Poster Abstract: On-Device Training from Sensor Data on Batteryless Platforms

Bahsima Islam, Yubo Luo, Seulki Lee and Shahriar Nirjon
University of North Carolina at Chapel Hill
{bashima,yubo,seulki,nirjon}@cs.unc.edu

ABSTRACT

In this paper, we argue that the fusion of machine learning (ML) and batteryless computing systems enables true lifelong learning in mobile devices. The lack of learning from experience in current batteryless systems makes them ignorant of changes in their operating environment. Due to high communication cost, latency, privacy, and dependency issues of offloading computation to an edge device, on-device training is a solution for batteryless systems to learn and adapt in dynamically changing environments. Combining batteryless systems and ML is however a challenging task. Sporadic energy supply and limited resources in a batteryless system cause execution-discontinuity and data-constraints in ML processes. To understand these challenges, we identify suitable ML tasks for such systems and study the energy producers, i.e., harvesters, and consumers, i.e., intermittently executable tasks in a ML pipeline. Using a trace-driven simulation, we demonstrate the feasibility of on-device training of a batteryless learner.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Computer systems organization** → **Embedded hardware**; • **Hardware** → **Sensor devices and platforms**; **Renewable energy**; *PCB design and layout*.

1 INTRODUCTION

With the growth of the Internet of Things (IoT), the number of connected embedded systems around us is increasing exponentially. A large number of IoT devices are battery-powered due to the need of mobility. Recharging or replacing batteries of a large number of devices is inefficient and inconvenient. To address this energy issue, batteryless systems have been proposed. Unlike battery-powered devices, batteryless systems can operate, in principle, forever – as long as the harvesting conditions are met. The burden of battery-related maintenance being lifted, such systems can be deployed in many applications where a battery-powered system is not feasible, e.g., monitoring wildlife, remote surveillance, and implantable devices. The majority of existing batteryless devices are used in sensor networks. They merely sense data and send them to another system via a network. Some of these systems perform slightly advanced computational tasks, e.g., analyzing sensor data and performing an action accordingly. Recently, on-device deep-inference has been demonstrated in a batteryless system [3]. However, none of these systems adapt to the changes in the environment.

We define *learning* for batteryless systems as their ability to recognize patterns, anomalies, or otherwise useful information (e.g., application-specific knowledge) from observed data over a long period by monitoring the surrounding environment. Our goal is to develop batteryless systems embodying this definition of learning and to show its feasibility through experiments. The fusion of machine learning (ML) and batteryless systems enables true lifelong learning in mobile devices. The capability of independently learning without any concern for energy and any assistance from an external device reduces the complexity of smart device deployment. For instance, wildlife monitoring devices are highly difficult to track and reprogram once they are deployed. By using batteryless learning devices, animal habits can be reliably learned throughout an animal's lifetime without requiring human interventions. Other examples include monitoring and pattern mining in harsh environments or places without an easy access.

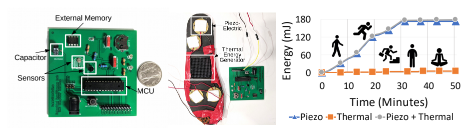

(a) (b) (c)

Figure 1: (a) Custom PCB (b) Prototype Learning Shoe (c) Power generated from Shoe

Combining ML with a batteryless system poses unique challenges. First, sporadic energy harvesting pattern enforces a ML algorithm to run intermittently and the system can miss important incoming data. Thus, the presence of both data occurrence and energy availability needs to be ensured. Second, in order to learn with constrained resources, a device needs to make necessary decisions by itself regarding which data to learn and when. Third, the type of learning algorithm a device can execute is limited since the supervision or ground truth is generally not available. Suitable options are online unsupervised, semi-supervised or reinforcement learning. A learning algorithm's computational complexity, learning time, or amount of data to learn should be carefully taken into account based on the energy budget.

In order to analyze the feasibility of batteryless learning systems, accurate measurements of energy produced by a harvester and consumed by a learner has to be attained. We study the amount of harvested energy from a variety of harvesters and measure the amount of consumed energy by different machine learning subtasks of an online unsupervised learning algorithm; e.g. K-means clustering; from both algorithmic and energy-level perspectives. We implement a prototype applications of batteryless learning, i.e.,

IPSN '19, April 16–18, 2019, Montreal, QC, Canada

ACM ISBN 978-1-4503-6284-9/19/04.
https://doi.org/10.1145/3302506.3312611

a learning shoe, that use a custom PCB for energy management. Finally, we use them in a trace-driven simulated learning application to demonstrate the feasibility of batteryless learning.

2 SYSTEM DESIGN

• **Custom PCB Hardware** We have developed an energy-harvesting platform which harvests energy from ambient sources(e.g. solar, piezo-electric) to execute machine learning operations. As shown in Figure 1(a), the board consists of a microcontroller, sensors, an external non-volatile memory, and a supercapacitor. The system follows Hibernus [1] design to save the system status during intermittence.

Model Training								Overhead			Total
Sample (%)	Sample Size	Accuracy (%)	Precision (%)	Recall (%)	Energy (mJ)	Step	Time (hh:mm)	Energy	Step	Time (hh:mm)	Time (hh:mm)
10	35	52	39.33	97.75	11	120	00:17	7	77	00:11	00:28
20	70	52	42	100	21.75	237	00:34	14	152	00:22	00:55
30	105	57	40.47	100	32.61	355	00:50	21	228	00:33	01:23
40	140	53	48	100	43.48	473	01:07	28	304	00:43	01:50
50	175	64.67	50.67	100	54.30	590	01:24	35	380	00:54	02:17
60	210	72.33	50	100	65.25	709	01:40	42	457	01:05	02:45
70	245	73.67	50.67	100	76	826	01:57	49	533	01:16	03:12
80	280	77.33	54.67	100	86.85	944	02:13	56	609	01:26	03:39
90	315	77.33	54.67	100	97.70	1062	02:30	63	685	01:37	04:07
100	350	80	60	100	108.54	1180	02:47	70	761	01:48	04:34

Table 1: The estimated energy, time, and performance of the learner for various amounts of training is shown.

• **Custom Shoe Prototype** We build a prototype *learning shoe* which harvests thermo-electric and piezo-electric energy. We include a thermal source to gather power from body heat gradient even when the wearer is not stepping. The shoe consists of six piezo disks and three thermoelectric generators (TEG) (connected in parallel), as shown in Figure 1(b). This prototype is envisioned to be used in applications such as gait analysis, activity recognition, and user identification. We conduct a detailed experiment to measure the power generated by the prototype shoe. We simulate five gait positions, i.e., walking, running, stepping stairs, standing and sitting, for 10 minutes each as shown in Figure 1(c).

3 ENERGY CONSUMPTION

Every machine learning task consists of the following methods – gathering data, preparing data (data selection, feature extraction), training a model, quality assessment, and prediction. On the other hand, every intermittently-powered system has some overhead, e.g., storing and retrieving program state. In this section, we study the power consumption of these methods to understand the characteristics of an intermittent learner.

Some common sensors used in batteryless devices are light sensor (2.34μW), temperature sensor (27.6μW), audio sensor (0.65mW), accelerometer (0.9 mW), pressure sensor (27μW), image sensor (4mW), and CO sensor (45μW). After collecting the data, there are four major methods in a machine learning task – extracting features, learning a model, evaluating the model and predicting the result. As features we calculate mean and standard deviation of data as features which consumes 22.25μ J. Due to the low energy budget, we avoid redundant data processing by measuring the diversity of a data sample using the standard deviation of the distances between consecutive samples(0.011 μJ). We implement online k-means clustering and divide it into training (updating cluster centroid) and predicting (predicting cluster) phases. These consume 108.5 mJ and 112.75 μJ respectively. We assess the quality of the model using weighted inter-intra cluster index (20.68 J). We consider a data buffer size of 350 samples and it requires 385 mJ and 397 μJ for writing and reading 1KB of system state in the internal EEPROM. By using external EEPROM this consumption changes to 15–241 mJ and 35 mJ for writing and reading respectively.

4 SIMULATED LEARNING SCENARIO

In order to simulate a learning shoe, we collect step counts of a user for 24 hours using a fitness tracker. We observe that a person takes around 7,000 steps daily, and about 213 steps per 30 minutes on average during active hours. We use a leg-mounted accelerometer dataset [2] to learn and predict different activities using k-means clustering algorithm. We choose two similar activities (running and playing basketball) to ensure overlapping clusters. We split the dataset into 70%-30% for training and validating the classifier and perform 10-fold cross validation.

Table 1 shows the model quality, the corresponding required energy, steps and time for different data sizes. Three major factors are to be noted in the table. First, by learning only 60% of the total data, we can learn a model with comparable accuracy in 2 hours 45 minutes instead of 4 hours 34 minutes. Thus by learning only a selected subset of data samples, we can learn a similar model using much less energy. Next, the significant overhead of intermittently powered system needs to be considered while designing such system. To minimize this overhead cost, different non-volatile memories (e.g., FRAM) needs to be exploited. Finally, it is evident that learning with harvested energy is possible. By using better energy harvesters, efficient memory, and optimized ML techniques, the learning time can be significantly minimized.

Next, we simulate a batteryless learner using the measurements from previous sections. We observe that learning a data point, including data gathering, selection, feature extraction (i.e., means and standard deviations), training online k-means clustering model and prediction require 108 mJ. Storing and retrieving data points (1.5 KB, 4-byte float) and program states (100 bytes) consume 120 mJ. To improve the system's energy efficiency, learning from fewer data samples is a possible solution. As mentioned previously, energy generated by the learning shoe prototype for each step is 92μJ. Therefore, approximately 1,181 steps are required to learn 350 data points and another 815 steps are required to compensate for the intermittence overhead. By learning only 50% of the data points, we reduce the total consumption to 1026 steps. The dependency of the result on data is a unique characteristic of machine learning and is absent in general computing algorithms. It introduces a trade-off between the quality of the model and required time/power.

5 CONCLUSION

We propose that the fusion of machine learning and batteryless technology to enable lifelong learning in mobile embedded systems and demonstrate the feasibility of batteryless learning systems.

REFERENCES

[1] Domenico Balsamo, Alex S Weddell, Geoff V Merrett, Bashir M Al-Hashimi, Davide Brunelli, and Luca Benini. 2015. Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems. *IEEE ESL* (2015).
[2] Billur Barshan and Murat Cihan Yüksek. 2014. Recognizing daily and sports activities in two open source machine learning environments using body-worn sensor units. *Comput. J.* 57, 11 (2014), 1649–1667.
[3] Graham Gobieski, Nathan Beckmann, and Brandon Lucia. 2019. Intelligence Beyond the Edge: Inference on Intermittent Embedded Systems. *ASPLOS* (2019).