BANDIT SAMPLING FOR FASTER NEURAL NETWORK TRAINING WITH SGD

Anonymous authors
Paper under double-blind review

ABSTRACT

Importance sampling is a valuable technique in deep learning that involves sampling useful training examples more frequently to improve learning algorithms. However, obtaining reliable sample importance estimates early on in training can be challenging, as existing importance sampling methods can be computationally expensive and slow to converge. In this work, we propose a novel sampling schemed based on Multi-arm bandits (MAB). The proposed sampler is able to achieve higher validation accuracies significantly earlier in the training, compared to baselines.

1 INTRODUCTION

Deep learning model training has become more efficient lately thanks to stochastic gradient-based optimizers like SGD and Adam, which have aided in improving learning algorithms. However, during training, uniformly sampling a mini-batch of training examples in each iteration may cause the learning algorithm to converge slowly, as not all examples contribute equally to the prediction task of the deep learning model Ganapathiraman et al.; Katharopoulos & Fleuret (2018); Alain et al. (2015). To address this issue, importance sampling techniques have been proposed, which involves sampling useful training examples more frequently based on their importance score. Importance sampling methods are particularly useful when training on a limited budget, such as during hyperparameter optimization or model selection.

One of the challenges in importance sampling for neural network training is to obtain reliable estimates of sample importance early on in training. Existing importance sampling algorithms rely on local metrics such as loss and gradient norms to estimate sample importance, which can be computationally expensive and often require several iterations to achieve significant improvements Katharopoulos & Fleuret (2017; 2018); Alain et al. (2015); Chang et al. (2017). This can lead to slower convergence of the learning algorithm and longer training times.

Reinforcement learning (RL) algorithms are famously well-suited for cases where training data is scarce, such as in recommender systems with the cold start problem, and are a natural candidate for importance sampling. RL can help alleviate the issue of unreliable sample importance estimates by providing a more data-efficient approach to sample selection. Recently, Liu et al. (2020) proposed a bandit-baed sampling strategy for ADAM with promising performance. In the proposed RL-based sampler, the Multi-Arm Bandit Sampler (MABSampler), loss is used as the reward to update and improve the sampler at each iteration. MABSampler has shown to consistently outperform other importance sampling techniques such as SGD and p-SGD Chang et al. (2017), achieving maximum accuracy in significantly fewer iterations than other methods

2 METHOD

The training architecture using the proposed MABSampler is shown in Figure 2. There are 2 main components:

1. **Importance sampler.** The sampler has b one-arm bandits, where b is the batch size. During training, bandit MAB_i samples one arm, which corresponds to one training example from a pool of N/b examples, which feeds into the *Batch Sampler*. The training data is split randomly into chunks of size N/b among the bandits, with *class stratification* to ensure that each bandit has access to the same data distribution; minimizing initial sampling bias.

2. **Loss Binarizer.** The bandits are updated using a reward function, that is based on the training loss of a batch. The Loss Binarizer module takes in a per-sample training loss and binarizes it using the following function:$r(x_a) = \begin{cases} 1.0, & \text{if } \ell(x_a) \geq m_l \\ 0, & \text{Otherwise} \end{cases}$, where m_l is a moving average of the training loss and $\ell(x_a)$ is the training loss of x_a.

3 EXPERIMENTS

The bandits were updated using the Thompson sampling algorithm, which gave the best results in our experiments, compared to other methods such as ϵ-greedy, softmax action selection. We also tried parametric and contextual bandits, but they were really slow to train. Figure 2 (a & b) show validation accuracies on image classification benchmarks. MABSampler achieves significantly higher performance much earlier in the training compared to baselines such as p-SGD and c-SGD.

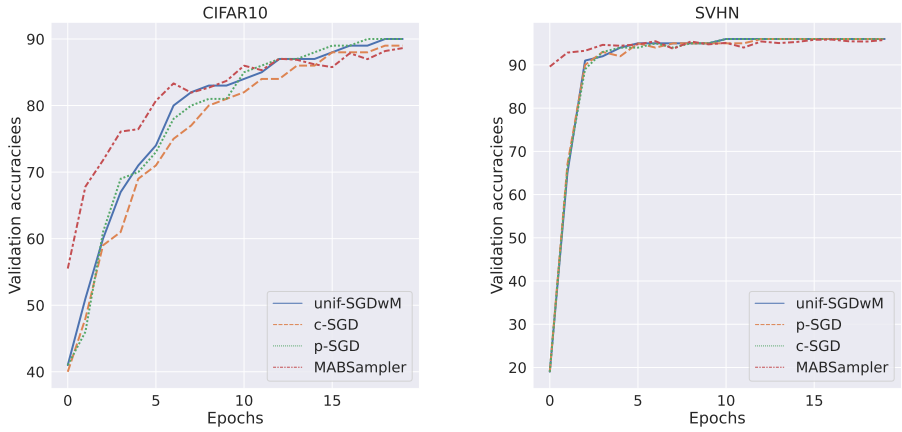

(a) (b)

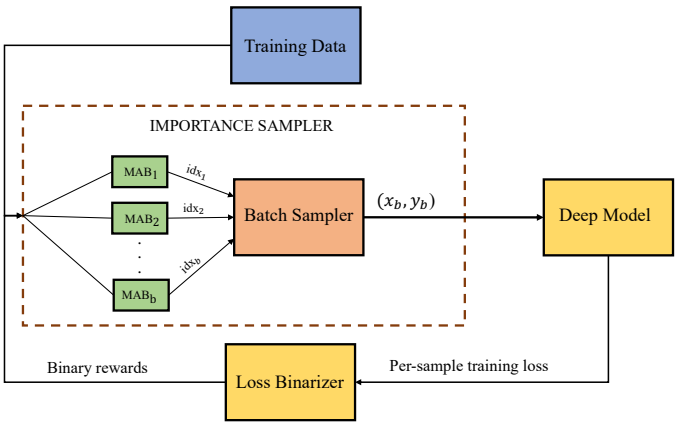

(c)

Figure 1: (a) & (b) Validation accuracies of MAB sampler along with popular importance sampling algorithms on popular image classification benchmarks. (c) Training pipeline of the proposed Bandit sampler for SGD.

4 CONCLUSION

In this work, we proposed a novel bandit-based non-uniform sampling technique for training neural networks with SGD. Empirical studies show that the proposed model outperformed existing baselines significantly in earlier batches, offering great benefit for budgeted training scenarios.

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2023 Tiny Papers Track.

REFERENCES

Guillaume Alain, Alex Lamb, Chinnadhurai Sankar, Aaron Courville, and Yoshua Bengio. Variance reduction in sgd by distributed importance sampling. *arXiv preprint arXiv:1511.06481*, 2015.

Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. Active bias: Training more accurate neural networks by emphasizing high variance samples. *Advances in Neural Information Processing Systems*, 30, 2017.

Vignesh Ganapathiraman, Francisco Calderon Rodriguez, and Anila Joshi. Impon: Efficient importance sampling with online regression for rapid neural network training. In *Has it Trained Yet? NeurIPS 2022 Workshop*.

Angelos Katharopoulos and François Fleuret. Biased importance sampling for deep neural network training. *arXiv preprint arXiv:1706.00043*, 2017.

Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. In *International conference on machine learning*, pp. 2525–2534. PMLR, 2018.

Rui Liu, Tianyi Wu, and Barzan Mozafari. Adam with bandit sampling for deep learning. *Advances in Neural Information Processing Systems*, 33:5393–5404, 2020.

A ADDITIONAL PLOTS FOR IMAGE CLASSIFICATION

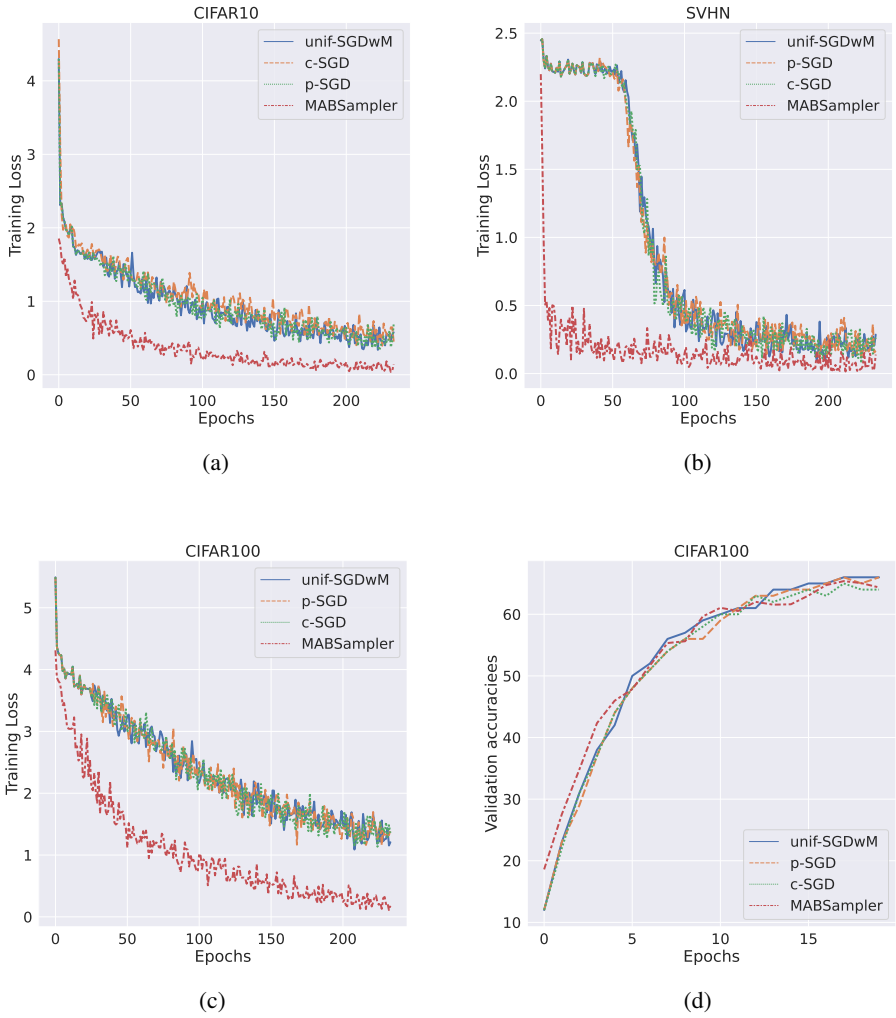

Figure 2: (a, b, c) Training loss of the proposed MAB-based sampler along with baselines on popular image classification benchmark datasets. (d). Validation accuracy on CIFAR100 dataset (additional plot from the validation accuracies plot from the main paper.)