Towards A Performance Evaluation for Federated Averaging on Edge Devices

Asfia Kawnine, Hung Cao
Analytics Everywhere Lab, Faculty of Computer Science
University of New Brunswick

Abstract

Federated learning is a distributed learning approach that enables edge devices to train machine learning models without sharing their data. Federated averaging is a popular algorithm for federated learning that aggregates local machine learning model updates from edge devices to train a global model. We are working towards an evaluation of the performance of federated averaging algorithm on edge devices and compare it to a centralized approach.

Introduction

Edge devices, such as smartphones and IoT devices, are becoming increasingly popular for machine learning applications. Due to limited computing and storage resources, federated learning can be used in edge devices to train models without the need to share their data with a central server. Federated Averaging (FedAvg) is a popular algorithm for federated learning that aggregates local model updates to train a global model. The objective of this research would be to evaluate and compare the performance of the FedAvg algorithm in the context of edge devices with streaming data using limited computational and communication resources. The research aims to enable collaborative machine learning while getting continuous data from the IoT devices and preserving data privacy.

Proposed Experiment

Our potential experiment is to evaluate the performance of the FedAvg algorithm on edge devices with limited computational and communication resources. The experiment would use streaming data that are collected from IoT devices or sensors, and evaluate the convergence rate, accuracy, and communication efficiency of the algorithm.

Literature Review

Previous research has shown that FedAvg can achieve comparable accuracy to centralized learning while preserving the privacy of edge devices. Also, FedAvg has fewer hyperparameters to tune than other algorithms, which can make it easier to use and deploy. The communication cost is relatively low compared to some other Federated Learning algorithms because only model updates are sent between the server and clients during training, not the entire dataset. Below a comparison is shown of different federated aggregated models.

	FedAvg [1,2]	**FedAvgM [3,4]**	**FedAdaGrad [5,6]**	**FedYogi [7,8]**	**FedAdam [9,10]**
Mechanism	Averaging of local model updates from all participating nodes	Averaging of local model updates with momentum	Averaging of local model updates with adaptive learning rate	Averaging of local model updates with adaptive learning rate	Averaging of local model updates with adaptive learning rate
Advantages	Simple and easy to implement	Faster convergence than FedAvg	Fast convergence and can handle non-i.i.d data	Handles non-i.i.d data and noisy gradients well	Fast convergence, handles non-i.i.d data, and noisy gradients
Limitations	Convergence may be slow due to communication bottleneck	Requires additional hyperparameters to be tuned	Requires additional hyperparameters to be tuned	Requires additional hyperparameters to be tuned	Requires additional hyperparameters to be tuned

Proposed Method

We are aiming to conduct an experiment on a set of edge devices, each running a simple classification task. The devices communicate with a central server using the federated averaging algorithm to collaboratively train a machine learning model. In the diagram below, our proposed the dataflow is presented as well as the overall system infrastructure.

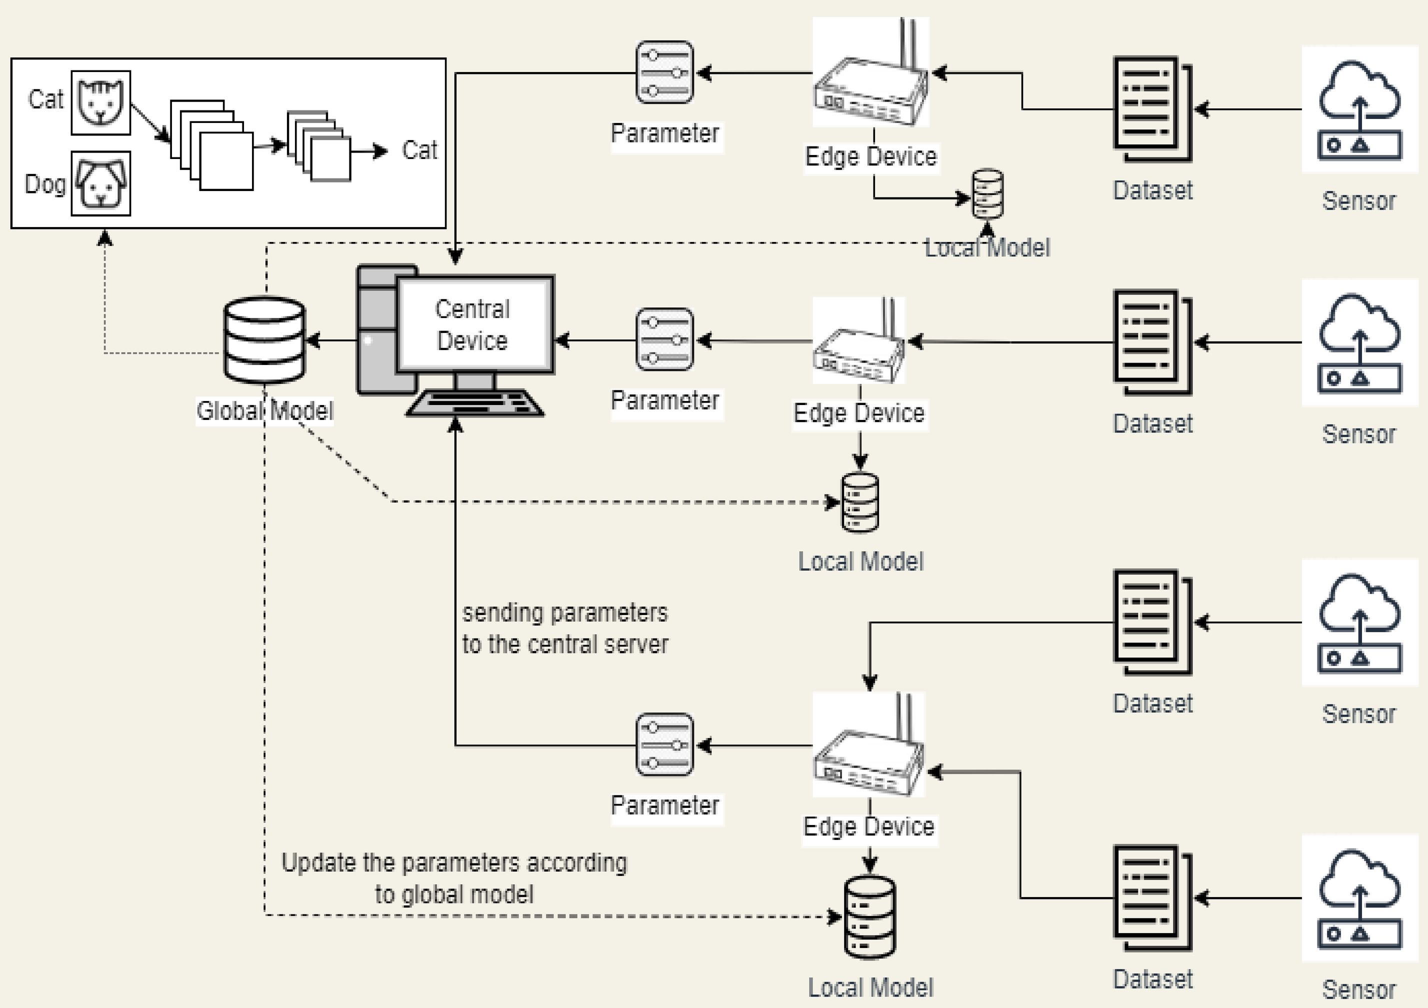

Expected Results

The results of the experiment may show that FedAvg performs well in homogeneous edge device environments where data is distributed in an i.i.d manner. Additionally, the experiment may reveal new optimization techniques or communication protocols that can improve the performance of FedAvg on edge devices, further advancing the field.

Conclusion

Our proposed system would ensure data privacy, alongside would enable collaborative ML. Also, would have diverse applications, such as in healthcare, smart homes, autonomous vehicles, and various other sectors where edge computing is used.

References

1. H. B. McMahan et al. "Communication-efficient learning of deep networks from decentralized data." AISTATS, 2017.
2. B. Konečnỳ et al. "Federated learning: Strategies for improving communication efficiency." arXiv:1610.05492, 2016.
3. S. Sattler et al. "Robust and communication-efficient federated learning from non-iid data." AISTATS, 2021.
4. W. Chen et al. "Federated meta-learning for recommendation." arXiv:2103.07454, 2021.
5. P. Wen et al. "Fedadagrad: Learning decentralized weight-adaptive learning rates for scalable federated learning." ICML, 2021.
6. T. Chen et al. "Federated learning of heterogeneous models." arXiv:1902.00146, 2019.
7. S. Sattler et al. "Yogi: An optimizer for federated learning." NeurIPS, 2020.
8. X. Li et al. "Federated learning on non-iid data silos: An experimental study." arXiv:1905.08841, 2019.
9. S. Reddi et al. "Adaptive federated optimization." NeurIPS, 2020.
10. L. Yu et al. "Fedadam: A robust optimization method for federated learning with non-iid data." IEEE Transactions on Cybernetics, 2021.