

Efficient Node Selection in Private Personalized Decentralized Learning

Edvin Listo Zec^{1,3}, Johan Östman², Olof Mogren¹, Daniel Gillblad²
edvin.listo.zec@ri.se

Abstract—Personalized decentralized learning is a promising paradigm for distributed learning, enabling each node to train a local model on its own data and collaborate with other nodes to improve without sharing any data. However, this approach poses significant privacy risks, as nodes may inadvertently disclose sensitive information about their data or preferences through their collaboration choices. We propose Private Personalized Decentralized Learning (PPDL), a novel approach that combines secure aggregation and correlated adversarial multi-armed bandit optimization to protect node privacy while facilitating efficient node selection. By leveraging dependencies between different arms, represented by potential collaborators, we demonstrate that PPDL can effectively identify suitable collaborators solely based on aggregated models. Additionally, we show that PPDL surpasses previous non-private methods in model performance on standard benchmarks under label and covariate shift scenarios.

I. INTRODUCTION

Collaborative machine learning is a technique in which a group of actors collaboratively trains a joint model while preserving the privacy of their individual datasets [7]. Two prevalent approaches to collaborative machine learning are federated learning (FL) and decentralized learning. FL has several inherent limitations and risks, including the difficulty in finding a trustworthy third party to coordinate the training process, and the need for large institutions to maintain autonomy over their data. Furthermore, FL’s scalability can be restricted, and it may also have a single-point-of-failure. Decentralized learning, on the other hand, eliminates the need for a central server by directly communicating model parameters among peers in the learning setup using a communication protocol, such as gossip learning [5], [4]. However, this approach may not be appropriate for non-iid settings where several distinct learning objectives may be present.

The idea of each node identifying useful peers in the network to train a personalized model was proposed in [9]. A score-based method, decentralized adaptive clustering (DAC), was presented in [6] where each node scores its neighboring peers based on the the empirical loss. While DAC manages to find beneficial nodes and identifies clusters in the network, model parameters are still communicated over the network in plain text and the peers receiving the updates must hence be trusted. As such, DAC is vulnerable to inference attacks. Since a node only receives an aggregate of the parameter updates of M nodes at a given point in time, it cannot infer a score on

the similarity of any one of the peers in the aggregate (as in DAC); such a score can only be computed for the aggregate. Instead, our solution exploits dependencies between different group selections and makes use of adversarial multi-armed bandit optimization to efficiently find the subsets of peers that are beneficial for collaboration. Our solution has a communication efficiency and performance similar to that of previous methods, but adds a higher level of privacy.

II. DECENTRALIZED LEARNING BY FINDING USEFUL COLLABORATIONS

We consider several decentralized learning tasks over a network consisting of K nodes, where each node $i \in [K]$ has a *private* local data distribution \mathcal{D}_i over the input features $x \in \mathcal{X}$ and the corresponding label $y \in \mathcal{Y}$. Each node $i \in [K]$ is equipped with a machine learning model f_i with model parameters $w_i \in \mathbb{R}^d$ and a real-valued loss function $\ell(f_i(w_i; x), y) : \mathbb{R}^d \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$. The objective of each node i is to solve its own task as well as possible (referred to as *personalized learning*) by minimizing the expected loss over the local data distribution,

$$w_i^* = \arg \min_{w_i \in \mathbb{R}^d} \mathbb{E}_{(x,y) \sim \mathcal{D}_i} [\ell(f_i(w_i; x), y)]. \quad (1)$$

Personalized decentralized learning faces the challenge of determining if nodes should collaborate based on similar data distributions. Collaboration may hurt performance if distributions are dissimilar, but collaboration with similar nodes can help towards the goal in (1). Privacy concerns make it difficult to reveal node data. We propose a private method to identify nodes with similar local datasets, similar to [6], [8]. Nodes aim to identify nodes in \mathcal{M} whose parameters can help approach w_i^* in (1), achieved through decentralized federated averaging over T rounds.

Privacy. Although federated averaging is commonly advertised as being private, recent results have demonstrated attacks able to recover training data from the models [1]. To protect the nodes from such attacks, we utilize secure aggregation to ensure that a node who queried multiple model parameters from a subset of its neighbors only get to observe an aggregate of those models.

Multi-armed bandits. Choosing collaborative nodes from a large pool in limited training rounds is difficult, so a node must evaluate a sampled group with a local measure, such as local accuracy. However, local accuracy is stochastic and non-stationary due to dependence on node selection and training of others in the network. Therefore, we frame the group-selection problem as an adversarial multi-armed bandit

¹ RISE Research Institutes of Sweden

² AI Sweden

³ KTH Royal Institute of Technology

problem. In the problem at hand, some groups will have overlapping member nodes. To leverage this idea, we utilize the framework of pseudo-rewards, as presented in [2]. Let the different groups available to node i be indexed from $1, \dots, C_i$ and, w.l.o.g., let the reward from choosing group $j \in [C_i]$ at time t satisfy $r_j^{(t)} \in [0, 1]$. We define the pseudo-rewards $s_{i,j}^{(t)}(\alpha_j^{(t)}) \in [0, 1]$ as an upper bound on the expected reward on $r_i^{(t)}$ given that we observe $r_j^{(t)}$ for $j \in [C_i]$ and $l \in [C_i] \setminus \{j\}$. This is mathematically represented as $\mathbb{E}[r_l^{(t)} | r_j^{(t)} = \alpha_j^{(t)}] \leq s_{i,j}^{(t)}(\alpha_j^{(t)})$. Let $u_{l,j} \in \{0, \dots, M-1\}$ denote the number of overlapping nodes between group l and group j . We consider pseudo-rewards of the form

$$s_{i,j}^{(t)}(\alpha_j^{(t)}) = \min \left\{ \alpha_j^{(t)} + \frac{q(t)}{u_{l,j}}, 1 \right\} \quad (2)$$

III. PRIVATE MULTI-ARMED BANDITS FOR NODE SELECTION

Here, we present our bandit algorithm for a given node $i \in \mathcal{N}$. For ease of notation, we exclude the node index in the sequel. Let $k^{(t)} \in [C_i]$ denote the group chosen at time t and let $n_{k^{(t)}}(t)$ denote the number of times group $k^{(t)}$ has been chosen after t rounds. The empirical reward from choosing group $j \in [C_i]$ is defined as $\mu_j(t) = \frac{\sum_{\tau=1}^t \mathbf{1}\{k^{(\tau)}=j\} r_j^{(\tau)}}{n_j(t)}$ and the empirical pseudo-reward for group $l \in [C_i] \setminus \{j\}$ when group $j \in [C_i]$ is selected, is given by $\phi_{l,j}(t) = \frac{\sum_{\tau=1}^t \mathbf{1}\{k^{(\tau)}=j\} s_{l,j}^{(\tau)}(r_j^{(\tau)})}{n_j(t)}$. [2] reduced the size of the multi-armed bandit problem by only selecting arms that are empirically competitive, i.e., arms whose minimum empirical pseudo-rewards exceeds the maximum empirical reward. To this end, we define the set of significant arms as $\mathcal{S}_i^{(t)} = \{j \in [C_i] : n_j(t) > t/N\}$ and let $\bar{k}^{(t)} = \arg \max_{l \in \mathcal{S}_i^{(t)}} \mu_l(t)$. The set of empirically competitive arms is defined as $\mathcal{A}_i^{(t)} = \{j \in [C_i] : \min_{l \in \mathcal{S}_i^{(t)}} \phi_{j,l}(t) \geq \mu_{\bar{k}^{(t)}}(t)\} \cup \{\bar{k}^{(t)}\}$. We use the *Tsallis-Inf* algorithm to consider adversarial rewards, which achieves optimal scaling pseudo-regret [10]. The value of $q(t)$ in (2) determines the size of empirical pseudo-rewards, affecting the competitive set. Large $q(t)$ encourages exploration, while small $q(t)$ encourages exploitation.

IV. EXPERIMENTS

We experiment with various cluster configurations and use CIFAR-10 and Fashion-MNIST datasets commonly used in literature for decentralized learning evaluations (Section 3.1 [3]). The results of our label shift experiment with two clusters (animals and vehicles) on CIFAR-10 are presented in Table I. We observe that both PPDL and DAC perform well, with PPDL achieving superior results. The highest accuracy is achieved with PPDL-var, in which $q(t)$ is exponentially decayed with respect to the number of communication rounds. It is worth noting that Random performs worse than local training without collaboration, likely due to model poisoning caused by nodes communicating with incorrect clusters. The node models learn different representations for the different clusters, and when merging models from two distinct

clusters, the resulting model is inferior due to the significant dissimilarity between the models, a phenomenon known as *client drift*. Future research could explore aggregation methods for models trained on different datasets to enhance node robustness and understand the impact of the number of nodes participating in secure aggregation on privacy. We plan to extend our algorithm to allow for group sizes of arbitrary sizes at each node by leveraging scaling bandits.

TABLE I: CIFAR-10 label shift test accuracy with 60 nodes in 'animal' cluster and 40 nodes in 'vehicle' cluster.

Method	Vehicles	Animals	Mean
PPDL	51.86	36.31	43.81
PPDL-var	52.86	36.33	44.60
DAC	52.78	33.87	43.32
Random	44.79	30.00	37.40
Local	51.10	35.11	43.11
Oracle	57.17	39.74	48.45

REFERENCES

- [1] DIMITROV, D. I., BALUNOVIC, M., KONSTANTINOV, N., AND VECHEV, M. Data leakage in federated averaging. *Transactions on Machine Learning Research* (2022).
- [2] GUPTA, S., CHAUDHARI, S., JOSHI, G., AND YAGAN, O. Multi-armed bandits with correlated arms. *IEEE Transactions on Information Theory* 67, 10 (2021), 6711–6732.
- [3] KAIROUZ, P., MCMAHAN, H. B., AVENT, B., BELLET, A., BENNIS, M., BHAGOJI, A. N., BONAWITZ, K., CHARLES, Z., CORMODE, G., CUMMINGS, R., D'OLIVEIRA, R. G. L., EICHNER, H., ROUAYHEB, S. E., EVANS, D., GARDNER, J., GARRETT, Z., GASCÓN, A., GHAZI, B., GIBBONS, P. B., GRUTESER, M., HARCHAQUI, Z., HE, C., HE, L., HUO, Z., HUTCHINSON, B., HSU, J., JAGGI, M., JAVIDI, T., JOSHI, G., KHODAK, M., KONEČNÝ, J., KOROLOVA, A., KOUZHANFAR, F., KOYEJO, S., LEPOINT, T., LIU, Y., MITTAL, P., MOHRI, M., NOCK, R., ÖZGÜR, A., PAGH, R., RAYKOVA, M., QI, H., RAMAGE, D., RASKAR, R., SONG, D., SONG, W., STICH, S. U., SUN, Z., SURESH, A. T., TRAMÈR, F., VEPAKOMMA, P., WANG, J., XIONG, L., XU, Z., YANG, Q., YU, F. X., YU, H., AND ZHAO, S. Advances and open problems in federated learning, 2019.
- [4] KEMPE, D., DOBRA, A., AND GEHRKE, J. Gossip-based computation of aggregate information. In *44th Annual IEEE Symposium on Foundations of Computer Science* (2003), pp. 482–491.
- [5] LIAN, X., ZHANG, C., ZHANG, H., HSIEH, C.-J., ZHANG, W., AND LIU, J. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems* (2017), vol. 30.
- [6] LISTO ZEC, E., EKBLUM, E., WILLBO, M., MOGREN, O., AND GIRDIJIAUSKAS, S. Decentralized adaptive clustering of deep nets is beneficial for client collaboration. *FL-IJCAI'22: International Workshop on Trustworthy Federated Learning* (2022).
- [7] MCMAHAN, B., MOORE, E., RAMAGE, D., HAMPSON, S., AND ARCAS, B. A. Y. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (2017), vol. 54, pp. 1273–1282.
- [8] SUI, Y., WEN, J., LAU, Y., ROSS, B. L., AND CRESSWELL, J. C. Find your friends: Personalized federated learning with the right collaborators, 2022.
- [9] ZANTEDESCHI, V., BELLET, A., AND TOMMASI, M. Fully decentralized joint learning of personalized models and collaboration graphs. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics* (2020), pp. 864–874.
- [10] ZIMMERT, J., AND SELDIN, Y. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *J. Mach. Learn. Res.* 22, 1 (jul 2022).