Locality Sensitive Teaching

Anonymous Author(s)
Affiliation
Address
email

Abstract

The emergence of the Internet-of-Things (IoT) sheds light on applying the machine teaching (MT) algorithms for online personalized education on home devices. This direction becomes more promising during the COVID-19 pandemic when in-person education becomes infeasible. However, as one of the most influential and practical MT paradigms, iterative machine teaching (IMT) is prohibited on IoT devices due to its inefficient and unscalable algorithms. IMT is a paradigm where a teacher feeds examples iteratively and intelligently based on the learner's current status. In each iteration, current IMT algorithms greedily traverse the whole training set to find the sample for the learner, which is computationally expensive in practice. We propose a novel teaching framework, Locality Sensitive Teaching (LST), based on locality sensitive sampling, to overcome these challenges. LST has provable near-constant time complexity, which is exponentially better than the existing baseline. With at most $425.12\times$ speedups and 99.76% energy savings over IMT, LST is the first algorithm that enables energy and time efficient machine teaching on IoT devices. Owing to LST's substantial efficiency and scalability, it is readily applicable in real-world education scenarios.

1 Introduction

During the COVID-19 pandemic, there is an increasing demand for learning at home. Computer-based personalized education (CBPE) [1, 2, 3] on Internet-of-Things (IoT) devices become essential as it improves the accessibility of students to the educational resources and reduces the potential privacy risks. Due to the popularity of Coursera, Duolingo, and EdX, online on-device education becomes increasingly more important nowadays. To better understand and improve the CBPE approaches in this context, we consider machine teaching (MT) [4] as a simplified yet helpful paradigm. Specifically, MT defines the problem where a machine teacher constructs a minimal set of examples that allow a student to learn a target concept repeatedly. However, MT on the IoT device is still prohibitive for two primary reasons: (1) MT technique does not allow real-time interaction between teachers and students. It is designed to work over a static set and cannot incorporate real-time student feedback, (2) MT can not support on-device training due to the expensive costs of constructing this minimal teaching set.

As a result, a more practical paradigm – iterative machine teaching (IMT) [5, 6] has been proposed to achieve state-of-the-art teaching performance. In IMT, the teacher interacts with the student in every iteration and aims to teach the target concept with a minimum number of iterations. Although IMT reduces both teaching set size and teaching iterations, it is still not efficient enough. The time and energy required by IMT are prohibitively expensive on IoT devices. There are two primary reasons: (1) IMT is not designed for real-time teaching. Current IMT uses a deterministic greedy algorithm, which traverses the entire dataset to find the optimal teaching example during each iteration. This linear time scan costs excessive time in real-world datasets, preventing more frequent interactions between machine teachers and students. (2) IMT has poor energy efficiency because the greedy scan

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

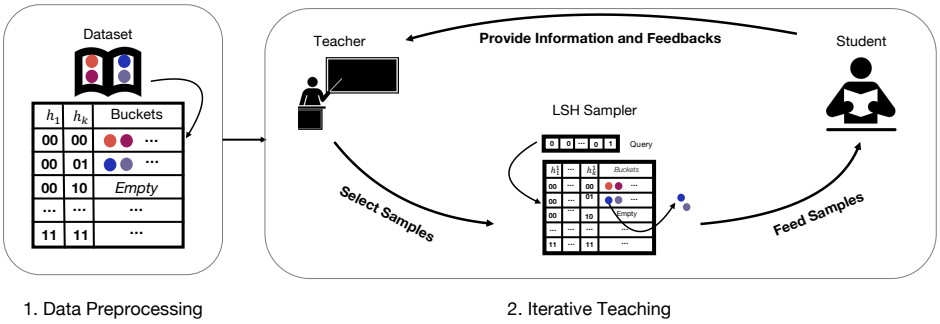

Figure 1: **A brief overview of the proposed LST algorithm. There are two stages: 1) during data preprocessing, the ids of data samples are stored in LSH hash tables. 2) during the teaching stage, in each iteration, the machine teacher uses student's information to query a sample from hash tables and feeds to the student.**

requires scoring and ranking on the training examples. The energy budget of IoT devices such as integrated CPU-GPU System-on-Chips (SoCs) cannot handle this expensive energy consumption in real-world teaching. Clearly, there is a gap between IMT and its deployment on IoT devices.

Contemporary IMT literature focuses on developing and analyzing theories, models, or paradigms in teaching. Few studies are addressing the trade-offs between teachability and efficiency. Although some IMT algorithms manage their potential applications in crowd-sourcing [7, 8, 9], personal education [4], or online data poison [10, 11], they do not consider the edge computing and IoT settings and ignore the substantial potential impact behind it.

Our work proposes a time and energy-efficient IMT algorithm, namely Locality Sensitive Teaching (LST), that enables machine teaching on IoT devices. A real-life teaching observation inspires LST: a human teacher usually teaches a target concept to the student efficiently based on his or her knowledge store in the brain rather than goes over all the exact course materials. Similarly, an efficient IMT algorithm should provide optimal examples by looking up from a data structure rather than scan over the example set. Our LST algorithm effectively pre-indexes the teaching examples in hash tables. Then, given the student's current state at each iteration, we regard it as a query and generate examples through efficient sampling on the hash tables. Moreover, we significantly improve the time and energy efficiency by reducing the expensive linear scan into a lightweight lookup in hash tables. As a result, LST could be compatible with edge computing on IoT devices.

Our proposal, LST, comes with three key questions: (1) how to reformulate the IMT problem as an efficient sampling problem that involves efficient hash table data structures? (2) how to maintain the identical teachability of IMT when utilizing the efficient hash table structures? (3) how to provide an efficient LST implementation on IoT devices that reduce energy and time consumption?

At a high level, we tackle these challenges as follows:

- We reformulate IMT as an adaptive inner product sampling problem. We partition the teaching materials into two parts: (1) query: current state of student and the optimal model possessed by the machine teacher, (2) data: teaching examples denoted as feature-label pairs. We apply an asymmetric transformation that projects both parts as vectors. Then, we argue that the original IMT formula can be formulated as an adaptive inner product sampling. Given the query vector, the task is to sample data vectors with large inner products. Therefore, locality-sensitive sampling can be introduced for efficient teaching examples generation.
- We demonstrate, both theoretically and empirically, that LST preserves the teachability of IMT. Theoretically, we prove the LST can achieve exponential teachability with high probability. Empirically, the experiments on real-world teaching indicate that LST matches or even exceeds the teachability of original IMT.
- We provide a novel LST system design on integrated CPU-GPU SoC platforms that performs time and energy-efficient teaching in real-world settings. For efficient locality sensitive sampling, we re-partition the components as GPU friendly random projection and CPU friendly hash table lookups. Then, we exploit the benefits of fast matrix multiplications on GPU and efficient hash

table lookups implemented on CPUs. Therefore, our LST system takes full advantage of the memory-constrained CPU-GPU SoCs.

Through extensive experiments in real-world teaching scenarios, we demonstrate that LST performs exponential teachability that matches or exceeds IMT while achieving at most 425.12× speedups and 99.76% energy savings on IoT devices. On server-based evaluation, LST achieves more than 2000× speedups over IMT.

2 Related Work and Preliminaries

2.1 Iterative Machine Teaching

Machine teaching defines a task where we need to find a minimal training set given a learner and a target concept. [4, 12] describes a general teaching framework that has nice connections to curriculum learning [13]. [14] considers Bayesian learners in exponential family and formulates the machine teaching as an optimization problem over teaching examples that balance the future loss of the learner and the teacher's effort. [15] gives the teaching dimension of linear learners. Machine teaching has been found useful in cybersecurity [16], human-computer interaction [17], and human education [18]. [19, 6, 20] study the teaching scenario where the teacher aims to teach multiple concepts to a forgetful learner. [21, 6] assume that the learner is black-box and study how the teaching should be performed. [22, 6] discuss how the teacher can teach multiple different learners. [23, 24, 25] connect the machine teaching problem with inverse reinforcement learning by studying an inverse reinforcement learner. [26, 2] considers the machine teaching from an interpretable perspective. [27, 28] discuss the teaching model with a version space learner. [3] extends machine teaching to human-in-the-loop settings. [29, 30, 31] study the machine teaching in theory.

Previous machine teaching works focus on the improvement of teachability. [5] proposes the iterative teaching paradigm and an omniscient teaching model where the teacher knows almost everything about the learner and provides training examples based on the learner's status. [32] further generalizes this teaching scenario and introduces an optimal control approach to address sequential machine teaching. Built upon the same framework as [5], [6] looks into the IMT problem when the learner becomes black-box. However, they all ignore the fact that a student model is typically optimized by an iterative algorithm (*e.g.*, SGD) in practice. In our work, we focus more on how fast a student can learn from the teacher.

This paper follows the same teaching setting as [5]. Without loss of generality, we mainly discuss the omniscient teaching model where the teacher knows the learner's full information. We apply our LST to speed up the omniscient teacher. We note that it is straightforward to improve the teaching quality and efficiency of more advanced teaching models using our proposed method.

2.2 Hash Based Sampling

This section briefly describes the recent development of using locality sensitive hashing [33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44] for sampling and estimation. The high-level idea of LSH is to quantize similar items into the same bucket of a hash table via random projection [45]. For details of LSH, we refer to Appendix A.

Denote p as the collision probability of LSH function. Therefore, an item that is retrieved by a (K, L) LSH algorithm with probability p. The precise form of p is determined by the LSH family. This sampling view of LSH has been used in a wide range of applications, such as neural networks training [46, 47], kernel density estimation [48], outlier detection [49], and optimization [50].

Recent advances in fast inner product search using LSH has made it possible to sample large inner products. More specifically, in this paper, given a collection $\mathcal{C}$ of input and a query Q, candidates S are drawn from $\mathcal{C}$, in which each $x_i \in \mathcal{C}$ is sampled with $p_i \leq 1$, using LSH algorithm with special designed hashing [51]. Note here p_i is a monotonically increasing function of $Q \cdot x_i$. Therefore, once a one-time linear cost of preprocessing $\mathcal{C}$ into hash tables is paid, any further adaptive sampling for query Q only requires a few hash lookups. The LSH sampler component in Figure 1 demonstrates the sampling process using LSH. Similar to what has been mentioned in [50], the sampling scheme is not a valid distribution, i.e., $\sum_{x_i \in \mathcal{C}} p_i \neq 1$. Given a query, the probability of sampling x_i is not independent of the probability of sampling x_j $(i \neq j)$. However, we can still use it for unbiased

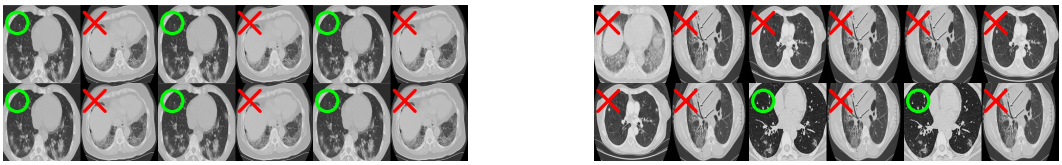

(a) IMT (b) **Proposed** LST

Figure 2: **A motivating example on COVID-19 Teaching via CT images. The first 12 examples selected by IMT and LST examples are presented on (a) and (b). We mark each COVID image by green circle and each regular image by red cross. We observe that our proposed LST selects diverse, confusing images while IMT keeps selecting the same 2 images.Details are shown in Appendix D.**

estimation. p_i is a monotonic function of $Q \cdot x_i$ because $p_i = (1 - (1 - g(Q \cdot x_i))^K)^L$, where $g(Q \cdot x_i)$ is the collision probability.

3 Locality Sensitive Teaching

This section first shows that finding a teaching example in the current IMT can be reformulated as an adaptive inner product sampling (AIPS) problem. Then we propose our LST algorithm to effectively and efficiently perform this adaptive sampling via Locality Sensitive Sampling (LSS).

3.1 A Motivating Example

To showcase the significance of LST, we conduct an experiment on teaching student to recognize COVID-19 from CT images. We apply both LST and IMT on a binary-class CT image dataset [52] (labeled as normal or COVID-19, details are in Appendix D). After these 12 examples, LST achieves less loss than IMT. Fig. 2(a) and (b) shows the comparison of teaching examples. One can observe that LST's teaching examples yield more diversity because of the randomness it introduces. In contrast, IMT simply repeats two examples without giving a broader view of the other teaching set examples. This phenomenon suggests LST's potential to achieve more intuitively plausible teaching.

3.2 Reformulating IMT as Adaptive Sampling

Following similar notation from [5], we denote N as the number of samples, and $S = \{(x_i, y_i)|0 < i \leq N\}$ as the input samples along with their labels . Then, we construct a student model with loss function ℓ and learning rate η. Given the model, we denote w_t as the parameter of the student model at iteration t. w^* is the optimal parameter obtained by the omniscient teacher. In each iteration, IMT selects an example by performing the following greedy algorithm:

Problem 1 (**Greedy Teaching Algorithm (GTA)).** *Given a large set S, the optimal parameter w^* and model parameter w_t, we want to solve the following optimization problem:*

$$x^t, y^t = \arg\min_{(x,y)\in S} \eta^2 \|\nabla l\|_2^2 - 2\eta\Big\langle w^t - w^*, \nabla l\Big\rangle, \tag{1}$$

where we have $\nabla l = \frac{\partial \ell(\langle w^t, x\rangle, y)}{\partial w_t}$.

This deterministic greedy algorithm is usually sub-optimal. For illustration, we present a toy example. We random sample 4 points from the synthetic dataset used by [5] and Section E. The data visualization is shown in Figure 3(a). Given the w^* generated in [5], we perform logistic regression using IMT with GTA for 5 steps. We also brutal force all possible choices in each iteration to find the optimal path in the first 5 steps. Train loss versus iterations is plotted in Figure 3(b). All methods are rerun for 10 times and take their average. We observe that given the optimal path "AAAAA", GTA always chooses "CCCCC".

To tackle the issue, we reformulate problem 1 to problem 2 via transformation functions f, g. We take linear regression and logistic regression classification as examples and derive corresponding f, g (shown in Appendix B). For other losses, we can derive f, g similarly when they are locally linearized. AIPS introduces randomness to help escape from local optima so that AIPS performs

close to the optimal path on average. Here the parameter c is determined by the LSH. In toy example, AIPS chooses among "AAAAA", "AAAAC" and "AACAA", which are near-optimal path. Therefore, the loss of AIPS in each iteration outperforms GTA in Figure 3(b).

Problem 2 (**Adaptive Inner Product Sampling (AIPS)**). *Given a large set S and a query point $q = g(w^*, w^t)$, and $0 < c < 1$, we aim to sample an example $(x, y) \subset S$ that for any $(x', y') \in S$, $f(x, y)^\top g(w^*, w^t)$ satisfies:*

$$f(x, y)^\top g(w^*, w^t) \geq c \left(f(x', y')^\top g(w^*, w^t) \right). \tag{2}$$

where f and g are transformation functions with the same dimension.

3.3 Locality Sensitive Sampling as the Key Ingredient

Our algorithm proposes to use LSS as an efficient sampler for problem 2. Figure 1 shows the complete workflow of the proposed LST. Note that, w_t changes in every iteration, and for every iteration, $\mathcal{O}(N)$ computations is typically required if we want to solve problem 1. This time complexity is infeasible for large-scale teaching problems. Therefore, we propose an efficient algorithm to perform IMT by sampling the transformed data vectors that produce large inner products $f(x, y)^\top g(w^*, w^t)$ at the cost of $\mathcal{O}(1)$.

Our algorithm guarantees to find an example for the student model that produces the maximum inner product of $f(x, y)^\top g(w^*, w^t)$ with high probability. The LSS process is designed to sample from a weighted distribution which is a monotonic function of the resulting inner product. Before we describe the efficient sampling process's details, we first argue that such adaptive sampling should perform much better than random sampling. It is because the probability of selecting the correct sample is $\frac{1}{N}$ for random sampling. On the other hand, for adaptive sampling, this probability is much larger than $\frac{1}{N}$, because by definition the probability is monotonic to the value of the inner product.

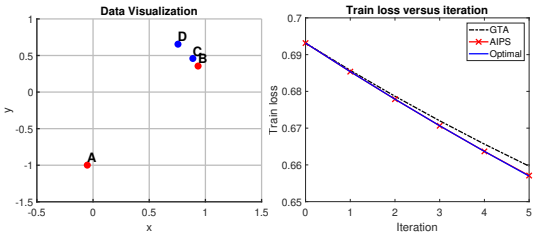

(a) (b)

Figure 3: **A toy example. (a) visualization of data from 2 class, red means 0 and blue is 1. (b) Performance after 5 iteration of example path chosen by (1) GTA, (2) AIPS, (3) Optimal.**

Moreover, for any monotonic function, the sampling probability distribution changes according to the updates on w. Due to monotonicity, the probability of choosing the sample which produces the maximum inner product in each iteration is always the highest. The key insight here is that there are two qualities in the inner product, $f(x, y)$ and $g(w^*, w^t)$. With successive iteration, $g(w^*, w^t)$ will change while $f(x, y)$ is always fixed. Therefore, we only need to perform one time pre-processing to obtain $f(x, y)$ from all x, y pairs in S and put $f(x, y)$ into the LSH tables. Then we can use $g(w^*, w^t)$ as the query for efficient LSS. Then in the next iteration, only the query changes to $g(w^*, w^{t+1})$, but the hash tables stay unchanged. Therefore, a few hash table lookups are usually sufficient to perform the sampling. The guarantee of the query time is theoretically analyzed in section 4. The detailed sampling process is presented in Algorithm 2, and the overall LST algorithm is put in Algorithm 1. In conclusion, after we pay a one-time pre-processing effort of building the hash tables, we just need a few lookups (typically just one) in each iteration to get a high-quality teaching example, which is far better than a random example and may outperform GTA example.

There are few technical subtleties due to the square of the gradient term. We present it in equation 1. In Appendix B, we derive the full terms of $f(x, y)$ and $g(w^*, w^t)$ which is similar to a quadratic kernel. Both f and g are the corresponding feature expansion transformation. To increase the efficiency of LST, we apply asymmetric transforms (details also in Appendix B) to deal with the quadratic terms and demonstrate its validated performance in practice (see experiments in section 6).

Algorithm 1: Locality Sensitive Teaching (LST)

Result: Model w
Input: $D = \{x,\ y\}, w^*, w^t, \eta$;
for i *in* π **do**
  $h_1 ... h_L \leftarrow$
    $H_1(f(x_i, y_i)) ... H_L(f(x_i, y_i))$;
  insert i (id) in L hash tables;
end
$l \leftarrow 0$;
$\pi \leftarrow permute(1, L)$;
while ***not converged*** **do**
  $j \leftarrow$ LSS($f(w^*, w^t)$) (Algorithm 2) ;
  $w \leftarrow w - \eta(\nabla L(x_j, y_j))$
end

Algorithm 2: Locality Sensitive Sampling (LSS)

Result: Sample id, **Sampling Probability** p
Input: Query q;
$l \leftarrow 0$;
$\pi \leftarrow permute(1, L)$;
for i *in* π **do**
  compute $H(q)$ for hash table i;
  if *Bucket B for $H(q)$ is not Empty* **then**
    $|B| \leftarrow$ Size of B;
    $id \leftarrow$ random(examples in B);
  else
    $l \leftarrow l + 1$;
  end
end

4 Theoretical Insights and Discussions

As shown in section 3.2, we reformulate IMT as an adaptive inner product sampling (AIPS) and provide LST. Here we provide comparison between LST and the current IMT. As shown in Appendix B, GTA used in IMT can be regarded as $\arg\max_i g(w^*, w)^\top f(x_i, y_i)$, which is an exact maximum inner product search (MIPS) problem. [5, 6] only show improvements over SGD, suggesting that GTA is a local optimal solution.

Our AIPS algorithm introduces randomness to GTA. If both $\|f(x_i, y_i)\|$ and $\|g(w^*, w)\|$ are assumed to be the same constant, then from [51], we see GTA and nearest neighbor search (NNS) problem are equivalent, formally,

$$\arg\max_i g(w^*, w)^\top f(x_i, y_i) \approx \arg\min_i \|g(w^*, w) - f(x_i, y_i)\|$$

Instead of doing exact MIPS, AIPS random samples a $c - approximate$ nearest neighbor defined above. The efficiency of AIPS is directly quantified by $\rho = \frac{\log p_1}{\log p_2} < 1$, where p_1, p_2 are defined in Section 2.2. The space complexity grows as $O(n^{1+\rho})$, while the query time grows as $O(n^\rho \log n)$, where n is the size of the dataset. Thus, smaller ρ indicates better theoretical performance.

Theorem 1. *Given a family of (S_0, cS_0, p_1, p_2)-sensitive hash function $\mathcal{H}$, one can construct a data structure for AIPS with $O(n^\rho \log n)$ query time and space $O(n^{1+\rho})$, where $\rho = \frac{\log p_1}{\log p_2}$.*

Theorem 1 shows the time and space complexity for AIPS. In section F, we provide an ablation study of efficiency and accuracy trade-off of our algorithm. Furthermore, asymmetric transformations can be used in accelerating AIPS on samples with large inner products. A similar trade-off with proofs is shown in [51].

Theorem 2 (**Exponential Teachability of LST**). *Following the theoretical setting in the synthesis-based teaching [5], the teacher is assumed to be able to provide any samples from*

$$\mathcal{X} = \{x \in \mathbb{R}^d, \|x\| \leq R\},$$
$$\mathcal{Y} = \mathbb{R} \text{ (Regression) or } \{-1, 1\} \text{ (Classification).}$$

Assume we have a student with fixed learning rate $\eta \neq 0$, if the loss function $\ell(\cdot, \cdot)$ satisfies that for any $w \in \mathbb{R}^d$, there exists $\gamma \neq 0, \|\gamma\| \leq \frac{R}{\|w - w^*\|}$ such that while $\hat{x} = \gamma(w - w^*)$ and $\hat{y} \in \mathcal{Y}$, we have that

$$0 < \gamma \nabla_{\langle w, \hat{x} \rangle} \ell(\langle w, \hat{x} \rangle, \hat{y}) \leq \frac{1}{\eta}. \tag{3}$$

Suppose for a (S_0, cS_0, p_1, p_2)-sensitive hashing function family $\mathcal{H}$ where $S_0 = \max_i f(\hat{x}, \hat{y})^\top g(w^*, w^i)$ and

$$\hat{x} = \gamma \left(w^t - w^*\right) \text{ and } \hat{y} \in \mathcal{Y}, \tag{4}$$

then LST can achieve exponential teachability with probability at least p_1 in a single iteration.

Theorem 2 shows the condition and probability of LST to achieve exponential teachability in each iteration. There is a probability p_1 to perfectly imitate the omniscient teacher. More importantly, it

introduces some randomness into the teaching algorithm and therefore makes the teaching examples more diverse and plausible for humans. Note that, all Lipschitz smooth and strongly convex loss functions can satisfy Eq. 3. Using Theorem 2, we can easily extend LST's exponential teachability to the other teaching scenarios such as combination-based teaching [5], (rescaled) pool-based teaching [5] and black-box teaching [6]. Detailed proofs are shown in Appendix C.

Time Complexity Analysis: Theorem 1 and 2 provide the analysis for the LSS algorithm to achieve the exponential teachability with provably sub-linear time complexity. In practice, we need to pay one-time pre-processing time $\mathcal{O}(NKL)$ to build the hash tables, where N is the number of samples, K is the number of hash functions in each hash table, and L is the number of hash tables. During inference, we usually only use a few lookups (typically just one), *i.e.*, $\mathcal{O}(1)$ complexity, in each iteration to get a high-quality teaching example. We also conduct an empirical experiment to show the dedicated trade-off between the accuracy and efficiency using different K and L. Therefore, LST reduces the time complexity for each teaching iteration from $\mathcal{O}(N)$ to $\mathcal{O}(1)$ by hash table lookups.

5 Locality Sensitive Teaching on Integrated CPU-GPU SoC

In this section, we present the LST system implemented on integrated CPU-GPU SoC platforms. In recent machine learning based IoT applications such as autonomous driving, smart home, and intelligent robots, devices with integrated CPU-GPU architecture (e.g., NVIDIA's Jetson [53] series) is gaining increasing preference. Therefore, the smart co-optimization[54, 55] on CPU and GPU resources determine the overall performance of machine learning algorithms on IoT deployment. In this work, our system design is based on two observations: (1) The CPU resources could be more utilized on the integrated CPU-GPU SoC. For instance, a dual-core NVIDIA Denver 2 64-bit CPU and a quad-core ARM Cortex-A57 are equipped on NVIDIA TX2 [56] SoC. Vision-based machine learning applications such as object detection [54] do not regard on-chip CPU as a computation resource. In our problem, if we can share some of the computation to those powerful CPUs, the overall efficiency can be improved. (2) Hashing is cheap on CPU. Most GPUs suffer memory constraints. Although GPU based LSH hash tables are 1.5x faster over CPU [57] in nearest neighbor search, it stores the hash table entirely on GPU memory. Therefore, the availability of GPU for other machine learning operations is limited. To tackle the issue, hybrid CPU-GPU LSH systems [57, 58] become powerful in both academics and industry.

In our LST systems on integrated CPU-GPU SoC, we implement LSH by separating the random projection and hash table lookups into GPU and CPU. We first generate hash codes of data vectors by GPU based random matrix multiplication via CuPy [59] and compiled CUDA kernels. Then, we insert the hash codes to hash tables built on the CPU. The query vector is transformed into hash codes in the query phase using the same random matrix in GPU. Then, the query hash codes are sent to the CPU and perform hash table lookups. We provide Cython wrapping for the implementation to make it PyTorch friendly. We believe this implementation would be more beneficial for the community as it can be easily plugged into any deep learning model.

6 Experiments

In this section, we demonstrate the effectiveness and efficiency of LST in various scenarios. We first perform the algorithmic level evaluation of LST on different datasets. Specifically, we target to answer the following questions: (1) does the LST performs better or the same as the current IMT in teaching the learner model towards faster convergence in iteration wise? (2) compared to the IMT, does the LST algorithm reduce the computation time when achieving the same performance? (3) is the proposed LST algorithm scalable in the dataset that the IMT is almost infeasible in practice? Next, we conduct experiments to demonstrate the scalability of LST on IoT devices. We would like to answer the following questions: (1) does the LST's teachability robust to IoT devices? (2) compared to IMT, is the LST energy efficient on IoT devices? Note that we present the results on regression in this section. For classification results, please refer to Appendix E.

Datasets. In the experiments, we use four regression datasets to demonstrate the performance of our LST. First, we use abalone, space_ga and mg dataset from LIBSVM [60] and UCI dataset [61]. The abalone dataset contains 4,177 8 dimensional feature vectors. The space_ga and mg datasets contain 3107 and 1385 7 dimensional vectors. Then we test LST's performance on large scale regression datasets to validate its scalability. We use slice dataset from UCI dataset [61]. Slice dataset contains 53500 training samples and 42800 testing samples. Each sample is a 74 dimension vector. We

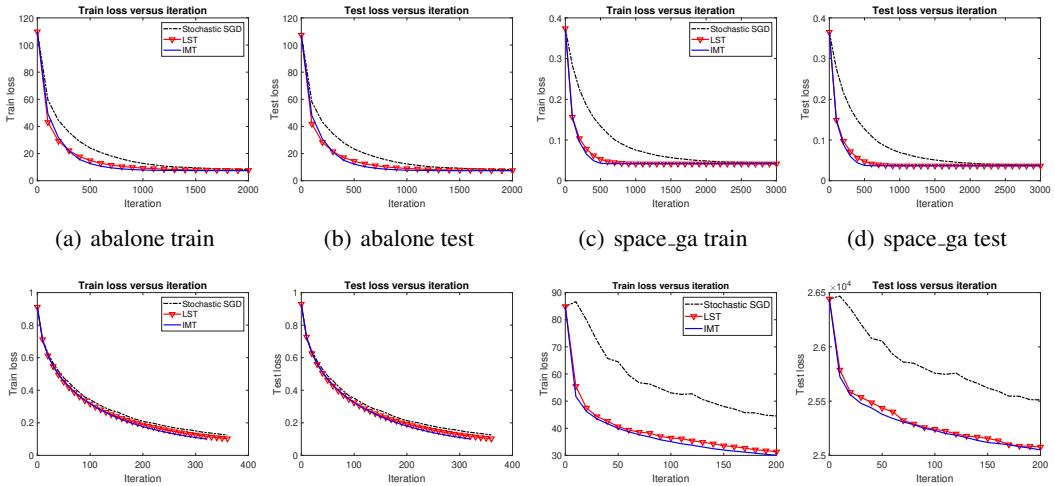

(a) abalone train (b) abalone test (c) space_ga train (d) space_ga test

(e) mg train (f) mg test (g) slice train (h) slice test

Figure 4: **Average train and test loss versus iteration of: abalone in (a) and (b), space_ga in (c) and (d), mg in (e) and (f), slice in (g) and (h). Results have demonstrated that both LST and the current IMT is superior to stochastic SGD in guiding the learner model towards faster convergence. Meanwhile, LST has similar performance with IMT that their train and test descents in almost the same way.**

use Slice only for algorithm level evaluation as it causes memory exhaustion on IoT devices. We randomly split 30% of samples in each dataset as a test set while others are training set. All datasets are under MIT license.

6.1 Evaluation on Locality Sensitive Teaching

In this section, we present the performance of our proposed LST algorithm in teaching regression models towards fast convergence. The evaluation is on a server with 1 Nvidia Tesla V100 GPU and two 20-core/40-thread processors (Intel Xeon(R) E5-2698 v4 2.20GHz). Given the four regression datasets, we set the learner model as a linear regression model with Mean Square Error (MSE) loss to compare LST's performance, the IMT, and stochastic SGD. Here, we consider the stochastic SGD as a random teaching approach. In each step, the three algorithms above choose a sample from the dataset and feed it into the learner. The learner model then processes the sample and performs AdaGrad[62] to update the weights accordingly. We use early stopping on the learner model to prevent overfitting. Details of LST for classification are in Appendix E.

Convergence. In Figure 4, we plot the train and test loss versus iteration of three teaching approaches mentioned above on the 4 regression dataset. From the (a), (b), (c) and (d), we observe that the LST and IMT [5] algorithm teaches the learner models that converge to optimum with much less iteration than stochastic SGD. This performance shows that both LST and the IMT leads to faster iteration wise convergence of learner models in train and test set. Moreover, in each step, the learners' training and test losses taught by LST are close or even better than IMT. In sub-figure (c) and (d), we show that LST also performs a good alteration to IMT in teaching towards faster convergence when the gap of train losses and test losses between IMT and stochastic SGD are small. This result demonstrates that LST's approximation error to the optimal path is affordable even in some settings that IMT only has few advantages towards stochastic SGD. In sub-figure (e) and (f), we demonstrate that LST and IMT share similar performance in large scale regression sets on average. Moreover, we observe that about 30% LST instances outperform IMT. The effectiveness experiment answers the first question. LST can achieve the same fast convergence effect in iteration wise when compared to IMT. With non-negligible probability, LST's teachability outperforms IMT in teaching regression towards convergence with fewer iterations.

Speed. In Figure 5, we plot the train loss versus train time (seconds in log scale) for LST and IMT in four regression sets. In each iteration, we record the time for teaching algorithms to select a sample and the weight update time for the learners given this sample. In the four sub-figures, we observe a consistent gap between two approaches in time. At the same loss level during training, our LST teaches 2000× faster than IMT on average. For a large scale slice dataset, it takes 404.95 seconds

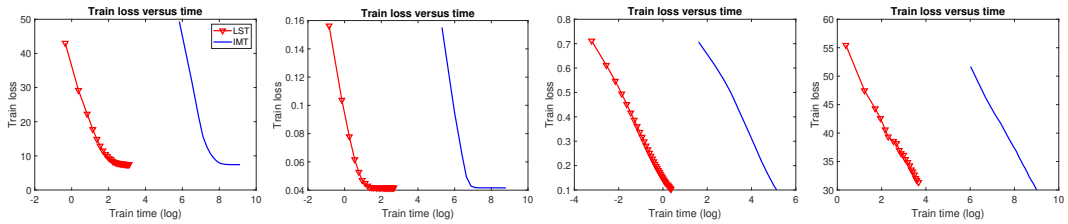

(a) abalone (b) space_ga (c) mg (d) slice

Figure 5: **Average train loss versus the train time (seconds in log scale) of abalone in (a), space_ga in (b), mg in (c), slice in (d). The time unit is second. Results have shown that LST is 2000x times faster than the current IMT to achieve the same loss at each step.**

on average for IMT to finish one iteration. This long-time prevents IMT from online deployment at scale. Our LST tackles the issue by reducing the time to 0.196 seconds while preserving the same performance in iteration wise. This significant acceleration makes LST scalable in large scale teaching scenarios. This experiment answers the second and third questions. LST is 2000$\times$ faster than IMT while achieving the same iteration wise fast convergence. Moreover, LST is efficient in teaching via large scale datasets. We also perform a parameter study in Appendix F.

6.2 Scalability on IoT Devices

In this section, we compare LST and IMT on Nvidia TX2 devices. For detailed setting, we refer to Appendix G. We first observe that the iteration wise performance remains identical to the evaluation on the server. This phenomenon answers the first question: LST's teachability is robust in IoT devices by achieving exactly the same iteration wise convergence as server-based experiments. Then, in Table 1, we present the energy savings and speedups of LST over IMT on four datasets. According to the results, we observe that LST achieves at least 98.73% energy savings and 117.20$\times$ speedups over IMT. In the synthetic classification dataset, LST is 306.12$\times$ faster than IMT with 99.67% energy savings. This performance gain is fewer than server-based experiments due to the less available resources on TX2 SoCs (e.g., 105 w power budget CPU in sever setting vs. 7.5w power budget in IoT setting). However, this improvement is still promising and answer the second question: LST can gain at most 99.76% energy savings with over 100$\times$ speedups compared with IMT on IoT devices.

Table 1: **Efficiency of LST over IMT on Nvidia TX2**

Dataset	Energy Savings	Speedups
abalone	99.76%	425.12$\times$
space_ga	99.34%	149.07$\times$
mg	98.73%	117.20$\times$

7 Negative Societal Impact

This paper discusses the energy reduction of current machine teaching algorithms. Thus, this paper does not contain the negative societal impact in the environment. Moreover, our paper improves personalized education privacy as we keep learners' information on their devices.

8 Concluding Remarks

The inefficiency of iterative machine teaching (IMT) prevents it from being used in IoT based teaching scenarios at home. There is no sub-linear time iterative teaching algorithm that enables teaching in large-scale real-world settings before our work. The paper proposes a practical and fast teaching approach. By reformulating IMT's optimization problem to an adaptive sampling problem, we propose a Locality Sensitive Teaching (LST) algorithm. LST achieves exponential teachability and reduces the time complexity of each iteration from $\mathcal{O}(N)$, where N is the number of teaching examples, to few $\mathcal{O}(1)$ lookups in hash tables. Moreover, we design an LST system that fully exploits the resources of integrated CPU-GPU SoCs on IoT devices. We demonstrate, both theoretically and empirically, LST's teachability improvements, speed acceleration, and tremendous energy reductions over IMT during experiments on IoT devices. The development of LST makes teaching algorithms feasible on more on-device personalized education applications..

References

[1] Christos Athanasiadis, Enrique Hortal, Dimitrios Koutsoukos, Carmen Zarco Lens, and Stylianos Asteriadis. Personalized, affect and performance-driven computer-based learning. In *CSEDU (1)*, pages 132–139, 2017.

[2] Oisin Mac Aodha, Shihan Su, Yuxin Chen, Pietro Perona, and Yisong Yue. Teaching categories to human learners with visual explanations. In *CVPR*, 2018.

[3] Edward Johns, Oisin Mac Aodha, and Gabriel J Brostow. Becoming the expert-interactive multi-class machine teaching. In *CVPR*, 2015.

[4] Xiaojin Zhu. Machine teaching: An inverse problem to machine learning and an approach toward optimal education. In *AAAI*, 2015.

[5] Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda B. Smith, James M. Rehg, and Le Song. Iterative machine teaching. In *ICML*, 2017.

[6] Weiyang Liu, Bo Dai, Xingguo Li, Zhen Liu, James M Rehg, and Le Song. Towards black-box iterative machine teaching. *arXiv preprint arXiv:1710.07742*, 2017.

[7] Yao Zhou, Fenglong Ma, Jing Gao, and Jingrui He. Optimizing the wisdom of the crowd: Inference, learning, and teaching. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3231–3232, 2019.

[8] Yikun Ban and Jingrui He. Generic outlier detection in multi-armed bandit. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 913–923, 2020.

[9] Yao Zhou, Arun Reddy Nelakurthi, Ross Maciejewski, Wei Fan, and Jingrui He. Crowd teaching with imperfect labels. In *Proceedings of The Web Conference 2020*, pages 110–121, 2020.

[10] Scott Alfeld, Xiaojin Zhu, and Paul Barford. Data poisoning attacks against autoregressive models. In *AAAI*, 2016.

[11] Xuezhou Zhang, Xiaojin Zhu, and Laurent Lessard. Online data poisoning attacks. In *Learning for Dynamics and Control*, pages 201–210, 2020.

[12] Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N Rafferty. An overview of machine teaching. *arXiv preprint arXiv:1801.05927*, 2018.

[13] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, 2009.

[14] Xiaojin Zhu. Machine teaching for bayesian learners in the exponential family. In *NIPS*, 2013.

[15] Ji Liu, Xiaojin Zhu, and H Gorune Ohannessian. The teaching dimension of linear learners. In *ICML*, 2016.

[16] Shike Mei and Xiaojin Zhu. Using machine teaching to identify optimal training-set attacks on machine learners. In *AAAI*, 2015.

[17] Christopher Meek, Patrice Simard, and Xiaojin Zhu. Analysis of a design pattern for teaching with features and labels. *arXiv preprint arXiv:1611.05950*, 2016.

[18] Faisal Khan, Bilge Mutlu, and Xiaojin Zhu. How do humans teach: On curriculum learning and teaching dimension. In *NIPS*, 2011.

[19] Anette Hunziker, Yuxin Chen, Oisin Mac Aodha, Manuel Gomez Rodriguez, Andreas Krause, Pietro Perona, Yisong Yue, and Adish Singla. Teaching multiple concepts to a forgetful learner. In *NeurIPS*, 2019.

[20] Yao Zhou, Arun Reddy Nelakurthi, and Jingrui He. Unlearn what you have learned: Adaptive crowd teaching with exponentially decayed memory learners. In *KDD*, 2018.

[21] Sanjoy Dasgupta, Daniel Hsu, Stefanos Poulis, and Xiaojin Zhu. Teaching a black-box learner. In *ICML*, 2019.

[22] Teresa Yeo, Parameswaran Kamalaruban, Adish Singla, Arpit Merchant, Thibault Asselborn, Louis Faucon, Pierre Dillenbourg, and Volkan Cevher. Iterative classroom teaching. In *AAAI*, 2019.

[23] Luis Haug, Sebastian Tschiatschek, and Adish Singla. Teaching inverse reinforcement learners via features and demonstrations. In *NeurIPS*, 2018.

[24] Sebastian Tschiatschek, Ahana Ghosh, Luis Haug, Rati Devidze, and Adish Singla. Learner-aware teaching: Inverse reinforcement learning with preferences and constraints. *arXiv preprint arXiv:1906.00429*, 2019.

[25] Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, and Adish Singla. Interactive teaching algorithms for inverse reinforcement learning. In *IJCAI*, 2019.

[26] Yuxin Chen, Oisin Mac Aodha, Shihan Su, Pietro Perona, and Yisong Yue. Near-optimal machine teaching via explanatory teaching sets. In *AISTATS*, 2018.

[27] Yuxin Chen, Adish Singla, Oisin Mac Aodha, Pietro Perona, and Yisong Yue. Understanding the role of adaptivity in machine teaching: The case of version space learners. In *NeurIPS*, 2018.

[28] Mohamadreza Ahmadi, Bo Wu, Yuxin Chen, Yisong Yue, and Ufuk Topcu. Barrier certificates for assured machine teaching. In *ACC*, 2019.

[29] Ziyuan Gao, Hans Ulrich Simon, and Sandra Zilles. On the teaching complexity of linear sets. In *ALT*, 2015.

[30] Sandra Zilles, Steffen Lange, Robert Holte, and Martin Zinkevich. Teaching dimensions based on cooperative learning. In *COLT*, 2008.

[31] Rahim Samei, Pavel Semukhin, Boting Yang, and Sandra Zilles. Algebraic methods proving sauer's bound for teaching complexity. *Theoretical Computer Science*, 558:35–50, 2014.

[32] Laurent Lessard, Xuezhou Zhang, and Xiaojin Zhu. An optimal control approach to sequential machine teaching. *arXiv preprint arXiv:1810.06175*, 2018.

[33] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *STOC*, 1998.

[34] Piotr Indyk and David Woodruff. Polylogarithmic private approximations and efficient matching. In *Theory of Cryptography*. 2006.

[35] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry (SoCG)*, pages 253–262, 2004.

[36] Alexandr Andoni, Piotr Indyk, Huy L Nguyen, and Ilya Razenshteyn. Beyond locality-sensitive hashing. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1018–1028. SIAM, 2014.

[37] Alexandr Andoni, Piotr Indyk, TMM Laarhoven, Ilya Razenshteyn, and Ludwig Schmidt. Practical and optimal lsh for angular distance. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1225–1233. Curran Associates, 2015.

[38] Alexandr Andoni and Ilya Razenshteyn. Optimal data-dependent hashing for approximate near neighbors. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing (STOC)*, pages 793–801, 2015.

[39] Piotr Indyk and Tal Wagner. Approximate nearest neighbors in limited space. In *Conference On Learning Theory*, pages 2012–2036. PMLR, 2018.

[40] Alexandr Andoni, Thijs Laarhoven, Ilya Razenshteyn, and Erik Waingarten. Optimal hashing-based time-space trade-offs for approximate near neighbors. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 47–66. SIAM, 2017.

[41] Alexandr Andoni, Piotr Indyk, and Ilya Razenshteyn. Approximate nearest neighbor search in high dimensions. *arXiv preprint arXiv:1806.09823*, 7, 2018.

[42] Yihe Dong, Piotr Indyk, Ilya Razenshteyn, and Tal Wagner. Learning space partitions for nearest neighbor search. In *International Conference on Learning Representations*, 2019.

[43] Hao Chen, Ilaria Chillotti, Yihe Dong, Oxana Poburinnaya, Ilya Razenshteyn, and M Sadegh Riazi. {SANNS}: Scaling up secure approximate k-nearest neighbors search. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 2111–2128, 2020.

[44] Xiaoyun Li and Ping Li. Rejection sampling for weighted jaccard similarity revisited. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021.

[45] Xiaoyun Li and Ping Li. Random projections with asymmetric quantization. *Advances in Neural Information Processing Systems*, 32:10858–10867, 2019.

[46] Ryan Spring and Anshumali Shrivastava. Scalable and sustainable deep learning via randomized hashing. *arXiv preprint arXiv:1602.08194*, 2016.

[47] Beidi Chen, Tharun Medini, and Anshumali Shrivastava. Slide: In defense of smart algorithms over hardware acceleration for large-scale deep learning systems. *arXiv preprint arXiv:1903.03129*, 2019.

[48] Benjamin Coleman and Anshumali Shrivastava. Sub-linear race sketches for approximate kernel density estimation on streaming data. In *Proceedings of The Web Conference 2020*, pages 1739–1749, 2020.

[49] Chen Luo and Anshumali Shrivastava. Arrays of (locality-sensitive) count estimators (ace) anomaly detection on the edge. In *Proceedings of the 2018 World Wide Web Conference*, pages 1439–1448, 2018.

[50] Beidi Chen, Yingchen Xu, and Anshumali Shrivastava. Lsh-sampling breaks the computational chicken-and-egg loop in adaptive stochastic gradient estimation. 2018.

[51] Anshumali Shrivastava and Ping Li. Asymmetric lsh (alsh) for sublinear time maximum inner product search (mips). In *NIPS*, 2014.

[52] Jinyu Zhao, Yichen Zhang, Xuehai He, and Pengtao Xie. Covid-ct-dataset: a ct scan dataset about covid-19. *arXiv preprint arXiv:2003.13865*, 2020.

[53] Nathan Otterness, Ming Yang, Sarah Rust, Eunbyung Park, James H Anderson, F Donelson Smith, Alex Berg, and Shige Wang. An evaluation of the nvidia tx1 for supporting real-time computer-vision workloads. In *2017 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 353–364. IEEE, 2017.

[54] Waqar Ali and Heechul Yun. Protecting real-time gpu applications on integrated cpu-gpu soc platforms. *ArXiv e-prints*, 2017.

[55] Soroush Bateni, Zhendong Wang, Yuankun Zhu, Yang Hu, and Cong Liu. Co-optimizing performance and memory footprint via integrated cpu/gpu memory management, an implementation on autonomous driving platform. In *2020 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 310–323. IEEE, 2020.

[56] NVIDIA Inc. NVIDIA Jetson TX2. https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-tx2/, accessed 2020-09-01.

[57] Yiqiu Wang, Anshumali Shrivastava, Jonathan Wang, and Junghee Ryu. Flash: Randomized algorithms accelerated over cpu-gpu for ultra-high dimensional similarity search. *arXiv preprint arXiv:1709.01190*, 2017.

[58] Nicholas Meisburger and Anshumali Shrivastava. Tera-slash: A distributed energy-efficient mpi based lsh system for tera-scale similarity search. *arXiv preprint arXiv:2008.03260*, 2020.

[59] Ryosuke Okuta, Yuya Unno, Daisuke Nishino, Shohei Hido, and Crissman Loomis. Cupy: A numpy-compatible library for nvidia gpu calculations. In *NIPS LearningSys Workshop*, 2017.

[60] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.

[61] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.

[62] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.

[63] Ryan Spring and Anshumali Shrivastava. A new unbiased and efficient class of lsh-based samplers and estimators for partition function computation in log-linear models. *arXiv preprint arXiv:1703.05160*, 2017.

[64] Moses S Charikar. Similarity estimation techniques from rounding algorithms. In *STOC*, 2002.

[65] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

[66] Jan-Mark Geusebroek, Gertjan J Burghouts, and Arnold WM Smeulders. The amsterdam library of object images. *IJCV*, 2005.

[67] Siddharth Gopal. Adaptive sampling for sgd by exploiting side information. In *ICML*, 2016.

[68] Chaojian Li, Tianlong Chen, Haoran You, Zhangyang Wang, and Yingyan Lin. Halo: Hardware-aware learning to optimize. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2020.

[69] Diana Wofk, Fangchang Ma, Tien-Ju Yang, Sertac Karaman, and Vivienne Sze. Fastdepth: Fast monocular depth estimation on embedded systems. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6101–6108. IEEE, 2019.

[70] Patrick Mochel and Mike Murphy. sysfs - The filesystem for exporting kernel objects. https://www.kernel.org/doc/Documentation/filesystems/sysfs.txt, accessed 2019-11-21.

[71] Texas Instruments Inc. INA3221 Triple-Channel, High-Side Measurement, Shunt and Bus Voltage Monitor. http://www.ti.com/product/INA3221, accessed 2019-11-21.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] We discuss the trade-off between pre-indexing time and query time of hash tables in Section 4
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] Yes, we read and follow the guidelines and make sure we conforms them.
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] We state the full assumptions.
   (b) Did you include complete proofs of all theoretical results? [Yes] We include the complete proofs in the supplementary material.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We include the data other experimental information. We would release the code under the permission of all stakeholders.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We specify the hyperparameters and perform a parameter study.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] We perform the average learning curves instead of error bars.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We specify these conditions
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] It does not have any personal information
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]