductor's Power Profiler kit [34] as an ammeter in series with the board. Table 1 shows the power consumption of each stage of the pipeline. At 1.9 V input voltage, the microcontroller consumed 932 μW during ADC sampling, 1.3 mW while running inference on the data, and 902 μW when it was backscattering the inference result. These power measurements are in line with those expected from the microcontroller's datasheet. Since prior work has demonstrated the potential to harvest up to few milliwatts from underwater acoustic and ultrasonic signals [12], these results demonstrate that it would be possible for our node to operate entirely based on harvested energy.

Stage	Power (μW)	Time (s)	Energy (mJ)
ADC Sampling	932	1.6	1.49
Inference	1300	3.0	3.91
Backscatter	902	0.012	0.0108
Total Energy Consumption: 5.40 mJ			

Table 1: On-board power and energy analysis of the battery-free inference prototype with a data rate of 1 kbps and a sampling rate of 330 samples/s.

Since the power budget for ADC sampling and backscatter communication is similar, one could argue that it may be more efficient to send the entire data to the base station rather than performing the on-board inference. However, recall from §2.2 that this approach will be highly inefficient because of the limited throughput of the underwater channel. To demonstrate this inefficiency, consider the following example: For a backscatter node transmitting at 1 kbps (which is standard throughput for underwater acoustic modems), it will take around 6.14 seconds to send 512 samples of raw data (total # of bits = 512×12). The ADC consumes 1.49 mJ of energy to sample 512 samples at a sampling rate of 330 samples/s. This translates to 7.03 mJ of required energy. On the other hand, performing on-board inference, which typically takes 3 seconds (consuming 3.91 mJ) then transmitting the 12-bit-long inference result (consuming 0.0108 mJ)

would require a total of 5.40 mJ, which makes sending raw data 30.19% more power hungry than sending inference data. This result demonstrates that battery-free edge inference is feasible, and that it makes on-board inference both more efficient and faster, even with this preliminary inference model. We envision that the gains can be significantly improved as the research evolves (see §5).

4 Related Work

Battery-free underwater sensing. The past few years have seen multiple advances in underwater battery-free sensing and computing [8, 11, 15, 16]. These advances have been propelled by three key trends. First is the downward trend in the power consumption of electronics, which has resulted in ultra-low-power processors (micro-controllers and FPGAs) that can operate in the micro-Watt regime [23, 39]. The second trend is the development of ultra-low-power underwater communication mechanisms, specifically piezo-acoustic backscatter [8, 16], which operates at net-zero power. Third is the continued improvement in the efficiency of energy harvesting from ambient underwater sources (sound, movements [2], and temperature gradients [20]). The combination of these factors has led to underwater battery-free sensors that can compute and communicate while powering up entirely from harvested energy. Our work builds on this line of research and extends it via battery-free inference, paving the way for a new generation of underwater battery-free devices capable of learning and inference.

Low-power machine learning. The past few years have also seen significant advances in low-power machine learning models. Researchers have developed techniques to reduce the size and power consumption of advanced learning models to make them suitable for running on mobile and edge devices. These techniques include knowledge distillation, model pruning, quantization, or replacing a model’s structure with a suitable one through network architecture searching [6, 14, 25]. Motivated by these advances, multiple research projects have explored the potential to perform machine learning and inference on battery-free devices [10, 18, 26, 29]. However, none of these past systems are suitable for underwater environments since they rely on RF and/or large solar panels to harvest energy - neither of which are available in the deep sea. Our research is motivated by a similar desire to enable battery-free machine learning and inference, and is the first to bring such capabilities to underwater environments.

Machine learning for underwater applications. Finally, there has been significant interest in applying machine learning for various underwater problems, such as recognition of marine animals using their sound or images [22, 43], classifying seabeds [7], robotic navigation [13], and image enhancement [40]. However, past systems for underwater ML required instruments with dedicated energy sources (typically underwater drones). Our research builds on this line of work and is the first to demonstrate the potential for battery-free underwater inference and machine learning.

5 Open Problems and Opportunities

In this paper, we proposed and prototyped battery-free ML inference in underwater environments as an exciting challenge for the community. We have shown that a lightweight CNN model can achieve decent accuracy while its on-board energy requirements can be supported by these devices. Our initial results indicate

feasibility of performing sensing and inference. Future potential directions include:

Battery-free on-device personalization. Our design assumed that the same low-level model would be deployed on all battery-free devices deployed underwater. However, when implementing battery-free ML at a large scale, we expect that the sensed data by different underwater devices may differ at different locations (e.g., animal sounds near or far from a coast) and with environmental factors (e.g., temperature, pressure, multipath all impact sounds). Therefore, models on the devices would need to be locally tuned to fit different testing data distributions. Previous research [24] proposes to use early exit during offline training so that an end-device can personalize a part of the model to improve its testing accuracy. To enable such on-device personalization, light-weight model training algorithms that can satisfy the computational constraints on battery-free devices are needed.

Model-optimized hardware/software design. Currently there are limited compilers that can build DNN models on ultra-low power MCUs (e.g., the MSP430 class) such as [36]. Co-design of task-specific hardware components alongside the ability to compile advanced models for lower-power MCUs can enable a wealth of new environmental monitoring applications. Specifically, one direct approach is to adopt existing optimization techniques for online inference on a target platform, such as post-training quantization [19, 28]. It can reduce the sizes of on-board models and increase computational efficiency by avoiding floating operations. Another “engineering” solution is to optimize the implementations of online inference for different target platforms by using hardware-specific instructions and libraries that offer faster or more energy-efficient operations. An auto-design pipeline for such individualization would be worthwhile considering the diversity of IoT devices and used low-power MCUs.

Battery-free distributed ML training. Given the potentially large scale deployment of low-cost and low-pollution battery-free nodes in the ocean, another future direction is to continuously sense data and train ML models in a decentralized manner among participating nodes. This will enable life-long ML with up-to-date data that reflects the changing nature of underwater environments, which is critical for many application domains including climate change and biodiversity monitoring. Similar to on-device personalization, the key challenge is to have light-weight ML training algorithms on the devices. Another challenge is the difficulty of obtaining labelled data from underwater environments. One option is to use unsupervised ML to only train feature extractors (e.g., autoencoders) that can help train classifiers on a cloud server [42]; another option is to generate pseudo labels on local data through data augmentation [17]. In addition, distributed/federated ML solutions that can coordinate devices, stations, and cloud servers are also needed to manage the local training tasks, to make sure that they can satisfy the energy constraints of battery-free underwater devices.

Acknowledgments. We thank our shepherd, the anonymous reviewers, and the Signal Kinetics group for their feedback. This research is supported by the Office of Naval Research, the MIT Media Lab, the Doherty Chair in Ocean Utilization, and the MIT-Imperial Seed Fund.

References

- [1] Mohamed R Abdelhamid, Ruicong Chen, Joonhyuk Cho, Anantha P Chandrakasan, and Fadel Adib. 2020. Self-reconfigurable Micro-implants for Cross-tissue Wireless and Batteryless Connectivity. In *MobiCom '20*. 1–14.
- [2] JJ Allen and AJ Smits. 2001. Energy Harvesting Eel. *Journal of fluids and structures* 15, 3–4 (2001), 629–640.
- [3] Ranya Aloufi, Hamed Haddadi, and David Boyle. 2020. Privacy-Preserving Voice Analysis via Disentangled Representations. In *CCSW '20*.
- [4] Aquarian Audio. 2019. H2a Hydrophone. <http://www.aquarianaudio.com/h2a-hydrophone.html>.
- [5] Rory Conlin, Keith Erickson, Joseph Abbate, and Egemen Kolemen. 2021. Keras2c: A Library for Converting Keras Neural Networks to Real-time Compatible C. *Engineering Applications of Artificial Intelligence* 100 (2021), 104182.
- [6] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2019. Neural Architecture Search: A Survey. *JMLR* 20, 1 (2019), 1997–2017.
- [7] Christina Frederick, Soledad Villar, and Zoi-Heleni Michalopoulou. 2020. Seabed Classification Using Physics-based Modeling and Machine Learning. *The Journal of the Acoustical Society of America* 148, 2 (2020), 859–872.
- [8] Reza Ghaffarivardavagh, Sayed Saad Afzal, Osvy Rodriguez, and Fadel Adib. 2020. Ultra-wideband Underwater Backscatter via Piezoelectric Metamaterials. In *SIGCOMM '20*. 722–734.
- [9] Reza Ghaffarivardavagh, Sayed Saad Afzal, Osvy Rodriguez, and Fadel Adib. 2020. Underwater Backscatter Localization: Toward a Battery-Free Underwater GPS. In *HotNets '20*. 125–131.
- [10] Graham Gobieski, Brandon Lucia, and Nathan Beckmann. 2019. Intelligence Beyond the Edge: Inference on Intermittent Embedded Systems. In *ASPLOS '19*. 199–213.
- [11] Raffaele Guida, Emrehan Demirors, Neil Dave, and Tommaso Melodia. 2020. Underwater Ultrasonic Wireless Power Transfer: A Battery-less Platform for the Internet of Underwater Things. *IEEE TMC* (2020).
- [12] Raffaele Guida, Emrehan Demirors, Neil Dave, and Tommaso Melodia. 2020. Underwater Ultrasonic Wireless Power Transfer: A Battery-less Platform for the Internet of Underwater Things. *IEEE Transactions on Mobile Computing* (2020).
- [13] Peter Gunnarson, Ioannis Mandralis, Guido Novati, Petros Koumoutsakos, and John Dabiri. 2021. Robotic Implementation of Online Deep Reinforcement Learning for Autonomous Underwater Navigation. *Bulletin of the American Physical Society* (2021).
- [14] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. [arXiv:1503.02531 \[stat.ML\]](https://arxiv.org/abs/1503.02531)
- [15] Guanjie Hu, Junliang Lin, Gongpu Wang, Ruisi He, and Xusheng Wei. 2021. Sparse Reconstruction Based Channel Estimation for Underwater Piezo-Acoustic Backscatter Systems. In *IEEE VTC '21*. 1–5.
- [16] Junsu Jang and Fadel Adib. 2019. Underwater Backscatter Networking. In *SIGCOMM '19*. 187–199.
- [17] Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. 2021. Federated Semi-Supervised Learning with Inter-Client Consistency & Disjoint Learning. In *ICLR '21*.
- [18] Seulki Lee, Bashima Islam, Yubo Luo, and Shahriar Nirjon. 2019. Intermittent Learning: On-device Machine Learning on Intermittently Powered System. *ACM IMWUT* 3, 4 (2019), 1–30.
- [19] Zhenhua Liu, Yunhe Wang, Kai Han, Wei Zhang, Siwei Ma, and Wen Gao. 2021. Post-training quantization for vision transformer. *Advances in Neural Information Processing Systems* 34 (2021).
- [20] Zhesong Ma, Yanhui Wang, Shuxin Wang, and Yanan Yang. 2016. Ocean Thermal Energy Harvesting with Phase Change Material for Underwater Glider. *Applied Energy* 178 (2016), 557–566.
- [21] Mohammad Malekzadeh, Richard Clegg, Andrea Cavallaro, and Hamed Haddadi. 2021. DANA: Dimension-Adaptive Neural Architecture for Multivariate Sensor Data. *ACM IMWUT* 5, 3, Article 120 (Sept. 2021).
- [22] David K Mellinger and Christopher W Clark. 2006. MobySound: A Reference Archive for Studying Automatic Recognition of Marine Mammal Sounds. *Applied Acoustics* 67, 11–12 (2006), 1226–1242.
- [23] MicroSemi. 2019. IGLoo nano FPGA. <https://www.microsemi.com/product-directory/fpgas/1689-igloo#igloo-nano>.
- [24] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. 2021. PPFL: Privacy-preserving Federated Learning with Trusted Execution Environments. In *MobiSys '21*. 94–108.
- [25] Fan Mo, Ali Shahin Shamsabadi, Kleomenis Katevas, Soteris Demetriou, Ilias Leontiadis, Andrea Cavallaro, and Hamed Haddadi. 2020. Darknetz: Towards Model Privacy at the Edge Using Trusted Execution Environments. In *MobiSys '20*. 161–174.
- [26] Alessandro Montanari, Manuja Sharma, Dainius Jenkus, Mohammed Alloulah, Lorena Qendro, and Fahim Kawsar. 2020. ePerceptive: Energy Reactive Embedded Intelligence for Batteryless Sensors. In *SensSys*. 382–394.
- [27] Camilo Mora, Derek P Tittensor, Sina Adl, Alastair GB Simpson, and Boris Worm. 2011. How Many Species Are There on Earth and in The Ocean? *PLoS biology* 9, 8 (2011), e1001127.
- [28] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. 2020. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*. PMLR, 7197–7206.
- [29] Matteo Nardello, Harsh Desai, Davide Brunelli, and Brandon Lucia. 2019. Camaroptera: A Batteryless Long-range Remote Visual Sensing System. In *ENSys'19*. 8–14.
- [30] Seyed Ali Osia, Ali Shahin Shamsabadi, Sina Sajadmanesh, Ali Taheri, Kleomenis Katevas, Hamid R. Rabiee, Nicholas D. Lane, and Hamed Haddadi. 2020. A Hybrid Deep Learning Architecture for Privacy-Preserving Mobile Analytics. *IEEE IoT Journal* 7, 5 (2020), 4505–4518.
- [31] Nasir Saeed, Abdulkadir Celik, Tareq Y Al-Naffouri, and Mohamed-Slim Alouini. 2018. Energy harvesting hybrid acoustic-optical underwater wireless sensor networks localization. *Sensors* 18, 1 (2018), 51.
- [32] Sayed Saad Afzal, Osvy Rodriguez, Reza Ghaffarivardavagh, JunSu Jang, Jose Muguira Iturralde, Fadel Adib. 2020. Underwater Backscatter Open Source. <https://github.com/saadafzal24/Underwater-Backscatter>
- [33] Laela Sayigh, Mary Ann Daher, Julie Allen, Helen Gordon, Katherine Joyce, Claire Stuhlmann, and Peter Tyack. 2016. The Watkins Marine Mammal Sound Database: An Online, Freely Accessible Resource. *POMA* 27, 1 (2016), 040013.
- [34] NORDIC Semiconductor. 2019. NORDIC Semiconductor Power Profiler Kit II. <https://www.nordicsemi.com/Products/Development-hardware/Power-Profiler-Kit-2>. PPK2.
- [35] SIGLENT. 2019. SIGLENT SDG1032X function generator. <https://siglentna.com/product/sdg1032x/>.
- [36] Benjamin Sliwa, Nico Piatkowski, and Christian Wietfeld. 2020. LIMITS: Lightweight machine learning for IoT systems with resource limitations. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*. IEEE, 1–7.
- [37] B Sai Srujana, Princy Mathews, VP Harigovindan, et al. 2015. Multi-source energy harvesting system for underwater wireless sensor networks. *Procedia Computer Science* 46 (2015), 1041–1048.
- [38] Steminc. 2019. Piezo Ceramic Cylinder 54.1x47x40mm 17 KHz. <https://www.steminc.com/PZT/en/piezo-ceramic-cylinder-541x47x40mm-17-khz> Part No. SMC5447T40111.
- [39] TI. 2021. MSP430FR5994. <https://www.ti.com/product/MSP430FR5994>.
- [40] Hao-Hsiang Yang, Kuan-Chih Huang, and Wei-Ting Chen. 2021. Laffnet: A Lightweight Adaptive Feature Fusion Network for Underwater Image Enhancement. *arXiv preprint arXiv:2105.01299* (2021).
- [41] Xu Yang, Xiaojie Yu, Chi Zhang, Sen Li, and Qiang Niu. 2021. MineGPS: Battery-free Localization Base Station for Coal Mine Environment. *IEEE Communications Letters* 25, 8 (2021), 2579–2583.
- [42] Yuchen Zhao, Hanyang Liu, Honglin Li, Payam Barnaghi, and Hamed Haddadi. 2021. Semi-supervised Federated Learning for Activity Recognition. [arXiv:2011.00851](https://arxiv.org/abs/2011.00851) [arXiv: 2011.00851](https://arxiv.org/abs/2011.00851).
- [43] Peiqin Zhuang, Linjie Xing, Yanlin Liu, Sheng Guo, and Yu Qiao. 2017. Marine Animal Detection and Recognition with Advanced Deep Learning Models. In *CLEF (Working Notes)*.